Journal of Monetary Economics 59 (2012) 738–750

Contents lists available at SciVerse ScienceDirect

Journal of Monetary Economics

journal homepage: www.elsevier.com/locate/jme

Robust animal spirits

R.M. Bidder [a,*], M.E. Smith [b]

[a] Economic Research Department, Federal Reserve Bank of San Francisco, 101 Market Street, San Francisco, CA 94105, USA
[b] Division of Monetary Affairs, Federal Reserve Board of Governors, 20th Street and Constitution Avenue, N.W., Washington, DC 20551, USA

ARTICLE INFO

Article history:
Received 27 December 2011
Received in revised form
17 October 2012
Accepted 18 October 2012
Available online 29 October 2012

ABSTRACT

In a real business cycle model, an agent's fear of model misspecification interacts with stochastic volatility to induce time varying worst case scenarios. These time varying worst case scenarios capture a notion of animal spirits where the probability distributions used to evaluate decision rules and price assets do not necessarily reflect the fundamental characteristics of the economy. Households entertain a pessimistic view of the world and their pessimism varies with the overall level of volatility in the economy, implying an amplification of the effects of volatility shocks. By using perturbation methods and Monte Carlo techniques we extend the class of models analyzed with robust control methods to include the sort of nonlinear production-based DSGE models that are popular in academic research and policymaking practice.

© 2012 Elsevier B.V. All rights reserved.

1. Introduction

The concept of 'animal spirits' has long proved fascinating to students of human behavior.[1] In modern macroeconomics, they have been closely associated with explanations of business cycles based on spontaneous fluctuations in underlying preferences or movements between different sets of self-confirming beliefs (Farmer, 2010). In this paper, we show that the interaction of stochastic volatility and a fear of model misspecification yields a novel representation of 'animal spirits' as endogenous fluctuations in agents' subjective probabilities and that these fluctuations induce business cycle variation. The particular problem we examine is based on a standard stochastic growth model with certain modifications. The most important of these modifications are that, first, households have multiplier preferences that express a desire for policies that are robust to misspecification (Hansen and Sargent, 2008) and, second, that the innovations to technology growth feature time varying volatility.

In our model, agents form decision rules that are robust to misspecification by envisaging pessimistic distortions to their model that would imply lower growth and elevated volatility. Notably, the pessimism varies with the level of volatility in the sense that when volatility is high (low), the pessimism is more (less) intense. Since changes in volatility are persistent, so too are the swings in sentiment they induce. Thus, the interaction between robust preferences and stochastic volatility induces a form of 'animal spirits'. To connect these 'animal spirits' to actions, we demonstrate that volatility shocks induce declines in consumption, investment, output and hours and that the interaction of stochastic volatility with robustness can help explain a small but nontrivial fraction of business cycle fluctuations. We argue that these effects are

* Corresponding author. Tel.: +1 415 974 2530; fax: +1 415 974 2168.
E-mail addresses: rhys.bidder@sf.frb.org (R.M. Bidder), matthew.e.smith@frb.gov (M.E. Smith).

[1] According to Keynes (1936), '*Most, probably, of our decisions to do something positive, the full consequences of which will be drawn out over many days to come, can only be taken as the result of animal spirits - a spontaneous urge to action rather than inaction, and not as the outcome of a weighted average of quantitative benefits multiplied by quantitative probabilities.*'

0304-3932/$ - see front matter © 2012 Elsevier B.V. All rights reserved.
http://dx.doi.org/10.1016/j.jmoneco.2012.10.017

due to the fact that the volatility shock induces fears of persistent economic weakness, even though these phenomena are not necessarily implied by the data generating process for the economy. Thus the effects of increased risk are amplified by interaction with model uncertainty.

Robust control provides a natural technology for analyzing situations in which agents doubt their model, or fear misspecifications in a form that cannot be expressed by a unique probability distribution over alternatives. Briefly, an agent with robust preferences is endowed with a 'benchmark' model of the economy but fears that it is misspecified. The agent expresses his doubts of his model by considering alternative distributions or 'models' that capture possible misspecifications.[2] In designing his policy to be robust against misspecification, the agent balances the damage that an implicit misspecification would cause him, against the plausibility of the misspecification. The distribution that emerges from this problem can be thought of as a 'worst case model'. By evaluating policy under this particular model the agent can bound its performance and thereby formulate a policy that is robust to misspecification.

The worst case model is expressed in relation to the benchmark by a particular likelihood ratio, following the approach described in Hansen and Sargent (2008). In our environment, it is not possible to compute the worst case distribution exactly and approximations are necessary. Bidder and Smith (2011) show how to approximate the conditional probability density function of the worst case model in general nonlinear, non-Gaussian settings and then design Monte Carlo algorithms to draw from that distribution. In this paper, we build on these techniques to allow the consideration of models with many state variables. Perturbation methods are used to approximate the value function and worst case likelihood ratio associated with the agent's problem. In addition, we approximate the distorted conditional moments of the structural shocks. One can then use these approximated distorted moments to characterize statistical properties of the worst case model's moments, or to improve the sampling algorithms described in Bidder and Smith (2011). This extension is particularly useful for application to the sort of medium scale models used by policymakers, where closed form expressions for the worst case distortions are unavailable and the number of states renders global approximation methods infeasible.

The main contribution of the paper, however, is that we capture 'animal spirits' in the time varying worst case conditional distribution implied by the solution to the agents' robust control problem and show that these fluctuations have a significant impact upon agents' behavior.

2. The model

In this section, we lay out an abstract robust control problem. We then explain the sense in which we use robust control to capture animal spirits and, finally, map a production economy with stochastic volatility into this abstract framework.

2.1. Seeking robustness

A robust agent is endowed with a 'benchmark' model but fears that it is misspecified. He is concerned that the world is actually described by a model that is similar to the benchmark but distorted in some way. The agent expresses his doubts of his model by considering alternative distributions that are distorted versions of the distribution implied by his benchmark model. In order to construct a robust policy the agent considers adverse distributions and balances the damage that an implicit misspecification would cause him, against the plausibility of the misspecification. The distribution that emerges from this problem can be thought of as a 'worst case distribution' that encodes these concerns and allows insight into the fears that guide the agent's decisions. Thus, the agent derives a policy that, while typically suboptimal in the context of his benchmark model, protects him against suspected misspecifications in that benchmark. We now formalize this intuition in the context of an abstract model, following the methodology described in Hansen and Sargent (2008).

Let us suppose that the robust agent entertains a benchmark model in which the state, control and innovation sequences are related according to the (possibly nonlinear) vector valued equation

$$x_{t+1} = g(x_t, u_t, \epsilon_{t+1}) \tag{1}$$

where x_t is the state vector, u_t is a vector of controls and $\{\epsilon_t\}$ is a sequence of random variates. Given a control law, $u_t = u(x_t)$, and a density, $p_\epsilon(\epsilon_{t+1}|x_t)$, for ϵ_{t+1}, Eq. (1) implies a benchmark transition density $p(x_{t+1}|x_t)$. It is convenient to partition the state, x_t into elements unknown on entering the period, which we identify with ϵ_t, and those elements that are predetermined, denoted s_t. We capture the dependence of s_t on the state prevailing in the previous period by the function f, such that $s_t = f(x_{t-1})$. The function f is determined by Eq. (1) and the relationship between the control and the state that captures the agent's behavior. With this decomposition we have $p(x_{t+1}|x_t) = p_\epsilon(\epsilon_{t+1}|x_t)\delta_{f(x_t)}(s_{t+1})$.[3]

We adopt multiplier preferences as a way of representing the agent's fear of model misspecification (Hansen and Sargent, 2008). In this case, the optimization problem of the agent takes the form of a particular two-player zero-sum

[2] The distorted models are perturbations of the probability distribution implied by the benchmark, rather than well defined economic structures. Thus the model uncertainty considered here is 'unstructured' (see Williams, 2008).

[3] Note that the x_t may contain ϵ_t as an element of the state so that an identity mapping is implicit in g. Note also that $\delta_{f(x_t)}(\cdot)$ takes the value of unity at $f(x_t)$ and zero elsewhere.

game between the robust agent (the maximizer) and a metaphorical ‘evil agent’ (the minimizer)

$$\max_{\{u_t\}} \min_{\{m_{t+1}\}} \sum_{t=0}^{\infty} E[\beta^t M_t\{h(x_t,u_t)+\beta\theta E(m_{t+1}\log m_{t+1}|\Im_t)\}|\Im_0] \tag{2}$$

where $h(\cdot,\cdot)$ is the robust agent's period payoff function and the problem is subject to Eq. (1), $M_{t+1}=m_{t+1}M_t$, $E[m_{t+1}|\Im_t]=1$, $m_{t+1}\geq 0$ and $M_0=1$.[4] Thus, $\{m_{t+1},t\geq 0\}$ is a sequence of martingale increments that recursively define a non-negative martingale $M_t=M_0\prod_{j=1}^{t} m_j$. The martingale defines Radon–Nikodym derivatives that twist the measures implicit in the benchmark model so as to obtain absolutely continuous measures that represent alternative distributions considered by the agent. Under these twisted measures one can form objects interpretable as expectations taken in the context of a distorted alternative model. This can be seen by defining a distorted conditional expectation operator to be

$$\tilde{E}[b_{t+1}|\Im_t]\equiv E[m_{t+1}b_{t+1}|\Im_t] \tag{3}$$

for some $\Im_{t+1}$ measurable random variable b_{t+1} given $\Im_t$. The conditional relative entropy associated with the twisted conditional distribution is given by the term $E[m_{t+1}\log m_{t+1}|\Im_t]$, which is a measure of how different the distorted measure is from the benchmark.

The agent's desire for robustness is reflected in the minimization over the sequence of martingale increments (chosen by the ‘evil’ player) that twist the distributions used to evaluate continuation values toward realizations of the state that are painful to the robust agent. The degree of robustness is controlled by the penalty parameter, $\theta>0$, that enters the objective by multiplying the conditional relative entropy associated with a given distortion. The penalty reflects our earlier intuition that the agent considers models that, although different, are somehow ‘near’ the benchmark.

We seek a recursive expression of the problem and, invoking results in Hansen and Sargent (2008), obtain a value function of the following form

$$\begin{aligned} V(\epsilon_t,s_t)=\max_{u_t} \min_{m(\epsilon_{t+1},s_{t+1})} h(x_t,u_t)+\beta\int & m(\epsilon_{t+1},s_{t+1})V(\epsilon_{t+1},s_{t+1})p_\epsilon(\epsilon_{t+1}|x_t) \\ & +\theta m(\epsilon_{t+1},s_{t+1})\log m(\epsilon_{t+1},s_{t+1})p_\epsilon(\epsilon_{t+1}|x_t)d\epsilon_{t+1} \end{aligned} \tag{4}$$

subject to $\int m(\epsilon_{t+1},s_{t+1})p(\epsilon_{t+1}|x_t)d\epsilon_{t+1}=1$ for all values of s_{t+1}. If one solves the inner minimization problem (interpretable as that of the ‘evil’ agent) one obtains the minimizing martingale increment, which has the form

$$m(\epsilon_{t+1},s_{t+1})=\frac{e^{-(V(\epsilon_{t+1},s_{t+1})/\theta)}}{E[e^{-(V(\epsilon_{t+1},s_{t+1})/\theta)}|\epsilon_t,s_t]} \tag{5}$$

If one substitutes this solution into the original problem, then we obtain the following expression for the Bellman equation

$$V(x_t)=\max_{u_t} h(x_t,u_t)-\beta\theta\log E\left[\exp\left(-\frac{V(x_{t+1})}{\theta}\right)\Big|x_t\right] \tag{6}$$

Note that the expectation in Eq. (6) is with respect to the benchmark transition density. This Bellman equation takes the same form as that of an agent with risk sensitive preferences (see Hansen and Sargent, 1995 and Tallarini, 2000). Under the risk sensitivity interpretation, however, θ reflects sensitivity to well defined, quantifiable risk whereas here it reflects the degree to which the agent fears model misspecification.

2.2. Worst case distributions and ‘animal spirits’

The martingale M_t from the solution of the agent's problem is a ratio of joint densities, $(\tilde{p}(x_{1:t})/p(x_{1:t}))$, where $\tilde{p}$ denotes the density implied by the worst case model while p denotes the benchmark model's density. The martingale increment, $m(x_{t+1})$, is a ratio of conditional densities, $(\tilde{p}(x_{t+1}|x_t)/p(x_{t+1}|x_t))$. Thus we have $\tilde{p}(x_{t+1}|x_t)=m(x_{t+1})p(x_{t+1}|x_t)$.

While $\tilde{p}$ is not directly interpretable as the conditional ‘beliefs’ of the agent, the fact that it differs from p emphasizes that, unlike under rational expectations, more than one distribution plays a role in the equilibrium. If the agent's benchmark model happens to be correct then the agent's behavior will seem inconsistent with full trust in the probabilities implied by the benchmark. In the next section we construct a production economy in which one can obtain a notion of ‘animal spirits’ by appealing to this state dependence of the worst case conditional distribution.

2.3. Planning problem

We will solve our production economy model via the problem of a planner operating under the benchmark model, whose objective is identified with that of the representative robust agent, obtained after solving the inner minimization problem for the worst case conditional likelihood ratio.

[4] We assume that the robust agent's information set, $\Im_t$ contains the entire history of states.

We use a Cobb–Douglas specification with variable capital utilization and labor-augmenting technology, driven by a random walk with drift and stochastic volatility in its innovations.[5] We also posit investment adjustment costs and variable capital utilization, with deprecation dependent on utilization. The resulting resource constraint, capital evolution, and exogenous state evolutions are given by

$$C_t = (u_{k,t}K_t)^{\alpha}(e^{z_t}l_t)^{1-\alpha} - I_t \tag{7}$$

$$K_{t+1} = (1-\delta(u_{k,t}))K_t + I_t\left(1-\frac{\kappa_I}{2}\left(\frac{I_t}{I_{t-1}} - e^{\Lambda_z}\right)^2\right) \tag{8}$$

$$z_{t+1} = \Lambda_z + z_t + e^{vol_{t+1}}\sigma_z\epsilon_{z,t+1} \tag{9}$$

$$vol_{t+1} = \rho_{vol}vol_t + \sigma_{vol}\epsilon_{vol,t+1} \tag{10}$$

where C_t, I_t, l_t, K_t and $u_{k,t}$ are consumption, investment, labor, capital and capital utilization in t, respectively. The technology process is given by z_t and vol_t is the component driving stochastic volatility. We assume that ϵ_{z_t} and $\epsilon_{vol,t}$ are iid standard Normal innovations, so that $p_\epsilon(\epsilon_{t+1}|x_t) = p_\epsilon(\epsilon_{t+1})$.

We follow Baxter and Farr (2005) in positing a depreciation rate for capital that is dependent on the rate of capital utilization according to

$$\delta(u_{k,t}) = \delta_0 + \frac{\delta_1}{1+\delta_2}u_{k,t}^{1+\delta_2} \tag{11}$$

We specify the period payoff function to be

$$h(x_t,u_t) = \log(C_t - \xi C_{t-1} - \eta_0 l_t^{\eta_1} J_t) \tag{12}$$

$$J_t = C_t^{\nu} J_{t-1}^{1-\nu} \tag{13}$$

where, following Jaimovich and Rebelo (2009), J_t, evolves to adjust the disutility of labor to ensure that the preferences are consistent with balanced growth and can admit varying degrees of wealth effects, according to the calibration of ν. We also include habits in consumption, controlled by the parameter ξ. The endowment of time is normalized to unity. Let $V(x_t)$ denotes the value function of the planner, then the planner's problem is

$$V(x_t) = \max_{u_t} h(x_t,u_t) - \beta\theta \log E_t\left[\exp\left(-\frac{V(x_{t+1})}{\theta}\right)\right]$$

s.t.

$$\begin{aligned}
&x_{t+1} = g(x_t,u_t,\epsilon_{t+1})\\
&u_t \equiv (C_t,I_t,l_t,u_{k,t},J_t,K_{t+1})'\\
&x_t \equiv (\epsilon_t,s_t)'\\
&\epsilon_t \equiv (\epsilon_{z,t},\epsilon_{vol,t})'\\
&s_t \equiv (vol_{t-1},K_t,I_{t-1},C_{t-1}J_{t-1})
\end{aligned}$$

where g is characterized by Eqs. (7)–(10) and (13). Thus we have mapped our model into a specific case of the abstract setup considered in Section 2.1.

2.4. Decentralization

The allocation that solves the planning problem can be decentralized in a familiar way by considering a competitive economy with complete markets where a representative, competitive firm operates the production technology, hiring labor and renting capital from the household. Using the shadow prices and allocations derived from the planning problem one can derive expressions for various objects of interest associated with such an equilibrium. Define Λ_t to be the stochastic discount factor of a representative household with a desire for robustness, which we decompose into two components: Λ_t^R, which takes the familiar form of the stochastic discount factor obtained under expected utility and Λ_t^U, which has conditional expectation of unity, the minimizing Martingale increment associated with the robust agent's problem.[6]

$$\Lambda_{t+1} = \Lambda_{t+1}^R \Lambda_{t+1}^U \tag{14}$$

[5] Various studies have identified evidence of time varying standard deviations of structural shocks within DSGE models (see, for example, Justiniano and Primiceri, 2008 and Fernandez-Villaverde et al., 2010).

[6] Note that with $V_{t+1} = V(x_{t+1})$ we denote a particular realization of the value function at time $t+1$, under the optimal solution. The dependence on the state is similarly suppressed for other variables.

$$\Lambda^R_{t+1} = \beta \frac{\lambda_{t+1}}{\lambda_t} \tag{15}$$

$$\Lambda^U_{t+1} = \frac{e^{-(V_{t+1}/\theta)}}{E_t[e^{-(V_{t+1}/\theta)}]} \tag{16}$$

where λ_t is the marginal utility of a unit of consumption to the representative household in the optimum. When a concern for model misspecification vanishes, $\theta \to \infty$, Λ^U_{t+1} becomes identically unity and we recover the stochastic discount factor implied by expected utility preferences. The conditional Market Price of Risk is given by $\sigma_t(\Lambda_{t+1})/E_t[\Lambda_{t+1}]$ and define $\sigma_t(\Lambda^U_{t+1})$ as the conditional Market Price of Model Uncertainty. We will denote these objects by MPR_t and MPR^U_t respectively.

Finally, the risk free rate, $r_{f,t}$ is given by

$$r_{f,t} = \frac{1}{E_t[\Lambda_{t+1}]} \tag{17}$$

2.5. Distorted moments

In addition to defining shadow prices associated with the planning problem, we can also define distorted conditional moments of functions of the state, x_{t+1}. In particular, one can define the distorted conditional mean and higher moments of the worst case conditional distributions of the innovations as follows:

$$\mu^1_t = E_t[\epsilon_{t+1}\Lambda^U_{t+1}] \tag{18}$$

$$\mu^j_t = E_t[(\epsilon_{t+1}-\mu^1_t)^j \Lambda^U_{t+1}] \tag{19}$$

2.6. Model solution

To obtain the system of equations that characterize the equilibrium objects to be analyzed, we combine the equations describing the shock processes, (9) and (10), the equations characterizing the solution to the planning problem (listed in the appendix), the equations defining prices under the associated decentralization, (14)–(17) and the equations characterizing distorted moments, (18) and (19). Together with the distributional assumptions regarding ϵ_t, we are left with a system that must hold in equilibrium for all values of the state. The system is solved by a particular set of 'equilibrium functions' for endogenous variables that relate the variables to the state, where the endogenous variables featured in the aforementioned structural equations are

$$\{V_t, C_t, I_t, l_t, u_{k,t}, J_t, K_t, \Lambda_t, \Lambda^U_t, \Lambda^R_t, r_{f,t}, \mu^1_t, \mu^j_t\} \tag{20}$$

As is typically the case in DSGE models, the equilibrium functions are unknown. Nevertheless, following the work of Aruoba et al. (2006) and Caldara et al. (2012), we can adopt a perturbation approach to solving an approximation to the equilibrium. However, due to the presence of the drifting random walk, z_t, certain variables listed above are nonstationary in equilibrium: $\{V_t, C_t, I_t, J_t, K_t\}$. Thus, in order to use perturbation methods one must re-express the structural equations in terms of 'detrended' versions of the variables that have constant nonstochastic steady states, around which we can approximate (using lower case to denote the transformation).[7] Thus the transformed list of variables is given by

$$\{v_t, c_t, l_t, j_t, i_t, u_{k,t}, k_t, \Lambda_t, \Lambda^U_t, \Lambda^R_t, r_{f,t}, \mu^1_t, \mu^2_t\} \tag{21}$$

We approximate the functions using a third-order solution and let $\hat{var}$ denote that a variable var has been approximated. It is worth noting that we can use standard perturbation techniques to obtain approximations to the equilibrium functions. In particular, Rudebusch and Swanson (2012) and Caldara et al. (2012) show how to obtain an approximation to the value function of an agent with recursive preferences and, mechanically, the methods we use are the same although the interpretation of the value function being approximated is very different from a robust control perspective. Indeed there is an algebraic equivalence between the recursive definition of utility in our case and the recursion obtained under a special case of Epstein and Zin (1989) preferences where the intertemporal elasticity of substitution equals unity. Under the alternative, Epstein-Zin interpretation θ (partly) reflects the agent's aversion to well defined risks.[8] However, we will regard θ simply as a parameter that controls desire for robustness, rather than aversion to risk.

[7] For example the variable, c_t can be related to its detrended version, c_t by $c_t \equiv (C_t/e^{z_t})$. It is also sometimes more convenient to approximate the log, rather than the level, of a variable.

[8] See Swanson (2012) for a discussion of the calculation of aversion to risk in the presence of labor in the utility function.

3. Characterizing worst case distributions

In order to completely analyze the robust control problem and gain insight into the robust agent's decisions, we must be able to compute the worst case model that agents use to evaluate consumption streams and decision rules. In particular, it is vital to obtain an approximation to the value function and, by extension, associated objects such as the minimizing martingale increment since these underpin the distorted conditional density of the worst case model. In linear-Gaussian, linear-quadratic and some continuous time models, one can directly compute these objects. However, in general nonlinear models, this tractability is lost.

3.1. An algorithm to draw from the worst case

As noted in Bidder and Smith (2011), an approximation to the value function is sufficient to be able to design Monte Carlo algorithms that sample asymptotically from the worst case distribution in general nonlinear models since it allows the construction of the approximation of the minimizing martingale increment and, consequently, an approximation to the worst case conditional density functions. This can be seen by recalling the nature of the distorted conditional density of the innovations, given by[9]

$$\tilde{p}_\epsilon(\epsilon_{t+1}|x_t) = m(\epsilon_{t+1},s_{t+1})p_\epsilon(\epsilon_{t+1}|x_t) = \frac{e^{-(V(\epsilon_{t+1},s_{t+1})/\theta)}}{E_t[e^{-(V(\epsilon_{t+1},s_{t+1})/\theta)}]}p_\epsilon(\epsilon_{t+1}|x_t)$$

Once we have an approximation to $\tilde{p}(x_{t+1}|x_t)$ we can then use it as a target density in a sampling algorithm. Then, by combining draws from the distorted conditional distributions with the agent's policy functions and other laws of motion in the economy, we can construct draws (approximately) from the worst case distribution over sequences entertained by the agent using the algorithm below.

Drawing from the Worst Case Distribution.

- Given some x_{t-1}
- draw $\epsilon_t^i \sim q(\epsilon_t|x_{t-1}) = N(\hat{\mu}_{t-1}^1,\hat{\mu}_{t-1}^2)$, $i=1,\ldots,N$
- assign weight $(W_t^i = (\hat{m}(\epsilon_t^i,\hat{f}(x_{t-1}))p(\epsilon_t^i|x_{t-1})/q(\epsilon_t^i|x_{t-1}))$
- resample with probability $\propto W_t^i$, call new set of draws $\tilde{\epsilon}_t$
- now $\tilde{\epsilon}_t^i \sim \tilde{p}(\epsilon_t|x_{t-1})$, $i=1,\ldots,N$
- draw once from $\{\epsilon_t^i\}_{i=1}^N$ and store - if a draw is needed for simulation

where we write $\hat{f}(x_{t-1})$ to convey the fact that our perturbation solution will imply an approximation to the evolution of the predetermined components of the state.[10]

Finally, it is worth emphasizing here that the fact that we retain the policy functions and other deterministic components of the law of motion for the state, implied under the benchmark, when simulating the worst case by feeding the economy innovations drawn from the distorted conditional distributions, does not mean that the only dimension in which the agent fears misspecification is in relation to the innovation distributions. The simulation approach we have taken, when looped over time, induces draws from the worst case joint distribution over *sequences* of all variables in the economy and the distortions implicit in this distribution could reflect a wide variety of misspecifications, in many dimensions of the economy other than specifically the distribution of innovations. This again emphasizes the unstructured nature of model uncertainty implicit in the robust control approach.[11]

3.2. Clarifying the role of perturbation methods and the worst case

It is useful to clarify at what point in our analysis the characterization of the worst case occurs and the role perturbation plays in our approach. Regarding the former issue, one can solve for allocations and shadow prices without explicitly specifying the nature of the worst case model or even including Eqs. (18) and (19) that define distorted moments in the system to be solved. After approximating the value function, one can then use the sampling methodology described above to elicit the properties of the worst case model that is implicit in the solution. This offers an interesting contrast with other approaches to capturing behavior under model uncertainty, in which one must be explicit about the structural nature of misspecification or ambiguity at the point of solution (see Ilut and Schneider, 2012, for example). The fact that one simply

[9] With our state written as $x_t = (\epsilon_t,s_t)$ and given that we have restricted to considering distortions that are absolutely continuous with respect to the benchmark model, we have that $\tilde{p}(s_{t+1}|x_t) = p(s_{t+1}|x_t) = \delta_{f(x_t)}(s_{t+1})$. Thus the 'evil agent' does not distort the transition function of the predetermined elements of the state.

[10] The moments of the distorted innovation distributions reported later in the paper are obtained using importance sampling (IS), where we simply use the normalized importance weights to calculate estimates of moments.

[11] See Bidder and Smith (2011) for an example of how this unstructured approach to model uncertainty can still yield insights into what sort of structural misspecifications might concern an agent. In that model, it appears that any structural misspecifications that would imply 'Long Run Risk' and 'disasters' would worry the agent.

needs to calculate an approximation to the value function to undertake worst case analysis means that misspecifications captured by complicated state dependence and correlations among shocks can be easily handled in a wide variety of models and using a broad class of solution methods.

Regarding the role of perturbation methods, we identify two main features. First, as discussed above, perturbation is used to solve the model and obtain approximations to the value function and other equilibrium objects. This reflects the fact that such methods are well suited to nonlinear DSGE models of the type we consider here, although our analysis and insights would remain intact if one were able to solve the model and obtain an approximation to the value function by other techniques. The second aspect of the role of perturbation in this paper is that one can use it to obtain approximations to the conditional distorted moments defined in Eqs. (18) and (19). These might be thought intrinsically interesting and they will underpin worst case impulse response analysis below, but they can also be used to enhance the proposal densities used in the Sampling Importance Resampling (SIR) algorithm described above.

By using a proposal distribution that is augmented by an approximation to the conditional moments under the worst case model we can improve the effective sample size of the SIR algorithm relative to proposing directly from the benchmark model.[12] Indeed, if one draws a proposal, ϵ_t^i, from the benchmark model, $\epsilon_t^i \sim p_\epsilon(\epsilon_t | x_{t-1}^i)$, then the associated importance weight is $W_t^i = \hat{m}(\epsilon_t^i, \hat{f}(x_{t-1}^i)) \propto e^{-\hat{v}_t^i/\theta}$. Consequently, when a concern for robustness is of particular interest and the worst case deviates meaningfully from the benchmark the benchmark distribution will deteriorate as a proposal density, increasing the benefit one gains from using a proposal augmented with $\hat{\mu}_{t-1}^j$.

In fact, when the effective sample size using this augmented proposal density is very close to the total number of samples, then it is a signal that the proposal accurately approximates the worst case conditional distribution and thus can itself be used for simulation directly, forgoing the resampling step. Obtaining approximate draws from the worst case model in this manner is substantially faster than when using the SIR algorithm although certain characteristics of the worst case distortions may not be fully captured by $\hat{\mu}_{t-1}^j$.

4. Results

We will calibrate our model in relation to US data by comparing moments with those generated under the benchmark model featuring robustness. We will then characterize the properties of the worst case distribution and, using these insights, demonstrate how and why prices and quantities respond to volatility shocks under the benchmark model. We show how robustness contributes to movements in prices and quantities and can partially explain economic fluctuations via persistent waves of relative optimism and pessimism, that is, ‘animal spirits’.

4.1. Calibration

We calibrate the model to match US quarterly data from 1953Q4 to 2011Q4 and list our parameterization in Table 1. A full description of the data series and calibration are contained in the appendix. Our parameter choices are fairly standard, but the parameter that governs the agent’s desire for robustness, θ, warrants further discussion. We choose a value of θ that induces a conditional market price of risk that is approximately 0.25 on average. However, we also relate our choice of θ to detection error probabilities, which are advocated in Hansen and Sargent (2008) as a way of assessing the plausibility of the calibration of a desire for robustness. Detection error probabilities express how difficult it is, with a limited amount of data, for an agent to distinguish between the worst case and benchmark model. If an agent has difficulty distinguishing between the two models then the misspecification implicit in the distorted model is regarded as one that an agent might plausibly seek robustness against. If, however, the agent can easily distinguish between the two models then the implicit misspecification is one against which an agent is unlikely to seek robustness since he is likely to have dismissed the possibility of such misspecifications, so contrary are they to the data he observes. Formalizing ‘plausibility’ in this way restricts the modeling freedom arising from introducing the additional parameter, θ.

One can obtain detection error probabilities by applying likelihood ratio tests, combined with a prior over the benchmark and worst case model, and calculating the probability with which the agent will mis-identify the model that generated the data. Using the methodology described in Bidder and Smith (2011), our benchmark calibration with robustness implies a detection error probability of approximately 2.5%.[13] This means that there is about a 2.5% chance of making an error in statistically discriminating between the worst case and benchmark models when the data is drawn from each models’ unconditional distribution, owing largely to a shift in the means of the respective ergodic distributions (see Table 2). This value is perhaps on the edge of plausibility. However, simulating data conditional on a fixed initial state (the deterministic steady state), detection error probabilities are about 15% for 50 quarters of data, and about 5% for 150 quarters. While unconditionally our calibration yields low detection errors, shorter data samples yield much more

[12] In our baseline calibration, using the tuned proposal distribution raises the effective sample size from 92.4% of the number of draws, to 99.0%

[13] The detection error probabilities were calculated using investment growth, consumption growth, hour growth, the risk free rate and the implied return on capital as observables. We applied zero mean measurement errors with standard deviations of approximately 10% of the corresponding variables’ unconditional standard deviations and a swarm of 160,000 particles. 1200 draws of series were made from the benchmark and worst case models, with each series being 250 periods long.

Table 1
Parameter values.

Parameter	β	θ	ξ	ν	η_1	l_{ss}	$\delta(1)$
Value	0.992	8	0.45	0.005	1.4	0.2305	0.03
Parameter	δ_2	α	κ	Λ_z	σ_z	ρ_{vol}	σ_{vol}
Value	0.85	0.36	4	0.004	0.01094	0.9	0.15

Table 2
Comparing moments. Simulations were made for 10,000 periods with the first 1000 being dropped as a burn-in. For the simulations under the worst case we employed a Sampling Importance Resampling algorithm (see Section 3) with 500,000 proposals from $N(\hat{\mu}^1_{t-1},\hat{\mu}^2_{t-1})$.

	Data	Benchmark (R.C.)	Worst case (R.C.)
$\mu(\Delta \log Y_t)$	0.415	0.418	−0.001
$\sigma(\Delta \log Y_t)$	0.959	0.978	1.344
$\mu(\Delta \log C_t)$	0.497	0.417	−0.001
$\sigma(\Delta \log C_t)$	0.539	0.753	1.046
$\mu(\Delta \log I_t)$	0.372	0.418	−0.001
$\sigma(\Delta \log I_t)$	2.360	1.531	2.329
$\mu(\Delta \log l_t)$	−0.029	−0.000	−0.000
$\sigma(\Delta \log l_t)$	0.879	0.697	1.475
$\mu(r_{f,t})$	0.403	0.612	0.925
$\sigma(r_{f,t})$	0.506	0.596	0.659
$\rho(\Delta \log C_t,\Delta \log Y_t)$	0.554	0.992	0.994
$\rho(\Delta \log I_t,\Delta \log Y_t)$	0.757	0.989	0.980
$\rho(\Delta \log l_t,\Delta \log Y_t)$	0.737	0.749	0.697

plausible results. In addition, although the our detection error probabilities are somewhat low, they still seem more plausible than the implied level of risk aversion associated with $\theta = 8$, were we to reinterpret preferences from a risk-sensitive perspective.

In Table 2 we show unconditional moments associated with simulations of the benchmark and worst case models. The moments denoted 'Benchmark (R.C.)' are the result of simulations where we assume an agent who fears his model is misspecified, but where the data generating process is in fact the benchmark model. Thus, the misspecifications that the agent fears are 'all in his head'. The moments denoted 'Worst Case (R.C.)' retain the policy functions that generated the moments in the Benchmark (R.C.) case (and other deterministic components of the law of motion for the state), but innovations are drawn from the (state dependent) worst case conditional distribution. The thought experiment here is to examine what the economy would look like if the agent fears his model is misspecified and the worst case model is the data generating process. If we compare the second and third columns of data we observe substantial reductions in the growth rate of the economy under the worst case, relative to the benchmark. The growth rates of the main aggregate quantities are also more volatile. Here we see why we obtain fairly low detection error probabilities. Given the rather large difference between the economy's growth rates under the benchmark and worst case, one might imagine that they would be easy to distinguish, even with relatively short spans of data.

4.2. 'Animal spirits'

In our robust control environment, an agent's actions are underpinned by a state dependent worst case conditional distribution. This state dependence gives rise to our notion of 'animal spirits'. Since the worst case distribution does not have a closed form, we will sample from it using the techniques described in Section 3 and compute conditional moments. Understanding the properties of these distributions is a precursor to understanding how an agent behaves in reaction to such 'animal spirits'. We show how these distributions depend, in particular, on the volatility state and then discuss how they lead to phenomena such as waves of sentiment and business cycle fluctuations.

4.2.1. Computing 'animal spirits'

Under the benchmark model, the innovations are iid standard Normal random variates. In Table 3 we display the distorted moments of the innovations to the technology and volatility processes when the economy is at its nonstochastic steady state, and when vol_{t-1} is ± 2 unconditional standard deviations of its mean. We observe a negative mean shift in ϵ_z, which signifies that the agent fears misspecifications that would imply lower growth. Reflecting the fact that volatility is undesirable, we see that the mean of ϵ_{vol} is distorted in the positive direction, implying a worst case distribution that exhibits elevated risk. Under the worst case, the innovations are negatively correlated.[14] Although the magnitude of

[14] The variances appear undistorted at unity so the covariances are essentially the same as the correlations.

Table 3
Distorted conditional moments of ϵ_z and ϵ_{vol}. We use importance sampling to compute the distorted conditional moments of ϵ_z and ϵ_{vol} where we use 500,000 draws from $N(\hat{\mu}^1,\hat{\mu}^2)$ as proposals. We initialize from the nonstochastic steady state and also with the lagged volatility element perturbed ± 2 standard deviations.

vol_{-1}	μ_{ϵ_z}	$\mu_{\epsilon_{vol}}$	$cov_{\epsilon_z,\epsilon_{vol}}$
$+2sd$	-0.447	0.111	-0.058
0	-0.263	0.094	-0.040
$-2sd$	-0.159	0.087	-0.016

the distortion is small, the negative sign of the covariance is intuitive in that ‘good’ (‘bad’) technology innovations are associated with ‘good’ (‘bad’) volatility shocks. The unconditional averages for ϵ_z and ϵ_{vol} underpinning the worst case simulations referred to in Section 4.1 are -0.307 and 0.115, respectively and their unconditional covariance is -0.053, which largely explains the lower growth and elevated volatility exhibited by the worst case simulations.

Most importantly for our characterization of ‘animal sprits’, we see that the conditional distortions vary with the volatility state. The pessimistic distortions are exaggerated when the volatility state is high and moderated when it is low. Since the positive distortion to the mean of ϵ_{vol} appears to be increasing in the level of volatility, the persistence of variations in volatility is slightly increased under the worst case model. We obtain, via a different specification from Ilut and Schneider (2012), time variation in conditional technological growth rates under the worst case, without directly restricting this to be the nature of the distortion feared by the households.

4.2.2. Waves of sentiment

Modeling ‘animal spirits’ as endogenous fluctuations in worst case conditional probability distributions might initially seem to be a restriction beyond what Keynes had in mind, since the term is often used to connote waves of optimism, as well as pessimism.[15] In order to reconcile our apparent emphasis on pessimism with the broader sense in which the term ‘animal spirits’ is typically used, one must clarify the probability measure with respect to which pessimism is defined. The pessimism we have discussed is in relation to the benchmark model we assume the agents possess. In our simulations, we also treat this benchmark as the data generating process and use it as a laboratory in which to examine the behavior of robust agents. However, in contrast with rational expectations analysis, the data generating process need not coincide with a robust agent’s benchmark model. Thus, what appears as pessimism with respect to the data generating process under our assumption that the benchmark model is the true description of the economy, may not appear as pessimism with respect to some other potentially unknown or unknowable data generating process.

While it is true that our robust agent entertains a pessimistic perspective relative to his benchmark model, the time varying volatility of technology shocks induces waves of optimism and pessimism when compared with the average degree to which he distorts the benchmark distribution. In low (high) volatility periods, the distorted conditional mean of the technology innovation is less (more) negative than on average. Thus, with a renormalization, one can reinterpret a period of relatively little pessimism as the sort of regime that might anecdotally be referred to as one of optimism. To emphasize this interpretation, we simulate 200 time periods under the approximating model with $\theta = 8$ and plot the distorted mean of the technology innovation and the conditional market price of model uncertainty. The distorted mean process has been standardized to have zero mean and unit variance. We see in Fig. 1 that the model generates relatively optimistic periods (induced by lower volatility) in which the market price of uncertainty and innovation mean distortion are lower than average, and relatively pessimistic periods (induced by higher volatility) in which the opposite is true.

4.2.3. ‘Animal spirits’ and recessions

The behavior of a robust agent is informed by concerns about misspecification that are reflected in the conditional worst case density. When this density varies with the state, so too do consumption, investment and hours desired by our agent. Thus, unlike in the endowment economy of Bidder and Smith (2011), a fear of model misspecification can feed back on the evolution of the state through agents’ actions.

How do these ‘animal spirits’ affect our agent’s actions? To answer this question, we examine the effect of a positive unit standard deviation shock to the vol_t process using generalized impulse response functions, conditional on the economy initially being at its nonstochastic steady state (see Koop et al., 1996). In Fig. 2, we observe the responses of labor, consumption, investment, output and capital utilization under Benchmark (R.C.) and Worst Case (R.C). For comparison, we also plot the responses for an agent with Expected Utility.[16] The impulse responses for labor and utilization express the gap between the shocked and unshocked scenarios as a percentage of the corresponding nonstochastic steady state values. The other impulse responses are the differences in logs (multiplied by 100) under the shocked and baseline scenarios.

First, consider the Benchmark (R.C.) case. On the impact of a volatility shock consumption, hours, capital utilization and investment all decline. The decrease in utilization and hours implies that output also falls. In addition, the fact that depreciation

[15] We thank an anonymous referee for encouraging us to clarify this issue.
[16] For expected utility we set $\theta = 10\,000$.

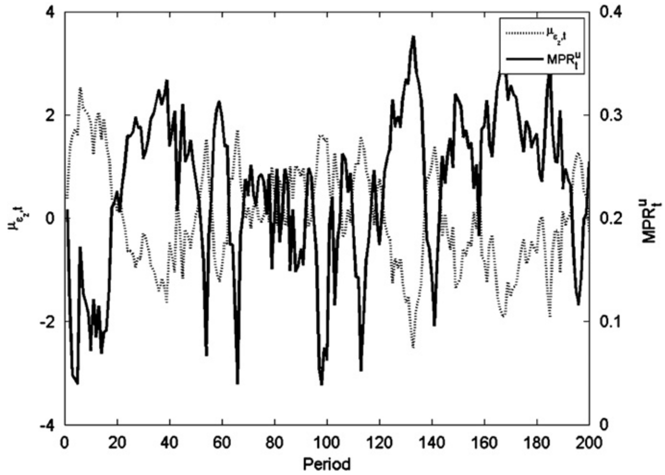

Fig. 1. Fluctuations in sentiment (as captured by MPR_t^U and distorted technology innovation mean, $\mu_{\epsilon_z,t}$). We use the first 200 periods (post burn-in) of the longer simulations used to obtain the moments discussed in Section 4.1.

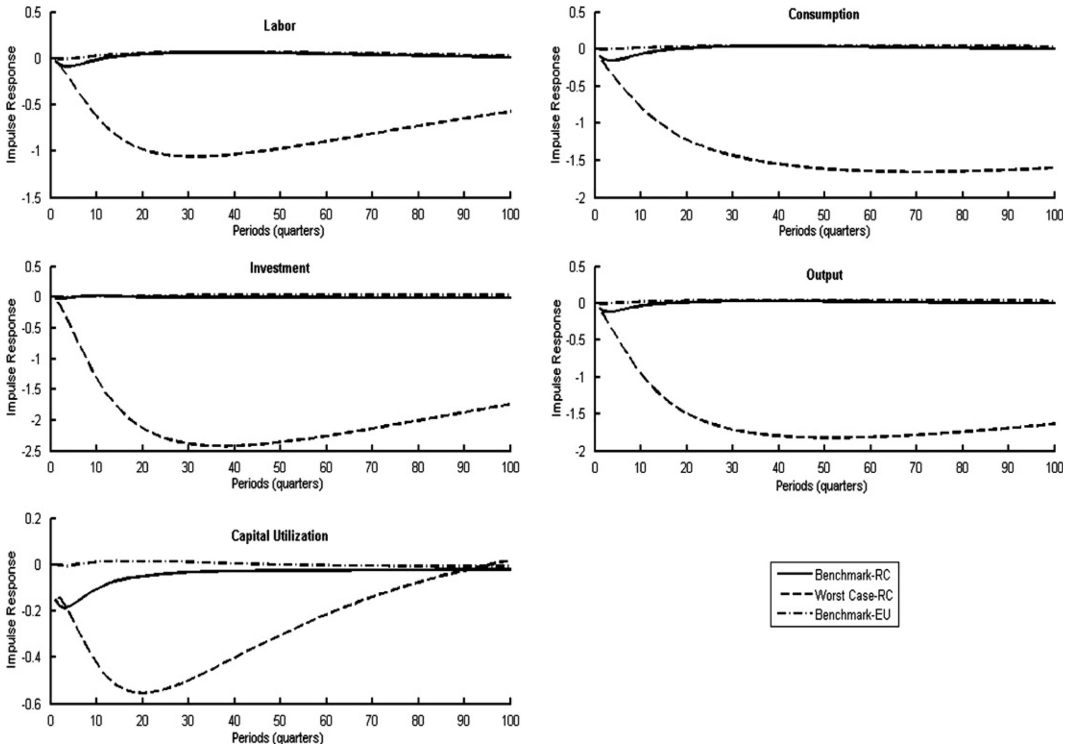

Fig. 2. Impulse responses to a positive volatility shock. In period $t=0$ we set s_0 to its nonstochastic steady state value and $\epsilon_{z,0}=0$ under both the shocked and unshocked case. In the shocked case we set $\epsilon_{vol,0}=1$ and in the unshocked case we set $\epsilon_{vol,0}=0$. For $t>1$ under the benchmark impulse responses we draw ϵ_t from $N(0,I)$ while under the worst case impulse responses we draw from $N(\hat{\mu}^1_{t-1},\hat{\mu}^2_{t-1})$.

declines as utilization drops explains why we are able to obtain a decline in investment, relative to the unshocked baseline, despite accumulation of capital resulting from powerful precautionary and wealth effects that we will discuss further below. It is therefore important that our Jaimovich and Rebelo (2009) preference specification largely insulates the agent's desired labor supply from wealth effects so that hours ultimately decline along with capital utilization.[17]

Continuing to analyze Fig. 2, it is instructive to analyze what might be going through the head of the robust households following a volatility shock. The 'worst case impulse response functions' plotted in the figure help us ascertain the worst

[17] See Bloom et al. (2012) and Basu and Bundick (2012) for alternative methods of obtaining a contraction and comovement among macroeconomic aggregates in response to volatility shocks.

case fears induced by the shock. These impulse responses are constructed by drawing innovations from the state dependent worst case conditional distribution after the initial period which the volatility shock hits.[18] In the Benchmark (R.C.) case above these distributions are state independent and standard Normal. However, in the worst case impulse responses the distributions are state dependent and, in particular, feature means that are nonzero and depend on the level of volatility since they reflect the pessimism of the households.

Owing largely to the pessimistic twists that accompany the volatility shock, the household fears an equilibrium path of the economy that is persistently depressed, relative to the unshocked case. This implies a much more dramatic decline in the real side of the economy, particularly in the consumption path. Thus, one can explain the behavior of the benchmark series under robustness by appealing to the worst case responses. In particular, the substantial contraction and negative wealth effect implicit in the worst case impulse responses suggest why households choose to scale back consumption, since these worst case fears interact with the household's desire for consumption smoothing. As the volatility shock dies away (implying reduced pessimism) they scale back on capital accumulation and allow consumption to recover, as shown in Fig. 2. Interestingly, the impact of 'animal spirits' exhibits some similarities with a news shock which never comes to fruition. After all, the exaggerated pessimism induced by the volatility shock and reflected in the worst case impulse responses reflects only the agent's fears rather than the true data generating process, which we here assume to be the Benchmark (R.C.) case. This again hints at why Jaimovich and Rebelo (2009) preferences play a useful role in our environment since part of their popularity has come through use in models with news shocks.

On the impact of a volatility shock, output drops by about 0.1% relative to the unshocked baseline. In contrast, under expected utility, we see from Fig. 2 that the impulse responses would be orders of magnitude smaller. Thus a fear of model misspecification is important in enhancing the effects of fluctuations in risk on the behavior of the economy. Allowing for robust preferences and stochastic volatility not only implies swings in pessimism, as discussed above, but these fluctuations in 'animal spirits' also seem to amplify the real effects of volatility shocks. Thus, we see an example of the 'urges' referred to by Keynes being translated into actions. When a volatility shock strikes, risk increases but the conditional worst case distribution in the head of a robust agent also changes. This additional effect would appear like a shock to preferences or an unmotivated shift in beliefs if an econometrician were to require that the agent fully trusts his model. However, in our framework this wedge reflects the state dependent twisting of conditional distributions for the purposes of deriving a robust policy. Furthermore, this expression of 'animal spirits' does not require the robust agent to 'believe' the worst case scenario, but simply that it (and other possible data generating processes) be plausible enough to guard against. It is often suggested that fluctuations in sentiment or 'animal spirits' are disconnected from fundamentals and that they and the actions they induce are therefore somehow unjustified. With a robust control approach, the connection of beliefs with 'fundamentals' is subtle and what might otherwise be regarded as unjustified, may in fact be sensible.

4.2.4. A driver of business cycles?

While we've seen that such 'animal spirits' are able to cause recessions, we have not yet shown how much of the business cycle is driven by such fluctuations in sentiment. Here we attempt to quantify this impact.

We will compare how our Benchmark (R.C.) economy behaves when we individually set the technology and volatility innovation equal to their average level (zero). Specifically, we simulate the model first imposing $\epsilon_{z,t} = 0\ \forall t$. Note that we assume the agent continues to envisage a sequence of standard Normal innovations under the benchmark but simply happens to get the mean realization in every period. Thus the variation in the endogenous quantities is entirely driven by animal spirits since $\epsilon_{vol,t}$ is the only innovation that actually hits the economy with a nonzero and random value in each period. We then compute the unconditional standard deviation of a selection of quantity and price variables, and divide it by the standard deviation of the fully stochastic model (contained in Table 2). The output of this calculation is in the first column of numbers in Table 4. We interpret this number as the fraction of total volatility generated by 'animal spirits' alone. It is a residual, once we have removed other sources of realized randomness. We then do the analogous exercise setting $\epsilon_{vol,t} = 0\ \forall t$.[19]

For the case when $\epsilon_{z,t} = 0\ \forall t$, we see large drops in overall volatility. However, it seems that 'animal spirits' alone can generate between approximately 1% and 16% of the total standard deviation in quantities without any variation due to random technology realizations. This variation is due to changes in the mean of the worst case conditional distribution of the technology shock, even though we simulate the economy so that the agent always experiences benign realizations of $\epsilon_{z,t} = 0\ \forall t$. Under expected utility, all of the comparable numbers for quantities are less than 1% and we omit them here due to space constraints. Hence volatility shocks alone without a fear of model misspecification do not generate much variability in quantities.

Simulating the economy with $\epsilon_{vol,t} = 0\ \forall t$ leads to a much smaller drop in volatility. The numbers for quantities suggest that for our simple model, technology shocks are still the primary drivers of business cycles even with 'animal spirits'. We

[18] In our third-order approximation, $\hat{\mu}_{t-1}^2$ is the unit vector. In addition, the distortion to the mean of ϵ_{vol} that we capture with our SIR algorithm is not captured by the direct perturbation approximation, so that our worst case impulse responses likely understate the effect of the volatility shock.

[19] Note that due to the nonlinearities inherent in our model, we cannot do an exact variance decomposition. For example, setting $\epsilon_{vol,t} = 0$ also decreases the unconditional variance of the $e^{vol_t}\epsilon_{z,t}$ term. Also, see Fernandez-Villaverde et al. (2010) for a similar exercise.

Table 4
Volatility residual. To compute the volatility residual when $\epsilon_{z,t}=0$, we simulate the benchmark economy setting $\epsilon_{z,t}=0\ \forall t$ and compute the sample moments. We then divide those moments by the same moments computed when $\epsilon_{z,t}\sim N(0,1)$ and report the ratio. The volatility residual when $\epsilon_{vol,t}=0$ is computed in an analogous way.

	$\epsilon_{z,t}=0$	$\epsilon_{vol,t}=0$
$\sigma(\Delta\log Y_t)$	0.092	0.886
$\sigma(\Delta\log C_t)$	0.158	0.874
$\sigma(\Delta\log I_t)$	0.018	0.895
$\sigma(\Delta\log l_t)$	0.086	0.895
$\sigma(u_t)$	0.118	0.902
$\sigma(r_{f,t})$	0.435	0.807

suspect that richer models where stochastic volatility plays a larger role would yield more scope for the 'animal spirits' to drive business cycles.

5. Conclusions

Applying new methods to analyze robust control problems within nonlinear DSGE models, we propose a novel formulation of 'animal spirits'. We have taken a simple production model with stochastic volatility and introduced a fear of model misspecification. Working with a form of preferences that express these doubts we demonstrated that the worst case model entertained by the robust agents differs from the benchmark in interesting ways.

Most importantly, we used the agents' benchmark model to illustrate interesting interactions between robust preferences and stochastic volatility, with variations in risk inducing fluctuations in the worst case scenarios considered by robust agents. When the economy is in a volatile state, pessimism is exaggerated. This interaction between time varying risk and robustness provides an amplification mechanism for volatility shocks, giving rise to a phenomenon interpretable as animal spirits. In addition we show that these animal spirits can help explain a small but nontrivial fraction of macroeconomic volatility in our benchmark model.

To obtain our results we employ a novel technology for characterizing the households' worst case scenarios. These methods extend the rapidly growing robust control literature into the domain of medium scale macroeconomic modeling.

Acknowledgments

The authors would like to thank Tom Sargent, Susanto Basu, Tim Cogley, Mark Gertler, Sydney Ludvigson, Virgiliu Midrigan, Filippo Occhino, Anna Orlik, Eric Swanson, Martin Schneider, Botao Wu, Paul Ho, Klaus Adam, Robert G. King and an anonymous referee for helpful advice and comments. The authors also wish to thank seminar participants at the University of Warwick, Cardiff University, Royal Holloway, the Bank of England, the Federal Reserve Bank of Chicago, the System Committee on Business and Financial Analysis conference (2011) and the European Meeting of the Econometric Society (2012). This paper was previously entitled 'Robust Control in a Nonlinear DSGE Model'. The views expressed in this paper, and all errors and omissions, should be regarded as those solely of the authors, and are not necessarily those of the parties listed above, the Federal Reserve Bank of San Francisco, the Federal Reserve Board of Governors or the Federal Reserve System.

Appendix A. Supplementary material

Supplementary data associated with this article can be found in the online version at http://dx.doi.org/10.1016/j.jmoneco.2012.10.017.

References

Aruoba, S.B., Fernandez-Villaverde, J., Rubio-Ramirez, J.F., 2006. Comparing solution methods for dynamic equilibrium economies. Journal of Economic Dynamics and Control 30, 2477–2508.
Basu, S., Bundick, B., 2012. Uncertainty shocks in a model of effective demand. Working Paper 18420. National Bureau of Economic Research.
Baxter, M., Farr, D.D., 2005. Variable capital utilization and international business cycles. Journal of International Economics 65, 335–347.
Bidder, R., Smith, M.E., 2011. Doubts and variability. Unpublished paper.
Bloom, N., Floetotto, M., Jaimovich, N., Saporta-Eksten, I., Terry, S.J., 2012. Really uncertain business cycles. Working Paper 18245. National Bureau of Economic Research.
Caldara, D., Fernandez-Villaverde, J., Rubio-Ramirez, J.F., Yao, W., 2012. Computing DSGE models with recursive preferences and stochastic volatility. Review of Economic Dynamics 15, 188–206.

Epstein, L., Zin, S., 1989. Substitution risk aversion and the temporal behavior of consumption and asset returns: a theoretical framework. Econometrica 57, 937–968.
Farmer, R.E.A., 2010. Animal spirits. UCLA Economics, Unpublished paper.
Fernandez-Villaverde, J., Guerron-Quintana, P., Rubio-Ramirez, J.F., 2010. Fortune or virtue: time-variant volatilities versus parameter drifting in U.S. data. Working Paper 15928. National Bureau of Economic Research.
Hansen, L.P., Sargent, T.J., 1995. Discounted linear exponential quadratic Gaussian control. IEEE Transactions on Automatic Control 40, 968–971.
Hansen, L.P., Sargent, T.J., 2008. Robustness. Princeton University Press, Princeton, NJ.
Ilut, C., Schneider, M., 2012. Ambiguous business cycles. Working Paper 17900. National Bureau of Economic Research.
Jaimovich, N., Rebelo, S., 2009. Can news about the future drive the business cycle? American Economic Review 99, 1097–1118.
Justiniano, A., Primiceri, G.E., 2008. The time-varying volatility of macroeconomic fluctuations. American Economic Review 98, 604–641.
Keynes, J.M., 1936. The General Theory of Employment, Interest and Money. Macmillan, London, UK.
Koop, G., Pesaran, M.H., Potter, S.M., 1996. Impulse response analysis in nonlinear multivariate models. Journal of Econometrics 74, 119–147.
Rudebusch, G.D., Swanson, E.T., 2012. The bond premium in a DSGE model with long-run real and nominal risks. American Economic Journal: Macroeconomics 4, 105–143.
Swanson, E.T., 2012. Risk aversion and the labor margin in dynamic equilibrium models. American Economic Review 102, 1663–1691.
Tallarini, T.D., 2000. Risk-sensitive real business cycles. Journal of Monetary Economics 45, 507–532.
Williams, N., 2008. Robust control. In: Durlauf, S.N., Blume, L.E. (Eds.), The New Palgrave Dictionary of Economics, Palgrave Macmillan, Basingstoke, UK.