The Price of Oil Risk

Steven D. Baker,* Bryan R. Routledge,†

September, 2011
[March 16, 2015]

Abstract We solve a Pareto risk-sharing problem with heterogeneous agents with recursive utility over multiple goods. We use this optimal consumption allocation to derive a pricing kernel and the price of oil and related futures contracts. This gives us insight into the dynamics of risk premia in commodity markets for oil. As an example, in a calibrated version of our model we show how rising oil prices and falling oil risk premium are an outcome of the dynamic properties of the optimal risk sharing solution. We also compute portfolios that implement the optimal consumption policies and demonstrate that large and variable open interest is a property of optimal risk sharing.

*McIntire School of Commerce, University of Virginia; sdb7e@virginia.edu.
†Tepper School of Business, Carnegie Mellon University; routledge@cmu.edu.

1 Introduction

The spot price of crude oil, and commodities in general, experienced a dramatic price increase in the summer of 2008. For oil, the spot price peaked in early July 2008 at $145.31 per barrel (see Figure 1). In real-terms, this price spike exceeded both of the OPEC price shocks of 1970's and has lasted much longer than the price spike at the time of the Iraq invasion of Kuwait in the summer of 1990. This run-up in the price of oil begins around 2004. Büyükşahin et. al. (2010) and **?** (?) document a structural break in the behavior of oil prices around 2004. This 2004 to 2008 time period also coincides with a large increase trading activity in commodities by hedge funds and other financial firms as well as a growing popularity of commodity index funds (best documented in Büyükşahin et. al. (2010)). In fact, there is much in the popular press that lays the blame for higher commodity prices, food in particular, on the "financialization" of commodities.[1] Others point out that since these new traders in futures do not end up consuming any of the spot commodity, the trading can have little (if any) effect on spot prices.[2] Resolving this debate requires modeling the equilibrium relationship between spot and futures prices. How do spot and futures prices respond to a change in, say, speculative demand from a hedge fund as opposed to the hedging demand of a firm in the oil market? To address this question, however, we need a clearer understanding of hedging and speculation. To do this, we look directly at the risk-sharing Pareto problem in an economy with heterogenous agents and multiple goods and solve for equilibrium risk premia.

Our intuition about the use and pricing of commodity futures contracts is often expressed with hedgers and speculators. This dates back to **?** (?) and his discussion of "normal backwardation" in commodity markets. The term backwardation is used in two closely related contexts. Often it is used to refer to a negatively sloped futures curve (where the one-year futures price is below the current spot price).[3] Here, Keynes use of the term "normal backwardation" (or "natural") refers to the situation where the current one-year

[1]See for example "The food bubble:How Wall Street starved millions and got away with it" by Frederick Kaufman, Harpers July 2010 `ttp://`arpers.org/archive/2010/07/0083022

[2]The clearest argument along this lines is by James Hamilton `ttp://www.econbrowser.com/arc`ives/2011/08/fundamentals_sp.html. See also **?** (?) and **?** (?)

[3]Often, backwardation refers to the contemporaneous slope of the futures curve. In oil markets – we focus entirely on crude oil in this paper – the forward price is typically below the spot price. In our data, of 1990 through 2010, the 12 month forward is smaller than the 1 month forward (as a proxy for the spot price), a negatively sloped forward curve called backwardation, 61% of the time. This fact is an important input to many derivative-pricing models in commodities. Typically, the slope of the forward curve is a (exogenous) stochastic factor capturing the "connivence yield" to owning the physical good over a financial contract (see **?** (?)). Alternatively, the dynamics of storage or production can be modeled directly to capture the contemporaneous relationship between spot and futures price (**?** (?), (?), Titman (2011), and others).

futures price is below the expected spot price in one-year. This you will recognize as a risk premium for bearing the commodity price risk. This relation is "normal" if there are more hedgers than there are speculators. Speculators earn the risk premium and hedgers benefit from off-loading the commodity price risk. First, there is no reason to assume that hedgers are only on one side of the market. Both oil producers (Exxon) and oil consumers (Southwest Air) might hedge oil. It happens to turn out that in oil markets in the 2004 to 2008 period there was a large increase on the long-side by speculators suggesting the net "commercial" or hedging demand was on the short side. This is documented in Büyükşahin et. al. (2010) who use proprietary data from the CFTC that identifies individual traders. For many reasons it is interesting to see who is trading what. Second, if we are interested in risk premia in equilibrium we need to look past the corporate form of who is trading. We own a portfolio that includes Exxon, Southwest Air, and a commodity hedge fund and consume goods that, to varying degrees, depend on oil. Are we hedgers or speculators?

It is hard to look at risk premium directly. However, it is easy to look at realized excess returns to get a sense of things. Figure 2 plots the one-month holding period expected excess returns. **?** (**?**) estimate a structural break around 2004, so we split the data into pre-2004 and post-2004 samples. In the top panel we see changes in the term structure of excess returns, which are downward sloping pre-2004 but upward sloping to flat post-2004. Realized excess returns on the nearest to maturity contracts, which are the most heavily traded, decline after 2004. This change is linked to a more frequently upward sloping futures curve after 2004. In the bottom panel we also condition on a positively or negatively sloping futures curve. The figure suggests substantial time-variation in the risk premium across all contract maturities, coinciding with changes to the slope of the futures curve. In addition the post-2004 sample has higher average excess returns for all contracts after conditioning on the slope of th futures curve. Of course, time variation in risk premia is not surprising in modern asset pricing. We see it in equity returns (**?** (**?**), **?** (**?**), **?** (**?**)) and bond returns **?** (**?**)).

Equilibrium risk premia properties depend on preferences, endowments, technologies, and financial markets. In this model we focus on complete and frictionless financial markets. We also leave aside production for the moment. Both of these are important aspects to consider in future research. In this paper we look at an endowment economy with two goods, one of which we calibrate to capture the salient properties of oil the other we think of as composite good akin to consumption in the macro data. We consider two agents with heterogenous preferences over the two goods as well as with different time and risk aggregators. Preference heterogeneity is a natural explanation for portfolio heterogeneity we see in commodity markets. Here we start with complete and frictionless markets, focus on "perfect" risk sharing, and solve for the Pareto optimal consumption allocations. From this solution, we can infer the "representative agent" marginal rates of substitution and calculate asset prices and the implied risk premia.

It is important in our model, to allow for a rich preference structure and so we start with the recursive preference structure of **?** (**?**) and **?** (**?**). Preference heterogeneity can be over the time aggregator, the risk aggregator,and the goods aggregator that modulates the trade off between oil and the numeraire consumption good. This is important for several reasons. First, as we know from recent research into the equity premium, recursive preferences are a necessary component to generating the observed dynamics properties of the equity premium. For example, in **?** (**?**) the long-run risk component and the stochastic volatility of the consumption growth process are not sufficient to generate a realistic equity premium. The recursive preference structure is need to generate a non-zero price impact of these components. Second, and more directly related to our interests here, we are interested in understanding the role of commodity futures prices to manage risk and their related risk premium. To get at this issue carefully, since this is a multi-good economy, we need to be careful with our intuition about "risk aversion" (e.g., **?** (**?**)). An oil futures contract might hedge direct future oil "consumption," future consumption in general, or future continuation-utility. Since portfolio choice is fundamentally a decision about intertemporal multi-good consumption lotteries, all of these characteristics are important.

The bulk of the paper explores our a calibration of model that we solve numerically. The example demonstrates how dynamic risk sharing between agents with different preferences generates wide variation in prices, return volatility, risk premia, and open interest over time. Each agent holds a pareto-optimal portfolio, but realized returns may increase the wealth of one agent versus the other. Although shocks to oil consumption may cause transitory changes to the oil futures curve, gradual shifts in the wealth distribution produce long-run changes in the typical behavior of futures markets. Because they also occur endogenously, changes in the wealth distribution also provide an alternative (or complementary) explanation for persistent changes in oil futures markets that does not rely upon exogenously imposed structural breaks, such as permanent alterations to the consumption growth process, or the changing access to financial markets.

Depending on the endogenous wealth distribution in our model, the annualized risk premium to the long side of the nearest futures contract varies from 2% to over 8%, the spot price of oil averages the equivalent of under $50 per barrel or over $90 per barrel, quarterly futures return volatility goes from 10% to 13%, and open interest in oil futures contracts may be negligable or many times the value of total oil consumption. The impact of the wealth distribution is not limited to the oil market, rather changes in futures risk premia coincide with changes in the equity premium and interest rates.

Therefore changes in the wealth distribution can have a large impact upon both the goods (spot) market and asset prices, without any changes in aggregate consumption dynamics. In fact a sufficiently volatile wealth distribution, with sufficiently diverse preferences, could account for all of the recently observed changes in oil market dynamics. A natural question is whether such changes in the wealth distribution could occur rapidly enough, in a model with sensible parameters, to match the fluctuations observed in oil markets over the last decade. We demonstrate that endogenous evolution of the wealth distribution leads to rising spot prices, futures return volatility, and open interest, and declining risk premium to the long side of the nearest futures contract. However in our calibration the wealth distribution shifts too slowly to account for the variation observed in the data. Open interest in oil futures is very sensitive to the wealth distribution, but wealth-driven trends in asset prices are slow moving. A more plausible explanation for rapid changes in the spot price or the risk premium is a change in the relative aggregate consumption growth of oil, which can swiftly double the spot price or futures risk premium in our model.

There are many related papers to mention. We mentioned some of the oil and commodity papers above. We also build on many papers that look at risk sharing and models with heterogenous agents. We are most closely building on on **?** (**?**) and (**?**). Foundational work in risk sharing with recursive preferences includes **?** (**?**), **?** (**?**), and more recently **?** (**?**). There are also several recent papers that are related, such as Croce and Colacito (2010) and **?** (**?**). In our model, we focus on a risk sharing problem in an endowment setting. This sets aside the many interesting properties of oil production and prices that come from modeling the extraction problem. Interesting examples includes **?** (**?**), **?** (**?**), and **?** (**?**). These papers all document and model important properties of commodity prices; particularly the volatility structure of futures prices.

2 Facts

We are interested in the empirical properties of the time variation in the expected excess returns to holding a long position in oil futures. Since a futures contract is a zero-wealth position, we define the return as the fully collateralized return as follows. Define $F_{t,n}$ as the futures price at date t for delivery at date $t+n$, with the usual boundary condition that the $n=0$ contract is the spot price of oil; $F_{t,0}=P_t$. The fully collateralized return involves purchasing $F_{t,n}$ of a one-month bond and entering into the $t+h$ futures contract with agreed price $F_{t,n}$ at date t. Cash-flows at date $t+1$ come from the risk-free rate and the change in futures prices $F_{t+1,n-1}-F_{t,n}$. So,

$$r_{t+1}^n = \log\left(\frac{F_{t+1,n-1}-F_{t,n}+(F_{t,n}(\exp r_{t+1}^f))}{F_{t,n}}\right) \tag{1}$$

We are interested in risk premiums, so will look at the return in excess of the risk-free rate $\log r_{t+1}^f$ Note the dating convention: that is the return earned from date t to date $t+1$. For the risk-free rate, this is a constant known at date t. Defining things this way means that excess returns are approximately equal to the log-change in futures prices.

$$r_{t+1}^n - r_{t+1}^f \approx \log F_{t+1,n-1} - \log F_{t,n} \tag{2}$$

The futures prices we use are for the one-month to the sixty-month contracts for light-sweet crude oil traded at NYMEX.[4] To generate monthly data, we use the price on the last trading day of each month. The liquidity and trading volume is higher in near-term contracts. However, oil has a reasonably liquid market even at the longer horizons, such as out to the 60 month contract.

All of the production-based or storage-based models of oil point to the slope of the futures curve as an important (endogenous) state variable (e.g., **?** (?), **?** (?), and **?** (?) **?** (?)). Table 1 highlights that, indeed, this state variable is an important component to the dynamic properties of the the risk premium associated with a long position in oil. Across all the various sub-samples and contracts, when the slope of the futures curve at date t is negative, the expected excess returns to a long position in oil is higher. We can see two changes across the sub-periods, slitting the sample at 2004. First, the frequency of a negatively slopped forward curve is much less in the post-2004 period. Backwardation occurs 68% of the time pre-2004 and only 41% of the months 2004 and beyond (the full sample has a 60% frequency of backwardation). Despite this change, excess returns are higher post-2004 on all but the nearest to maturity contract. Returns on all maturities are higher post-2004 when we condition on the sign of the slope of the oil futures curve.[5]

The fact that oil seems to command a risk premium suggests its price is correlated with economic activity (or, by definition, the pricing kernel). Of course, oil is an important

[4]Data was aggregated by Barcharts Inc. All the contract details are at `http://www.cmegroup.com/trading/energy/crude-oil/light-sweet-crude\_contract\_specifications.html`

[5]**?** (?) use the closely retated fact that the negative slope is highly correlated to a high spot price of oil.

commodity directly to economic activity. **?** (**?**) documents that nine out of ten of the U.S. recessions since World War II were preceded by a spike up in oil prices with the "oil-shocks" of the 1970's as the most dramatic examples. Even the most recent recession follows the dramatic spike in oil prices. The NBER dates the recession as December 2007 to June 2009. The peak oil price was the summer of 2008 – right before the collapse of Lehman Brothers. However, by December of 2007, the WTI spot price was $91.73 per barrel. Table 2 looks at the one-month excess returns on holding US Treasury bonds over the same time-frame and conditioning information as Table 1. Notice in the upper-left hand panel the familiar pattern that excess returns on bonds are increasing in maturity. As you would expect from **?** (**?**), the evidence suggests that bond risk premia are time varying (again, subject to the caveat we are measuring these with ex-post realized returns). Over the time-subsamples we use here, there is little variation in the excess returns. What is an interesting characteristic (and perhaps even new), is that the risk-premia depend on the slope of the oil futures curve. When the oil curve is negatively sloped, excess returns on bonds are larger. The effect is strongest for longer-horizon bonds. Unconditionally, however, the correlations of the excess returns across the bonds and oil futures are small (slightly negative).

The classic empirical method to explore risk premiums is by way of a forecasting regression of **?** (**?**) (and many related papers in non-commodities). To remind us of the basic idea, write define $\phi_{t,n}$ as $\phi_{t,n} = E_t[P_{t+n}] - F_{t,n}$. Predictable movement in prices reflect risk premium and just a bit of algebra formally relates the $\phi_{t,n}$ to the covariance with the pricing kernel.[6] We can run the following regression

$$\begin{aligned} P_{t+n} - P_t &= a + b\,(F_{t,n} - P_t) + \epsilon_{t+n} \\ &= a_n + b_n\,(E_t[P_{t+n} - P_t] - \phi_{t,n}) + \epsilon_{t+n} \end{aligned}$$

And note that if $\phi_{t,n}$ is a constant, then b_n should be one. Table 3 shows that the coefficient b_n is broadly decreasing in maturity, and is significantly less than one for long-dated contracts. This suggests more variation in the risk premium associated with the longer-dated oil futures contracts.

There is certainly more work to do here, but we think we have established an interesting fact worth pursing: The risk premium in oil is time varying with interesting connections to economic activity. So, perhaps, if we want to understand the oil markets in the 2004 to 2008 period and the increased "financialization" of commodity markets, getting a handle on the source of the time variation is a good place to start.

[6] It is easiest to describe this with futures price and the future spot price (using $F_{t+n,0} = P_{t+n}$). However, implementing this empirically we use a near-term contract with $\phi_{t,n,k} = E_t[F_{t+n-k,k}] - F_{t,n}$ with $n < k$. In Table 3, $k = 1$ and we also transform things by log.

3 Exchange Economy - Two Goods and Two Agents

We model an exchange or endowment economy as in (?). We specify a stochastic process for the endowment growth. Specifically, we will have one good x_t we think of as the "numeraire" or composite commodity good. Our second good, which we calibrate to be oil, we denote y_t. This is the usual tree structure where here we have two trees, each with it's own type of fruit. We use the short-hand notation subscript-t to indicate conditional on the history to date t. Similarly, we use E_t and μ_t to indicate expectations and certainty equivalents conditional on information to date-t. Heterogeneity in our set-up will be entirely driven by preference parameters. Beliefs across all agents are common.

We are interested in the Pareto optimal allocation or "perfect" risk sharing solution. So with complete and frictionless markets, we focus on the social planner's Pareto problem. This means, for now, we need not specify the initial ownership of the endowment; we treat x_t and y_t as resource constraints. However, we can use this solution to characterize portfolio policies that implement the optimal consumption policies allowing us to investigate open interest in oil futures contracts. The preferences, which we allow to differ across our two agents, are recursive as in **?** (?) and **?** (?). They are characterized by three "aggregators" (see **?** (?)) First, a goods aggregator determines the tradeoff between our two goods. This is, of course, a simplification since oil is not directly consumed. But the heterogeneity across our two agents will capture that some of us are more reliant on or more flexible with respect to the consumption of energy-intensive products. The other two aggregators are the usual time aggregator and risk aggregator that determine intertemporal substitution and risk aversion. Finally, the familiar time-additive expected utility preferences are a special case of this set up.

3.1 Single Agent, Two Goods

To get started, consider a single-agent economy with two goods. In this representative agent setting, optimal consumption is simply to consume the endowment x_t and y_t each period. We model utility from consumption of the "aggregated good" with a constant elasticity of substitution (CES) aggregator: $A(x_t, y_t) = \left((1-\gamma)x_t^\eta + \gamma y_t^\eta\right)^{1/\eta}$ with $\gamma \in [0,1]$ and $\eta \leq 1$. Intertemporal preferences over the aggregate good are represented with an Epstein-Zin recursive preference structure

$$\begin{aligned} W_t = W(x_t, y_t, W_{t+1}) &= \left[(1-\beta)A(x_t, y_t)^\rho + \beta\mu_t(W_{t+1})^\rho\right]^{1/\rho}, \\ \mu_t(W_{t+1}) &= E_t\left[W_{t+1}^\alpha\right]^{1/\alpha}, \end{aligned}$$

The endowment of (x_t, y_t) follows a finite-state Markov process such that both the growth rates of the goods and the ratio of their levels (y_t/x_t) is stationary. Combined with the CES aggregator, this ensures that the price of good y_t in units of numeraire x_t will be stationary. Denote the current state of the Markov process s_t, with the probability of transitioning to next period state s_{t+1} given by $\pi(s_t, s_{t+1})$, for $s_t, s_{t+1} \in S$ with S finite. Growth in the numeraire good is $g_{t+1} = g(s_{t+1}) = x_{t+1}/x_t$, and the goods ratio is $h_{t+1} = h(s_{t+1}) = y_{t+1}/x_{t+1} > 0$.

With a little algebra, we can write intertemporal marginal rate of substitution, written here as a pricing kernel or stochastic discount factor.

$$m_{t+1} = \frac{\partial W_t/\partial x_{t+1}}{\partial W_t/\partial x_t}(\pi_{t+1})^{-1} = \beta \left(\frac{x_{t+1}}{x_t}\right)^{\eta-1} \left(\frac{A_{t+1}}{A_t}\right)^{\rho-\eta} \left(\frac{W_{t+1}}{\mu_t(W_{t+1})}\right)^{\alpha-\rho}. \tag{3}$$

Note this is denominated in terms of the numeraire good (x). So we can use m_{t+1} to compute the price at t of arbitrary numeraire-denominated contingent claims that pay-off at $t+1$. Claims to oil good y at t are converted to contemporaneous numeraire values using the "spot price" of oil,

$$P_t = \frac{\partial W_t/\partial y_t}{\partial W_t/\partial x_t} = \frac{\gamma}{1-\gamma} h_t^{\eta-1}. \tag{4}$$

The pricing kernel and spot price can be used in combination to price arbitrary contingent claims to either good.

We scale utility by current consumption of the numeraire, x_t, similar to **?** (**?**). This normalized utility process is stationary, due to the assumed stationarity of y_t/x_t and homogeneity of the time and state aggregators. Define, $\hat{W}_t = W_t/x_t$ and $\hat{A}_t = A_t/x_t$:

$$\hat{W}_t = \hat{W}(s_t) = \left[(1-\beta)\hat{A}_t^\rho + \beta\left(\mu_t(\hat{W}_{t+1} g_{t+1})\right)^\rho\right]^{1/\rho} \tag{5}$$

$$\hat{A}_t = \hat{A}(s_t) = \left((1-\gamma) + \gamma\,(h_t)^\eta\right)^{1/\eta}. \tag{6}$$

Substituting $\hat{W}_t$ and $\hat{A}_t$ into the pricing kernel, we have

$$\begin{aligned} m_{t+1} &= \beta\,(g_{t+1})^{\eta-1} \left(g_{t+1}\frac{\hat{A}_{t+1}}{\hat{A}_t}\right)^{\rho-\eta} \left(\frac{g_{t+1}\hat{W}_{t+1}}{\mu_t(g_{t+1}\hat{W}_{t+1})}\right)^{\alpha-\rho} \\ &= \beta\,(g_{t+1})^{\alpha-1} \left(\frac{\hat{A}_{t+1}}{\hat{A}_t}\right)^{\rho-\eta} \left(\frac{\hat{W}_{t+1}}{\mu_t(g_{t+1}\hat{W}_{t+1})}\right)^{\alpha-\rho} \end{aligned} \tag{7}$$

The pricing kernel depends on the current state s_t, and on next period state s_{t+1}.

The price of oil depends on the relative levels of the two goods (h_t), and on share parameter γ and elasticity of substitution parameter η. Changes in the price of oil only depend upon η and changes in the goods ratio:

$$\frac{P_{t+1}}{P_t} = \left(\frac{h_{t+1}}{h_t}\right)^{\eta-1}.$$

Therefore a single parameter η influences the volatility of the price of oil, but any time-variation in its dynamics flows directly from the goods ratio h_t, which follows an exogenous Markov process s_t.

The advantages and limitations of a single agent single good representative agent model are quite well known. For example, with a thoughtfully chosen consumption growth process one can capture many salient feature of equity and bond markets (**?** (**?**)). Alternatively, one can look at more sophisticated aggregators or risk to match return moments (**?** (**?**)). Presumably, one could take a similar approach to extend to a two-good case to look at oil prices and risk premia (see **?** (**?**) as a nice example). It would require some work in our specific set up, since dynamics of the oil spot price (and all derivatives) will simply depend on the current growth state s_t. Instead, we extend this model to a second (but similar) agent. The dynamics of the risk sharing problem we discuss next will provide us a new and endogenous state variable, to generate realistic time variation in the oil risk premium. This also lets us look at the portfolios and trades the two agents choose to make. Lastly, note that the single agent case in this section corresponds to the boundary cases in the two-agent economy where one agent receives zero Pareto weight or has no wealth.

3.2 Two Agents, Two Goods

Next we consider our model with two agents. The two-good endowment process and recursive preference structure are unchanged. What is new is we allow the two agents to have differing parameters for their goods, risk, and time aggregators. Denote the two agents “1” and “2” (these subscripts will denote the preference heterogenity and the endogenous goods allocations). The risk sharing or Pareto problem for the two agents is to allocate consumption of the two goods across the two agents, such that $c^x_{1,t} + c^x_{2,t} = x_t$ and $c^y_{1,t} + c^y_{2,t} = y_t$. Agent one derives utility from consumption of the aggregated good $A_1(c^x_{1,t}, c^y_{1,t}) = \left((1-\gamma_1)(c^x_{1,t})^{\eta_1} + \gamma_1(c^y_{1,t})^{\eta_1}\right)^{1/\eta_1}$. The utility from the stochastic stream of this aggregated good has the same recursive form as previously:

$$\begin{aligned} W_t = W(c^x_{1,t}, c^y_{1,t}, W_{t+1}) &= \left[(1-\beta)A_1(c^x_{1,t}, c^y_{1,t})^{\rho_1} + \beta\mu_{1,t}(W_{t+1})^{\rho_1}\right]^{1/\rho_1}, \\ \mu_{1,t}(W_{t+1}) &= E_t\left[W^{\alpha_1}_{t+1}\right]^{1/\alpha_1}. \end{aligned} \tag{8}$$

Agent 2 has similar preference structure with $A_2(c^x_{2,t}, c^y_{2,t}) = \Big((1-\gamma_2)(c^x_{2,t})^{\eta_2} + \gamma_2(c^y_{2,t})^{\eta_2}\Big)^{1/\eta_2}$ and recursive preferences

$$\begin{aligned} V_t = V(c^x_{2,t}, c^y_{2,t}, V_{t+1}) &= \Big[(1-\beta)A_2(c^x_{2,t}, c^y_{2,t})^{\rho_2} + \beta\mu_{2,t}(V_{t+1})^{\rho_2}\Big]^{1/\rho_2}, \\ \mu_{2,t}(V_{t+1}) &= E_t\left[V_{t+1}^{\alpha_2}\right]^{1/\alpha_2}. \end{aligned} \tag{9}$$

The idea is that the two agents can differ about the relative importance of the oil good, willingness to substitute non-oil for oil, risk aversion over the "utility lotteries", or the inter-temporal smoothing. Recall that with recursive preferences all of these parameters will determine the evaluation of a lifetime consumption bundle. "Oil risk" does not just depend on the γ and η parameters since it involves an inter-temporal, risky consumption lottery. Note we give the two agents common rate of time preference β.[7]

The two-agent Pareto problem is a sequence of consumption allocations for each agent $\{c^x_{1,t}, c^y_{1,t}, c^x_{2,t}, c^y_{2,t}\}$ that maximizes the weighted average of date-0 utilities subject to the aggregate resource constraint which binds at each date and state:

$$\begin{aligned} \max_{\{c^x_{1,t}, c^y_{1,t}, c^x_{2,t}, c^y_{2,t}\}} \lambda W_0 &+ (1-\lambda)V_0 \\ \text{s.t. } c^x_{1,t} + c^x_{2,t} &= x_t \quad \text{and} \\ c^y_{1,t} + c^y_{2,t} &= y_t \quad \text{for all } s^t \end{aligned}$$

where λ determines the relative importance (or date-0 wealth) of the two agents. Note that even though each agent has recursive utility, the objective function of the social planner is not recursive (except in the case of time-additive expected utility). We can rewrite this as a recursive optimization problem following, **?** (**?**), and **?** (**?**):

$$\begin{aligned} J(x_t, y_t, V_t) = \max_{c^x_{1,t}, c^y_{1,t}, V_{t+1}} &\Big[(1-\beta)A_1(c^x_{1,t}, c^y_{1,t})^{\rho_1} + \beta\mu_{1,t}(J(x_{t+1}, y_{t+1}, V_{t+1}))^{\rho_1}\Big]^{1/\rho_1} \\ \text{s.t. } &V(x_t - c^x_{1,t}, y_t - c^y_{1,t}, V_{t+1}) \geq V_t. \end{aligned} \tag{10}$$

The optimal policy involves choosing agent one's date-t consumption, $c^x_{1,t}$, $c^y_{1,t}$ and the resource constraint pins down agent two's date-t bundle. In addition, at date-t, we solve for a vector of date-$t+1$ "promised utility" for agent two. Making good on these promises at date $t+1$ means that V_{t+1} is an endogenous state variable we need to track.[8] That is,

[7] Differing β's are easy to accommodate, but seem an unnecessary added complexity. Heterogeneous time discount factors often lead to less interesting dynamics, since the agent with the larger β quickly dominates the optimal allocation (see e.g., **?** (**?**)).

[8] In practise it is convenient to track an alternative state variable that maps monotonically to normalized $\hat{V}_t$.

optimal consumption at date t depends on the aggregate supply of goods x_t and y_t and the previously promised utility V_t. Finally, note that the solution to this problem is "perfect" or optimal risk sharing. Since we consider complete and frictionless markets, there is no need to specify the individual endowment process.

Preferences are monotonic, so the utility-promise constraint will bind. Therefore with optimized values, we have $W_t = W(c^x_{1,t}, c^y_{1,t}, W_{t+1}) = J(x_t, y_t, V_t)$ and $V_t = V(x_t - c^x_{1,t}, y_t - c^y_{1,t}, V_{t+1})$. The first order and envelope conditions for the maximization problem with date-t-dependent Lagrange multiplier λ_t[9] are

$$
\begin{aligned}
W_x(c^x_{1,t}, c^y_{1,t}, W_{t+1}) &= \lambda_t V_x(x_t - c^x_{1,t}, y_t - c^y_{1,t}, V_{t+1}) && (11)\\
W_y(c^x_{1,t}, c^y_{1,t}, W_{t+1}) &= \lambda_t V_y(x_t - c^x_{1,t}, y_t - c^y_{1,t}, V_{t+1}) && (12)\\
W_V(c^x_{1,t}, c^y_{1,t}, W_{t+1}) &= -\lambda_t V_V(x_t - c^x_{1,t}, y_t - c^y_{1,t}, V_{t+1}) && (13)\\
J_V(x_t, y_t, V_t) &= -\lambda_t, && (14)
\end{aligned}
$$

Rearranging these optimality conditions implies, not surprisingly, that the marginal utilities of agent 1 and agent 2 are aligned across goods and inter-temporally. These equations imply that

$$
\begin{aligned}
m_{t+1} &= \beta \left(\frac{c^x_{1,t+1}}{c^x_{1,t}}\right)^{\eta_1 - 1} \left(\frac{A_{1,t+1}}{A_{1,t}}\right)^{\rho_1 - \eta_1} \left(\frac{W_{t+1}}{\mu_{1,t}(W_{t+1})}\right)^{\alpha_1 - \rho_1} \\
&= \beta \left(\frac{c^x_{2,t+1}}{c^x_{2,t}}\right)^{\eta_2 - 1} \left(\frac{A_{2,t+1}}{A_{2,t}}\right)^{\rho_2 - \eta_2} \left(\frac{V_{t+1}}{\mu_{2,t}(V_{t+1})}\right)^{\alpha_2 - \rho_2}.
\end{aligned} \tag{15}
$$

Recall that beliefs are common across the two agents so probabilities drop out. Note that we can use this marginal-utility process as a pricing kernel. Optimality implies agents agree on the price of any asset. Similarly, the first-order conditions imply agreement about the intra-temporal trade of the numeraire good for the oil good. Hence the spot price of oil:

$$
P_t = \frac{\gamma_1}{1-\gamma_1}\left(\frac{c^y_{1,t}}{c^x_{1,t}}\right)^{\eta_1 - 1} = \frac{\gamma_2}{1-\gamma_2}\left(\frac{c^y_{2,t}}{c^x_{2,t}}\right)^{\eta_2 - 1}. \tag{16}
$$

As in the single agent model, it is helpful to use the homogeneity to scale things to be stationary. Here is the analogous scaling in the two-agent setting. Define

$$
\begin{aligned}
\hat{c}^x_{1,t} &= \frac{c^x_{1,t}}{x_t} \quad , \quad \hat{c}^y_{2,t} = \frac{c^x_{2,t}}{x_t} = 1 - \hat{c}^x_{1,t} \\
\hat{c}^y_{1,t} &= \frac{c^y_{1,t}}{y_t} \quad , \quad \hat{c}^y_{2,t} = \frac{c^y_{2,t}}{y_t} = 1 - \hat{c}^y_{1,t}
\end{aligned}
$$

[9] Many papers directly characterize the "stochastic Pareto weight" process. E.g, **?** (?), **?** (?), **?** (?).

The $\hat{c}$'s are consumption shares of the two goods. Scale utility values by the aggregate supply of good x:

$$\hat{W}_t = \frac{W_t}{x_t} \quad , \quad \hat{V}_t = \frac{V_t}{x_t}.$$

Notice we scale the utilities by the total supply of good x, and not just the agent's share. This has the advantage of being robust if one agent happens to (optimally) get a declining share of consumption over time. Plugging these into the equation (15), and we can state the pricing kernel as

$$m_{t+1} = \beta g_{t+1}^{\alpha_1 - 1} \left(\frac{\hat{c}^x_{1,t+1}}{\hat{c}^x_{1,t}} \right)^{\eta_1 - 1} \left(\frac{\hat{A}_{1,t+1}}{\hat{A}_{1,t}} \right)^{\rho_1 - \eta_1} \left(\frac{\hat{W}_{t+1}}{\mu_t(g_{t+1}\hat{W}_{t+1})} \right)^{\alpha_1 - \rho_1} . \tag{17}$$

(or equivalently from the perspective of agent 2). In the one-agent case, the pricing kernel depends on the current growth state s_t and the future growth state s_{t+1}. In the two agent case, the pricing kernel depends on both the growth state and the (scaled) level of utility of agent 2, $(s_t, \hat{V}_t)$, and their date-$t+1$ values $(s_{t+1}, \hat{V}_{t+1})$. The new state variable $\hat{V}_t$, which has endogenous dynamics determined by the optimality conditions, is analogous to the wealth share of agent 2. Because $\hat{V}_t$ affects the consumption sharing rule $(\hat{c}^x_{1,t}, \hat{c}^y_{1,t})$, the current spot price is also a function of s_t and $\hat{V}_t$:

$$P_t = \frac{\gamma_1}{1 - \gamma_1} \left(\frac{\hat{c}^y_{1,t}}{\hat{c}^x_{1,t}} h_t \right)^{\eta_1 - 1} . \tag{18}$$

3.3 Financial Prices

As is standard in an exchange economy, we can now use the pricing kernel to price assets and calculate their returns. To start, we can look at the value of the agents' consumption streams. Note it is easier here to use the kernel defined in equation (15) for discounting. Define the cum-dividend price of a claim to agent 1's optimal consumption streams (i.e., agent 1's wealth inclusive current period consumption):

$$C_{1,t} = c^x_{1,t} + P_t c^y_{1,t} + \beta E_t \left[C_{1,t+1} \left(\frac{c^x_{1,t+1}}{c^x_{1,t}} \right)^{\eta_1 - 1} \left(\frac{A_{1,t+1}}{A_{1,t}} \right)^{\rho_1 - \eta_1} \left(\frac{W_{t+1}}{\mu_t(W_{t+1})} \right)^{\alpha_1 - \rho_1} \right] . \tag{19}$$

To derive an expression for wealth in terms of date t variables, a slightly different normalization of aggregated consumption will be useful:

$$\hat{H}_{1,t} = \left(\frac{A_1(c^x_{1,t}, c^y_{1,t})}{c^x_{1,t}} \right)^{\eta_1} = 1 - \gamma_1 + \gamma_1 \left(\frac{c^y_{1,t}}{c^x_{1,t}} \right)^{\eta_1} . \tag{20}$$

Guess that agent 1 wealth is

$$C_{1,t} = \frac{c^x_{1,t}\hat{H}_{1,t}W_t^{\rho_1}}{(1-\beta)(1-\gamma_1)A_{1,t}^{\rho_1}}, \tag{21}$$

and verify the guess:

$$\begin{aligned}
C_{1,t} &= c^x_{1,t} + P_t c^y_{1,t} + \beta E_t\left[\frac{c^x_{1,t+1}\hat{H}_{1,t+1}W_{t+1}^{\rho_1}}{(1-\beta)(1-\gamma_1)A_{1,t+1}^{\rho_1}}\left(\frac{c^x_{1,t+1}}{c^x_{1,t}}\right)^{\eta_1-1}\left(\frac{A_{1,t+1}}{A_{1,t}}\right)^{\rho_1-\eta_1}\left(\frac{W_{t+1}}{\mu_t(W_{t+1})}\right)^{\alpha_1-\rho_1}\right]\\
&= \frac{(1-\gamma_1)c^x_{1,t}}{1-\gamma_1} + \frac{\gamma_1 c^x_{1,t}}{1-\gamma_1}\left(\frac{c^y_{1,t}}{c^x_{1,t}}\right)^{\eta_1} + \frac{\beta c^x_{1,t}\hat{H}_{1,t}\mu_t(W_{t+1})^{\rho_1-\alpha_1}E_t\left[W_{t+1}^{\alpha_1}\right]}{(1-\beta)(1-\gamma_1)A_{1,t}^{\rho_1}}\\
&= \frac{c^x_{1,t}\hat{H}_{1,t}}{1-\gamma_1} + \frac{\beta c^x_{1,t}\hat{H}_{1,t}\mu_t(W_{t+1})^{\rho_1}}{(1-\beta)(1-\gamma_1)A_{1,t}^{\rho_1}}\\
&= \frac{(1-\beta)c^x_{1,t}\hat{H}_{1,t}A_{1,t}^{\rho_1} + \beta c^x_{1,t}\hat{H}_{1,t}\mu_t(W_{t+1})^{\rho_1}}{(1-\beta)(1-\gamma_1)A_{1,t}^{\rho_1}}\\
&= \frac{c^x_{1,t}\hat{H}_{1,t}W_t^{\rho_1}}{(1-\beta)(1-\gamma_1)A_{1,t}^{\rho_1}}.
\end{aligned} \tag{22}$$

The analogous series of derivations yield agent 2's wealth:

$$C_{2,t} = \frac{c^x_{2,t}\hat{H}_{2,t}V_t^{\rho_2}}{(1-\beta)(1-\gamma_2)A_{2,t}^{\rho_2}}. \tag{23}$$

The price of the stock, or aggregate wealth, is simply

$$C_t = C_{1,t} + C_{2,t}. \tag{24}$$

Given these processes for wealth, we can calculate the return to a claim on these assets. The return to aggregate wealth is

$$r^C_{t+1} = \log C_{t+1} - \log(C_t - x_t - P_t)$$

In comparing to the the data, we map the above return on aggregate wealth to the return on (unlevered) equity.

Bond (and the risk-free rate) all follow from the pricing kernel in the usual way. Define the price of a zero-coupon bond recursively as

$$B_{t,n} = E_t[m_{t+1}B_{t+1,n-1}], \tag{25}$$

where $B_{t,n}$ is the price of a bond at t paying a unit of the numeraire good at period $t+n$ with the usual boundary condition that $B_{t,0} = 1$. Rates follow as $r^n_{t+1} = -n^{-1}\log(B_{t,n})$ (with $n = 1$ as the risk-free rate used to compute excess returns).

Define the futures price of the oil good, y, is defined as follows. $F_{t,n}$ is the price agreed to in period t for delivery n period hence. Futures prices satisfy

$$
\begin{aligned}
0 &= E_t[m_{t+1}(F_{t+1,n-1} - F_{t,n})] \\
F_{t,n} &= (B_{t,1})^{-1} E_t[m_{t+1}F_{t+1,n-1}], \\
\log F_{t,n} &= -\log B_{t,1} + \log E_t\left[\exp\left\{\log m_{t+1} + \log F_{t+1,n-1}\right\}\right],
\end{aligned} \tag{26}
$$

with the boundary condition $F_{t,0} = P_t$. Recall, from our discussion of futures returns earlier we focus on the fully collateralized one-period holding period returns. In particular, $r^n_{t+1} - r^f_{t+1} \approx \log F_{t+1,n-1} - \log F_{t,n}$.

3.4 Portfolios and open interest

One interesting feature of a multi-agent model is we can look directly at the role of financial markets in implementing the optimal allocations. We are solving for Pareto allocation of the two resources, so implementing this in a decentralized economy generally requires complete markets. In particular, we are interested in how oil futures can be used to implement the optimal consumptions. We defer that specific question to our numerical calibration of the model since we lack analytical expressions for futures prices.

However, under one simplifying assumption we can look analytically at how "equity" claims can implement the optimal allocations. Let $\eta_1 = \eta_2 \to 0$, such that each agent has Cobb-Douglas consumption aggregator $A_i(c^x_{i,t}, c^y_{i,t}) = (c^x_{i,t})^{1-\gamma_i}(c^y_{i,t})^{\gamma_i}$, $i \in 1, 2$. Under this assumption convenient expressions are available for separate claims to numeraire (good x) and oil (good y) consumption. Agent 1's claim to numeraire consumption good x has value

$$
C^x_{1,t} = \frac{c^x_{1,t} W_t^{\rho_1}}{(1-\beta)A^{\rho_1}_{1,t}}, \tag{27}
$$

and similarly agent 2's numeraire consumption has value

$$
C^x_{2,t} = \frac{c^x_{2,t} V_t^{\rho_2}}{(1-\beta)A^{\rho_2}_{2,t}}. \tag{28}
$$

Claims to each agent's oil consumption streams are priced proportionally to the prices of non-oil consumption:

$$C^y_{1,t} = \frac{\gamma_1 C^x_{1,t}}{1-\gamma_1}, \tag{29}$$

$$C^y_{2,t} = \frac{\gamma_2 C^x_{2,t}}{1-\gamma_2}. \tag{30}$$

$$\tag{31}$$

Just as aggregate wealth can serve as a proxy for a broad equity index, claims to aggregate consumption streams of goods x and y can proxy for non-oil sector and oil sector equity, respectively. Normalizing outstanding shares for each sector to 1, non-oil equity shares have price

$$C^x_t = C^x_{1,t} + C^x_{2,t}, \tag{32}$$

and for the oil sector,

$$C^y_t = \frac{\gamma_1 C^x_{1,t}}{1-\gamma_1} + \frac{\gamma_2 C^x_{2,t}}{1-\gamma_2}. \tag{33}$$

Suppose agent 1 has portfolio of $\phi^x_{1,t}$ shares numeraire and $\phi^y_{1,t}$ shares in oil. It turns out that these two assets are sufficient for agent 1 (and, in similar fashion, agent 2) to implement his optimal wealth/consumption process. In fact the optimal portfolio consists of a constant number of shares in each asset,

$$\begin{aligned} \phi^x_{1,t} &= \frac{-\gamma_2}{\gamma_1 - \gamma_2} \\ \phi^y_{1,t} &= \frac{1-\gamma_2}{\gamma_1 - \gamma_2}. \end{aligned} \tag{34}$$

Portfolios for agent 2 are straightforward to derive via market clearing.

As we will see in the numerical section in a moment, optimal consumption for the two agents in this setting and the implied prices and asset returns have many interesting dynamic properties. However, the homogeneity of the preference structure means portfolio policies are "buy and hold." This deceptively simple allocation of shares follows because each agent holds a constant proportion of his wealth in each sector. This is a similar result to the portfolio separation results with single good and time-additive CRRA utility.

While this result is interesting, perhaps, it might not be all that practical. Most of the aggregate consumption of oil is not captured in an easily traded claim. There is a large production in state-owned enterprises (e.g., Saudi Aramco, PDVSA, and indirect state claims from oil royalties and well-head taxes). In the numerical section, next, we look at portfolio policies that implement the optimal consumption using oil futures contracts. This also gives us a perspective on open-interest dynamics.

4 Calibrated Numerical Example

The recursive Pareto problem is hard to characterize analytically, so we look at a numerical example. The first ingredient is a reasonable process for the supply of the two goods. We estimate a vector autoregression (VAR) in logs for g_t (growth in good x) and h_t (the goods ratio y_t/x_t),

$$\begin{bmatrix}\log g_t \\ \log h_t\end{bmatrix} = M + \Lambda \begin{bmatrix}\log g_{t-1} \\ \log h_{t-1}\end{bmatrix} + \Sigma \begin{bmatrix}\epsilon_{g,t} \\ \epsilon_{h,t}\end{bmatrix} \tag{35}$$

where $\epsilon_{g,t}$ and $\epsilon_{h,t}$ are standard normal. We fit the model at a quarterly frequency. From the BEA NIPA tables, we use growth in real, non-durable personal consumption expenditures less energy consumption as a measure of g_t. Since we treat "oil" as a final consumption good in our model, we also use personal energy consumption data as a proxy for consumption of oil-derived products, rather than using more direct alternatives (e.g., total crude oil supplied to the US market, from the EIA). Following **?** (**?**), the NIPA data corresponds more closely to movements in the price oil than does data on crude oil supply from the EIA, particularly when adjusted to account for changes in fuel economy. Therefore we also scale real personal energy consumption expenditures to reflect changes in fleet fuel economy, as estimated by the EIA. We estimate h_t using the ratio of adjusted real personal energy consumption to real non-durable consumption less energy.

Results of the estimation are in Table 4. We make some simplifying assumptions prior to discretizing the VAR for our model solution. Estimates of the conditional covariance terms $\Sigma_{1,2} = \Sigma_{2,1}$ are small, although statistically significant at the 10% level; we treat innovations to g_t and h_t are uncorrelated. The cross-correlation term $\Lambda_{2,1}$ for g_{t-1} as a predictor of h_t is not statistically significant, and is also set to zero. Autocorrelation terms are positive and significant for both g_t ($\Lambda_{1,1}$) and h_t ($\Lambda_{2,2}$). In fact h_t has close to unit root. For the spot price process to be stationary (indeed bounded), h_t cannot have unit root. Whether the spot price process is or, from a theoretical standpoint, ought to be stationary for an exhaustible resource such as oil is a matter of debate, which depends upon a variety of assumptions. For our purposes it is convenient to enforce stationarity.

The estimated parameter $\Lambda_{2,1}$ is positive and significant at the 5% level, implying that relatively high oil consumption (and low spot price of oil) would coincide with strong numeraire growth in the model. However the high persistence of h_t and sensitivity to corrections for fuel economy complicate estimation of the relationship between the relative level of oil consumption and the growth rate of non-oil consumption g_t. Historically positive excess returns to the long side of futures contracts (even short dated contracts) imply that spot prices are negatively correlated with the pricing kernel. This suggests that relatively high oil consumption (low spot price) coincides with low growth. Given our focus on futures returns, we choose slightly negative $\Lambda_{2,1} = -0.003$ in our model implementation, which allows for a range of futures risk premia depending upon preferences parameters. To solve the model, we discretize the VAR process using the method of **?** (**?**), with 25 states (5 unique states for each of g_t and h_t).

Our objective is a simple growth process that preserves numerical tractability, yet allows for comparisons with asset prices observed from 1990 through 2010. In particular, since our main aim is to study the time-variation in risk-premia that occurs *endogenously* through dynamic risk-sharing, we wish to avoid divergent growth trends or "structural breaks" that would amount to exogenously imposed trends or shifts in risk premia. In this context our stationary growth process seems reasonable.

We choose preference parameters to capture a few key characteristics of asset prices. Numerical values for each parameter are given in Table 5. Historically, oil consumption has represented around 4% of US GDP (**?** (**?**)), whereas personal energy consumption expenditures are around 7% of the value of personal nondurables and services consumption. In our model this characteristic is governed chiefly by the choice of goods aggregation parameters, which we set to value oil consumption in the $4 - 7\%$ range, conditional on the distribution of wealth among agents. Choices for γ_i affect the mean value of oil consumption relative to numeraire consumption, whereas values of η_i affect the mean price of oil and especially its volatility. In Table 1 we see that the volatility of futures returns (which relates to spot price volatility) is higher in the post-2004 period, when the real spot price was also higher on average. We set $\gamma_1 = \gamma_2$ but $\eta_1 > \eta_2$. Heterogeneous values of η_i allow for endogenous stochastic mean and volatility of spot and futures prices: both will tend to rise if agent 2 becomes more wealthy.The dynamic properties of the risk-sharing model will generate the variation.

For oil risk premium, our goal is to generate variation consistent with what we infer from the data. **?** (**?**), for example, suggests a lower oil-risk premium in recent years. Figure 2 also suggests a different term structure of risk premia, with mean excess returns rising or flat with time to delivery in the post 2004 sample, but falling with time to delivery in the pre 2004 sample. We set $\alpha_1 > \alpha_2$, yielding a generally higher risk premium to the long side of the futures contract when agent 1 is wealthy. However we also set $\rho_1 > \rho_2$, which results in higher risk premia on long dated contracts relative to short dated contracts when agent 2 is wealthy. We aim to match the large level of the equity premium, although with our definition of equity as aggregate wealth, expected excess returns average $0.5-1.5\%$ annually, depending upon the wealth distribution. Redefining equity as a leveraged consumption claim (similar to **?** (**?**)) would increase the equity premium. Finally, we set the common time-preference parameter, β such that the average risk-free rate is around 3%.

Because we assume infinitely lived agents and a Markov process for aggregate consumption dynamics, the state of our economy at any point in time is fully characterized by output level for numeraire x_t, the Markov state s_t, and the wealth and consumption distribution, captured by promised agent 2 utility V_t in our derivations. In our numerical analysis we substitute for V_t a more convenient state variable, agent 2's numeraire consumption share $\hat{c}^x_{2,t}$, which has domain $[0, 1]$. We drop time subscripts and describe the economy as a function of the state variables. As a standing assumption we set $x_t = 1$; as our model is homogeneous in x most results are invariant to its level. We further focus our analysis on the key state variable $\hat{c}^x_{2,t}$, by taking expectations over growth states s_t according to their stationary distribution where necessary. Section 4.1 shows how $\hat{c}^x_{2,t}$ governs wealth and consumption sharing between agents. Section 4.2 demonstrates the importance of shifts in wealth for asset prices and risk premia, whereas Section 4.3 relates these effects to trade in financial markets. Section 4.4 discusses the models dynamics, and suggests that endogenous changes in $\hat{c}^x_{2,t}$ in response to growth shocks offer a possible explanation for observed changes in price levels, volatility, risk premia, and open interest.

4.1 Wealth and consumption

The key of our model is the presence of two agents with different preferences, who interact to determine prices in competitive markets. This interaction is governed by the distribution of wealth and consumption among the two agents. Fortunately these are monotonically related, so we select a single state variable, agent 2's numeraire consumption share $\hat{c}^x_2$, to summarize the distribution of consumption and wealth. This state variable is analogous to the consumption share state variable used in heterogeneous agent models with only one good (see e.g. **?** (**?**) or **?** (**?**)). The numeraire consumption share suffices as a state variable because of the complementarity of the two goods: in equilibrium, agent 2's share oil consumption (and wealth) increases if and only if his share of numeraire consumption increases. Numeraire consumption share is more convenient than the wealth share because the divisor (aggregate numeraire consumption x_t) is a single exogenous state variable, whereas aggregate wealth is a function of x_t, s_t, and the wealth distribution itself.

Monotonicity suffices for $\hat{c}_2^x$ to serve as a state variable. However the relationship between numeraire consumption, oil consumption, and wealth shares is nonlinear. Figure 3 illustrates the relationship, plotting oil consumption share (top), wealth share (middle) and utility (bottom) for each agent versus our state variable $\hat{c}_2^x$. The plots show averages over growth state s_t computed using its stationary distribution. Conditional on relatively low availability of oil (low h_t), agent 2 receives a higher share of oil consumption, due to his lower elasticity of substitution over goods ($\eta_2 < \eta_1$). The opposite occurs in states where oil is relatively abundant. Hence oil and consumption shares are approximately linearly related unconditional of the growth state.

By contrast, the wealth share and utility of each agent is markedly nonlinear in relation to consumption shares. What is shown unconditional of the growth state in the plots generally holds for individual growth states also. When agent 2 consumes half of the numeraire, his net worth is less than 40% of aggregate wealth. Although aggregate consumption is exogenous, each agent still faces an individual consumption-saving decision. Agent 2 consumes relatively more out of his wealth than agent 1, chiefly because of his lower elasticity of intertemporal substitution ($\rho_2 < \rho_1$). So the wealth distribution and the consumption distribution are related, but distinct. Utility, up to scale, is closely tied to wealth share. By contrast the spot price of oil is a direct function of the consumption distribution. Given the importance of spot price dynamics to our analysis, $\hat{c}_2^x$ seems an appropriate state variable. Finally, note that $\hat{c}_2^x = c_2^x$ (the consumption share is equivalent to the consumption level) under standing assumption $x = 1$.

4.2 Prices and risk premia

Changes in the consumption distribution impact asset prices and the goods market. We first study these effects unconditional of the growth state s. Figure 4 summarizes the model response to the consumption distribution, showing the risk-free rate, the spot price of oil, the equity premium, and the risk premium to a long position in the nearest futures contract. Consistent with our choice of preference parameters, a larger agent 2 consumption share corresponds to a higher average spot price and a lower risk premium. An economy dominated by agent 2 would spend over 50% more on oil than an economy dominated by agent 1, and excess returns on oil futures would be only 0.5% per quarter, versus over 2% per quarter under agent 1. Although our focus is on the oil market, the channels that drive risk aversion in the futures market also manifest in the equity premium and risk free rate. Specifically, the lower risk aversion coefficient α_2 of agent 2 implies a lower equity premium coincides with a lower futures premium. Although the range of futures risk premium corresponds reasonably well to the data, the equity premium is low. Equity is defined as an unlevered claim to aggregate consumption, so the volatility of dividends is also lower than observed in the data. Adding 4.5 fold leverage along the lines of **?** (**?**) would yield a reasonable annualized equity premium of over 6% under agent 1, but an economy where agent 2 has a substatial consumption share would still have an equity premium below the historical mean even with leverage. This reflects the difficulty of disentangling variation in the futures risk premium from variation in the equity premium.

The risk-free rate also increases as consumption share varies from agent 1 to agent 2. This change relates to agent 2's lower elasticity of intertemporal substitution, which is also seen in the term structure of futures prices and returns. For oil futures, Figure 5 shows (from top to bottom) prices, the risk premium to the long position, and the volatility of returns for contracts of increasing maturity. Results are shown conditional on three values of the consumption distribution. The consumption distribution affects the price curve mainly through the mean spot price: the curve shifts upward when agent 2 has a higher consumption share. However the term structure of the futures risk premium also flattens as agent 2 gains consumption share, in addition to shifting downward. The term structure of return volatility is downward sloping, but the curve shifts upward and flattens somewhat when agent 2's consumption share increases.

We have in mind agent 1 as the dominant agent pre-2004, with a higher agent 2 consumption share post-2004. Overall the statistics in Figure 5 compare reasonably well with those in Table 1, except that the level of the risk premium under agent 2 is too low for most maturities. Replicating the steeply upward sloping term structure of risk premia for near to maturity contracts seen after 2004 is difficult in the model. We match the flatter term structure of long horizon returns, with the level of the risk premia closer to those of the heavily trade near-horizon contracts.

Averaging results over Markov state s understates the range of outcomes that is possible in the model, therefore Figure 6 shows oil futures prices, risk premia, and return volatility conditional on one of three states s. Each state corresponds to a different goods ratio $h(s)$, but the same aggregate growth rate $g(s)$.[10] We focus on the 1-year futures contract, and vary agent 2 consumption share along the x axis. The risk premium varies substantially with s,

[10]The aggregate growth rate $g(s)$ has comparatively minor impact on oil futures moments.

but far more when agent 1 has a large consumption share. In contrast the volatility of returns is little affected by s, but is greatly affected by the consumption share. Therefore variation in the risk premium over states s is not driven merely by the conditional volatility of returns; in fact return volatility and the risk premium are inversely related in this calibration. The results for volatility also highlight variation in the consumption share (driven by risk-sharing in financial markets) as an endogenous source of persistent changes in volatility that are in some sense decoupled from consumption volatility.

Another way of indirectly conditioning on the Markov state is to condition on the slope of the futures curve, which facilitates comparison with the data. Table 6 shows futures return moments conditional on the slope of the futures curve and the consumption distribution. Model results are similar to the data in Table 1 along some dimensions, with higher excess futures returns conditional on a downward slope for example, and volatility decreasing with maturity. (Note that the model table is at a quarterly, rather than monthly, frequency.)

For completeness, Figure 7 shows the term structure of zero coupon bond yields, the maturity or bond risk premium, and bond return volatility. Results are conditional on one of three consumption share values. Similar statistics are compiled conditional on the slope of the futures curve in Table 7. Comparing these results with Table 2, the model does not match the data very well, except for the increase in return volatility with maturity. **?** (**?**) analyze the term structure of interest rates in a model with Epstein-Zin preferences, and suggest that the interaction of monetary policy with the short rate is important for matching the (nominal) term structure. It would be difficult to generate an upward sloping real term structure of interest rates without implausible assumptions regarding the growth process.

4.3 Portfolios and open interest

A key advantage of modeling multiple agents is the ability to show portfolios and trade. Since the model has 25 states of nature, we would require 25 non-redundant assets to complete markets. The resultant portfolios would be difficult to visualize, so instead we project the optimal wealth process of each agent onto a limited space of four assets: equity, a 1-period bond, and the first two futures contracts. This allows us to focus on portfolios for interesting and (in the data) heavily traded assets. In fact it is possible to replicate wealth dynamics fairly well with only these four assets. Adding (for example) bonds and futures of increasing maturity to complete the market does not dramatically alter trade in the four assets chosen for study.

For each agent, Figure 8 shows the numeraire value of investments in each asset conditional on the consumption share and unconditional of the Markov state s. Results conditional on s show a similar pattern. The supply of equity nets to one share outstanding, whereas other assets are in zero net supply. Combined with the monotonic relationship between consumption share and wealth share, this implies that if either agent has 100% of the consumption share, then he holds all outstanding equity and there is no trade in other assets. However equity holdings are remarkably nonmonotonic in the interior of $\hat{c}_2^x$. Also noteworthy is the fact that agents straddle futures contract maturities: it is not the case that agent 2, who is less flexible as regards oil consumption than agent 1, insures himself by taking a long futures position generally. Nor does the more risk tolerant agent (also agent 2, if measured by parameter α_2) always take a long futures position to earn the risk premium. Some intuition comes from the result, under the simplifying assumption of Cobb-Douglas goods aggregators, that portfolios would be constant *if* separate claims to oil and non-oil consumption streams were traded. Even with Cobb-Douglas preferences, portfolios are similar to those shown in Figure 8 if markets are instead completed with futures contracts and bonds. Therefore the futures positions likely reflect a dynamic replication strategy for a claim to the (untraded) oil consumption stream.

Nevertheless much is made of changes to open interest in oil futures, so we highlight this in Figure 9, computed as the absolute value of agent 1's futures contracts. If we instead plotted the number of contracts outstanding (rather than their value), the results would remain similar in shape: what we see is not merely a reflection of changes in the value of contracts, rather it reflects changes in the number of outstanding contracts. Open interest differs dramatically depending upon the consumption share, and may be negligible, or orders of magnitude larger than the total value of oil consumed in the economy. It is also non-monotonic, peaking around $\hat{c}_2^x = 0.7$. For comparison, the right panel of Figure 9 shows the impact of $\hat{c}_2^x$ on the spot price of oil, which is monotonic, peaking when agent 2 is dominant in the economy ($\hat{c}_2^x \to 1$), reflecting his lower elasticity of substitution over goods. Open interest is not a sufficient statistic to determine the price of oil, even if we know the current growth state. Directional change in open interest does not unambiguously relate to spot prices, neither does it reflect an increase in "speculation" on the part of the agents in the economy. Nevertheless increased open interest will coincide with an increased spot price for much of the domain of $\hat{c}_2^x$. As we show in the next section, so long as $\hat{c}_2^x$ is sufficiently low initially the model would produce decades of data suggesting positive correlation between spot prices and open interest, despite the lack of a robust relationship in the long term.

4.4 Dynamics

The value of $\hat{c}_2^x$ has a strong impact upon spot prices, risk premia, and open interest. As the economy evolves, so does $\hat{c}_2^x$, reflecting changes in the consumption distribution brought about by realizations of the exogenous growth process. This section examines how $\hat{c}_2^x$ and key economic variables change over time. We show that the economy will tend toward higher spot prices, higher open interest, higher future return volatility, and a lower futures risk premium.

To illustrate the range of possible outcomes in our economy, we have thus far allowed for three widely dispersed values of $\hat{c}_2^x$. Rather than consider the possible paths of our economy from each of three starting points, we select one initial value to roughly match historical data. That data suggests lower spot prices and open interest in the past. Although risk premia are difficult to estimate with precision, at least one study **?** (**?**) suggests that futures risk premia were larger in the past. These facts recommend $\hat{c}_{2,0}^x = 0.1$, a state in which most economic wealth belongs to agent 1, as a reasonable starting value for $\hat{c}_2^x$ in our economy. Figure 10 shows how the probability density of $\hat{c}_2^x$ evolves, conditional on our chosen starting value. Over time, the possible values of $\hat{c}_2^x$ become dispersed, allowing for a good deal of variation over time in the economic variables driven by $\hat{c}_2^x$. The second feature is the drift; $\hat{c}_2^x$ tends to increase over time. However it is extremely unlikely that either agent will have a dominant position in the terminal period.

To directly relate changes in $\hat{c}_2^x$ over time to economic outcomes, we turn to the mean path of the economy. Figure 11 shows the average 50-year path of the economy, computed using Monte Carlo simulation. As before, we choose $\hat{c}_{2,0}^x = 0.1$. The initial growth state, s_0, is chosen from the stationary distribution. In the top panel, we see that $\hat{c}_2^x$ is expected to increase from 0.1 to about 0.15 over 50 years. In expectation this coincides with an increased spot price, decreased risk premium, and increased volatility. However the magnitude of the expected increases is small. An exception is open interest, which is sufficiently sensitive to the consumption share that doubling of open interest is expected.

5 Conclusions

We have focused on heterogeneous exposure to oil risk as an important driver of oil-risk premium dynamics, volatility, and trade evident in the data. To attack this question, we look at the optimal Pareto consumption sharing problem with two agents with different attitudes towards consumption risk and, specifically, the oil-component of consumption. The solution lets us look at consumption and wealth paths and the implications for risk premia. In one calibrated example, we can generate rising oil prices, decreasing risk premia, increasing futures volatility, and increasing open interest. However the calibrated example suggests that consumption or wealth share dynamics are unlikely to account for rapid variation oil futures market dynamics, except perhaps for open interest.

Table 1: **Monthly Excess returns on oil futures contracts**

		ALL			Pre 2004			Post 2004		
	Horizon	Obs	Mean	St. Dev.	Obs	Mean	St. Dev.	Obs	Mean	St. Dev.
	1	241	0.68	9.36	168	0.74	9.26	73	0.54	9.65
	3	240	1.03	8.23	167	0.93	7.85	73	1.28	9.08
	6	240	1.06	7.09	167	0.85	6.37	73	1.56	8.53
ALL	12	237	0.98	5.85	164	0.63	4.73	73	1.76	7.79
	18	189	1.25	5.36	116	0.87	3.57	73	1.86	7.34
	24	203	0.96	5.05	130	0.43	3.40	73	1.90	7.03
	36	153	0.78	5.04	88	0.01	3.06	65	1.81	6.75
	48	97	0.87	5.72	37	-0.02	3.53	60	1.41	6.69
	60	91	0.86	5.87	31	-0.24	3.83	60	1.43	6.64
	Horizon	Obs	Mean	St. Dev.	Obs	Mean	St. Dev.	Obs	Mean	St. Dev.
	1	95	-0.61	9.44	53	-0.38	9.16	42	-0.89	9.89
	3	95	-0.17	8.36	53	-0.29	7.65	42	-0.03	9.28
	6	95	-0.03	7.31	53	-0.24	6.17	42	0.23	8.60
Slope +	12	93	0.02	6.23	51	-0.35	4.67	42	0.47	7.74
	18	68	0.45	6.08	26	0.17	3.88	42	0.62	7.15
	24	77	0.25	5.41	35	-0.33	3.27	42	0.73	6.70
	36	60	0.30	5.34	20	-0.84	2.45	40	0.86	6.26
	48	43	0.41	5.68	6	-0.38	2.74	37	0.54	6.04
	60	44	0.26	5.60	7	-1.60	3.49	37	0.61	5.89
	Horizon	Obs	Mean	St. Dev.	Obs	Mean	St. Dev.	Obs	Mean	St. Dev.
	1	145	1.52	9.27	114	1.26	9.34	31	2.49	9.12
	3	144	1.84	8.09	113	1.51	7.94	31	3.05	8.65
	6	144	1.80	6.89	113	1.37	6.46	31	3.36	8.21
Slope -	12	143	1.61	5.54	112	1.09	4.73	31	3.51	7.61
	18	121	1.70	4.88	90	1.07	3.47	31	3.53	7.39
	24	126	1.39	4.78	95	0.71	3.42	31	3.48	7.26
	36	93	1.09	4.84	68	0.26	3.19	25	3.34	7.36
	48	54	1.23	5.78	31	0.05	3.70	23	2.82	7.56
	60	47	1.42	6.12	24	0.16	3.90	23	2.74	7.66

Holding period returns are monthly (shown as percent per month) on fully collateralized futures position in oil. The "Slope+" and "Slope-" correspond to the sign of $F_{t,18} - F_{t,1}$ (the 18 month futures contract price less the one month price) at the date the position is initiated (i.e., date $t+1$ return conditional on date t slope). The "pre 2004" is the period 1990-2003. The "post 2004" is 2004 to 2010.

Table 2: **Monthly Excess returns on US Treasury Bonds**

		ALL			Pre 2004			Post 2004		
	Horizon	Obs	Mean	St. Dev.	Obs	Mean	St. Dev.	Obs	Mean	St. Dev.
ALL	6	240	0.03	0.06	168	0.04	0.05	72	0.02	0.08
	12	240	0.07	0.17	168	0.09	0.16	72	0.05	0.20
	18	240	0.11	0.32	168	0.13	0.31	72	0.05	0.34
	24	240	0.13	0.47	168	0.16	0.46	72	0.08	0.48
	30	240	0.16	0.61	168	0.19	0.60	72	0.10	0.64
	36	240	0.19	0.75	168	0.21	0.74	72	0.13	0.78
	42	240	0.21	0.89	168	0.24	0.87	72	0.14	0.92
	48	240	0.22	1.02	168	0.25	1.00	72	0.15	1.07
	54	240	0.24	1.15	168	0.26	1.13	72	0.18	1.21
	60	240	0.23	1.25	168	0.26	1.23	72	0.17	1.31
	120	240	0.26	1.45	168	0.29	1.37	72	0.19	1.62
	>121	240	0.33	2.54	168	0.37	2.32	72	0.25	3.03
	Horizon	Obs	Mean	St. Dev.	Obs	Mean	St. Dev.	Obs	Mean	St. Dev.
Slope +	6	94	0.03	0.05	53	0.03	0.05	41	0.02	0.05
	12	94	0.05	0.15	53	0.06	0.16	41	0.05	0.14
	18	94	0.06	0.28	53	0.08	0.29	41	0.05	0.26
	24	94	0.08	0.41	53	0.08	0.44	41	0.08	0.38
	30	94	0.09	0.54	53	0.09	0.58	41	0.10	0.50
	36	94	0.11	0.68	53	0.09	0.72	41	0.13	0.64
	42	94	0.12	0.81	53	0.10	0.85	41	0.14	0.76
	48	94	0.11	0.96	53	0.09	0.98	41	0.14	0.94
	54	94	0.12	1.12	53	0.09	1.11	41	0.17	1.13
	60	94	0.11	1.21	53	0.07	1.21	41	0.16	1.23
	120	94	0.10	1.59	53	0.07	1.46	41	0.14	1.77
	>121	94	0.10	2.89	53	0.07	2.28	41	0.15	3.55
	Horizon	Obs	Mean	St. Dev.	Obs	Mean	St. Dev.	Obs	Mean	St. Dev.
Slope -	6	145	0.04	0.07	114	0.04	0.05	31	0.03	0.10
	12	145	0.09	0.19	114	0.10	0.16	31	0.05	0.25
	18	145	0.14	0.34	114	0.16	0.31	31	0.06	0.42
	24	145	0.17	0.50	114	0.20	0.46	31	0.08	0.60
	30	145	0.21	0.65	114	0.25	0.61	31	0.10	0.80
	36	145	0.25	0.79	114	0.28	0.74	31	0.12	0.95
	42	145	0.28	0.93	114	0.32	0.88	31	0.14	1.11
	48	145	0.30	1.05	114	0.34	1.00	31	0.16	1.23
	54	145	0.32	1.17	114	0.36	1.12	31	0.18	1.32
	60	145	0.33	1.27	114	0.36	1.22	31	0.19	1.43
	120	145	0.37	1.34	114	0.40	1.31	31	0.25	1.44
	>121	145	0.51	2.27	114	0.54	2.30	31	0.39	2.22

The data is the Fama Bond Portfolio's from CRSP. These are the one-month holding period return of an equally weighted portfolio of bonds of similar maturity. For example, horizon 18 is bonds of maturity 13-18 months, the 120 is bonds from 61-121 months and >120 is all bonds of a longer horizon that 121 months or more. All returns are excess of the one-month risk-free rate. The "Slope+" and "Slope-" is from the oil futures process. It correspond to the sign of $F_{t,18} - F_{t,1}$ (the 18 month futures contract price less the one month price) at the date the position is initiated (i.e., date $t+1$ return conditional on date t slope). The "pre 2004" is the period 1990-2003. The "post 2004" is 2004 to 2010.

Table 3: **Predictive regression for crude oil**

	a		b			
Horizon (n)	a.	t ($a = 0$)	b	t ($b = 1$)	R^2	nobs
3	0.0197	1.6519	**1.0692**	-0.2111	0.0431	238
6	0.0462	2.7311	**0.8285**	0.7193	0.0493	235
12	0.0848	3.9861	**0.8573**	0.7656	0.0853	229
18	0.1377	5.7044	**0.9315**	0.3922	0.1141	223
24	0.1929	7.4039	**0.7321**	1.5744	0.0832	206
36	0.2645	7.3578	**0.0789**	4.8220	0.0012	146
48	0.5638	9.7593	**0.4254**	2.3378	0.0491	60
60	0.7670	13.0886	**0.3373**	3.0781	0.0517	47

Crude oil futures data for 1990-2011. Regression of:

$$\log F_{t+n,1} - \log F_{t,1} = a + b\,(\log F_{t,n} - \log F_{t,1}) + \epsilon_{t+n}$$

Note the t-stat shown for a is for a different from zero and for b is the t-stat reflects b different from one.

Table 4: **Aggregate Consumption Growth Process**

$$\begin{bmatrix} \log g_t \\ \log h_t \end{bmatrix} = M + \Lambda \begin{bmatrix} \log g_{t-1} \\ \log h_{t-1} \end{bmatrix} + \Sigma \begin{bmatrix} \epsilon_{g,t} \\ \epsilon_{h,t} \end{bmatrix} \tag{36}$$

Parameter	Estimate		Implementation	
M	0.004		0.004	
	0.000		0.000	
Λ	0.430	0.011	0.430	-0.003
	0.064	0.996	0.000	0.996
$10^3 * \Sigma$	0.034	-0.015	0.048	0.000
	-0.015	0.381	0.000	0.381

Growth process characteristics. M(1) matches the sample estimate of unconditional mean non-energy consumption growth, which is 0.72 % per quarter. Without loss of generality, since the relative importance of oil as part of consumption bundles is determined by preference preference parameters, $\log h_t$ is normalized to mean 0, hence M(2) = 0. Remaining parameters for the de-meaned process are estimated by maximum likelihood. Autocorrelation terms $\Lambda(1,1)$ and $\Lambda(2,2)$ are statistically significant at the 1% level, $\Lambda(1,2)$ at the 5% level, whereas $\Lambda(2,1)$ is not statisically significant.

Table 5: **Parameters**

Parameter	Value	Description
α_1	-20	risk aversion, agent 1
α_2	-10	risk aversion, agent 2
ρ_1	0.2	intertemporal substitution, agent 1
ρ_2	-0.2	intertemporal substitution, agent 2
γ_1	0.03	oil preference, agent 1
γ_2	0.03	oil preference, agent 2
η_1	-4	oil preference, agent 1
η_2	-6	oil preference, agent 2
β	0.997	impatience, agents 1,2

Preference parameters used in numerical examples.

Table 6: **Model-implied expected excess returns on oil futures contracts (%)**

		$c_2^x = 0.10$		$c_2^x = 0.50$		$c_2^x = 0.90$	
	Horizon (Q)	Mean	St. Dev.	Mean	St. Dev.	Mean	St. Dev.
	1	2.029	10.047	1.511	11.389	0.853	13.130
	2	2.018	9.938	1.507	11.299	0.854	13.055
	4	1.997	9.723	1.500	11.121	0.856	12.907
	6	1.975	9.512	1.492	10.945	0.857	12.759
	8	1.953	9.303	1.484	10.771	0.858	12.613
All	12	1.907	8.895	1.467	10.428	0.860	12.324
	16	1.861	8.499	1.450	10.093	0.861	12.041
	20	1.813	8.116	1.431	9.765	0.861	11.762
	30	1.690	7.213	1.379	8.977	0.858	11.088
	40	1.563	6.388	1.323	8.234	0.850	10.443
	Horizon (Q)	Mean	St. Dev.	Mean	St. Dev.	Mean	St. Dev.
	1	0.643	9.993	0.563	11.176	0.418	13.025
	2	0.640	9.902	0.565	11.099	0.421	12.956
	4	0.631	9.722	0.564	10.947	0.425	12.820
	6	0.622	9.544	0.560	10.797	0.428	12.685
	8	0.613	9.368	0.556	10.648	0.430	12.552
Slope +	12	0.595	9.022	0.548	10.353	0.435	12.288
	16	0.577	8.683	0.539	10.063	0.438	12.028
	20	0.558	8.353	0.530	9.778	0.440	11.773
	30	0.513	7.562	0.505	9.089	0.443	11.152
	40	0.469	6.822	0.479	8.430	0.442	10.556
	Horizon (Q)	Mean	St. Dev.	Mean	St. Dev.	Mean	St. Dev.
	1	2.121	10.050	1.574	11.403	0.888	13.139
	2	2.110	9.941	1.570	11.312	0.889	13.063
	4	2.088	9.723	1.562	11.132	0.890	12.913
	6	2.065	9.510	1.554	10.955	0.892	12.765
	8	2.042	9.299	1.546	10.779	0.893	12.618
Slope -	12	1.995	8.886	1.529	10.433	0.894	12.327
	16	1.946	8.487	1.510	10.095	0.895	12.042
	20	1.897	8.100	1.491	9.764	0.895	11.762
	30	1.768	7.190	1.438	8.969	0.891	11.082
	40	1.636	6.359	1.379	8.221	0.883	10.433

Quarterly holding period mean and standard deviation of returns on fully collateralized oil futures in excess of the 1-period bond rate, in percent. Results are shown for different horizons to maturity and values of state variable c_2^x. Increasing values of c_2^x correspond to an increasing consumption and wealth share for the agent with lower risk aversion and lower elasticity of substitution over goods. Slope is defined as the difference between the 1 quarter and 6 quarter futures contracts.

Table 7: **Model-implied expected excess returns on bonds (%)**

	Horizon (Q)	$c_2^x = 0.10$		$c_2^x = 0.50$		$c_2^x = 0.90$	
	Horizon (Q)	Mean	St. Dev.	Mean	St. Dev.	Mean	St. Dev.
	1	-0.000	0.000	-0.000	0.000	-0.000	0.000
	2	-0.056	0.247	-0.054	0.289	-0.049	0.345
	4	-0.089	0.400	-0.087	0.470	-0.080	0.560
	6	-0.093	0.430	-0.093	0.507	-0.086	0.603
	8	-0.091	0.439	-0.095	0.518	-0.088	0.614
All	12	-0.088	0.448	-0.096	0.530	-0.090	0.625
	16	-0.086	0.455	-0.097	0.542	-0.092	0.636
	20	-0.084	0.463	-0.099	0.554	-0.094	0.648
	30	-0.084	0.478	-0.105	0.581	-0.099	0.680
	40	-0.086	0.488	-0.113	0.606	-0.106	0.712
	Horizon (Q)	Mean	St. Dev.	Mean	St. Dev.	Mean	St. Dev.
	1	-0.000	0.000	-0.000	0.000	-0.000	0.000
	2	-0.054	0.237	-0.051	0.279	-0.048	0.337
	4	-0.084	0.368	-0.081	0.438	-0.077	0.533
	6	-0.090	0.389	-0.086	0.465	-0.083	0.567
	8	-0.091	0.393	-0.089	0.472	-0.085	0.575
Slope +	12	-0.093	0.396	-0.091	0.479	-0.088	0.583
	16	-0.094	0.399	-0.094	0.487	-0.091	0.592
	20	-0.095	0.403	-0.096	0.496	-0.093	0.602
	30	-0.098	0.412	-0.102	0.523	-0.100	0.635
	40	-0.101	0.422	-0.107	0.553	-0.106	0.673
	Horizon (Q)	Mean	St. Dev.	Mean	St. Dev.	Mean	St. Dev.
	1	-0.000	0.000	-0.000	0.000	-0.000	0.000
	2	-0.057	0.248	-0.054	0.290	-0.049	0.346
	4	-0.089	0.402	-0.088	0.472	-0.081	0.562
	6	-0.093	0.433	-0.094	0.510	-0.087	0.605
	8	-0.091	0.442	-0.095	0.521	-0.088	0.617
Slope -	12	-0.088	0.451	-0.096	0.534	-0.090	0.628
	16	-0.085	0.459	-0.097	0.546	-0.092	0.640
	20	-0.084	0.467	-0.099	0.558	-0.094	0.652
	30	-0.083	0.482	-0.106	0.585	-0.099	0.684
	40	-0.085	0.493	-0.114	0.610	-0.106	0.716

Annual returns on zero-coupon bonds in excess of the 1-period bond rate, in percent. Results are shown for different horizons to maturity and values of state variable $\hat{c}_2^x$. Increasing values of $\hat{c}_2^x$ correspond to an increasing wealth share for the agent with lower risk aversion and higher preference for oil consumption. Slope is defined as the difference between the spot price and the 18-month futures price.

Figure 1: Time series, crude oil spot and futures prices

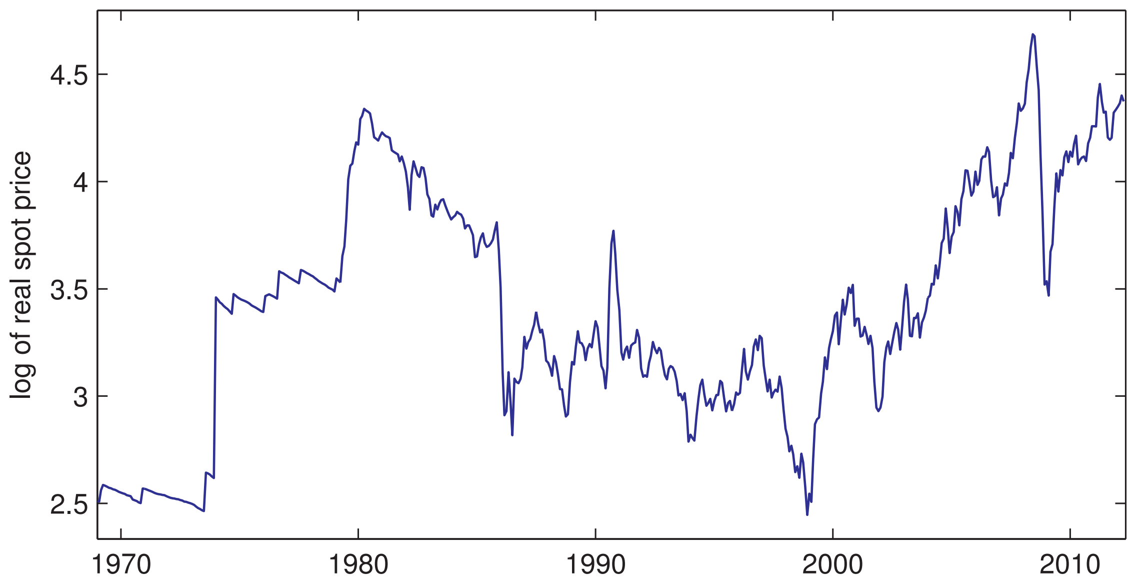

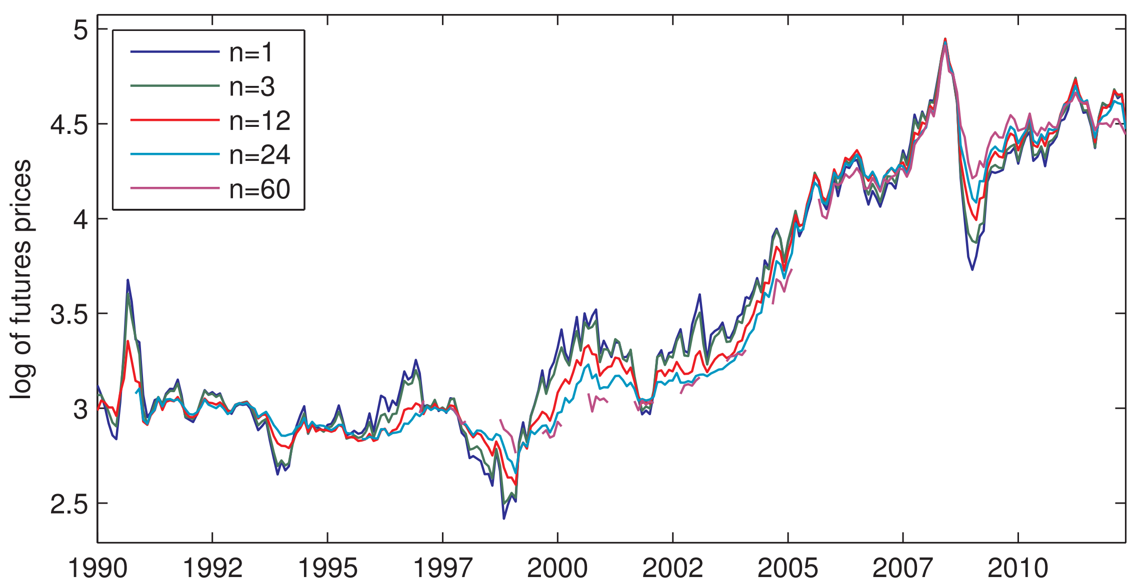

Crude oil spot price is from St. Louis Fed, deflated by chain-weighted price index (indexed at 2001/01). Futures prices are nominal NYMEX WTI crude oil futures prices.

Figure 2: Term structure of excess returns, crude oil futures

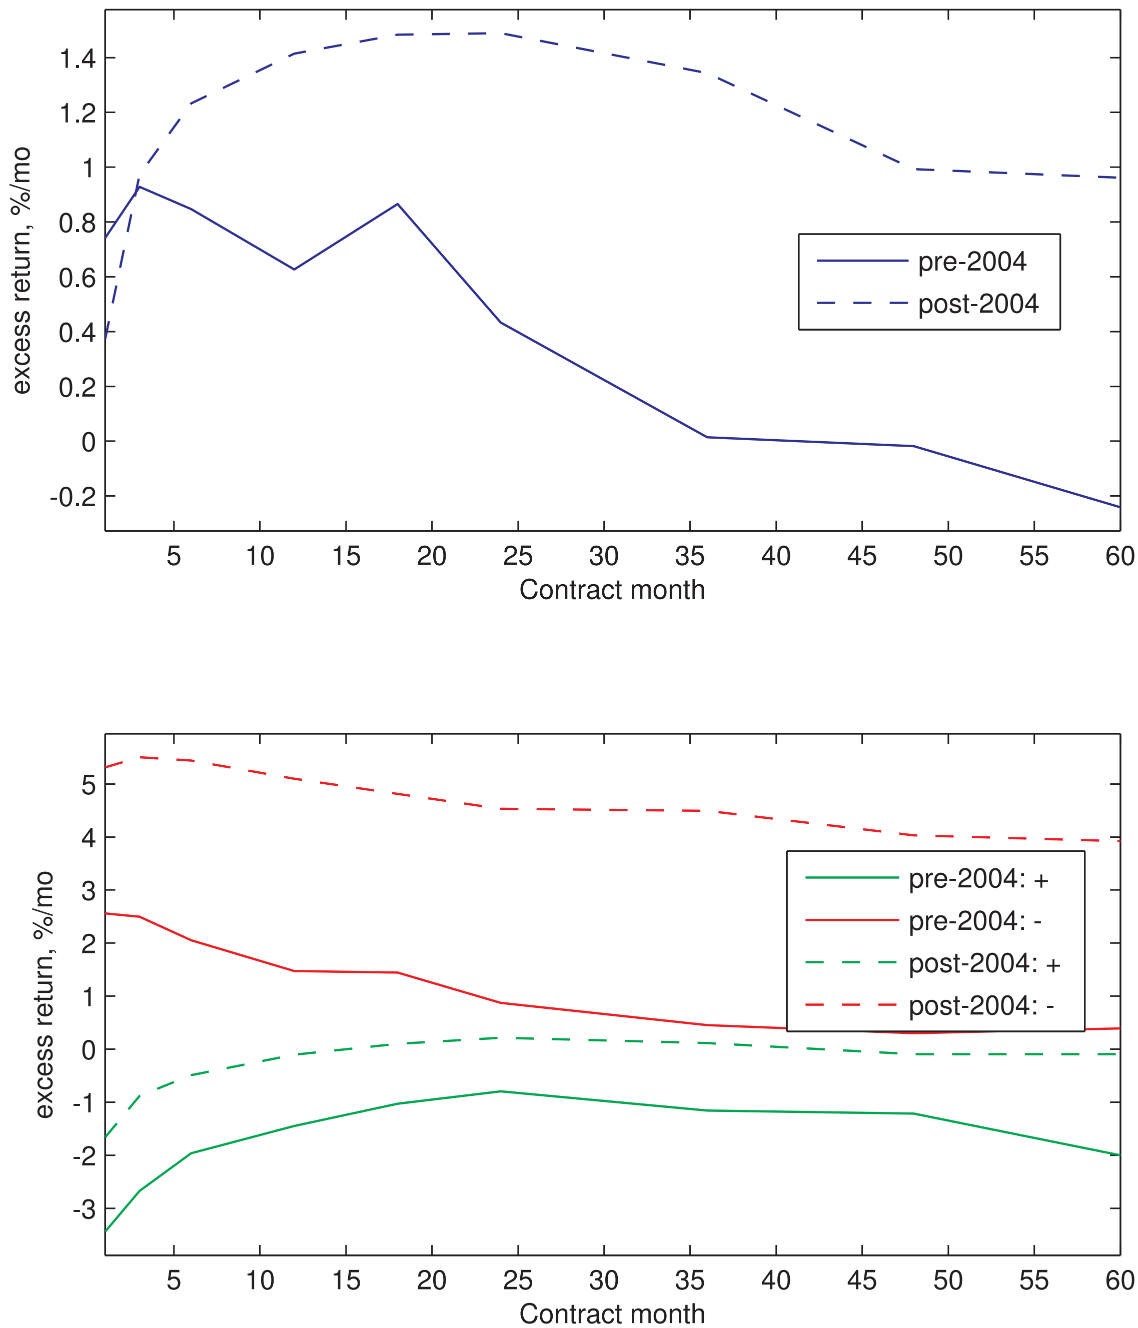

One month expected holding excess returns on a long fully-collateralized futures position. The excess return is calculated as the mean excess return realized in sample. The bottom panel shows averages conditional on the sign of the slope of the futures curve, defined as the difference between the 18 month and the 1 month contract prices.

Figure 3: Wealth and consumption shares

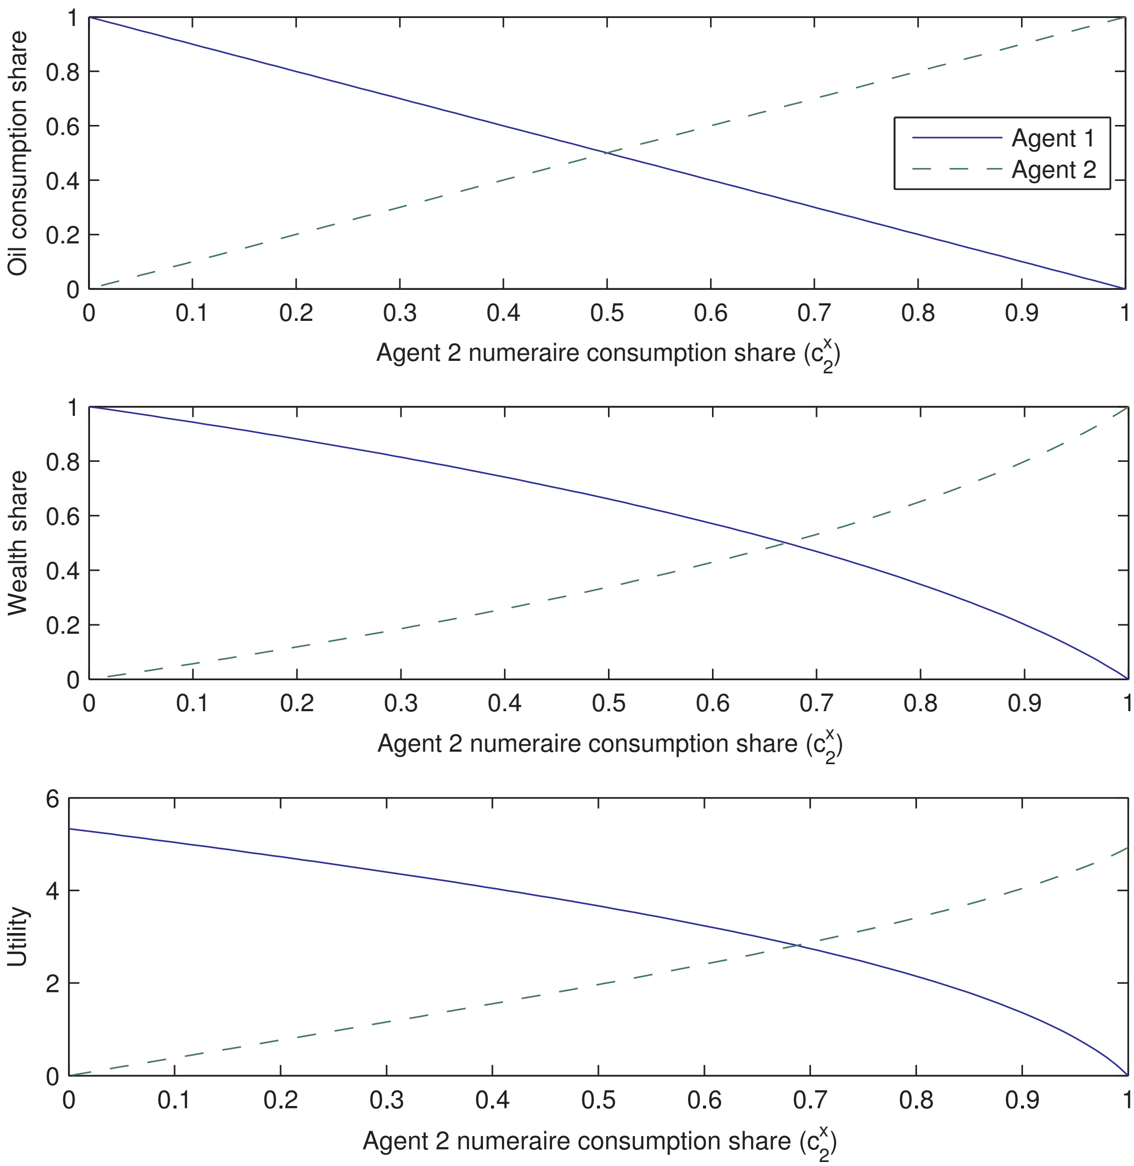

Wealth and consumption shares of each agent, relative to state variable c_2^x. Results are averages over the stationary distribution of growth states (s). However results conditional on a particular s are generally similar to the mean.

Figure 4: Prices, returns and the consumption distribution

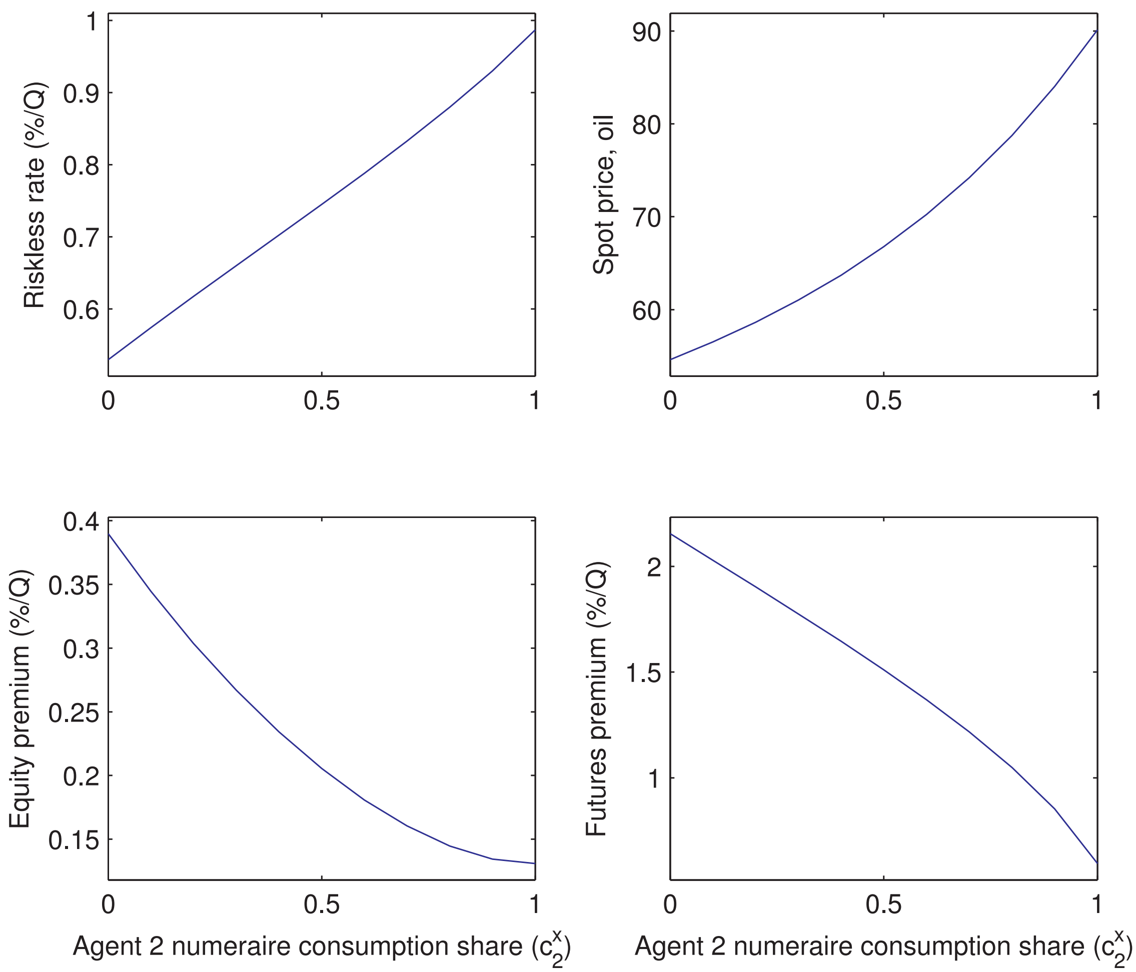

The above figure summarizes the model response to the consumption distribution, showing the risk-free rate, the spot price of oil, the equity premium, and the risk premium to a long position in the nearest futures contract. All values are expectations taken over the stationary distribution of growth states s. Rates and returns are quarterly. Spot prices are rescaled to dollar-per-barrel equivalents in the figure.

Figure 5: Futures characteristics versus maturity

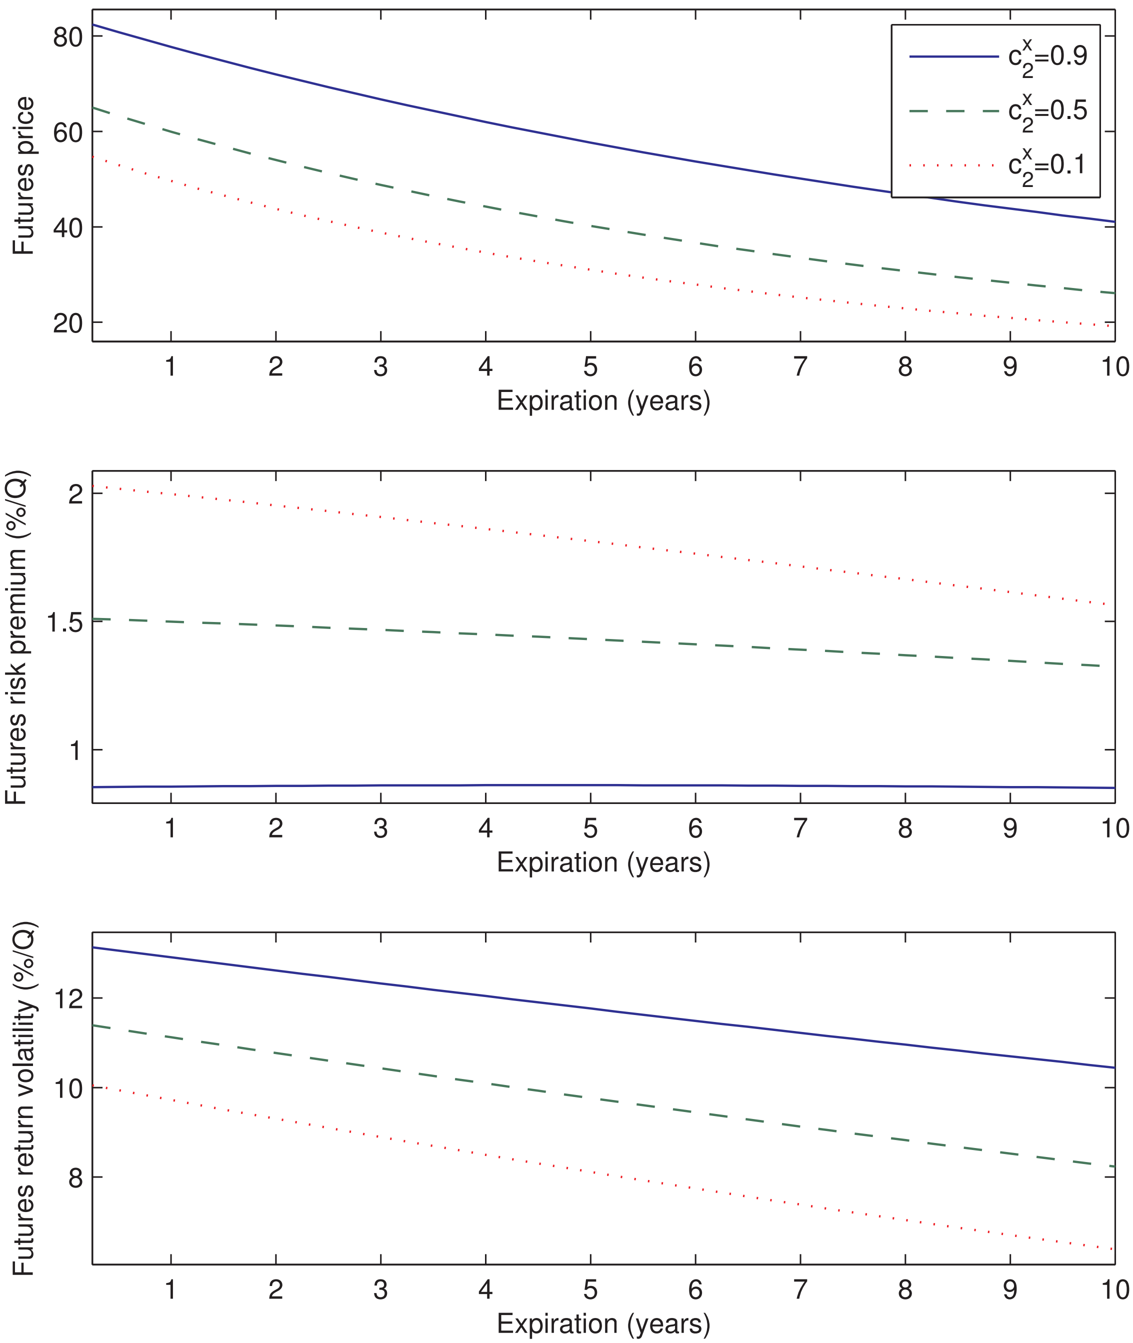

The figure shows prices, the quarterly risk premium and volatility of returns to a long oil futures position, averaged over the stationary distribution of growth states. Results are shown for increasing time until expiration of the futures contract (x-axis), and for varying consumption distribution.

Figure 6: Futures characteristics versus consumption share and growth state

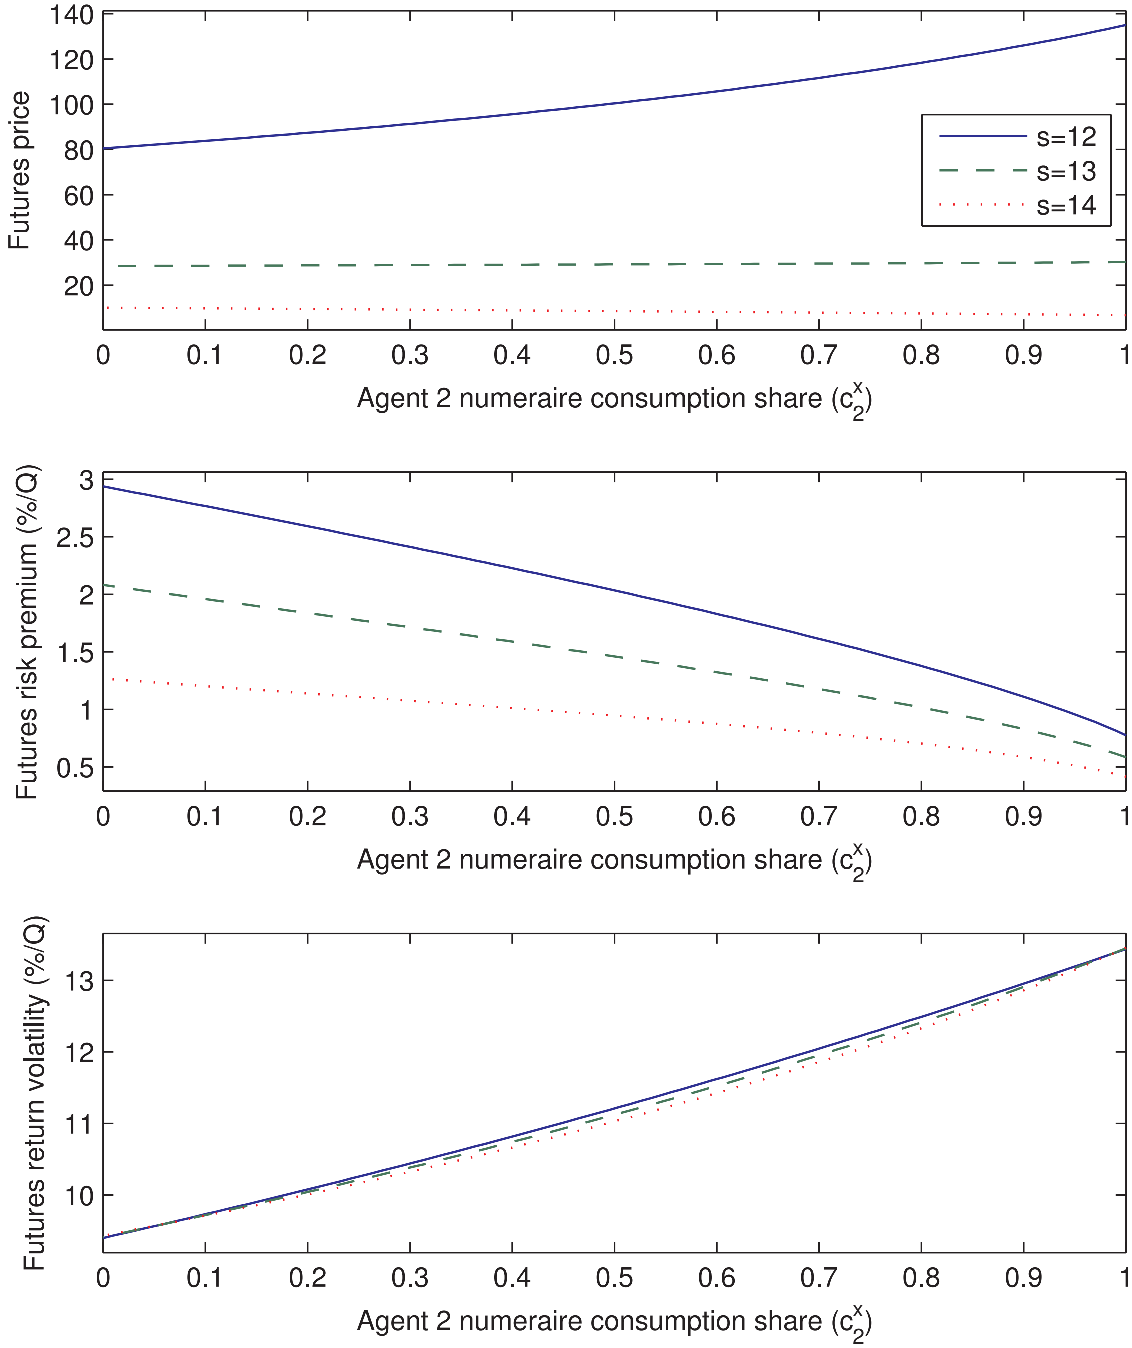

Prices, risk premium, and volatility on 1-year oil futures, shown for each of three growth states. The selected states have the same aggregate growth rate $g(s)$ but different goods consumption ratios $h(s)$.

Figure 7: Zero-coupon bond characteristics versus maturity

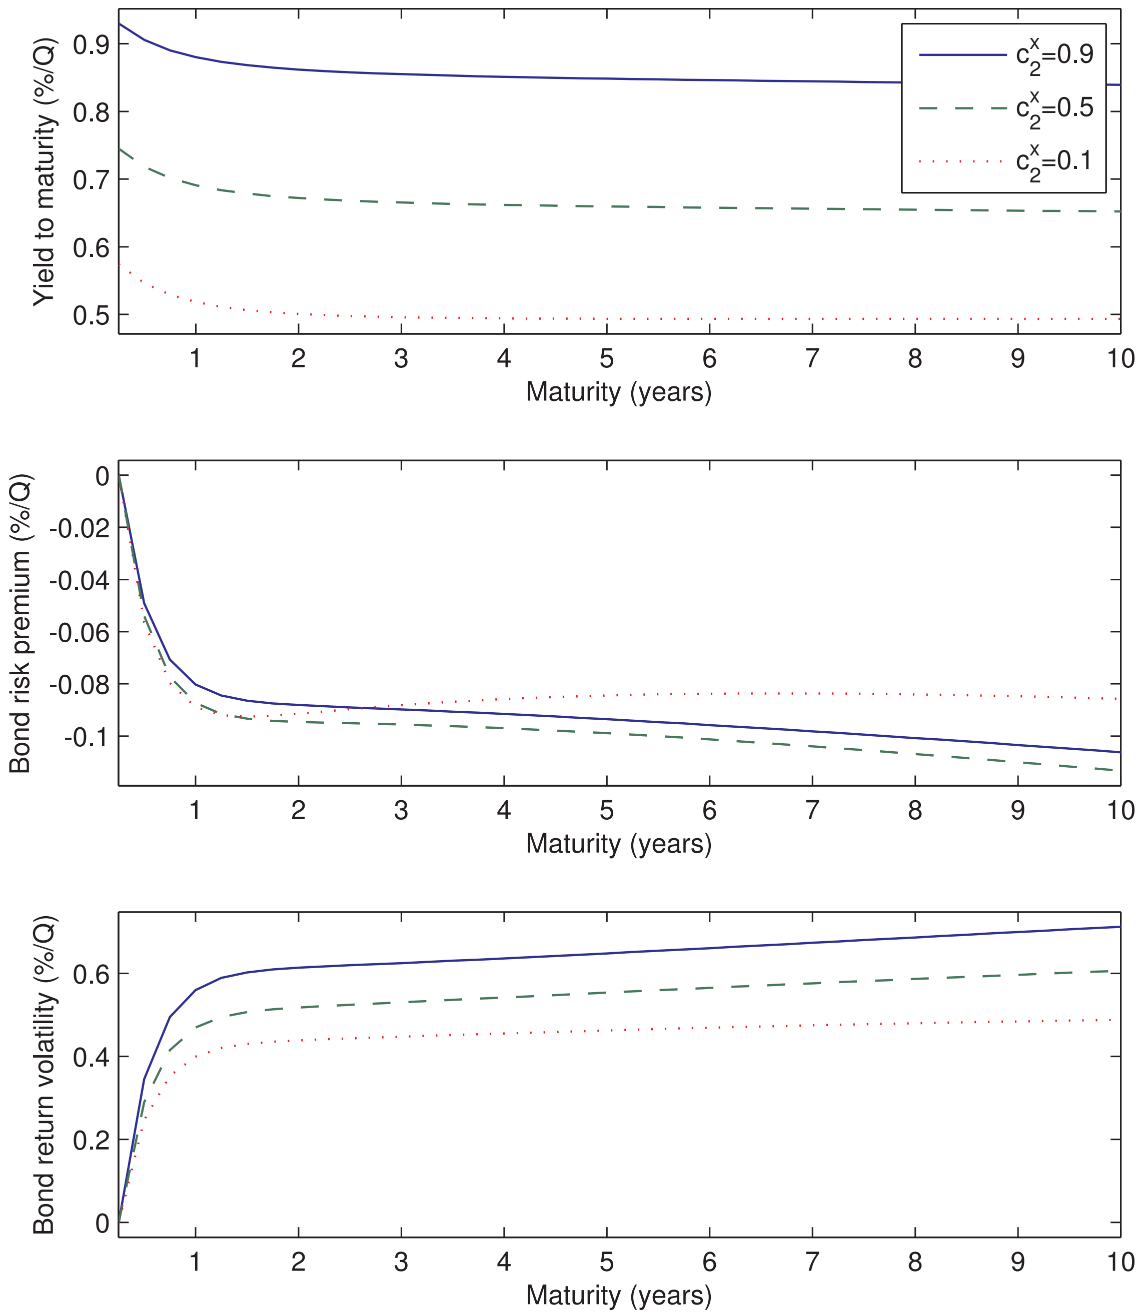

Yields to maturity, risk premium, and volatility of quarterly returns on zero coupon bonds, averaged over the stationary distribution of growth states.

Figure 8: Portfolios

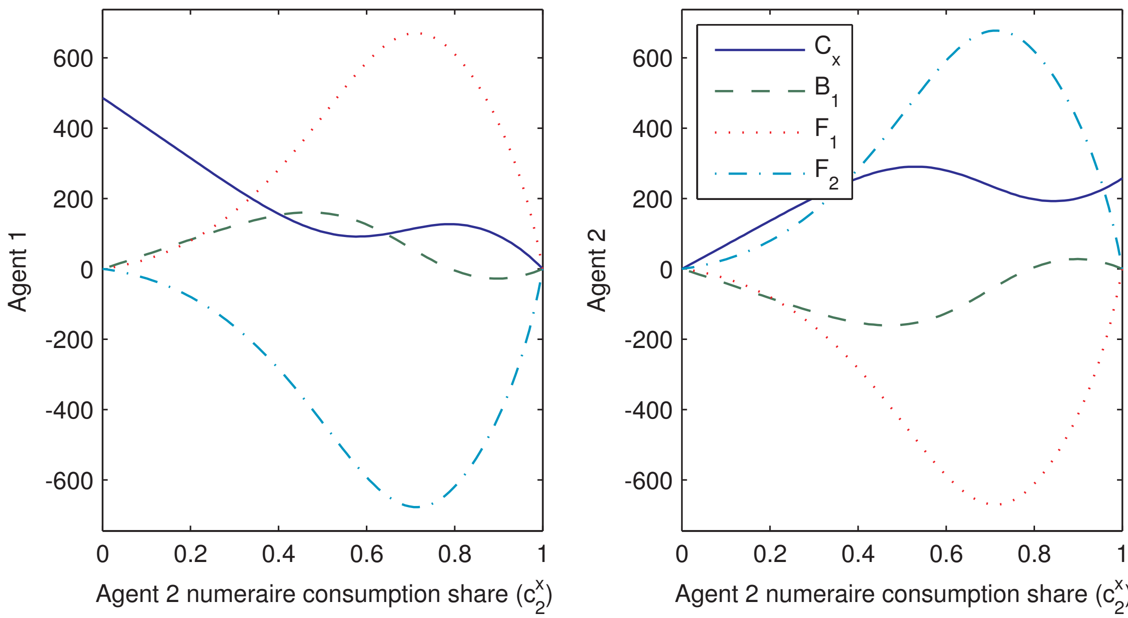

Portfolios for each agent, in terms of numeraire value of investment in each asset, versus c_2^x. Plots are averages over growth states using the stationary distribution.

Figure 9: Open interest in oil futures and spot prices

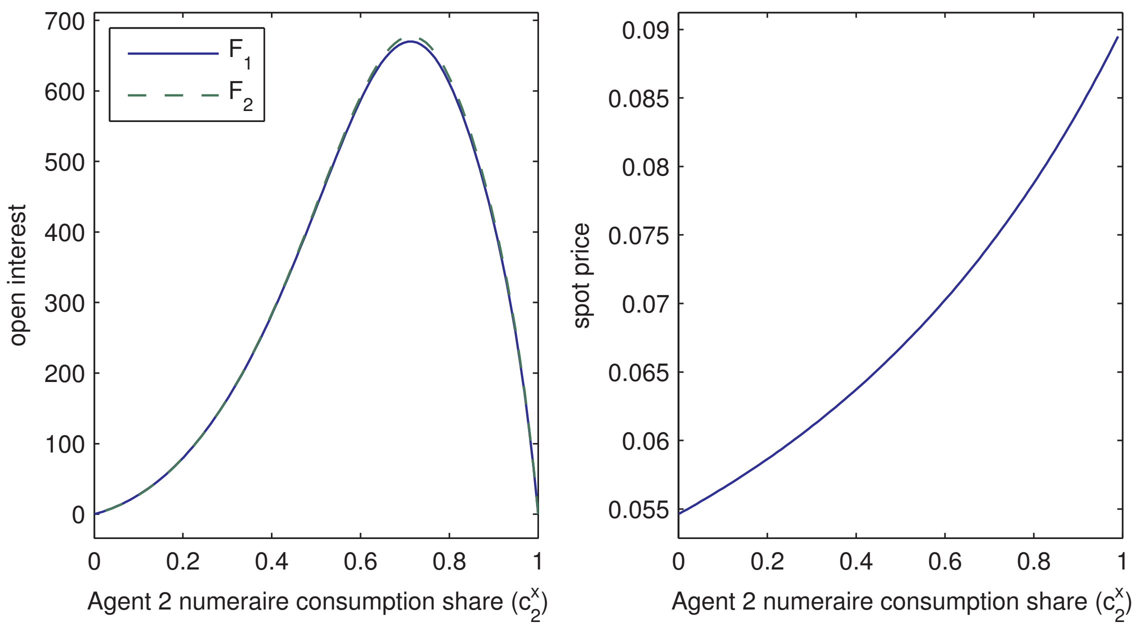

The left panel shows open interest in 3 and 6 month oil futures contracts, expressed as the numeraire value of the contracts, versus consumption distribution c_2^x. The right panel shows spot price of oil versus c_2^x. Results are averaged over the stationary distribution of growth states s.

Figure 10: Probability density of $c^x_{2,t}$

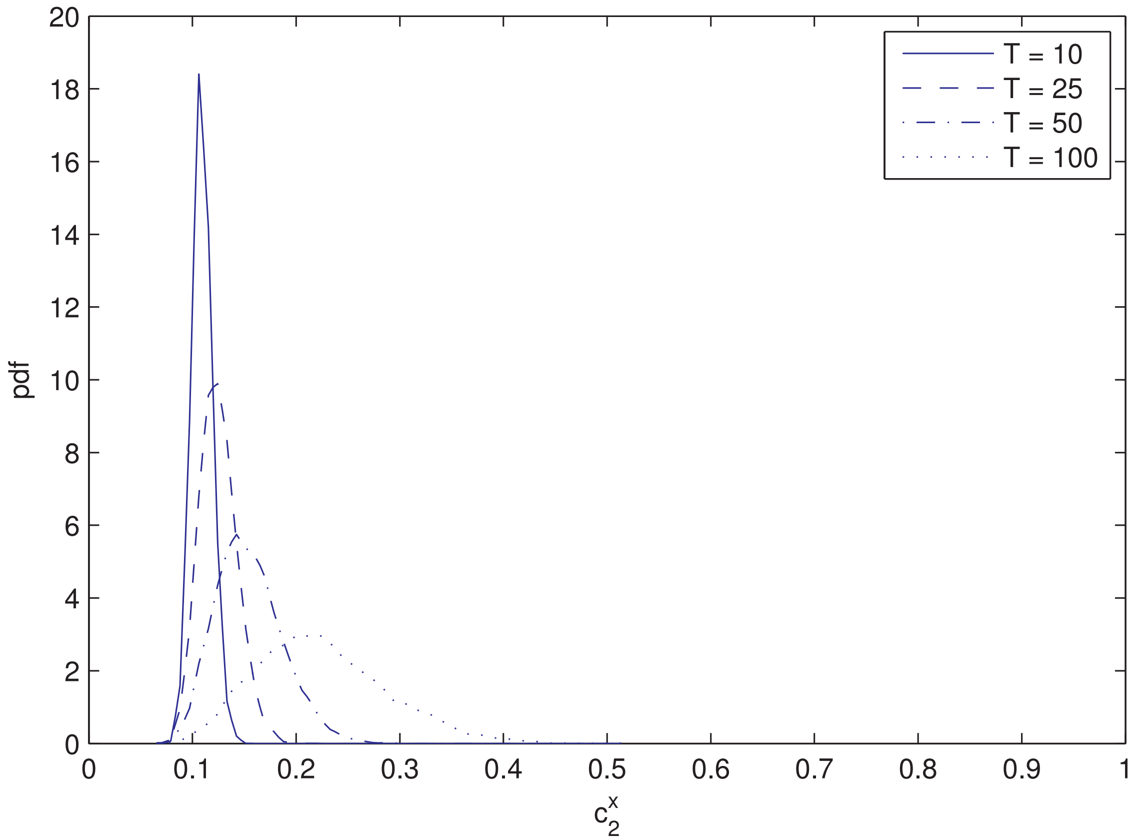

The probability density of $c^x_{2,t}$ is plotted at increasing horizons, of 10, 25, 50, and 100 years, conditional on initial value $c^x_{2,0} = 0.1$. Although $c^x_{2,t}$ has a discrete distribution conditional on $c^x_{2,0}$, we plot a continuous analog to the probability mass function for ease of visualization. The resulting plot has two main features: (1) $c^x_{2,t}$ exhibits an upward drift, such that values $c^x_{2,t} > c^x_{2,0}$ become very likely at longer horizons and (2) the probability mass becomes more dispersed, such that the range of probable values for $c^x_{2,t}$ becomes much wider for larger t. Results are computed using Monte Carlo simulation with 10000 paths.

Figure 11: Evolution over time

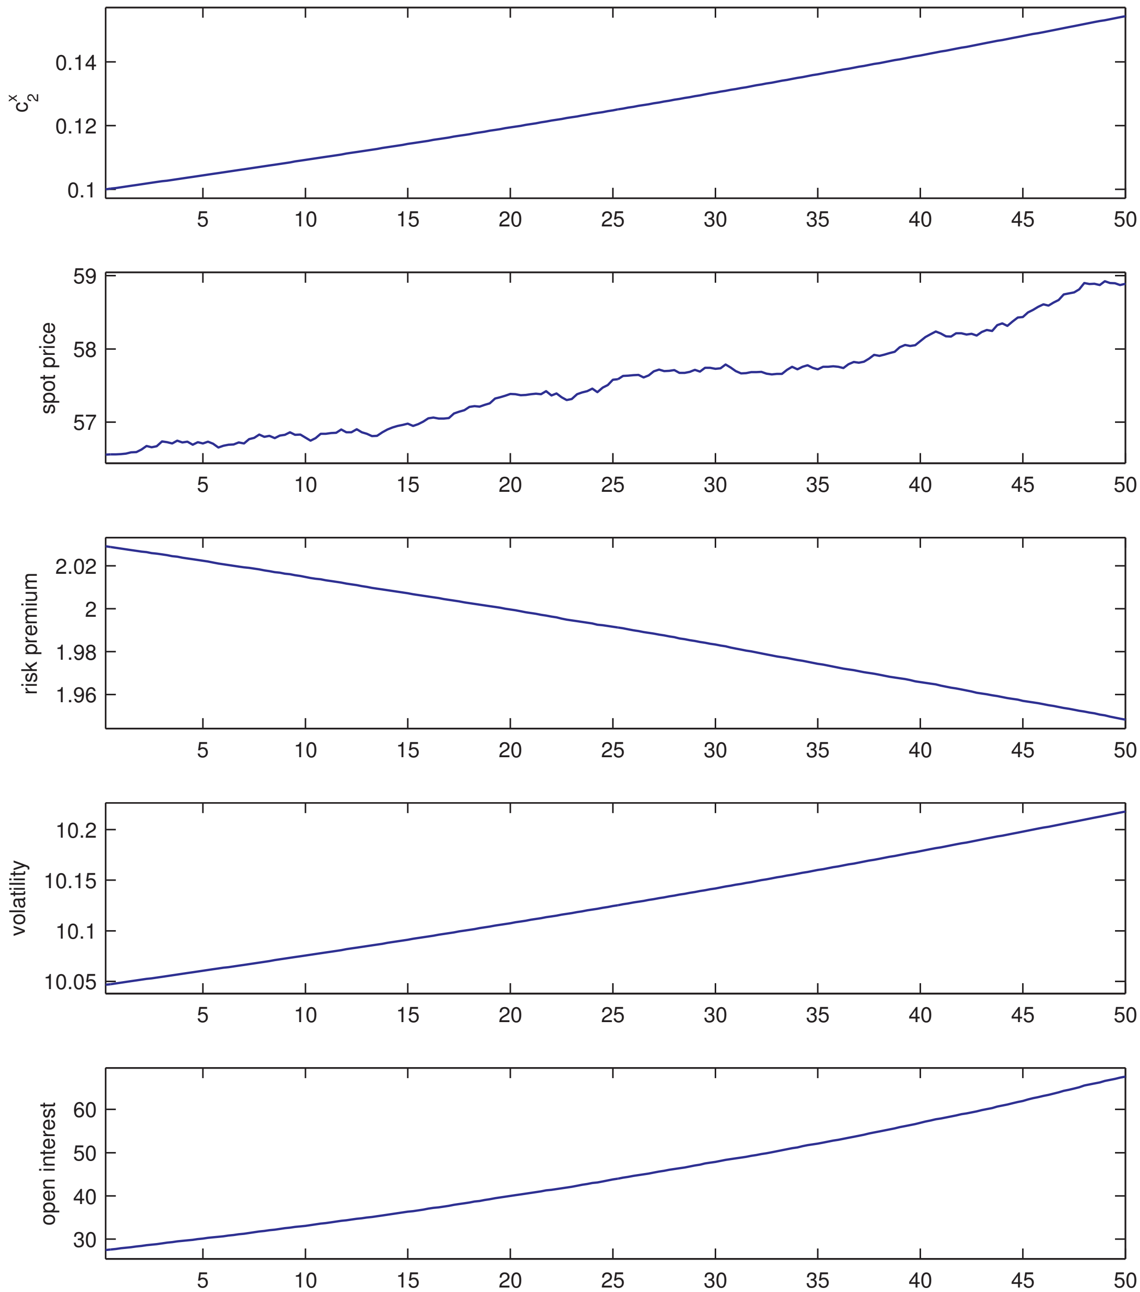

The plots illustrate the average path of the economy over a 50-year period, conditional initial $c^x_{2,0} = 0.1$. The initial growth state is selected independently for each path according to the stationary distribution. From top to bottom, the panels show $c^x_{2,t}$, the oil spot price, risk premium to the long side of the nearest (3-month) oil futures contract, volatility of returns to the nearest oil futures contract, and open interest in the nearest futures contract. Over time, the economy is likely to exhibit a rising spot price, increasing open interest, decreasing futures risk premium, and increasing oil price volatility. Results are computed using Monte Carlo simulation with 50000 paths.