The Price of Oil Risk*

Steven D. Baker,[†] Bryan R. Routledge,[‡]

September 2011
[December 20, 2012]

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.

* Thanks to seminar participants at the Carnegie Mellon University, Federal Reserve Board, University of Texas at Austin, and Alberta Finance Institute Conference: Speculation, Risk Premiums, and Financing Conditions in Commodity Markets at the University of Calgary for comments.

† Tepper School of Business, Carnegie Mellon University; sdbaker@cmu.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 Hamilton and Wu (2011) 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 Keynes (1936) 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

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

[2] The clearest argument along this lines is by James Hamilton http://www.econbrowser.com/archives/2011/08/fundamentals_sp.html. See also Hamilton (2009) and Wright (2011)

[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 Schwartz (1997)). Alternatively, the dynamics of storage or production can be modeled directly to capture the contemporaneous relationship between spot and futures price (Routledge, Seppi, and Spatt (2000), (2001), Titman (2011), and others).

"normal backwardation" (or "natural") refers to the situation where the current one-year 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. At all maturities, you can see from Figure 2(a), excess return or risk premium is much higher in the post-2004 sample. In the time series, Figure 2(b) you can see the variation across the subsample is reflecting the steady increase in excess returns over the 2000-2004 period particularly in the longer-dated contracts. Hamilton and Wu (2011) estimate the time variation in the risk premium as a structural break around 2004. Of course, time variation in risk premia is not surprising in modern asset pricing. We see it in equity returns (Campbell and Cochrane (1999), Bansal and Yaron (2005), Routledge and Zin (2010)) and bond returns Cochrane and Piazzesi (2005)).

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 Epstein and Zin (1989) and Kreps and Porteus (1978). 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 Bansal and Yaron (2005) 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., Kihlstrom and Mirman (1974)). 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, 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 transitive changes to the oil futures curve, gradual shifts in the wealth distribution produce long-run changes in the typical behavior of futures markets. Depending on the endogenous wealth distribution, the oil futures curve may be upward sloping 80% of the time, or only 6% of the time. The expected risk premium on oil futures (averaged over a period of 10 years, say) may be over 3%, or less than -4%. Open interest in futures markets may be trivial, or orders of magnitude larger than the value of aggregate oil consumption. And the impact is not limited to oil futures markets: we show that the interaction of the two agents may amplify the equity premium, and also causes it to fluctuate over time.

Our dynamic analysis of the model shows that large changes in the behavior of asset prices are not only possible, but likely to occur over the span of a few decades. Contingent upon an initial wealth distribution, our economy is expected to produce a near doubling of the spot price, a four-fold increase in open interest in oil futures markets, and roughly a

1% decrease in the average risk premium on oil futures. These represent persistent changes that occur in addition to short term fluctuations brought about by temporary oil supply "crises" or "booms", which can cause the risk premium to spike to 11% or plunge to -2% immediately. In contrast to the changes brought about by temporary shocks, changes in the wealth distribution are felt in the long term. Because they also occur endogenously, they 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.

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 Backus, Routledge, and Zin (2009) and (2008). Foundational work in risk sharing with recursive preferences includes Lucas and Stokey (1984), Kan (1995), and more recently Anderson (2005). There are also several recent papers that are related, such as Colacito and Croce (2011), (2012) and Gârleanu and Panageas (2010). 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 Carlson, Khokher, and Titman (2007), Casassus, Collin-Dufresne, and Routledge (2007), and Kogan, Livdan, and Yaron (2009). 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,

$$r^n_{t+1} \quad = \quad \log\left(\frac{F_{t+1,n-1}-F_{t,n}+(F_{t,n}(\exp r^f_{t+1}))}{F_{t,n}}\right) \tag{1}$$

We are interested in risk premiums, so will look at the return in excess of the risk-free rate $\log r^f_{t+1}$ 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.

$$r^n_{t+1} - r^f_{t+1} \quad \approx \quad \log F_{t+1,n-1} - \log F_{t,n} \tag{2}$$

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.

To get a feel for excess returns, Figure 2 plots the realized returns for one-month return for a long position in an crude oil future at different maturities. Figure 2(a) realized excess return, which is a noisy proxy for risk premium, is much higher in the post-2004 sample. Figure 2(b) plots the average realized excess return for a rolling 60 month window (with the convention that the plot at date t is the mean excess return from t to t+60). Here you can see the the variation across the subsample is reflecting the steady increase in expected returns over the 2000-2004 period, particularly in the longer-dated contracts. The estimation is simplistic here. We are just look at average realized excess returns. Hamilton and Wu (2011) estimate this more carefully with a VAR. Their paper, estimates a structural break, concludes that the risk premium properties are quite different after 2004. In particular, there is much greater variation in the risk premium post 2004.

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., Carlson, Khokher, and Titman (2007), Casassus, Collin-Dufresne, and Routledge (2007), and Kogan, Livdan, and Yaron (2009) Routledge, Seppi, and Spatt (2000)). 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

[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`

this change, excess returns are higher post 2004. They are higher in both cases where 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 commodity directly to economic activity. Hamilton (2008) 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 Cochrane and Piazzesi (2005), 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 Fama and French (1987) (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 regressions (one for each futures horizon, n).

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

And note that if $\phi_{t,n}$ is a constant, then b_n should be one. Table 3 confirms that the b_n , particularly at longer horizons is significantly less than one. Interestingly, the coefficient is

[5] Casassus and Collin-Dufresne (2005) use the closely retated fact that the negative slope is highly correlated to a high spot price of oil.

[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.

also broadly decreasing in maturity suggesting 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.

3 Exchange Economy - Two Goods and Two Agents

We model an exchange or endowment economy as in (Lucas 1978). 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. We specify both of these endowment processes to have finite state stationary Markov growth rates. This is the usual tree structure where here we have two trees. 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 Epstein and Zin (1989) and Kreps and Porteus (1978). They are characterized by three "aggregators" (see Backus, Routledge, and Zin (2005)) 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 Cobb-Douglas aggregator: $A(x_t, y_t) = x_t^{1-\gamma} y_t^{\gamma}$ with $\gamma \in [0, 1]$. Intertemporal preferences over the aggregate good are represented with an Epstein-Zin recursive preference structure

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

For the finite-state Markovian growth process for endowment of (x_t, y_t), denote the state s_t, and 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). We will denote growth in the numeraire good as in $f_{t+1} = f(s_{t+1}) = x_{t+1}/x_t$ and similar for the oil-good $g_{t+1} = g(s_{t+1}) = y_{t+1}/y_t$. With a little algebra, we can write intertemporal marginal rate of substitution, written here as a pricing kernel or stochastic discount factor.

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

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,

$$P_t = \frac{\partial W_t/\partial y_t}{\partial W_t/\partial x_t} = \frac{\gamma x_t}{(1-\gamma)y_t}. \tag{4}$$

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

The homogeneity of the Cobb-Douglas aggregator along with the standard homogeneity of the time and risk aggregators allow us to rescale things so utility is stationary (as in Hansen, Heaton, and Li (2008)). Define, $\hat{W}_t = W_t/A(x_t, y_t)$.

$$\begin{aligned} \hat{W}_t &= \left[(1-\beta) + \beta\mu_t\left(\hat{W}_{t+1}\frac{A(x_{t+1}, y_{t+1})}{A(x_t, y_t)}\right)^{\rho}\right]^{1/\rho} \\ &= \left[(1-\beta) + \beta\mu_t\left(\hat{W}_{t+1}A(f_{t+1}, g_{t+1})\right)^{\rho}\right]^{1/\rho} \end{aligned}$$

Note this uses the Cobb-Douglas property that $\frac{A(x_{t+1},y_{t+1})}{A(x_t,y_t)} = A(f_{t+1}, g_{t+1})$. Written, in this form, note that $\hat{W}_t$ is stationary and is a function only of the current state s_t. Similarly, substituting $\hat{W}_t$ into the pricing kernel, we have

$$m_{t+1} = \beta \left(f_{t+1}\right)^{-1} \left(A(f_{t+1}, g_{t+1})\right)^{\rho} \left(\frac{A(f_{t+1}, g_{t+1})\hat{W}_{t+1}}{\mu_t(A(f_{t+1}, g_{t+1})\hat{W}_{t+1})} \right)^{\alpha-\rho}. \tag{5}$$

The pricing kernel depends on the current state s_t, via the conditional expectation in the risk aggregator, and $t+1$ growth state s_{t+1}.

The price of oil depends on the relative levels of the two goods. However, changes in the oil price will depend only on the relative growth rates:

$$\frac{P_{t+1}}{P_t} = \frac{\gamma}{(1-\gamma)} \frac{f_{t+1}}{g_{t+1}}.$$

This implies the change in price depends only on the growth state s_{t+1}. To keep the level of the price rate plausible requires a joint assumption about the growth rates of the two goods (e.g., cointegrated).

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 (Bansal and Yaron (2005)). Alternatively, one can look at more sophisticated aggregators or risk to match return moments (Routledge and Zin (2010)). Presumably, one could take a similar approach to extend to a two-good case to look at oil prices and risk premia (see Ready (2010) as a nice example). It would require some work in our specific set up, since oil prices (and all derivatives) will simply depend on the current growth state s_t in combination with constant preference parameters γ. 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 second state variable, besides s_{t+1}, 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 is 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 $A_1(c^x_{1,t}, c^y_{1,t}) = (c^x_{1,t})^{1-\gamma_1}(c^y_{1,t})^{\gamma_1}$. The utility from the stochastic stream of this aggregated good has the same recursive form as above.

$$\begin{aligned} W_t = W(c^x_{1,t}, c^y_{1,t}, W_{t+1}) &= \left[(1-\beta)A_1(c^x_{1,t}, c^y_{1,t})^{\rho_1} + \beta\mu_{1,t}(W_{t+1})^{\rho_1}\right]^{1/\rho_1}, \\ \mu_{1,t}(W_{t+1}) &= E_t\left[W^{\alpha_1}_{t+1}\right]^{1/\alpha_1}, \end{aligned} \tag{6}$$

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

$$\begin{aligned} V_t = V(c^x_{2,t}, c^y_{2,t}, V_{t+1}) &= \left[(1-\beta)A_2(c^x_{2,t}, c^y_{2,t})^{\rho_2} + \beta\mu_{2,t}(V_{t+1})^{\rho_2}\right]^{1/\rho_2}, \\ \mu_{2,t}(V_{t+1}) &= E_t\left[V^{\alpha_2}_{t+1}\right]^{1/\alpha_2}. \end{aligned} \tag{7}$$

The idea is that the two agents can differ about the relative importance of the oil good, 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 consumption bundle. "Oil risk" does not just depend on the γ parameter since it involves a 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:

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

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, Lucas and Stokey (1984), and Kan (1995):

$$\begin{aligned} J(x_t, y_t, V_t) = \max_{c^x_{1,t}, c^y_{1,t}, V_{t+1}} &\left[(1-\beta)A_1(c^x_{1,t}, c^y_{1,t})^{\rho_1} + \beta\mu_{1,t}(J(x_{t+1}, y_{t+1}, V_{t+1}))^{\rho_1}\right]^{1/\rho_1} \\ \text{s.t. } &V(x_t - c^x_{1,t}, y_t - c^y_{1,t}, V_{t+1}) \geq V_t. \end{aligned} \tag{8}$$

[7] Differing β's are easy to accommodate but lead to uninteresting models since the agent with the larger β quickly dominates the optimal allocation. E.g., Yan (2010)

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. Note this promised utility is a vector since we choose one for each possible growth state s_{t+1} at date $t+1$. Making good on these promises at date $t+1$ means that V_{t+1} is an endogenous state variable we need to track. That is, optimal consumption at date t depends on the exogenous growth state s_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[8] are

$$
\begin{aligned}
W_x(c^x_{1,t}, c^y_{1,t}, W_{t+1}) &= \lambda_t V_x(x_t - c^x_{1,t}, y_t - c^y_{1,t}, V_{t+1}) && (9)\\
W_y(c^x_{1,t}, c^y_{1,t}, W_{t+1}) &= \lambda_t V_y(x_t - c^x_{1,t}, y_t - c^y_{1,t}, V_{t+1}) && (10)\\
W_V(c^x_{1,t}, c^y_{1,t}, W_{t+1}) &= \lambda_t V_V(x_t - c^x_{1,t}, y_t - c^y_{1,t}, V_{t+1}) && (11)\\
J_V(x_t, y_t, V_t) &= -\lambda_t, && (12)
\end{aligned}
$$

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

$$
\begin{aligned}
m_{t+1} &= \beta \left(\frac{c^x_{1,t+1}}{c^x_{1,t}}\right)^{-1} \left(\frac{A_{1,t+1}}{A_{1,t}}\right)^{\rho_1} \left(\frac{W_{t+1}}{\mu_{1,t}(W_{t+1})}\right)^{\alpha_1-\rho_1} \\
&= \beta \left(\frac{c^x_{2,t+1}}{c^x_{2,t}}\right)^{-1} \left(\frac{A_{2,t+1}}{A_{2,t}}\right)^{\rho_2} \left(\frac{V_{t+1}}{\mu_{2,t}(V_{t+1})}\right)^{\alpha_2-\rho_2}.
\end{aligned} \tag{13}
$$

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:

$$
P_t = \frac{\gamma_1 c^x_{1,t}}{(1-\gamma_1)c^y_{1,t}} = \frac{\gamma_2 c^x_{2,t}}{(1-\gamma_2)c^y_{2,t}}. \tag{14}
$$

[8] Many papers directly characterize the "stochastic Pareto weight" process. E.g, Basak and Cuoco (1998), Basak (2005), Dumas, Kurshev, and Uppal (2009).

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

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

The $\hat{c}$'s are consumption shares of the two goods. Scale utility values by their respective aggregated goods.

$$\hat{W}_t = \frac{W_t}{A_1(x_t, y_t)} \quad , \quad \hat{V}_t = \frac{V_t}{A_2(x_t, y_t)}$$

Notice we scale the utilities by the total available goods (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 (13), and we can state the pricing kernel as

$$m_{t+1} = \beta \left(\frac{f_{t+1}\hat{c}^x_{1,t+1}}{\hat{c}^x_{1,t}} \right)^{-1} \left(\frac{A_1(f_{t+1}\hat{c}^x_{1,t+1}, g_{t+1}\hat{c}^y_{1,t+1})}{A_1(\hat{c}^x_{1,t}, \hat{c}^y_{1,t})} \right)^{\rho_1} \left(\frac{A_1(f_{t+1}, g_{t+1})\hat{W}_{t+1}}{\mu_1(A_1(f_{t+1}, g_{t+1})\hat{W}_{t+1})} \right)^{\alpha_1 - \rho_1} \tag{15}$$

(or equivalently from the perspective of agent 2). In the one-agent case, the pricing kernel depends on the current growth state s_t (via conditional exceptions) and the future growth state s_{t+1}. Now, in the two agent case, the pricing kenel depends on both the growth state and the level of utility (scaled) of agent 2, $(s_t, \hat{V}_t)$. The current state shows up in expectations and the promised utility influences the allocations of the goods across agent one and two. In addition, the pricing kernel depends on the date-$t+1$ values, $(s_t, \hat{V}_{t+1})$ realized.

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 (i.e., a measure of their wealth). Agent 1's claim to numeraire consumption good x has value at date-t denoted $C^x_{1,t}$

$$C^x_{1,t} = E_t \left[\sum_{\tau=t}^{\infty} m_\tau c^x_{1,\tau} \right]. \tag{16}$$

Note, by convention this is the "*cum* dividend" value including current consumption. To solve this, conjecture that $C^x_{1,t} = \frac{c^x_{1,t} W_t^{\rho_1}}{(1-\beta) A_{1,t}^{\rho_1}}$, and verify. Note it is easier here to use the kernel defined in equation (13)

$$\begin{aligned}
C^x_{1,t} &= c^x_{1,t} + \beta E_t \left[\left(\frac{c^x_{1,t+1}}{c^x_{1,t}} \right)^{-1} \left(\frac{A_{1,t+1}}{A_{1,t}} \right)^{\rho_1} \left(\frac{W_{t+1}}{\mu_{1,t}(W_{t+1})} \right)^{\alpha_1 - \rho_1} \frac{c^x_{1,t+1} W_{t+1}^{\rho_1}}{(1-\beta) A_{1,t+1}^{\rho_1}} \right] \\
&= c^x_{1,t} + \frac{\beta c^x_{1,t} \mu_{1,t}(W_{t+1})^{\rho_1 - \alpha_1} E_t \left[W_{t+1}^{\alpha_1} \right]}{(1-\beta) A_{1,t}^{\rho_1}} \\
&= \frac{\left[(1-\beta) A_{1,t}^{\rho_1} + \beta \mu_{1,t}(W_{t+1})^{\rho_1} \right]}{(1-\beta) A_{1,t}^{\rho_1}} \\
&= \frac{c^x_{1,t} W_t^{\rho_1}}{(1-\beta) A_{1,t}^{\rho_1}} \\
&= \left(\frac{\hat{c}^x_{1,t} \hat{W}_t^{\rho_1}}{(1-\beta) A_1(\hat{c}^x_{1,t}, \hat{c}^y_{1,t})^{\rho_1}} \right) x_t.
\end{aligned} \tag{17}$$

Note that the price of claims to numeraire consumption is independent of the level of oil consumption (it does depend on ratio of oil to numeraire good.).; that is the ratio $C^x_{1,t}/x_t$ is stationary and depends only on our state variables s_t and $\hat{V}_t$. To price the claim to the oil consumption good, we use the oil price to convert to units of numeraire good.

$$C^y_{1,t} = E_t \left[\sum_{\tau=t}^{\infty} m_\tau P_\tau c^y_{1,\tau} \right].$$

From the spot price, (14), we can write $P_t c^y_{1,t} = \frac{\gamma_1 c^x_{1,t}}{1-\gamma_1}$, so

$$\begin{aligned}
C^y_{1,t} &= E_t \left[\sum_{\tau=t}^{\infty} m_\tau c^x_{1,t} \frac{\gamma_1}{1-\gamma_1} \right] \\
&= \frac{\gamma_1}{1-\gamma_1} C^x_{1,t}
\end{aligned}$$

Again, note the level of oil does not play a role and the ratio $C^y_{1,t}/x_t$ is stationary and depends on the state variables s_t and V_t (alternatively we could scale by y_t). Lastly, summing the value of the numeraire and oil claim we calculate the total wealth of agent one.

$$C_{1,t} = C^x_{1,t} + C^y_{1,t} = \frac{1}{1-\gamma_1} C^x_{1,t} \tag{18}$$

By equivalent logic, the values of agent 2's consumption claims are

$$
\begin{aligned}
C^x_{2,t} &= \left(\frac{\hat{c}^x_{2,t}\hat{V}^{\rho_2}_t}{(1-\beta)A_2(\hat{c}^x_{2,t},\hat{c}^y_{2,t})^{\rho_2}}\right)x_t. \\
C^y_{2,t} &= \frac{\gamma_2 C^x_{2,t}}{1-\gamma_2} \\
C_{2,t} &= \frac{C^x_{2,t}}{1-\gamma_2}
\end{aligned}
$$

With the wealth of each agent, we can calculate the wealth of both sectors of the economy, numeraire and oil, and the overall wealth of the economy. Wealth in the numeraire sector,

$$C^x_t = C^x_{1,t} + C^x_{2,t}, \tag{19}$$

in the oil sector,

$$C^y_t = \frac{\gamma_1 C^x_{1,t}}{1-\gamma_1} + \frac{\gamma_2 C^x_{2,t}}{1-\gamma_2}, \tag{20}$$

and the overall wealth in the economy,

$$C_t = \frac{C^x_{1,t}}{1-\gamma_1} + \frac{C^x_{2,t}}{1-\gamma_2}. \tag{21}$$

Given these processes for wealth, we can calculate the return to a claim on these assets. We define the equity return as a claim to the numeraire stream of consumption, so

$$r^x_{t+1} = \log C^x_{t+1} - \log(C^x_t - x_t)$$

(We could also compute the equity return including the entire wealth in the economy. Given our calibration in the next section, the difference is minor).

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

$$B_{t,n} = E_t[m_{t+1}B_{t+1,n-1}], \tag{22}$$

where $B_{t,n}$ is the price of a bond at t paying a unit of the numeraire good at period $t+n$ period 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

$$\begin{aligned} 0 &= E_t[m_{t+1}(F_{t+1,n-1} - F_{t,n})] \\ F_{t,n} &= (B_{t,1})^{-1} E_t[m_{t+1} F_{t+1,n-1}], \\ \log F_{t,n} &= -\log B_{t,1} + \log E_t \left[\exp\left\{\log m_{t+1} + \log F_{t+1,n-1}\right\}\right], \end{aligned} \tag{23}$$

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 who 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, we can look analytically at how "equity" claims can implement the optimal allocations.

Recall that C^x_t is the value of a claim to the stream of numeraire good and C^y_t is the value of the claim to a stream of the oil good. We think of these as (unlevered) claims to the equity in the numeraire and oil sectors and normalize the shares outstanding in each sector to be one. Suppose these were traded claims in the economy, what portfolio of $\phi^x_{1,t}$ shares numeraire and $\phi^y_{1,t}$ shares in oil generate optimal consumption? It turns out this is easy to solve since all we need to do is replicate the wealth processes that represents optimal consumptions of the two goods for agent one (and, analogously agent two). The agents' budget constraint are:

$$\begin{aligned} C_{1,t} &= \phi^x_{1,t} C^x_t + \phi^y_{1,t} C^y_t \\ C_{2,t} &= \phi^x_{2,t} C^x_t + \phi^y_{2,t} C^y_t \end{aligned}$$

Substitute in the definition of the aggregate value of the numeraire sector in equation (19) and oil sector in equation 20). The key here is that for each agent, the value of the oil consumption stream is proportional to the value of the numeraire stream, i.e.,

$$\frac{C^y_{1,t}}{C^x_{1,t}} = \frac{\gamma_1}{1-\gamma_1} \quad , \quad \frac{C^y_{2,t}}{C^x_{2,t}} = \frac{\gamma_2}{1-\gamma_2}$$

This all implies, for agent one:

$$C^x_{1,t} = \frac{(1-\gamma_1)\left((-\gamma_2 C^x_t + (1-\gamma_2)C^y_t\right)}{\gamma_1 - \gamma_2} \tag{24}$$

and

$$\begin{aligned} \phi^x_{1,t} &= \frac{-\gamma_2}{\gamma_1 - \gamma_2} \\ \phi^y_{1,t} &= \frac{1-\gamma_2}{\gamma_1 - \gamma_2}. \end{aligned} \tag{25}$$

(And similarly for agent two). Note the right hand side is a constant. 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 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. We calibrate our example so one of the goods matches oil consumption and match basic moments of the observed risk premia. To compare the model's implications with the data, we study an example with a four-state Markov growth process loosely calibrated to annual moments. We assume that the numeraire (x) and oil (y) processes are cointegrated, with unconditional (i.e., long-run) mean growth rates of 2% per year. Unconditional standard deviations of numeraire and oil consumption growth are 3% and 6%, respectively, reflecting the higher variability of oil relative to aggregate consumption. To obtain a reasonable equity premium in this setting, we need either small highly persistent risks to consumption growth as in Bansal and Yaron (2005), or rare disaster-like risk as in Barro (2009). To make our numerical computations feasible, we have a four state Markov structure. This gives our calibration the flavor of a disaster model, and allows for the possibility of a sharp drop in either x or y, with the likelihood of such occurrences substantially lower than that of moderate, positive growth in both x and y. Table 4 gives the full specification of the growth process. Of course a simple 4-state Markov cannot realistically replicate all aspects of household oil-derived and non-oil consumption but does seem to generate some interesting results. For a more thorough investigation of empirical issues calibrating a model with household consumption of oil products, see Ready (2010).[9] 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 Hamilton (2008). In our model this characteristic is governed chiefly by the choice of goods aggregation parameters. We set γ_1 and $\gamma_2 = 2\gamma_1$ to give agent 2 a substantially greater preference for oil consumption while keeping oil within a plausible range of the 4% historical average. Risk aversion (α) and intertemporal substitution (ρ) parameters address the risk premia on equity and oil futures. We aim to match the large level of the equity premium. For oil risk premium, our goal is to generate variation consistent with what we infer from the data. Hamilton and Wu (2011), for example, suggests a range 4 % to -3% over the 1990-2010 period. Specifically, that paper suggests a lower oil-risk premium in recent years. Our parameters imply a positive oil risk premium in an economy dominated by agent 1, and a negative risk premium under agent 2. The dynamic properties of the risk-sharing model with generate the variation. We set the common time-preference parameter, β such that the average risk-free rate is around 2%. Finally, we choose initial consumption levels x_0 and y_0 to put the level of spot prices in the ballpark of those observed; the choice has no impact upon returns or the dynamics of the model.

Because of our assumptions of infinitely lived agents and a Markovian growth process, the state of our economy at any point in time is fully characterized by output levels for each good x_t and y_t, the growth state s_t, and promised agent 2 (normalized) utility $\hat{V}_t$. Therefore

[9] One issue we omit here is technological progress. A gallon of gasoline in 1975 is consumed much differently than the same gallon in 2012. This is a central issue in resource economics that more carefully considers the long run implications of "peak oil."

we drop time subscripts and describe the economy as a function of the state variables. As a standing assumption we set x and y to the values in Table 5 unless otherwise specified. We further focus our analysis on the key state variable $\hat{V}$, by taking expectations over growth states according to their stationary distribution where necessary, given in Table 4. Section 4.1 shows how $\hat{V}$ 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 V in response to growth shocks offer a possible explanation for observed changes in prices, 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. The novel state variable governing this interaction is $\hat{V}$, the utility promised by agent 1 to agent 2. It is interesting to see how this state variable maps to observable indicators of an agent's relative importance in the economy: his wealth and consumption shares. Recall $\hat{V}$ is bounded and stationary. To facilitate comparison, we re-normalize things so that this state variable is on the domain $[0, 1]$.[10] Roughly, a value for the V of 0.25 corresponds to agent 2 owning 25% of the wealth. To see this, Figure 3 relates V to the wealth and consumption shares of each agent. The plots show expectations over growth states $s \in S$ taken according to stationary distribution $\bar{\pi}$.[11] In the top panel, we see that V relates to wealth in a nearly linear, one-to-one mapping. Therefore V is a close proxy for agent 2's share of aggregate wealth. Similarly, the center panel shows that agent 2's share of numeraire consumption rises almost (but not exactly) linearly with his wealth, and agent 1's share declines correspondingly. The consumption sharing rule for oil is substantially different. The bottom panel illustrates agent 2's higher preference for oil consumption relative to agent 1; as V increases, his share of oil consumption increases more rapidly than his wealth share, such that he consumes roughly 50 % of oil when he holds around 33 % of the wealth. The results for consumption follow directly from the goods

[10] Although an agent's utility may grow without bound as the economy expands, we have observed previously that $\hat{V}_t = \frac{V_t}{A_2(x_t, y_t)}$ is bounded for any x_t and y_t. However the domain of $\hat{V}$ is determined in equilibrium based on the model parameters. Its minimum value is 0, and its maximum $\hat{V}_{\max}$ corresponds to the case where agent 2 consumes the aggregate output of each good, i.e. $c^x_{2,t} = x_t$, $c^y_{2,t} = y_t$ in perpetuity. Therefore we further normalize by dividing $\hat{V}$ by $\hat{V}_{\max}$, to obtain a state variable between zero and 1.In fact there is a further nuance, since $\hat{V}_{\max}$ depends upon the current growth state; to be precise, we normalize $\hat{V}$ by its conditional maximum.

[11] Although results do differ for each $s \in S$, to the naked eye they are all very similar to the mean values shown.

aggregation parameters: because $\gamma_2 = 2\gamma_1$, agent 2 is twice as inclined to spend on oil as agent 1. However since oil represents a relatively small fraction of expenditures for either agent (γ_1 is small), shares of numeraire consumption are closely related to wealth. The conclusion (given our calibration) is that V can safely be interpreted as a measure of wealth distribution, but not consumption distribution.

4.2 Prices and risk premia

Changes in the wealth distribution have a dramatic impact upon the level, term structure, and risk premia of financial assets. This section first examines the effects of V unconditionally of the growth state s. Later we condition our analysis on the growth state, as felt through its effect upon the slope of the futures curve.

Figure 4 shows the average term structure of oil futures prices conditional on a given V (effectively wealth distribution). The most obvious impact is on the level of prices, which increase dramatically with higher V, doubling prices at short maturities and almost tripling prices at longer horizons. The slope of the futures curve changes from moderately downward sloping at low V to sharply upward sloping for high V. Although the affects of preferences towards risk and consumption goods cannot be cleanly separated (recall Kihlstrom and Mirman (1974)), the difference in level at the short end is attributable to different preferences over consumption baskets (γ), whereas changes in slope are strongly impacted by preferences towards risk and intertemporal substitution (α and ρ). To the extent that our two agents represent a fair approximation to the broad range of consumer preferences relevant for international oil markets, Figure 4 illustrates that large changes in the average level or shape of the futures curve needn't imply a structural change in oil output, the imposition or removal of market frictions, or other stimuli: they may simply represent an endogenous change in the "tastes" of the representative agent driven by trade, and the resulting shifts in wealth. Large variations in V can and do occur within our model, as we will see in Section 4.4. Changes in V are analogous to endogenous "demand shocks", even though they do not represent an additional source of randomness. For example, the increasing *absolute* wealth of China represents an additional demand for oil commonly offered as a partial explanation for the spike in oil prices in 2007-2008 Hamilton (2011). It seems reasonable to suppose that the increasing *relative* wealth of China would also have an impact upon oil markets, through different attitudes towards energy consumption and risk. We do not explicitly impose the role of countries on our agents, but this is one possible interpretation.

The changes observed in the slope of the futures curve imply changes in risk premia, which are illustrated in Figure 5. On average, an economy dominated by agent 1 (low V) implies a positive and hump-shaped term structure of oil futures risk premia, with the highest risk premium on the 2-year contract. As the share of wealth given over to agent 2 increases, the level of risk premium decreases across maturities, with the term structure first flattening and then turning concave for high V. For example, for small V, the 2-year contract offers a positive risk-premium that is higher than that of the 1-year contract. But for large V, the two-year contract has a negative risk-premium that is lower than that of the 1-year contract. Figure 6 offers another view of the effect of V on the oil risk premium, plotting the 2-year contract premium continuously versus V. The premium declines monotonically but non-linearly, from almost 3 % when agent 1 is wealthy to just above -2 % when agent 2 is wealthy. Shown on the figure are the risk-premium conditional on the economic growth state. Recall that our parameterization puts the economy most frequently in state 2 and 3.

In Table 1 we see a negatively sloping oil futures curve implies higher excess returns. Table 6 summarizes expected excess holding returns on oil futures conditional on a positive or negative slope in our calibration. We define the slope as the difference between the spot price and the 2-year futures price. In the model, conditioning on slope in addition to V amounts to conditioning on the occurrence of certain growth states (s). For all contract maturities and all V, a negative slope implies a higher risk premium, consistent with the empirical results in Table 1. In particular, conditional on a negative slope, the 1-year contract always offers a positive risk premium in the model. The model diverges from empirical estimates in that the risk premium is not generally increasing with contract maturity, and risk premia may be sharply negative, particularly when the slope is positive. The evidence is that the premium at the short end is both smaller and less volatile. Our model is giving a more uniform pattern. Table 8 reproduces the Fama predictive regressions from Table 3. Note the time variation of the risk premia on oil is seen in the slope coefficient (b) being less than one. While the b's are decreasing with horizon, the effect is not as pronounced as in the data.

We can also examine relationships between the oil futures curve and the term structure of interest rates. Figure 7 shows the average term structure for bonds for different V, which exhibit wide variation in level and slope. The case where agent 1 is dominant (low V) produces the highest short rate of any curve (around 2.5 %), but due to a sharp downward slope, it also leads to the lowest long rate (nearly 0). For larger V the term structure flattens, then becomes upward sloping when agent 2 is dominant. The impact of V at the short end of the term structure is very different than at the long end: the short rate is non-monotonic in V, first decreasing then increasing, whereas the long rate is increasing in V. The net result is a positive risk premium on long bonds given high V, but a negative one given low V. On average, bond risk premia move opposite oil risk premia in relation to V. The situation is more nuanced when we condition on the slope of the futures curve. We saw in Table 2 that negatively sloping oil futures implied higher excess returns on long bonds. Although that finding is not robust in the model, a negative futures slope does imply higher excess returns on long-horizon bonds, except when V is very small.

The non-monotonicity evident in the risk-free rate is also present in the equity premium, per Figure 8, which plots the average equity premium relative to V. The smallest values of V correspond approximately to an economy populated only by agent 1, whereas the largest V imply an economy nearly dominated by agent 2. What is surprising is that the equity premium peaks around $V = 0.5$, when each agent has a similar share of the wealth. The opportunity for risk-sharing between agents actually *increases* the equity premium! A shift in wealth towards an agent who would individually demand a lower equity premium says nothing regarding the direction of the equilibrium effect: depending on the initial wealth distribution, the equity premium may either increase or decrease.

4.3 Portfolios and open interest

As we saw in Section 3.3, when claims to aggregate consumption of the numeraire and oil are traded in financial markets, then the agents can implement their optimal consumption plans with a constant portfolio. In our calibrated example, agent 1 would hold 2 shares of the non-oil (numeraire) stock and roughly -32.3 shares of the oil stock. Since we normalize net supply of each stock to 1 share, agent 2 holds -1 and 33.3 shares, respectively.[12] Even if the agents are allowed to trade oil futures, they find it unnecessary to do so. However, in practice investors are unable to trade a claim to aggregate consumption of oil, and close proxies may not exist in the stock market: multinational oil companies produce a small fraction of global output, with much of production due to state owned enterprises (e.g., Saudi Aramco, PDVSA) that are not publicly traded. And, it is perhaps not feasible to have such large short positions in an equity claim. An alternative and a direct way for investors to manage their exposure to oil is through oil futures contracts. We approximate this situation by allowing agents to trade in markets that are complete, but lack a directly tradable claim to aggregate oil consumption. Instead agents may trade the numeraire stock,

[12] When one agent or the other is dominant in the economy, that agent must hold 1 share of the x stock and one share of the y stock to clear markets. A little algebra confirms that the portfolios specified are equivalent to the dominant agent holding one share of each stock when $V \to 0$ or $V \to 1$.

a one-period bond, and "collateralized" 1 and 2 year oil futures contracts (labeled F_1 and F_2, respectively). Figure 9 shows each agent's optimal portfolios for a range of V. Since agent 2's holdings of futures contracts are approximately the mirror image of agent 1's, we describe only agent 1's position, in the top panel. We plot portfolios averaged over growth states according the the stationary distribution; however results for individual growth states are very similar to the mean. In contrast, changing V has a large effect on portfolios, with the value of agent 1's exposure to F_1 ranging from roughly 0 to -1000, and his exposure to F_2 from roughly 0 to 500. Despite his simple objective - to have constant exposure to oil consumption - his replicating portfolio may change substantially. Furthermore the direction of his exposure to oil futures is not the same for contracts of different maturity - he is short the 1-year contract and long the 2-year contract. What might appear to be a hedging strategy that "bets" on a change in the spread between 1 and 2 year contracts is actually a reflection of a simple preference characteristic: agent 1 desires less exposure to oil than agent 2.

Since much is made of changes to open interest in oil futures, we highlight this in the top panel of Figure 10, 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 V, and may be negligible, or orders of magnitude larger than the total value of oil consumed in the economy. It is also non-monotonic in V, peaking around $V = 0.35$. For comparison, the bottom panel of Figure 10 shows the impact of V on the spot price of oil, which is monotonic and almost linear in V, peaking when agent 2 is dominant in the economy ($V \to 1$), reflecting his higher preference for oil consumption. Obviously open interest is not a sufficient statistic to determine the price of oil, even if we know the current growth state. Ceteris paribus, if V increased from 0.1 to 0.3, we would observe a dramatic increase in open interest and a roughly 20 % increase in the spot price, whereas a decrease in V from 0.6 to 0.4 would also produce a significant increase in open interest, but accompanied by a roughly 20 % *decline* in the spot price. Therefore the directional change in open interest does not relate to spot prices, neither does it reflect an increase in "speculation" on the part of the agents in the economy.

4.4 Dynamics

The value of V has a strong impact upon spot prices, risk premia, and open interest. As the economy evolves, so does V, reflecting changes in the wealth distribution brought about by realizations of the exogenous growth process. This section examines how V and key economic variables change over time. We argue that the economy will tend toward higher spot prices, higher open interest, and a lower futures risk premium.

To illustrate the range of possible outcomes in our economy, we have thus far allowed for four widely dispersed values of V. Rather than consider the possible paths of our economy from each of four starting points, we select one initial V_0 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 Hamilton and Wu (2011) suggests that futures risk premia were larger in the past. These facts recommend $V_0 = 0.05$, a state in which most economic wealth belongs to agent 1, as a reasonable starting value for V in our economy. Figure 11 shows how the probability density of V evolves, conditional on our chosen starting value. Over time, the possible values of V become dispersed, allowing for a good deal of variation over time in the economic variables driven by V. The second feature is evidence of drift; V tends to increase over time. Although lower values remain possible, after 50 years V is likely to be above 0.1, and after 100 years it is extremely unlikely that V will be near its small initial value of 0.05. However V is unlikely to take values much larger than 0.4, even after 100 years. A typical 100 year path through the economy involves wide variation in V, spot prices, the term structure of bonds and futures, and related risk premia. There is a strong tendency for agent 2's wealth to increase, but neither agent will have a dominant position in the terminal period.

To directly relate changes in V over time to economic outcomes, we turn to the mean path of the economy. Figure 12 shows the average 50-year path of the economy, computed using Monte Carlo simulation. As before, we choose $V_0 = 0.05$. The initial growth state, s_0, is chosen from the stationary distribution. In the top panel, we see that V is expected to increase from 0.05 to more than 0.12 over 50 years. Interestingly the spot price, in the second panel, shows a much larger expected increase than would be implied directly by the mean change in V: it increases by more than 2/3, from around 32 to more than 55. If we had statically adjusted V_0 to a value of 0.12, the corresponding spot price at $t = 0$ would only be around 35. The surprisingly large expected increase in the spot price is the result of the joint distribution of x, y, s, and V, of which the spot price is a nonlinear function. Despite the expectation that x and y grow at the same rate, the spot price is expected to increase dramatically.

The third panel of Figure 12 shows the expected futures risk premium, which is quite volatile, and therefore leads to somewhat noisy estimate even after 10000 simulated paths. However there is a clear downward trend, with the premium decreasing from over 2%

initially to less than 1.5%. This occurs despite a likely shift in wealth towards agent 2, for whom oil is a more important commodity, and an increase in the spot price. The decreasing risk premium is also accompanied by a large increase in open interest in oil futures markets, shown in the fourth panel. Although analysts confronted with these results might be tempted to reason that increasing open interest and a decreasing risk premium represented evidence of "increased demand from buyers of futures pushing down the risk premium," the connection is illusory. As discussed earlier in the analysis of Figure 10, the relationship between the open interest and other variables - including risk premia - is nonmonotonic. However, given the initial state of our economy, increasing open interest is very likely to coincide with a decreasing futures risk premium. In fact the increase in open interest is quite dramatic, with a four-fold rise expected over 50 years. Some of this change results from growth in the economy, as the size of each agent's portfolio, and hence of open interest, grows with aggregate wealth. To isolate the impact of changes in V (or the wealth *distribution*) on open interest, the final panel of Figure 12 presents the average open interest over time as a fraction of aggregate wealth. Even with this normalization, open interest is expected to more than double over 50 years.

5 Conclusions

There is, of course, much to do. Our model has abstracted from the complex dynamics of oil exploration, development, storage, and refinement. We have also abstracted from all the production decisions and technological innovations surrounding oil consumption. Instead we have focussed on heterogenous exposure to oil risk as an important source that drives the complicated oil-risk premium dynamics 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 as well as capturing salient properties of financial returns in general. A nice feature of this set up is that we can look directly at optimal positions in futures markets to also note that rising open-interest is a natural consequence of the risk sharing outcome.

References

ANDERSON, E. W. (2005): "The dynamics of risk-sensitive allocations," *Journal of Economic Theory*, 125(2), 93 – 150.

BACKUS, D. K., B. R. ROUTLEDGE, AND S. E. ZIN (2005): "Exotic Preferences for Macroeconomists," in *NBER Macroeconomics Annual 2004*, ed. by M. Gertler, and K. Rogoff, vol. 19, pp. 319–391. MIT Press, Cambridge, MA.

——— (2008): "Recursive Risk Sharing: Microfoundations for Representative-Agent Asset Pricing," Carnegie Mellon Working Paper.

——— (2009): "Who Holds Risky Assets?," Carnegie Mellon Working Paper.

BANSAL, R., AND A. YARON (2005): "Risks For the Long Run: A Potential Resolution of Asset Pricing Puzzles," *Journal of Finance (forthcoming)*, Wharton Working Paper.

BARRO, R. J. (2009): "Rare Disasters, Asset Prices, and Welfare Costs," *The American Economic Review*, 99(1), pp. 243–264.

BASAK, S. (2005): "Asset pricing with heterogeneous beliefs," *Journal of Banking and Finance*, 29(11), 2849–2881.

BASAK, S., AND D. CUOCO (1998): "An Equilibrium Model with Restricted Stock Market Participation," *The Review of Financial Studies*, 11(2), pp. 309–341.

CAMPBELL, J. Y., AND J. H. COCHRANE (1999): "By Force of Habit: A Consumption-Based Explanation of Aggregate Stock Market Behavior," *Journal of Political Economy*, 107, 205–251.

CARLSON, M., Z. KHOKHER, AND S. TITMAN (2007): "Equilibrium Exhaustible Resource Price Dynamics," *Journal of Finance*, 62(4), 1663–1703.

CASASSUS, J., AND P. COLLIN-DUFRESNE (2005): "Stochastic Convenience Yield Implied from Commodity Futures and Interest Rates," *Journal of Finance*, 60(5), 2283–2331.

CASASSUS, J., P. COLLIN-DUFRESNE, AND B. ROUTLEDGE (2007): "Equilibrium Commodity Prices with Irreversible Investment and Non-linear Technologies," Hass Working Paper.

COCHRANE, J. H., AND M. PIAZZESI (2005): "Bond Risk Premia," *The American Economic Review*, 95(1), pp. 138–160.

COLACITO, R., AND M. M. CROCE (2011): "Risks for the Long-Run and the Real Exchange Rate," *Journal of Political Economy*, 119, 153–181.

——— (2012): "International Asset Pricing with Recursive Preferences," *Journal of Finance*, forthcoming.

DUMAS, B., A. KURSHEV, AND R. UPPAL (2009): "Equilibrium Portfolio Strategies in the Presence of Sentiment Risk and Excess Volatility," *The Journal of Finance*, 64(2), 579–629.

EPSTEIN, L. G., AND S. E. ZIN (1989): "Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Framework," *Econometrica*, 57(4), 937–969.

FAMA, E. F., AND K. R. FRENCH (1987): "Commodity Futures Prices: Some Evivedence on Forecast Power, Premiums and the Theory of Storage," *Journal of Business*, 60(1), 55–73.

GÂRLEANU, N., AND S. PANAGEAS (2010): "Young, old, conservative and bold: The implications of heterogeneity and finite lives for asset pricing," *Unpublished Manuscript, Haas School of Business, University of California at Berkeley.*

Hamilton, J. (2009): “Causes and Consequences of the Oil Shock of 2007-08,” *NBER Working Paper.*

——— (2011): “Historical oil shocks,” Discussion paper, National Bureau of Economic Research.

Hamilton, J., and J. Wu (2011): “Risk Premia in Crude Oil Futures Prices,” Working Paper.

Hamilton, J. D. (2008): “Oil and the Macroeconomy,” in *The New Palgrave Dictionary of Economics*, ed. by S. N. Durlauf, and L. E. Blume. Palgrave Macmillan.

Hansen, L., J. Heaton, and N. Li (2008): “Consumption Strikes Back? Measuring Long-Run Risk,” *Journal of Political Economy*, 116(2), 260–302.

Kan, R. (1995): “Structure of Pareto Optima When Agents Have Stochastic Recursive Preferences,” *Journal of Economic Theory*, 66(2), 626–631.

Keynes, J. M. (1936): *The general theory of employment, interest and money.* Macmillan and Co., London.

Kihlstrom, R. E., and L. J. Mirman (1974): “Risk aversion with many commodities,” *Journal of Economic Theory*, 8(3), 361–388.

Kogan, L., D. Livdan, and A. Yaron (2009): “Futures Prices in a Production Economy with Investment Constraints,” *Journal of Finance*, 64, 13451375, The Wharton School, University of Pennsylvania, Working Paper.

Kreps, D. M., and E. L. Porteus (1978): “Temporal Resolution of Uncertainty,” *Econometrica*, 46, 185–200.

Lucas, R. E., and N. L. Stokey (1984): “Optimal Growth with Many Consumers,” *Journal of Economic Theory*, 32(1), 139–171.

Lucas, Jr., R. E. (1978): “Asset Prices in an Exchange Economy,” *Econometrica*, 46, 1426–14.

Ready, R. (2010): “Oil Prices and Long-Run Risk,” *SSRN eLibrary.*

Routledge, B. R., D. J. Seppi, and C. S. Spatt (2000): “Equilibrium Forward Curves for Commodities,” *Journal of Finance*, 60, 1297–1338.

——— (2001): “The Spark Spread: An Equilibrium Model of Cross-Commodity Price Relationships in Electricity,” Carnegie Mellon University Working Paper.

Routledge, B. R., and S. E. Zin (2010): “Generalized Disappointment Aversion and Asset Prices,” *Journal of Finance*, 65.

Schwartz, E. S. (1997): “The Stochastic Behavior of Commodity Prices: Implications for Valuation and Hedging,” *Journal of Finance*, 52, 923–973.

Wright, B. D. (2011): “The Economics of Grain Price Volatility,” *Applied Economic Perspectives and Policy*, 33, 32–58.

Yan, H. (2010): “Natural selection in financial markets: Does it work?,” *Management Science*, 54(11), 1935.

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.
ALL	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
	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.
Slope +	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
	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.
Slope -	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
	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.
	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
ALL	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.
	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
Slope +	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.
	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
Slope -	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	.0187209	1.720852	1.069224	-.2222553	.0427325	266
6	.0452831	2.967716	.8309402	.7552679	.0501505	263
12	.0879378	4.721614	.9119217	.5278193	.1048359	257
18	.1394924	6.779568	.9735516	.1796232	.1493467	251
24	.1836817	7.941727	.8189865	1.215011	.1152461	234
36	.2582488	8.407877	.2833659	4.402931	.0173169	174
48	.4259285	9.551207	.2318959	3.420123	.0122456	88
60	.6186527	13.91229	.0719869	4.764505	.0020346	69

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

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

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

$$\begin{bmatrix} s \in S \\ f(s) \\ g(s) \end{bmatrix} = \begin{bmatrix} 1 & 2 & 3 & 4 \\ 0.99 & 1.03 & 1.05 & 0.93 \\ 0.90 & 1.04 & 1.06 & 1.07 \end{bmatrix}$$

$$\pi = \begin{bmatrix} 0.80 & 0.10 & 0.05 & 0.05 \\ 0.05 & 0.85 & 0.05 & 0.05 \\ 0.05 & 0.18 & 0.72 & 0.05 \\ 0.05 & 0.05 & 0.63 & 0.27 \end{bmatrix}$$

$$\bar{\pi} = \begin{bmatrix} 0.20 & 0.48 & 0.26 & 0.06 \end{bmatrix}$$

Growth process characteristics. The first matrix shows possible growth outcomes for the numeraire ($f(s)$) and oil ($g(s)$) for each growth state (s). In matrix π, entry $\pi_{i,j}$ is the probability of transitioning from current growth state i to next period state j. The stationary (long-run) probability of being in a given growth state is shown in $\bar{\pi}$.

Table 5: **Parameters**

Parameter	Value	Description
α_1	-20	risk aversion, agent 1
α_2	-12.6	risk aversion, agent 2
ρ_1	-1.12	intertemporal substitution, agent 1
ρ_2	0.754	intertemporal substitution, agent 2
γ_1	0.03	oil preference, agent 1
γ_2	0.06	oil preference, agent 2
β	0.96	impatience, agents 1,2
x_0	1000	initial aggregate consumption, numeraire
y_0	1	initial aggregate consumption, oil

Preference parameters and initial aggregate consumption levels used in numerical examples.

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

	Horizon (yrs)	$V = 0.05$	$V = 0.35$	$V = 0.65$	$V = 0.95$
All	1	2.128	0.336	-0.379	-0.704
	2	2.975	0.348	-1.154	-1.882
	3	2.923	-0.208	-2.065	-2.963
	4	2.665	-0.789	-2.820	-3.790
	5	2.401	-1.256	-3.375	-4.376
	6	2.185	-1.595	-3.760	-4.774
	7	2.021	-1.827	-4.018	-5.037
Slope +	1	-0.136	-1.232	-0.686	-0.962
	2	-0.400	-1.958	-1.650	-2.283
	3	-0.751	-2.577	-2.676	-3.446
	4	-1.049	-3.034	-3.506	-4.323
	5	-1.275	-3.355	-4.110	-4.939
	6	-1.442	-3.576	-4.528	-5.357
	7	-1.563	-3.725	-4.807	-5.633
Slope -	1	2.683	0.721	4.129	3.085
	2	3.804	0.914	6.120	3.995
	3	3.825	0.373	6.906	4.117
	4	3.576	-0.238	7.242	4.018
	5	3.304	-0.740	7.403	3.886
	6	3.075	-1.109	7.494	3.775
	7	2.901	-1.361	7.553	3.694

Annual holding 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 V. Increasing values of V 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 2-year futures price. Although the magnitude of the effect varies, a negative slope always implies a higher risk premium on holding the contract. For the two smallest values of V, the slope is negative 80% of the time. For the two largest V, the slope is negative 6% of the time. The stark contrast in these percentages is an artifact of our simple, 4-state growth process.

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

	Horizon (yrs)	$V = 0.05$	$V = 0.35$	$V = 0.65$	$V = 0.95$
All	1	0.000	0.000	0.000	0.000
	2	-1.078	-0.166	0.009	0.063
	3	-1.656	-0.244	0.041	0.128
	4	-1.997	-0.279	0.082	0.190
	5	-2.218	-0.291	0.124	0.248
	6	-2.369	-0.294	0.165	0.299
	7	-2.480	-0.291	0.202	0.344
Slope +	1	0.000	0.000	0.000	0.000
	2	-0.279	-0.196	-0.019	0.039
	3	-0.520	-0.365	-0.008	0.086
	4	-0.727	-0.512	0.017	0.135
	5	-0.907	-0.640	0.047	0.182
	6	-1.062	-0.752	0.078	0.226
	7	-1.198	-0.849	0.108	0.264
Slope -	1	0.000	0.000	0.000	0.000
	2	-1.274	-0.159	0.428	0.416
	3	-1.934	-0.214	0.765	0.741
	4	-2.309	-0.221	1.035	0.998
	5	-2.539	-0.206	1.254	1.205
	6	-2.690	-0.181	1.435	1.373
	7	-2.795	-0.154	1.585	1.512

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 V. Increasing values of V 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 2-year futures price. The connection between the slope of the oil futures curve and the bond risk premium depends upon V. For $V = 0.35$, the risk premium is somewhat higher given a positive slope. All other values of V imply higher excess returns given a negative slope. For the two smallest values of V, the slope is negative 80% of the time. For the two largest V, the slope is negative 6% of the time. The stark contrast in these percentages is an artifact of our simple, 4-state growth process.

Table 8: **Predictive regressions (model)**

V	Horizon (yrs)	a	b	R^2
0.05	1	0.009	0.600	0.225
	2	0.018	0.526	0.207
	3	0.025	0.480	0.181
0.35	1	-0.001	0.640	0.225
	2	-0.006	0.585	0.204
	3	-0.016	0.546	0.176
0.65	1	-0.005	0.688	0.264
	2	-0.017	0.646	0.240
	3	-0.036	0.623	0.204
0.95	1	-0.007	0.741	0.304
	2	-0.024	0.727	0.281
	3	-0.051	0.728	0.251

On simulated data from our model, regress:

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

Figure 1: Time Series Oil Prices

(a) Real Spot Price (1970)

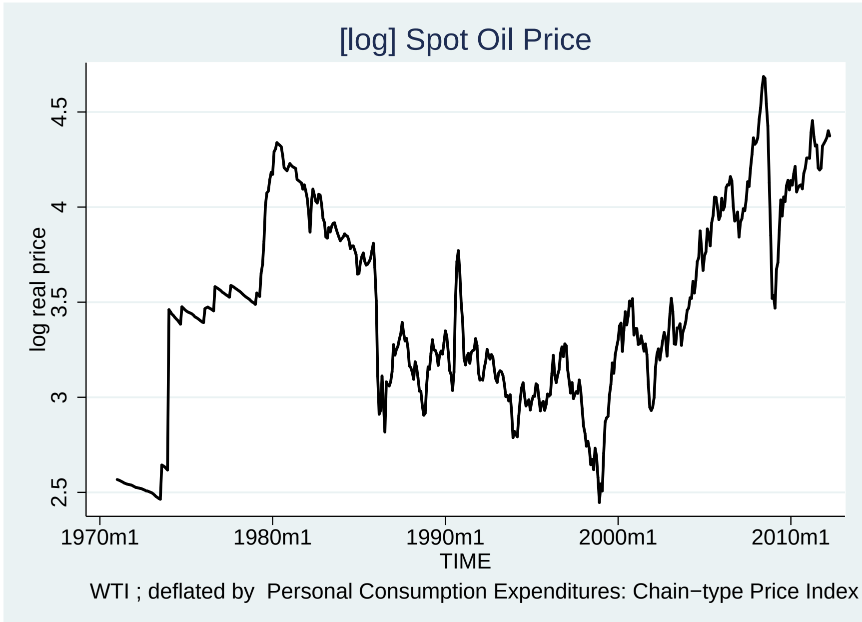

(b) Futures Prices (from 1990)

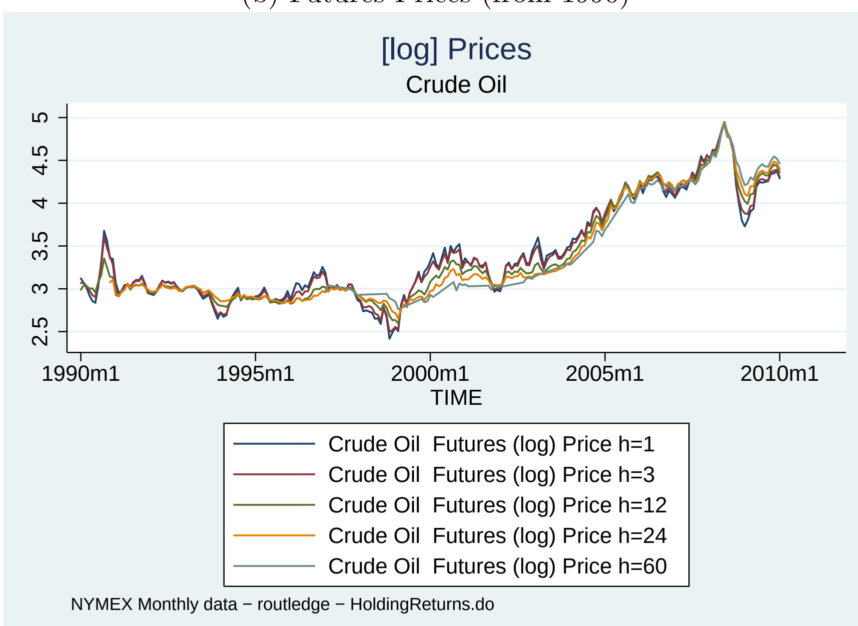

Spot Price is from St. Louis Fed. Deflated by chain-weighted price index (Indexed at 2001/01). Futures prices are from nominal

Figure 2: Term Structure of Holding Period Returns

a. By Horizon

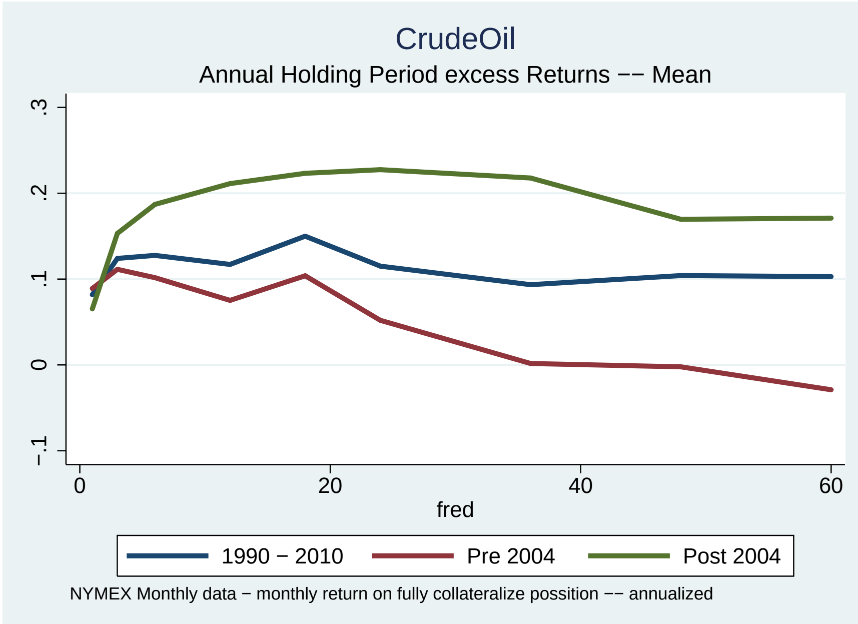

b. Time Series

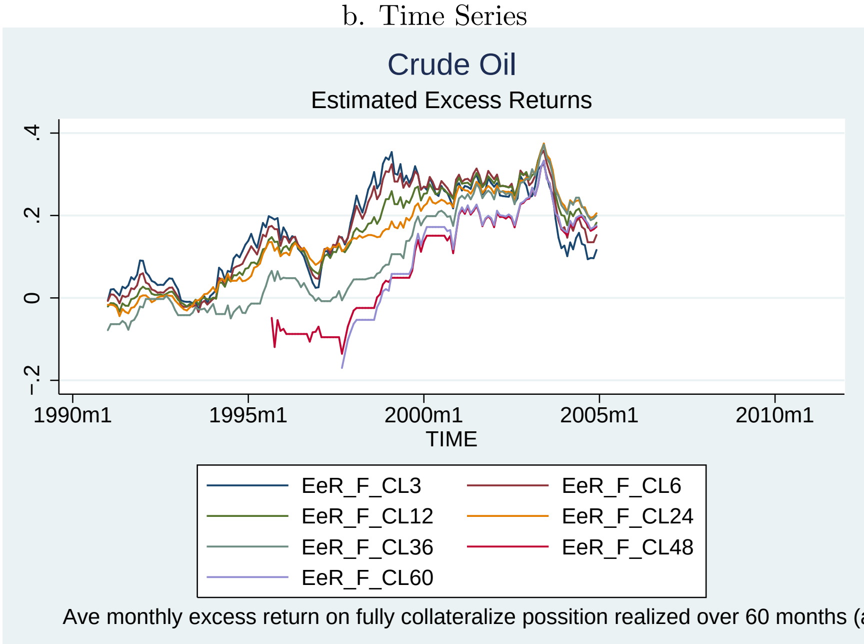

One month expected holding excess returns on fully-collaterialized futures position (annualized). (a) The excess return is calculated as the mean excess return realized in sample. (b) The excess return is calculated as the mean realized excess return in the subsequent 60 months

Figure 3: Wealth and Consumption Shares

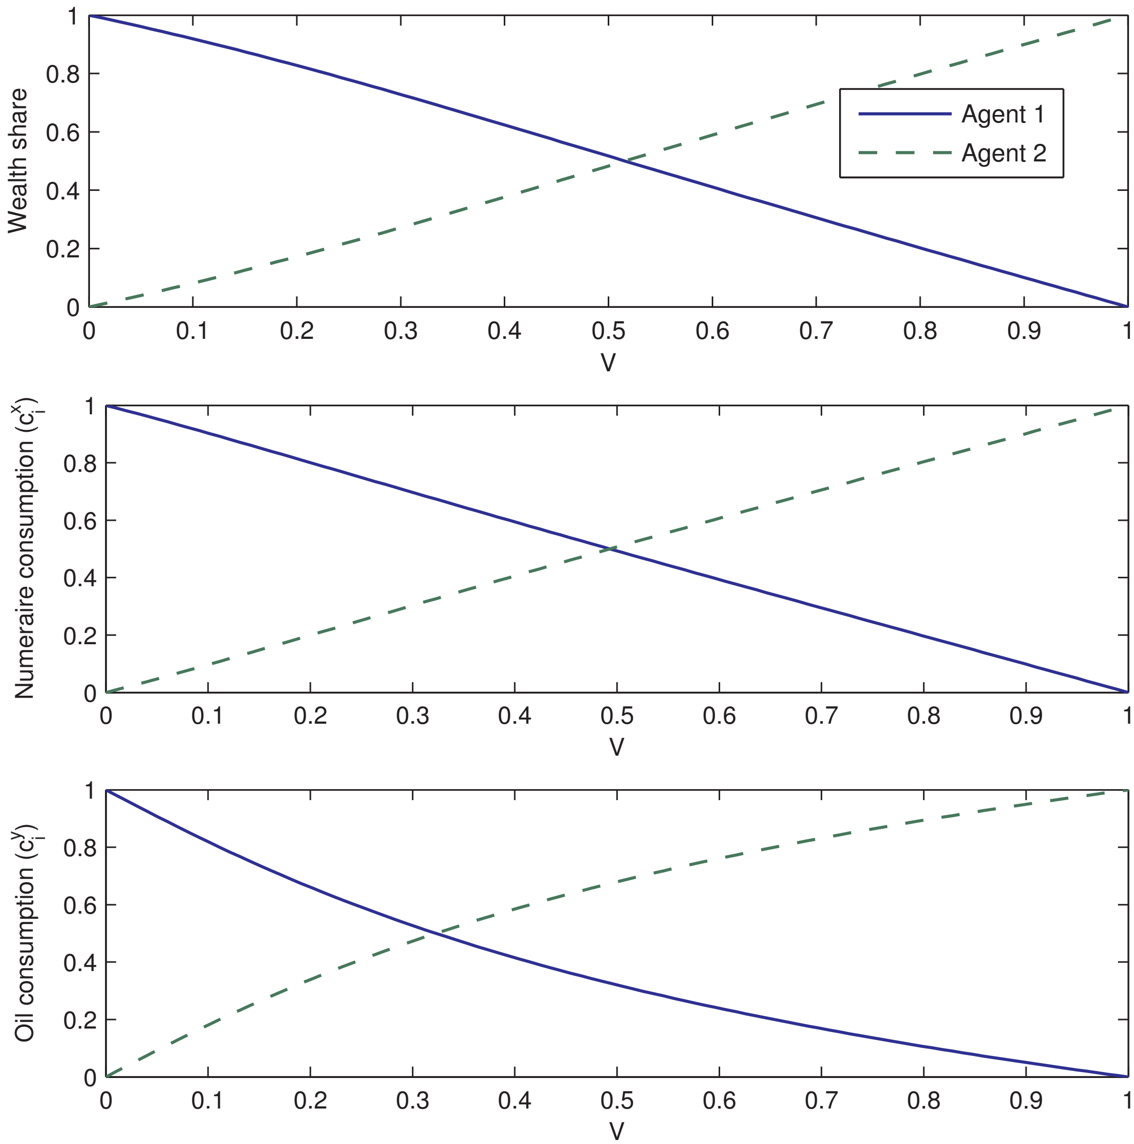

Wealth and consumption shares of each agent, relative to state variable V. Results are averages over the stationary distribution of growth states (s). However results conditional on a particular s are similar to the mean. The mapping from V to the wealth share of agent 2 is almost linear and 1-to-1. However consumption shares, particularly the distribution of oil, are very nonlinear in V, reflecting the different preferences for oil consumption. Therefore V can safely be interpreted as a measure of wealth distribution, but not consumption distribution.

Figure 4: Oil futures prices

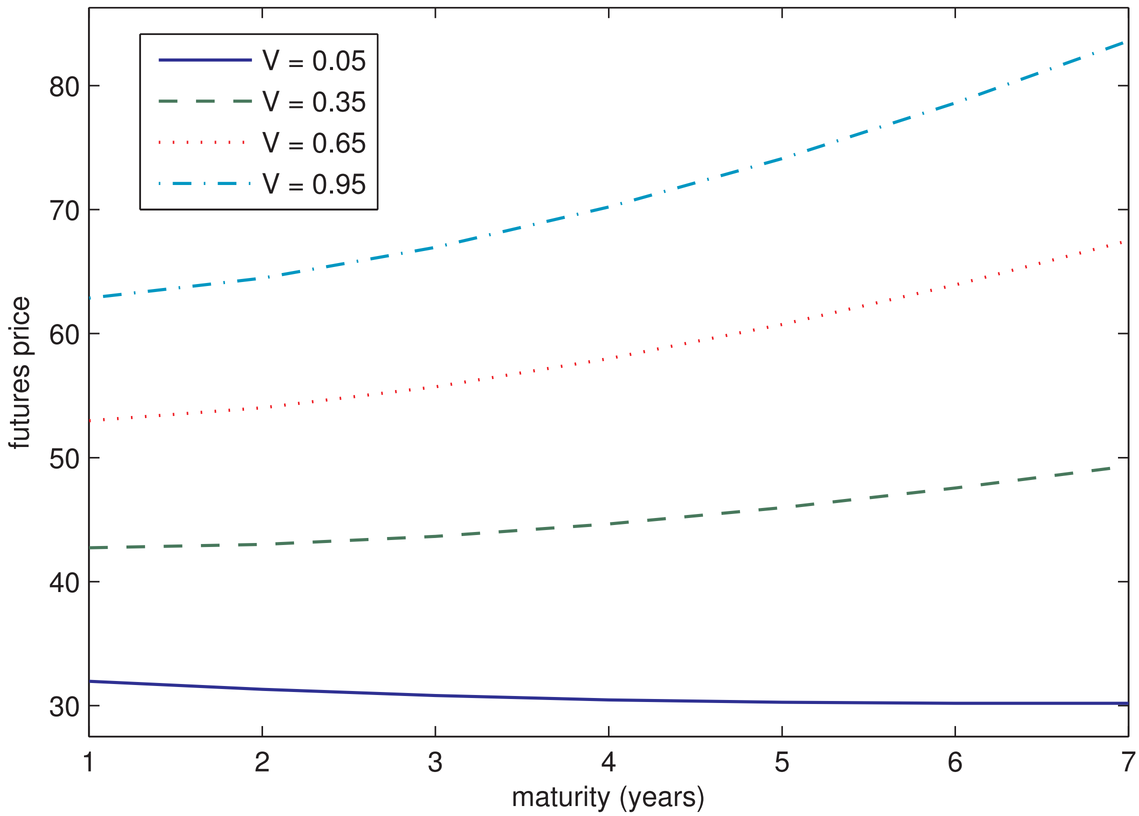

Futures prices, averaged over the stationary distribution of growth states. Higher V correspond to much higher average price levels. In addition, the slope changes: futures are downward sloping on average for small V, but upward sloping on average for large V. Although we show averages over growth state s to emphasize the role of V (i.e., wealth distribution), the growth state has a large (but relatively transitive) impact. For example, conditional on state $s = 1$, all futures curves are downward sloping, whereas conditional on state $s = 4$, all curves are upward sloping.

Figure 5: Oil futures risk premia

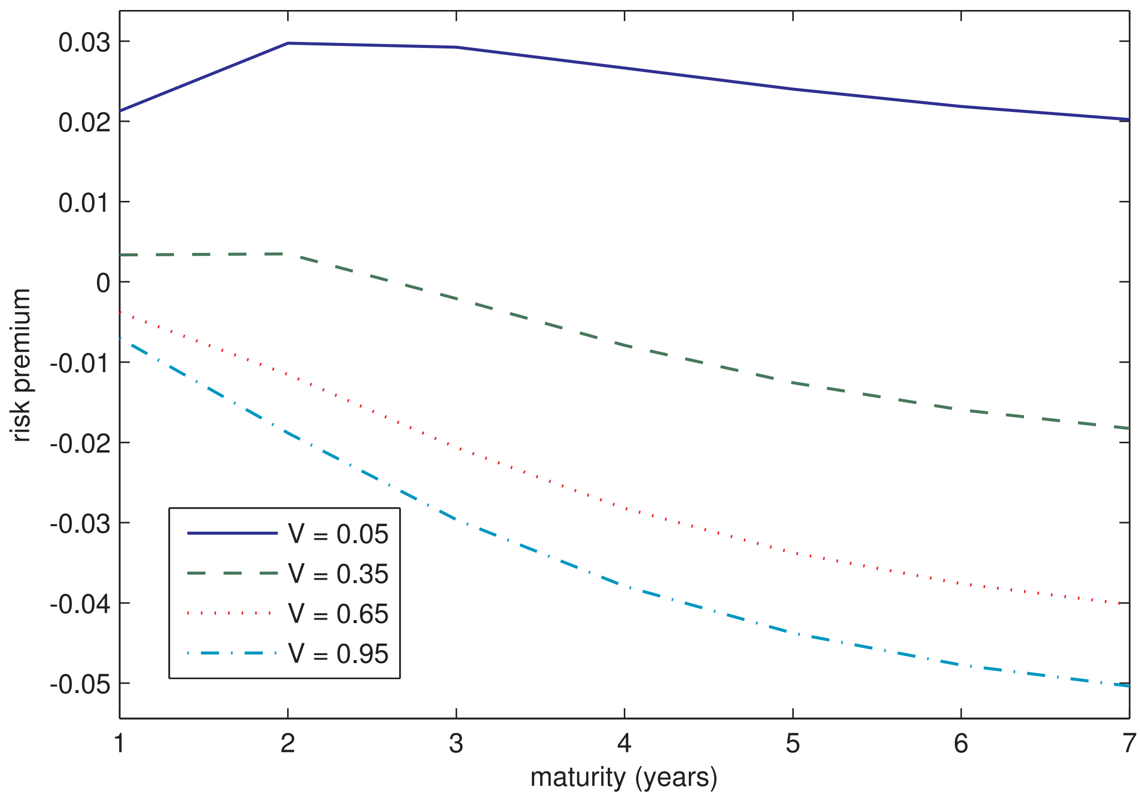

Annual risk premia (excess returns) on oil futures, averaged over the stationary distribution of growth states. V influences the level and slope of risk premia. Low V (low agent 2 wealth) produce a positive and hump-shaped term structure of risk premia, whereas high V (high agent 2 wealth) produce downward sloping and negative risk premia. Although we emphasize the role of V by averaging over growth states s, the growth state is important for risk premia. For example, conditional on state $s = 1$, risk premia are positive for all V, whereas for $s = 4$, risk premia are negative for all V.

Figure 6: Futures risk premium (2-year contract)

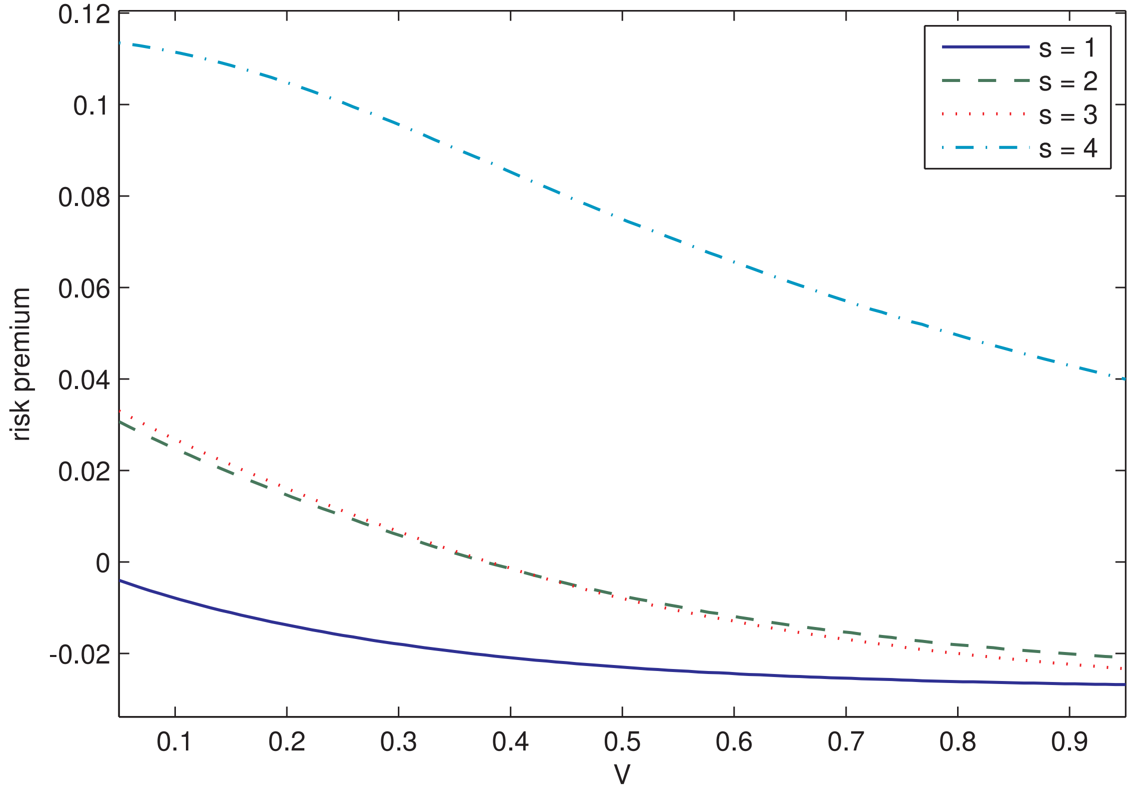

Annual risk premium on 2-year oil futures (shown for each growth state). Note the average across all growth states is approximately equal to that of state 2 and 3. The plot provides intuition as to the magnitude and direction of changes in risk premia due to V, which evolves endogenously. For our calibration, the average risk premium is monotonically decreasing in V.

Figure 7: Zero-coupon bond rates

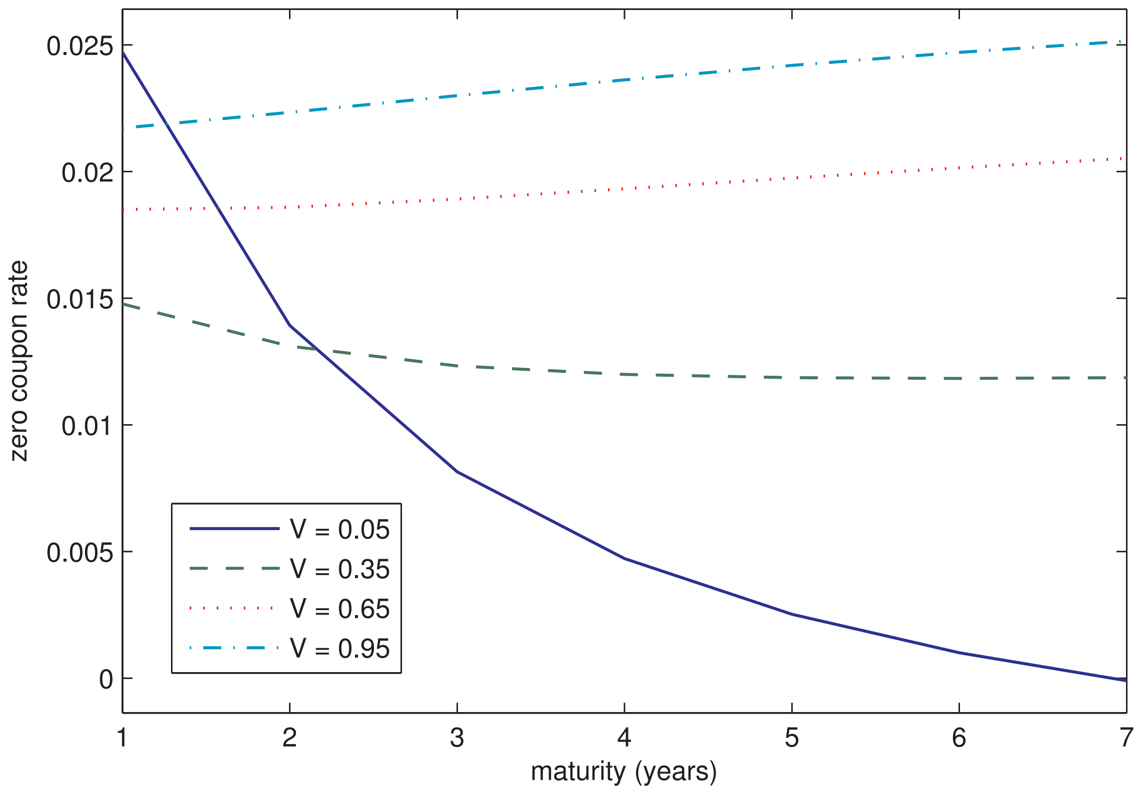

Annual interest rates on zero coupon bonds, averaged over the stationary distribution of growth states. Depending on V, the term structure may be downward sloping (low V) or upward sloping (high V). In addition, the risk-free rate (rate on the 1-year bond) is non-monotonic in V. The interaction of these effects leads to bond risk premia that may be positive or negative, and increasing or decreasing with maturity.

Figure 8: Equity premium

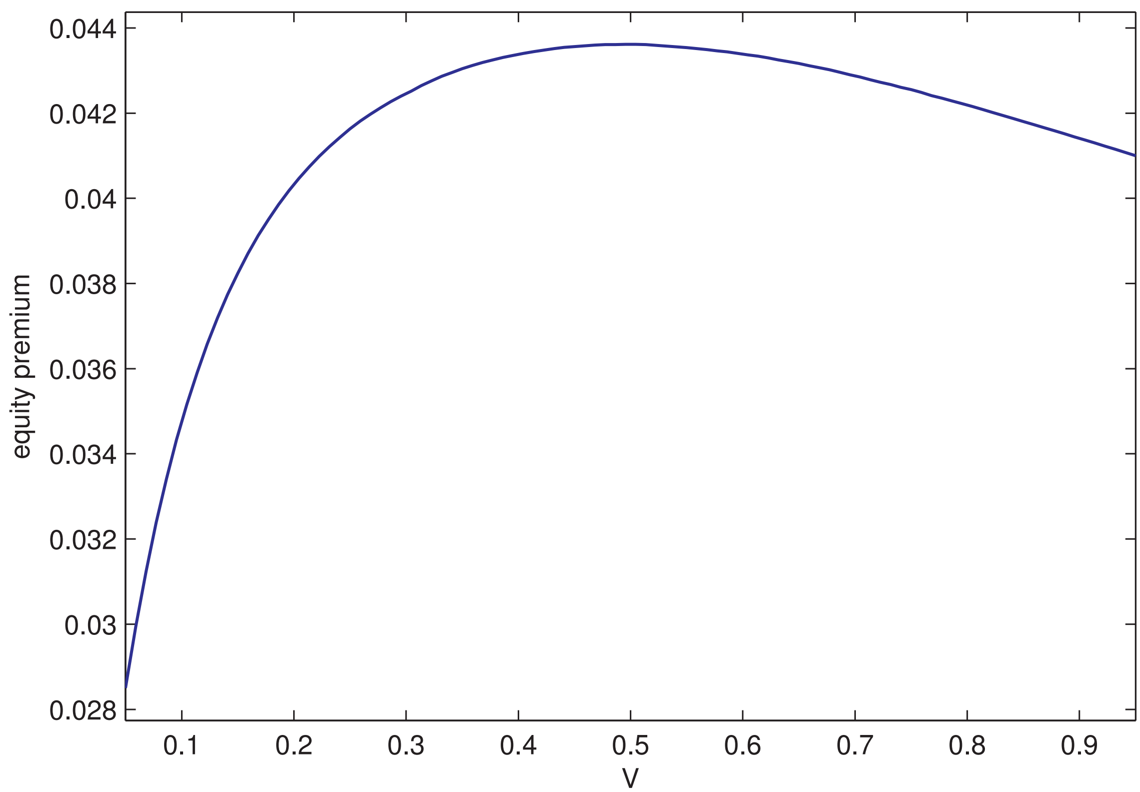

Annual equity premium, averaged over the stationary distribution of growth states. The equity premium is non-monotonic in V. Strikingly, the maximum equity premium occurs when the two agents have roughly equal wealth shares - around $V = 0.5$ - rather than when one agent dominates the economy. The equity premium may be higher in the multi-agent economy than in an economy populated by either agent 1 or agent 2 alone.

Figure 9: Portfolios

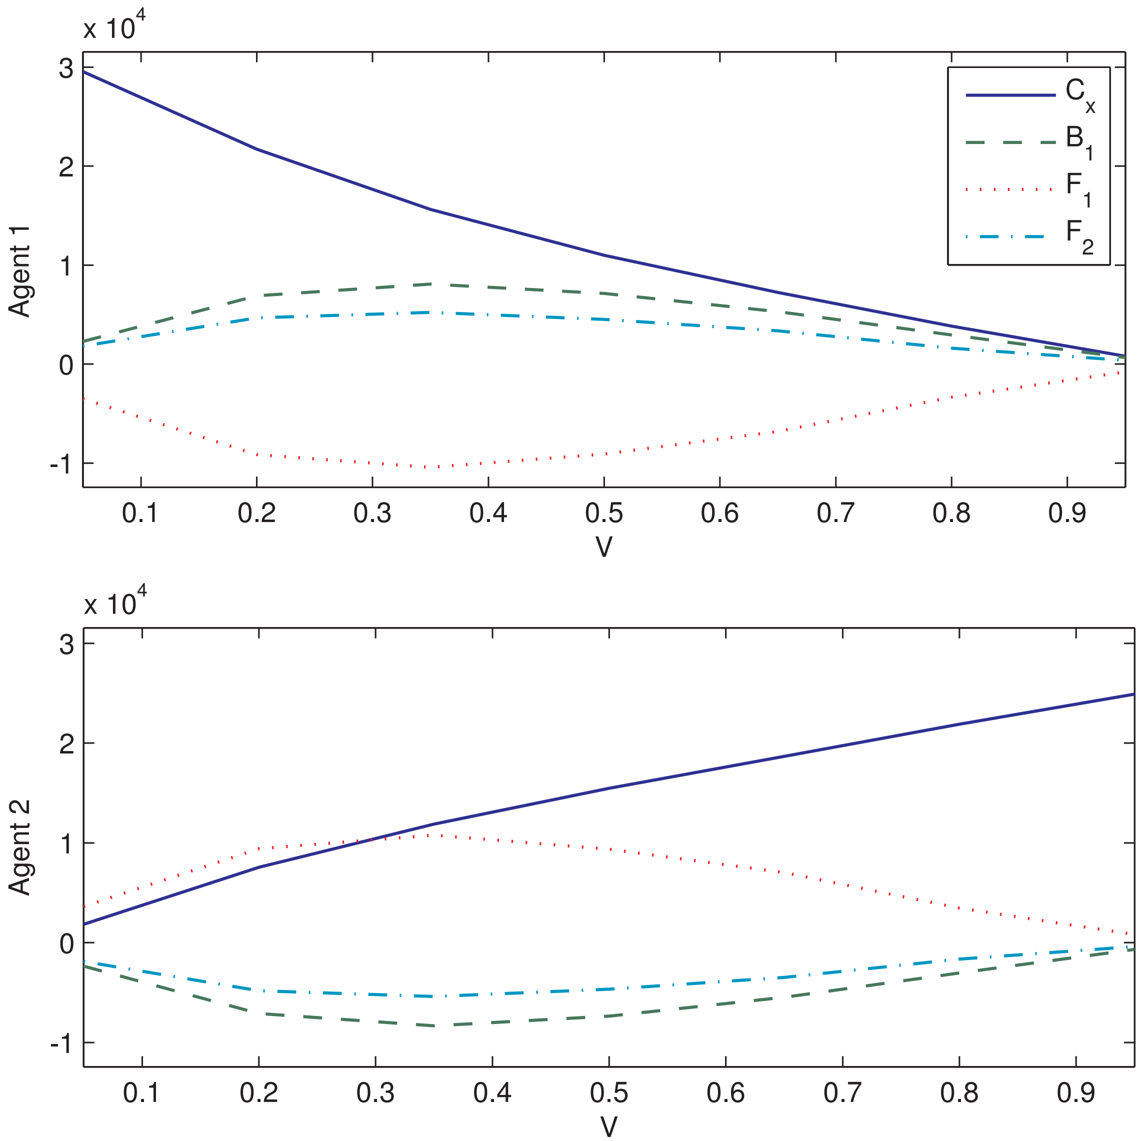

Portfolios for each agent, in terms of numeraire value of investment in each asset, versus V. Markets are completed using the numeraire stock, one-period bonds, and fully collateralized 1 and 2 year futures contracts. Plots are averages over growth states using the stationary distribution; however results for each growth state are very similar. Holdings of the x-stock are monotonic in V, approximately following changes in wealth. Holdings of bonds and futures are non-monotonic in V, more reflective of the level of trade between agents for different wealth distributions. Positions in bonds and futures for agent 1 are approximately the negative of those for agent 2.

Figure 10: Open interest in oil futures and spot prices

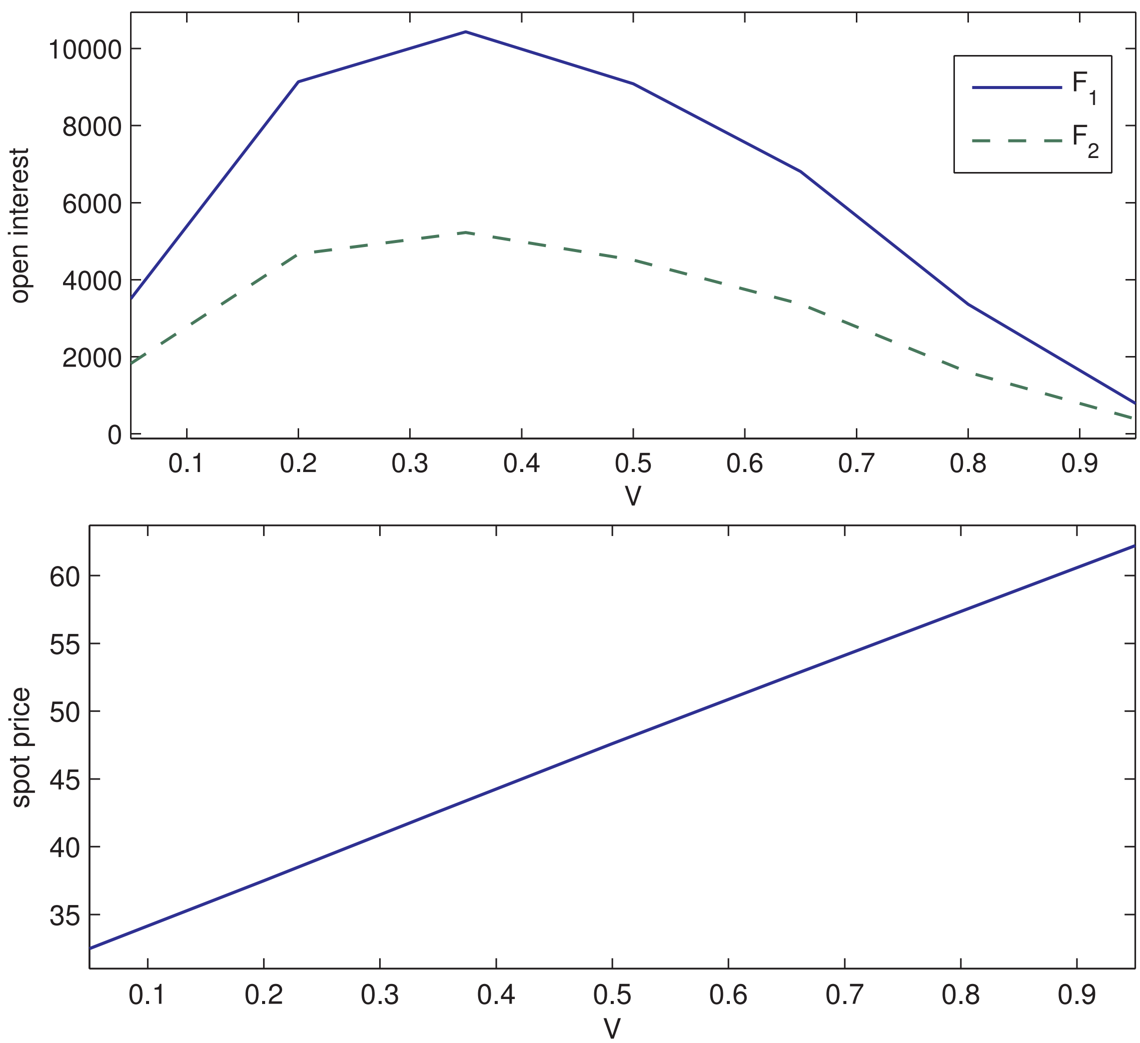

The top panel shows open interest in 1 and 2 year oil futures contracts, expressed as the numeraire value of the contracts. The bottom panel shows spot price of oil for various V. Results are averaged over the stationary distribution of growth states (s). However results conditional on a particular s are similar to the mean. Open interest varies widely and nonmonotonically with V, from nearly 0 to orders of magnitude more than the total value of oil consumed in the economy. Although the spot price of oil varies also, it does so monotonically with V. Therefore an increase in open interest has no clear implication for the spot price of oil, contrary to claims in popular news outlets.

Figure 11: Probability density of V

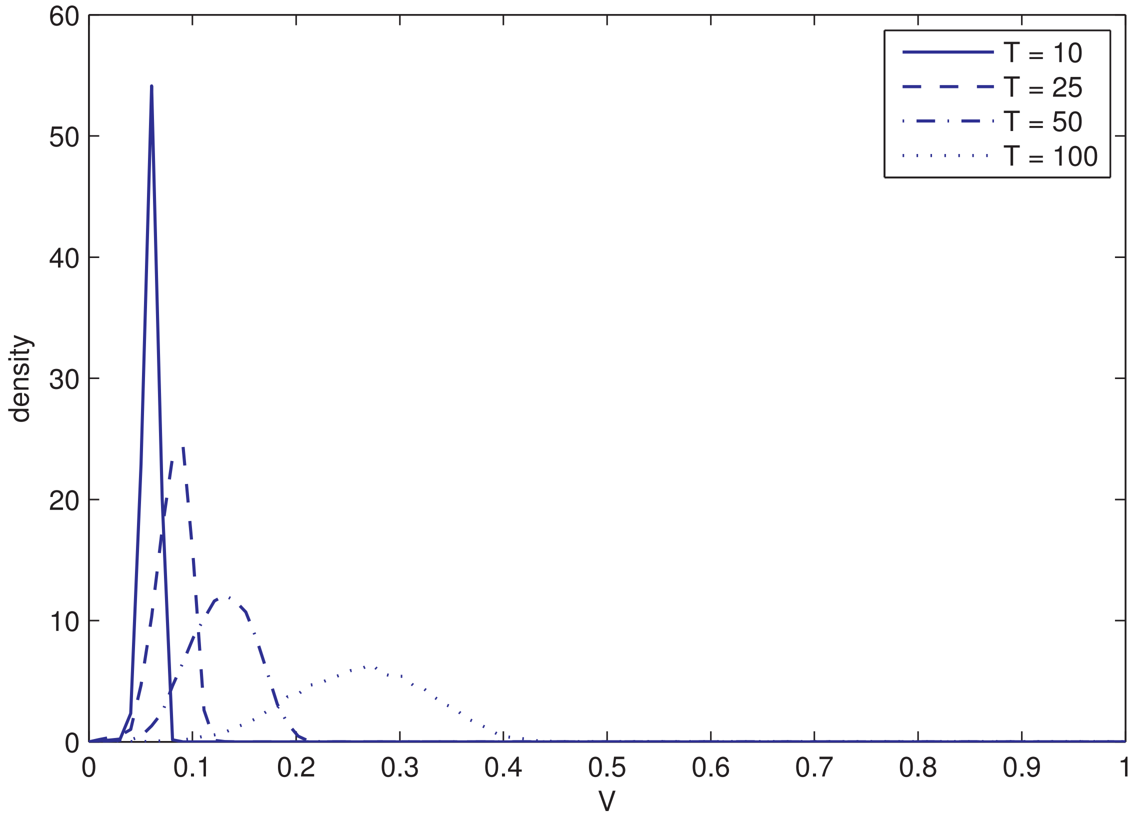

The probability "density" of V is plotted at increasing horizons, of 10, 25, 50, and 100 years, conditional on initial value $V_0 = 0.05$. Although V has a discrete distribution conditional on V_0, we plot a continuous analog to the probability mass function for ease of visualization. The resulting plot has two main features: (1) V exhibits an upward drift, such that values $V_t > V_0$ become very likely at longer horizons and (2) the probability mass becomes more dispersed, such that the range of probable values for V_t becomes much wider for larger t. Results are computed using Monte Carlo simulation with 10000 paths.

Figure 12: Evolution over time

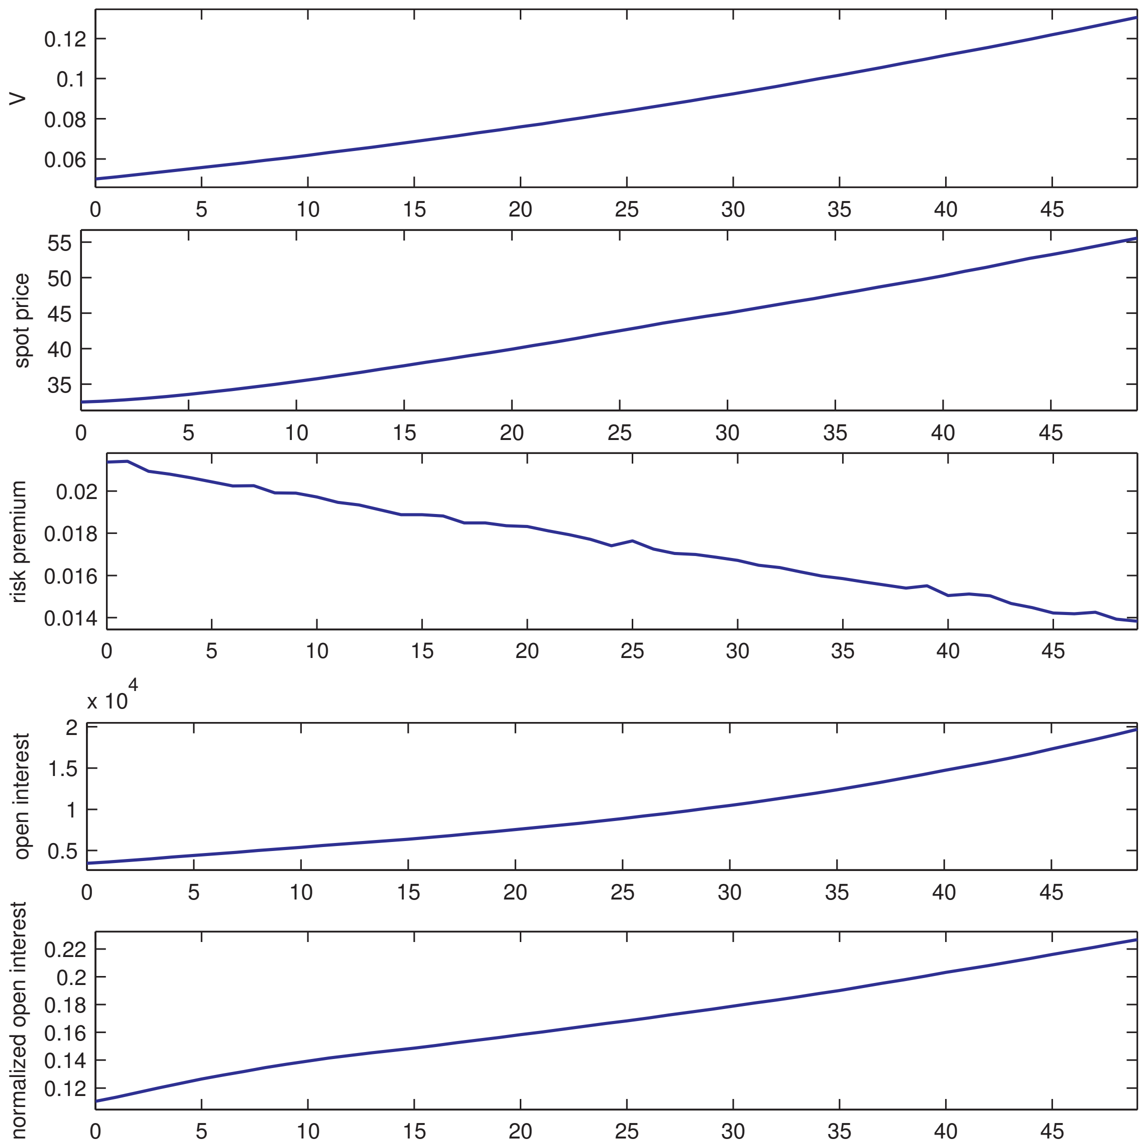

The plots illustrate the average path of the economy over a 50-year period, conditional initial $V_0 = 0.05$. The initial growth state is selected according to the stationary distribution. From top to bottom, the panels show V (indicative of agent 2's wealth share), the oil spot price, risk premium on one-year oil futures, open interest on 1-year futures, and open interest normalized by aggregate wealth. Over time, the economy is likely to exhibit a rising spot price, increasing open interest, and decreasing futures risk premium. Results are computed using Monte Carlo simulation with 10000 paths.