

Fear of model misspecification and the robustness premium*

Konstantinos Angelopoulos
University of Glasgow

James Malley
University of Glasgow and CESifo

September 17, 2010

Abstract

Robust decision making implies welfare costs or robustness premia when the approximating model is the true data generating process. To examine the importance of these premia at the aggregate level we employ a simple two-sector dynamic general equilibrium model with human capital and introduce an additional form of precautionary behavior. The latter arises from the robust decision maker's ability to reduce the effects of model misspecification through allocating time and existing human capital to this end. We find that the extent of the robustness premia critically depends on the productivity of time relative to that of human capital. When the relative efficiency of time is low, despite transitory welfare costs, there are gains from following robust policies in the long-run. In contrast, high relative productivity of time implies misallocation costs that remain even in the long-run. Finally, depending on the technology used to reduce model uncertainty, we find that while increasing the fear of model misspecification leads to a net increase in precautionary behavior, investment and output can fall.

*We would like to thank Efram Castelnuovo, Martin Ellison, Campbell Leith, John Tsoukalas, Apostolis Philippopoulos, Ulrich Woitek, Raf Wouters, participants at the Workshop on "New Developments in Dynamic Stochastic General Equilibrium Models", University of Nottingham, Dec. 2009 and seminar participants at the Athens University of Economics and Business Dec. 2009 and University of Zurich, Aug. 2010 for helpful comments and suggestions. The usual disclaimer applies.

I hope for nothing. I fear nothing. I am free. [N. Kazantzakis, *epitaph*]

1 Introduction

The path breaking research of Hansen and Sargent (e.g. 2001, 2008) and Hansen *et al.* (2006) on robust control theory has focused the attention of the economics profession on the importance of model uncertainty and its effects on decision making. This approach starts by acknowledging that economic agents make choices in an environment where they must rely on approximations to the true model generating the observed data. These choices, depending on the extent of the gap between the correct and approximating models *ex post*, in turn imply that resources will be less than optimally allocated. Hence, a robust decision maker aims to make choices that will guarantee a minimum level of efficiency or welfare conditional on the degree to which he fears that his approximating model is misspecified. To hedge against the fact that his model might not be correct, he engages in precautionary behavior by increasing his buffer stocks of non-human and human wealth (see Hansen and Sargent (2008) and Hansen, Sargent and Tallarini (1999)). These actions and their associated opportunity costs are obviously increasing in the degree to which the agent is uncertain about his model.

Since robust choices come at a cost, when fears of model misspecification do not materialize *ex post*, the decision maker has to decide, *ex ante*, how conservative to be. If he is not conservative enough, he risks being exposed to big misspecification errors. On the other hand, if he is too conservative, he risks paying unnecessary costs or "robustness premia". Understanding the implications of making more conservative robust choices is highly relevant for many important decisions that a society must take, for instance in designing how to react to climate change or the potential spread of disease or banking crises. In fact, the main argument usually raised against (more conservative) robust policies is that the opportunity costs associated with these are too high.

In the robustness literature to date, fear of model misspecification or model uncertainty is generally determined exogenously. In particular, it is defined as the discounted life-time sum of the conditional relative entropy of the true (but unknown) model relative to the approximating model that the economic agent must use.¹ The two models (true and approximating) are typically restricted to be close in the statistical sense that the life-time

¹Conditional relative entropy in this context is defined as the expected log-likelihood ratio of the true relative to the approximating model, evaluated with respect to the true model.

conditional relative entropy cannot be larger than an exogenous constant. This constant can be calibrated so that given a finite amount of data, a decision maker would find it difficult to statistically distinguish members of a set of alternative models. In contrast to this approach, we endogenize the entropy constraint and allow the robust decision maker to protect himself against model uncertainty by earmarking resources to mitigate the potential negative effects of model misspecification. This allows us to obtain a balanced assessment of the premia associated with following robust policies since buffer stocks of productive resources, although not "optimal" when the approximating is close to the true model, are nevertheless useful for the economy and can, under the right circumstances, increase productivity in the long-run.

To implement our ideas in a parsimonious and transparent framework, we extend a simple two-sector dynamic general equilibrium model with human capital. The model is a variant of the two-sector real business cycle models in e.g. Mankiw et al. (1992), Perli and Sakellaris (1998) and DeJong and Ingram (2001). We augment this model by adding model uncertainty and a technology using existing human capital and time to create knowledge to reduce the effects of model misspecification. Time in this context complements human capital and is employed at the expense of its alternative uses. The complementarity between these two inputs implies that more educated individuals (or societies) will understand more about their economic environment if they put the same effort time as less educated individuals (or societies).

To obtain the robustness premium, we compare outcomes obtained from making robust decisions, to those from making optimal choices when the approximating model governs the data, i.e. when fears of model misspecification are unfounded. Hence, in this scenario, effort time to restrict the effects of model misspecification implies a misallocation of resources. On the other hand, efforts to mitigate model misspecification by accumulating human capital will have a productive effect on the economy.

Our main findings are as follows: (i) Despite more conservative robust decisions leading to higher robustness premia, when the mitigation of model misspecification depends on a low productivity of time relative to human capital, robust choices increase welfare in the long-run; (ii) If the technology used to mitigate the effects of model misspecification relies on a low productivity of human capital relative to time then the misallocation of resources under the approximating model is so strong that more conservative robust policies imply welfare costs which remain even in the long-run; and (iii) Increasing the fear of model misspecification leads to a net increase in precautionary behavior, however, depending on the technology used to reduce

model uncertainty, investment and output can fall.

2 A misspecified human capital model

To illustrate the main ideas in this paper we use a standard two-sector RBC model with human capital in the presence of model uncertainty. In contrast to approaches which study optimal choices under the correct model specification, following Hansen and Sargent (2001, 2008), we allow for particular types of misspecification in the dynamic processes of the model.

The general equilibrium solution we derive in this environment consists of a system of dynamic relations specifying the paths of output, Y_t , consumption, C_t , investment in physical capital, I_t , investment in human capital, I_t^h , physical capital stock, K_{t+1} , human capital stock, H_{t+1} , knowledge to mitigate model misspecification, h_t^i , and the fractions of time allocated to work, n_t , education, e_t and reducing misspecification, q_t . The model's stochastic exogenous processes include total factor productivity, \tilde{a}_t , human capital-specific productivity, φ_t and investment-specific productivity, τ_t .

In this section we first discuss the non-linear setup for the stationary model, under model misspecification, followed by the linear-quadratic (LQ) approximation of the latter. We then solve the stationary LQ representation to obtain decision or policy rules which are robust to model misspecification.

2.1 Model setup

Taking the stationary paths of the stochastic exogenous processes, $\{\tilde{a}_t, \tau_t, \varphi_t\}_{t=0}^{\infty}$, as given, the representative agent chooses the stationary variables $\{y_t, c_t, i_t, i_t^h, n_t, e_t, q_t, k_{t+1}, h_{t+1}\}_{t=0}^{\infty}$ to maximize expected discounted lifetime utility:

$$\max_{\{y_t, c_t, i_t, i_t^h, n_t, e_t, q_t, k_{t+1}, h_{t+1}\}_{t=0}^{\infty}} E_0 \sum_{t=0}^{\infty} \tilde{\beta}^t \frac{\left[(c_t)^\mu (l_t)^{(1-\mu)} \right]^{1-\varsigma}}{1-\varsigma} \quad (1a)$$

subject to:

$$l_t = 1 - n_t - e_t - q_t \quad (1b)$$

$$c_t + i_t + i_t^h = y_t \quad (1c)$$

$$gh_{t+1} = (1 - \delta^h)h_t + \kappa (e_t h_t)^\gamma (i_t^h)^{1-\gamma} + \varphi_t, h_0 > 0 \quad (1d)$$

$$gk_{t+1} = (1 - \delta^k)k_t + i_t + \tau_t, k_0 > 0 \quad (1e)$$

$$y_t = \exp(\tilde{a}_t) (k_t)^\alpha (h_t)^\zeta (n_t)^{1-\alpha-\zeta} \quad (1f)$$

where, E_0 is the conditional expectations operator; $0 < (\check{\beta} \equiv \beta g^{\mu(1-\varsigma)}) < 1$ is the time discount factor; g is the gross growth rate of exogenous labor augmenting technical progress;² $\varsigma > 1$ is the relative risk aversion parameter; $0 < \mu < 1$ is the weight given to consumption relative to leisure, l_t , in the utility function; $0 < \delta^k, \delta^h < 1$ are the physical and human capital depreciation rates; $\kappa > 0$ and $0 < \gamma < 1$ capture respectively, the level and productivity of knowledge in new human capital production; and $0 < \alpha, \zeta < 1$, $\alpha + \zeta < 1$ measure the productivity of physical and human capital respectively in the Cobb-Douglas production function.

There are three stochastic exogenous technology processes in the model. The first is total factor productivity (TFP), and is given by

$$\tilde{a}_{t+1} = (1 - \rho^a)\tilde{a} + \rho^a\tilde{a}_t + \sigma^a(\varepsilon_{t+1}^a + w_{t+1}^a) \quad (1g)$$

where, $\tilde{a}_t = \log(A_t)$; $A_0 > 0$; $0 < \rho^a < 1$ is a first order autocorrelation coefficient; \tilde{a} is a constant; and ε_{t+1}^a is a Gaussian random variable distributed identically and independently through time, with zero mean and unit variance. In addition, human capital- and investment-specific productivity processes are given by

$$\varphi_{t+1} = \rho^\varphi\varphi_t + \sigma^\varphi(\varepsilon_{t+1}^\varphi + w_{t+1}^\varphi) \quad (1h)$$

$$\tau_{t+1} = \rho^\tau\tau_t + \sigma^\tau(\varepsilon_{t+1}^\tau + w_{t+1}^\tau) \quad (1i)$$

where again $0 < \rho^\varphi, \rho^\tau < 1$ are first order autocorrelation coefficients; and $\varepsilon_{t+1}^\varphi, \varepsilon_{t+1}^\tau$ are Gaussian random variables distributed identically and independently through time, with zero means and unit variances.

Model uncertainty in this setup is captured by the unknown stochastic perturbations, $w_{t+1}^a, w_{t+1}^\varphi$ and w_{t+1}^τ , to the exogenous processes (1g – 1i). In particular, these may be non-linear time dependent functions, ϑ_t , of the past states, i.e.

$$w_{t+1}^j = \vartheta_t^j(x_t, x_{t-1}, \dots) \quad (2)$$

for $j = a, \varphi, \tau$, where $\{\vartheta_t^j\}$ are sequences of measurable functions; x_t is a vector including the state variables of the model (defined below); and $\sigma^a, \sigma^\varphi, \sigma^\tau > 0$ scale the variances of the ε_{t+1}^j and w_{t+1}^j shocks.

²Following the standard RBC literature, labour augmenting technical progress Z_t , is deterministic and grows according to the gross rate g , such that $\frac{Z_{t+1}}{Z_t} = g$ and $Z_0 > 0$. The lower case stationary magnitudes for $\{y_t, c_t, i_t, i_t^h, k_t, h_t\}$ denote that their non-stationary values have been detrended by Z_t .

2.1.1 Model misspecification

As pointed out above, the decision maker cannot use the correct representation for the exogenous processes given by (1g – 1i) since the w_{t+1}^j are unknown. Instead, a misspecified approximating model for these must be employed, i.e.

$$\tilde{a}_{t+1} = (1 - \rho^a)\tilde{a} + \rho^a\tilde{a}_t + \sigma^a\tilde{\varepsilon}_{t+1}^a \quad (3a)$$

$$\varphi_{t+1} = \rho^\varphi\varphi_t + \sigma^\varphi\tilde{\varepsilon}_{t+1}^\varphi \quad (3b)$$

$$\tau_{t+1} = \rho^\tau\tau_t + \sigma^\tau\tilde{\varepsilon}_{t+1}^\tau \quad (3c)$$

where $\tilde{\varepsilon}_{t+1}^j \sim iid N(0, 1)$. Following Hansen and Sargent (2008), by introducing model uncertainty in this manner, we essentially bound the model the decision maker employs, i.e. (1a – 1f) and (3a – 3c), with a set of alternative models of the form given by (1a – 1i). The latter clearly allow for more general forms of the approximation errors than those implied by $\tilde{\varepsilon}_{t+1}^j$. The correct model in this context can then be viewed as a distorted or perturbed version of the approximating model since the conditional means of w_{t+1}^j are non-zero and as suggested above can influence the evolution of the states.

To operationalize the notion that the approximating model using (3a – 3c) provides good predictions when (1g – 1i) actually generates the data, we follow Hansen and Sargent (2008) and restrict w_{t+1}^j as follows. First, write the perturbed model in matrix form as:

$$\tilde{z}_{t+1} = \tilde{z}_0 + \rho\tilde{z}_t + \sigma(\varepsilon_{t+1} + w_{t+1}), \text{ where} \quad (4)$$

$$\begin{aligned} \tilde{z}_t &= [\tilde{a}_t \quad \varphi_t \quad \tau_t]'; \quad \tilde{z}_0 = [(1 - \rho^a)\tilde{a}]; \\ w_{t+1} &= [w_{t+1}^a \quad w_{t+1}^\varphi \quad w_{t+1}^\tau]'; \quad \varepsilon_{t+1} = [\varepsilon_{t+1}^a \quad \varepsilon_{t+1}^\varphi \quad \varepsilon_{t+1}^\tau]'; \\ \rho &= \begin{bmatrix} \rho^a & 0 & 0 \\ 0 & \rho^\varphi & 0 \\ 0 & 0 & \rho^\tau \end{bmatrix}; \quad \sigma = \begin{bmatrix} \sigma^a & 0 & 0 \\ 0 & \sigma^\varphi & 0 \\ 0 & 0 & \sigma^\tau \end{bmatrix}. \end{aligned}$$

Second, let f_a denote the one-step transition density associated with the approximating model and f the one-step transition density associated with the correct or perturbed model. In the present setting, the assumptions employed imply that the transition density for the approximating model is:

$$f_a(\tilde{z}_{t+1} | \tilde{z}_t) \sim N(\tilde{z}_0 + \rho\tilde{z}_t, \sigma\sigma') \quad (5a)$$

while, the transition density for the perturbed model is:³

$$f(\tilde{z}_{t+1} | \tilde{z}_t) \sim N(\tilde{z}_0 + \rho\tilde{z}_t + \sigma w_{t+1}, \sigma\sigma'). \quad (5b)$$

However, the approximating and perturbed models are assumed not to be too different statistically. To measure the discrepancy between the two models in terms of the transition from \tilde{z}_t to \tilde{z}_{t+1} , we use the notion of conditional relative entropy defined as the expected log-likelihood ratio of the two models, evaluated with respect to the perturbed model:

$$I(f_a, f)(\tilde{z}) = \int \log \left(\frac{f(\tilde{z}_{t+1} | \tilde{z}_t)}{f_a(\tilde{z}_{t+1} | \tilde{z}_t)} \right) f(\tilde{z}_{t+1} | \tilde{z}_t) d\tilde{z}_{t+1}. \quad (6a)$$

It can be shown that (see Appendix A):

$$I(f_a, f)(\tilde{z}) = I(w_{t+1}) = 0.5w'_{t+1}w_{t+1} \quad (6b)$$

Therefore, the total size of model misspecification in each time period is given by $I(f_a, f)(\tilde{z})$. Model misspecification is typically restricted by requiring that the discounted lifetime sum of (6b) is a constant, e.g.

$$\begin{aligned} 2E_0 \sum_{t=0}^{\infty} \tilde{\beta}^{t+1} I(f_a, f)(\tilde{z}) &\leq \tilde{\eta} \\ E_0 \sum_{t=0}^{\infty} \tilde{\beta}^{t+1} \left[(w_{t+1}^a)^2 + (w_{t+1}^p)^2 + (w_{t+1}^r)^2 \right] &\leq \tilde{\eta}. \end{aligned} \quad (6c)$$

In reality, the potential effects of model misspecification on economic outcomes can depend on the actions of the economic agent. In particular, the potential negative effects of model misspecification can be reduced by purposeful economic activity which seeks to understand the form of the unknown exogenous processes and their potential effects on the economy. For example resources are diverted to learning about the channels through which disease is spread, the existence of terrorist networks and their targets, risks associated with defaults on loans, etc.. Here, in the context of our business cycle model with human capital, we examine the case where knowledge creation can reduce the potential negative effects of model misspecification. In particular, one way in which economic agents might reduce the effects of model misspecification is to spend time informing themselves about the economic

³Note that (5b) allows for misspecifications that occur only as a perturbation to the conditional mean of the innovation to the state, \tilde{z}_{t+1} (and hence x_{t+1} as well) in (5a) and leaves the conditional volatility of the shock, as parametrized by σ , unchanged. We provide further information on this point in Appendix A.

environment. In practice, a part of effort and leisure time is spent by economic agents in such information-acquiring activities. This includes, for example, reading newspapers and books in leisure time, going to meetings and catching up with developments and news during work and education time. These activities do not increase the productivity of future labor *per se* but, instead, help the individual have a better understanding of the economic and sociopolitical system and thus create knowledge that can be useful in mitigating the potential effects of model misspecification. In other words, the extent of model misspecification, $I(f_a, f)(\tilde{z})$, depends on knowledge creation for this purpose. To capture this, we can write the within-period entropy constraint as

$$(w_{t+1}^a)^2 + (w_{t+1}^\varphi)^2 + (w_{t+1}^\tau)^2 = \frac{\eta_0}{h_t^i} - \eta_1 \quad (7a)$$

where knowledge to mitigate model misspecification h_t^i implies an information technology such that

$$h_t^i = q_t^{\eta_2} h_t^{(1-\eta_2)} \quad (7b)$$

and $\eta_0, \eta_1 > 0$ and $0 < \eta_2 < 1$.

In other words, the creation of knowledge about the economic environment, or else information acquisition, is a constant returns to scale production function, which, analogous to human capital creation, has existing human capital, h_t , and effort time (in this case q_t) as inputs. This setup encapsulates the idea that knowledge for model misspecification requires time, which comes at the expense of other uses of time, but also depends positively on existing human capital in a society. The complementarity between q_t and h_t implies that more educated individuals (or societies) will understand more about their economic environment if they put the same effort time as less educated individuals (or societies). The parameter $0 < \eta_2 < 1$ determines the relative importance of information time. Similar to using education time e_t for the creation of human capital, using q_t implies a trade-off: more knowledge is created, at the expense of less time available for leisure, work and education.

It is important to note that in (7a), the agent *chooses* how much to restrict model misspecification, in the sense that by choosing h_t^i he chooses the constraint on model misspecification $I(f_a, f)(\tilde{z})$.⁴ A key assumption made in

⁴Note that (7a) implies that when $\eta_0 = \eta_1 = 0$, the true or perturbed model given by (1a – 1i) is equivalent to the approximating model, (1a – 1f) and (3a – 3c) and thus robust and optimal policies are the same. When $q_t \rightarrow 0$, model misspecification is effectively unrestricted, while when $q_t = \left(\frac{1}{\eta_1 h_t^{(1-\eta_2)}}\right)^{1/\eta_2}$, model misspecification is eliminated. Finally, if $\frac{1}{h_t^i} = 1$ and $\eta_1 = 0$ then the per period entropy constraint would be exogenous.

(7b) is that knowledge used to reduce the effects of model misspecification, h_t^i , depends both on a time input, q_t , and the stock of existing human capital, h_t . As will be discussed below, when the approximating model governs the data, efforts to mitigate model misspecification by acquiring information h_t^i via increases in q_t imply an unproductive use of time. Hence, when fears of model misspecification are unfounded, effort time q_t is not associated with a trade-off, but instead implies a misallocation of resources. On the other hand, efforts to increase information acquisition h_t^i by accumulating human capital, will have a positive productive effect on the economy, even if fears of model misspecification are unfounded.

To summarize, the decision maker believes that the exogenous processes follow (1g) – (1i). However, the approximating processes (3a – 3c) must be used since the w_{t+1} process is unknown. Despite not being able to make optimal decisions using the true model, the agent desires to make good decisions over a set of models satisfying the entropy constraint which has been specified to allow the agent to choose to reduce the effects of model misspecification by using time and existing human capital to this end.⁵ However, the benefits accruing to such activities have to be weighed against the opportunity costs of reducing time spent on other utility enhancing activities. The decision rules or policy functions arising from the solution of this model can then be characterized as being robust to misspecification of the approximating model.

2.2 Robust policies

As discussed above, the conservative decision maker would like to obtain decision rules that are robust against model misspecification, in the sense that they provide good results even in the presence of unfavorable w_{t+1} shocks. To ensure a lower bound of utility in the least favorable environment, the agent makes choices as if the w_{t+1} process follows the worst-case scenario. In particular, it is assumed that w_{t+1} is chosen by a malevolent agent with a view to minimizing the objective of the decision maker. By planning against such a worst-case scenario, the agent, in effect, designs a decision rule that performs well under a set of perturbed models. In other words, the agent uses the malevolent agent as a device to achieve robustness.

⁵Note that while the agent's choices are generally good and not optimal in this framework such good decisions are optimal under a particular realization of the w_{t+1} process, i.e. the worst-case shock. Hence, such a good decision rule is undominated, in the sense that it is optimal for some model. This allows a Bayesian interpretation of the decision rule as it is optimal for a particular belief about the shocks (see Hansen and Sargent (2008, pp. 37, 142 and 158 for more details)).

The problem for the decision maker can be now be stated as:

$$\max_{\{y_t, c_t, i_t, i_t^h, n_t, e_t, q_t, k_{t+1}, h_{t+1}\}_{t=0}^{\infty}} \min_{\{w_{t+1}^a, w_{t+1}^\varphi, w_{t+1}^\tau\}_{t=0}^{\infty}} E_0 \sum_{t=0}^{\infty} \check{\beta}^t \frac{[(c_t)^\mu (l_t)^{(1-\mu)}]^{1-\varsigma}}{1-\varsigma} \quad (8)$$

subject to (1b)-(1i) and (7a)-(7b). We will solve this problem by using dynamic programming methods in a linear quadratic approximation around the non-stochastic steady state of the solution to the non-linear problem (see Appendix B for details).

2.2.1 LQ representation

The first step is to eliminate non-linearities from the constraints of the above problem. To this end, we start by defining a new variable:

$$m_t = \kappa (e_t h_t)^\gamma (i_t^h)^{1-\gamma} \quad (9a)$$

which implies

$$i_t^h = \left(\frac{m_t}{\kappa (e_t h_t)^\gamma} \right)^{\frac{1}{1-\gamma}}. \quad (9b)$$

We next substitute (9b) and (1f) into (1c) for i_t^h and y_t respectively and the resulting expression for c_t into the objective function (8). Then we substitute (7a) and (7b) in (1b) for q_t and the resulting expression for l_t into equation (8). These steps result in the following objective function:

$$\begin{aligned} & \max_{\{u_t\}_{t=0}^{\infty}} \min_{\{w_{t+1}\}_{t=0}^{\infty}} E_0 \sum_{t=0}^{\infty} \check{\beta}^t \left\{ \frac{1}{1-\varsigma} \left[\left(a_t (k_t)^\alpha (h_t)^\zeta (n_t)^{1-\alpha-\zeta} - i_t - \Omega_t^{\frac{1}{1-\gamma}} \right)^\mu \right. \right. \\ & \left. \left. \times \left(1 - n_t - e_t - \left(\frac{\eta_0}{(\tilde{\Omega}_{t+1}) h_t^{(1-\eta_2)}} \right)^{\frac{1}{\eta_2}} \right)^{(1-\mu)} \right]^{1-\varsigma} \right\} \quad (10) \end{aligned}$$

where $a_t = \exp(\tilde{a}_t)$, $\Omega_t = \frac{m_t}{\kappa (e_t h_t)^\gamma}$ and $\tilde{\Omega}_{t+1} = (w_{t+1}^a)^2 + (w_{t+1}^\varphi)^2 + (w_{t+1}^\tau)^2 + \eta_1$. Thus, the representative agent now chooses $u_t = \{i_t, m_t, n_t, e_t\}_{t=0}^{\infty}$ to maximize (10), assuming that the malevolent agent chooses $w_{t+1} = \{w_{t+1}^a, w_{t+1}^\varphi, w_{t+1}^\tau\}_{t=0}^{\infty}$ to minimize (10), subject to:

$$gh_{t+1} = (1 - \delta^h)h_t + m_t + \varphi_t, h_0 > 0 \quad (11a)$$

$$gk_{t+1} = (1 - \delta^k)k_t + i_t + \tau_t, k_0 > 0 \quad (11b)$$

$$\tilde{a}_{t+1} = (1 - \rho^a)\tilde{a} + \rho^a\tilde{a}_t + \sigma^a(\varepsilon_{t+1}^a + w_{t+1}^a) \quad (11c)$$

$$\varphi_{t+1} = \rho^\varphi\varphi_t + \sigma^\varphi(\varepsilon_{t+1}^\varphi + w_{t+1}^\varphi) \quad (11d)$$

$$\tau_{t+1} = \rho^\tau\tau_t + \sigma^\tau(\varepsilon_{t+1}^\tau + w_{t+1}^\tau). \quad (11e)$$

Letting \tilde{h} , n^{\max} and n^{\min} equal the number of states, $(1, k_t, h_t, \tilde{a}_t, \varphi_t, \tau_t)$, the number of controls, (i_t, m_t, n_t, e_t) , for the maximizing agent and the number of controls, $(w_{t+1}^a, w_{t+1}^\varphi, w_{t+1}^\tau)$, for the minimizing agent, we can now write the linear constraints (11a) – (11e) in matrix form as:⁶

$$x_{t+1} = Ax_t + Bu_t + C(\varepsilon_{t+1} + w_{t+1})$$

where

$$\begin{aligned} x_t &= [1 \quad k_t \quad h_t \quad \tilde{a}_t \quad \varphi_t \quad \tau_t]'; & u_t &= [i_t \quad n_t \quad m_t \quad e_t]'; \\ w_{t+1} &= [w_{t+1}^a \quad w_{t+1}^\varphi \quad w_{t+1}^\tau]'; & \varepsilon_{t+1} &= [\varepsilon_{t+1}^a \quad \varepsilon_{t+1}^\varphi \quad \varepsilon_{t+1}^\tau]'; \\ A_{(\tilde{h}x\tilde{h})} &= \begin{bmatrix} 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & \frac{1-\delta^k}{g} & 0 & 0 & 0 & 1 \\ 0 & 0 & \frac{(1-\delta^h)}{g} & 0 & 0 & 0 \\ (1-\rho^a)\tilde{a} & 0 & 0 & \rho^a & 0 & 0 \\ 0 & 0 & 0 & 0 & \rho^\varphi & 0 \\ 0 & 0 & 0 & 0 & 0 & \rho^\tau \end{bmatrix} \\ B_{(\tilde{h}xn^{\max})} &= \begin{bmatrix} 0 & 0 & 0 & 0 \\ \frac{1}{g} & 0 & 0 & 0 \\ 0 & 0 & \frac{1}{g} & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix}; & C_{(\tilde{h}xn^{\min})} &= \begin{bmatrix} 0 & 0 & 0 \\ 0 & 0 & 0 \\ \sigma^a & 0 & 0 \\ 0 & \sigma^\varphi & 0 \\ 0 & 0 & \sigma^\tau \end{bmatrix} \end{aligned}$$

We approximate the objective function in (10) to quadratic form around the non-stochastic steady state of the true model given by (10) – (11e).⁷ We

⁶Note that we have added the coefficient 1 as the first component of the state vector to capture constant terms in the law of motion as in Anderson *et al.* (1996).

⁷The non-stochastic steady state is obtained by setting $\sigma = 0$. Note that while this removes the stochastic effects of w_t in the exogenous productivity processes, w_t still remains in the entropy constraint. Thus the optimal outcomes of the game played between the maximizing and minimizing agents are for steady-state w and q choices to be zero and non-zero respectively implying that eliminating model uncertainty comes at a cost. See Appendix B for details of the steady-state solution of the non-stochastic problem and Appendix C for details on the quadratic approximation.

can write then the linear-quadratic problem as:

$$\max_{\{u_t\}_{t=0}^{\infty}} \min_{\{w_{t+1}\}_{t=0}^{\infty}} E_0 \sum_{t=0}^{\infty} \check{\beta}^t \left\{ \left[\begin{pmatrix} \tilde{u}_t \\ x_t \end{pmatrix}' \begin{pmatrix} R & W' \\ W & Q \end{pmatrix} \begin{pmatrix} \tilde{u}_t \\ x_t \end{pmatrix} \right] \right\} \quad (12a)$$

subject to:

$$x_{t+1} = Ax_t + \tilde{B}\tilde{u}_t + C\epsilon_{t+1}, \quad (12b)$$

where $\tilde{u}_t = (u_t' \ w_{t+1}')'$ and $\tilde{B}_{(\tilde{h} \times \tilde{n})} = [B \ C]$, $\tilde{n} = n^{\max} + n^{\min}$ and R , W and Q are the appropriate partitioned matrices of the respective M matrix that, as explained in Appendix C, defines the second-order approximation of (10). In particular, $R_{(\tilde{n} \times \tilde{n})} = M_{22}$, $W_{(\tilde{h} \times \tilde{n})} = M_{12}$ and $Q_{(\tilde{h} \times \tilde{h})} = M_{11}$.

For the problem to be well defined, the objective function needs to be concave with respect to u_t and convex with respect to w_{t+1} . To check that these conditions are met, we can write the problem in (12a) as:

$$\begin{aligned} \max_{\{u_t\}_{t=0}^{\infty}} \min_{\{w_{t+1}\}_{t=0}^{\infty}} E_0 \sum_{t=0}^{\infty} \check{\beta}^t \{ & u_t' R_u u_t + w_{t+1}' R_w w_{t+1} + 2u_t' R_{wu}' w_{t+1} + x_t' Q x_t + \\ & + 2x_t' W_u u_t + 2x_t' W_w w_{t+1} \} \end{aligned}$$

subject to:

$$x_{t+1} = Ax_t + \tilde{B}\tilde{u}_t + C\epsilon_{t+1}$$

where $R = \begin{pmatrix} R_u & R_{wu}' \\ R_{wu} & R_w \end{pmatrix}$ and R_u is $(n^{\max} \times n^{\max})$, R_w is $(n^{\min} \times n^{\min})$ and R_{wu} is $(n^{\min} \times n^{\max})$; and $W = \begin{pmatrix} W_u & W_w \end{pmatrix}$ and W_u is $(\tilde{h} \times n^{\max})$ and W_w is $(\tilde{h} \times n^{\min})$. We require R_u to be negative definite, Q to be negative semi-definite and R_w to be positive definite (see also Anderson *et al.* (1996) and Hansen and Sargent (2008) for assumptions regarding the coefficient matrices for linear quadratic problems).

2.2.2 Model Solution

We make use of a type of *certainty equivalence*, which applies to the class of linear quadratic games relevant here (see e.g. Hansen and Sargent, 2008, ch. 2). In particular, the decision rules for u_t and w_{t+1} in the *maxmin* game (12a)-(12b) are the same in a particular non-stochastic version of the problem, i.e. where $\epsilon_{t+1} = 0$. Therefore, we focus on the problem

$$\max_{\{u_t\}_{t=0}^{\infty}} \min_{\{w_{t+1}\}_{t=0}^{\infty}} \sum_{t=0}^{\infty} \check{\beta}^t \left\{ \left[\begin{pmatrix} \tilde{u}_t \\ x_t \end{pmatrix}' \begin{pmatrix} R & W' \\ W & Q \end{pmatrix} \begin{pmatrix} \tilde{u}_t \\ x_t \end{pmatrix} \right] \right\} \quad (13a)$$

subject to:

$$x_{t+1} = Ax_t + \tilde{B}\tilde{u}_t. \quad (13b)$$

We will solve for a Markov-perfect equilibrium (see Hansen and Sargent, 2008, Definition 7.4.1).⁸ Since this is a dynamic zero-sum game the order of optimization does not affect the solution. In other words the *Bellman-Isaacs* condition holds, see e.g. Hansen and Sargent (2008, chapter 7); see also Basar and Olsder (1999, ch. 5 and 6) for dynamic, zero-sum open-loop and feedback Nash equilibria). Therefore, to solve the problem in (13a)-(13b), we can stack the first order conditions of the maximizing and the minimizing agents.

To implement this solution, we follow Hansen and Sargent (2008, ch. 2) and note that, in (13a)-(13b), the first order conditions of the maximizing agent with respect to u_t and of the minimizing agent with respect to w_{t+1} , are the same with the first order conditions of an ordinary (i.e. non-robust) optimal linear regulator (OLR) who chooses \tilde{u}_t . Hence, we write the *extremization*⁹ problem in (13a)-(13b) as:

$$\underset{\{\tilde{u}_t\}_{t=0}^{\infty}}{ext} \sum_{t=0}^{\infty} \tilde{\beta}^t \left\{ \left[\begin{pmatrix} \tilde{u}_t \\ x_t \end{pmatrix}' \begin{pmatrix} R & W' \\ W & Q \end{pmatrix} \begin{pmatrix} \tilde{u}_t \\ x_t \end{pmatrix} \right] \right\} \quad (14a)$$

subject to:

$$x_{t+1} = Ax_t + \tilde{B}\tilde{u}_t. \quad (14b)$$

To solve the above OLR problem, we next eliminate the cross-products between states and controls in the objective function (see e.g. Hansen and Sargent, 2008, ch. 7). If we define:

$$\begin{aligned} \tilde{Q} &= Q - WR^{-1}W', \\ \tilde{A} &= A - \tilde{B}R^{-1}W', \\ \tilde{v}_t &= \tilde{u}_t + R^{-1}W'x_t. \end{aligned}$$

the problem can then be written as:

$$\max_{\{\tilde{v}_t\}_{t=0}^{\infty}} E_0 \sum_{t=0}^{\infty} \beta^t \begin{pmatrix} \tilde{v}_t \\ x_t \end{pmatrix}' \begin{pmatrix} R & 0 \\ 0 & \tilde{Q} \end{pmatrix} \begin{pmatrix} \tilde{v}_t \\ x_t \end{pmatrix}, \quad (15a)$$

subject to:

$$x_{t+1} = \tilde{A}x_t + \tilde{B}\tilde{v}_t. \quad (15b)$$

⁸The definition of the *Markov Perfect Equilibrium* discussed here describes equilibria that are, under certainty, consistent with memoryless feedback equilibria (see e.g. Basar and Olsder, 1999, ch. 6).

⁹Following Whittle (1990), extremization denotes joint maximization and minimization.

The problem in (15a)-(15b) can be solved as a standard OLR problem (see e.g. Anderson *et al.*, 1996) to give the policy functions:

$$\tilde{v}_t = -\tilde{F}x_t$$

and the state evolution equation:

$$x_{t+1} = (\tilde{A} - \tilde{B}\tilde{F})x_t \quad (16a)$$

where

$$\tilde{F} = \check{\beta}(R + \check{\beta}\tilde{B}'P^*\tilde{B})^{-1}\tilde{B}'P^*\tilde{A}$$

and P^* is obtained by iterating on:

$$P^* = \tilde{Q} + \check{\beta}\tilde{A}'P^*\tilde{A} - \check{\beta}^2\tilde{A}'P^*\tilde{B}(R + \check{\beta}\tilde{B}'P^*\tilde{B})^{-1}\tilde{B}'P^*\tilde{A}.$$

The stability of the solution is guaranteed if the eigenvalues of $(\tilde{A} - \tilde{B}\tilde{F})$ are all less than one in absolute value.¹⁰ Working backwards, we can obtain the solution for \tilde{u}_t as

$$\tilde{u}_t = -\left(\tilde{F} + R^{-1}W'\right)x_t$$

and then, by appropriately partitioning of the matrix $(\tilde{F} + R^{-1}W')$ we get the policy function for u_t and the worst case scenario for w_{t+1} :

$$u_t = -Fx_t \quad (16b)$$

$$w_{t+1} = Kx_t. \quad (16c)$$

The value of the game or the welfare under the worst case scenario, is given by:

$$v(x_0) = x_0'P^*x_0 + d \quad (17)$$

where $d = \check{\beta}(1 - \check{\beta})^{-1}\text{tr}(P^*CC')$; and P^* is the value of P at convergence.

3 The robustness premium and precautionary behavior

The importance of the costs associated with robust decision making has been acknowledged in the literature. Hansen and Sargent (2010, p. 13), for example, warn against recommendations where "conservative behavior is so

¹⁰Of course, except for the one corresponding to the constant term, which will be exactly unity.

cautious that it makes the robust decision rule look silly". The more conservative the decision maker, the higher the robustness premium he is expected to pay if his fears of model misspecification are unfounded. However, at the same time, the more he is protected against model misspecification.

In the model uncertainty literature, particular interest has also been paid to whether robust policy making leads to more precautionary behavior. This behavior in a robust context is different from the conventional form of precautionary savings which results from a convex marginal utility function and a positive variance for the productivity process (see e.g. Leland (1968) and Sandmo (1970)). In contrast, the precautionary behavior emerging from robust decision making is a result of fear of misspecifying the conditional means of the model's stochastic state variables and does not require a convex marginal utility of consumption. However, similar to the conventional case, the precautionary behavior evolving from robust decision making is increasing in the variance of the stochastic process.¹¹ To assess the cost and precautionary implications of robust policy rules below we follow Hansen and Sargent (2008) and Dennis (2010) and compare robust to optimal model outcomes when the approximating model governs the data.

3.1 Calibrating the model parameters

We next calibrate the model parameters using data and empirical evidence from econometric studies. Where neither is available, we calibrate the remaining parameters so that the solution of the model in the long-run reproduces the stylized facts regarding the U.S. data. In particular, we aim to obtain reasonable values of the so-called "great ratios" such as the components of expenditure as shares of output, physical and human capital to output ratios, etc.

For the human capital model we examine here, we use standard values in the literature for calibrating the model using U.S. data for the post-war period (see e.g. King and Rebelo (1999)). These are reported in Table 1. We set the value of $(1 - \alpha)$ equal to labor's share in income using compensation of employees data. The value for β is set such that $1/\beta$ is equal to 1 plus the *ex-post* real interest rate. The values for δ^h and δ^k are from Jorgenson and Fraumeni (1989). The gross growth rate g is set to the average per labour input growth rate. Following Kydland (1995, ch. 5, p. 134), we set μ to

¹¹These findings have also been confirmed in other applications using robust control methods. Model uncertainty has also been linked to precaution in monetary policy design (see e.g. Brainard, 1967). However, in several recent applications in monetary policy, an increased preference for robustness can imply more aggressive policy behavior (see e.g. the papers and the discussion in Leitmo and Soderstrom (2008)).

be equal to the average hours of work versus leisure time, which is obtained using data on hours worked.¹² We use a value for $1/\zeta$ that is common in the DSGE literature (i.e. $\zeta = 2$), as micro evidence suggests (although not conclusively) that $(1/\zeta)$ is less than one.

We calibrate those parameters we cannot retrieve from the data or previous empirical studies, in particular the exponents in the production function for human capital, i.e. (γ, ζ) and the constants in goods and human capital production functions, i.e. \tilde{a} and κ respectively, so that the solution of the model in the long-run reproduces the stylized facts regarding the U.S. data.

Table 1: Model calibration (benchmark)

definition	parameter	value
productivity of capital	α	0.330
constant term in \tilde{a}_t	\tilde{a}	0.820
time discount factor	β	0.970
constant term in human capital production	κ	0.900
productivity of human capital investment	γ	0.700
depreciation rate on physical capital	δ^k	0.049
depreciation rate on human capital	δ^h	0.018
productivity of human capital	ζ	0.300
labor augmenting technical progress	g	1.025
weight on consumption in utility	μ	0.350
coefficient of relative risk aversion	ς	2.000
elasticity of h_t^i with respect to q_t	η_2	0.400
constant term in entropy constraint	η_1	0.275
standard deviation of $\tilde{a}_t, \varphi_t, \tau_t$	$\sigma^a, \sigma^\varphi, \sigma^\tau$	0.010
persistence of $\tilde{a}_t, \varphi_t, \tau_t$	$\rho^a, \rho^\varphi, \rho^\tau$	0.990

Regarding the knowledge for misspecification sector, we normalize the value of η_0 to unity since increases in η_0 have the same effect as decreases in η_1 . We then calibrate η_1 and η_2 so that they imply a time for misspecification, q_t , that would be considered reasonable and at the same time make the approximating and the distorted model difficult to distinguish statistically. Given the lack of data on q_t , it is obviously hard to identify a target value for q_t . We consider as a benchmark a calibration that implies values for q_t around 1%, but as we discuss later we also examine calibrations that imply higher values for q_t . Moreover, following Hansen and Sargent (2008, 2010),

¹²To obtain this we divide total hours worked by total hours available for work or leisure, following e.g. Ho and Jorgenson (2000). For example, they assume that there are 14 hours available for work or leisure on a daily basis with the remaining 10 hours accounted for by physiological needs.

we restrict our interest to circumstances in which the decision maker, faced with a finite amount of data, finds it difficult to statistically distinguish the approximating from the perturbed model. We achieve this by applying their detection error probabilities approach to measure the extent of or fear of model misspecification. In particular, if a likelihood ratio test is used to distinguish the approximating from the worst-case perturbed model, there would be a detection error probability, because the two models can differ due to genuine randomness. Hansen and Sargent suggest that it is reasonable to design robust policies that would allow for at least a 10% probability of making such an error.¹³ We thus calibrate the parameters η_1 and η_2 so that robust policies do not imply detection error probabilities that are lower than 10%.¹⁴

3.2 Results

To quantitatively assess the cost and precautionary implications of robust decision making, we start by simulating the approximating model under optimal (see Table 2) and robust policy rules (see Table 3). We further explore these by changing the parameters in the entropy constraint of the robust model. In particular, we first increase, η_2 i.e. the relative productivity of time, holding the size of the entropy constraint constant (see Table 4). Second, we decrease η_1 holding the other parameters constant, so that the entropy constraint changes (see Table 5).

In the Tables below we present (i) the model averages for the main economic variables;¹⁵ (ii) the extent of feared model misspecification, $\eta = \frac{\eta_0}{\bar{q}\eta_2\bar{h}^{(1-\eta_2)}} - \eta_1$; (iii) welfare under the optimal policy rule, $v^o(x_0)$; (iv) welfare under the robust rule, $v^r(x_0)$; and (v) the robustness premium, ϕ calculated as the percent difference between welfare in the robust versus the optimal policies (see e.g. Dennis (2010)):

$$\phi = \frac{v^r(x_0) - v^o(x_0)}{v^o(x_0)}.$$

¹³In the context on an exogenous entropy constraint, the smaller (larger) the model detection error probability, the larger (smaller) the difference between the two models and hence the greater (less) the extent of feared of model misspecification. We return to a modified interpretation of the model detection error probability, under an endogenous entropy constraint, below.

¹⁴Following the literature, these probabilities are calculated using 100,000 simulations.

¹⁵The time horizon for the simulations is 350 years. Later periods do not affect outcomes, because of discounting.

3.2.1 Optimal versus robust policies

To simulate the approximating model under optimal policy rules reported in Table 2 we set $\eta_0 = \eta_1 = 0$ in equation (7a).¹⁶ This implies that the decision maker designs optimal policy rules when the approximating model is the correct one. The results in Table 2 provide the benchmark we require to assess the cost and precautionary implications of robust decision making in the remaining Tables.

Table 2: Optimal policies

\bar{y}	2.581
\bar{c}	1.761
\bar{i}	0.596
\bar{i}^h	0.224
\bar{k}	8.059
\bar{h}	33.454
\bar{e}	0.110
\bar{n}	0.201
$v^o(x_0)$	-34.854

To simulate the approximating model under robust policy rules reported in Table 3, we use the same initial conditions as in Table 2 and set $K = 0$ in equations (16a – 16c). The latter implies that the approximating model is the correct one and governs the evolution of the states but the agent solves for robust decision rules (see e.g. Hansen and Sargent, 2008, p. 39). In other words, the decision maker’s fears of model misspecification are unfounded.

Since the simulations in Table 3 are initialized as in Table 2, there is a transition period before the stochastic steady-state is reached. The figures in the first column of Table 3 are calculated including this transition path. Thus they refer to the lifetime of the agent. In contrast, the transition path is excluded in the second column so that the model averages and welfare refer exclusively to the stochastic steady-state.

¹⁶Note that the initial values for the simulation reported in Table 2 correspond to the non-stochastic steady-state of this model.

Table 3: Robust policies

	lifetime	steady-state
\bar{y}	4.285	4.518
\bar{c}	2.759	2.965
\bar{i}	1.017	1.069
\bar{i}^h	0.509	0.484
\bar{k}	13.501	14.438
\bar{h}	104.699	113.534
\bar{e}	0.139	0.134
\bar{n}	0.197	0.197
\bar{q}	0.012	0.010
η	0.133	0.132
$v^r(x_0)$	-35.086	-29.770
ϕ	0.007	-0.146

Comparing $v^r(x_0)$ in Table 3 to $v^o(x_0)$ in Table 2 using ϕ , it can be seen that robust policies imply a small welfare cost for the agent, if the fears of model misspecification are unfounded. However, as indicated by ϕ in the second column, this cost is temporary, since welfare is 14.6% higher than $v^o(x_0)$. In other words, an economy that chooses to pay the robustness premium of 0.7% is guaranteed to have a higher welfare in the long-run.

Examining the effects of precautionary behavior under the robust decision rule helps to explain the welfare results described above. Compared to the averages in Table 2, robust decisions imply precautionary increases in investment in physical, \bar{i} , and human capital, \bar{i}^h , which result in increases in output, \bar{y} , and consumption, \bar{c} . These effects are more pronounced in the long-run. Robust policies also imply an increase in education time, \bar{e} , and time to create knowledge to reduce model uncertainty, \bar{q} . Hence, leisure time, $(1 - \bar{e} - \bar{n} - \bar{q})$, falls. Despite this, the net effect on welfare, due to higher consumption, is positive in the long-run.

As is standard in the robust control literature, precautionary behavior on the part of the robust decision maker implies that he is increasing his stocks of physical and human capital as buffers against negative effects of model misspecification. In the long-run, these provide a benefit, as they increase productivity (and consumption). In the short-run, however, these increases in investment necessitate a sharp decrease in consumption. This is captured by an overall increase in the volatility of consumption, when the transition path is included in evaluating outcomes and welfare. For the model in Table 3, the volatility of consumption including the transition path is 0.49, whereas

for the model under certainty in Table 2, volatility is 0.17.¹⁷ This increase in volatility hurts welfare and is predominantly associated with the transition path, as the volatility of consumption in the steady-state returns to 0.20.

To summarize, there are benefits and costs from following robust policies when the approximating model governs the data. The benefits take the form of increased productivity due to precautionary accumulation of physical and human capital. This effect is intensified in our model because human capital can also be used to reduce the effects of model uncertainty. Hence, the endogenous entropy constraint provides an additional motive for precautionary accumulation of human capital. The costs include a misallocation of time resources to reduce the effects of model misspecification and a temporary cost of increased volatility of consumption which results from the need to finance the increase in investment in physical and human capital along the transition path.

3.2.2 Increasing the relative productivity of time

The welfare effects of robust decision making crucially depend on the technology used to reduce model uncertainty. We demonstrate this below, by considering the effects of increasing the relative efficiency of time, η_2 . In our experiment, we increase the size of η_2 to 0.6 and 0.8 and to hold the detection error probability constant at $p = 10\%$ we also increase η_1 to 0.625 and 1.43 respectively.

When η_2 increases from 0.4 to 0.6, both \bar{q} and \bar{h} increase relative to their values in Table 3 leading to robustness premia, ϕ , of 7.68% and -12.9% for the lifetime and steady-state respectively.¹⁸ In contrast, when η_2 increases to 0.8, while \bar{q} increases, \bar{h} decreases relative to their values in Table 3 leading to robustness premia of 27.8% and 8.4% for the lifetime and steady-state respectively.¹⁹ In other words, the misallocation of resources in this case is dramatic and outweighs the benefits of precautionary increases in output and consumption displayed for lower values of η_2 .

¹⁷The volatility for leisure does not change significantly between the different policies, increasing from 0.009 for the solution in Table 2, to 0.014 for the solution in Table 3.

¹⁸In this case, the volatility of consumption including the transition path is 0.65, returning to 0.29 when the stochastic steady state is considered only.

¹⁹In this case, the volatility of consumption including the transition path is 0.34, returning to 0.22 when the stochastic steady state is considered only.

Table 4: Robust policies

	$\eta_2 = 0.6$		$\eta_2 = 0.8$	
	lifetime	steady-state	lifetime	steady-state
\bar{y}	4.668	5.025	3.420	3.708
\bar{c}	2.927	3.242	2.157	2.373
\bar{i}	1.107	1.187	0.810	0.877
\bar{i}^h	0.634	0.597	0.454	0.458
\bar{k}	14.656	16.029	10.801	11.847
\bar{h}	138.569	156.588	82.720	94.728
\bar{e}	0.146	0.140	0.130	0.127
\bar{n}	0.184	0.186	0.158	0.160
\bar{q}	0.068	0.061	0.198	0.190
η	0.126	0.126	0.123	0.123
$v^r(x_0)$	-37.515	-30.360	-44.527	-37.770
ϕ	0.076	-0.129	0.278	0.084

To summarize, increases in η_2 or the return to q_t relative to h_t in the knowledge production function, leads to an increase in the robustness premium. In the extreme, this increase can even reverse the long-run gains of robust policies, when fear of model misspecification is unfounded.

3.2.3 Increasing the fear of model misspecification

Ceteris paribus, a fall in η_1 in the entropy constraint given by (7a) implies *ex ante* that it is more difficult to mitigate the effects of model misspecification. In other words, this reflects that the economic agent fears model misspecification more and will choose more conservative policies. In Table 5, we present the effects of a fall in η_1 for low and for high values of η_2 , in particular for η_2 equal to 0.4 and 0.8 as in Tables 3 and 4 respectively. In both cases, we examine a fall in η_1 that implies a detection error probability of 20%.²⁰ Finally, we again simulate the solution in the approximating model under the robust decision rule, using the same initial conditions as in Table 2.

²⁰As pointed out in footnote 13, under an exogenous entropy constraint, an increase in the fear of model misspecification would imply, by construction, a fall in the detection error probability, p . In our model, a more conservative robust decision maker chooses to mitigate model misspecification using q_t and h_t . Hence, when the fear of model misspecification, $\eta = \frac{\eta_0}{\bar{q}^{\eta_2} \bar{h}^{(1-\eta_2)}} - \eta_1$, rises, due to a fall in η_1 , we observe a fall in η *ex post* due to offsetting changes in \bar{q} and \bar{h} and hence a corresponding increase in p . In this sense, higher detection error probabilities indicate more and not less conservative behavior.

Table 5: Robust policies

	$\eta_1 = 0.22, \eta_2 = 0.4$		$\eta_1 = 1.21, \eta_2 = 0.8$	
	lifetime	steady-state	lifetime	steady-state
\bar{y}	4.589	4.854	3.387	3.703
\bar{c}	2.923	3.174	2.104	2.344
\bar{i}	1.078	1.137	0.794	0.868
\bar{i}^h	0.587	0.543	0.489	0.491
\bar{k}	14.285	15.352	10.595	11.716
\bar{h}	125.611	136.613	88.516	102.738
\bar{e}	0.144	0.138	0.130	0.127
\bar{n}	0.194	0.195	0.148	0.150
\bar{q}	0.023	0.019	0.249	0.238
η	0.054	0.054	0.056	0.056
$v^r(x_0)$	-35.578	-29.556	-47.994	-39.943
ϕ	0.021	-0.152	0.377	0.146

For both cases considered, more conservative policies imply more precautionary behavior. However, the way this behavior is reflected differs. In particular, for $\eta_2 = 0.4$, the robust decision maker increases his buffer stocks of physical and human capital relative those reported in Table 3, which implies increases in output and consumption. In this case, he also increases q_t , trying to restrict model misspecification. This increase in q_t represents an increase in the misallocation of resources when the approximating model governs the data. In addition, the volatility of consumption increases, so that the costs in the transition period increase even more.²¹ Comparing ϕ in Tables 3 and 5, the outcome of the new trade-off under more conservative policies is that the robustness premium increases in the lifetime of the agent, from 0.7% to 2.1%, but the gains from robust decision making increase in the long-run, from 14.6% to 15.2%.

However, when $\eta_2 = 0.8$, so that the mitigation of the effects of model misspecification depends predominantly on q_t , precautionary investment in physical capital disappears. The incentives to engage in activities to mitigate the potential effects of model misspecification are so high that the robust decision maker channels his resources to this effort (e.g. \bar{q} rises to nearly 25% in Table 5 from 19% in Table 4), so that investment, output and consumption are reduced when robust policies are more conservative. Comparing ϕ in Tables 4 and 5, the outcome of the new trade-off under more conservative policies is that the robustness premium increases both in the lifetime of the agent, from 27.8% to 37.7% and in the long-run, from 8.4% to 14.6%.

²¹In particular, the volatility of consumption is 0.62 and 0.27 in this case, for the lifetime and long-run results respectively.

To summarize, more conservative policies, resulting from an increased fear of model misspecification, imply more precautionary behavior. However, the form in which this behavior is realized again depends critically on the relative productivity of time and human capital in the knowledge production function.

4 Conclusions

In this paper we have examined the importance of the premia associated with robust decision making at the aggregate level using a two-sector real business cycle model. To achieve this we introduced an additional form of precautionary behavior by allowing the robust decision maker to reduce the effects of model misspecification using time and existing human capital. To obtain the robustness premium, we compared outcomes obtained from making robust decisions, to those from making optimal choices when the approximating model governed the data. In this environment the decision maker had the incentive to pay a robustness premium to protect himself from the uncertainty associated with not knowing the true model.

We found that despite more conservative robust decisions leading to higher robustness premia, when the mitigation of model misspecification depends on a low productivity of time relative to human capital, robust choices increase welfare in the long-run. Thus, unfounded conservative policies imply an opportunity cost for this generation but they will be beneficial for future generations. In stark contrast, if the technology used to mitigate the effects of model misspecification relies on a low productivity of human capital relative to time then the misallocation of resources under the approximating model is so strong that more conservative robust policies imply welfare costs which remain even in the long-run. Hence, unfounded conservatism will hurt not only the current, but also future generations. Finally, depending on the technology used to reduce model uncertainty, we found that while increasing the fear of model misspecification leads to a net increase in precautionary behavior, investment and output can fall. Of course the relative productivity of human capital and time used to protect against model misspecification is an open issue which remains to be resolved empirically.

References

- [1] Anderson, E., Hansen, L., McGrattan, E. and T. Sargent (1996): “Mechanics of forming and estimating dynamic linear economies”, in H.

- Amman, D. Kendrick and J. Rust (eds.), Handbook of Computational Economics, vol. 1., Amsterdam: North Holland.
- [2] Barillas, F., Hansen, L. and T. Sargent (2009); “Doubts or variability?”, Journal of Economic Theory 144 (6), 2388-2418.
- [3] Basar, T. and G. J. Older (1999): Dynamic Noncooperative Game Theory, SIAM (2nd ed).
- [4] Brainard, W. (1967): “Uncertainty and the effectiveness of policy”, American Economic Review, 57, 411-425.
- [5] Dennis, R. (2010): “How robustness can lower the cost of discretion”, Journal of Monetary Economics, doi:10.1016/j.moneco.2010.06.003.
- [6] DeJong, D. and B. Ingram (2001): “The cyclical behavior of skill acquisition”, Review of Economic Dynamics, 4, 536-561.
- [7] Hansen, L. and T. Sargent (2001): “Robust control and model uncertainty”, American Economic Review, 91, 60-66.
- [8] Hansen, L. and T. Sargent (2008): Robustness, Princeton University Press, Princeton, NJ.
- [9] Hansen, L. and T. Sargent (2010): “Wanting robustness in macroeconomics”, *mimeo*, forthcoming in the Handbook of Monetary Economics, edited by B. Friedman and M. Woodford.
- [10] Hansen, L., Sargent, T. and T. Tallarini (1999): “Robust permanent income and pricing,” Review of Economic Studies, 66, 873-907.
- [11] Hansen, L., Sargent, T., Turmuhambetova, G. and N. Williams (2006): “Robust control, min-max expected utility and model misspecification”, Journal of Economic Theory, 128, 45-90.
- [12] Ho, M. and D. Jorgenson (2000): “Educational policies to stimulate growth”, in Using Dynamic General Equilibrium Models for Policy Analysis, edited by G. W. Harrison, S. E. H. Jensen, L. H. Pedersen, and T. H. Rutherford. North-Holland, Amsterdam.
- [13] Jorgenson, D. and B. Fraumeni (1989): “The accumulation of human and nonhuman capital, 1948-1984, in R.E. Lipsey and H.S. Tice (eds.), The Measurement of Saving, Investment and Wealth, Studies in Income and Wealth,” Vol. 52, Chicago, University of Chicago Press.

- [14] King, R. and S. Rebelo (1999): “Resuscitating real business cycles,” in Handbook of Macroeconomics, vol. 1B, edited by J. Taylor and M. Woodford, North Holland.
- [15] Kydland F. (1995): “Business cycles and aggregate labor market fluctuations”, in Frontiers of Business Cycle Research, edited by T. Cooley, Princeton University Press.
- [16] Leitemo, K. and U. Soderstrom (2008): “Robust monetary policy in a small open economy”, Journal of Economic Dynamics and Control, 32, 3218-3252.
- [17] Leland, H. (1968): “Saving and Uncertainty: The Precautionary demand for saving,” Quarterly Journal of Economics, 82, 465–473.
- [18] Ljungqvist, L. and T. Sargent (2000): Recursive Macroeconomic Theory, The MIT Press, Cambridge, MA.
- [19] Mankiw, G., D. Romer and D. Weil (1992): “A contribution to the empirics of economic growth”, Quarterly Journal of Economics, 107, 407-437.
- [20] Perli, R. and P. Sakellaris (1998): “Human capital formation and business cycle persistence,” Journal of Monetary Economics, 42, 67-92.
- [21] Sandmo, A. (1970): “The effect of uncertainty on saving decisions,” The Review of Economic Studies, 37, 353-360.
- [22] Whittle P. (1990): “Risk Sensitive Optimal Control”, New York, Wiley.

5 Appendices

5.1 Appendix A

Using the probability densities for the approximating and the distorted model²²

$$f_a(\tilde{z}_{t+1} | \tilde{z}_t) \sim N(\tilde{z}_0 + \rho\tilde{z}_t, \sigma\sigma')$$

²²Note that $f(\tilde{z}_{t+1} | \tilde{z}_t)$ allows for misspecification that occurs only as perturbations to the conditional means of the innovations to the states, \tilde{z}_{t+1} in $f_a(\tilde{z}_{t+1} | \tilde{z}_t)$ and leaves the conditional volatilities of the shocks, as parametrized by $\sigma\sigma'$, unchanged. As Hansen and Sargent (2008) show (see ch. 2.6, 3 and 7), as long as we stay within the LQ framework with Gaussian distributions for the approximating model $f_a(\tilde{z}_{t+1} | \tilde{z}_t)$, allowing for a more general class of misspecification does not change important results. In particular, even when the minimizing agent is allowed to distort the variance of w_{t+1} , it is still optimal to choose a normal density with the same mean. However, the minimizing agent would choose to increase the variance of the distribution if given the chance. This would only

$$f(\tilde{z}_{t+1} | \tilde{z}_t) \sim N(\tilde{z}_0 + \rho\tilde{z}_t + \sigma w_{t+1}, \sigma\sigma')$$

we can calculate the log-likelihood ratio in

$$I(f_a, f)(\tilde{z}) = \int \log \left(\frac{f(\tilde{z}_{t+1} | \tilde{z}_t)}{f_a(\tilde{z}_{t+1} | \tilde{z}_t)} \right) f(\tilde{z}_{t+1} | \tilde{z}_t) d\tilde{z}_{t+1}$$

as:

$$\begin{aligned} \log \left(\frac{f(\tilde{z}_{t+1} | \tilde{z}_t)}{f_a(\tilde{z}_{t+1} | \tilde{z}_t)} \right) &= \log f(\tilde{z}_{t+1} | \tilde{z}_t) - \log f_a(\tilde{z}_{t+1} | \tilde{z}_t) \\ &= \frac{1}{2} w'_{t+1} w_{t+1} + \frac{1}{2} w'_{t+1} \sigma^{-1} (\tilde{z}_{t+1} - \tilde{z}_0 - \rho\tilde{z}_t - \sigma w_{t+1}) + \\ &\quad + \frac{1}{2} (\tilde{z}_{t+1} - \tilde{z}_0 - \rho\tilde{z}_t - \sigma w_{t+1})' (\sigma')^{-1} w_{t+1}. \end{aligned}$$

Hence, the conditional relative entropy in $I(f_a, f)(\tilde{z})$ is given as:

$$\begin{aligned} I(f_a, f)(\tilde{z}) &= \int \left[\frac{1}{2} w'_{t+1} w_{t+1} + \frac{1}{2} w'_{t+1} \sigma^{-1} (\tilde{z}_{t+1} - \tilde{z}_0 - \rho\tilde{z}_t - \sigma w_{t+1}) + \right. \\ &\quad \left. + \frac{1}{2} (\tilde{z}_{t+1} - \tilde{z}_0 - \rho\tilde{z}_t - \sigma w_{t+1})' (\sigma')^{-1} w_{t+1} \right] f(\tilde{z}_{t+1} | \tilde{z}_t) d\tilde{z}_{t+1} \end{aligned}$$

where the expectation is evaluated with respect to the true, or distorted model, which implies that:

$$I(f_a, f)(\tilde{z}) = \frac{1}{2} w'_{t+1} w_{t+1}.$$

5.2 Appendix B

The problem of the robust decision maker when $\sigma = 0$, is to choose $\{i_t, m_t, n_t, e_t, k_{t+1}, h_{t+1}\}_{t=0}^{\infty}$ to maximize

$$\begin{aligned} &\sum_{t=0}^{\infty} \check{\beta}^t \left\{ \frac{1}{1-\zeta} \left[\left(a_t (k_t)^\alpha (h_t)^\zeta (n_t)^{1-\alpha-\zeta} - i_t - \Omega_t^{\frac{1}{1-\gamma}} \right)^\mu \times \right. \right. \\ &\quad \left. \left. \times \left(1 - n_t - e_t - \left(\frac{1}{(\tilde{\Omega}_{t+1}) h_t^{(1-\eta_2)}} \right)^{\frac{1}{\eta_2}} \right)^{(1-\mu)} \right]^{1-\zeta} \right\} \end{aligned}$$

hurt the maximizing agent through the constant term, d , in the value function, $v(x_0)$ (see eq. 17), but would otherwise not affect choices. In other words, if we allowed for more general distortions along these lines there would be a fall in the value of the game reflected by the smaller constant term in the value function, but the policy functions chosen by the maximizing and minimizing agents would be the same.

where $a_t = \exp(\tilde{a}_t)$, $\Omega_t = \frac{m_t}{\kappa(e_t h_t)^\gamma}$, and $\tilde{\Omega}_{t+1} = (w_{t+1}^a)^2 + (w_{t+1}^\varphi)^2 + (w_{t+1}^\tau)^2 + \eta_1$ subject to:

$$\begin{aligned} g h_{t+1} &= (1 - \delta^h) h_t + m_t + \varphi_t, h_0 > 0 \\ g k_{t+1} &= (1 - \delta^k) k_t + i_t + \tau_t, k_0 > 0 \\ \tilde{a}_{t+1} &= (1 - \rho^a) \tilde{a} + \rho^a \tilde{a}_t \\ \varphi_{t+1} &= (1 - \rho^\varphi) \varphi + \rho^\varphi \varphi_t \\ \tau_{t+1} &= (1 - \rho^\tau) \tau + \rho^\tau \tau_t \end{aligned}$$

and the paths $\{w_{t+1}^a, w_{t+1}^\varphi, w_{t+1}^\tau\}_{t=0}^\infty$ are chosen by a malevolent agent with the objective to minimize the agent's lifetime utility, subject to the above constraints.

To solve this problem, we use the *Bellman-Isaacs* condition that holds for this game, which indicates that the first order conditions for *extremizing*, i.e. for simultaneously maximizing and minimizing the objective with respect to $i_t, m_t, n_t, e_t, k_{t+1}, h_{t+1}$ and $w_{t+1}^a, w_{t+1}^\varphi, w_{t+1}^\tau$ respectively, match those of an ordinary optimization problem with joint control processes $\{i_t, m_t, n_t, e_t, k_{t+1}, h_{t+1}, w_{t+1}^a, w_{t+1}^\varphi, w_{t+1}^\tau\}$. Hence, we can stack the first order conditions of the zero-sum two-player dynamic game above and solve them simultaneously (see e.g. Hansen and Sargent (2008, ch. 7)). The steady-state solution of this set of non-linear equations gives the results in Table C for the calibration reported in Table 1.

Table C: Steady-state

c/y	0.662
i/y	0.231
i^h/y	0.107
k/y	3.122
h/y	24.324
y	4.173
c	2.762
i	0.964
i^h	0.447
k	13.028
h	101.506
e	0.132
q	0.025
n	0.195
w^a	0.000
w^φ	0.000
w^τ	0.000

5.3 Appendix C

We express the objective function (10) in implicit form as $r(z_t)$ where $z_t = [1 \ k_t \ h_t \ \tilde{a}_t, \varphi_t, \tau_t \ i_t \ m_t \ n_t \ e_t, w_{t+1}^a, w_{t+1}^\varphi, w_{t+1}^\tau]'$. Our aim is then to replace the non-linear function $r(z_t)$ by a quadratic one of the form $z_t' M z_t$. To this end we approximate $r(z_t)$ around a fixed point z , which is the steady state solution obtained above in Appendix B, using the first two terms of a Taylor series expansion, i.e.

$$r(z_t) \approx \hat{r}(z_t) = r(z) + (z_t - z)' \frac{\partial r}{\partial z_t} \Big|_{ss} + \frac{1}{2} (z_t - z)' \frac{\partial^2 r}{\partial z_t \partial z_t'} \Big|_{ss} (z_t - z) \quad (D1)$$

where, the partial derivatives in (D1) are evaluated at the steady-state z (see also Anderson *et al.* (1996) for more details).

Let j be a $((\tilde{h} + \tilde{n}) \times 1)$, where $\tilde{n} = n^{\max} + n^{\min}$, zero vector, except for a 1 in the row that corresponds to the entry of unity in the z_t vector, so that $j' z_t = 1$. We can now rewrite (D1) as:²³

$$\begin{aligned} \hat{r}(z_t) &= z_t' M z_t, \text{ where} \\ M &= j \left[r(z) - \left(\frac{\partial r}{\partial z_t} \Big|_{ss} \right)' z + \frac{1}{2} z' \frac{\partial^2 r}{\partial z_t \partial z_t'} \Big|_{ss} z \right] j' + \\ &\quad \frac{1}{2} \left[\frac{\partial r}{\partial z_t} \Big|_{ss} j' - j z' \frac{\partial^2 r}{\partial z_t \partial z_t'} \Big|_{ss} - \frac{\partial^2 r}{\partial z_t \partial z_t'} \Big|_{ss} z z' + j \left(\frac{\partial r}{\partial z_t} \Big|_{ss} \right)' \right] \\ &\quad + \frac{1}{2} \frac{\partial^2 r}{\partial z_t \partial z_t'} \Big|_{ss}. \end{aligned}$$

We next partition $M_{(\tilde{h}+\tilde{n}) \times (\tilde{h}+\tilde{n})}$ as $\begin{pmatrix} M_{11} & M_{12} \\ M_{21} & M_{22} \end{pmatrix}$, where M_{11} is $(\tilde{h} \times \tilde{h})$; M_{12} is $(\tilde{h} \times \tilde{n})$; M_{21} is $(\tilde{n} \times \tilde{h})$; and M_{22} is $(\tilde{n} \times \tilde{n})$. Finally, we can now obtain the desired quadratic form for $r(z_t)$, i.e.

$$z_t' M z_t = \begin{pmatrix} x_t \\ \tilde{u}_t \end{pmatrix}' \begin{pmatrix} M_{11} & M_{12} \\ M_{21} & M_{22} \end{pmatrix} \begin{pmatrix} x_t \\ \tilde{u}_t \end{pmatrix}.$$

²³See Ljungvist and Sargent (2000, ch. 4) Appendix B for details.