Risk preferences, intertemporal substitution, and business cycle dynamics

Jim Dolmas*

February 15, 2012

Abstract

This paper examines the implications of alternative specifications of risk preferences—including preferences displaying first-order risk aversion (FORA)—together with alternative assumptions regarding individuals' elasticities of intertemporal substitution (EIS), for the behavior of a technology shock-driven business cycle model. The most general version of the model I consider also includes external habit formation and capital adjustment costs. I solve the model by applying a version of the Chebyshev collocation method described by Caldara *et al.* [4]. Risk preferences matter for some first moments, because of precautionary capital accumulation, though they have little impact on model-implied average asset returns. In terms of the models' second moment predictions, the assumed EIS and the presence or absence of habits matters a great deal, while the impact of alternative risk preferences is negligible. Some curious outcomes obtain in cases where the EIS is bigger than one or habits are present—two assumptions that have become more common in the literature—including countercyclical consumption in the former case and countercyclical hours worked in the latter case.

Despite their negligible impact on the model dynamics or asset returns, risk preferences matter a great deal for the perceived welfare cost of aggregate volatility. Under the FORA specification I use, which I argue has some empirical plausibility, costs are as high as 1.3% of aggregate consumption.
JEL CODES: E32, D90, G12
PRELIMINARY & INCOMPLETE.

*Federal Reserve Bank of Dallas, 2200 North Pearl Street, Dallas, TX 75201. E-mail: jim@jimdolmas.net. URL: http://www.jimdolmas.net/economics. This is a substantially revised version of an earlier paper that circulated under the title "Real business cycle dynamics under first-order risk aversion". For comments on that earlier paper, I thank, without implicating, Erwan Quintin, Eric Young, Chris Otrok, Stan Zin and Chris Telmer. I thank Anton Cheremukhin and Ananth Ramanarayanan for useful discussions during the development of this revision. Disclaimer: The views expressed herein are those of the author and do not necessarily reflect the views of the Federal Reserve Bank of Dallas or the Federal Reserve System. Typeset in LaTeX, using Kile 2.1 from KDE.

1 Introduction

In this paper, I take as a starting point intertemporal preferences that allow a separation of attitudes towards intertemporal substitution (for deterministic paths) from attitudes towards risk (for timeless gambles). I then consider three different specifications of risk preferences, and ask whether the choice matters for business cycle dynamics in variants of the real business cycle (RBC) model that lies at the core of many current dynamic stochastic general equilibrium (DSGE) models.

Since business cycle models are neither deterministic nor timeless, I also consider variation in individuals' willingness to substitute intertemporally, looking for potential interactions between risk preferences and intertemporal substitution.

The intertemporal preferences that result encompass the standard, time- and state-separable expected utility form, and two variants of non-time-separable preferences of the Kreps-Porteus [23] variety. One variant—the more standard, what one commonly thinks of as 'Epstein-Zin (EZ) preferences'—collapses to constant relative risk aversion expected utility in timeless settings (gambles over constant levels of consumption and leisure). The other, described in detail below, embodies non-state-separability and first-order risk aversion (FORA), and also owes to Epstein and Zin [12]. The latter two cases are calibrated to higher levels of risk aversion than is common in the business cycle literature.[1]

The business cycle model, in the most general form I consider, includes both external habit formation in consumption and capital adjustment costs. I also examine a stripped-down version of the model in which those frictions are absent. The model is driven by a single shock to total factor productivity (TFP).

The model behavior I am interested in includes implications for first moments, second moments and impulse responses, and the welfare costs associated with aggregate volatility. In terms of first moments, the two non-EU preference specifications generate noticeable amounts of precautionary capital accumulation, especially the specification incorporating first-order risk aversion. In contrast to results obtained in endowment economies ([12]) or production economies with fixed labor supply ([20]), risk preferences, intertemporal substitution and the presence or absence of habits make little difference for the average equity risk premia the model generates.

The model's second moment implications are essentially independent of the specification of risk preferences, consistent with the previous finding of Tallarini [35] in a more specialized model. Impulse responses are, in fact, virtually indistinguishable across the three risk preference specifications I use. The choices that do matter, significantly, for second moments are the assumed elasticity of intertemporal substitution (EIS) and the presence or absence of habit formation. The use of an EIS greater than one is now common—and necessary—in models that seek to resolve asset pricing puzzles by allowing for

[1] I argue in section A.2 that the first-order risk averse calibration is, nevertheless, empirically plausible.

long-run consumption risk (*e.g.*, [1] or [10]), but has not been much-analyzed in business cycle models. In the stripped-down version of the model in this paper—the standard neoclassical stochastic growth model—assuming an EIS in excess of one can lead to countercyclicality of consumption. The model with habits, for EIS both greater or less than one, can generate countercyclical behavior in labor hours. The habit model also produces much less amplification of TFP shocks, as compared to amplification in the no-habit model.

In terms of the welfare cost of aggregate volatility, they range from nil under expected utility to 1.3% of consumption for the FORA specification of risk preferences. The latter number is robust to variation in the EIS and the presence or absence of habits.

Other papers have examined the business cycle properties of models featuring non-standard risk preferences. Campanale *et al.*, for example, look at preferences that exhibit "disappointment aversion" (as in Gul [16]) in a production economy with capital adjustment costs. A version of their model, calibrated to match experimental evidence on the degree of disappointment aversion and aggregate data on the volatility of output and consumption growth, is able to match well the first and second moments of asset return data. Campanale *et al.* do not incorporate elastic labor effort, however, and the results below suggest this feature is important for determining whether a production economy can generate plausible business cycle behavior and plausible asset returns.

The paper most akin to this one is that of Tallarini [35]. Working in a standard RBC framework, but with recursive preferences of the Epstein-Zin [11] form, Tallarini shows that when the EIS is equal to one, the intertemporal preferences that result are formally analogous to the type of objective functions used in the literature on risk-sensitive optimal control (see [39])—*i.e.*, the social planner's problem that results takes the form of a discounted linear exponential quadratic Gaussian control problem. In contrast to standard linear-quadratic control, risk-sensitive control does not entail certainty equivalence—decision rules, and not just value functions, may depend on the variances of the model's shocks. Tallarini uses methods developed by Hansen and Sargent [18] to calculate approximate solutions to the model. He then shows that the model's business cycle predictions—relative volatilities and correlations—are little changed when the coefficient of relative risk aversion is increased from 1 to 100, though this has a significant impact on the model's welfare implications. The analysis here differs from Tallarini's in a couple important ways. First, since I solve the model using Chebyshev collocation, I am able to consider EIS values different from one; as noted above, variation in the elasticity of intertemporal substitution has important effects on the model's dynamics. Also, the model here incorporates habit formation and capital adjustment costs, features which have become common in many DSGE models.

2 Model

The model is the standard real business cycle core of many recent DSGE models.[2] I will describe the model in the most general form I employ—including external habit formation and capital adjustment costs—though for several of the numerical exercises, I will shut those channels down.

A unit measure of identical agents have preferences over consumption and leisure given by the recursive form

$$U_t = [(1-\beta)((C_t - \phi H_t)^{\psi}(1-N_t)^{1-\psi})^{\rho} + \beta \mu_t(U_{t+1})^{\rho}]^{1/\rho}$$

where C_t is consumption at date t, H_t is the consumption habit stock, N_t is the fraction of the unit time endowment devoted to work, and $\mu_t(\,\cdot\,)$ is a linearly homogeneous certainty equivalent operator conditional on information at date t. This expression holds for $\rho < 1$ and $\rho \neq 0$. When $\rho = 0$, the recursion takes the form

$$U_t = ((C_t - \phi H_t)^{\psi}(1-N_t)^{1-\psi})^{1-\beta} \mu_t(U_{t+1})^{\beta}.$$

The habit stock is assumed to evolve (externally) according to

$$H_{t+1} = (1-\delta_h)H_t + C_t^a.$$

In the last expression, we understand C_t^a to denote aggregate consumption—identical to individual consumption in equilibrium, but taken as given by individual agents. (The artifice here is that our representative agent is assumed atomistic, but identical to all other atomistic agents.)

The parameter ρ in the agent's utility function governs attitudes towards intertemporal substitution along constant paths of consumption and leisure. In particular, $\epsilon \equiv 1/(1-\rho)$ is the elasticity of intertemporal substitution (EIS) with respect to the composite good ("felicity") $(C - \phi H)^{\psi}(1-N)^{1-\psi}$. This differs from the elasticity of intertemporal substitution in consumption, which is given by $1/(1-\rho\psi)$. For a given value of $\psi \in (0, 1)$, the EIS for felicity and the EIS for consumption are both equal to 1 at $\rho = 0$, less than 1 for $\rho < 0$ and greater than 1 for $\rho > 0$. The parameter ψ governs the steady state allocation of time between labor effort and leisure. Note that the agent's utility will be homogeneous of degree ψ in (equilibrium) consumption.

Outside of the case where the EIS is unity, intratemporal preference that are non-separable in consumption and leisure are essential for consistency with balanced growth—see [21] or [2].

The certainty equivalent operator μ_t takes the form

$$\mu_t(U_{t+1}) = \hat{\mathbb{E}}_t[U_{t+1}^{1-\theta}]^{1/(1-\theta)} \tag{1}$$

where $\theta \geq 0$. In the case of $\theta = 1$, this expression becomes $\mu_t(U_{t+1}) = \exp[\hat{\mathbb{E}}_t \ln(U_{t+1})]$. The hat over the expectations operator $\mathbb{E}_t$ denotes a possible

[2] See, for example, the New Keynesian DSGE models of Goodfriend and King [14], Erceg *et al.* [13], or Smets and Wouters [32].

distortion in the subjective probabilities used (relative to the objective probabilities implied by the stochastic process for U_t). In particular, we allow for the sort of distorted probabilities associated with both Quiggin's [27] "anticipated utility" theory and Yaari's [40] "dual theory" of choice under uncertainty.

Quiggin's and Yaari's theories incorporate rank dependence—their certainty equivalents depend on the ranking of outcomes from worst to best. Suppose that a random variable ξ takes on values in $\{\xi(1), \xi(2), \ldots, \xi(k)\}$, $\xi(1) \leq \xi(2) \leq \cdots \leq \xi(k)$, with (objective) probabilities $p(1), p(2), \ldots, p(k))$. Then, for $\gamma \in (0, 1]$, the Quiggin-Yaari expectation of ξ is given by

$$\hat{\mathrm{E}}[\xi] = \sum_{i=1}^{k} \{ (\sum_{h=1}^{i} p(h))^{\gamma} - (\sum_{h=1}^{i-1} p(h))^{\gamma} \} \xi(i)) \tag{2}$$

where $\sum_{h=1}^{-1} p(h) \equiv 0$. For the case of two outcomes, for example, this expectation in effect replaces the probability of the worse outcome, $p(1)$, with $p(1)^{\gamma}$—which exceeds $p(1)$ if $\gamma < 1$—and the probability of the better outcome, $p(2) = 1 - p(1)$ with $1 - p(1)^{\gamma}$.

Risk preferences of the Quiggin/Yaari form fall under the more general heading of risk preferences that display 'first-order risk aversion': the risk premia they generate, over timeless gambles, are proportional to the gamble's standard deviation, rather than its variance. In the context of a Lucas tree economy, Epstein and Zin [12] explored the extent to which preferences of this sort provide a resolution to the equity premium puzzle.

Section A.1 of the appendix provides some background on this form for risk preferences.

When $\theta \neq 0$ in (1) and $\gamma \neq 1$ in (2), $\mu_t(U)$ incorporates aspects of preferences featuring both first-order risk aversion and more standard constant relative risk aversion (CRRA). If $\gamma = 1$, we are in the case of Epstein-Zin, or Epstein-Zin/Weil, preferences. If, further, $\theta = 1/\epsilon$—*i.e.*, if $\theta = 1 - \rho$—we obtain the case of expected utility (EU).

Agents face a sequence of budget constraints of the form

$$W_t N_t + R_t K_t \geq C_t + X_t$$

where X_t denotes gross investment—forgone consumption allocated to capital accumulation—and W_t and R_t are the real wage and rental rate. An agent's stock of capital evolves according to

$$K_{t+1} = (1 - \delta_k) K_t + K_t g(X_t / K_t),$$

where $g(\,\cdot\,)$ is a concave adjustment cost function (to be described in more detail below).

A representative firm operates a constant returns to scale production technology of the Cobb-Douglas form

$$Y_t = A_t K_t^{\alpha} N_t^{1-\alpha},$$

where Y_t is output available for consumption or gross investment, and A_t is a stochastic process for total factor productivity.[3] The competitive equilibrium wage and rental rates will be given by the marginal products of labor and capital, $W_t = (1-\alpha)A_t K_t^{\alpha} N_t^{-\alpha}$ and $R_t = \alpha A_t K_t^{\alpha-1} N_t^{1-\alpha}$.

The economy's resource constraint is

$$Y_t \geq C_t + X_t.$$

The TFP process A_t is assumed to contain a deterministic trend component—*i.e.*, $A_t = \eta^t a_t$. Quantities (apart from labor hours) and the real wage rate will grow, on average, at the rate η. To render the model stationary, we deflate those variables by $\{\eta^t\}$, and denote the deflated variables by lower-case letters—*i.e.*, $c_t = C_t/\eta^t$, $x_t = X_t/\eta^t$, *etc.*

With these definitions, the recursive representation of the agent's utility process takes the form

$$u = [(1-\beta)[(c-\phi h)^{\psi}(1-N)^{1-\psi}]^{\rho} + \hat{\beta}\mu(u' \mid \Omega)^{\rho}]^{1/\rho}$$

for $\rho \neq 0$, or

$$u = ((c-\phi h)^{\psi}(1-N)^{1-\psi})^{1-\beta}\mu(u' \mid \Omega)^{\beta}$$

for $\rho = 0$, where $\hat{\beta} = \beta\eta^{\psi\rho}$ and Ω denotes the agent's conditioning information. The budget constraint becomes

$$wN + Rk \geq c + x,$$

while the equations describing the evolutions of the capital and habit stocks are modified as follows

$$\eta k' = (1-\delta_k)k + kg(x/k) \tag{3}$$

$$\eta h' = (1-\delta_h)h + c^a. \tag{4}$$

Finally, the economy-wide production constraint is

$$y = ak^{\alpha}N^{1-\alpha} \geq c + x. \tag{5}$$

3 Equilibrium: Characterization and solution

3.1 Characterization

The model's equilibrium paths solve a restricted social planner's problem—maximizing utility subject to the economy's technology and resource constraints, and the evolution equation for the physical capital stock, taking the evolution of the habit stock as given. After obtaining the intratemporal and intertemporal first-order conditions for that problem, equality between c_t and c_t^a is imposed.

[3] Here and throughout, where there is no possibility of confusion, we suppress the superscript a on aggregate quantities.

The intratemporal first-order condition is

$$\frac{1-\psi}{\psi}\frac{c-\phi h}{1-N} = (1-\alpha)ak^{\alpha}N^{-\alpha} \tag{6}$$

Let $mu_c \equiv (c-\phi h)^{\psi\rho-1}(1-N)^{(1-\psi)\rho}$, which is the marginal utility of current consumption (divided by $(1-\beta)U^{1-\rho}$); $mp_k \equiv \alpha a k^{\alpha-1}N^{1-\alpha}$, the marginal product of capital; and $z \equiv x/k$, the investment rate. Also, let

$$Dg(z) \equiv \partial G(z)/\partial z,$$

as we reserve the prime symbol ($'$) to denote next-period values of the various model variables. With this notation, the intertemporal first-order condition is an Euler equation of the form

$$\eta\frac{mu_c}{Dg(z)} = \hat{\beta}\mu(v' \mid \Omega)^{\theta+\rho-1}\times$$
$$\hat{\mathbb{E}}[(v')^{1-\rho-\theta}mu_c'(mp_k' + \frac{1-\delta_k+g(z')-z'Dg(z')}{Dg(z')}) \mid \Omega]. \tag{7}$$

In the last expression, v denotes the agent's value function, evaluated at equilibrium quantities, which we could write more explicitly as $v(a,k,h)$.

v satisfies the following Bellman equation

$$v(a,k,h) = [(1-\beta)[(c-\phi h)^{\psi}(1-N)^{1-\psi}]^{\rho} + \hat{\beta}\mu(v(a',k',h') \mid \Omega)^{\rho}]^{1/\rho}, \tag{8}$$

when c, x and N are optimally chosen, given the state (a,k,h), the economy's production constraint, and the evolution equation for k, and when h' is given by $\eta h' = (1-\delta_h)h + c$.

A solution to the model consists of a value function $v(a,k,h)$, and decision rules for the labor input, $N(a,k,h)$, consumption, $c(a,k,h)$, and gross investment, $x(a,k,h)$, such that at all (a,k,k): $N(a,k,h)$ and $c(a,k,h)$ satisfy (6); $N(a,k,h)$, $c(a,k,h)$, and $x(a,k,h)$ satisfy the resource constraint (5); and, when k' and h' are derived from (3) and (4), given $c(a,k,h)$ and $x(a,k,h)$, these mappings satisfy the Euler equation (7) and the Bellman equation (8).

If the economy had a conditionally risk-free asset, we could calculate its rate of return, $R^f(a,k,h)$, in standard fashion:

$$\frac{1}{1+R^f(a,k,h)} = \eta^{-1}\hat{\beta}\mu(v' \mid \Omega)^{\theta+\rho-1}\hat{\mathbb{E}}[(v')^{1-\rho-\theta}\frac{mu_c'}{mu_c} \mid \Omega]$$

When $\theta = 1-\rho$ and $\gamma = 1$ (so $\hat{\mathbb{E}} = \mathbb{E}$), this reduces to the familiar expression relating the risk-free rate to the expected growth rate of the marginal utility of consumption.

For comparison with the risk-free rate, we take the model's "equity" return to be the return on a marginal unit of gross investment; the expected risky return is then

$$1+R^e(a,k,h) = \mathbb{E}[Dg(z)(mp_k' + \frac{1-\delta_k+g(z')-z'Dg(z')}{Dg(z')}) \mid \Omega]$$

The model's equity premium is then $R^e(a,k,h) - R^f(a,k,h)$ at state (a,k,h).

3.2 Solution method

Our solution method follows the Chebyshev polynomial collocation method described by Caldara *et al.* [4].

As in their approach, consideration of the conditions describing a solution to the model reveals that a solution really consists in two mappings, a value function $v(a,k,h)$ and a labor supply decision rule $N(a,k,h)$. Given an arbitrary $N(a,k,h)$, decision rules for consumption and gross investment can be backed out of the purely intratemporal conditions (6) and (5). Those decision rules—call them $c(a,k,h)$and $x(a,k,h)$—imply transition laws for the two endogenous states, k and h, according to

$$k'(a,k,h) = (1-\delta_k)k + kg(x(a,k,h)/k)$$

and

$$h'(a,k,h) = (1-\delta_h)h + c(a,k,h).$$

These decision rules and transition laws are feasible and satisfy the agent's intratemporal optimality condition. The rules that actually solve the model are then found by further requiring that the Euler equation and Bellman equation hold at all (a,k,h).

Our solution method approximates the labor decision rule and value function with Chebyshev polynomials. The polynomial coefficients solve the system of nonlinear equations determined by the Euler equation and Bellman equation, when those equations are evaluated at a finite number of polynomial roots (the collocation points).

As in Caldara *et al.*, the exogenous TFP process a is assumed to be given by a finite-state Markov chain, and the collocation points along the a dimension are assumed fixed at the values of the Markov chain states. That is, for all t, we assume $a_t \in \{a_1, a_2, \ldots, a_{\#a}\}$, and there is a $\#a \times \#a$ matrix P with

$$\Pr\{a_{t+1} = a_j \mid a_t = a_i\} = P_{ij}.$$

I then treat the labor decision rule and value function as vector-valued functions at each point (k,h), $N_i(k,h) = N(a_i,k,h)$ and $v_i(k,h) = v(a_i,k,h)$. For each i, $N_i(k,h)$ and $v_i(k,h)$ are approximated as tensor products of Chebyshev polynomials in k and h. Precisely,

$$N_i(k,h) \approx \sum_{l=0}^{O_k} \sum_{m=0}^{O_h} D_N^i(l,m) T_l(\iota(k)) T_m(\iota(h)) \equiv \mathcal{N}(k,h;D_N^i),$$

where $T_l(\,\cdot\,)$ is the lth order Chebyshev polynomial of the first kind; $\{D_N^i(l,m) : l = 0,1,\ldots O_N, m = 0,1,\ldots O_h\}$ are the polynomial coefficients; and $\iota(\,\cdot\,)$ maps the domains of k and h into $[-1,1]$. Likewise,

$$v_i(k,h) \approx \sum_{l=0}^{O_k} \sum_{m=0}^{O_h} D_v^i(l,m) T_l(\iota(k)) T_m(\iota(h)) \equiv \mathcal{V}(k,h;D_v^i).$$

While in principle, the polynomial orders O_k and O_h could differ (and differ across the approximations $\mathcal{N}$ and $\mathcal{V}$), in practice, I use a single polynomial order, O, for both the k and h dimensions, for both $\mathcal{N}$ and $\mathcal{V}$.[4]

While ocular tests aren't necessarily conclusive, the model's impulse responses (described below) changed only negligibly in moving from a polynomial order of $O = 6$ to $O = 7$, while the time required to solve the model did begin to increase steeply at the point. Hence, I opted for a choice of $O = 6$ in all the numerical experiments.[5] A tensor product of 6th-order polynomials contains 7^2 terms, so an approximation with $O = 6$ entails determining $\#a \times 49$ coefficients for each of $\mathcal{N}$ and $\mathcal{V}$. With $\#a = 9$ (a choice described in the next section), this means determining $2 \times 9 \times 49 = 882$ coefficients for the full model (when habit formation is present).

With the approximating decision rules in hand, sequences of model quantities and prices can be produced simply by feeding in TFP sequences together with initial values for k and h.

4 Parameter values

I assign parameter values through a mix of calibration and direct specification (with some appeal to existing literature). The parameters governing intertemporal substitution and risk aversion, ρ, γ and θ, are all set directly, without a view toward matching any particular moments, since our numerical experiments will consist mostly in varying these three parameters.

Capital's share in the Cobb-Douglas production function is set to $\alpha = 0.4$ and the capital depreciation rate, δ_k is set directly to 10 percent per year. The growth rate of the deterministic trend in TFP, η, is set to 1.8 percent per year. (The model's period is one quarter.)

The habit formation parameters δ_h and ϕ are set somewhat arbitrarily. The habit depreciation rate is set to $\delta_h = 1$, so $h_t = c_{t-1}$. I set $\phi = 0.5$. One sees higher values in the literature, though so far as I can tell only in models solved via linear approximation methods. For values much above $\phi = 0.5$, our Chebyshev collocation algorithm typically could not find coefficient matrices to simultaneously satisfy the Euler and Bellman equations at all collocation points.[6] In any case, $\phi = 0.5$ seems in line with the available evidence, for example Naik and Moore [25].

I also consider the no-habit case of $\phi = 0$.

Capital adjustment costs of the form used here, where k' depends on $kg(x/k)$, have been used by Baxter and Crucini [3], Jermann [20], and others. In some

[4] All the MATLAB codes for this paper are available at `http://www.jimdolmas.net/economics`.

[5] The robustness of all the results to this choice should be a subject of further study.

[6] This obviously needs further analysis. Since I don't have an existence result, I do not wish to make too much out of this finding, but on the face of it, it would not be that surprising that a model where utility depends on $c_t - \phi h_t$, where the technology is subject to TFP shocks, and where investment is not reversible, might fail to have a solution for ϕ too large.

of those models, solved by linear approximation, calibration only requires an elasticity of the first derivative of the adjustment cost function (which is the inverse of the elasticity of the investment rate z with respect to Tobin's Q). Our approach here requires a functional form for g. The desiderata are that g should be (weakly) concave, and it should be possible to calibrate it so that there are no adjustment costs in the model's deterministic steady state—*i.e.*, that $g(z) = z$ and $Dg(z) = 1$ should be possible for an appropriate choice of parameters. To that end, following Jermann [20], I specify

$$g(z) = (b_0/b_2)z^{b_2} - b_1$$

for $b_0, b_1, b_2 > 0$, $b_2 \leq 1$. The two steady state conditions $g(z) = z$ and $Dg(z) = 1$ cannot, of course, pin down all three parameters. I therefore set the curvature parameter b_2 and, given the steady state investment rate, call it $\bar{z}$, set $b_0 = \bar{z}^{1-b_2}$ and $b_1 = \bar{z}/b_2$.

I consider two cases for b_2, $b_2 = 1$ (no adjustments costs, in conjunction with no habit formation) and $b_2 = 0.5$ (high adjustment costs). The latter value implies an elasticity of the rate of investment with respect to Tobin's Q of 2.

Given these choices, there remain two parameters to select—ψ, the relative weight of consumption in felicity, and β, the utility discount factor. These are chosen to give a plausible consumption share of output in the deterministic steady state (73 percent) and a plausible allocation of time to work effort (30 percent).[7] Note that in the deterministic steady state,

$$\bar{h} = \frac{\bar{c}}{\eta - 1 + \delta_h}$$

so that

$$\bar{c} - \phi\bar{h} = \left(1 - \frac{\phi}{\eta - 1 + \delta_h}\right)\bar{c} \equiv \Phi\bar{c}.$$

The agent's intratemporal first-order condition (6) then implies

$$\frac{1-\psi}{\psi} = (1-\alpha)\frac{1-\bar{N}}{\bar{N}}\frac{1}{\Phi s_c},$$

so there is a unique value of ψ consistent with given targets for $\bar{N}$ and s_c (and the previous parameter choices). Likewise, given that there are no adjustment costs in the steady state, a target for s_c implies a target for the steady state capital-output ratio. The steady state version of the agent's Euler equation—

$$\eta = \beta\eta^{\rho\psi}\left(\frac{\alpha}{\bar{k}/\bar{y}} + 1 - \delta_k\right)$$

—then determines a unique value of β consistent with these targets (and the other parameter choices).

[7] $\bar{N} = 0.3$ is a standard choice, and $\bar{c}/\bar{y} \equiv s_c = 0.73$ is consistent with the average GDP share of private consumption of nondurables and services plus government consumption for the US over the last 50 years.

The value of β depends on the assumed elasticity of intertemporal substitution, ϵ. In the numerical experiments I report on below, I consider two values for ϵ, $\epsilon = 0.5$ and $\epsilon = 15$. While the latter value may seem high, it yields an intertemporal elasticity of substitution in consumption on the order of 1.3 in the habit model.

Table 1 contains a summary of the parameter values I employ.

	Value	Remarks
Technology parameters:		
α	0.4	Standard
δ_k	0.0127	10% annual
η	1.0045	1.8 % annual
Habit formation parameters:		
ϕ	0.5	Habit strength
δ_h	1	$h_t = c_{t-1}$
Capital adjustment cost function:		
b_0	0.1312	$g(\bar{z}) = \bar{z}, Dg(\bar{z}) = 1$
b_1	0.0172	$g(\bar{z}) = \bar{z}, Dg(\bar{z}) = 1$
b_2	0.5	Q elasticity of $z = 2$
Parameters set indirectly:		
ψ	0.2075	Gives $\bar{N} = 0.3, s_c = 0.73$
$\beta\,(\epsilon = 0.5/\epsilon = 15)$	0.9927/0.9910	Gives $\bar{k}$ such that $s_c = 0.73$
If no habit or adjustment costs:		
ψ	0.3427	Gives $\bar{N} = 0.3, s_c = 0.73$
$\beta\,(\epsilon = 0.5/\epsilon = 15)$	0.9933/0.9904	Gives $\bar{k}$ such that $s_c = 0.73$

Table 1: A summary of model parameter values

Lastly, I calibrate the parameters of the Markov chain describing the stationary part of TFP using Rouwenhorst's [30] method to approximate an AR(1) process, in logarithms, with persistence parameter 0.95 and conditional standard deviation 0.07 (the numbers used by Cooley and Prescott [9]).[8] I assume the Markov chain has nine states. After all other parameters are set, I fix the unconditional mean of a so that output is equal to one unit in the deterministic steady state (which replaces a_t with $\mathbb{E}(a_t)$ at all dates).

Note that, set in this manner, the mean level of stationary TFP a does not depend on the parameter choices for intertemporal substitution, risk aversion, habit or capital adjustment costs. Experiments where any of these parameters are varied will thus take place using technology with the same average level of TFP.

[8]Kopecky and Suen [22] show that Rouwenhorst's method is superior to the more-common method of Tauchen [36] for approximating highly persistent AR processes.

5 Results

5.1 Features of interest

There are a number of features of the model output we could examine. Except for the impulse responses, the construction of which I describe momentarily, all the characterizations I will make are based on a draw of 10,100 Markov chain states consistent with the calibrated transition matrix P. Physical capital and the habit stock (when present) are initialized at their deterministic steady state values, but I discard the first 100 observations of all simulated series. The sequence of Markov chain states is the same in all the numerical experiments (again, apart from the impulse responses).

5.1.1 Means

First moments typically aren't that interesting in models of this sort—means of variables in the stochastic economies rarely depart by much from their values in the deterministic steady state, values which the models have typically been calibrated to reproduce.

This is not the case when first-order risk aversion is present, as agents engage in a significant amount of precautionary capital accumulation. This behavior is, in fact, one of a small number of aspects in which the model's behavior under first-order risk aversion is distinguished from risk aversion which is very high, but second-order ($\gamma = 1$ but θ very large).

From an asset-pricing standpoint, this is not a positive for first-order risk averse preferences: the higher long-run capital stocks they generate push down the rental rate on physical capital, hence also the average equity return.

5.1.2 Second moments

Obviously, we are interested in the standard second-moment measures, which I calculate after applying the Hodrick-Prescott filter to the simulated, logged model data (after restoring deterministic trends to previously deflated variables).

I calculate the standard measures of volatility and volatility relative to output, as well as volatility relative to the TFP shock process. Cross-correlations of variables with output are of particular interest, since these—while generally not sensitive to risk aversion—are sensitive with respect to assumptions about the elasticity of intertemporal substitution (a high enough EIS makes consumption countercyclical) and the presence or absence of habits (habit can make labor effort countercyclical).

For brevity's sake, I omit results on the autocorrelation of output (or phase relations between output and other variables). As in the standard RBC model, output here, regardless of assumptions about risk aversion, simply inherits the dynamics of the TFP shock (see, for example, [38] on the standard RBC model).

5.1.3 Welfare

Since one can calculate the welfare cost of volatility in different ways, it is important to be precise about whatever calculation is being made. In this paper, I use two welfare measures. Let $(\bar{k}, \bar{h})$ denote physical capital and the habit stock in the deterministic steady steady, and $(\hat{k}, \hat{h})$ their mean values in the stochastic economy. Recall that the deterministic economy is identical to the stochastic economy, with the exception of the TFP process $\{a_t\}$ being replaced with its mean value.

I then compare $\mu[v(a, \hat{k}, \hat{h})]$, the unconditional certainty equivalent lifetime value of an agent in the stochastic economy, who holds the long-run average capital stock and habit for the stochastic economy, to two values: one is the lifetime value of an agent in the deterministic economy, holding $(\hat{k}, \hat{h})$—call it $v_D(\hat{k}, \hat{h})$—and the other is the lifetime value of an agent in the deterministic economy holding $(\bar{k}, \bar{h})$, $v_D(\bar{k}, \bar{h})$.

$\mu[v(a, \hat{k}, \hat{h})]$ is the certainty equivalent value of an agent who parachutes into the stochastic economy, holding average stocks for the stochastic economy, but unaware of the current TFP state. I then ask what across-the-board increase in consumption, at all dates and states, would make the agent indifferent between making that parachute drop and residing in the deterministic economy, holding either the average stocks appropriate to the stochastic economy or the steady state stocks of the deterministic economy. Precisely, because lifetime utility is homogeneous of degree ψ in (equilibrium) consumption, I calculate the values λ_1 and λ_2 that satisfy:

$$\left(1 + \frac{\lambda_1}{100}\right)^{\psi} \mu[v(a, \hat{k}, \hat{h})] = v_D(\hat{k}, \hat{h})$$

and

$$\left(1 + \frac{\lambda_2}{100}\right)^{\psi} \mu[v(a, \hat{k}, \hat{h})] = v_D(\bar{k}, \bar{h}).$$

5.1.4 Impulse responses

I report a few impulse responses to TFP shocks, for different variants of the model, below. Given that the shock process driving the model is discrete, the calculation of impulse responses merits some discussion.

Similar to the approach in Campanale *et al.* [5], I run 10,000 simulations of 120 periods each, each simulation starting from the average stocks $\hat{k}$ and $\hat{h}$, and from the lowest TFP level, $a(1)$. I then calculate the average path for each variable, over the 10,000 runs. For any variable b, letting $\{b_t : t = 1, \ldots 120\}$ denote the average path and $\hat{b}$ the variable's long-run mean value, the numbers I plot are $\{100 \log(b_t/\hat{b}) : t = 1, \ldots 120\}$.

Of particular interest is the impulse response of hours worked in some variants of the habit model: hours can actually rise at impact in response to a negative shock. With regard to the role of risk preferences, though, we'll see that the model's impulse responses change only negligibly as we make significant changes in the agent's risk attitudes.

5.2 No habits, no adjustment costs

Some features of the model can be illustrated readily, abstracting from habit formation and capital adjustment costs (setting $\phi = 0$ and $b_2 = 1$).

I consider two values of the elasticity of intertemporal substitution, $\epsilon = 1/(1-\rho)$, $\epsilon = 0.5$ and $\epsilon = 15$. This range may seem wide, but the former implies an EIS in consumption of roughly two-thirds, the latter an EIS in consumption of roughly 1.3 (or 1.5, if no habit is present).

Why consider EIS's greater than one? Models with an EIS greater than one have, until recently, been little-studied—most likely because justifying an EIS in consumption even as big as one has proven difficult, empirically.[9] Also, to the extent that earlier work assumed expected utility—so that $\theta = 1/\epsilon$—higher values of the EIS necessitated lower levels of risk aversion.

The assumption of an EIS greater than one has become standard, though, in at least one strand of the literature, that studying models of "long-run risk" (for example [1] or [10]). In fact, for those models an EIS in consumption around 1.5 is standard, though to my knowledge no one has studied the ramifications of this assumption in a long-run risk model with elastic labor supply. While our model here is not one featuring long-run risk in consumption or TFP, its behavior when the EIS in consumption is around 1.5 may still be of interest to that literature.

In combination with those two values for ϵ, I consider three cases for the parameter pair (θ, γ) governing risk aversion. Case 1 is expected utility (EU): $\theta = 1/\epsilon$ and $\gamma = 1$. Case 2 is high constant relative risk aversion, as in Tallarini [35]: $\theta = 100$ and $\gamma = 1$. Finally, case 3 incorporates first-order risk aversion: $\gamma = 0.9$ and $\theta = 1$.

Table 2 presents some first moment results for the no-habit economy, in particular a measure of precautionary capital accumulation, and average returns on physical capital and the hypothetical riskless asset, defined above in section 3.1. Precautionary capital accumulation is noticeably higher under first-order risk aversion, particularly in the low-ϵ case, where the long-run average capital stock exceeds the deterministic steady state stock by 1.4%.

Note that average asset returns vary only very slightly with changes in risk aversion—consistent with the findings in Tallarini [35]—or, for that matter, intertemporal substitution. Equity risk premia are either nil or even slightly negative in all specifications. Even under high risk aversion, the ability to smooth the effects of shocks through capital accumulation—and to smooth variations in marginal utility through variation in leisure hours—essentially undoes the results, for an endowment economy, found in Epstein and Zin [12].

[9] See Guvenen [17] for a discussion of issues in the measurement of the EIS, and a potential resolution to the conflict between studies finding near-zero values of the EIS, using aggregate consumption data, and business cycle models that typically assume an EIS around 1.0.

	$\epsilon = 0.50$		
	$\%(k/\bar{k})$	R^e	R^f
EU, $\theta = 1/\epsilon, \gamma = 1$	-0.0675	5.2054	5.2019
High CRRA, $\theta = 100, \gamma = 1$	0.7372	5.1614	5.2063
FORA, $\theta = 1, \gamma = 0.9$	1.4267	5.1560	5.2019

	$\epsilon = 15$		
	$\%(k/\bar{k})$	R^e	R^f
EU, $\theta = 1/\epsilon, \gamma = 1$	-0.0386	5.2037	5.2007
High CRRA, $\theta = 100, \gamma = 1$	0.5568	5.1711	5.1965
FORA, $\theta = 1, \gamma = 0.9$	0.8359	5.1560	5.2007

Table 2: Some first moments, no-habit model. Column 1, $\%(k/\bar{k})$, gives the excess, in percent, of average capital in the stochastic model compared to the deterministic steady state capital stock. Average asset returns are in annualized percent.

Some second moment implications are given in Table 3. As noted above, for these calculations, I reintroduce the linear time trend η^t to the deflated model variables, then apply the Hodrick-Prescott filter to the natural logarithms of all the simulated series.

The important features to note in Table 3 are, firstly, the robustness of the numbers (for a given ϵ) across the different specifications of risk preferences, and secondly, the dramatic difference in the moments (for any risk preference configuration) across the two values of ϵ. Of particular note is the negative contemporaneous correlation between logged output and consumption, in the case where ϵ is bigger than one.

The countercyclical behavior of consumption in the high EIS case is apparent in the model's impulse responses for consumption, shown in the two panels of figure 1. The top panel plots the impulse responses of consumption in the low EIS case for all three risk preference cases—the paths overlap to the point of being indistinguishable. The lower panel of the figure shows the impulses responses in the high EIS case. In that case, consumption is above its long-run average for two quarters, before ultimately falling to a level nearly 5 percent below its long-run average, then converging gradually back.[10]

We're also interested in the extent to which the model provides any amplification of the shocks fed into it. For the lower EIS value, the standard deviation of logged output is roughly 1.40 times the standard deviation of the TFP shock, for all risk parameter settings. For the larger EIS value—consistent with the

[10] The non-separability of utility in consumption and leisure appears to be important here. Using the parameters ψ and ρ, and the steady state labor effort $\bar{N}$, one can combine the impulse responses for consumption and hours (not shown) to obtain an approximate impulse response for the marginal utility of consumption. Despite the increase in consumption during the first two quarters, the marginal utility of consumption—which depends positively on leisure—definitely rises in response to the shock.

		$\epsilon = 0.50$						
		$\ln(Y)$	$\ln(C)$	$\ln(X)$	$\ln(N)$	$\ln(W)$	$ln(1+R^e)$	$\ln(1+R^f)$
EU, $\theta = 1/\epsilon, \gamma = 1$	$100\sigma(\cdot)$	1.2936	0.4152	3.7317	0.6274	0.6737	0.0304	0.0487
	$\sigma(\cdot)/\sigma(Y)$	1.0	0.3210	2.8848	0.4850	0.5208	0.0235	0.0376
	$corr(\cdot, Y)$	1.0	0.9710	0.9965	0.9938	0.9946	0.9834	0.9910
High CRRA, $\theta = 100, \gamma = 1$	$100\sigma(\cdot)$	1.2941	0.4141	3.7225	0.6282	0.6374	0.0303	0.0486
	$\sigma(\cdot)/\sigma(Y)$	1.0	0.3200	2.8766	0.4854	0.5203	0.0234	0.0376
	$corr(\cdot, Y)$	1.0	0.9708	0.9965	0.9938	0.9946	0.9835	0.9910
FORA, $\theta = 1, \gamma = 0.9$	$100\sigma(\cdot)$	1.2982	0.4083	3.7436	0.6352	0.6706	0.0303	0.0488
	$\sigma(\cdot)/\sigma(Y)$	1.0	0.3145	2.8836	0.4893	0.5166	0.0233	0.0376
	$corr(\cdot, Y)$	1.0	0.9694	0.9964	0.9939	0.9945	0.9834	0.9910
		$\epsilon = 15$						
		$\ln(Y)$	$\ln(C)$	$\ln(X)$	$\ln(N)$	$\ln(W)$	$ln(1+R^e)$	$\ln(1+R^f)$
EU, $\theta = 1/\epsilon, \gamma = 1$	$100\sigma(\cdot)$	1.7051	0.3700	6.8650	1.3143	0.4774	0.0388	0.0503
	$\sigma(\cdot)/\sigma(Y)$	1.0	0.2170	4.0261	0.7708	0.2800	0.0228	0.0295
	$corr(\cdot, Y)$	1.0	−0.3801	0.9871	0.9832	0.8649	0.9686	0.9805
High CRRA, $\theta = 100, \gamma = 1$	$100\sigma(\cdot)$	1.7071	0.3716	6.8599	1.3176	0.4762	0.0388	0.0503
	$\sigma(\cdot)/\sigma(Y)$	1.0	0.2177	4.0185	0.7718	0.2790	0.0227	0.0295
	$corr(\cdot, Y)$	1.0	−0.3880	0.9872	0.9833	0.8641	0.9687	0.9806
FORA, $\theta = 1, \gamma = 0.9$	$100\sigma(\cdot)$	1.7145	0.3780	6.9098	1.3300	0.4729	0.0389	0.0507
	$\sigma(\cdot)/\sigma(Y)$	1.0	0.2205	4.0302	0.7757	0.2758	0.0227	0.0296
	$corr(\cdot, Y)$	1.0	−0.4090	0.9871	0.9834	0.8598	0.9686	0.9806

Table 3: Select second moments, no-habit model. Simulated data have been logged and HP-filtered, as described in the text. Entries for $corr(\cdot, Y)$ are variables' contemporaneous correlations with output.

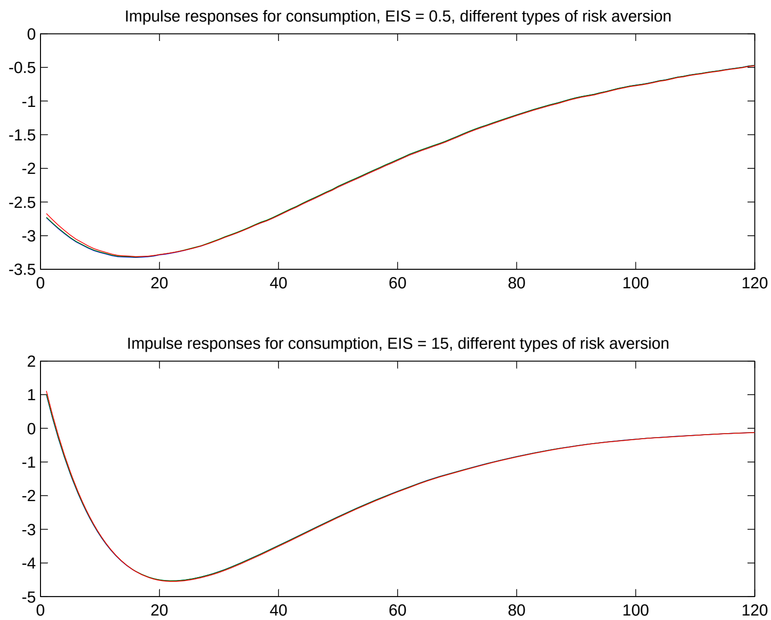

Figure 1: Impulse responses for consumption in the no-habit model. See section 5.1.4 for details on the calculations involved.

greater volatility of hours reported in table 3—we see more amplification, with the standard deviation of output roughly 1.85 times the standard deviation of the TFP process.

While attitudes toward risk have only a negligible impact on the model's second moments, they have a very significant effect on the perceived cost of aggregate volatility. Table 4 reports the welfare measures λ_1 and λ_2 defined in section 5.1.3. Under expected utility—as we've come to expect from Lucas's [24] calculations—the agent's welfare gain from eliminating volatility is essentially nil. In the high CRRA case, the gains are in the range of $0.3 - 0.4\%$ of consumption, which would be on the order of $40 billion annually, if we consider current US consumption expenditures of around $10 trillion annually.

	$\epsilon = 0.50$	
	λ_1	λ_2
EU, $\theta = 1/\epsilon, \gamma = 1$	0.01	0.02
High CRRA, $\theta = 100, \gamma = 1$	0.42	0.29
FORA, $\theta = 1, \gamma = 0.9$	1.30	1.05

	$\epsilon = 15$	
	λ_1	λ_2
EU, $\theta = 1/\epsilon, \gamma = 1$	-0.01	-0.00
High CRRA, $\theta = 100, \gamma = 1$	0.41	0.31
FORA, $\theta = 1, \gamma = 0.9$	1.30	1.14

Table 4: Welfare cost of fluctuations in the no-habit model. The compensation measures λ_1 and λ_2 are percentage increases in consumption at all dates and states, defined in section 5.1.3.

In the FORA case, the gains are roughly triple in size, for either EIS value. Again using $10 trillion as the value of aggregate consumption, the compensating differentials shown in table 4 for the FORA cases would be in the range of $110 − $130 billion, or nearly $1,000 per person per year.

Does the FORA specification with $\gamma = 0.9$ represent a ridiculously extreme level of risk aversion? In the appendix, section A.2, I give some calibration examples which argue that $\gamma = 0.9$ (with $\theta = 1$) is, in fact, roughly consistent with some of the observed behavior of individuals in insurance markets.

5.3 The full model with habits and adjustment costs

I now consider the full model with external habit formation ($\phi = 0.5, \delta_h = 1$) and capital adjustment costs ($b_2 = 0.5$). Table 5, analogous to Table 2, gives some select first moments from the simulated model data.

Similar to the no-habit version of the model, precautionary capital accumulation is greater when the elasticity of intertemporal substitution is set to its lower value, and at that setting, FORA risk preferences generate the greatest

	$\epsilon = 0.50$		
	$\%(k/\bar{k})$	R^e	R^f
EU, $\theta = 1/\epsilon, \gamma = 1$	−0.1130	5.2661	5.2064
High CRRA, $\theta = 100, \gamma = 1$	0.0586	5.2572	5.2078
FORA, $\theta = 1, \gamma = 0.9$	0.4659	5.2364	5.2064
	$\epsilon = 15$		
	$\%(k/\bar{k})$	R^e	R^f
EU, $\theta = 1/\epsilon, \gamma = 1$	−0.0998	5.2498	5.2029
High CRRA, $\theta = 100, \gamma = 1$	−0.0652	5.2480	5.2014
FORA, $\theta = 1, \gamma = 0.9$	−0.0800	5.2488	5.2029

Table 5: Some first moments, model with habit formation and capital adjustment costs. Column 1, $\%(k/\bar{k})$, gives the excess, in percent, of average capital in the stochastic model compared to the deterministic steady state capital stock. Average asset returns are in annualized percent.

amount of such accumulation. Nevertheless, with the long-run average capital stock in the stochastic economy just about 0.5% bigger than the deterministic steady state stock, the precautionary effect under habits is much weaker. For the higher EIS value of $\epsilon = 15$, long-run average capital stocks under all risk specifications are slightly lower than the deterministic steady state stock.

As far as asset returns are concerned, the second and third columns of Table 5 show that, again, neither risk preferences nor intertemporal substitution have much of an effect on either the expected capital return or the risk-free rate. Equity risk premia are negligible, in contrast to Jermann's [20] results in a fixed-labor-supply model otherwise similar to the model here.

Movements in labor effort are a key feature of the model with habits and adjustments, as illustrated in Table 6, which replicates for the full model the second-moment calculations for the no-habit model reported in Table 3. Several salient features of the model's second moments are worth noting:

- The model with habit formation and adjustment costs has much weaker amplification than the does the model without those features, as is apparent from the standard deviations of $\ln(Y)$ across all specifications. In fact, where the standard deviation of logged output was $1.40 - 1.85$ times the standard deviation of logged TFP in the stripped-down model, here the ratio is essentially one across all specifications of attitudes towards risk and intertemporal substitution.

- The volatility of labor hours—much lower in the habit model—is probably key to the first observation. That hours should be less volatile in the model with habits is, at least, intuitive: a period with low TFP may be a bad time to work, but hours cannot fall too much if consumption is to be maintained above the habit level.

- Hours are now only weakly correlated with output, with contemporaneous correlations of just over 0.5. (Below, we'll see an impulse response plot that suggests this contemporaneous correlation would actually be negative if we were characterizing linearly de-trended or first-differenced data, rather than HP-filtered data.)
- The specification of risk attitudes again has a negligible impact on the model second moments.
- Now, in contrast to what we saw in the stripped-down model, variation in the EIS also has a negligible impact on the results.

Figure 2 shows the model impulse responses for labor effort, for our two values of ϵ and all three specifications of risk preferences. As was the case in the stripped-down model, here again the series for each of the cases of risk preferences lie virtually on top of one another in each panel of the figure.

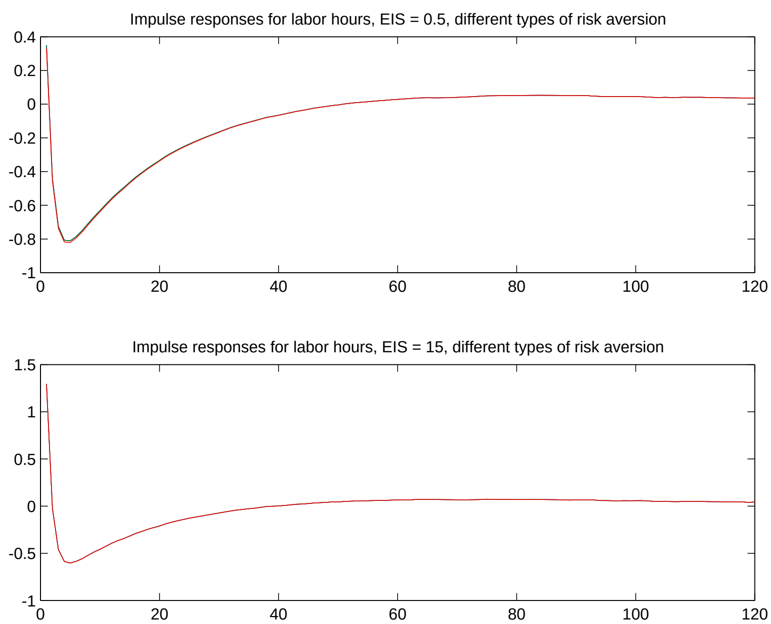

Figure 2: Impulse responses for labor effort in the model with habit formation and capital adjustment costs. See section 5.1.4 for details on the calculations involved.

Perhaps more striking, though, is the behavior of labor hours at impact: hours actually rise in responses to the shock. Note that the impulse responses plot, in effect, linearly de-trended data. While hours lack a time trend, given

					$\epsilon = 0.50$			
		$\ln(Y)$	$\ln(C)$	$\ln(X)$	$\ln(N)$	$\ln(W)$	$ln(1+R^e)$	$\ln(1+R^f)$
	$100\sigma(\cdot)$	0.9641	0.6954	1.8047	0.1318	0.8965	0.7926	0.1014
EU, $\theta = 1/\epsilon, \gamma = 1$	$\sigma(\cdot)/\sigma(Y)$	1.0	0.7213	1.8719	0.1367	0.9299	0.8221	0.1052
	$corr(\cdot, Y)$	1.0	0.9703	0.9678	0.5634	0.9926	0.9778	−0.3648
	$100\sigma(\cdot)$	0.9642	0.6952	1.8045	0.1318	0.8964	0.7926	0.1014
High CRRA, $\theta = 100, \gamma = 1$	$\sigma(\cdot)/\sigma(Y)$	1.0	0.7210	1.8715	0.1367	0.9297	0.8220	0.1052
	$corr(\cdot, Y)$	1.0	0.9703	0.9678	0.5647	0.9926	0.9778	−0.3648
	$100\sigma(\cdot)$	0.9652	0.6940	1.8089	0.1323	0.8957	0.7949	0.1011
FORA, $\theta = 1, \gamma = 0.9$	$\sigma(\cdot)/\sigma(Y)$	1.0	0.7191	1.8741	0.1370	0.9280	0.8236	0.1047
	$corr(\cdot, Y)$	1.0	0.9702	0.9681	0.5753	0.9927	0.9779	−0.3645
					$\epsilon = 15$			
		$\ln(Y)$	$\ln(C)$	$\ln(X)$	$\ln(N)$	$\ln(W)$	$ln(1+R^e)$	$\ln(1+R^f)$
	$100\sigma(\cdot)$	0.9204	0.7386	1.4710	0.1794	0.9275	0.6805	0.0433
EU, $\theta = 1/\epsilon, \gamma = 1$	$\sigma(\cdot)/\sigma(Y)$	1.0	0.8025	1.5983	0.1949	1.0077	0.7394	0.0470
	$corr(\cdot, Y)$	1.0	0.9874	0.9768	0.0576	0.9812	0.9778	−0.2568
	$100\sigma(\cdot)$	0.9204	0.7386	1.4711	0.1794	0.9275	0.6805	0.0433
High CRRA, $\theta = 100, \gamma = 1$	$\sigma(\cdot)/\sigma(Y)$	1.0	0.8025	1.5984	0.1949	1.0077	0.7394	0.0470
	$corr(\cdot, Y)$	1.0	0.9874	0.9768	0.0578	0.9812	0.9778	−0.2586
	$100\sigma(\cdot)$	0.9204	0.7385	1.4713	0.1794	0.9274	0.6806	0.0433
FORA, $\theta = 1, \gamma = 0.9$	$\sigma(\cdot)/\sigma(Y)$	1.0	0.8024	1.5986	0.1949	1.0076	0.7395	0.0471
	$corr(\cdot, Y)$	1.0	0.9874	0.9768	0.0580	0.9812	0.9778	−0.2586

Table 6: Select second moments, model with habit formation and capital adjustment costs. Simulated data have been logged and HP-filtered, as described in the text. Entries for $corr(\cdot, Y)$ are variables' contemporaneous correlations with output.

the high persistence of the shock process it is no doubt the case that HP-filtered hours (the data underlying Table 6) behave somewhat differently, as of course does HP-filtered output (which does have a trend). The countercyclical impulse response of hours in Figure 2 is thus not necessarily inconsistent with the weak positive correlation shown in the table.

As discussed above, the countercyclical response of hours here presumably reflects the agent's desire to maintain consumption above the habit level, in spite of the poor production possibilities in the impact period. After impact, hours fall to a level slightly below their long-run mean before quickly transitioning back and, in fact overshooting very slightly.

With 100% depreciation assumed for the habit stock—and a habit-strength parameter ϕ much less than one—the habit level can quickly adjust downward in response to a persistent negative TFP shock. One has to assume that with less than complete depreciation, the period of sharply increased hours, currently limited to the period of impact, would lengthen, though this is an experiment which I still need to run.

The odd behavior of hours in the presence of habit formation has been noted by Graham [15], though he does not obtain hours that are actually countercyclical. The intratemporal preferences Graham uses—like those used by Smets and Wouters [32]—impose a form of separability between consumption and hours worked that is inconsistent with balanced growth. The complementarity between consumption and leisure in the preferences I employ, necessary for consistency with balanced growth, may thus contribute to the countercyclicality observed in Figure 2.

Finally, what of the welfare costs of volatility in the full model? Table 7 contains the results. Compared to the no-habit results, the costs of volatility in the model with habit formation are slightly higher under EU (though still minuscule); lower by about 0.1% of consumption, in general, for the high CRRA case; and essentially unchanged for the FORA specification. The λ_2 values are higher in that case, but this simply reflects the fact that there is much less precautionary accumulation of capital in the full model.[11] The cost of volatility is still on the order of \$120 billion annually, if current US annual consumption spending is our reference.

[11] Recall that both λ_1 and λ_2 envision a switch from a deterministic economy to the stochastic economy. Lifetime utility in the stochastic economy is the same for both comparisons—it is certainty equivalent lifetime utility in the stochastic economy, holding the the stochastic economy's long-run average stocks. The two measures differ, though, in their deterministic benchmarks. For λ_1, the agent is holding the stochastic economy's long-run stocks, while for λ_2 the agent is holding the deterministic economy's long-run (steady state) stocks. The deterministic benchmark in λ_1 is more attractive than that in λ_2 to the extent that the stochastic steady state stocks are higher—hence the agent requires more compensation to switch out of λ_1's deterministic environment. That effect is very pronounced in the stripped-down model, but is close to negligible in the model with habits.

	$\epsilon = 0.50$	
	λ_1	λ_2
EU, $\theta = 1/\epsilon, \gamma = 1$	0.04	0.06
High CRRA, $\theta = 100, \gamma = 1$	0.28	0.27
FORA, $\theta = 1, \gamma = 0.9$	1.31	1.33

	$\epsilon = 15$	
	λ_1	λ_2
EU, $\theta = 1/\epsilon, \gamma = 1$	0.03	0.05
High CRRA, $\theta = 100, \gamma = 1$	0.28	0.29
FORA, $\theta = 1, \gamma = 0.9$	1.31	1.33

Table 7: Welfare cost of fluctuations in the model with habit formation and capital adjustment costs.

6 Conclusions and directions for future work

Risk preferences matter a great deal for the cost of aggregate volatility. In a technology shock-driven business cycle model with production and elastic labor effort, it's doubtful that they matter for anything else, including the average level of the risk-free rate.

Intertemporal substitution matters a great deal as does the presence of habit formation. With elastic labor supply and intratemporal preferences constrained to be consistent with balanced growth, values of the EIS greater than one can lead, somewhat counterintuitively, to countercyclical behavior for consumption, and habits, together with capital adjustment costs, can produce countercyclical behavior for labor effort. Habit does not contribute noticeably to the model's average equity premium, in contrast to results from models with fixed labor supply.

At various points in the paper, I noted the need for additional work—not extensions of the current paper, but improvements of it. Those include: analyzing the robustness of the solutions to the order of the approximating polynomials; understanding exactly why the method fails to find solutions to the model when habits are sufficiently strong; and considering different calibrations of the habit and adjustment cost parameters, including less-than-complete depreciation of the habit stock.

A Appendix: First-order risk aversion

A.1 Background

The FORA risk preferences I employ in this paper are similar to those used by Epstein and Zin [12] in studying the equity premium puzzle. Epstein and Zin's specification, in turn, is based on the non-expected utility formulations of Yaari [40] and Quiggin [27]. Risk preferences of this sort are referred to in the literature variously as "rank-dependent expected utility," "expected utility with rank-dependent probabilities," or "anticipated utility" (the latter following the language of Quiggin [27]). They can be derived under various sets of axioms (see Wakker [37], and the references therein). A key feature of these preferences—like many other alternatives to expected utility—is that they are non-linear in probabilities. Among the aims of the authors who originally formulated risk preferences of this form was to elaborate models of choice under risk capable of rationalizing the apparent fact that individuals often make choices that are inconsistent with the independence axiom of expected utility—for example, the Allais paradox or the common ratio effect documented by Kahneman and Tversky.[12] The FORA risk preferences used in this paper, like other alternatives to expected utility, can be parametrized to be consistent with the choices generally made by individuals in the Allais paradox and are consistent with the common ratio effect.

The fact that risk preferences of this form are non-linear in probabilities gives them another attractive feature: the ability to at least partially divorce agents' attitudes towards risk from their attitudes towards wealth.[13] Under expected utility, aversion to risk is equivalent to diminishing marginal utility of wealth, and the intimate connection between the two concepts has been shown to be problematic for the EU model. For example, Chetty [6] has shown that estimates of labor supply elasticity (and the degree of complementarity between consumption and leisure) can put sharp bounds on admissible coefficients of relative risk aversion, since both values are linked to the curvature of agents' von Neumann-Morgenstern utilities over consumption. Chetty finds that the mean coefficient of relative risk aversion implied by 33 studies of labor supply elasticity is roughly unity, which would mean that the EU model is incapable of rationalizing both observed labor supply behavior and the degrees of risk aversion observed in many risky choice settings, many of which imply double-digit coefficients of relative risk aversion.

One of the most attractive features of these preferences, though, from the standpoint of empirical plausibility, is the fact that they can be parametrized to give a reasonable amount of risk aversion for both large and small gambles. This is in contrast to the standard expected utility specification. In the CRRA class, for example, if the coefficient of risk aversion is calibrated so that an agent with those preferences gives plausible answers to questions about large gambles, the agent will be roughly risk neutral for small gambles. If, on the other hand, the coefficient of risk aversion is set sufficiently large that the agent gives plausible answers to questions about small gambles, he will appear extremely risk averse when confronted with large gambles.[14] This is because

[12]Starmer [33] is an excellent recent survey of this literature.

[13]As Yaari [40] puts it: "At the level of fundamental principles, risk aversion and diminishing marginal utility of wealth, which are synonymous under expected utility, are horses of different colors." In Yaari's theory the divorce of the two concepts is complete.

[14]Note that my claim here is much more modest than that in Rabin [28]. Indeed, as Safra and Segal show in a recent paper [31], almost all common alternatives to expected utility are susceptible

the standard expected utility specification with constant relative risk aversion is "smooth at certainty"—the agent's indifference curves between consumption in different states of nature are smooth and tangent (at the certainty point) to the indifference curves of a risk neutral agent.

The risk preferences I use introduce a kink into agents' indifference curves at the certainty point; the kink is what allows for a plausible calibration of risk aversion for small gambles.[15] The parameter γ—which makes outcome rankings matter—is the source of the kink. The parameter θ, analogous to the risk aversion coefficient in CRRA preferences, governs curvature away from the certainty point and allows for a plausible calibration of risk aversion for large gambles.

A.2 Some calibration exercises

A few calculations can illustrate the sense in which the FORA risk preferences just described can be parametrized to give an empirically plausible degree of risk aversion over a broader range of gamble sizes than is possible with the EU specification.

In each case, I evaluate binary lotteries of the form $\{\tilde{w}; p\}$, where wealth $\tilde{w}$ takes on the value w_L with probability p or $w_H > w_L$ with probability $1 - p$, using

$$\mu(\tilde{w}) = [p^{\gamma} w_L^{1-\theta} + (1 - p^{\gamma}) w_H^{1-\theta}]^{1/(1-\theta)} \tag{9}$$

and compare the certainty equivalents to the mean of wealth as a measure of willingness to pay to avoid the gamble. I consider the FORA case of $\gamma = 0.9$ and $\theta = 1$, and compare it to the EU/CRRA case that sets $\gamma = 1$ and arbitrary $\theta \geq 0$.

Consider first a very small risk. Suppose an agent with initial wealth of \$30,000 faces a 0.00477 probability of losing \$55. This is a small risk—the standard deviation of the lottery $\{\tilde{w}; p\} = \{(29945, 30000); (0.00477, 0.99523)\}$, as a percent of mean wealth, is about 0.013%. If the agent has FORA preferences with $\gamma = .9$ and $\theta = 1$, he would be willing to pay just under 45 cents to insure against this risk. Is that a lot? Apparently not: while the initial wealth level of \$30,000 is purely hypothetical, the 0.00477 probability and \$55 loss are averages from Cicchetti and Dubin's [7] data on phone wire insurance: repair charges averaged \$55 per claim and the average probability of a claim was 0.00477 per month. The average price of phone wire insurance was 45 cents per month (nearly two times the expected loss), and 57% of the customers in

to a Rabin-like critique. The one exception noted by Safra and Segal is Yaari's dual theory of choice under risk, which is a special case of the preferences I employ here (holding when $\theta = 0$). The calculations I make below follow (albeit to a somewhat different end) the spirit of Palacios-Huerta, Serrano and Volij [26]—"[I]t is more useful not to argue whether expected utility is literally true (we know that it is not, since many violations of its underpinning axioms have been exhibited). Rather, one should insist on the identification of a useful range of empirical applications where expected utility is a useful model to approximate, explain, and predict behavior."

[15] See figure 1 in [12]. The "disappointment aversion" preferences used by Campanale *et al.* [5] share this feature.

their sample purchased phone wire insurance. Coaxing a willingness to pay 45¢for this insurance out of an EU/CRRA agent with the same initial wealth would require a coefficient of relative risk aversion of 550.[16]

Now consider a more modest-sized risk. Suppose the agent with wealth equal to $30,000 faces a 0.245 probability of losing $182. The standard deviation of this gamble, as percent of mean wealth, is 0.26%. A FORA agent with $\gamma = .9$ and $\theta = 1$ (as in the last example), would be willing to pay about $51 to insure against this risk. The loss and loss probability again come from an empirical study: Cohen and Einav's [8] analysis of the choice of auto insurance deductibles in a large sample of Israeli drivers. The $51 the FORA agent would pay is in the right ballpark—in Cohen and Einav's data, the average deductible-premium menu offers savings on deductible of $182 (in the event of claims, which have an average frequency of 0.245) at a price of $55. About 18% of the individuals in the sample chose higher premiums in exchange for a lower deductible. Coaxing a willingness to pay $51 for this insurance out of an EU/CRRA agent with the same initial wealth of $30,000 would require a coefficient of relative risk aversion of about 50.[17]

Finally, consider a large risk. Suppose the agent, again with initial wealth of $30,000, faces a 7% probability of suffering a $5,000 loss. This represents a gamble with a standard deviation equal to 4.3% of mean wealth. A FORA agent with $\gamma = .9$ and $\theta = 1$ (as in the previous two examples) would be willing to pay $495 to insure against this risk. The 7% probability and $5,000 loss are roughly the US average homeowners' multi-peril insurance claim rate and claim intensity for the period 2000–2004, according to the Insurance Information Institute.[18] $495 is low compared to the US average premium, in 2004, of over $600, but it's in the general vicinity. Coaxing a willingness to pay $495 for insurance against this risk from an EU/CRRA agent with the same initial wealth is easier here than in the smaller-risk examples—a coefficient of relative risk aversion of about 4 will work. Of course, the market for homeowners' insurance is complex—the industry is regulated, homeowners with mortgages have little choice as to whether to insure or not, and the average figures mask considerable heterogeneity. The point of this example, though, together with the two previous examples, is simply to show that a FORA agent with risk preferences that depart modestly from EU/CRRA (which is $\gamma = 1$) will be in the ballpark of empirical plausibility in all three cases. To achieve the same for the EU/CRRA specification meant re-calibrating the coefficient of relative risk

[16] Note that my interpretation of Cicchetti and Dubin's data differs from their own, as they conclude that the data are consistent with EU with only a modest coefficient of risk aversion. My interpretation is more akin to that of Rabin and Thaler [29].

[17] Cohen and Einav estimate a structural model taking account of adverse selection and allowing for heterogeneity in individual risk and risk aversion. Using average annual Israeli income as a proxy for wealth, they obtain an average relative risk aversion coefficient of 81 in their benchmark specification. Sydnor [34] presents a similar example using data on deductible choices in the market for homeowners insurance and finds implied relative risk aversion coefficients in the triple digits.

[18] `http://www.iii.org/media/facts/statsbyissue/homeowners/`. [This link was active as of 5/11/2011. *These numbers need to be brought up to date.*]

aversion from 550 for the very small risk, to 50 for the modest risk and finally to 4 for the large risk. Choose only one of those numbers and apply it to all three examples, and the EU/CRRA agent will be far out of the ballpark in two out of three cases.

There is, of course, also a large literature in experimental economics that seeks to adduce individuals' attitudes towards risk, though I confess to finding the results in that literature difficult to interpret, as it is generally assumed that the consequences to the subjects of the choices they make are simply the payoffs they receive in the experiment—initial wealth or income are "checked at the door", so to speak. The recent study by Holt and Laury [19] is a good case in point, in particular because it seems so well-done. In Holt and Laury's experiments, subjects are given a choice between pairs of binary gambles, a safer gamble (call it A) that pays \$2.00 with probability p and \$1.60 with probability $1 - p$ and a riskier gamble (B) that pays \$3.85 with probability p and \$0.10 with probability $1 - p$. When the probability of the higher outcomes, p, is near zero, gamble A has a higher expected value than B; the opposite is true when p is near one. They are interested in the effects that payoff size has on their experimental results, so in addition to the gambles just described, Holt and Laury also perform experiments where the payoffs are scaled up by a factors of 20, 50 and 90. They conduct experiments where the payoffs are real (a subject can really go home with \$346.50, say) and experiments where the payoffs are hypothetical.

Holt and Laury's experiments ask subject to choose between the A and B gambles at different values of p, the probability of the higher outcomes; in particular, they record the point at which subjects switch from the safer A gamble to the riskier B gamble as p is increased from $1/10$ to 1 in increments of $1/10$. An expected income maximizer, for example, would choose A up to $p = 4/10$, then switch to B. Among the results Holt and Laury report are the proportions of subjects' choices that are consistent with maximizing $\mathbb{E}(x^{1-\theta})/(1-\theta)$ for a CRRA coefficient θ in various ranges.[19] While they find significantly more risk aversion than had been found previously in the experimental literature—especially when the payoffs are real and high—only a very small fraction of the choices are consistent with maximizing $\mathbb{E}(x^{1-\theta})/(1-\theta)$ for θ as big as even 1.37. But this is under the assumption that the relevant x in the subjects' minds is simply the payoff from the experiment. If one allows that the consequences are final wealth levels—$w + x$ where w is a subject's initial wealth and x is income earned in the experiment—then even a modest value of w will blow up the the implied most-common values of θ considerably.[20] For example, in their '$\times$20 real' treatment—when payoffs are 20 times the values given in the previous paragraph, and are real rather than hypothetical—the most common behavior is sticking with A up to $p = 6/10$, then switching to B. If an individual

[19] This coefficient is r in Holt and Laury's notation.

[20] And if one doesn't factor in initial wealth, then the gambles they consider—in terms of their standard deviation—are extremely large, with percent standard deviations ranging from about 6 to over 200 (excluding degenerate gambles). In that case it's no surprise that the CRRA coefficients that can be ascribed to most subjects must be quite small.

is indifferent between A and B at $p = 6/10$, and is an EU/CRRA maximizer who regards the consequences of the experiment as simply the payoffs from the experiment, then he must have a value of $\theta \cong 0.41$. If the subject has even \$100 in initial wealth, though, and views the consequences as being values of final wealth, he must have $\theta \cong 2$; if initial wealth is \$1,000, this value becomes $\theta \cong 15$.[21]

References

[1] Ravi Bansal and Amir Yaron. Risks for the long run: A potential resolution of asset-pricing puzzles. *Journal of Finance*, 59(4):1481–1509, 2004.

[2] Susanto Basu and Miles Kimball. Long-run labor supply elasticity and the elasticity of intertemporal substitution for consumption. Unpublished manuscript, October 2002.

[3] Marianne Baxter and Mario Crucini. Explaining savings-investment correlations. *American Economic Review*, 83(3):416–436, 1993.

[4] Dario Caldara, Jesus Fernandez-Villaverde, Juan F. Rubio-Ramirez, and Wen Yao. Computing DSGE models with recursive preferences. Unpublished manuscript, November 2009.

[5] Claudio Campanale, Rui Castro, and Gian Luca Clementi. Asset pricing in a production economy with chew-dekel preferences. *Review of Economic Dynamics*, 13(2):379–402, 2010.

[6] Raj Chetty. A new method of estimating risk aversion. *American Economic Review*, 96:1821–1834, 2006.

[7] Charles Cicchetti and Jeffrey Dubin. A microeconometric analysis of risk aversion and the decision to self-insure. *Journal of Political Economy*, 102:169–186, 1994.

[8] Alma Cohen and Liran Einav. Estimating risk preferences from deductible choice. *American Economic Review*, pages 745–788, 2007.

[9] Thomas F. Cooley and Edward C. Prescott. Economic growth and business cycles. In Thomas F. Cooley and Edward C. Prescott, editors, *Frontiers of Business Cycle Research*, pages 1–38. Princeton University Press, Princeton, NJ, 1995.

[21] Cohen and Einav [8] estimate a CARA utility model for the Holt and Laury subjects who participated in the highest-payoff experiment, assuming a lognormal distribution for the subjects' CARA parameters, and report a point estimate of the mean of this distribution equal to 0.032. Taking average US disposable income in 2002 as a proxy for wealth, they calculate an average relative risk aversion coefficient of 865.75.

[10] Mariano Massimiliano Croce. Long-run productivity risk: A new hope for production-based asset pricing. Working paper, University of North Carolina, March 2010.

[11] Larry Epstein and Stanley Zin. Substitution, risk aversion and the temporal behavior of asset returns: A theoretical analysis. *Econometrica*, 57:937–969, 1989.

[12] Larry Epstein and Stanley Zin. f́irst-orderŕisk aversion and the equity premium puzzle. *Journal of Monetary Economics*, 26:387–407, 1990.

[13] Christopher Erceg, Luca Guerrieri, and Christopher Gust. SIGMA: A new open economy model for policy analysis. *International Journal of Central Banking*, 2(1):1–50, 2006.

[14] Marvin Goodfriend and Robert G. King. The new neoclassical synthesis and the role of monetary policy. *NBER Macroeconomics Annual 1997*, pages 231–283, 1997.

[15] Liam Graham. Consumption habits and labor supply. *Journal of Macroeconomics*, 30(1):382–395, 2008.

[16] Faruk Gul. A theory of disappointment aversion. *Econometrica*, 59(3):667–686, 1991.

[17] Fatih Guvenen. Reconciling conflicting evidence on the elasticity of intertemporal substitution: A macroeconomic perspective. *Journal of Monetary Economics*, 53(7):1451–1472, 2006.

[18] Lars Peter Hansen and Thomas Sargent. Discounted linear exponential quadratic Gaussian control. *IEEE Transactions on Automatic Control*, 40:968–971, 1995.

[19] Charles Holt and Susan Laury. Risk aversion and incentive effects. *American Economic Review*, 92:1644–1655, 2002.

[20] Urban Jermann. Asset pricing in production economies. *Journal of Monetary Economics*, 41(2):257–275, 1998.

[21] Robert G. King, Charles Plosser, and Sergio Rebelo. Production, growth and business cycles, I: The basic neoclassical model. *Journal of Monetary Economics*, 21:195–232, 1988.

[22] Karen A. Kopecky and Richard M. H. Suen. Finite state markov-chain approximations to highly persistent processes. *Review of Economic Dynamics*, 13:701–714, 2010.

[23] David Kreps and Evan Porteus. Temporal resolution of uncertainty and dynamic choice theory. *Econometrica*, 46(1), 1978.

[24] Robert E. Lucas, Jr. *Models of Business Cycles*. Basil Blackwell, New York, NY, 1987.

[25] Narayan Y. Naik and Michael J. Moore. Habit formation and intertemporal substitution in individual food consumption. *The Review of Economics and Statistics*, 78(2):321–328, 1996.

[26] Ignacio Palacios-Huerta, Roberto Serrano, and Oscar Volij. Rejecting small gambles under expected utility. *Economics Letters*, 91:250–259, 2004.

[27] John Quiggin. A theory of anticipated utility. *Journal of Economic Behavior and Organization*, 3:323–343, 1982.

[28] Matthew Rabin. Risk aversion and expected-utility theory: A calibration result. *Econometrica*, 68:1281–1292, 2000.

[29] Matthew Rabin and Richard Thaler. Anomalies: Risk aversion. *Journal of Economic Perspectives*, 15:219–232, 2001.

[30] K. Geert Rouwenhorst. Asset pricing implications of equilibrium business cycle models. In Thomas F. Cooley and Edward C. Prescott, editors, *Frontiers of Business Cycle Research*, pages 294–330. Princeton University Press, Princeton, NJ, 1995.

[31] Zvi Safra and Uzi Segal. Calibration results for non-expected utility theories. *Econometrica*, 76:1143–1166, 2008.

[32] Frank Smets and Raf Wouters. An estimated dynamic stochastic general equilibrium model of the euro area. *Journal of the European Economic Association*, 1(5):1123–1175, 2003.

[33] Chris Starmer. Developments in non-expected utility theory: The hunt for a descriptive theory of choice under risk. *Journal of Economic Literature*, 38:333–382, 2000.

[34] Justin Sydnor. (Over)insuring modest risks. *American Economic Journal: Applied Economics*, 2(4):177–199, 2010.

[35] Thomas Tallarini. Risk-sensitive real business cycles. *Journal of Monetary Economics*, 45:507–532, 2000.

[36] George Tauchen. Finite state markov chain approximations to univariate and vector autoregressions. *Economics Letters*, 20:177–181, 1986.

[37] Peter Wakker. Separating marginal utility and probabilistic risk aversion. *Theory and Decision*, 36:1–44, 1994.

[38] Yi Wen. Where's the beef? The trivial dynamics of real business cycle models. Working Paper 2005-039A, Federal Reserve Bank of St. Louis, 2005.

[39] Peter Whittle. *Risk-Sensitive Optimal Control*. Wiley Interscience Series in Systems and Optimization. John Wiley & Sons, New York, NY, 1990.

[40] Menachem Yaari. The dual theory of choice under risk. *Econometrica*, 55:95–115, 1987.