Model Uncertainty, Precautionary Learning and Welfare*
(Preliminary Draft)

Konstantinos Angelopoulos
University of Glasgow

James Malley
University of Glasgow and CESifo

November 20, 2009

Abstract

Using a simple two-sector real business cycle model and robust control theory, this paper first establishes that the welfare costs of model uncertainty can be significant and are increasing in the variance of shocks to productivity. Model uncertainty creates incentives for precautionary behavior in the form of increased savings and education effort. Based on these findings, we allow for the possibility that agents can attempt to reduce model uncertainty by optimally diverting part of their time allocation to learn more about the economy/model. This is achieved by endogenizing the entropy constraint of the robust decision maker. In the extended model, in addition to increased savings in response to a rise in the volatility of productivity shocks, precautionary learning or the time allocated to education and information acquisition effort also increases. These changes in turn cause a fall in model uncertainty and hence an improvement in welfare relative to the model without precautionary learning.

*We would like to thank the Center for Economic Studies, University of Munich for financial support when conducting this research.

Standard control theory tells a decision maker how to make optimal decisions when his model is correct. Robust control theory tells him how to make good decisions when his model approximates a correct one. [Hansen and Sargent (2008)]

1 Introduction

The path breaking research of Hansen and Sargent (1995, 2001 and 2008) on robust control theory has focused the attention of the economics profession on the importance of model uncertainty and its effects on decision making. The starting point of their approach assumes that economic agents make choices in an environment where they are uncertain about the model that generates the data and thus must rely on approximations to it. In such a setup, robust control tools can help to make choices that will guarantee a minimum level of welfare even in the face of very unfavorable realizations of the stochastic state variable(s).

Two main findings of this research programme to date are first that model uncertainty leads to precautionary behavior and second that it implies welfare costs for agents populating the model. As discussed in Hansen and Sargent (2008, ch.10) and Hansen, Sargent and Tallarini (1999), the robust decision maker prefers to increase non-human and human wealth to protect against model misspecification. This type of precautionary behavior 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 shocks 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. Under model uncertainty, welfare is reduced in the sense that a robust decision maker would be willing to pay to eliminate it (see e.g. Barillas, Hansen and Sargent (2008)).

In the robustness literature to date, the size of model uncertainty is generally determined exogenously. In particular, model uncertainty 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. The two models (true and approximating) are restricted to be close in the statistical sense that the life-time conditional relative entropy cannot be larger than an exogenous constant. Usually, this constant is 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. If an agent selects a particular model over an alternative, there is a non-zero probability of making a detection error since the models can differ due to randomness. The lower this probability, the easier it is to distinguish between models, as the approximating and the true model are very different. In other words, the agent in this case acknowledges that the two models can be very different, or else that the entropy constraint is not very tight and the agent's behavior will be more conservative (or more cautious). Hence, a particular constraint on model uncertainty implies a degree of fear of model misspecification.

In contrast to determining the size of model uncertainty exogenously, we propose to allow the entropy constraint to be endogenous to the decision maker's actions. In particular, we assume that the agent uses resources to acquire information to reduce model uncertainty. This reflects the idea that economic agents spend time informing themselves about the economic environment to "learn" the model. In practice, a part of effort and leisure time is spent by economic agents in information-acquiring activities. This includes, for example, reading newspapers and books in leisure time, going to meetings and catching up with developments and news in work or education time. These activities do not increase the productivity of future labor but, instead, help the individual have a better understanding of the economic and socio-political system.

The above form of information acquisition activities result in a type of "precautionary learning", where the agent increases educational activities (for productive and information gathering purposes) due to fears of model misspecification. This "precautionary learning" is again increasing in the variance of the stochastic process. The outcome of such learning is a reduction in model uncertainty and an improvement in welfare relative to the situation where the entropy constraint is exogenous.

To illustrate our ideas in a parsimonious and transparent fashion we use a simple two-sector dynamic general equilibrium model where human capital enhances productivity. The model is a variant of the two-sector real business cycle models in e.g. Perli and Sakellaris (1998) and DeJong and Ingram (2001) and allows for uncertainty relating to the process generating total factor productivity. In the benchmark case of an exogenous entropy constraint on model misspecification, model uncertainty leads to precautionary savings and precautionary education activities, as the robust decision maker tries to create a buffer stock of physical and human capital, to protect against potential very bad realizations of technology. This precautionary behavior is more pronounced when the variance of the stochastic process is increasing, as the economic agent becomes more exposed to model uncertainty. Model uncertainty also reduces welfare compared to the case where the agent knows the true model. Moreover, the welfare cost is increasing in the variance of the stochastic process. Consistent with the findings in Barillas, Hansen and Sargent (2008), increases in economic volatility hurt economic agents since this volatility also increases model uncertainty.

The above findings also maintain qualitatively when we allow for an additional use of time in the form of information acquiring activities which aims to reduce model misspecification. In the extended setup, the robust decision maker combines human capital with time used to acquire information to learn about the economy/model. Therefore, there are now two types of knowledge created: (i) productive knowledge that enhances productivity and hence increases production; (ii) knowledge about the model which reduces the agent's exposure to model misspecification. Time for information, however, comes at a cost, as it reduces the total time available for leisure, work and education effort. Given this trade-off, the robust decision maker can obtain welfare gains by reducing model uncertainty. The amount by which model uncertainty falls depends, of course, on the productivity of information acquisition, but, also on the variance of the stochastic process. In other words, the robust decision maker has an incentive to decrease model uncertainty by more when economic uncertainty is greater. This follows since it is in these environments that model uncertainty can cause the greatest losses in welfare. Therefore, a prediction of this model is that information acquiring activities will be intensified in more volatile economic environments.

The rest of the paper is organized as follows. In Section 2 we quantify the precautionary behavior induced by and the welfare costs resulting from model uncertainty. In Section 3 we introduce information acquisition in a way that makes the entropy constraint endogenous to the actions of the economic agent. Finally Section 4 contains the preliminary conclusions and some discussion of our on going work.

2 Robust business cycle model

As discussed in the introduction, to conduct our analysis we use a standard two-sector RBC model in the presence of model uncertainty. In contrast to approaches which assume the model is correct, following Hansen and Sargent (2007), we employ an approximating model which allows for particular types of model misspecification. The general equilibrium solution we initially 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, human capital stock, H_t, and the fractions of time allocated to work, n_t, and education, e_t.[1]

In this section we start by setting out the non-linear setups for the robust non-stationary and stationary models followed by the linear-quadratic (LQ) approximation of the latter. We then solve the stationary LQ representation to obtain decision rules which are robust to model misspecification. Finally, we calibrate the model to quantify the effects of model uncertainty on precautionary behavior and welfare.

2.1 Non-stationary setup

Taking the stationary path of technology, $\{\widetilde{a}_{t+1}\}_{t=0}^{\infty}$, as given, the representative agent chooses $\{Y_t, C_t, I_t, I_t^h, n_t, e_t, K_{t+1}, H_{t+1}\}_{t=0}^{\infty}$ to maximize expected lifetime utility:

$$E_0 \sum_{t=0}^{\infty} \beta^t \frac{\left[(C_t)^{\mu} (l_t)^{(1-\mu)} \right]^{1-\varsigma}}{1-\varsigma} \tag{1a}$$

subject to:

$$l_t = 1 - n_t - e_t \tag{1b}$$

$$C_t + I_t + I_t^h = Y_t \tag{1c}$$

$$H_{t+1} = (1-\delta^h) H_t + \kappa \left(e_t H_t \right)^{\gamma} \left(I_t^h \right)^{1-\gamma}, H_0 > 0 \tag{1d}$$

$$K_{t+1} = (1-\delta^k) K_t + I_t, K_0 > 0 \tag{1e}$$

$$Y_t = \exp\left(\widetilde{a}_t\right) (K_t)^{\alpha} (H_t)^{\zeta} (Z_t n_t)^{1-\alpha-\zeta} \tag{1f}$$

where, E_0 is the conditional expectations operator; $0 < \beta < 1$ is the time discount factor; $1/\varsigma > 0$ is the intertemporal elasticity of substitution; $0 < \mu < 1$ is the weight given to consumption relative to leisure, l_t, in the utility function;[2] $0 < \delta^k, \delta^h < 1$ 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 two exogenous technology processes in the model. The first is the Harrod-neutral, labour augmenting technical progress Z_t, which grows

[1] When we endogenize the constraint for model uncertainty or the entropy constraint below, we also solve for the paths of information investment, H_t^i and information time, q_t.

[2] The utility function implies that there are exactly offsetting income and substitution effects of wage changes on labour supply. This guarantees that in the steady state the labour input does not grow when productivity grows (see e.g. King, Plosser and Rebelo (1988) and King and Rebelo (1999) for more details and regularity conditions that the utility function satisfies).

according to the gross rate g, such that $\frac{Z_{t+1}}{Z_t} = g$ and $Z_0 > 0$. The second process is stochastic total factor productivity (TFP), A_t, given by

$$\widetilde{a}_{t+1} = (1-\rho)\widetilde{a} + \rho\widetilde{a}_t + \sigma(\epsilon_{t+1} + w_{t+1}) \tag{1g}$$

where, $\widetilde{a}_t = \log(A_t)$; $A_0 > 0$; $0 < \rho < 1$ is a first order autocorrelation coefficient; $\widetilde{a}$ is a constant; and ϵ_{t+1} is a Gaussian random variable distributed identically and independently through time, with zero mean and unit variance. To capture fears that there are unknown stochastic perturbations to the above model, an extra process, w_{t+1}, is incorporated. In particular, this shock may be a non-linear time dependent function, φ_t, of the past states, i.e.

$$w_{t+1} = \varphi_t(x_t, x_{t-1}, ...)$$

where $\{\varphi_t\}$ is a sequence of measurable functions; x_t (to be defined later) is a vector including the state variables of the model; and $\sigma > 0$ scales the variances of the ϵ_{t+1} and w_{t+1} shocks. Thus, this specification implies that the agent cannot use the true model to make optimal choices, since the exogenous w_{t+1} process is unknown.

2.2 Stationary representation

To solve the model requires that we approximate the stationary objective function around the non-stochastic steady-state. However, in the above setup, all endogenous variables, except n_t and e_t, grow in the steady-state at the exogenous gross rate of labour augmenting technical progress, g. Thus to render the remaining variables stationary we define $\pi_t = \Pi_t/Z_t$, for $\Pi_t = \{Y_t, C_t, I_t, I_t^h, K_t, H_t\}$. The transformed setup can be written as:

$$\max_{\{y_t,c_t,i_t,i_t^h,n_t,e_t,k_{t+1},h_{t+1}\}_{t=0}^{\infty}} E_0 \sum_{t=0}^{\infty} \breve{\beta}^t \frac{\left[(c_t)^{\mu}(l_t)^{(1-\mu)}\right]^{1-\varsigma}}{1-\varsigma} \tag{2a}$$

subject to:

$$l_t = 1 - n_t - e_t \tag{2b}$$

$$c_t + i_t + i_t^h = y_t \tag{2c}$$

$$gh_{t+1} = (1-\delta^h)h_t + \kappa\,(e_t h_t)^{\gamma}\left(i_t^h\right)^{1-\gamma}, h_0 > 0 \tag{2d}$$

$$gk_{t+1} = (1-\delta^k)k_t + i_t, k_0 > 0 \tag{2e}$$

$$y_t = \exp\left(\widetilde{a}_t\right)(k_t)^{\alpha}(h_t)^{\zeta}(n_t)^{1-\alpha-\zeta} \tag{2f}$$

$$\widetilde{a}_{t+1} = (1-\rho)\widetilde{a} + \rho\widetilde{a}_t + \sigma(\epsilon_{t+1} + w_{t+1}) \tag{2g}$$

where, $0 < \left(\breve{\beta} \equiv \beta g^{\mu(1-\varsigma)}\right) < 1$.[3]

[3] See Appendix A for the derivation of $(2a)$.

2.3 Stationary LQ representation

2.3.1 Transformed objective function

We next to eliminate non-linearities from the constraints of the stationary setup. To this end, we start by defining a new variable:

$$m_t = \kappa \left(e_t h_t\right)^{\gamma} \left(i_t^h\right)^{1-\gamma} \tag{3a}$$

which implies

$$i_t^h = \left(\frac{m_t}{\kappa \left(e_t h_t\right)^{\gamma}}\right)^{\frac{1}{1-\gamma}}. \tag{3b}$$

Substituting $(3b)$ and $(2f)$ into $(2c)$ for i_t^h and y_t respectively and the resulting expression for c_t as well as $(2b)$ into the objective function $(2a)$ gives:

$$\max_{\{u_t\}_{t=0}^{\infty}} E_0 \sum_{t=0}^{\infty} \breve{\beta}^t \frac{\left[\left(a_t \left(k_t\right)^{\alpha} \left(h_t\right)^{\zeta} \left(n_t\right)^{1-\alpha-\zeta} - i_t - \Omega_t^{\frac{1}{1-\gamma}}\right)^{\mu} \left(1 - n_t - e_t\right)^{(1-\mu)}\right]^{1-\varsigma}}{1-\varsigma} \tag{4a}$$

where $a_t = \exp\left(\widetilde{a}_t\right)$ and $\Omega_t = \frac{m_t}{\kappa(e_t h_t)^{\gamma}}$. Thus, the representative agent now chooses $u_t = \{i_t, m_t, n_t, e_t\}_{t=0}^{\infty}$ to maximize $(4a)$ subject to:

$$g k_{t+1} = (1-\delta^k)k_t + i_t, k_0 > 0 \tag{4b}$$

$$g h_{t+1} = (1-\delta^h)h_t + m_t, h_0 > 0 \tag{4c}$$

$$\widetilde{a}_{t+1} = (1-\rho)\widetilde{a} + \rho\widetilde{a}_t + \sigma(\epsilon_{t+1} + w_{t+1}). \tag{4d}$$

Letting $\widetilde{h}$ and $\widetilde{n}$ equal the number of states $(1, k_t, h_t, \widetilde{a}_t)$ and controls (i_t, m_t, n_t, e_t) respectively, we can now write the linear constraints $(4b) - (4d)$ in matrix form as:[4]

$$x_{t+1} = Ax_t + Bu_t + C(\epsilon_{t+1} + w_{t+1}) \tag{5}$$

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

where

$$x_t = \left[\begin{array}{cccc} 1 & k_t & h_t & \widetilde{a}_t \end{array} \right]'; \qquad u_t = \left[\begin{array}{cccc} i_t & m_t & n_t & e_t \end{array} \right]';$$

$$A_{(\widetilde{h}x\widetilde{h})} = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & \frac{1-\delta^k}{g} & 0 & 0 \\ 0 & 0 & \frac{\left(1-\delta^h\right)}{g} & 0 \\ (1-\rho)\widetilde{a} & 0 & 0 & \rho \end{bmatrix};$$

$$B_{(\widetilde{h}x\widetilde{n})} = \begin{bmatrix} 0 & 0 & 0 & 0 \\ \frac{1}{g} & 0 & 0 & 0 \\ 0 & \frac{1}{g} & 0 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix}; \quad \text{and } C_{(\widetilde{h}x1)} = \begin{bmatrix} 0 \\ 0 \\ 0 \\ \sigma \end{bmatrix}.$$

2.3.2 LQ approximation of the transformed objective function

To simplify notation we next express the objective function in $(4a)$ at $t = 0$ in implicit form as $r(z_t)$ where $z_t = [1\ k_t\ h_t\ \widetilde{a}_t\ i_t\ m_t\ n_t\ e_t]'$. 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,[5] using the first two terms of a Taylor series expansion, i.e.

$$r(z_t) \approx \widehat{r}(z_t) = r(z) + (z_t - z)' \frac{\partial r}{\partial z_t}|_{ss} + \frac{1}{2}(z_t - z)' \frac{\partial^2 r}{\partial z_t \partial z_t'}|_{ss}(z_t - z) \qquad \text{(6a)}$$

where, the partial derivatives in $(6a)$ are evaluated at the steady-state z (see also Anderson *et al.* (1996) for more details).

Let j be a $((\widetilde{h} + \widetilde{n}) \times 1)$ 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 $(6a)$ as:[6]

$$\begin{aligned} \widehat{r}(z_t) &= z_t' M z_t, \text{ where} && \text{(6b)} \\ M &= j \left[r(z) - \left(\frac{\partial r}{\partial z_t}|_{ss} \right)' z + \frac{1}{2} z' \frac{\partial^2 r}{\partial z_t \partial z_t'}|_{ss} z \right] j' + \\ & \frac{1}{2} \left[\frac{\partial r}{\partial z_t}|_{ss} j' - j z' \frac{\partial^2 r}{\partial z_t \partial z_t'}|_{ss} - \frac{\partial^2 r}{\partial z_t \partial z_t'}|_{ss} z j' + j \left(\frac{\partial r}{\partial z_t}|_{ss} \right)' \right] \\ & + \frac{1}{2} \frac{\partial^2 r}{\partial z_t \partial z_t'}|_{ss}. \end{aligned}$$

[5] The fixed point z is the steady-state of the non-stochastic robust control problem. See Appendix B, Table 7 and Appendix F, Table 8 for the steady-states associated with the models incorporating exogenous and endogenous entropy constraints respectively.

[6] See Ljungvist and Sargent (2000, ch. 4) Appendix B for details.

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

$$z_t' M z_t = \begin{pmatrix} x_t \\ u_t \end{pmatrix}' \begin{pmatrix} M_{11} & M_{12} \\ M_{21} & M_{22} \end{pmatrix} \begin{pmatrix} x_t \\ u_t \end{pmatrix}. \tag{6c}$$

2.3.3 Stationary LQ problem

Given the above, we can re-express the stationary non-linear setup given by $(2a)$-$(2g)$ approximately as:

$$\max_{\{u_t\}_{t=0}^{\infty}} E_0 \sum_{t=0}^{\infty} \breve{\beta}^t \begin{pmatrix} u_t \\ x_t \end{pmatrix}' \begin{pmatrix} R & W' \\ W & Q \end{pmatrix} \begin{pmatrix} u_t \\ x_t \end{pmatrix} \tag{7a}$$

subject to:

$$x_{t+1} = Ax_t + Bu_t + C(\epsilon_{t+1} + w_{t+1}) \tag{7b}$$

where $R = M_{22}$, $W = M_{12}$ and $Q = M_{11}$.[7]

2.4 Model uncertainty

As pointed out above, the decision maker cannot use the correct model $(7b)$ since the w_{t+1} process is unknown.[8] Instead, a misspecified approximating model is employed:

$$\widetilde{a}_{t+1} = (1-\rho)\widetilde{a} + \rho\widetilde{a}_t + \sigma\breve{\varepsilon}_{t+1} \tag{8}$$

where $\breve{\varepsilon}_{t+1} \sim iid\ N(0,1)$. Following Hansen and Sargent (2007), by introducing model uncertainty, we essentially bound the model the decision maker employs, i.e. (8), with a set of alternative models of the form given by $(4d)$. These models allow for more general forms of the approximation errors than those implied by $\breve{\varepsilon}_{t+1}$ in (8). The correct model in this context can then be viewed as a distorted or perturbed version of the approximating model

[7] Note that this setup is consistent with the form used for the discounted stochastic regulator problem (see e.g. Anderson *et al.* (1996)).

[8] As in Hansen and Sargent (2007) we also assume, at least initially, that the agent forgoes learning the correct model since this model and the approximating model are close in a statistical sense. This working hypothesis is predicated on the assumption that a moderate-sized time series of data on the past history of the state is available. See Hansen and Sargent (2008, pp. 17 and 27 and ch. 9 for more details).

since the conditional mean of w_{t+1} is non-zero and as suggested above can influence the evolution of the states.

To operationalize the notion that the approximating model (8) is a good approximation when ($4d$) generates the data, we follow Hansen and Sargent (2007) and restrict w_{t+1} as follows. First we 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:

$$f_a(\widetilde{a}_{t+1} \mid \widetilde{a}_t) \sim N((1-\rho)\widetilde{a} + \rho\widetilde{a}_t, \sigma^2) \tag{9a}$$

while, the transition density for the perturbed model is[9]:

$$f(\widetilde{a}_{t+1} \mid \widetilde{a}_t) \sim N((1-\rho)\widetilde{a} + \rho\widetilde{a}_t + \sigma w_{t+1}, \sigma^2). \tag{9b}$$

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 $\widetilde{a}_t$ to $\widetilde{a}_{t+1}$, 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:

$$I(f_a, f)(\widetilde{a}) = \int \log\left(\frac{f(\widetilde{a}_{t+1} \mid \widetilde{a}_t)}{f_a(\widetilde{a}_{t+1} \mid \widetilde{a}_t}\right) f(\widetilde{a}_{t+1} \mid \widetilde{a}_t) d\widetilde{a}_{t+1}. \tag{10a}$$

It is then straightforward to show that (see Appendix C):

$$I(f_a, f)(\widetilde{a}) = I(w_{t+1}) = 0.5w_{t+1}^2. \tag{10b}$$

The corresponding intertemporal measure of the size model misspecification, $S(w)$, is defined as

$$S(w) = E_0 \sum_{t=0}^{\infty} \breve{\beta}^{t+1}\left(w_{t+1}' w_{t+1}\right) \tag{11a}$$

where E_0, conditioned on x_0, is evaluated with respect to the perturbed model. Moreover, $S(w)$ is constrained to satisfy what is termed as the entropy constraint:

$$S(w) \leq \eta_0 \tag{11b}$$

[9] Note that ($4d$) allows for misspecifications that occur only as a perturbation to the conditional mean of the innovation to the state, $\widetilde{a}_{t+1}$ (and hence x_{t+1} as well) in (8) and leaves the conditional volatility of the shock, as parametrized by σ, unchanged. We provide further information on this point in Appendix F.

where $\eta_0 \geq 0$. This constraint suggests that for larger values of η_0, "fear of model misspecification" is greater.

To summarize, the decision maker believes that the TFP data are generated by the model given by $(4d)$. However, the approximating model (8) must be used since the w_{t+1} process is unknown. Despite not being able to make optimal decisions using the correct model, the agent desires to make good decisions over a set of models $(4d)$ satisfying the entropy constraint $(11b)$.[10] Such decisions can then be characterized as being robust to misspecification of the approximating model.

2.5 Model Solution

As suggested above, the agent would like to obtain decision rules that are robust to model misspecification, in the sense that they provide good results even in the presence of unfavorable w_{t+1} shocks. For the decision rule to ensure the agent a lower bound of utility in the least unfavorable environment, the agent makes choices as if the w_{t+1} process follows a 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.[11]

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

$$\max_{\{u_t\}_{t=0}^{\infty}} \min_{\{w_{t+1}\}_{t=0}^{\infty}} E_0 \sum_{t=0}^{\infty} \breve{\beta}^t \begin{pmatrix} u_t \\ x_t \end{pmatrix}' \begin{pmatrix} R & W' \\ W & Q \end{pmatrix} \begin{pmatrix} u_t \\ x_t \end{pmatrix} \tag{12a}$$

subject to the perturbed model $(7b)$ and the entropy constraint $(11b)$ for a given η_0 satisfying $\overline{\eta} > \eta_0 \geq 0$. Here, $\overline{\eta}$ measures the largest set of perturbations against which it is possible to attain robustness and E_0, conditioned on x_0, is evaluated with respect to the perturbed model.

[10] Note 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 (i.e. for the worst-case w_{t+1} process; see Hansen and Sargent (2007, pp. 37, 142 and 158 for more details)).

[11] In this context Hansen and Sargent (2007) argue that the representative agent is not being too cautious, by putting so much weight on a very unlikely scenario. Given the entropy constraint (11b), the perturbed models are difficult to distinguish statistically from the approximating model with the amount of data at hand.

The problem described by (12a) is known as the constraint problem. As discussed in Hansen and Sargent (2007) it is equivalent to the multiplier problem given by:

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

where the extremization is subject to the perturbed model (7b).[12] Here, $\underline{\theta} \leq \theta \leq +\infty$ is a penalty parameter restraining the minimizing choice of the w_{t+1} sequence. The lower bound, $\underline{\theta}$, is the so-called breakdown point beyond which it is fruitless to seek more robustness because the minimizing agent is sufficiently unconstrained and can push the criterion function to $-\infty$, despite the best response of the maximizing agent.[13]

The constraint and the multiplier problems in the context of our application have the same solution under four different timing protocols:[14] (i-ii) a static Stackelberg game, where the maximizing (minimizing) player commits at time 0 to an entire sequence for $\{u_t\}_{t=0}^{\infty}$, after which the minimizing (maximizing) player commits to a sequence for $\{w_{t+1}\}_{t=0}^{\infty}$; (iii-iv) and a Markov-perfect equilibrium in a recursive game, where the two players move sequentially, with the maximizing (minimizing) player playing first.[15] Here we solve for a Markov-perfect equilibrium where the maximizing player moves first. The solution to this problem, described in detail in Appendix E, is given by the policy functions[16]:

$$\begin{aligned} u_t &= -F x_t \\ w_{t+1} &= K x_t \end{aligned} \quad \text{(13a)}$$

[12]Following Whittle (1990), extremization denotes joint maximization and minimization.

[13]Essentially, values of θ below the lower bound imply that the objective is not convex in w_{t+1}, so that the minimizing agent can push the value of the game to $-\infty$. See Appendix D on how to relate θ to $\overline{\eta}$ so that the two problems have equivalent outcomes. We shall also discuss below how to calibrate θ and test whether $\theta > \underline{\theta}$.

[14]That all timing protocols imply the same solution results from the fact that the two-player games are zero-sum and the agents in those have the same objective. This implies that the two players' preferences are perfectly misaligned (see Hansen and Sargent, 2007, ch. 7 for more details).

[15]See Hansen and Sargent (2007, ch. 7.7) who use Markov perfect equilibria as candidates for the solution of the Stackelberg games and verify that they provide the solution to these games.

[16]To implement the solution to the multiplier problem in (12b) we replace $r(z_t)$ in (4a) by the linear-quadratic form in (12b) and evaluate the second-order Taylor approximation around the non-stochastic steady-state solution of the multiplier problem in (12b) (see Appendix B for more details).

and the state evolution equation:

$$x_{t+1} = (A - BF + CK)x_t \tag{13b}$$

where

$$F = (R + \check{\beta}B'D(P)B)^{-1} \left(\check{\beta}B'D(P)A + W'\right);$$
$$K = -(\theta I + C'PC)^{-1}C'P\,[A - BF];$$

and P is obtained by iterating on the composite operator:

$$\begin{aligned} D(P) &= P - PC(\theta I + C'PC)^{-1}C'P \\ P &= Q + \check{\beta}A'D(P)A - \left(\check{\beta}A'D(P)B + W\right) \times \\ &\quad (R + \check{\beta}B'D(P)B)^{-1} \left(\check{\beta}B'D(P)A + W'\right). \end{aligned}$$

The stability of the solution is guaranteed if the eigenvalues of $(A - BF + CK)$ are all less than one in absolute value.[17] The value of the game or welfare is given by:

$$v(x_0) = x_0' P^* x_0 + d \tag{14}$$

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

2.6 Quantitative analysis

To implement the above solution quantitatively, we require numeric values for the parameters of the model. We start with the calibration of the structural parameters followed by the calibration of θ. Our general strategy is to use empirical evidence from econometric studies to calibrate some of the structural parameters and then choose 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. With the calibrated model in hand, we are then in a position to calculate the welfare costs associated with model uncertainty.

2.6.1 Calibrating the structural parameters

For the basic RBC model we examine here, in Table 1 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) and Angelopoulos *et al.* (2008)). We

[17] Except, of course, for the one corresponding to the constant term, which will be exactly unity.

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 be equal to the average hours of work versus leisure time, which is obtained using data on hours worked.[18] We use a value for $1/\varsigma$ that is common in the DSGE literature (i.e. $\sigma = 2$), as micro evidence suggests (although not conclusively) that $(1/\sigma)$ is less than one.

As our aim is to use the model to evaluate welfare around the steady-state, it is important that the calibrated parameters imply a steady-state that matches the long-run data averages in the real economy.[19] As mentioned above, this provides the criterion for choosing those parameters we cannot retrieve from the data or previous empirical studies, especially the exponents in the production function for human capital, i.e. $(\gamma,\ \zeta)$ and the constants in goods and human capital production functions, i.e. a and κ respectively.

Table 1: Model calibration

definition	parameter	value
productivity of capital	α	0.33
constant term in the TFP process	a	0.82
time discount factor	β	0.97
constant term in human capital production	κ	0.90
productivity of human capital investment	γ	0.70
depreciation rate on physical capital	δ^k	0.049
depreciation rate on human capital	δ^h	0.018
productivity of human capital	ζ	0.30
labor augmenting technical progress	g	1.025
weight on consumption in utility	μ	0.35
intertemporal elasticity of substitution	$1/\varsigma$	0.50

2.6.2 Calibrating θ in the entropy constraint

As discussed above, $\theta = +\infty$ (or equivalently $\eta_0 = 0$) corresponds to the case where there is no fear of model misspecification. In other words the robust solution is the same as the standard solution without model uncertainty.

[18] To 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.

[19] See Appendices B and F for more details.

Starting from $\theta = +\infty$, lowering θ increases the fear of misspecification, by increasing the loss to the utility that the minimizing agent can cause through the choice of w_{t+1}. There is a lower bound on θ, associated with the largest set of alternative models, as measured by relative entropy, against which it is feasible to seek a robust rule. For values of θ below this bound the malevolent agent is penalized so little that the criterion function can be driven to $-\infty$.[20]

It is obviously important to find an "appropriate" calibration for θ, as this will measure the degree of fear of (and also of the robustness to) misspecification in the model. Hansen and Sargent (2008, see ch. 9) suggest a method based on detection error probabilities. The basic idea is to set θ so that, given the finite amount of data available, a decision maker would find it difficult to statistically distinguish members of a set of alternative models against which robustness is sought.

The method proceeds as follows. First, assume that the approximating model given by (8) is the true data generating process. Hence, the true error process is $\breve{\varepsilon}_{t+1}$. However, the observed data cannot distinguish between $\breve{\varepsilon}_{t+1}$ and $\varepsilon_{t+1} + w_{t+1}$ (i.e. the error process under the perturbed model, $(4d)$) with certainty, as in reality the observed error process is ε_{t+1}. In other words, the agent observes ε_{t+1}, but does not know whether this has been generated by $\breve{\varepsilon}_{t+1}$ or by $\varepsilon_{t+1} + w_{t+1}$. To help select between the two models, a likelihood ratio test is formed. Assume that, given the observed ε_{t+1}, the likelihood of observing this process using the approximating model is L_A, while the likelihood using the perturbed model is L_B. Also note that the likelihood under the perturbed model is a function of θ (as different values for θ will give different solutions for the perturbed model). The log-likelihood ratio is $\log \frac{L_A}{L_B}$. If this is larger than zero, the agent would select the approximating model. Hence, if the approximating model is the true data generating process (DGP) and the agent uses $\log \frac{L_A}{L_B}$ to decide between the two models, a model detection error is made when the ratio is negative. This follows since, in this case, the perturbed model is selected when the true DGP is the approximating model. Thus the probability of a model detection

[20] Formally, for the inner minimization in the multiplier problem (12b) to be well defined (i.e. so that a minimum exists and the infimum is not $-\infty$) the objective function, J, must be convex with respect to w_{t+1} (see Appendix D). By examining the minimization problem, we can calculate $\frac{\partial^2 J}{\partial w_{t+1}} = \theta I + C'PC$. Hence, if $\theta I + C'PC$ is positive definite, the objective J is convex. Therefore, to check whether the calibrated value of θ we use is above the lower bound we need to ensure that the eigenvalues of $\theta I + C'PC$ are all positive. A failure of this condition to hold might imply that a fixed point on the iterations does not exist. Therefore we need to check whether this condition holds in each iteration on the composite operator, $D(P)$, described above.

error is:

$$p_A = Prob\left(\log\frac{L_A}{L_B} < 0 \mid A\right). \tag{15a}$$

Next assume that the perturbed model is the true data generating process. Again, the agent observes ε_{t+1}, but does not know whether this has been generated by $\breve{\varepsilon}_{t+1}$ or by $\varepsilon_{t+1} + w_{t+1}$. To help select between the two models, the likelihood ratio test $\log\frac{L_A}{L_B}$ is again formed. If this is larger than zero, the agent would select the approximating model. Accordingly, if the perturbed model is the true DGP and the agent uses $\log\frac{L_A}{L_B}$ to decide between the two models, a model detection error is made when the ratio is positive since in this case the approximating model is selected when the true DGP is the perturbed model. Hence, the probability of a model detection error is:

$$p_B = Prob\left(\log\frac{L_A}{L_B} > 0 \mid B\right). \tag{15b}$$

Notice that since L_B (and thus $\log\frac{L_A}{L_B}$) is a function of θ, so are the two detection error probabilities. Finally, the overall probability of model detection error is formed by averaging p_A and p_B with prior probabilities 0.5:

$$p(\theta) = \frac{1}{2}(p_A + p_B). \tag{15c}$$

The idea then is to choose θ so that this probability is "reasonable", in the sense that it is not easy for the agent to select between the two models given the available observations on ε_{t+1}. In other words, select θ so that $p(\theta)$ is high enough, since if it is very low then the agent could easily select correctly between the two models, with a very low probability of error. Note that in the limiting case where $\theta \to +\infty$, implying that the agent does not fear misspecification, the detection error probability is 0.5 as the model under fear for misspecification (i.e. the perturbed model) is essentially the same as the approximating model. Hence, lowering θ results in falls in the detection error probability. In other words, the agent is trading off: (i) rules that are optimal under a particular model, but imply a high chance of using a wrong model; against (ii) rules that secure robust outcomes under a variety of potential models, but imply a lower chance of using a wrong model.

How fast the detection error probability falls when θ decreases depends on the particular models and how many observations on ε_{t+1} are available. The more observations available, the easier it is to distinguish between models, and hence the detection error probability should, *ceteris paribus*, fall faster. Hansen and Sargent (2007) suggest that a detection error probability of 0.1 is reasonable. The higher the detection error probability that the calibrated value of θ implies, the more difficult it is to distinguish between models and hence the less conservative (or cautious) the agent is.

To operationalize the above approach, we work as follows. First, assume that the approximating model is the true data generating process. Solve and simulate the model for a time period T as if there is no fear for misspecification, i.e. solve the "standard" RBC model and simulate a series of length T for x_t (denote this as x_t^A). Then, calculate the worst-case scenario process w_{t+1} that a malevolent agent would choose for the x_t^A process to minimize the agent's welfare, under the perturbed model. For this, use the optimal policy of the minimizing agent from the robust solution of the problem, $w_{t+1}^A = -Kx_t^A$. We can then calculate the log-likelihood ratio assuming that the approximating model is the true data generating process. As shown in Hansen and Sargent (2007), p. 216, this is equal to:

$$r \mid A = \frac{1}{T}\sum_{t=0}^{T-1}\left\{\frac{1}{2}w_{t+1}^{A\prime}w_{t+1}^{A} - w_{t+1}^{A\prime}\breve{\varepsilon}_{t+1}\right\}. \tag{16a}$$

We simulate the model 100,000 times and calculate $r \mid A$ as above and obtain the probability p_A by counting the fraction of realizations for which $r \mid A$ is negative.

Working similarly, we calculate p_B. In particular, we now assume that the perturbed model is the true data generating process. We then: (i) solve the "robust" RBC model and simulate a series of length T for x_t^B; (ii) calculate the worst-case scenario process w_t that a malevolent agent chooses for that x_t process, $w_{t+1}^B = -Kx_t^B$; and (iii) calculate the log-likelihood ratio assuming that the perturbed model is the true data generating process. Again as shown in Hansen and Sargent (2007), p. 217, this is equal to:

$$r \mid B = -\frac{1}{T}\sum_{t=0}^{T-1}\left\{\frac{1}{2}w_{t+1}^{B\prime}w_{t+1}^{B} + w_{t+1}^{B\prime}\varepsilon_{t+1}\right\}. \tag{16b}$$

We simulate the model and calculate $r \mid B$ as above to obtain the probability p_B by counting the fraction of realizations for which $r \mid B$ is positive. Finally, we obtain the model detection error probability by using (15c).

2.6.3 Results: exogenous model uncertainty

To help contextualize the differences in welfare due to model uncertainty, we set θ to a very large number in (13a) and (13b) and report the results for welfare given by (14) in Table 2. Note that in the case of $\theta = \infty$, the robust decision maker makes choices that are identical to those of a representative agent who trusts the model and thus we term the results obtained in Table 2 as "standard model" outcomes. We start with a baseline value of $\sigma = 0.01$ and then decrease (increase) it by half a percentage point.[21] As expected, since certainty equivalence holds in this setup, net welfare, $v(x_0) - d$, remains the same when uncertainty changes and the values of the endogenous variables reported in Table 2 do not change. Thus the small changes in $v(x_0)$ reported in Table 2 reflect the effect of changes in the constant term, d which incorporates σ (see the definition of d under (14)).

To examine the precautionary behavior due to model uncertainty, we report the model averages for the most relevant endogenous variables in Tables 2 and 3. In particular, we concentrate on the two types of goods investment, i and i^h, and the time investment in human capital or education time, e. We define the barred values in the Tables below as the averages obtained from simulating the model 1,000 times.[22]

Table 2: Standard Model

σ	θ	$v(x_0)$	$\bar{i}$	$\bar{i}^h$	$\bar{e}$
0.005	∞	-34.844	0.598	0.225	0.110
0.01	∞	-34.848	0.598	0.225	0.110
0.015	∞	-34.854	0.598	0.225	0.110

We next use the robust solution under model uncertainty for $\sigma = 0.01$ (see the middle rows of Tables 3) and calibrate θ using the detection error probability approach set out in subsection 2.6.2. Note that for these experiments we follow Hansen and Sargent (2008) and assume $T = 50$ and employ 100,000 simulations. The value of θ consistent with a $p = 10\%$ error detection probability implies a value of η_0 of about 4, obtained from $(d6)$ in Appendix D.[23]

Table 3: Exogenous Model Uncertainty

σ	θ	p	η_0	$\bar{i}$	$\bar{i}^h$	$\bar{e}$	$v(x_0)$	ϕ
0.005	0.096	0.100	4.189	0.643	0.241	0.112	-35.480	0.90
0.010	0.198	0.100	4.061	0.684	0.254	0.114	-36.117	1.79
0.015	0.307	0.100	3.934	0.719	0.264	0.116	-36.753	2.65

[21] Note that this range for the standard deviation of the TFP shocks encompasses values reported in both econometric and calibration studies. This also applies to the value employed for AR(1) parameter in the TFP process, i.e. $\rho = 0.99$.

[22] Note that in the certainty-equivalent standard model, these are equal to the steady-state values for these variables.

[23] This has been obtained by simulating the model 1,000 times and using the model generated values for w_{t+1}.

Table 3 indicates that welfare is reduced when there is model uncertainty, i.e. when η_0 changes from zero (the value implied in Table 2) to 4.061. This happens for all values of σ. In particular, we recalibrate θ for the different levels of σ to keep $p = 10\%$. Thus, in the face of a change in the size of the TFP shock, model uncertainty is held constant which implies that the set of alternative models with which we surround our approximating model does not change. This implies that the the restriction on model uncertainty, i.e. η_0 is basically the same, approximately 4, for all rows in Table 3.[24]

To find the welfare costs associated with model uncertainty, we compare the welfare figures in Table 2 with those from Table 3. Following Lucas (1990), to calculate these costs we assume a compensating consumption supplement at each date t that is proportional by ϕ to private consumption under the alternative assumptions regarding σ and makes $v(x_0)^j = v(x_0)^k$:

$$\phi = \frac{1}{(1-\varsigma)} \log \left[\frac{v(x_0)^j}{v(x_0)^k} \right] \times 100 \tag{17}$$

where the j, k superscripts refer to welfare under different values of σ. For $\sigma = 0.05, 0.01, 0.015$, the welfare costs are respectively $\phi = 0.9\%, 1.79\%, 2.65\%$. The results in Table 3 for ϕ suggest that the welfare costs of model uncertainty are nontrivial and increase with σ. Note that this increase in the welfare costs, when σ rises, is not due to the increased discounted lifetime uncertainty, $\eta_0 = \sum_{t=0}^{\infty} \check{\beta}^{t+1} \left(w'_{t+1} w_{t+1} \right)$, which, as pointed out above, is held constant in these experiments. Instead, they are explained by the fact that the true model can be further away from the approximating one in particular time periods, and this can on balance hurt the robust decision maker more. In other words, increased economic uncertainty, captured by a higher σ, allows the malevolent agent to choose a bigger w_{t+1} for particular time periods (see ($7b$)) and hurt the maximizing agent more, even if, over the lifetime, the entropy constraint is the same. Hence, higher economic uncertainty implies higher model uncertainty and higher welfare costs.

In Tables 2 and 3, we also report the model averages for the variables that are directly related to precautionary behavior, i.e. i, i^h and e. These results suggest that, consistent with previous work on robust choices, model uncertainty induces precautionary behavior. In this model, precautionary behavior

[24] Note that since both the calculations of p and η_0 are approximate and obtained by simulations of the model, it is difficult to match both exactly. Following e.g. Barillas *et al.* (2009), we prefer to calibrate θ so that p is 10% for all models, and then we calculate η_0 residually. We have also calibrated θ so that η_0 is exactly the same when σ changes and this gives very similar results regarding the welfare and precautionary behaviour presented here. In this case, p is approximately 10% for all cases.

takes the form of increases in both investment and education time. In particular, as the agent fears model misspecification, buffer stocks of physical and human capital are created to protect against potential very bad realizations of technology.

To obtain a quantitative notion of the magnitudes involved, in Table 4 we calculate the percentage increase in i, i^h and e in Table 3 relative to Table 2. Two results are noteworthy. First, increases in precautionary behavior are larger the bigger the volatility of TFP. As discussed above, increases in economic uncertainty, as captured by increases in σ, essentially imply that the economic agent is more exposed to model uncertainty, even if, in a lifetime sense, model uncertainty has not changed. As a result, the agent will want to increase precautionary physical and human capital accumulation. Second, the quantitative effect of precautionary behavior is more pronounced in physical investment, as opposed to education time investment, although precautionary learning in the form of education time is not insignificant.

Table 4: Precautionary Behavior
(% change)

σ	0.005	0.01	0.015
$\bar{i}$	7.53	14.38	20.23
$\bar{i}^h$	7.11	12.89	17.33
$\bar{e}$	1.82	3.64	5.45

3 Learning as information acquisition

The above results suggest that, despite using a model without distortions and low to moderate values of σ, model uncertainty can have nontrivial effects on welfare. However, the previous analysis treats model uncertainty as exogenous to the robust decision maker's actions. As Hansen and Sargent (2008, e.g. pp. 27) suggest, when the alternative models are statistically close so that it is not easy to statistically distinguish between them, such an assumption might be plausible. In the previous section we examined a case where the approximating and the distorted model were statistically close, in the sense that a likelihood ratio test used to select one over the other would leave a detection error probability of 10%. However, as we saw, the welfare costs associated with such model uncertainty are significant and would provide the robust decision maker with an incentive to try to reduce model uncertainty, if given the opportunity.

One way in which economic agents might reduce model uncertainty is to spend time informing themselves about the economic environment and thus expose themselves to less model misspecification. In practice, a part of effort and leisure time is spent by economic agents in information-acquiring activities. This includes, for example, reading newspapers and books in leisure time, going to meetings and catching up with developments and news in work or education time. These activities do not increase the productivity of future labor but, instead, help the individual have a better understanding of the economic and socio-political system.

3.1 Modified setup

The above discussion suggests that a share of (effort and leisure) time not used in productive-related activities, but in acquiring information, will help to reduce model uncertainty and hence increase welfare. However, allocating time to reducing model uncertainty carries the opportunity costs associated with work, education and leisure time. By allowing the entropy constraint to depend on this additional type of learning activity, we can model this trade-off to reveal the net benefits related to acquiring information and how much time needs to be devoted to such activities. Hence, we extend the previous model by accounting for an alternative use of time, q_t, which involves accumulating information about the economy with the view to reducing model uncertainty.

3.1.1 The entropy constraint with information acquisition

To see how information acquisition can reduce model uncertainty, we assume that the robust decision maker acknowledges that there is model uncertainty of the size of η_0 and tries to use (time) resources to restrict it. Hence, model uncertainty is a function of information acquisition, h_t^i. In particular, the entropy constraint becomes:

$$E_0 \sum_{t=0}^{\infty} \breve{\beta}^{t+1} \left(w'_{t+1} w_{t+1} \right) = \widetilde{\eta}(\eta_0, h_t^i) \tag{18a}$$

where $\widetilde{\eta}(\eta_0, h_t^i)$ satisfies the following properties: $\partial \widetilde{\eta}(\eta_0, h_t^i)/\partial \eta_0 > 0$, with $\lim\limits_{\eta_0 \longrightarrow 0} \widetilde{\eta}(\eta_0, h_t^i) = 0$ and $\partial \widetilde{\eta}(\eta_0, h_t^i)/\partial h_t^i < 0$, with $\lim\limits_{h_t^i \longrightarrow 0} \widetilde{\eta}(\eta_0, h_t^i) = \eta_0$. Hence, η_0 measures the extent of model uncertainty that would exist if there were no endogenous efforts to reduce it. A functional form for the entropy constraint that satisfies these requirements is:

$$E_0 \sum_{t=0}^{\infty} \breve{\beta}^{t+1} \left(w'_{t+1} w_{t+1} + \breve{\beta}^{-1} \eta_1 h_t^i \right) = \eta_0. \tag{18b}$$

In this specification, η_1, is an efficiency parameter that measures the effectiveness of information acquisition in restricting model uncertainty.

Information acquisition h_t^i implies an information production function such that:

$$h_t^i = q_t^{\eta_2} h_t^{(1-\eta_2)} \tag{18c}$$

and q_t denotes effort time for information, such that

$$l_t = 1 - n_t - e_t - q_t. \tag{18d}$$

In this specification, knowledge about the economic environment, is a constant returns to scale production function, which, analogous to the productive 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 information acquisition 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 individual (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 in this information production function. Hence, the entropy constraint becomes:

$$E_0 \sum_{t=0}^{\infty} \check{\beta}^{t+1} \left[w'_{t+1} w_{t+1} + \check{\beta}^{-1} \eta_1 q_t^{\eta_2} h_t^{(1-\eta_2)} \right] = \eta_0. \tag{18e}$$

The information production function, $q_t^{\eta_2} h_t^{(1-\eta_2)}$, effectively makes the entropy constraint for model uncertainty tighter. As the robust decision maker learns more about the economic environment, model uncertainty, or the set of actions allowed to the malevolent agent in the *maxmin* game, is reduced. Note that when $\eta_1 \longrightarrow 0$, this model reduces to that of exogenous entropy constraint, discussed in the previous sections.

3.1.2 The model with information acquisition

Formally, the model (in stationary form) is now:

$$\max_{\{y_t, c_t, i_t, i_t^h, n_t, e_t, q_t, k_{t+1}, h_{t+1}\}_{t=0}^{\infty}} E_0 \sum_{t=0}^{\infty} \check{\beta}^t \frac{\left[c_t^{\mu} (l_t)^{(1-\mu)} \right]^{1-\varsigma}}{1-\varsigma} \tag{19a}$$

subject to:

$$l_t = 1 - n_t - e_t - q_t \tag{19b}$$

$$c_t + i_t + i_t^h = y_t \tag{19c}$$

$$gh_{t+1} = (1-\delta^h)h_t + \kappa \left(e_t h_t\right)^{\gamma} \left(i_t^h\right)^{1-\gamma}, h_0 > 0 \tag{19d}$$

$$gk_{t+1} = (1-\delta^k)k_t + i_t, k_0 > 0 \tag{19e}$$

$$y_t = \exp(\widetilde{a}_t) \left(k_t\right)^{\alpha} \left(h_t\right)^{\zeta} \left(n_t\right)^{1-a-\zeta} \tag{19f}$$

$$\widetilde{a}_{t+1} = (1-\rho)\widetilde{a} + \rho \widetilde{a}_t + \sigma(\varepsilon_{t+1} + w_{t+1}) \tag{19g}$$

and the entropy constraint in ($18e$).

The time constraint ($19b$) shows that there are two educational/ information activities available to the agent, one that increases human capital and in turn the agent's productivity and a second that can be used to increase understanding about the economic environment. Therefore, the agent decides how to allocate the time endowment between work, n_t, education for the two types of knowledge, e_t and q_t, and leisure, $l_t = 1 - n_t - e_t - q_t$. The benefit from allocating time for information acquisition, derives from, as we saw in ($18e$), making the entropy constraint tighter.

A scheme like the above implies that, given the robust decision maker's fear of model misspecification (as measured by η_0), it will be optimal to try to allocate some resources to reducing model uncertainty. Working as above, provided that the multiplier θ satisfies the condition $\sum\limits_{t=0}^{\infty} \breve{\beta}^{t+1}[w'_{t+1}w_{t+1} + \breve{\beta}^{-1}\eta_1 q_t^{\eta_2} h_t^{(1-\eta_2)}] = \eta_0$, we can write the multiplier game for a robust decision maker for this problem as:

$$\max_{\{\widetilde{u}_t\}_{t=0}^{\infty}} \min_{\{w_{t+1}\}_{t=0}^{\infty}} E_0 \sum_{t=0}^{\infty} \breve{\beta}^t \{ \frac{1}{1-\varsigma} [\left(a_t \left(k_t\right)^{\alpha} \left(h_t\right)^{\zeta} \left(n_t\right)^{1-\alpha-\zeta} - i_t - \Omega_t^{\frac{1}{1-\gamma}} \right)^{\mu}$$
$$\times \left(1 - n_t - e_t - q_t\right)^{(1-\mu)}]^{1-\varsigma} +$$
$$\breve{\beta}\theta \left(w'_{t+1} w_{t+1} + \breve{\beta}^{-1} \eta_1 q_t^{\eta_2} h_t^{(1-\eta_2)} \right) \}$$

or equivalently

$$\max_{\{\widetilde{u}_t\}_{t=0}^{\infty}} \min_{\{w_{t+1}\}_{t=0}^{\infty}} E_0 \sum_{t=0}^{\infty} \beta^t \left\{ \left[\begin{pmatrix} \widetilde{u}_t \\ x_t \end{pmatrix}' \begin{pmatrix} \widetilde{R} & \widetilde{W}' \\ \widetilde{W} & \widetilde{Q} \end{pmatrix} \begin{pmatrix} u_t \\ x_t \end{pmatrix} \right] + \breve{\beta}\theta w_{t+1}^2 \right\} \tag{20a}$$

subject to:

$$x_{t+1} = \widetilde{A} x_t + \widetilde{B} \widetilde{u}_t + \widetilde{C}(\varepsilon_{t+1} + w_{t+1}) \tag{20b}$$

where now

$$x_t = \left[\begin{array}{cccc} 1 & k_t & h_t & \widetilde{a}_t \end{array}\right]'; \qquad \widetilde{u}_t = \left[\begin{array}{ccccc} i_t & m_t & n_t & e_t & q_t \end{array}\right]';$$

$$\widetilde{A} = \left[\begin{array}{cccc} 1 & 0 & 0 & 0 \\ 0 & \frac{1-\delta^k}{g} & 0 & 0 \\ 0 & 0 & \frac{\left(1-\delta^h\right)}{g} & 0 \\ (1-\rho)\widetilde{a} & 0 & 0 & \rho \end{array}\right];$$

$$\widetilde{B} = \left[\begin{array}{ccccc} 0 & 0 & 0 & 0 & 0 \\ \frac{1}{g} & 0 & 0 & 0 & 0 \\ 0 & \frac{1}{g} & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \end{array}\right]; \quad \widetilde{C} = \begin{bmatrix} 0_{3\times 1} \\ \sigma \end{bmatrix}$$

and $\widetilde{R}$, $\widetilde{W}$ and $\widetilde{Q}$ are the appropriate partitioned matrices of the respective $\widetilde{M}$ matrix that defines the second-order approximation of

$$\frac{\left[\left(a_t\left(k_t\right)^{\alpha}\left(h_t\right)^{\zeta}\left(n_t\right)^{1-\alpha-\zeta} - i_t - \Omega_t^{\frac{1}{1-\gamma}}\right)^{\mu}\left(1-n_t-e_t-q_t\right)^{(1-\mu)}\right]^{1-\varsigma}}{1-\varsigma}$$
$$+\theta\eta_1 q_t^{\eta_2} h_t^{(1-\eta_2)}$$

around the non-stochastic steady state solution of $(19a)-(19g)$ and $(18e)$, which is described in Appendix F.

For the problem to be well defined, we require that the objective is convex with respect to w_{t+1} (which will give the familiar condition that $\theta I + \widetilde{C}'P\widetilde{C}$ be positive definite as explained previously) and concave with respect to $\widetilde{u}_t$, which in turn requires that $\widetilde{R}$ is negative definite. As can be seen from the maximization problem of the robust decision maker, for the problem with to be well defined we need also $\widetilde{R} + \breve{\beta}B'D(P)B$ to be negative definite so that the objective of the maximizing agent, taking into account the optimal responses of the minimizing agent, remains concave. A necessary condition for the negative definiteness of $\widetilde{R}$ is that there are decreasing returns in $\theta\eta_1 q_t^{\eta_2} h_t^{(1-\eta_2)}$ with respect to q_t, which is satisfied when $0 \leq \eta_2 < 1$.

Note also that when $\theta \to \infty$, $\widetilde{R}$ is not concave. Intuitively, the maximizing agent can push the value of the game to infinity by choosing a non-zero q_t, despite the efforts of the minimizing agent. This also implies that when $\eta_0 \to 0$, the problem is not well defined, because when $\eta_0 \to 0$, $\sum_{t=0}^{\infty} \breve{\beta}^{t+1} w_{t+1}' w_{t+1} \to 0$. Hence, there is an upper bound on θ (or alternatively, a lower bound on η_0) above (below) which the problem is not well defined. Therefore, this problem restricts η_0 to be bounded as $0 < \eta^l \leq \eta_0 \leq \eta^u < \infty$, or, alternatively, θ to be bounded as $0 < \theta^l \leq \theta \leq \theta^u < \infty$, where the lower bound on θ corresponds to the upper bound on η_0 and *vice-versa*. When we solve the model below, we examine solutions that imply that the error detection probability is $p = 10\%$. This condition implies values for model uncertainty that are consistent with the above restrictions on θ and η_0. Thus, in this context, error detection probabilities, help to discipline the choices of both the minimizing and the maximizing agents.

3.1.3 Results: endogenous model uncertainty

To make the evaluation of the model with information acquisition comparable to previous analysis of the model with exogenous entropy constraints, we solve the model in (20a)-(20b) by requiring that the implied error detection probability is 10%. In other words, we impose on the robust decision maker (the malevolent agent) the restriction that model uncertainty cannot be reduced (increased) by more than what is implied by $p = 10\%$. Recall that as p increases, model uncertainty is decreased, as the true and approximating model are closer in a likelihood ratio sense, so that selecting one over the other implies a bigger error detection probability. Therefore, we solve for an economy where the existing model uncertainty associated with $p = 10\%$ is the outcome of both exogenous circumstances to and endogenous actions of the robust decision maker. The questions we pose are does the economic agent decide to allocate resources to learning with a view to reducing model uncertainty and what are the welfare implications of this precautionary behaviour?

To answer these questions, we solve the model in $(20a) - (20b)$, for θ so that $p = 10\%$. Given that the results will depend on the parameters that govern information creation, we provide a sensitivity analysis for η_1 and η_2. We first fix $\eta_2 = 0.5$ and examine the effects of increasing η_1. We start by examining a value of $\eta_1 = 0.1$ to obtain the results reported in Table 5a and then we examine the sensitivity of these results when η_1 is halved or doubled. Recall that when $\eta_1 \longrightarrow 0$, this model reduces to that of an exogenous entropy constraint.[25]

[25] Note that these figures implied reasonable values for the steady-state value of time allocated to information acquisition activities, q, and for the remaining model variables. For further details see Appendix F.

Table 5a: Endogenous Model Uncertainty

	$(\eta_1 = 0.1, \eta_2 = 0.5)$		
σ	0.005	0.01	0.015
θ	0.096	0.200	0.314
p	0.100	0.100	0.100
η_0	4.805	5.522	6.497
η_w	4.189	4.061	3.917
$\bar{i}$	0.639	0.689	0.740
$\bar{i}^h$	0.240	0.257	0.271
$\bar{e}$	0.113	0.116	0.120
$\bar{q}$	0.001	0.005	0.012
$v(x_0)$	-35.475	-36.096	-36.693
$v^{ex}(x_0)$	-35.532	-36.343	-37.326
ϕ_{ee}	0.08	0.34	0.86

The middle column of Table 5a presents results for $\sigma = 1\%$. These suggest that when the robust decision maker faces an error detection probability of 10%, the implied model uncertainty against which the agent is covered is given by $\eta_0 = 5.522$. However, for this to be achieved, resources have been used to restrict model uncertainty. For example, notice that information time, $\bar{q}$, is non-zero in this model and also that the resources allocated to human capital (education time, $\bar{e}$ and investment in human capital, $\bar{i}^h$) have also increased relative to the model with exogenous model uncertainty. We will return to quantify this form of precautionary behaviour, which we term "precautionary learning", below. Alternatively, to quantify the effects of h^i on reducing model uncertainty, we calculate $\eta_w = \sum_{t=0}^{\infty} \breve{\beta}^{t+1} w_{t+1}^2$, i.e. the remaining model uncertainty after taking into account "precautionary learning". Thus, the reduction in model uncertainty is $\eta_0 - \eta_w$.

To illustrate the incentives to the robust decision maker to restrict model uncertainty, we compare the welfare of this model with endogenous entropy constraint ($v(x_0)$), to the welfare of the model with the same entropy constraint, η_0, but with no information acquisition ($v^{ex}(x_0)$). The last column in Table 5a suggests that there are welfare gains $\phi_{ee} = 0.34\%$ for the robust decision arising from diverting some time and income resources to reducing model uncertainty.

Analogous results are obtained when we examine the cases of $\sigma = 0.5\%$ and $\sigma = 1.5\%$. In particular, the maximizing agent has an incentive engage in precautionary learning activities to decrease model uncertainty. Note that such precautionary learning behavior is stronger (and pays off more, in terms of welfare gains) when economic uncertainty is higher. This is not surprising, given that as discussed earlier, when economic uncertainty increases, the robust decision maker is more exposed to model uncertainty and thus the gains from reducing it are higher.

Table 5b: Endogenous Model Uncertainty

	$(\eta_1 = 0.05, \eta_2 = 0.5)$		
σ	0.005	0.01	0.015
θ	0.096	0.199	0.310
p	0.101	0.100	0.100
η_0	4.333	4.420	4.545
η_w	4.181	4.071	3.951
$\bar{i}$	0.635	0.674	0.713
$\bar{i}^h$	0.238	0.251	0.261
$\bar{e}$	0.112	0.115	0.117
$\bar{q}$	0.000	0.001	0.003
$v(x_0)$	-35.479	-36.116	-36.751
$v^{ex}(x_0)$	-35.482	-36.178	-36.902
ϕ_{ee}	0.00	0.09	0.21

Table 5c: Endogenous Model Uncertainty

	$(\eta_1 = 0.2, \eta_2 = 0.5)$		
σ	0.005	0.01	0.015
θ	0.097	0.206	0.332
p	0.100	0.101	0.101
η_0	6.785	10.768	16.266
η_w	4.177	3.950	3.782
$\bar{i}$	0.655	0.738	0.817
$\bar{i}^h$	0.248	0.274	0.288
$\bar{e}$	0.114	0.120	0.127
$\bar{q}$	0.004	0.021	0.059
$v(x_0)$	-35.456	-35.974	-36.331
$v^{ex}(x_0)$	-35.656	-36.953	-38.888
ϕ_{ee}	0.28	1.34	3.40

In Tables 5b and 5c above we repeat the above analysis for $\eta_1 = 0.05$ and $\eta_1 = 0.2$ respectively. As expected, when $\eta_1 = 0.05$ the model with an endogenous entropy constraint becomes very similar to the model with an exogenous entropy constraint. In particular, given the very low returns to information acquisition, precautionary learning is very limited, since q_t is almost zero, and the values for $\bar{i}$, $\bar{i}^h$ and $\bar{e}$ are very similar to the exogenous entropy constraint case. Hence, the reduction of model uncertainty $(\eta_0 - \eta_w)$ is very small.

On the other hand, when the returns to h^i increase, as in Table 5c, precautionary learning increases and, for higher levels of economic uncertainty, becomes a sizeable proportion of the time endowment. In these cases, the advantage of having access to a technology that reduces model uncertainty implies important welfare gains.

Table 5d: Endogenous Model Uncertainty

	$(\eta_1 = 0.1, \sigma = 0.01)$		
η_2	0.4	0.5	0.6
θ	0.206	0.200	0.199
p	0.101	0.100	0.100
η_0	11.236	5.522	4.271
η_w	3.922	4.061	4.089
$\bar{i}$	0.768	0.689	0.671
$\bar{i}^h$	0.282	0.257	0.249
$\bar{e}$	0.122	0.116	0.115→
$\bar{q}$	0.018	0.005	0.000
$v(x_0)$	-35.805	-36.096	-36.127
$v^{ex}(x_0)$	-36.991	-36.343	-36.152
ϕ_{ee}	1.63	0.34	0.03

In Table 5d above we examine the sensitivity of the results with respect to changes in η_2. For this experiment, we fix $\eta_1 = 0.1$ and $\sigma = 0.01$. Note that decreases in η_2 are expected to be beneficial for the maximizing agent, as reducing model uncertainty in this case depends more on a by-product of productive activities (h), rather than a diversion of resources (q). Indeed, as can be seen in Table 5d, lower values of η_2 imply more precautionary learning, a bigger reduction is model uncertainty and bigger welfare gains compared to a situation where the robust decision maker does not have access to a technology that reduces model uncertainty.

Table 6 presents the percentage increase in $\bar{i}$, $\bar{i}^h$ and learning activities, $\bar{e} + \bar{q}$ under endogenous model uncertainty, relative to the standard model. Tables 5 and 6 suggest that the main difference regarding precautionary behavior in this model is with respect to the learning activities. In particular, in addition to time allocated to information acquisition, q, which is increasing with σ, education time, e, has also increased in this model relative to the case of exogenous model uncertainty. The reason is that, in this environment, human capital also helps to decrease model uncertainty and this provides an additional incentive to increase education time.

Table 6: Precautionary Behavior (% change)

	$\eta_1 = 0.05$, $\eta_2 = 0.5$			$\eta_1 = 0.1$, $\eta_2 = 0.5$			$\eta_1 = 0.2$, $\eta_2 = 0.5$		
σ	0.005	0.01	0.015	0.005	0.01	0.015	0.005	0.01	0.015
$\bar{i}$	6.19	12.71	19.23	6.86	15.22	23.75	9.53	23.41	36.62
$\bar{i}^h$	5.78	11.56	16	6.67	14.22	20.44	10.22	21.78	28
$\bar{e} + \bar{q}$	1.82	5.45	9.09	3.63	10	20	7.27	28.18	69.09

4 Preliminary Conclusions and Ongoing Work

In this paper we examined a two-sector model with human capital which can be beneficial for two reasons. As is standard in the economics literature, human capital raises the productivity of labour and hence increases output. The additional benefit arises when we allow for model uncertainty. In particular, we explored the idea that knowledge-generating activities can also help the economic agent to learn about the economic environment and thus reduce model uncertainty. This reflects the assumption that learning in the society in not only productivity-motivated, but also information-motivated.

When a robust decision maker fears that the model he uses to make his choices is misspecified, he will, if given the opportunity, use some of his resources to decrease model uncertainty. This gives rise to a type of precautionary behavior, which we call "precautionary learning", which is beneficial to the economic agent. Since the robust decision maker is more exposed to model uncertainty when economic uncertainty is higher, such precautionary learning activities, which include information acquisition and productive knowledge, will increase when economic volatility is higher. This prediction is consistent with historical evidence that suggests that important scientific breakthroughs have occurred in turbulent socioeconomic times. Moreover, since economic downturns are generally associated with increased uncertainty, the well document countercyclical link between recessions and education in its broadest sense is also captured by the model. In ongoing work we are attempting to more thoroughly assess these sorts of predictions against the data for the US and UK.

References

[1] Angelopoulos, K., J. Malley and A. Philippopoulos (2008): "Macroeconomic effects of public education expenditure," CESifo Economic Studies, 54, 471-498.

[2] 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.

[3] Barillas, F., Hansen, L. and T. Sargent (2008); "Doubts or variability?," Working paper, University of Chicago and New York University.

[4] Basar, T. and P. Bernhard (1995): "H-Infinity-Optiml Control and Related Minimax Design Problems: A Dynamic Game Approach" (Systems and Control: Foundations and Applications), Birkhauser (2nd ed).

[5] DeJong, D. and B. Ingram (2001): "The cyclical behavior of skill acquisition," Review of Economic Dynamics, 4, 536-561.

[6] Hansen, L. and T. Sargent (1995): "Discounted linear exponential quadratic Gaussian control," IEEE Transactions on Automatic Control, 40, 968-971.

[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., Sargent, T. and T. Tallarini (1999): "Robust permanent income and pricing," Review of Economic Studies, 66, 873-907.

[10] Hansen, L., Sargent T., G. Turmuhambetova and N. Williams (2006): "Robust control, min-max expected utility and model misspecification", Journal of Economic Theory, 128, 45-90.

[11] Leland, H. (1968): "Saving and Uncertainty: The Precautionary demand for saving," Quarterly Journal of Economics, 82, 465–473.

[12] Ljungqvist, L. and T. Sargent (2000): "Recursive Macroeconomic Theory", The MIT Press, Cambridge, MA.

[13] 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.

[14] 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.

[15] 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.

[16] King, R., Plosser, C. and Rebelo, S., (1988): "Production, growth and business cycles: I. The basic neoclassical model," Journal of Monetary Economics, 21, 195-232.

[17] King, R., Plosser, C. and Rebelo, S., (1988): "Production, growth and business cycles: II. New directions," Journal of Monetary Economics, 21, 309-341.

[18] Kydland F. (1995): "Business cycles and aggregate labor market fluctuations", in Frontiers of Business Cycle Research, edited by T. Cooley, Princeton University Press.

[19] Perli, R. and P. Sakellaris (1998): "Human capital formation and business cycle persistence," Journal of Monetary Economics, 42, 67-92.

[20] Sandmo, A. (1970): "The effect of uncertainty on saving decisions," The Review of Economic Studies, 37, 353-360.

[21] Whittle P. (1990): "Risk Sensitive Optimal Control", New York, Wiley.

5 Appendices

5.1 Appendix A

To obtain $(2a)$, note that the objective function in $(1a)$ can be written as:

$$E_0 \sum_{t=0}^{\infty} (\beta)^t \frac{\left[(c_t Z_t)^{\mu} (l_t)^{(1-\mu)}\right]^{1-\varsigma}}{1-\varsigma}, \text{ or}$$
$$E_0 \sum_{t=0}^{\infty} (\beta)^t (Z_0 g^t)^{\mu(1-\varsigma)} \frac{\left[(c_t)^{\mu} (l_t)^{(1-\mu)}\right]^{1-\varsigma}}{1-\varsigma}. \tag{a1}$$

Normalizing $Z_0 = 1$ gives:

$$E_0 \sum_{t=0}^{\infty} (\beta g^{\mu(1-\varsigma)})^t \frac{\left[(c_t)^{\mu} (l_t)^{(1-\mu)}\right]^{1-\varsigma}}{1-\varsigma}. \tag{a2}$$

5.2 Appendix B

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

$$\sum_{t=0}^{\infty} \breve{\beta}^t \frac{\left[\left(a_t (k_t)^{\alpha} (h_t)^{\zeta} (n_t)^{1-\alpha-\zeta} - i_t - \Omega_t^{\frac{1}{1-\gamma}}\right)^{\mu} (1 - n_t - e_t)^{(1-\mu)}\right]^{1-\varsigma}}{1-\varsigma} \tag{b1}$$

subject to:

$$g k_{t+1} = (1 - \delta^k) k_t + i_t \tag{b2}$$
$$g h_{t+1} = (1 - \delta^h) h_t + m_t \tag{b3}$$
$$\widetilde{a}_{t+1} = (1-\rho)\widetilde{a} + \rho \widetilde{a}_t + \sigma w_{t+1} \tag{b4}$$

where $a_t = \exp(\widetilde{a}_t)$; $\Omega_t = \frac{m_t}{\kappa (e_t h_t)^{\gamma}}$; and the paths $\{w_{t+1}, \widetilde{a}_{t+1}\}_{t=0}^{\infty}$ are chosen by a malevolent agent[26] with the objective to minimize the agent's lifetime utility, subject to the above constraints and

$$R(w) \leq \eta_0. \tag{b5}$$

[26] Note that by choosing w_{t+1}, the malevolent agent is effectively choosing $\widetilde{a}_{t+1}$ as well.

Using the *multiplier problem* representation, we can write this extremization problem as:[27]

$$\max_{\{i_t,m_t,n_t,e_t,k_{t+1},h_{t+1}\}_{t=0}^{\infty}} \min_{\{w_{t+1},\widetilde{a}_{t+1}\}_{t=0}^{\infty}} L =$$

$$\sum_{t=0}^{\infty}\left\{\begin{array}{c} \breve{\beta}^t \frac{\left[\left(a_t(k_t)^{\alpha}(h_t)^{\zeta}(n_t)^{1-\alpha-\zeta}-i_t-\Omega_t^{\frac{1}{1-\gamma}}\right)^{\mu}(1-n_t-e_t)^{(1-\mu)}\right]^{1-\varsigma}}{1-\varsigma} \\ +\lambda_t^K\breve{\beta}^t\left[gk_{t+1}-(1-\delta)k_t-i_t\right] \\ +\lambda_t^H\breve{\beta}^t\left[gh_{t+1}-(1-\delta^h)h_t-m_t\right] \\ +\psi_t\breve{\beta}^t\left[\widetilde{a}_{t+1}-(1-\rho)\widetilde{a}-\rho\widetilde{a}_t-\sigma w_{t+1}\right] \\ +\breve{\beta}^{t+1}\theta w_{t+1}^2 \end{array}\right\}. \qquad \text{(b6)}$$

To solve this problem, we use the *Bellman-Isaacs* condition that holds for this game, which indicates that the first order conditions for *extremising*, 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} and $\widetilde{a}_{t+1}$ respectively, match those of an ordinary optimization problem with joint control processes $\{i_t, m_t, n_t, e_t, k_{t+1}, h_{t+1}, w_{t+1}, \widetilde{a}_{t+1}\}$. 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) and Basar and Bernhard (1995, ch. 2). The steady-state solution of this set of non-linear equations gives the following results (for the calibration reported in Table 1):

Table 7: Steady-state

	Standard Model	Robust Model		
σ	-	0.005	0.01	0.015
c/y	0.682	0.682	0.682	0.682
i/y	0.231	0.231	0.231	0.231
i^h/y	0.087	0.087	0.087	0.087
k/y	3.122	3.122	3.122	3.122
h/y	12.96	12.96	12.96	12.96
e	0.110	0.110	0.110	0.110
n	0.201	0.201	0.201	0.201
$\widetilde{a}$	-0.20	-0.207	-0.214	-0.220
w	n.a.	-0.014	-0.014	-0.013

[27]See the discussion in section 2.5 and Appendix D, regarding the equivalence between the constraint and the multiplier problem, which, for *fixed* θ, implies that the solution of the two problems is equivalent provided that the fixed value of θ implies a solution for w that satisfies $\sum_{t=0}^{\infty}\beta^{t+1}w_{t+1}^2=\eta_0$.

Note that the values of the ratios and endogenous variables reported in this Table match the data well and are consistent with other values used in the literature. Further note that we have also approximated $(2a)$ around a steady-state without model uncertainty. In this case, we can calculate the steady-state by solving the system of deterministic non-linear first-order conditions from $(2a) - (2g)$ where both $\widetilde{\varepsilon}_t$ and $w_{t+1}$are set to zero. For values of θ that imply a 10% error detection probability, the solution for the endogenous variables is very similar with and without model uncertainty in the steady state. In particular, consumption, investment and output are all higher in the robust solution, but the solution for the great ratios and the time shares reported in Table 7 does not change (this is because the malevolent agent's choices affect the long run value of TFP, which affects consumption and investment proportionately to output and does not affect the time shares).

5.3 Appendix C

Using the probability densities for the approximating and the distorted model[28]

$$f_a(\widetilde{a}_{t+1} \mid \widetilde{a}_t) \sim N((1-\rho)\widetilde{a} + \rho\widetilde{a}_t, \sigma^2) \tag{c1}$$

$$f(\widetilde{a}_{t+1} \mid \widetilde{a}_t) \sim N((1-\rho)\widetilde{a} + \rho\widetilde{a}_t + \sigma w_{t+1}, \sigma^2) \tag{c2}$$

we can calculate the log-likelihood ratio in

$$I(f_a, f)(\widetilde{a}) = \int \log\left(\frac{f(\widetilde{a}_{t+1} \mid \widetilde{a}_t)}{f_a(\widetilde{a}_{t+1} \mid \widetilde{a}_t}\right) f(\widetilde{a}_{t+1} \mid \widetilde{a}_t) d\widetilde{a}_{t+1}$$

as:

$$\log\left(\frac{f(\widetilde{a}_{t+1} \mid \widetilde{a}_t)}{f_a(\widetilde{a}_{t+1} \mid \widetilde{a}_t}\right) =$$

$$\log\left(f(\widetilde{a}_{t+1} \mid \widetilde{a}_t)\right) - \log(f_a(\widetilde{a}_{t+1} \mid \widetilde{a}_t), \text{ or}$$

[28] Note that $f(\widetilde{a}_{t+1} \mid \widetilde{a}_t)$ allows for misspecifications that occur only as a perturbation to the conditional mean of the innovation to the state, $\widetilde{a}_{t+1}$ in $f_a(\widetilde{a}_{t+1} \mid \widetilde{a}_t)$ and leaves the conditional volatility of the shock, as parametrized by σ^2, unchanged. As Hansen and Sargent (2007) show (see ch. 2.6, 3 and 7), as long as we stay within the LQ framework with Gaussian distributions for the approximating model $f_a(\widetilde{a}_{t+1} \mid \widetilde{a}_t)$, allowing for a more general class of misspecifications 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 hurt the maximizing agent through the constant term, d, in the value function, $v(x_0)$, 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.

$$-\frac{1}{2}\log\sqrt{2\pi\sigma^2}-\frac{1}{2\sigma^2}\left(\widetilde{a}_{t+1}-(1-\rho)\widetilde{a}-\rho\widetilde{a}_t-\sigma w_{t+1}\right)^2+$$
$$+\frac{1}{2}\log\sqrt{2\pi\sigma^2}+\frac{1}{2\sigma^2}\left(\widetilde{a}_{t+1}-(1-\rho)\widetilde{a}-\rho\widetilde{a}_t\right)^2,\text{ or}$$
$$-\frac{1}{2\sigma^2}\left(\widetilde{a}_{t+1}-(1-\rho)\widetilde{a}-\rho\widetilde{a}_t\right)^2-\frac{1}{2}w_{t+1}^2+\frac{1}{\sigma}w_{t+1}\left(\widetilde{a}_{t+1}-(1-\rho)\widetilde{a}-\rho\widetilde{a}_t\right)$$
$$+\frac{1}{2\sigma^2}\left(\widetilde{a}_{t+1}-(1-\rho)\widetilde{a}-\rho\widetilde{a}_t\right)^2,\text{ or}$$
$$-\frac{1}{2}w_{t+1}^2+\frac{1}{\sigma}w_{t+1}\left(\widetilde{a}_{t+1}-(1-\rho)\widetilde{a}-\rho\widetilde{a}_t\right),\text{ or}$$
$$-\frac{1}{2}w_{t+1}^2+\frac{1}{\sigma}w_{t+1}\left(\widetilde{a}_{t+1}-(1-\rho)\widetilde{a}-\rho\widetilde{a}_t-\sigma w_{t+1}+\sigma w_{t+1}\right),\text{ or}$$
$$-\frac{1}{2}w_{t+1}^2+\frac{1}{\sigma}w_{t+1}\left(\widetilde{a}_{t+1}-(1-\rho)\widetilde{a}-\rho\widetilde{a}_t-\sigma w_{t+1}\right)+w_{t+1}^2,\text{ or}$$
$$\frac{1}{2}w_{t+1}^2+\frac{1}{\sigma}w_{t+1}\left(\widetilde{a}_{t+1}-(1-\rho)\widetilde{a}-\rho\widetilde{a}_t-\sigma w_{t+1}\right).$$

Hence, the conditional relative entropy in $I(f_a,f)(x)$ is given as:

$$\begin{aligned}
I(f_a,f)(x) &= \\
&= \int\frac{1}{2}w_{t+1}^2+\frac{1}{\sigma}w_{t+1}\left(\widetilde{a}_{t+1}-(1-\rho)\widetilde{a}-\rho\widetilde{a}_t-\sigma w_{t+1}\right) \\
\times f(\widetilde{a}_{t+1} &\mid \widetilde{a}_t)d\widetilde{a}_{t+1},\text{ or} \\
&= E\left[\frac{1}{2}w_{t+1}^2+\frac{1}{\sigma}w_{t+1}\left(\widetilde{a}_{t+1}-(1-\rho)\widetilde{a}-\rho\widetilde{a}_t-\sigma w_{t+1}\right)\right].
\end{aligned}$$

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

$$\begin{aligned}
I(f_a,f)(x) &= \frac{1}{2}w_{t+1}^2+\frac{1}{\sigma}w_{t+1}E\left[\left(\widetilde{a}_{t+1}-(1-\rho)\widetilde{a}-\rho\widetilde{a}_t-\sigma w_{t+1}\right)\right],\text{ or} \\
&= \frac{1}{2}w_{t+1}^2+\frac{1}{\sigma}w_{t+1}E\left[\widetilde{a}_{t+1}-E(\widetilde{a}_{t+1})\right] \\
&= \frac{1}{2}w_{t+1}^2 \qquad \text{(c3)}
\end{aligned}$$

5.4 Appendix D

The constraint problem for the decision maker that fears model uncertainty can be stated as:

$$\max_{\{u_t\}_{t=0}^{\infty}}\min_{\{w_{t+1}\}_{t=0}^{\infty}} E_0\sum_{t=0}^{\infty}\breve{\beta}^t\left\{x_t'Qx_t+u_t'Ru_t+2x_t'Wu_t\right\} \qquad \text{(d1)}$$

given η_0 satisfying $\overline{\eta} > \eta_0 \geq 0$, where the extremization is subject to the perturbed model ($7b$) and the entropy constraint ($11b$). Also note that, E_0, conditioned on x_0, is also evaluated with respect to the perturbed model. Here $\overline{\eta}$ measures the largest set of perturbations against which it is possible to attain robustness. Provided we can relate θ to $\overline{\eta}$ appropriately, the outcomes (i.e. the same policy functions for u_t and w_{t+1} and the same value of the game, $v(x_0)$) for the constraint and multiplier problems are equivalent. The latter problem takes the form:

$$\max_{\{u_t\}_{t=0}^{\infty}} \min_{\{w_{t+1}\}_{t=0}^{\infty}} E_0 \sum_{t=0}^{\infty} \breve{\beta}^t \left\{ x_t^{'} Q x_t + u_t^{'} R u_t + 2x_t^{'} W u_t + \breve{\beta}\theta w_{t+1}^{\prime} w_{t+1} \right\} \qquad \text{(d2)}$$

given $\underline{\theta} \leq \theta \leq +\infty$, where the extremization is subject to the perturbed model ($7b$). Here θ is a penalty parameter restraining the minimizing choice of the w_{t+1} sequence.[29]

We next write the Lagrangian for the constraint problem, ($d1$), incorporating the entropy constraint:

$$\begin{aligned} \sup_{\theta \in [\underline{\theta}, +\infty]} \max_{\{u_t\}_{t=0}^{\infty}} \min_{\{w_{t+1}\}_{t=0}^{\infty}} \Lambda &= E_0 \{ \sum_{t=0}^{\infty} \breve{\beta}^t \left(x_t^{'} Q x_t + u_t^{'} R u_t + 2x_t^{'} W u_t \right) \\ &\quad + \theta \left(\sum_{t=0}^{\infty} \breve{\beta}^{t+1} w_{t+1}^{\prime} w_{t+1} - \eta_0 \right) \}. \qquad \text{(d3)} \end{aligned}$$

Note that this and all optimization problems discussed below are subject to the perturbed model ($7b$). Provided that θ is constrained to be $[\underline{\theta}, +\infty]$, so that the objective is convex with respect to w_{t+1} and hence the requirements of the Lagrange Theorem are met, we can interpret θ as the Lagrange multiplier attached to the entropy constraint in the Lagrangian for the constraint problem. Since the order of maximization is inconsequential to equilibrium outcomes we can postpone maximization over θ until the last step and first study:

$$\begin{aligned} \max_{\{u_t\}_{t=0}^{\infty}} \min_{\{w_{t+1}\}_{t=0}^{\infty}} \Lambda &= E_0 \{ \sum_{t=0}^{\infty} \breve{\beta}^t \left(x_t^{'} Q x_t + u_t^{'} R u_t + 2x_t^{'} W u_t \right) \\ &\quad + \theta \left(\sum_{t=0}^{\infty} \breve{\beta}^{t+1} w_{t+1}^{\prime} w_{t+1} - \eta_0 \right) \} \qquad \text{(d4)} \end{aligned}$$

[29] This is discussed in: (i) a static framework in Hansen and Sargent (2008, ch. 6); (ii) a LQ framework with a general class of distortions in Hansen and Sargent (2008, ch. 7.8); and (iii) a continuous time framework in Hansen *et al.* (2006, see especially pp. 58 – 62).

for a given θ. When θ is fixed, η_0 is a constant that is irrelevant to the optimal choices of the controls (u_t and w_{t+1}) and thus it can be omitted from the optimization.[30] Therefore, we can re-write the problem of uncovering the optimal paths for u_t and w_{t+1} as:

$$\begin{aligned} \max_{\{u_t\}_{t=0}^{\infty}} \min_{\{w_{t+1}\}_{t=0}^{\infty}} \Lambda &= E_0 \sum_{t=0}^{\infty} \check{\beta}^t \{x_t^{'} Q x_t + u_t^{'} R u_t + 2x_t^{'} W u_t \\ &\quad + \theta \check{\beta} w_{t+1}^{\prime} w_{t+1}\} \end{aligned} \tag{d5}$$

which is the multiplier problem. For the two problems (constraint and multiplier) then to have the same solutions, θ in the multiplier problem must related to η_0 in the constraint problem as the Lagrange multiplier associated with the entropy constraint in the latter. Noticing that $\Lambda = \Lambda(\theta) - \eta_0 \theta$, the relationship between η_0 and θ must be given by:

$$\begin{aligned} \frac{\partial \Lambda}{\partial \theta} &= 0 \\ &\Rightarrow \frac{\partial \Lambda(\theta)}{\partial \theta} = \eta_0 \\ &\Rightarrow E_0 \sum_{t=0}^{\infty} \check{\beta}^{t+1} w_{t+1}^{\prime} w_{t+1} = \eta_0 \end{aligned} \tag{d6}$$

where the solution for w_{t+1} will be a function of θ. Therefore, a known parameter η_0 will map via ($d6$) into a particular parameter for θ and in this case the two problems, constraint and multiplier will have equivalent outcomes (i.e. for the policy functions and value of the game). Hence, if we knew η_0, we could solve the multiplier problem for different values of θ, calculate $\sum_{t=0}^{\infty} \check{\beta}^{t+1} w_{t+1}^{\prime} w_{t+1}$ for each case and choose the value of θ that makes $\sum_{t=0}^{\infty} \check{\beta}^{t+1} w_{t+1}^{\prime} w_{t+1} = \eta_0$. Of course, in practice, we work the other way around, as we calibrate the parameter θ (see main text sub-section 2.6.2), solve the *maxmin* problem and then we can use $\sum_{t=0}^{\infty} \check{\beta}^{t+1} w_{t+1}^{\prime} w_{t+1}$ to calculate the entropy ball that gives the set of models against which the maximizing agent seeks robust actions. In particular, this will be given by:

$$W(\eta_0) = \left\{ w \in W : E_0 \sum_{t=0}^{\infty} \check{\beta}^{t+1} w_{t+1}^{\prime} w_{t+1} \leq \eta_0 \right\}. \tag{d7}$$

Hence, the calibration of θ maps to the calibration of η_0.

[30] Note that it can be included if we want to optimize over θ in the last step.

5.5 Appendix E

The multiplier problem we solve is:

$$\max_{\{u_t\}_{t=0}^{\infty}} \min_{\{w_{t+1}\}_{t=0}^{\infty}} E_0 \sum_{t=0}^{\infty} \breve{\beta}^t \left\{ \left[\begin{pmatrix} u_t \\ x_t \end{pmatrix}' \begin{pmatrix} R & W' \\ W & Q \end{pmatrix} \begin{pmatrix} u_t \\ x_t \end{pmatrix} \right] + \breve{\beta}\theta w'_{t+1} w_{t+1} \right\} \quad \text{(e1)}$$

subject to $\underline{\theta} \leq \theta \leq +\infty$ and

$$x_{t+1} = Ax_t + Bu_t + C(\epsilon_{t+1} + w_{t+1}). \quad \text{(e2)}$$

Alternatively, we can write the multiplier problem as:

$$\max_{\{u_t\}_{t=0}^{\infty}} \min_{\{w_{t+1}\}_{t=0}^{\infty}} E_0 \sum_{t=0}^{\infty} \breve{\beta}^t \left\{ x_t' Q x_t + u_t' R u_t + 2x_t' W u_t + \breve{\beta}\theta w'_{t+1} w_{t+1} \right\} \quad \text{(e3)}$$

subject to $\underline{\theta} \leq \theta \leq +\infty$ and

$$x_{t+1} = Ax_t + Bu_t + C(\epsilon_{t+1} + w_{t+1}). \quad \text{(e4)}$$

We require R to be negative definite and Q to be negative semi-definite. The Bellman equation for the multiplier problem is:

$$v(x) = \max_{u_t} \min_{w_{t+1}} \left\{ x_t' R x_t + u_t' Q u_t + 2x_t' W u_t + \breve{\beta} E\left[\theta w'_{t+1} w_{t+1} + v(\widetilde{x}) \right] \right\} \quad \text{(e5)}$$

where $v(x)$ is the value function, x is current period's state and $\widetilde{x}$ is next period's state. We have essentially exchanged the original problem of finding an infinite sequence of controls $\{u_t\}_{t=0}^{\infty}$ and $\{w_{t+1}\}_{t=0}^{\infty}$ that solve the multiplier problem, for the problem of finding the optimal value function $v(x)$ and a policy function that solves the continuum of extremization problems in the right hand side of $(e5)$. Hence, we solve the multiplier problem recursively, assuming that the maximizing agent plays first, as the order *max min* indicates.

We guess that the value function for the problem is quadratic:

$$v(x) = x' P x + d \quad \text{(e6)}$$

where P is a symmetric $(\widetilde{h} \times \widetilde{h})$ matrix and d is a constant to be determined. Substituting the guess function in $(e6)$ in the Bellman equation gives:

$$v(x) = \max_{u_t} \min_{w_{t+1}} \left\{ x_t' R x_t + u_t' Q u_t + 2x_t' W u_t + \breve{\beta} E\left[\theta w'_{t+1} w_{t+1} + \widetilde{x}' P \widetilde{x} + d \right] \right\}, \text{ or}$$

$$v(x) = \max_{u_t} \min_{w_{t+1}} \left\{ x_t' R x_t + u_t' Q u_t + 2x_t' W u_t + \breve{\beta} E\left[\theta w'_{t+1} w_{t+1} + \widetilde{x}' P \widetilde{x} \right] + \breve{\beta} d \right\}.$$

Then, use the transition law in $(e4)$ to eliminate next period's state:

$$
\begin{aligned}
v(x) &= \max_{u}\min_{w}\{x'Rx + u'Qu + 2x_t^{'}Wu_t + \\
&\quad \breve{\beta}E[\theta w'w + (Ax + Bu + C\epsilon + Cw)' \\
&\quad \times P\,(Ax + Bu + C\epsilon + Cw)] + \breve{\beta}d\}.
\end{aligned} \tag{e7}
$$

Calculating expectations, equation $(e7)$ in turn implies:[31]

$$
\begin{aligned}
v(x_t) &= x'Px + d = \max_{u_t}\min_{w_{t+1}}\{x_t^{'}Qx_t + u_t^{'}Ru_t + 2x_t^{'}Wu_t + \breve{\beta}\theta w_{t+1}^{\prime}w_{t+1} \\
&\quad \breve{\beta}x_t^{\prime}A^{\prime}PAx_t + \breve{\beta}x_t^{\prime}A^{\prime}PBu_t + \breve{\beta}x_t^{\prime}A^{\prime}PCw_{t+1} + \breve{\beta}u_t^{\prime}B^{\prime}PAx_t \\
&\quad +\breve{\beta}u_t^{\prime}B^{\prime}PBu_t + \breve{\beta}u_t^{\prime}B^{\prime}PCw_{t+1} + \breve{\beta}w_{t+1}^{\prime}C^{\prime}PAx_t + \\
&\quad \breve{\beta}w_{t+1}^{\prime}C^{\prime}PBu_t + \breve{\beta}w_{t+1}^{\prime}C^{\prime}PCw_{t+1} + \breve{\beta}E\left(\epsilon^{\prime}C^{\prime}PC\epsilon\right) + \breve{\beta}d\}.
\end{aligned} \tag{e8}
$$

Equating the constant terms we have that:

$$d = \breve{\beta}E(\epsilon'C'PC\epsilon) + \breve{\beta}d$$

Using a result from linear algebra that[32]: $E(\epsilon'C'PC\epsilon) = trE(\epsilon'C'PC\epsilon) = trE(C'PC\epsilon\epsilon') = trC'PCE(\epsilon\epsilon') = trC'PC = trPCC'$, this implies that:

$$d = \breve{\beta}(1-\breve{\beta})^{-1}trPCC'.$$

We solve the problem by backward induction. This means that we solve first for the inner minimization problem (as the minimizing player moves second) by taking the choices of the maximizing player u_t as given. We can write the inner minimization problem as:

$$
\begin{aligned}
J &= \min_{w_{t+1}}\{\breve{\beta}\theta w_{t+1}^{\prime}w_{t+1} + \breve{\beta}x_t^{\prime}A^{\prime}PAx_t + \breve{\beta}x_t^{\prime}A^{\prime}PBu_t + \breve{\beta}x_t^{\prime}A^{\prime}PCw_{t+1} + \\
&\quad \breve{\beta}u_t^{\prime}B^{\prime}PAx_t + \breve{\beta}u_t^{\prime}B^{\prime}PBu_t + \breve{\beta}u_t^{\prime}B^{\prime}PCw_{t+1} + \breve{\beta}w_{t+1}^{\prime}C^{\prime}PAx_t + \\
&\quad \breve{\beta}w_{t+1}^{\prime}C^{\prime}PBu_t + \breve{\beta}w_{t+1}^{\prime}C^{\prime}PCw_{t+1}\}
\end{aligned} \tag{e9}
$$

which can be simplified to

$$
\begin{aligned}
J &= \min_{w_{t+1}}\{w_{t+1}^{\prime}\left(\breve{\beta}\theta I\right)w_{t+1} + \breve{\beta}(Ax_t + Bu_t + Cw_{t+1})' \\
&\quad \times P(Ax_t + Bu_t + Cw_{t+1})\}
\end{aligned} \tag{e10}
$$

[31] Note that the remaining terms involving ϵ are equal to zero, given that they are linear in ϵ and $E\epsilon|x = 0$.

[32] The first equality holds because $E(\epsilon'C'PC\epsilon)$ is a scalar. The second because of the cyclic property of the trace, i.e. $tr(ABC) = tr(CAB) = tr(BCA)$, where we treat $C'PC$ as a single matrix. The third because $C'PC$ is non-stochastic. The fourth because $E(\epsilon\epsilon') = 1$. The fifth again because of the cyclic property of the trace.

where J is the value function for the inner minimization problem as defined in $(e10)$, so that it includes the elements of the original problem under the control of the minimizing agent.

The first order condition necessary for the minimum problem on the right side of $(e10)$ requires:[33]

$$\begin{aligned} 0 &= 2\check{\beta}\theta w_{t+1} + \check{\beta}C'P'Ax_t + \check{\beta}C'P'Bu_t + \check{\beta}C'PAx_t \\ &+ \check{\beta}C'PBu_t + \check{\beta}(C'PC + C'P'C)w_{t+1}. \end{aligned} \quad \text{(e11)}$$

Solving for w_{t+1} gives:

$$w_{t+1} = -(\theta I + C'PC)^{-1}C'P\left[Ax_t + Bu_t\right]. \quad \text{(e12)}$$

At time period t, the minimizing agent is choosing w_{t+1} taking as given the period t realization of the state variable x_t and the period t maximizing agent's choice, u_t. Hence, the state at period t for the minimizing agent is $Ax_t + Bu_t$. Equation $(e12)$ provides the best response of the minimizing agent to the choices of the maximizing agent, given by u_t, and the state of the game, as summarized by x_t. Alternatively, it gives the optimal choices for the minimizing agent conditional on its state, $Ax_t + Bu_t$.

Assume that the solution for the maximizing u_t takes the form $u_t = -Fx_t$, where F is a matrix to be determined. Then, the Markov perfect equilibrium solution for the worst case shock is:

$$w_{t+1} = Kx_t \quad \text{(e13)}$$

where $K = -(\theta I + C'PC)^{-1}C'P\left[A - BF\right]$.

The minimized value from this problem is obtained by substituting the minimizing w_{t+1} from $(e12)$ into the objective $(e10)$ from which we obtain:

$$\begin{aligned} J &= \check{\beta}(Ax_t + Bu_t)'(P - PC(\theta I + C'PC)^{-1}C'P) \\ &\times(Ax_t + Bu_t) \end{aligned} \quad \text{(e14)}$$

where $D(P)$, which gives the minimized value of J, satisfies:

$$D(P) = P - PC(\theta I + C'PC)^{-1}C'P. \quad \text{(e15)}$$

In order to proceed with the maximizing agent's choices, note that $D(P)$ is an operator that defines via $(e15)$ the next (i.e. $t+1$) period's value function for the maximizing agent, given the choices of the minimizing agent. The expression $Ax_t + Bu_t$ is essentially the $t + 1$ period state for the maximizing

[33]Note that below we use the fact that P is symmetric.

agent once the reaction function of the minimizing agent is accounted for by using the operator $D(P)$. Hence, J can be used to incorporate the behavior of the malevolent agent into the maximizing agent's problem. This implies that the Bellman equation for the maximizing agent can be written as:

$$\begin{aligned} v(x_t) &= \max_{u_t}\{x_t^{'}Qx_t + u_t^{'}Ru_t + 2x_t^{'}Wu_t + \check{\beta}x_t^{'}A'D(P)Ax_t + \\ &\quad \check{\beta}u_t^{'}B'D(P)Ax_t + \check{\beta}x_t^{'}A'D(P)Bu_t + \check{\beta}u_t^{'}B'D(P)Bu_t\}. \end{aligned} \tag{e16}$$

Essentially, in $(e16)$ we have substituted out the optimal behavior (i.e. the reaction function) of the minimizing agent in the original *maxmin* problem. We can now proceed to the second step of the solution process implied by the order of movements, which is to find the optimal policy for the maximizing agent. Note that using the $D(P)$ operator, the maximizing problem has a similar structure to the "standard" LQ problem where the agent trusts the model. The only difference is the $D(P)$ operator replaces the P matrix on the right hand side. This exactly incorporates the distortion in the valuation of the future value of the game, resulting from mistrusting the model. Note again that when $\theta \to +\infty$, $D(P) \to P$ for any matrix P and the Bellman equation for the maximizing agent collapses exactly to that of the "standard" LQ problem.

The first order condition necessary for the maximum problem on the right side of $(e16)$ requires:[34]

$$(R + \check{\beta}B'D(P)B)u_t = -\left(\check{\beta}B'D(P)A + W'\right)x_t. \tag{e17}$$

Equation $(e17)$ implies the following rule for u_t:

$$u_t = -(R + \check{\beta}B'D(P)B)^{-1}\left(\check{\beta}B'D(P)A + W'\right)x_t \tag{e18}$$

which can be written as:

$$u_t = -Fx_t \tag{e19}$$

where

$$F = (R + \check{\beta}B'D(P)B)^{-1}\left(\check{\beta}B'D(P)A + W'\right).$$

The next step in the solution is to calculate the value function $v(x_t)$ by substituting the optimal solution for u_t into the right side of the Bellman equation, then substituting the guess function in the left side and equating coefficients:

$$\begin{aligned} x_t'Px_t &= x_t'Qx_t + (Fx_t)'RFx_t - x_t'WFx_t - x_t'F'W'x_t + \\ &\quad \check{\beta}x_t'A'D(P)Ax_t - \check{\beta}x_t'A'D(P)BFx_t - \\ &\quad \check{\beta}(Fx_t)'B'D(P)Ax_t + \check{\beta}(Fx_t)'B'D(P)BFx_t \end{aligned} \tag{e20}$$

[34] Note that below we use the fact that R, P and $D(P)$ are symmetric.

which implies that:

$$\begin{aligned}
x_t' P x_t &= x_t' Q x_t + x_t' F' R F x_t - x_t' W F x_t - x_t' F' W' x_t + \\
&\quad x_t' \check{\beta} A' D(P) A x_t - x_t' \check{\beta} A' D(P) B F x_t - \\
&\quad x_t' \check{\beta} F' B' D(P) A x_t + x_t' \check{\beta} F' B' D(P) B F x_t. \qquad \text{(e21)}
\end{aligned}$$

If we next pre-multiply by $(x_t')^{-1}$, post-multiply by x_t^{-1} and re-arrange we obtain

$$\begin{aligned}
P &= Q + \check{\beta} A' D(P) A - \left(\check{\beta} A' D(P) B + W\right) \left(R + \check{\beta} B' D(P) B\right)^{-1} \\
&\quad \times \left(\check{\beta} B' D(P) A + W'\right) \qquad \text{(e22)}
\end{aligned}$$

where it can be shown that $F'RF - F'W' - \check{\beta} F' B' D(P) A + \check{\beta} F' B' D(P) B F = 0$.

The system of equations $(e15)$ and $(e22)$ can be solved by iterations. Once it has converged to a solution for P, denoted P^* in the main text, we can calculate F and hence the policy functions $u_t = -F x_t$ and $w_{t+1} = K x_t$ and the state evolution $x_{t+1} = (A - BF + CK) x_t$. We should also check whether the state process is stationary, which is guaranteed if the eigenvalues of $(A - BF + CK)$ are all less than one in absolute value (except for one, corresponding to the constant term, which will be exactly equal to one). To summarize, we find the symmetric matrix P by iterating on the composite operator:

$$\begin{aligned}
D(P) &= P - PC(\theta I + C'PC)^{-1} C' P \\
P &= Q + \check{\beta} A' D(P) A - \left(\check{\beta} A' D(P) B + W\right) \times \\
&\quad \left(R + \check{\beta} B' D(P) B\right)^{-1} \left(\check{\beta} B' D(P) A + W'\right)
\end{aligned}$$

and then we calculate F and K as:

$$F = \left(R + \check{\beta} B' D(P) B\right)^{-1} \left(\check{\beta} B' D(P) A + W'\right)$$

$$K = -(\theta I + C'PC)^{-1} C' P \left[A - BF\right].$$

Finally, the value of the game or welfare is given by:

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

where $d = \check{\beta}(1 - \check{\beta})^{-1} tr\left(P^* C C'\right)$; and P^* is the value of P at convergence. Note that the expressions for $D(P)$, P, F, K and $v(x_0)$ are those reported under equations $(13a)$ and $(13b)$ in the main text.

5.6 Appendix F

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

$$\sum_{t=0}^{\infty} \breve{\beta}^t \frac{\left[\left(a_t (k_t)^{\alpha} (h_t)^{\zeta} (n_t)^{1-\alpha-\zeta} - i_t - \Omega_t^{\frac{1}{1-\gamma}}\right)^{\mu} (1 - n_t - e_t - q_t)^{(1-\mu)}\right]^{1-\varsigma}}{1-\varsigma} \tag{f1}$$

where $a_t = \exp(\widetilde{a}_t)$ and $\Omega_t = \frac{m_t}{\kappa(e_t h_t)^{\gamma}}$, subject to:

$$gk_{t+1} = (1-\delta^k)k_t + i_t \tag{f2}$$

$$gh_{t+1} = (1-\delta^h)h_t + m_t \tag{f3}$$

$$\widetilde{a}_{t+1} = (1-\rho)\widetilde{a} + \rho\widetilde{a}_t + \sigma w_{t+1} \tag{f4}$$

$$\sum_{t=0}^{\infty} \breve{\beta}^{t+1} \left(w_{t+1}^2 + \breve{\beta}^{-1} \eta_1 q_t^{\eta_2} h_t^{(1-\eta_2)}\right) = \eta_0 \tag{f5}$$

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

Using the *multiplier problem* representation, we can write this extremization problem as:[35]

$$\max_{\{i_t,m_t,n_t,e_t,q_t,k_{t+1},h_{t+1}\}_{t=0}^{\infty}} \min_{\{w_{t+1},\widetilde{a}_{t+1}\}_{t=0}^{\infty}} L$$

$$= \sum_{t=0}^{\infty} \left\{ \begin{array}{c} \breve{\beta}^t \frac{\left[\left(a_t(k_t)^{\alpha}(h_t)^{\zeta}(n_t)^{1-\alpha-\zeta} - i_t - \Omega_t^{\frac{1}{1-\gamma}}\right)^{\mu} (1-n_t-e_t-q_t)^{(1-\mu)}\right]^{1-\varsigma}}{1-\varsigma} \\ +\lambda_t^K \breve{\beta}^t \left[gk_{t+1} - (1-\delta)k_t - i_t\right] \\ +\lambda_t^H \breve{\beta}^t \left[gh_{t+1} - (1-\delta^h)h_t - m_t\right] \\ +\psi_t \breve{\beta}^t \left[\widetilde{a}_{t+1} - (1-\rho)\widetilde{a} - \rho\widetilde{a}_t - \sigma w_{t+1}\right] \\ +\breve{\beta}^{t+1}\theta \left[w_{t+1}^2 + \breve{\beta}^{-1}\eta_1 q_t^{\eta_2} h_t^{(1-\eta_2)}\right] \end{array} \right\}. \tag{f6}$$

In order to solve this problem, we use the *Bellman-Isaacs* condition and stack the first order conditions of the zero-sum two-player dynamic game above and solve them simultaneously. The steady-state solution of this set of non-linear equations gives the following results for $\eta_1 = 0.1$ and $\eta_2 = 0.5$:

[35] See the discussion in section 3.1 and Appendix D, regarding the equivalence between the constraint and the multiplier problem, which, for *fixed* θ, implies that the solution of the two problems is equivalent provided that the fixed value of θ implies a solution for w and h_t^i that satisfy $\sum_{t=0}^{\infty} \breve{\beta}^{t+1} \left(w_{t+1}^2 + \breve{\beta}^{-1} \eta_1 q_t^{\eta_2} h_t^{(1-\eta_2)}\right) \leq \eta_0$.

Table 8: Steady-state

	Standard Model	Robust Model		
σ	-	0.005	0.01	0.015
c/y	0.682	0.682	0.682	0.682
i/y	0.231	0.231	0.231	0.231
i^h/y	0.087	0.087	0.087	0.087
k/y	3.122	3.122	3.122	3.122
h/y	12.96	12.94	12.87	12.74
e	0.110	0.110	0.110	0.109
n	0.201	0.201	0.200	0.200
$\widetilde{a}$	-0.20	-0.207	-0.214	-0.220
w	n.a.	-0.014	-0.013	-0.013
q	n.a.	0.001	0.003	0.007