

Modeling Financial Markets with Heterogeneous Agents

Doctoral Dissertation

Steven David Baker

Tepper School of Business
Carnegie Mellon University

April, 2013

Submitted in Partial Fulfillment of the Requirements for the Degree of
Doctor of Philosophy in Financial Economics

Contents

1	The Financialization of Storable Commodities	5
1.1	Introduction	5
1.2	The Model	9
1.2.1	Markets	10
1.2.2	Producers	11
1.2.3	Dealers	12
1.2.4	Households	14
1.3	Equilibrium	15
1.3.1	Characterization	17
1.3.2	Prices and Risk Premia	19
1.4	Data and Calibration	21
1.5	Results	24
1.5.1	Before Financialization	24
1.5.2	Financialization	26

1.5.3	Welfare	28
1.6	Conclusion	29
1.7	Additional Discussion and Extensions	30
1.7.1	Open Interest in Crude Oil Derivatives, and the Spot Price	30
1.7.2	Additional Shocks	31
1.7.3	Intermediate Production	33
1.7.4	Active Trade in Multi-period Futures Contracts	35
1.7.5	Characterization of Risk-aversion	38
1.7.6	Comments on General Equilibrium	39
1.8	References	41
1.9	Tables	44
1.10	Figures	49
2	The Price of Oil Risk (with Bryan R. Routledge)	69
2.1	Introduction	69
2.2	Facts	73
2.3	Exchange Economy - Two Goods and Two Agents	76
2.3.1	Single Agent, Two Goods	77
2.3.2	Two Agents, Two Goods	78
2.3.3	Financial Prices	81
2.3.4	Portfolios and open interest	84

2.4	Calibrated Numerical Example	85
2.4.1	Wealth and consumption	87
2.4.2	Prices and risk premia	88
2.4.3	Portfolios and open interest	90
2.4.4	Dynamics	92
2.5	Conclusions	93
2.6	Additional Discussion and Extensions	94
2.6.1	Risk Aversion with Recursive Preferences and Two Goods	94
2.6.2	Sources of Time-variation in the Risk Premium	95
2.6.3	Improving the Calibration with Transient Growth Shocks	96
2.7	References	98
2.8	Tables	101
2.9	Figures	108
3	Disagreement, Financial Markets, and the Real Economy	123
3.1	Introduction	123
3.2	The Model	127
3.2.1	Firms	127
3.2.2	Households	128
3.2.3	Markets	129
3.3	Equilibrium	130

3.3.1	Aggregate Saving and Capital Allocation	130
3.3.2	State Prices	131
3.3.3	Wealth and Budget Constraints	132
3.3.4	Asset Prices	133
3.3.5	Portfolios	134
3.4	An Economy with One Firm	135
3.4.1	Evolution of the Optimist's Consumption Share	137
3.4.2	Consumption	139
3.4.3	Asset Prices, Portfolios, and Perceptions	141
3.4.4	Long-run Impact of Overconsumption	145
3.5	An Economy with Two Firms	146
3.6	Conclusion	151
3.7	References	152
Appendix 3.A	Proofs	153
Appendix 3.B	Asymptotic Behavior of Optimist Consumption Share (ω)	154

Chapter 1

The Financialization of Storable Commodities

I construct a dynamic equilibrium model of storable commodities populated by producers, dealers, and households. When financial innovation allows households to trade in futures markets, they choose a long position that leads to lower equilibrium excess returns on futures, a more frequently upward-sloping futures curve, and higher volatility in futures and spot markets. The effect on spot price levels is modest, and extremely high spot prices only occur in conjunction with low inventories and poor productivity. Therefore the “financialization” of commodities may explain several recently observed changes in spot and futures market dynamics, but it cannot directly account for a large increase in spot prices.

1.1 Introduction

During the last decade, trade in commodity futures contracts has risen along with the popularity of commodity index investment.¹ New products, such as commodity exchange traded funds (ETFs), have made it practical for household consumers to trade commodity futures. Although empirical studies have documented a range of changes in spot and futures markets accompanying the financialization of commodities,

¹For example, a CFTC (2008) report notes “volume growth on futures markets has increased fivefold in the last decade.” See also Buyuksahin et al. (2011), Irwin, Sanders, Merrin, et al. (2009).

there is limited theoretical understanding of the dynamic linkages between spot and futures markets. Participation in futures markets alters the incentives of producers, storers, and consumers of the commodity, indirectly affecting the spot market. Commodity funds marketed to households typically represent a rolling long position in futures or similar derivatives. Risk-averse intermediaries or dealers implementing these funds can offset sales of futures via purchase and storage of the underlying commodity, which can effect spot prices. Or dealers can offset sales of futures to households by purchasing more futures from producers, presumably by offering better terms, i.e., a higher futures price. Because flows to commodity funds have been large and volatile, politicians and regulators wonder if financialization has caused observed changes in futures and spot markets, including unwelcome increases in volatility and price levels.

To address these questions, I develop a dynamic equilibrium model of commodity prices incorporating an active futures market, heterogeneous risk-averse participants, and storage. I analyze the effects of financialization by reducing the cost to household consumers of trading in the futures market, which is initially dominated by commercial producers and dealers. The experiment calibrates the model to approximate moments of prices and quantities *prior* to financialization, then adjusts only the household's transaction cost, to reflect observed retail trade. This procedure provides estimates of the magnitude of financialization's effects, in addition to explaining the economic channels driving these effects. Financialization accounts for an order of magnitude decrease in excess returns (from 2% to 0.2% per quarter) on the most commonly traded futures contracts, and a roughly 50% reduction in the frequency of a downward sloping futures curve (backwardation). Financialization also accounts for some of the increase in spot and futures price volatility: roughly 50% of the empirical increase when controlling for differences in mean prices during the pre (1990-2003) versus post (2004-2012) financialization periods. Financialization cannot be blamed, however, for an increase in mean spot price levels. This result obtains even though the model allows dealers to accumulate inventory to offset (or augment) their positions in the futures market. I argue that the increase in speculative storage - which occurs in the model - actually has a beneficial smoothing effect in the long run, resulting in a small improvement to household welfare after financialization.

Two main effects connect increased futures trade to spot market dynamics: amplification via household income, and smoothing via dealer inventory. As in the data, households choose a long position in futures once transaction costs are lowered, earning a risk-premium and hedging their consumption risk. Therefore the most direct effect of financialization is an income effect to households: they are insured against spikes

in the commodity spot price. This amplifies spot price volatility, as households have more to spend on the commodity precisely when it is scarce, and less to spend when it is abundant. If household demand for the commodity is quite inelastic, the insurance effect will be asymmetrical, and the mean spot price will increase. However a second, indirect effect of financialization works in the opposite direction, via inventory smoothing. The intermediary dealers who sell futures to households offset some of their futures sales through increased inventory accumulation. Because there is generally more inventory available to smooth supply disruptions, severe shortages - stockouts - are less likely after financialization than before. The reduction in stockouts reduces volatility and the mean spot price, offsetting the household income effect. Therefore a calibrated example is required to assess which of the two main effects dominates, and how they affect the risk premium, the slope of the futures curve, and household welfare.

Although the model is applicable to any storable commodity, I calibrate it to crude oil spot and futures markets, for several reasons. Retail investment in commodity derivatives is often done via index funds that span energy, metals, and agricultural commodities, but common reference indices such as the Goldman Sachs Commodity Index (GSCI) are tilted heavily towards oil. As of the year ended 2011, roughly 50% of the GSCI was crude oil, and energy commodities as a whole (including gasoline and heating oil derived from crude) comprised over 70 % of the index.² The large weight on oil is designed to reflect its significance to the global economy relative to other commodities. For example, oil prices are often used in forecasts of US GDP growth (Kilian and Vigfusson, 2012). Hamilton (2008) notes that “nine out of ten of the U.S. recessions since World War II were preceded by a spike up in oil prices,” an association that increases concern among regulators and the public about increased oil prices. Tang and Xiong (2011) document increasing correlation between oil futures and non-energy commodity futures concurrent with increased index investment. Therefore increased volatility in oil futures may explain some of the increased volatility in, for example, agricultural futures. The oil futures market is also one of the most liquid, with extensive trade in contracts up to 3 years from delivery and listings up to 9 years from delivery (against 4 years for agricultural commodities such as corn).

The calibrated model provides a good statistical approximation to oil prices in the pre-financialization period, matching futures volatility and risk premium, spot and futures price autocorrelation, and the term structure of futures prices. It also generates periods of backwardation and contango at reasonable frequen-

²See for example the GSCI fact sheet at

<http://www.standardandpoors.com/indices/articles/en/us/?articleType=PDF&assetID=1245186878016>

cies, and approximates the hedging behavior of producers. The calibration implies reasonable macroeconomic properties for oil, which constitutes roughly 3% of total value of household consumption. The model is nevertheless tractable, and results are robust to small changes in the seven parameters used in calibration.

The model is an extension of the canonical commodity storage model developed in Williams and Wright (1991), which is analyzed empirically in Deaton and Laroque (1992) and (1996). Its appeal is simplicity coupled with an ability to produce autocorrelated spot prices and occasional, dramatic price “spikes” characteristic of the data. Routledge, Seppi, and Spatt (2000) extend the model to analyze forward prices, and conclude that the storage model performs surprisingly well when calibrated to crude oil futures, matching the shape of the mean futures curve and the unconditional term structure of futures volatility.³ However these models abstract from production, consumption, and the risk premium, focusing on a risk-neutral dealer or “speculator” with access to a storage technology. Futures prices are such that the dealer is indifferent to trade. I add producers and consumers alongside the dealer. Heterogeneous preferences and technological endowments motivate trade in spot and futures markets, and generate a time-varying risk premium. Transaction costs, which Hirshleifer (1990) identified and modeled as a barrier to consumer participation in futures markets, are incorporated to proxy for financial innovation.

Other recent empirical papers also focus on oil while analyzing the financialization of commodities. Singleton (2011) finds that investor flows have predictive power for excess holding returns on oil futures at longer horizons. Buyuksahin et al. (2011) document changes in the amount and composition of futures trade, and demonstrate associated changes in the cointegration of futures over the term structure. Bessembinder et al. (2012) provide a detailed description of how oil ETFs operate, and investigate transaction costs associated with a rolling futures position. Although they find transaction costs of roughly 30 basis points per roll (around 4% per year), they reject the hypothesis of “predation” of retail traders. Pan (2011) estimates semi-parametric and non-parametric state price densities (SPD) for crude oil derivatives, and relates futures volume to skewness in the SPD. Hamilton and Wu (2011) estimate a time-varying risk premium on oil futures using a vector autoregression (VAR) that incorporates the position of index traders. Broadening the scope to agricultural commodities, Brunetti and Reiffen (2011) model financialization as participation by uninformed index traders, and find that they reduce hedging costs in theory and in the data. Irwin and Sanders (2011) summarize additional literature on commodity financialization. My paper contributes to the

³In their model as in mine, forward and futures contracts are functionally equivalent. However futures data is more readily available.

literature by incorporating storage and consumer decision making. This allows me to analyze joint spot and futures market dynamics under financialization.

Several recent papers analyze structural models of oil markets without explicitly modeling financial innovation. In a frictionless DSGE model, Baker and Routledge (2012) show that changes in open interest and risk premia on oil futures arise endogenously as a result of heterogeneous risk-aversion. The same forces lead to persistent increases in spot price levels, as wealth drifts toward the more oil-loving agent in the economy. Ready (2012) demonstrates that changes in oil spot and futures price dynamics after 2003 can be jointly explained by a structural break in the oil consumption process. Alternatively, Caballero, Farhi, and Gourinchas (2008) suggest that oil prices increased due to the formation of a rational bubble, with oil replacing housing-related assets as a store of value. Two recent papers present static (two-period) versions of the storage model with active futures markets, and find empirical support for the models' predictions. Acharya, Lochstoer, and Ramadorai (2012) study the connection between managerial risk-aversion and hedging in oil markets, and find that empirical proxies for managerial risk-aversion forecast futures returns. Gorton, Hayashi, and Rouwenhorst (2012) document a connection between inventories and futures risk-premia in markets for many storable commodities. Finally, Arseneau and Leduc (2012) study a general equilibrium storage model with production and consumption. They abstract from derivatives markets to focus on connections between spot prices and the macroeconomy, and examine the effects of biofuel and food subsidies.

I organize the paper as follows. Section 2 describes the model. Section 3 defines equilibrium and describes the solution technique. Section 4 summarizes the available data and describes the model calibration. Section 5 presents results, and Section 6 concludes.

1.2 The Model

I model a dynamic, stochastic, infinite-horizon economy with two goods: a composite numeraire good, and a commodity. The economy is populated by three competitive price-taking agents: a commodity producer, a commodity dealer, and a household. The agents are distinguished by their endowments, preferences, and access to financial markets. The dealer and producer are commercial, whereas households represent consumers. Dealers and producers are concerned with numeraire profits, whereas the household maximizes utility over consumption of the numeraire and the commodity.

1.2.1 Markets

Before proceeding to detailed specifications of each agent, I describe the markets governing interaction among agents. All prices are real, and denominated in units of the numeraire. There is a frictionless spot market for the commodity. In each period t , any agent may buy or sell the commodity at spot price s_t per unit, which is determined in equilibrium. The commodity is always in positive net supply, and cannot be “sold short” on the spot market. There is also an incomplete financial market for commodity futures contracts in which only the front contract (with one period until maturity) is actively traded. The futures contract promises delivery of one unit of the commodity at time $t + 1$ for price f_t paid at $t + 1$. The futures price is chosen such that its date t value is zero: if ϕ_t contracts are bought today, no money changes hands initially, but the buyer pays (or receives) $\phi_t(s_{t+1} - f_t)$ at $t + 1$. Households pay a transaction cost at settlement: after buying ϕ_t contracts, he pays (or receives) $\phi_t(s_{t+1} - f_t) - \tau f_t \phi_t^2$.⁴ The transaction cost is dissipative. Futures contracts are in zero net supply, and all agents are free to take long or short positions - or both, in the case of the dealer. The dealer acts as an intermediary between producers and households, but earns no spreads: up to the household transaction cost, all agents face the same equilibrium futures price f_t . Although only the front contract is actively traded, I allow dealers to notionally trade longer term contracts “among themselves”, with futures prices determined such that the representative dealer’s position is zero in equilibrium.

The household settles transactions out of a numeraire endowment, whereas the dealer and producer have access to credit at fixed rate r . Therefore the household is subject to a budget constraint (reflecting aggregate wages and dividends, for example), whereas the commercial types face a cost of capital (reflecting access to liquid global credit markets). I abstract from active numeraire bond or equity markets. Financial markets are designed to enable hedging or speculation, but they do not allow intertemporal transfers.⁵ The model also abstracts from collateral constraints. For an analysis of intertemporal risk sharing with equity, bonds, and fully-collateralized futures contracts, see Baker and Routledge (2012).

⁴The choice of quadratic, rather than linear transaction costs is primarily one of numerical convenience, as households may be long or short futures. Payment of the transaction cost at settlement seems reasonable if it is interpreted as fund or ETF management fees, for example, and avoids introducing intertemporal transfers into the household problem.

⁵Abstracting from the source of producer and speculator capital makes the model easier to solve, and avoids issues of survival related to the wealth dynamics of the agents.

1.2.2 Producers

The representative producer seeks to maximize expected risk-adjusted profits through production and trade of the commodity and related futures contracts. The level of productivity is determined by a random variable a_t , a finite-state Markov chain, which proxies for various supply side disruptions or booms.⁶ The firm owns an oil well that produces a_t in period t . In addition the firm may adjust production intensity I_t upward, but this is done subject to a one-period lag, and decreasing returns to scale, producing $a_{t+1}I_t^{1/2}$ in period $t + 1$.⁷ The objective is to represent elementary features of oil production in an extremely simple way. Production is inelastic in the very short term. If existing capacity is underutilized then output can be increased with some delay, reflecting the fact that the product must often be shipped to market. The model abstracts from long-term opportunities to explore for and develop additional oil deposits. Obviously the setup is highly stylized. More realistic models of oil production with irreversible investment include Kogan, Livdan, and Yaron (2009) and Casassus, Collin-Dufresne, and Routledge (2009), whereas Carlson, Khokher, and Titman (2007) takes account of the fact that oil is an exhaustible resource.

The firm maximizes expected discounted profits given cost of capital r and a penalty for risk:

$$\begin{aligned} & \max_{\{I_t, \phi_t\}_{t=0}^{\infty}} \sum_{t=1}^{\infty} (1+r)^{-t} \left(E_0[p_t^p] - \frac{\theta}{2} \text{Var}_0[p_t^p] \right), \\ & \text{s.t. } p_{t+1}^p = s_{t+1}a_{t+1}(1 + I_t^{1/2}) - (1+r)I_t + \phi_t(s_{t+1} - f_t), \\ & \quad I_t \geq 0, \forall t \end{aligned} \tag{1.1}$$

for L_{-1} and ϕ_{-1} given. Producers have mean-variance preferences over profits, where the aversion to variance is a reduced-form reflection of bankruptcy costs, unmodelled owner preferences, management risk-aversion, etc. The variance of $t + 1$ profits can be decomposed into

$$\text{Var}_t[p_{t+1}^p] = \underbrace{\text{Var}_t[s_{t+1}a_{t+1}]}_{\sigma_{sa,t}^2} (1 + I_t^{1/2})^2 + 2 \underbrace{\text{Cov}_t[s_{t+1}a_{t+1}, s_{t+1}]}_{\sigma_{sas,t}} (1 + I_t^{1/2})\phi_t + \underbrace{\text{Var}_t[s_{t+1}]}_{\sigma_{s,t}^2} \phi_t^2 \tag{1.2}$$

⁶Throughout the paper, subscripts with respect to t denote measurability with respect to period t information.

⁷This is a gross simplification of a Cobb-Douglas production technology with stochastic TFP a_t and constant return to scale. Start with production function $f(K_t, L_t) = a_t K_t^\nu L_t^{1-\nu}$, $\nu \in [0, 1]$. I assume that aggregate labor supply is fixed at $L = 1$, so that we have simply $f(K_t, L_t) = a_t K_t^\nu$. (Assume N firms, each endowed with $L = 1/N$ labor, and let $N \rightarrow \infty$.) Let capital be the numeraire (equivalently that the numeraire can be converted into capital at the frictionless rate of 1). Assume K_0 is an exogenously specified constant. The firm decides how much capital to invest today for use in production tomorrow. Existing capital depreciates at rate $\rho \in (0, 1]$, so we have $K_t = I_{t-1} + (1 - \rho)K_{t-1}$. Setting $\nu = 1/2$ and shutting down capital accumulation with $\rho = 1$ obtains the result.

Taking first order conditions of the maximization problem and solving for the production intensity and futures portfolio gives solutions

$$\begin{aligned}\phi_t &= \frac{E_t[s_{t+1}] - f_t - \theta\sigma_{sas,t}(1 + I^{1/2})}{\theta\sigma_{s,t}^2} \\ I_t &= \left(\frac{\sigma_{s,t}^2 E_t[s_{t+1}a_{t+1}] - \sigma_{sas,t}(E_t[s_{t+1}] - f_t) - \theta(\sigma_{s,t}^2\sigma_{sa,t}^2 - \sigma_{sas,t}^2)}{\theta\sigma_{s,t}^2\sigma_{sa,t}^2 + \theta\sigma_{sas,t}^2 + 2\sigma_{s,t}^2(1 + r)} \right)^2\end{aligned}\quad (1.3)$$

1.2.3 Dealers

Dealers are intermediaries in the futures market. They neither produce nor consume the commodity, but they have access to a storage technology, through which they participate in the goods market. Like producers, dealers have mean-variance preferences over profits, although their risk-aversion parameter ρ may differ from that of producers (θ). Dealers are sometimes called “speculators” in traditional storage models, where they are assumed to be risk-neutral, and so accumulate inventory only in the hope that its value will appreciate (net of costs). My model nests this as a special case, with $\rho = 0$. However I focus on the case of risk-averse dealers ($\rho > 0$). With risk-aversion and an active futures market, some elements of dealer behavior may be viewed as “hedging”, so I avoid the term “speculator” as potentially misleading. Dealers may borrow at interest rate r , and have access to a commodity storage facility that can preserve a nonnegative quantity of the commodity at a cost: there is a numeraire storage fee of k per unit.⁸ In a given period t , the state variable Q_{t-1} , which represents inventory held over from the previous period, is a key determinant of present period spot and futures prices. The evolution of inventory is what makes the dynamic storage model interesting, as a buildup of inventory can lead to depressed and stable prices, whereas “stockouts” ($Q_t = 0$) are associated with higher and more volatile prices.

⁸An alternative cost to storage - depreciation or spoilage - is used in Deaton and Laroque (1992). Cafiero et al. (2011) compare the two costs, and argue that a numeraire cost improves the basic storage model’s ability to match asset prices, and is more realistic, as fees charged by storage facilities typically do not track spot prices. In an extended version of the model I implement both costs; for purposes of calibrating to oil the difference between the two seems modest. The choice potentially has important welfare implications, but these are tied up with the modelling assumption that the numeraire can be borrowed elastically at rate r , whereas the commodity is in finite supply.

The dealer's lifetime risk-adjusted profit maximization problem is

$$\begin{aligned} \max_{\{Q_t, \psi_t\}_{t=0}^{\infty}} \sum_{t=1}^{\infty} (1+r)^{-t} \left(E_0[p_t^d] - \frac{\rho}{2} \text{Var}_0[p_t^d] \right), \\ \text{s.t. } p_{t+1}^d = s_{t+1}Q_t - (1+r)(s_t+k)Q_t + \psi_t(s_{t+1} - f_t), \\ Q_t \geq 0, \forall t. \end{aligned} \tag{1.4}$$

for Q_{-1} and ψ_{-1} given. Although the *change* in inventory $\Delta Q_t = Q_t - Q_{t-1}$ is important for equilibrium prices, given the price-taking behavior of dealers we can still address each period's optimization problem individually.

When the nonnegative inventory constraint is nonbinding ($Q_t > 0$), first order conditions imply

$$\begin{aligned} Q_t &= \frac{E_t[s_{t+1}] - (1+r)(s_t+k)}{\rho\sigma_{s,t}^2} - \psi_t, \\ \psi_t &= \frac{E_t[s_{t+1}] - f_t}{\rho\sigma_{s,t}^2} - Q_t. \end{aligned} \tag{1.5}$$

In isolation the solution to the dealer's problem is indeterminate; the dealer's futures position is determined by market clearing in equilibrium. When the inventory constraint binds, then the solution to the dealer's problem is

$$\begin{aligned} Q_t &= 0 \\ \psi_t &= \frac{E_t[s_{t+1}] - f_t}{\rho\sigma_{s,t}^2}. \end{aligned} \tag{1.6}$$

The zero-inventory case highlights the speculative aspect of the dealer's behavior: when $Q_t = 0$, his futures position reflects only the direction and magnitude of the risk premium, scaled by his risk-aversion. However when $Q_t > 0$ as in Equation (1.5), a hedging effect is present. If the dealer takes a short futures position, it hedges some of his inventory risk, so he is inclined to buy more inventory. If on the other hand he is long futures, he reduces inventory, because he is already long the next period spot price. Similarly the FOC for futures reflects any exposure to next period's spot price due to inventory. Therefore the dealer's decisions reflect a combination of hedging and speculative motives. If the risk premium to the long futures position is sufficiently high, the dealer may choose to be long the physical commodity and the futures contract simultaneously.

1.2.4 Households

Households consume the commodity, purchased on the spot market. They finance their purchases out of an endowment of one unit of the numeraire each period. The amount of numeraire left over after purchasing the commodity is consumed also. Specifically, households enjoy utility over composite consumption described by CES aggregator

$$c_t = A(c_{x,t}, c_{y,t}) = \left[(1 - \gamma)c_{x,t}^\eta + \gamma c_{y,t}^\eta \right]^{1/\eta}, \quad (1.7)$$

with subscript x denoting the numeraire and y the commodity. A limiting case is $\eta \rightarrow 0$, which corresponds to Cobb-Douglas aggregator

$$c_t = A(c_{x,t}, c_{y,t}) = c_{x,t}^{1-\gamma} c_{y,t}^\gamma. \quad (1.8)$$

Utility is time additively separable, given in each period by

$$u(c_t) = \log(c_t). \quad (1.9)$$

Although many commodities are used as intermediate goods in the production of household items, the simplifying assumption that households consume the commodity directly seems reasonable, particularly in the case of oil. Most of the cost to producing refined products such as gasoline and heating oil is accountable to the crude input, so the spot prices of refined products track crude oil prices.

Accordingly, in each period households optimally choose to consume an amount $c_{y,t}$ of the commodity, which they purchase at the spot price s_t . To hedge future exposure to variation in the spot price they also buy a number of nearest-to-maturity futures contracts φ_t . Note that the price of these contracts at purchase is zero by construction, with any gains or losses to ownership settled upon delivery next period. However households incur a small transaction cost τ proportional to the face value of the contracts squared, payable at time of settlement. The adjusted endowment available to spend on current consumption is

$$\hat{x}_t = 1 + \varphi_{t-1}(s_t - f_{t-1}) - \tau \varphi_{t-1}^2. \quad (1.10)$$

The fact that the only financial asset available to households is a one-period zero-price contract is important to the determination of equilibrium, since it implies that households may hedge their risks but cannot smooth consumption through saving. Household portfolio choice reduces to a series of one-period problems.

Taking price processes as given, households solve lifetime utility maximization problem

$$\begin{aligned} \max_{\{c_{x,t}, c_{y,t}, \varphi_t\}_{t=0}^{\infty}} (1 - \beta) E_0 \sum_{t=0}^{\infty} \beta^t \log(A(c_{x,t}, c_{y,t})) \\ \text{s.t. } c_{x,t} + s_t c_{y,t} \leq \hat{x}_t, \\ \hat{x}_t = 1 + \varphi_{t-1}(s_t - f_{t-1}) - \tau f_{t-1} \varphi_{t-1}^2 \end{aligned} \quad (1.11)$$

Applying Walras Law, substituting for the budget constraint, and differentiating w.r.t. the choice variables yields the following optimality conditions:

$$\begin{aligned} c_{y,t} &= \hat{x}_t s_t^{-1} \left[1 + \left(\frac{\gamma}{(1 - \gamma) s_t^\eta} \right)^{\frac{1}{\eta-1}} \right]^{-1} \\ c_{x,t} &= \hat{x}_t - s_t c_{y,t} \\ 0 &= E_t \left[\frac{s_{t+1} - f_t - 2\tau f_t \varphi_t}{1 + \varphi_t (s_{t+1} - f_t) - \tau f_t \varphi_t^2} \right]. \end{aligned} \quad (1.12)$$

The subjective discount factor $\beta \in (0, 1)$ does not appear in the FOCs. Although the household policy for distributing spending over goods is closed-form, the futures position must be determined numerically.

1.3 Equilibrium

Taking policies as given, prices are determined by market clearing. Define production of the commodity as

$$y_t = a_t (1 + I_{t-1}^{1/2}). \quad (1.13)$$

In each period, commodity market clearing requires

$$y_t - c_{y,t} = Q_t - Q_{t-1} = \Delta Q_t \quad (1.14)$$

Expanding and reorganizing terms above, the equation that the current spot price s_t must satisfy to clear goods markets is

$$(1 + \varphi_{t-1}(s_t - f_{t-1}) - \tau f_{t-1} \varphi_{t-1}^2) s_t^{-1} \left(1 + \left(\frac{\gamma}{(1 - \gamma) s_t^\eta} \right)^{\frac{1}{\eta-1}} \right)^{-1} = y_t - \Delta Q_t, \quad (1.15)$$

which must be solved numerically. Fortunately finding the futures price that clears financial markets is simple given state-contingent spot prices. Futures positions must net to zero,

$$\phi_t + \varphi_t + \psi_t = 0. \quad (1.16)$$

Writing $\psi_t = -(\phi_t + \varphi_t)$ and substituting into dealer FOC Equation (1.5), the futures price is

$$f_t = E_t[s_{t+1}] - \rho\sigma_{s,t}^2(Q_t - (\phi_t + \varphi_t)) \quad (1.17)$$

Definition 1 (Equilibrium). *Equilibrium is a sequence of state-contingent prices and policies $\{s_t, f_t, Q_t, I_t, \phi_t, \varphi_t, \psi_t, c_{x,t}, c_{y,t}\}$ such that each agent's policy solves his maximization problem, and commodity spot and futures markets clear $\forall a_t, t \geq 0$, for $\{f_{-1}, Q_{-1}, I_{-1}, \phi_{-1}, \varphi_{-1}, \psi_{-1}\}$ given.*

The solution is characterized by equilibrium aggregate inventory function

$$Q_t = J(Q_{t-1}, f_{t-1}, \varphi_{t-1}, a_t), \quad (1.18)$$

spot price function

$$s_t = S(Q_{t-1}, f_{t-1}, \varphi_{t-1}, a_t), \quad (1.19)$$

and household portfolio function

$$\varphi_t = \Phi(f_t, s_{t+1}). \quad (1.20)$$

As specified above, equilibrium policies are defined over three continuous variables (f_{t-1} , Q_{t-1} and φ_{t-1}) and one discrete variable (a_t). However, if we restrict portfolios to lie on equilibrium values, then these functions can be reparameterized as

$$\begin{aligned} Q_t &= J(Q_{t-1}, a_{t-1}, a_t) \\ s_t &= S(Q_{t-1}, a_{t-1}, a_t) \\ \varphi_t &= \Phi(Q_t, a_t) \end{aligned} \quad (1.21)$$

This reduces the state space to one continuous variable and two discrete variables (or one discrete variable if we assume a_t is i.i.d.), which makes solving for equilibrium simpler. In practise I approximate state variable Q_{t-1} using a grid, and interpolate using a shape-preserving polynomial. Cafero et al. (2011) find, and I confirm, that the problem does not forgive use of a very coarse regular grid, and solving the problem with a fine 4-d grid would not be practical. Because futures are not collateralized, and so are costless at

purchase and don't allow the transfer of wealth between periods, the previous period's futures position does not constrain the household's choice for the next period. Therefore even if off-equilibrium futures positions were incorporated, households would adopt the equilibrium policy the following period.

I solve the model by policy iteration as follows:

1. Guess initial inventory and policy functions
2. Update household portfolio rule taking inventory policy as given
3. Update inventory rule taking household and future inventory policy as given
4. Repeat steps 2 and 3 until convergence

Although the basic solution technique is entirely standard, the problem is more computationally intensive than for the canonical dynamic storage model, chiefly because it involves solving a nested system of non-linear equations for state-contingent prices given policies. Details are available upon request.

1.3.1 Characterization

Since the model lacks a closed-form solution, I provide some intuition for how it functions using the baseline parameters given in Table 1.1. The exact choice of parameters is justified in subsequent sections, but the model has some basic characteristics that are independent of parameter choices. There is one endogenous state variable, the level of inventory, and many of the model's interesting features depend upon how inventory is accumulated and sold off by the dealer. Figure 1.4 shows optimal inventory policy for the baseline parameters. The dealer's end of period inventory Q_t depends on inventory carried over from the previous period (Q_{t-1} , given on the x-axis), and the productivity realization (a_t , shown by the different curves). Inventory policy is also affected by lagged productivity, a_{t-1} ; however the effect is modest. I show results for $a_{t-1} = 1$. The 45° line is shown in dashed red, with the region above that line indicating inventory accumulation and the one below indicating sell-offs. When inventory is low, the dealer will choose to accumulate inventory if productivity is high. Given a succession of high productivity realizations, he eventually reaches a point where no more inventory is purchased, even in the most productive states: an endogenous maximum inventory level. There are also low productivity states in which the producer always sells inventory, and a

succession of these low states will eventually lead to a stock-out ($Q_t = 0$). Therefore inventory is bounded, and given sufficient time, a stockout is guaranteed to occur (i.e., inventory is a regenerative process). It follows that the distribution of long-run prices does not depend upon current inventory. These basic characteristics are features of the classic storage model, and in simpler settings they are provable properties of equilibrium (see e.g. Routledge, Seppi, and Spatt (2000)). With endogenous production and active futures markets, the properties of equilibrium are difficult to prove, because the function giving spot prices in terms of inventory changes and productivity is itself an equilibrium construct. Nevertheless, the basic storage model characteristics hold for a range of parameter values, and I have found no equilibrium in which they were violated.

From an asset pricing perspective, inventory dynamics are interesting because they generate persistent spot prices, even if productivity has little persistence or is i.i.d. Storage also leads to moderate variance in spot prices most of the time, as inventory is usually available to smooth consumption and hence prices. But when inventory is exhausted and production is poor, prices jump; consumers are willing to pay a high price for a rare, but essential, commodity. These features are seen in Figure 1.5, which shows equilibrium spot prices versus inventory (on the x-axis) and productivity realizations (different curves). The top curve is the low-productivity state in which inventory is sold off, whereas the bottom curve is a high-productivity state in which inventory is accumulated, up to a point. Fluctuation in productivity will cause prices to bounce vertically between the curves, but little horizontal movement (change in inventory) will occur in the short term. The result is periods of high prices and high volatility when inventory is low, and periods of low volatility and low prices when inventory is plentiful.

Therefore desirable features of the classic storage model are preserved, but several new features are added that improve the utility of the model for asset pricing. By relaxing the traditional assumption that the dealer is risk-neutral, I introduce a risk-premium into futures prices, a key determinant of which is the volatility of the underlying spot price. With active trade in futures markets, the risk premium affects production, storage, and consumption decisions, which naturally feed back into spot prices. This leads to a rich interaction between spot and futures markets, the basic mechanics of which are explained in the next section.

1.3.2 Prices and Risk Premia

The one-period futures contract is the only actively traded contract in the model, and the only one with an impact upon spot prices. Therefore it is the focus of analysis. However generalizing definitions of futures returns and risk-premia to multi-period contracts is easy, and allows for testing the model's implications regarding the slope of the futures curve and the term structure of risk premia. I write $f_{t,n}$ for an n-period to delivery futures contract, with $f_t = f_{t,1}$ the one-period contract, and $f_{t,0} = s_t$ the deliverable contract.

The risk premium on a futures contract is the expected excess one-period holding return on that contract. This definition requires a notion of "returns" for a contract that is priced at zero by construction. I follow the usual approach of considering a fully-collateralized contract, with collateral in the form of a one-period bond. This has price $f_{t,n}$ at date t , and value $(1+r)f_{t,n} + f_{t+1,n-1} - f_{t,n}$ at $t+1$. The one-period holding return on the collateralized contract is⁹

$$\frac{(1+r)f_{t,n} + f_{t+1,n-1} - f_{t,n}}{f_{t,n}} - 1 = \frac{f_{t+1,n-1}}{f_{t,n}} - 1 + r, \quad (1.22)$$

and so the excess holding return is

$$\frac{f_{t+1,n-1}}{f_{t,n}} - 1. \quad (1.23)$$

Focus for now on the one-period contract. One can look at the risk premium (expected excess holding return) through at least two lenses: in terms of trade in futures markets, or in terms of fundamentals. In terms of futures positions, the risk premium is

$$\frac{E_t[s_{t+1}] - f_t}{f_t} = \frac{\rho\sigma_{s,t}^2(Q_t + \psi_t)}{f_t} \quad (1.24)$$

This suggests that the risk-premium will be larger when the dealer's net position in the futures market is positive, or if the dealer is very risk averse. It also suggests a large risk premium given lots of inventory, or high dealer risk aversion. However high inventory will tend to decrease spot-price variance, which will decrease the risk-premium, and high risk aversion will reduce the dealer's willingness to take a long futures position. Which of these effects dominates will depend upon the parameterization of the model. We can conclude, however, that the risk-premium must be positive if the dealer's futures position is positive.

⁹When comparing with estimates from the data it is convenient to use the continuously compounded excess return, which is $\log(f_{t+1,n-1}) - \log(f_{t,n})$. The risk premium is the expected excess return, $E_t[\log(f_{t+1,n-1})] - \log(f_{t,n})$.

Derive another expression for the risk premium that focuses on “fundamentals” from the market clearing condition for futures:

$$\begin{aligned}
-\frac{E_t[s_{t+1}] - f_t - \theta\sigma_{sas,t}(1 + I^{1/2})}{\theta\sigma_{s,t}^2} - \varphi_t &= \frac{E_t[s_{t+1}] - f_t}{\rho\sigma_{s,t}^2} - Q_t \\
\Rightarrow \frac{E_t[s_{t+1}] - f_t}{f_t} &= \frac{\rho\theta}{f_t(\rho + \theta)} \left(\sigma_{sas,t}(1 + I^{1/2}) + \sigma_{s,t}^2(Q_t - \varphi_t) \right).
\end{aligned} \tag{1.25}$$

Consider the situation prior to financialization (assume $\varphi_t = 0$):

$$\frac{E_t[s_{t+1}] - f_t}{f_t} = \frac{\rho\theta}{f_t(\rho + \theta)} \left(\sigma_{s,t}^2 Q_t + \sigma_{sas,t}(1 + I^{1/2}) \right). \tag{1.26}$$

This result implies a negative risk premium is possible before household participation in the futures market, depending upon whether price or quantity effects dominate in the producer’s profits. This will also determine whether the producer takes a long or short position in the futures market. Prior to financialization, a negative risk premium is only possible if the quantity effect dominates, in which case the producer must take a long position, per Equation (1.24), as $-\phi_t = \psi_t$. After financialization ($\varphi_t \neq 0$) the connection between the sign of the risk premium and the producer’s hedging motive is loosened.

To extend the results to multi-period contracts, define long-dated futures prices such that markets clear with dealers - the financial intermediaries of the model - as the only participants. Let $\psi_{t,n}$ be the dealer’s position in n -period to delivery futures, and solve his maximization problem for arbitrary n on the assumption that the only other contract available for trade is the 1-period contract. The futures price for the n -period contract is such that it is not traded in equilibrium (i.e., such that the market with only dealer participants clears).

The dealer’s augmented profit-maximization problem is

$$\begin{aligned}
\max_{Q_t, \psi_t} E_t[\hat{p}_{t+1}^d] - \frac{\rho}{2} Var_t[\hat{p}_{t+1}^d], \\
\text{s.t. } \hat{p}_{t+1}^d &= p_t^d + \psi_{t,n}(f_{t+1,n-1} - f_{t,n}), \\
p_{t+1}^d &= s_{t+1}Q_t - (1 + r)(s_t + k)Q_t + \psi_{t,1}(s_{t+1} - f_{t,1}).
\end{aligned} \tag{1.27}$$

Imposing $\psi_{t,n} = 0, \forall n > 1$ and determining prices to clear markets, the futures price at t for delivery at $t + n$ is

$$f_{t,n} = E_t[f_{t+1,n-1}] - \rho(Q_t + \psi_{t,1})Cov(s_{t+1}, f_{t+1,n-1}) \tag{1.28}$$

For any n , the price of the futures contract can be computed by recursing until $n = 1$, and using the price of the one-period contract.

1.4 Data and Calibration

Because the model features several counterbalancing forces, a realistic calibration is required to estimate the direction and magnitude of financialization's effects. I calibrate the model to crude oil. Because crude oil markets have received attention in several recent empirical papers, I highlight only a few aspects of the data that explain model design and calibration choices. I first determine the stochastic forcing process that drives most of the variation in quantities in the model. Given this process, I then choose agent parameter values to match moments for asset prices prior to financialization. This leaves one key parameter to vary in the results section: the household transaction cost parameter τ .

I use data on quantities from the Energy Information Administration (EIA)¹⁰. Annual world oil supply is available from 1970-2009. Monthly US oil consumption data ("U.S. Product Supplied") is available from Jan. 1963 - May 2011. US consumption data exhibits seasonalities that are particularly strong in the first two decades of the sample, but become modest in recent years. Given that world production data is only available on a monthly basis from 2001, I annualize domestic consumption to match the longer annual world production sample. Since I abstract from growth in the model, I estimate a linear time trend in each of the samples using OLS regression, and normalize the data to obtain fractional deviations from trend. The result is plotted in Figure 1.2. A natural alternative is to normalize US oil consumption by US GDP - at this coarse level of analysis the result is similar.¹¹ In the figure, we see that global supply and domestic consumption are highly correlated (coefficient of 0.85).

Despite the existence of frictions (in the form of national energy subsidy programs, embargoes, etc.), there is obviously extensive international trade in oil. Global production is more closely tied to US consumption than is US production. As seen in Figure 1.2, US production actually declines by roughly 25 % over the sample period, whereas US consumption increases around 25%. Therefore the interpretation of the model as having a "global producer" but a "US Consumer" seems reasonable. Figure 1.3 further emphasizes that variation in global production and US oil consumption are tightly linked. Although many factors may have an effect on oil markets, for increased tractability and ease of interpretation I choose to model only one exogenous shock: an oil productivity shock.

¹⁰www.eia.gov

¹¹GDP data is from the BEA.

I define a simple Markov process for oil productivity shocks that leads to equilibrium oil production and consumption sequences similar to the data. Global production and US consumption are characterized by persistent deviations from trend. I estimate AR1 parameters via the Yule-Walker method, which yields autocorrelation of 0.81 and conditional standard deviation of 0.031 for production, and autocorrelation of 0.86 and conditional standard deviation of 0.046 for consumption. I adjust the parameters to a quarterly rather than annual frequency; adjusted values are in Table 1.2. I choose a quarterly calibration because it implies households will hold a rolling position in the three-month futures contract, which is more consistent with typical fund strategies than a 1-year contract and annual rolls or a 1 month contract with monthly rolls. The monthly calibration would also be unappealing for the very high autocorrelation necessary to model consumption and production at that frequency. It is well-known that approximating an AR1 with close-to-unit root is difficult using a finite-state Markov process. Floden (2008) investigates various approximation schemes, and finds the method of Tauchen (1986) to be relatively robust when using a small state space. I use this method with 5 states. For comparison, Cafiero et al. (2011) and Deaton and Laroque (1996) use a 10-state approximation to independent normal shocks, and Routledge, Seppi, and Spatt (2000) use a 2-state Markov process. With only 5 states, there is a trade-off between “high-frequency” variation (i.e., non-zero quarterly changes) and high autocorrelation: with very high autocorrelation and few states, autocorrelation and conditional standard deviation are matched using large but infrequent changes in productivity. Since inventory will smooth shocks and increase persistence of consumption in equilibrium, I relax autocorrelation to 0.6 to allow smaller but more frequent changes in productivity, and set conditional standard deviation to 0.035.

Calibration of prices uses monthly spot prices from the St. Louis Fed, and monthly NYMEX futures prices from barchart.com. I focus on the period 1990-2011, for which futures data is available (spot prices are available from 1969). Consistent with the quarterly calibration, I use the 3-month futures contract as the front (one-period) contract, and the 6-month contract as the second (two-period) contract. The slope of the futures curve is the 6-month price less the 3-month price. I estimate autocorrelation and unconditional standard deviations of prices using quarter-on-quarter prices. Although I assume fundamentals remain unchanged over the sample period (and production data is inadequate for a split sample), I split prices into pre (1990-2003) and post (2004-2011) financialization samples. The break is consistent with Baker and Routledge (2012), and similar to Hamilton and Wu (2011), who split the sample at the beginning of 2005. Statistics are summarized in Table 1.2.

Detailed data on futures positions is not freely available, so I rely upon references for summary statistics. Acharya, Lochstoer, and Ramadorai (2012) conduct a survey of roughly 2,500 quarterly and annual reports of oil-sector firms since June of 2000. They find that roughly 70% of firms hedge at least 25 % of their production. Given that some firms hedge more than 25%, I assume that producers in general are short around 25% of production. Regarding households, I found no commodity funds targeting retail investors prior to 1996, when the Oppenheimer Real Asset Fund was established with the purpose of pursuing investments linked to the GSCI index.¹² The first commodity index ETF was the DB Commodity Index Tracking fund, established January 2006, with heavy weights on crude and heating oil futures selected from contracts under 13 months based on maximum “implied roll yield.”¹³ The oil-only ETF USO began trading in April, 2006. Although wealthy households may have participated in futures markets prior to 1996, most households could not have participated prior to the availability of retail funds, except indirectly through pension funds. For purposes of calibration, I assume that no households participated in futures markets during 1990-2003. For the 2004-2011 period, I use figures from Stoll and Whaley (2010), CFTC (2008) and the EIA, and estimate that households hedged around 20% of their exposure to crude oil.¹⁴

Consider the following thought experiment: suppose “fundamentals” - production and storage technologies and the firms that use them - remain the same over the sample period, such that the only exogenous structural change reflects financial innovation allowing household participation in futures markets. How well can the model match the data prior to household entry? As we relax transaction cost τ to reflect financial innovation, does the behavior of the model change in a way that is consistent with the data? To carry out the experiment, I set $\tau = \infty$ to reflect no household trade in futures, and choose the remainder of the model parameters to approximately match the data during the pre-Entry period (1990-2003). I refer to this as the “baseline calibration”. Subsequently I reduce transaction cost τ while leaving the other parameters unchanged, and compare the model’s “post-Entry” behavior with the 2004-2011 data. The results section argues that the changes in the model’s behavior are consistent with the data along several dimensions, despite

¹²See prospectus at <http://www.sec.gov/Archives/edgar/data/1018862/0001018862-97-000003.txt>

¹³See <http://www.sec.gov/Archives/edgar/data/1328237/000119312506118678/d424b3.htm>

¹⁴Quarterly open interest data from 2000-2012 is available from www.eia.gov/finance/markets/financial_markets.cfm, which lists Q1 open interest of around 1,400,000 contracts. From Stoll and Whaley (2010), over 30% of open interest is attributable to index investors, of which 50% corresponds to mutual funds, ETFs, etc., and a further 40% to institutional investors, including pension funds. Therefore I attribute 25% of open interest to households, around 350,000 contracts, corresponding to around 20% of quarterly household consumption in 2008. This is, of course, a very rough estimate. Since not all index funds rely upon commodity futures, the correct figure may be higher.

adjusting only one parameter.

Parameter values for the baseline calibration are given in Table 1.1. I set the risk-free rate r to the average real return on 90-day Treasury Bills from 1990-2012, roughly 0.1%.¹⁵ Storage costs are $k = 0.001$ per unit per quarter, which is around 3% of the average unit price of oil. Household goods aggregation parameters γ and η are used to match the value of oil consumption as a fraction of total consumption, and the sensitivity of spot prices to changes in oil consumption (essentially, spot price volatility). I use producer (θ) and dealer (ρ) risk aversion parameters to match producer hedging as a fraction of output (25% short position), and the futures risk premium, which is around 2% per quarter to the long position. In broad strokes, the levels of these risk-aversion parameters regulate the magnitude of the risk premium, whereas the difference between the parameters determines the size of the producer's futures position. The performance of the model is discussed in more detail in the results section.

1.5 Results

Given the simplicity of the model, the baseline calibration performs fairly well overall. Table 1.2 lists several summary statistics for data and the model, with the baseline case designed to match the 1990-2003 column without household trade in futures ($\tau = \infty$). I review the baseline case first.

1.5.1 Before Financialization

The baseline model is able to match basic statistical characteristics of spot and futures prices. Spot and futures price autocorrelation in the model match the data to a close approximation (coefficients of around 0.7). For ease of comparison (since the data and model use different units), standard deviations are given as a fraction of the mean price. Unconditional standard deviation of spot prices is too high in the model, at roughly 40% in the model to 25% in the data. However the model implies unconditional standard deviation of futures of 25%, around the same as the data. Mean quarterly excess holding returns are around 2% in both the model and the data. Although the futures curve in the model is not in backwardation as frequently

¹⁵Data is from the Federal Reserve Board (<http://www.federalreserve.gov/releases/h15/data.htm>) and the Federal Reserve Bank of Cleveland (<http://www.clevelandfed.org/research/data/us-inflation/chartsdata/>)

as in the data, backwardation is a frequent occurrence: 45% of the time, versus 70% in the data. In futures markets, producers hedge around 28% of their production, which is quite close to the (admittedly imprecise) 25% target.

The calibration matches asset pricing moments without implying production and consumption dynamics completely at odds with the data; given the simple production technology and reduction of the GDP process to a constant, a close match along all dimensions is not realistic. However oil consumption as a fraction of GDP is around 3% in both the model and the data, and unconditional standard deviation of oil production is also consistent, at around 5%. Consumption and production dynamics are influenced most by the choice of parameters for the AR1 productivity process, and fitting an AR1 process to the model's consumption and production series offers some insight into how the effects of the forcing process are transformed in equilibrium. Predictably, smoothing via inventory causes consumption to become more autocorrelated than the AR1 (0.7 versus 0.6), and also reduces conditional standard deviation to around 2%, versus the 3.5% for the forcing process. Meanwhile, production is less autocorrelated (down to 0.55), and slightly more volatile. In terms of matching the data, standard deviation of oil consumption is close conditionally, but not close unconditionally. As discussed previously, autocorrelation of production and consumption was sacrificed on the alter of tractability: it is too low in the model.

Overall this seems a good performance for a model with 7 free parameters. In fact it even matches some moments for futures over the term structure. Figure 1.6 shows mean futures prices over the term structure, with the price of the first contract normalized to one. By this metric, the model is a close fit to the data. Although magnitudes are not matched closely beyond the first contract, the shape of the term structure of standard deviation (volatility) is similar to the data, as shown in Figure 1.7. The ability of the storage model to match unconditional standard deviation over the term structure was emphasized in Routledge, Seppi, and Spatt (2000); these results verify that the storage model continues to do fairly well in this regard, even when constrained by increased realism in the modeling of production and consumption. In addition, the model implies a downward sloping term structure for the risk premium, shown in Figure 1.8, with a slope similar to that in the data.

1.5.2 Financialization

I explore the effects of financial innovation in the oil futures market via the relaxation of the household's transaction cost parameter, τ . Summary statistics for decreasing τ are given alongside the baseline calibration in Table 1.2. The results are quite interesting, although perhaps not surprising. Fundamentals change relatively little, but there are significant changes in asset prices, with the exception of the key value: mean spot prices are essentially unchanged. Across all other price moments there is at least some change, and in every case the direction of change is consistent with the data: decreased autocorrelation, increased standard deviation, decreased risk premium, decreased backwardation, and of course, increasing open interest in futures. With a value of around $\tau = 0.08$, the model matches the risk premium, percent backwardation, and approximate household futures position relative to consumption. A reasonable preliminary conclusion is that financialization explains several recent changes in spot and futures markets, but it had less impact upon fundamentals, and didn't increase price levels. The effect of financialization upon spot prices can be thought of as the net of changes to commercial behavior and changes to household behavior. Although a clean dichotomy is impossible in an equilibrium model, I proceed in such a fashion.

Dealers and Producers

The equilibrium effects of financialization are best understood through the lens of the dealer's optimization problem. The dealer has access to two investment opportunities, stored oil and a futures contract on oil, that offer identical per-unit gross payoffs next period: they will each be worth the spot price of oil, s_{t+1} . Although the inability to store negative amounts of oil differentiates the two investments, in most states of the world the dealer will choose to hold positive inventory; I focus on this case. When inventory is positive, a "no-arbitrage" condition links the futures price (f_t) with the price of buying a unit of oil today and storing it until tomorrow ($s_t + k$). This links financial markets to the goods market. Given that the per-unit cost of storage is constant, any change in the futures price must imply an identical change in the spot price, and visa versa. Therefore in any model with commodity storage and derivatives, one should assume that altering financial markets *will* alter goods markets.

Indeed goods markets are altered, and one of the main conduits is inventory policy. For any starting level of inventory, the dealer is expected to accumulate more inventory after financialization than before.

As shown in a histogram for inventory in Figure 1.10, expected inventory is higher after financialization, and there are fewer stockouts. This is because households wish to take a long futures position to hedge their exposure as consumers of oil, which implies that dealers take a short position. This nets out part of the dealer's futures position with producers, where he takes the long side of the contract. In order for the dealer's policy to remain optimal, he must increase his exposure to oil, either by buying more futures contracts from the producer, or by purchasing more inventory. In equilibrium, the linkage between futures prices and the prices of storage implies that he must do both.

When the dealer stores more inventory, it drives up the spot price today, reducing expected excess returns to storing the commodity. This implies that the dealer is also willing to buy futures at a higher price (with a reduced risk premium to the long side), which makes hedging more attractive to producers. Consequently producers sell more contracts to dealers. Figure 1.9 shows the net futures position of dealers in equilibrium, before and after financialization. Dealers have reduced net exposure to futures after financialization, but part of the reduction from the household's long position is offset by an increased producer short position.

The increased producer short position affects his optimization problem in turn: he has reduced exposure to oil risk, and so is willing to boost output. Figure 1.11 shows that production intensity increases after financialization, conditional on the level of inventory. On a fractional basis intensity increases as much as 10% for a given level of inventory, but the absolute change is small.¹⁶ Because financialization also increases inventory on average, unconditional production intensity increases only about 6% after financialization. Oil production intensity has not been precisely calibrated to capacity utilization data, so the value of the result is less in its magnitude than its direction: financialization boosts average production.

In summary, households provide a natural counterparty to producers, such that financial innovation leads to increased storage, increased production, and a reduced risk premium on futures. The inventory and production effects reduce spot price mean and volatility

¹⁶The main results related to financialization do not depend on the producer's ability to adjust output, based on experiments (not reported) with the restriction $I_t = 0$. The objective of including production intensity is to assess, in a simple way, whether financialization induces producers to increase or decrease output.

Households

The effects of household entry upon the commercial sector are almost directly offset by the effects upon households themselves. Figure 1.13 shows spot prices versus inventory and productivity. When oil production is good, spot prices after financialization are almost the same as before. Although households lose money on their futures position in high productivity states (which should reduce spot prices), dealers accumulate more inventory in high productivity states after financialization (which increases spot prices). The net effect is a wash. However when productivity is low, spot prices are higher after financialization than before, especially if there is a stockout. Households enjoy a windfall on their futures in low productivity states, and dealers cannot sell more than their entire inventory. Household trade in futures effectively makes their endowment positively correlated with spot prices, amplifying the price effects of productivity shocks. Results are summarized in Figure 1.12, which presents a histogram of spot prices before and after financialization. Although increased inventory accumulation after financialization makes stockouts less likely, when they occur, prices spike even higher than before. The result is more volatile spot prices, but very little change in the mean level.

1.5.3 Welfare

Public officials express concern that financialization hurts households by driving up spot prices.¹⁷ Although the model suggests that it has little effect on average spot prices, there are other costs to financialization. Figure 1.14 shows expected next-period household utility conditional on the level of inventory. On a conditional basis, financialization actually reduces expected household utility. Although futures allow households to hedge commodity price risk, they also make spot prices more volatile. The risk premium the long futures position equilibrates to almost zero after financialization, so households often expect negative excess returns after transaction costs. The net affect is a reduction in expected utility *given* a level of inventory. However, households have higher expected utility when inventory is high, both before and after financialization. Since

¹⁷For example, on April 17, 2012 the Wall Street Journal quotes President Obama as saying that “Rising gas prices means a rough ride for a lot of families... we can’t afford a situation where speculators artificially manipulate markets by buying up oil, creating the perception of a shortage and driving prices higher, only to flip the oil for a quick profit.” He suggests “[giving the] CFTC authority to raise margin requirements for oil futures traders.” Although the criticism implies market manipulation rather than mere speculation, the effect of his suggestion would be to reduce household access to futures markets by constraining intermediaries.

financialization also increases inventory, in a typical period the household's position on the utility curve is shifted to the right on the curve after financialization.

To answer whether households benefit or suffer from financialization I adopt the household's unconditional mean utility as a metric. The comparison is between a world in which financial innovation never occurs, and one in which it occurs and remains in place indefinitely. On this basis, households benefit slightly from financialization: they would be willing to pay around 0.02% of their endowment indefinitely to move to a world with financial innovation (with $\tau = 0.08$). Given the baseline parameters, increasing financialization (decreasing τ) generates increasing, but modest, benefits to households.

In a model with multiple goods and incomplete markets, it is not a foregone conclusion that moving markets closer to completeness leads to welfare improvement or greater efficiency. Hart (1975) establishes that, when there are multiple goods, opening new markets may make all agents worse off, so long as markets remain incomplete, as they do here. Examining the effects of financialization piecewise, it is not even obvious that the increase in storage caused by financialization will improve household welfare. Wright and Williams (1984) illustrate that the welfare impact of storage is sensitive to its costs, and to the elasticity of production and consumption. If storage costs were modeled via depreciation rather than a nominal cost, higher mean storage would imply greater depletion of commodity supplies through depreciation, which could lead to welfare losses. In summary, it seems unlikely that households would suffer large welfare losses due to financialization without substantially changing the assumptions of the model, but the conclusion that benefits outweigh costs is a fragile one.

1.6 Conclusion

I construct a model of storable commodities with producers, dealers, households, and active futures markets. I use the model to study how the financialization of commodities impacts spot and futures prices. When calibrated to oil markets, the model implies that financial innovation has essentially no impact on mean spot prices, and reduces the frequency of stockouts. But when stockouts occur, spot prices soar even higher than before financialization, implying higher spot price volatility due to financial innovation. In addition, the risk premium to holding a long position in futures decreases substantially. Although the reduced risk premium and increased volatility is undesirable to households, they benefit from increased oil production and higher

mean inventory. On balance, household welfare increases slightly due to financialization.

1.7 Additional Discussion and Extensions

This section offers additional analysis of the data and the model, and relaxes some assumptions via extensions to the baseline model.

1.7.1 Open Interest in Crude Oil Derivatives, and the Spot Price

I review the data on open interest in crude oil derivatives, and how it relates to the spot price of crude. Open interest in futures is available from the CFTC from 1986, and combined futures and options open interest is available from 1995. Open interest is measured in the number of contracts, for which the underlying asset (in the case of crude oil) is 1000 barrels of oil. I back out options from the combined data, and plot the two series for NYMEX crude oil in Figure 1.15. Both series increase gradually from the beginning of their sample periods, but open interest grows especially quickly in absolute terms starting around 2004. However, the series through 2008 are roughly linear in logs. Both series fall after the financial crisis (around 2008). Futures open interest soon rebounds, whereas open interest in options continues to decline through the present. The reason for apparent swings in preference between futures and options is unclear. Because I model futures and not options, the remaining analysis concerns only futures data.

Figure 1.16 shows open interest in NYMEX crude oil futures with the real spot price. Starting around 2004 (right of the thin red line), notice that both series rise sharply. They also become more correlated in levels. This is the essence of testimony given before the US Congress that raised concern about financialization among elected representatives.¹⁸ The increase in correlation is clearer in Figure 1.17. I split the data between 1990-2003 (top plot), and 2004-2012 (bottom plot), and plot open interest versus spot price. A simple OLS regression suggests no statistically significant relationship between the two variables in the early period, but there is a positive and significant relationship after 2004 (p-value of 0, $R^2 = 0.47$). For this analysis to be statistically valid, one must assume that both open interest and real prices are stationary in the data, which is

¹⁸See for example 2008 testimony of Michael Masters before the US Senate:
<http://www.hsgac.senate.gov/imo/media/doc/052008Masters.pdf>

a questionable assumption. However, the model does offer an explanation for these observations in a setting which, by construction, is stationary in levels.

We have already seen that financialization (decreased costs to consumer trade in futures) dramatically increases open interest in the model. However it also implies an increase in correlation between spot prices and open interest. Figure 1.18 shows expected open interest conditional on the spot price for the pre ($\tau = \infty$) and post ($\tau = 0.08$) financialization model calibrations. Open interest rises sharply with the spot price post-financialization, leveling out beyond a certain price, remarkably similar to the relationship in the data post 2004. Although open interest is also somewhat positively related to the spot price in the pre-financialization model (more so than in the data), the relationship is much weaker. The essential insight from the model, however, is that higher open interest coincides with higher spot prices, but does not cause them. In fact, the origin of the relationship in the model is inventory dynamics. Low inventory coincides with high and volatile spot prices, to which market participants respond by trading more futures. The futures market response is more dramatic once consumers are allowed to participate.

1.7.2 Additional Shocks

Section 1.4 explains the decision to focus on the productivity of the oil sector as the driving force behind oil price fluctuations. However additional sources of randomness affect oil prices. Strong output by the non-oil sector will increase demand for complementary energy goods such as oil. To model this effect, consider a household with a stochastic numeraire endowment, x_t , that is a Markov process. In a given period the household now has net income

$$\hat{x}_t = x_t + \varphi_{t-1}(s_t - f_{t-1}) - \tau f_{t-1} \varphi_{t-1}^2. \quad (1.29)$$

The model in the main paper is a special case of the above with $x_t = 1$ for all t .

Given the interpretation (and calibration) of the household as a US Consumer, a third source of risk might be foreign demand. A simple way of modeling this without explicitly introducing a fourth agent is to add a third Markov process \hat{a}_t , such that goods market clearing now requires

$$\hat{a}_t y_t - c_{y,t} = Q_t - Q_{t-1} = \Delta Q_t. \quad (1.30)$$

Fluctuations in \hat{a}_t will reduce or possibly increase the amount of oil available for US consumption.¹⁹ The model in the main text is a special case with $\hat{a}_t = 1$. One could restrict $\hat{a}_t < 1 \forall t$, so that foreign demand strictly reduces available supply. Or one could consider foreign nations that are on average self sufficient ($E[\hat{a}_t] = 1$), but may have net imports or exports from time to time, so that \hat{a}_t represents foreign *net* demand. I adopt the latter approach.

The three shocks have intuitive meaning as production (a_t), domestic non-oil (x_t) and foreign demand (\hat{a}_t) shocks. They also facilitate a decomposition of the effects of quantity versus price risk in the model. In the default calibration driven by a_t only, a low productivity outcome reduces the quantity of oil sold by the producer, but the equilibrium price of oil rises; the price and quantity effects work against each other. Consumers, meanwhile, face oil price risk, but the value of their endowment is constant by assumption. Alternatively, if we consider a model driven only by shocks to x_t , the direct benefit to consumers of a high endowment outcome is offset by a rise in equilibrium oil prices; consumers face offsetting quantity and price effects. Producers, however, are subject only to price risk. Finally we might consider the effects of \hat{a}_t in isolation. The foreign demand shock affects neither the producers' nor consumers' endowments, but it will affect the price of oil, so agents face only price risk.

I analyze quantity versus price risk by assigning the stochastic process used for a_t in the main calibration to each of the shocks in turn. This keeps the amount of intrinsic risk in the oil market roughly the same, but shifts the exposure to quantity and price risk. Table 1.3 shows summary statistics for each source of risk, with and without household participation in the futures market. Statistics from the data are shown at the left. In broad strokes the model results are quite similar across shocks, but differences in risk exposure manifest in the futures market. Moving from column a_t to column \hat{a}_t , we see that producers short more futures when they face only price effects (with \hat{a}_t). They also pay a higher risk premium on those contracts. The reason is that low price states are relatively worse for producers who face only price risk, since they are not accompanied by high oil output. From the standpoint of households, their endowment is uncorrelated with oil prices whether they face production (a_t) or foreign demand (\hat{a}_t) shocks. However the higher risk premium to the long side of the futures contract under stochastic \hat{a}_t leads households to go long more futures (comparing columns with $\tau = 0.08$). Therefore financialization generates a sharper drop in the risk premium

¹⁹This formulation also implicitly defines the dealer as a "US market" participant. The decision to make \hat{a}_t multiplicative rather than additive is primarily technical. An appropriately scaled additive shock to aggregate oil supply would have similar consequences for the model.

and frequency of backwardation than in the baseline model.

When the model is driven by numeraire (non-oil) consumer endowment shocks (x_t), producers again face only price risk, so they short futures heavily. But households have a reduced incentive to buy futures: the spot price is high when x_t is large, so buying futures hedges oil price risk but amplifies numeraire endowment risk. Under these circumstances it is not obvious that households would choose a long position in oil futures at all, since a short position would insure against bad endowment outcomes.²⁰ However for a sufficiently large risk premium households will take a long position. Indeed that is what occurs: the equilibrium risk premium remains substantial (1.2 % per quarter) with household participation ($\tau = 0.08$) in futures, and households take a more moderate long position than in scenarios with shocks to oil availability.

This constitutes a simple analysis of the multi-shock model that illustrates the shocks' effects in isolation. A combination of shocks could improve the empirical performance of the model, perhaps generating more realistic conditional asset price moments. For example, price volatility conditional on the slope of the futures curve is unrealistically low in the baseline calibration, because the slope reveals a great deal of information about the state of the model. With multiple sources of randomness the slope of the futures curve might be less revealing. The importance of one source of risk versus another may also vary across commodity markets. These may be promising directions for future work.

1.7.3 Intermediate Production

The previous section considered shocks that complicate the connection between commodity production and the market price of a commodity. Another factor that intercedes in commodity markets is intermediate production. Relatively little raw oil is consumed; rather it is refined into fuels, waxes and lubricants, or burned to produce electricity. This may alter the analysis for two main reasons: time-variation in the efficiency of intermediate production, and a delay between use of the raw commodity as an input and production of the consumable good.

If the efficiency of intermediate production is unpredictable, then it is possible for the spot price of oil to

²⁰This is one of the arguments for shutting down numeraire shocks in the baseline calibration. Financial assets such as equity derivatives or bonds might be more appropriate for insuring against numeraire risk than oil futures, but these assets are excluded from consideration. The case $x_t = 1$ is one where numeraire risk is “fully hedged” using assets that are outside of the model.

be unexpectedly low whereas (or perhaps because) the price of gasoline is unexpectedly high. This could be modeled with a shock similar to \hat{a}_t , but applied instead to the net supply of oil, as

$$\hat{b}_t \underbrace{(y_t - \Delta Q_t)}_{\hat{y}_t} = c_{y,t}. \quad (1.31)$$

Imagine the intermediate producer as a competitive risk-neutral agent who simply purchases \hat{y}_t units of oil at spot price s_t , and sells $\hat{b}_t \hat{y}_t$ units of gasoline at spot price \hat{s}_t . A zero-profit condition will apply, and proceeds from gasoline sales will pass through to the oil sector. The situation for the producer is essentially identical to the baseline model with productivity shock a_t , given the stochastic process is suitably altered to incorporate direct and intermediate productivity fluctuations. However households view futures markets differently, because oil futures insure against oil price risk, but their exposure is to gasoline price risk. Therefore a long position in oil futures would not be an ideal hedge, but it would offer a risk-premium. As shown in the previous section, under plausible assumptions households will take a long position to earn a risk premium even if it amplifies their intrinsic risk exposure. Household behavior would likely fall between the example with a_t random and the example with x_t random shown in Table 1.3. Perhaps the most novel aspect of intermediate productivity risk is that the dealer now faces de-facto quantity risk through \hat{b}_t . As with quantity risk to producers and households, high intermediate productivity \hat{b}_t would likely correlate with lower prices for the final good. This would dampen the incentive for speculative storage.

Intermediate production might introduce a delay between use of the raw commodity as an input, and production of the consumable good. Consider the baseline calibration with only a_t random, and for simplicity set $\hat{b}_t = 1$. If refining requires one period, then the aggregate resource constraint for oil becomes

$$y_t - \Delta Q_t = c_{y,t+1}. \quad (1.32)$$

Our risk-neutral producer would again pass gasoline sales proceeds through to the oil sector, subject to financing costs at the risk-free rate, assuming the household pays for gasoline upon delivery. Therefore the lag alone changes very little in the commercial oil sector. Household period utility for $t + 1$ is now observable at time t , although consumption doesn't occur until the next period. Therefore a long one-period oil futures position might appeal to households for its risk premium, but it would neither hedge nor amplify household risk. Households would still be long futures for the risk premium. The overall effects of delayed finished-good production on the model's behavior should be modest.

1.7.4 Active Trade in Multi-period Futures Contracts

This section relaxes the assumption that only the one period (3 month) futures contract is actively traded. The decision to focus on the 3 month contract was in part because it falls in the range of contracts with the most open interest. During the year 2000, 43% of crude oil open interest was in contracts for delivery in 3 months or less, and a further 31% of open interest was in contracts of 3 to 12 months, according to Buyuksahin et al. (2011). Early ETFs also focused on contracts maturing in one year or less. However the term structure of open interest later shifted somewhat toward contracts with greater maturity. By 2008, 37% of open interest was in contracts 3 months or less, and 33 % for 3 to 12 months. Contracts of 3 years or more had risen to 7% of open interest. Therefore I investigate what changes occur in the model when households may trade both 1 and 2 period (3 and 6 month) futures, focusing on how heavily households trade multiperiod contracts.

Following the convention established earlier for the dealer, the period t futures price for an n -period to delivery contract is $f_{t,n}$. Futures positions of the dealer, producer, and household are $\psi_{t,n}$, $\phi_{t,n}$, and $\varphi_{t,n}$, respectively, with $n \in \{1, 2\}$. Period $t + 1$ profits for the dealer and producer, and net income for the household, are

$$\begin{aligned} Q_{t+1}^d &= Q_t(s_{t+1} - (1+r)(s_t + k)) + \psi_{t,1}(s_{t+1} - f_{t,1}) + \psi_{t,2}(f_{t+1,1} - f_{t,2}), \\ Q_{t+1}^p &= s_{t+1}a_{t+1}(1 + I_t^{1/2}) + \phi_{t,1}(s_{t+1} - f_{t,1}) + \phi_{t,2}(f_{t+1,1} - f_{t,2}), \\ \hat{x}_{t+1} &= 1 + \varphi_{t,1}(s_{t+1} - f_{t,1}) - \tau f_{t,1}\varphi_{t,1}^2 + \varphi_{t,2}(f_{t+1,1} - f_{t,2}) - \tau f_{t,2}\varphi_{t,2}^2. \end{aligned} \quad (1.33)$$

Effectively the household is able to buy two “ETFs”, one with a rolling position in the 1-period contract, the other with a rolling position in the 2-period contract (with quarterly rolls). I assume the two ETFs have identical cost parameter τ , although this assumption would be easy to relax.

Leaving the dealer’s optimization problem unchanged but for the introduction of 2-period futures, his first order conditions imply:

$$\begin{aligned} Q_t &= \frac{E_t[s_{t+1}] - (1+r)(s_t + k)}{\rho \text{Var}_t(s_{t+1})} - \psi_{t,1} - \frac{\psi_{t,2} \text{Cov}_t(s_{t+1}, f_{t+1,1})}{\text{Var}_t(s_{t+1})}, \text{ for } Q_t > 0 \\ \psi_{t,1} &= \frac{E_t[s_{t+1}] - f_{t,1}}{\rho \text{Var}_t(s_{t+1})} - Q_t - \frac{\psi_{t,2} \text{Cov}_t(s_{t+1}, f_{t+1,1})}{\text{Var}_t(s_{t+1})} \\ \psi_{t,2} &= \frac{E_t[f_{t+1,1}] - f_{t,2}}{\rho \text{Var}_t(f_{t+1,1})} - \frac{(\psi_{t,1} + Q_t) \text{Cov}_t(s_{t+1}, f_{t+1,1})}{\text{Var}_t(f_{t+1,1})}. \end{aligned} \quad (1.34)$$

Earlier we saw that, when $Q_t > 0$, then the dealer's FOCs imply $E_t[s_{t+1}] - (1+r)(s_t + k) = E_t[s_{t+1}] - f_{t,1}$: buying and storing inventory is equivalent to buying futures. The other agents are not able to store the commodity, so their futures positions are determinate, and the dealer's position in one-period futures is pinned down by market clearing. However with two actively traded futures, a similar result holds for the two-period contract if $Q_t > 0$ and we know $Q_{t+1} > 0$. With a finite state space for a_t , current inventory can be sufficiently high that this is guaranteed to occur. Even if a_t were generalized to include extreme shocks, levels of inventory would accumulate such that a stockout would be very unlikely in the next period. Under these circumstances, the dealer's FOCs imply that all three of his assets have equilibrium prices and payoffs coupled to period t and $t + 1$ spot prices:

$$\begin{aligned} f_{t,2} &= (1+r)(f_{t,1} + k) = (1+r)((1+r)(s_t + k) + k), \\ f_{t+1,1} &= (1+r)(s_{t+1} + k). \end{aligned} \tag{1.35}$$

This poses a technical problem, because not only the dealer but also the producer chooses between redundant assets in high-inventory states. Even if the household's futures positions are determinate (and they are), the positions of his commercial counterparties are indeterminate. Rather than imposing ad-hoc conditions on commercial futures positions in certain states of the world, I restrict the producer to trade only the one-period contract at all times. In addition, I make the simplifying assumption that production intensity $I_t = 0$, which simplifies solving for equilibrium without dramatically changing the model's behavior. (I_t is always small given the calibration.)

In high inventory states, some closed form results for the household's futures positions are available. The household's first order conditions are

$$\begin{aligned} 0 &= E_t \left[\frac{s_{t+1} - f_{t,1} - 2\tau f_{t,1} \varphi_{t,1}}{\hat{x}_{t+1}} \right], \\ 0 &= E_t \left[\frac{f_{t+1,1} - f_{t,2} - 2\tau f_{t,2} \varphi_{t,2}}{\hat{x}_{t+1}} \right]. \end{aligned} \tag{1.36}$$

Conditional on $Q_t > 0$ and $Q_{t+1} > 0$, the FOC for $\varphi_{t,2}$ becomes

$$0 = E_t \left[\frac{s_{t+1} - f_{t,1} - 2\tau(f_{t,1} + k)\varphi_{t,2}}{\hat{x}_{t+1}} \right]. \tag{1.37}$$

This implies that one-and-two-period futures positions relate according to

$$\varphi_{t,2} = \frac{f_{t,1}}{f_{t,1} + k} \varphi_{t,1}. \tag{1.38}$$

Therefore in high inventory states, households will hold a strictly smaller position in the 2-period contract than in the 1-period contract. This relationship is peculiar to the choice of quadratic transaction costs. What would happen if transaction costs were, for example, linear, so that costs are $\tau f_{t,n} \varphi_{t,n}$?²¹ In high inventory states, next-period cash-flows per unit of 1-period and 2-period contracts will be, respectively,

$$\begin{aligned} s_{t+1} - f_{t,1}(1 + \tau) \\ s_{t+1} - f_{t,1}(1 + \tau) - \tau k. \end{aligned} \tag{1.39}$$

Conditional on high inventory in period t , period $t + 1$ cash flows from the 1-period contract are strictly higher than those for the 2-period contract in all states of the world. Therefore households would trade only the 1-period contract under linear transaction costs.

Although these conditional analytical results offer useful intuition, they paint an incomplete picture. Therefore I solve for equilibrium with active trade in 1-and-2-period contracts, imposing $\phi_{t,2} = I_t = 0$ on producers, and with quadratic transaction costs to households. The solution method is similar to that used for the basic model, but it is necessary to separate each iteration into two stages to keep the solution tractable. The key insight is to solve for market clearing futures prices in zero inventory states as a separate fixed-point problem, then to use the zero-inventory futures prices as an input to the next iteration of the dynamic programming problem for inventory policy. Details are available upon request.

Summary statistics are given in Table 1.4, with results from the baseline model with $I_t = 0$ listed for comparison. The household trades the first contract more heavily than the second. Although summary statistics are similar in both models, there is greater open interest in the multi-contract model. This results partly from the quadratic transaction cost faced by the household: splitting trade across two contracts reduces costs, since the sum of the squares is less than the square of the sum. Household futures positions conditional on inventory are shown in Figure 1.19. The plot shows household futures positions in 1-and-2-period contracts, and in 1-period contracts for the original model with only one actively traded contract. Households always trade the 1-period contract more heavily. The spread between the positions is larger in low-inventory states, when a stockout is possible in the next period. When next-period stockout is very unlikely, the ratio of the positions converges to $f_{t,1}/(f_{t,1} + k)$. The household's 1-period contract position is larger in the original model, but the household has a much larger combined long position when trading multiple contracts. This impacts the risk premium, which is shown in Figure 1.20. It plots the risk premium

²¹Assume w.l.o.g. that there is a risk premium to the long side of the contracts, and that households will want long positions.

(expected excess return) on actively traded 1-and-2-period contracts, and for the original model with only one actively traded contract. The risk premium is lower with multiple contracts trading. A larger combined long position in futures by households means the dealer has reduced long exposure to the oil, since he takes the short side of household futures. This leads to a reduction in equilibrium risk premia to the long side of the contracts.

In summary, active trade in multiple contracts is unlikely to change the core conclusions regarding financialization. Depending on how transaction costs scale with the size of the position taken by ETFs, spreading trade across maturities may reduce costs, leading to increased trade by households. This would further reduce the risk premium to being long futures.

1.7.5 Characterization of Risk-aversion

The presence of risk aversion, and the levels of risk-aversion among the agents in the model, is important to the size and direction of the futures risk premium. It also affects how agents trade in futures markets. Interpreting the calibrated parameter choices for risk-aversion is complicated by the different endowments of the agents. In addition, the household has a different preference specification than the other two agents. The household has log utility, a form of constant relative risk aversion (CRRA), whereas the commercial agents have mean-variance preferences, which are closer to constant absolute risk aversion (CARA). The household also consumes a composite good, and the parameters of the goods aggregator affect risk-aversion.

One tool to compare the effective risk-aversion of the three agents is to compute certainty equivalents: if each agent were given a constant (riskless) consumption level, how high would it need to be to provide the same average utility as the equilibrium (risky) consumption stream? Table 1.5 shows certainty equivalents in relative and absolute terms for each agent. In the case of households, results are for aggregated consumption. In absolute terms the agents would pay costs within a reasonable range of each other in exchange for certainty. Households would pay about 1/3 as much as producers, with dealers in the middle, paying around 2/3 what producers would. The fact that producers would pay more than dealers for certainty is related to the inherent riskiness of the producer's endowment. Dealers have no endowment, whereas households have a relatively large and riskless endowment. This makes the second row of the table, giving certainty equivalent consumption relative to expected consumption, less informative. Rather mechanically, households would

pay the smallest fraction of expected consumption for certainty, whereas dealers would pay the most. Since dealers could be given a large (but riskless) numeraire endowment without it altering their equilibrium behavior, the relative statistics should be taken with a grain of salt. The main insight is that oil sector risk introduces little fractional variation into the household's consumption stream.

Although the parameters to the household's goods aggregator are pinned down by the calibration, I make no attempt to determine an appropriate level for household aversion to aggregated consumption risk. Intentionally, there is no corresponding free parameter. Such an exercise would be improper without carefully calibrating the moments of household consumption generally, whereas I focus only on the moments of household oil consumption. Instead, the calibration determines household transaction costs, which fulfill a similar role to adjustable household risk aversion in determining futures positions. Given the motivating topic of financial innovation, adjusting transaction costs seems by far the more sensible approach. However this does leave the appropriate level of household risk-aversion indeterminate, which could be important for welfare analysis. One approach would be to calibrate transaction costs to match average futures positions for a range of household utility functions, then perform the welfare analysis on each of these calibrations to evaluate whether the results are robust. This is left for future work.

1.7.6 Comments on General Equilibrium

I conclude with an informal consideration of how a general equilibrium version of the model might differ from the present setup. There are two major concerns in moving to general equilibrium: clearing the numeraire good market, and firm ownership.

Taken purely as a technical requirement, enforcing numeraire market clearing could be done with minimal changes to the model. As mentioned in the previous section, the dealer could be given a constant numeraire endowment without altering his behavior. The same is true of the producer. Rather than assume a partial equilibrium credit facility with exogenous rate r , one could simply provide large enough endowments that borrowing would not be required, and shut down the bond market entirely. Individual and aggregate numeraire resource constraints would be satisfied without trade in bonds, and the numeraire would only change hands to settle oil-related transactions. Although there might be subtle changes to dealer behavior due to altered financing of inventory purchase, a suitable choice of subjective (time) discount factor for the

dealer, say $1/(1 + r)$, should produce a “general equilibrium” model equivalent to the existing model. This solution completely circumvents the interesting topic of how trade in oil futures might relate to trade in bond markets. Fortunately, Chapter 2 of this dissertation takes up this question.

Firm ownership is a more central issue in the current chapter. Although they are commercial in nature, the firms in this model are agents, not legal entities owned by households. Obviously this is a simplifying assumption. If the firms in question are owned by a distinct set of households who have limited direct interaction with the oil consuming households, then the “firms as agents” characterization sacrifices only an unnecessary layer of complexity. Several large oil producers are government controlled enterprises in developing countries. Maximizing shareholder value for oil-consuming households in developed countries is not their objective. Therefore the producers-as-agents assumption has some merit. The amalgam of financial intermediaries and commodity merchants represented by the dealers, however, is problematic. Many intermediaries in the oil futures and physical oil markets are western companies owned by oil-consuming western households: they are banks, swaps dealers, divisions of multinational oil companies, etc. If households own the intermediaries, how can they reasonably gain from trading futures with them? General equilibrium needn't force resolution of this question: the canonical complete-markets framework with value-maximizing firms will not explain why firms trade in futures markets as they do, as such trades are a value-neutral proposition. This leaves ample opportunity for further research.

1.8 References

- CFTC (Sept. 2008). *Staff Report on Commodity Swap Dealers and Index Traders with Commission Recommendations*. Tech. rep. Commodity Futures Trading Commission.
- Buyuksahin, B. et al. (2011). *Fundamentals, trader activity and derivative pricing*. Tech. rep. American University.
- Irwin, S.H., D.R. Sanders, R.P. Merrin, et al. (2009). “Devil or angel? The role of speculation in the recent commodity price boom (and bust)”. In: *Journal of Agricultural and Applied Economics* 41.2, pp. 377–391.
- Kilian, L. and R.J. Vigfusson (2012). *Do oil prices help forecast US real GDP? the role of nonlinearities and asymmetries*. Tech. rep. University of Michigan.
- Hamilton, James D. (2008). “Oil and the Macroeconomy”. In: *The New Palgrave Dictionary of Economics*. Ed. by Steven N. Durlauf and Lawrence E. Blume. Palgrave Macmillan.
- Tang, K. and W. Xiong (2011). *Index Investment and Financialization of Commodities*. Tech. rep. Princeton University.
- Williams, JC and BD Wright (1991). *Storage and Commodity Markets*. Cambridge University Press.
- Deaton, Angus and Guy Laroque (Jan. 1992). “On the Behaviour of Commodity Prices”. In: *Review of Economic Studies* 59, pp. 1–23.
- (Oct. 1996). “Competitive Storage and Commodity Price Dynamics”. In: *Journal of Political Economy* 104, pp. 896–923.
- Routledge, Bryan R., Duane J. Seppi, and Chester S. Spatt (June 2000). “Equilibrium Forward Curves for Commodities”. In: *Journal of Finance* 60, pp. 1297–1338.
- Hirshleifer, D. (1990). “Hedging pressure and futures price movements in a general equilibrium model”. In: *Econometrica: Journal of the Econometric Society*, pp. 411–428.
- Singleton, Kenneth J. (2011). “Investor Flows and the 2008 Boom/Bust in Oil Prices”. English. In: *SSRN eLibrary*.
- Bessembinder, Hendrik (. et al. (2012). “Predatory or Sunshine Trading? Evidence from Crude Oil ETF Rolls”. English. In: *SSRN eLibrary*.
- Pan, Xuhui (Nick) (2011). “Investor Beliefs and State Price Densities in the Crude Oil Market”. English. In: *SSRN eLibrary*.
- Hamilton, J.D. and J.C. Wu (2011). *Risk Premia in Crude Oil Futures Prices*. Tech. rep. UCSD.

- Brunetti, C. and D. Reiffen (2011). “Commodity index trading and hedging costs”. In: *Finance and Economics Discussion Series*.
- Irwin, S.H. and D.R. Sanders (2011). “Index funds, financialization, and commodity futures markets”. In: *Applied Economic Perspectives and Policy* 33.1, pp. 1–31.
- Baker, Steven D. and Bryan R. Routledge (2012). *The Price of Oil Risk*. Tech. rep. Carnegie Mellon University.
- Ready, R.C. (2012). *Oil Prices and Long-Run Risk*. Tech. rep. University of Rochester.
- Caballero, R.J., E. Farhi, and P.O. Gourinchas (2008). *Financial crash, commodity prices and global imbalances*. Tech. rep. National Bureau of Economic Research.
- Acharya, V.V., L.A. Lochstoer, and T. Ramadorai (2012). *Limits to arbitrage and hedging: Evidence from commodity markets*. Tech. rep. NYU.
- Gorton, G.B., F. Hayashi, and K.G. Rouwenhorst (2012). “The Fundamentals of Commodity Futures Returns”. In: *Review of Finance*.
- Arseneau, David M. and Sylvain Leduc (2012). *Commodity Price Movements in a General Equilibrium Model of Storage*. Tech. rep. Federal Reserve Board.
- Kogan, L., D. Livdan, and A. Yaron (2009). “Oil futures prices in a production economy with investment constraints”. In: *The Journal of Finance* 64.3, pp. 1345–1375.
- Casassus, J., P. Collin-Dufresne, and B.R. Routledge (2009). *Equilibrium commodity prices with irreversible investment and non-linear technology*. Tech. rep. National Bureau of Economic Research.
- Carlson, M., Z. Khokher, and S. Titman (2007). “Equilibrium exhaustible resource price dynamics”. In: *The Journal of Finance* 62.4, pp. 1663–1703.
- Cafiero, C. et al. (2011). “The empirical relevance of the competitive storage model”. In: *Journal of Econometrics* 162.1, pp. 44–54.
- Floden, M. (2008). “A note on the accuracy of Markov-chain approximations to highly persistent AR (1) processes”. In: *Economics Letters* 99.3, pp. 516–520.
- Tauchen, G. (1986). “Finite state Markov-chain approximations to univariate and vector autoregressions”. In: *Economics letters* 20.2, pp. 177–181.
- Stoll, H.R. and R.E. Whaley (2010). “Commodity index investing and commodity futures prices”. In: *Journal of Applied Finance* 20.1, pp. 7–46.

Hart, O.D. (1975). "On the optimality of equilibrium when the market structure is incomplete". In: *Journal of Economic Theory* 11.3, pp. 418–443.

Wright, B.D. and J.C. Williams (1984). "The welfare effects of the introduction of storage". In: *The Quarterly Journal of Economics* 99.1, pp. 169–192.

1.9 Tables

Table 1.1: **Parameters**

Parameter	Value	Purpose
γ	0.06	oil preference
η	-15	elasticity of substitution over goods
k	0.001	nominal storage cost/unit
r	0.0015	risk-free rate
ρ	60	dealer risk aversion
θ	40	producer risk aversion

Preference and technology parameters.

Table 1.2: **Summary Statistics**

Moment	Data		Model			
	1990-2003	2004-2011	$\tau = \infty$	$\tau = 0.1$	$\tau = 0.08$	$\tau = 0.05$
Mean Spot	1.000	2.571	1.000	1.002	1.000	0.999
Autocorrelation Spot	0.677	0.670	0.652	0.628	0.623	0.603
Std. Dev. Spot (Pct. of Mean)	0.226	0.278	0.412	0.438	0.443	0.453
Autocorrelation Futures	0.721	0.707	0.727	0.725	0.725	0.724
Std. Dev. Futures (Pct. of Mean)	0.225	0.300	0.238	0.264	0.268	0.276
Excess Return Futures	0.020	0.002	0.023	0.006	0.003	-0.002
Percent Backwardation	0.696	0.267	0.443	0.292	0.282	0.221
Mean Household Fut. (Pct. Consumption)	0.200	0.200	0.000	0.174	0.198	0.256
Mean Producer Fut. (Pct. Production)	-0.250		-0.281	-0.380	-0.393	-0.425
Mean Oil Expenditure (Pct. GDP)	0.028		0.032	0.032	0.032	0.032
Uncond. Std. Dev. Oil Production	0.053		0.046	0.046	0.046	0.046
Cond. Std. Dev. Oil Production	0.017		0.038	0.038	0.038	0.038
Autocorrelation Oil Production	0.948		0.555	0.550	0.550	0.548
Uncond. Std. Dev. Oil Consumption	0.091		0.033	0.031	0.031	0.031
Cond. Std. Dev. Oil Consumption	0.024		0.023	0.020	0.020	0.019
Autocorrelation Oil Consumption	0.963		0.708	0.757	0.764	0.776

The table shows summary statistics from data and the model. Asset price statistics are split between pre and post financialization periods. Mean spot prices are given relative to the 1990-2003 mean. Standard deviations are given relative to the mean price of the asset, i.e., the unconditional standard deviation of futures in the 2004-2011 period was roughly 30 % of the mean futures price during that period. All model statistics are computed using equivalent normalizations to those applied to the data. Data on household futures positions before and after financialization is a crude discretization due to limited data - see the main text for details. Data on the producer futures position is based on summary information from Acharya, Lochstoer, and Ramadorai (2012) that spanned the period of financialization. Model results are split into a baseline calibration before financialization ($\tau = \infty$), and results for increasing levels of financialization (decreased transaction costs, given by τ). The baseline calibration is designed to match 1990-2003 asset prices. Calibrations after financialization should be compared against 2004-2011 statistics. Statistics for quantities are from annual data, and so a single sample period is used. The baseline calibration was designed to match statistics for quantities as far as possible.

Table 1.3: **Summary Statistics: Alternative Shocks**

Moment	Data		Model: a_t		Model: \hat{a}_t		Model: x_t	
	'90-'03	'04-'11	$\tau = \infty$	$\tau = 0.08$	$\tau = \infty$	$\tau = 0.08$	$\tau = \infty$	$\tau = 0.08$
Mean Spot	1.000	2.571	1.000	1.000	1.000	1.003	1.000	1.002
Autocorrelation Spot	0.677	0.670	0.652	0.623	0.628	0.587	0.646	0.644
Std. Dev. Spot (Pct. of Mean)	0.226	0.278	0.412	0.443	0.421	0.472	0.397	0.414
Autocorrelation Futures	0.721	0.707	0.727	0.725	0.720	0.716	0.728	0.728
Std. Dev. Futures (Pct. of Mean)	0.225	0.300	0.238	0.268	0.225	0.267	0.229	0.254
Excess Return Futures	0.020	0.002	0.023	0.003	0.037	0.010	0.029	0.012
Percent Backwardation	0.696	0.267	0.443	0.282	0.489	0.232	0.478	0.410
Mean Household Fut. (Pct. Cons.)	0.000	0.200	0.000	0.198	0.000	0.264	0.000	0.192
Mean Producer Fut. (Pct. Prod.)	-0.250		-0.281	-0.393	-0.372	-0.516	-0.372	-0.481
Mean Oil Expenditure (Pct. GDP)	0.028		0.032	0.032	0.032	0.032	0.032	0.032
Uncond. Std. Dev. Oil Prod.	0.053		0.046	0.046	0.003	0.004	0.003	0.004
Cond. Std. Dev. Oil Prod.	0.017		0.038	0.038	0.002	0.003	0.002	0.003
Autocorrelation Oil Prod.	0.948		0.555	0.550	0.720	0.716	0.728	0.728
Uncond. Std. Dev. Oil Cons.	0.091		0.033	0.031	0.034	0.031	0.024	0.025
Cond. Std. Dev. Oil Cons.	0.024		0.023	0.020	0.025	0.020	0.022	0.023
Autocorrelation Oil Cons.	0.963		0.708	0.764	0.684	0.764	0.407	0.400

The table shows summary statistics from data, and for the model with different sources of randomness. Results are for oil production (a_t), domestic non-oil (x_t) and foreign demand (\hat{a}_t) shocks. I analyze quantity versus price risk by assigning the stochastic process used for a_t in the main calibration to each of the shocks in turn. This keeps the amount of intrinsic risk in the oil market roughly the same, but shifts the exposure to quantity and price risk. In broad strokes the model results are quite similar across shocks, but differences in risk exposure manifest in the futures market. Moving from column a_t to column \hat{a}_t , we see that producers short more futures when they face only price effects (with \hat{a}_t). They also pay a higher risk premium on those contracts. The reason is that low price states are relatively worse for producers who face only price risk, since they are not accompanied by high oil output. From the standpoint of households, their endowment is uncorrelated with oil prices whether they face production (a_t) or foreign demand (\hat{a}_t) shocks. However the higher risk premium to the long side of the futures contract under stochastic \hat{a}_t leads households to go long more futures (comparing columns with $\tau = 0.08$). Therefore financialization generates a sharper drop in the risk premium and frequency of backwardation than in the baseline model.

Table 1.4: **Summary Statistics: Multiperiod Futures**

Moment	Baseline	Multiple Contracts
Autocorrelation Spot	0.651	0.652
Std. Dev. Spot (Pct. of Mean)	0.445	0.442
Autocorrelation Futures	0.740	0.738
Std. Dev. Futures (Pct. of Mean)	0.278	0.280
Excess Return Futures	0.006	0.001
Percent Backwardation	0.264	0.204
Mean Household Fut. (Pct. Cons.)	0.199, 0.000	0.178, 0.110
Mean Producer Fut. (Pct. Prod.)	-0.379	-0.416
Mean Oil Expenditure (Pct. GDP)	0.037	0.037
Uncond. Std. Dev. Oil Prod.	0.048	0.047
Cond. Std. Dev. Oil Prod.	0.038	0.038
Autocorrelation Oil Prod.	0.595	0.593
Uncond. Std. Dev. Oil Cons.	0.032	0.032
Cond. Std. Dev. Oil Cons.	0.020	0.019
Autocorrelation Oil Cons.	0.794	0.801

The table shows summary statistics for the baseline model, and for a model with active trade in one and two period futures contracts. The two period contract is traded between dealers and consumers only. Production intensity is constrained to $I_t = 0$ in both the baseline and multi-contract models. Household futures positions relative to consumption of the commodity are formatted as (1-period contract, 2-period contract). The household trades the first contract more heavily than the second. Although summary statistics are similar in both models, there is greater open interest in the multi-contract model. This results partly from the quadratic transaction cost faced by the household: splitting trade across two contracts reduces costs, since the sum of the squares is less than the square of the sum. We also see that the risk premium for the 1-period contract is slightly lower with multiple contracts trading. (The risk premium is also somewhat lower for the 2-period contract, not shown.) A larger combined long position in futures by households means the dealer has reduced long exposure to the commodity, since he takes the short side of household futures. This leads to a reduction in equilibrium risk premia to the long side of the contracts.


Table 1.5: **Certainty Equivalent Consumption**

	Dealer	Producer	Household
Expected consumption - certainty equivalent	$1.2 * 10^{-4}$	$1.8 * 10^{-4}$	$5.9 * 10^{-5}$
Certainty equiv. as % of expected consumption	49.1	0.589	0.006

Above is a simple characterization of risk aversion in terms of certainty equivalent consumption (or profit) relative to expected consumption (or profit). In the case of households, consumption is aggregated consumption. In absolute terms the agents would pay prices within a reasonable range of each other in exchange for certainty. Households would pay about 1/3 as much as producers, with dealers in the middle. However in relative terms the agents are very different. This is chiefly because the value of their respective endowments is so different.


1.10 Figures

Figure 1.1: Real Spot Prices (WTI Oil)



The plot above shows real spot prices for WTI crude oil over time. Data is from the St. Louis Fed. I focus on the period from 1990-2012, for which futures prices are also available (marked with the solid vertical line). I split the sample into 1990-2003 and 2004-2012 (marked with the dashed vertical line), to investigate periods before and after financialization of crude oil futures markets.


Figure 1.2: Oil Consumption and Production



Annual US oil consumption is shown alongside worldwide and US production. I normalize each series by its initial value.

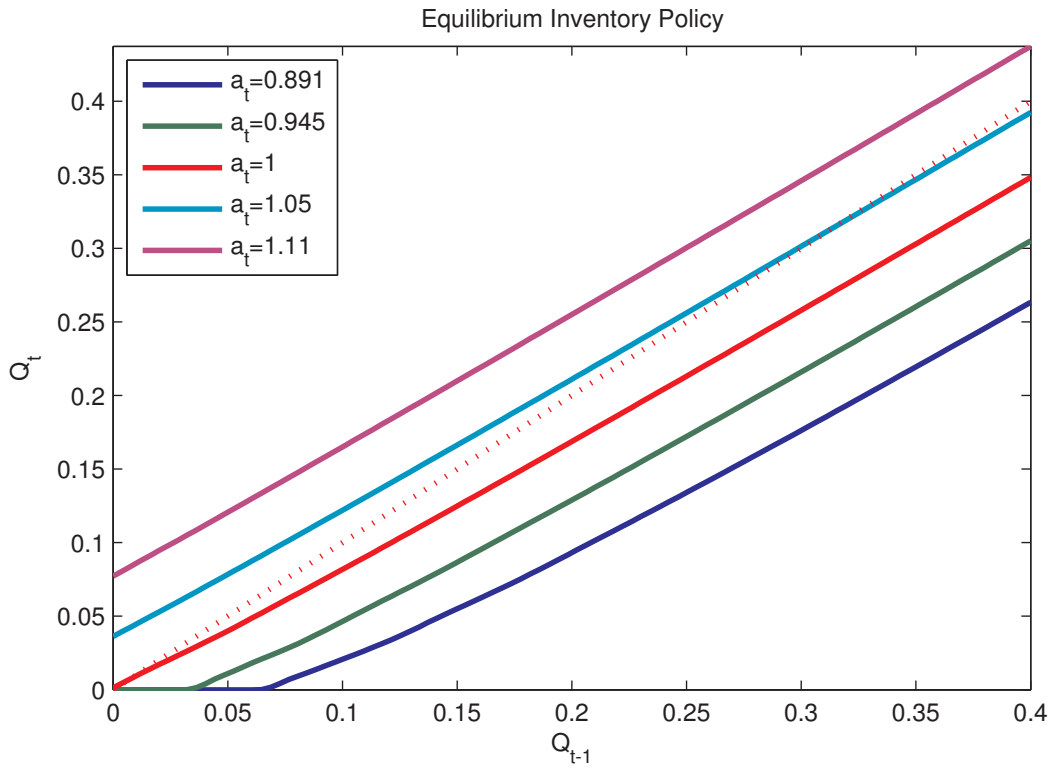
Because oil is a globally traded commodity, the path of US consumption follows global production, not US production.

Figure 1.3: Detrended Oil Consumption and Production




The figure shows detrended annual US consumption and global production. For each series I estimate a linear time trend via OLS, then divide each series by the trend. The two detrended series are very correlated (Pearson's coefficient of 0.85).

Figure 1.4: Inventory Policy




The plot shows optimal inventory policy for the baseline model. The dealer's end of period inventory Q_t depends on inventory carried over from the previous period (Q_{t-1} , given on the x-axis), and the productivity realization (a_t , shown by the different curves). Inventory policy is also affected by lagged productivity, a_{t-1} , as it affects investment; however the effect is modest. Results above are for $a_{t-1} = 1$. The 45° line is shown in dashed red, with the region above that line indicating inventory accumulation and the one below indicating sell-offs.

Figure 1.5: State-contingent Spot Prices




For the baseline calibration, spot prices are shown for entering inventories (Q_{t-1} , given on the x-axis), and each possible productivity realization (a_t , shown by the different curves). If there's plenty of inventory, enough will be sold in response to poor productivity that the price impact of the shock will be modest. If entering inventory is low, however, even selling the entire stock may not reduce the spot price to typical levels. The slope of the curve for low inventories is mostly governed by the household's elasticity of substitution; based on observed prices and consumption, demand for oil is fairly inelastic, so the slope is steep. However prices also reflect how much inventory the dealer is willing to sell in response to increased prices, and how much he will retain against the possibility of another period of low productivity tomorrow - which of course would mean more profits for him! The trade-off is governed by the cost of storage, risk aversion, and opportunities in the futures market.

Figure 1.6: Mean Futures Curve: 1990-2003




The plot shows the mean futures curve before financialization, in the data (1990-2003) and the model ($\tau = \infty$). Each curve is normalized by the price of its front contract. The model provides a good approximation to the shape of the mean futures curve: the relationship between futures of different maturities in the model is similar to that in the data.

Figure 1.7: Term Structure of Futures Volatility: 1990-2003




The plot shows the unconditional quarterly standard deviation of futures contracts before financialization, in the data (1990-2003) and the model ($\tau = \infty$). Standard deviations are expressed as percentages of the mean price of the front contract. The model matches the standard deviation of the front (3-month) contract and the shape of the term structure, but underestimates the volatility of long-dated contracts. This may reflect that fact that productivity is more persistent in the data than in the model, or that the simple mean-variance preferences of the dealer and producer are inadequate substitutes for pricing kernel dynamics that cause volatility at long horizons.

Figure 1.8: Expected Excess Returns on Futures: 1990-2003




The plot shows the unconditional expected excess quarterly holding returns (the risk premium) on futures contracts before financialization, in the data (1990-2003) and the model ($\tau = \infty$). The model does well at matching the risk premium of the front (3-month) contract, and approximates the slope of the curve for longer contracts. However the model fails to reproduce the “hump” between the 3-month and 6-month contracts. As a consequence, the model implies lower risk premia for long dated contracts than seen in the data.

Figure 1.9: Dealer Net Futures Position




Above we see the dealer's net position in the futures market before ($\tau = \infty$) and after ($\tau = 0.08$) financialization, plotted against inventory (Q_t) on the x-axis. Before financialization, the dealer is the only counterparty to the producer, who has a lot of exposure to oil spot price fluctuations. When inventory is low, the dealer has little exposure to oil, and is willing to take on some of the producer's risk for a premium: the dealer goes long, the producer short. When inventory is very high, the dealer is so exposed to spot price risk that he sells short to the producer: the "hedger" and "speculator" depends upon the relative exposure of dealers and producers to oil, which varies over time. After financial innovation, households take a long position in the futures market, to hedge their exposure as consumers of oil. The dealer, who intermediates, now has reduced net exposure to futures. He can take two actions: increase his oil exposure by increasing inventory, or increase his oil exposure by buying more futures from the producer. Either one of these actions will decrease the risk premium, which causes producers to increase their sales of futures, and consumers to decrease their purchases. The dealer's position - the net of producers and households - will increase again. In equilibrium, financialization reduces the net futures position of the dealer at all inventory levels, but because risk premiums drop and producers increase their short position, the magnitude of the drop in the dealer's position is moderated.

Figure 1.10: Distribution of Inventory (Q)




The histogram shows the unconditional distribution of inventory (Q) before and after financialization. Financialization reduces the frequency of very low inventory or stock-out conditions.

Figure 1.11: Financialization and Investment




Financialization increases investment by the commodity producing firm. Taken in isolation, this should increase supply of the commodity relative to the numeraire, and so decrease the spot price.

Figure 1.12: Distribution of Spot Prices




The histogram shows the unconditional distribution of the spot price (s) before and after financialization. The results reflect the offsetting effects of financialization on inventory dynamics and the household budget. Low-inventory states are less likely after financialization, which tends to reduce the level and volatility of spot prices. However price spikes become more extreme after financialization, because the households enjoy a windfall from their futures position precisely when the commodity is in short supply.

Figure 1.13: State-contingent Spot Prices and Financialization



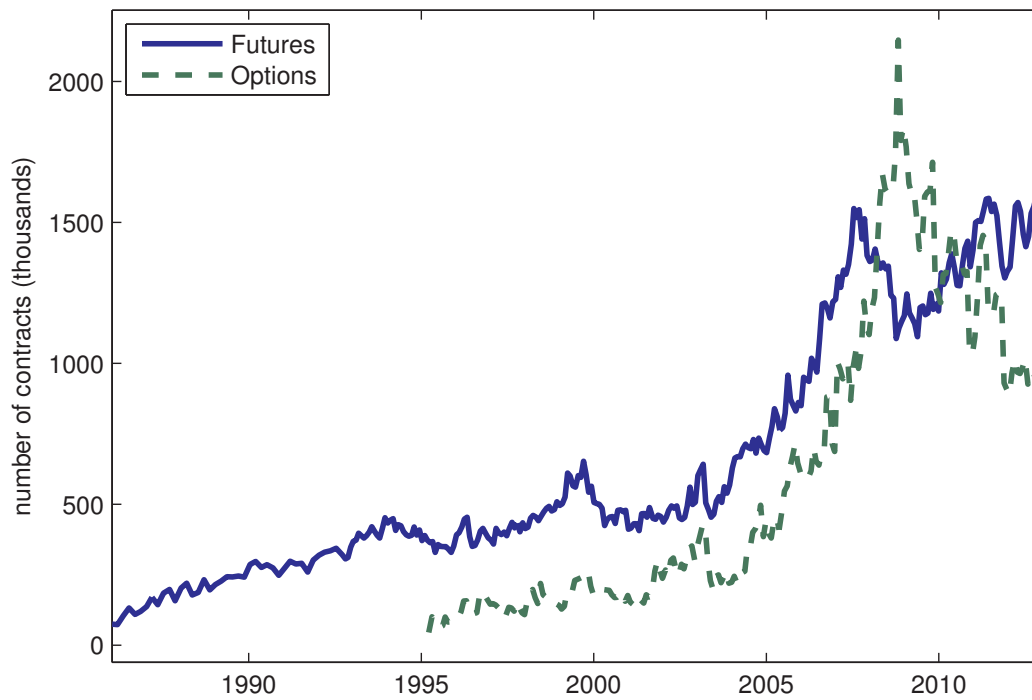
State-contingent spot prices are shown before and after financialization, for the most extreme productivity realizations. Although spot prices in high-productivity states are little changed, prices are elevated in low-productivity states after financialization. Stockouts lead to more extreme price spikes after financialization, although stockouts occur less often.

Figure 1.14: Expected Household Utility




The plot shows expected next-period household utility conditional on the level of inventory. On a conditional basis, financialization reduces household utility. Although futures allow households to hedge commodity price risk, they also make spot prices more volatile. The net affect is a reduction in expected utility *given* a level of inventory. However, households have higher expected utility when inventory is high, both before and after financialization. Since financialization also increases inventory, in a typical period the household's position on the utility curve is shifted to the right on the curve after financialization. On an unconditional basis, households have slightly higher mean utility after financialization, because the inventory smoothing effect dominates the volatility amplification caused by trade in futures.

Figure 1.15: Open Interest in Crude Oil Futures and Options (CFTC)




Open interest in futures is available from the CFTC from 1986, and combined futures and options open interest is available from 1995. Open interest is measured in the number of contracts. I back out options from the combined data, and plot the two series for NYMEX crude oil above. Both series show an upward trend from the beginning of their sample periods, but open interest grows especially quickly in absolute terms starting around 2004. Both series fall after the financial crisis (around 2008). However futures open interest rebounds, whereas open interest in options continues to decline through the present.

Figure 1.16: Crude Oil Open Interest and the Spot Price




The above plot shows open interest in NYMEX crude oil futures (left axis labels) with the real spot price (right axis labels). Open interest is in thousands of contracts, not notional value of contracts. Starting around 2004 (right of the thin red line), notice that both series rise sharply. They also become more correlated in levels.

Figure 1.17: Split Sample: Crude Oil Open Interest versus Spot Price




The above plots shows open interest in NYMEX crude oil futures (y-axis) with the real spot price (x-axis). Open interest is in number of contracts. Data is split between 1990-2003 (top plot), and 2004-2012 (bottom plot). A simple OLS regression suggests no statistically significant relationship between the two variables in the early period, but there is a positive and significant relationship after 2004 (p-value of 0, $R^2 = 0.47$).

Figure 1.18: Split Sample: Crude Oil Open Interest versus Spot Price




Results for open interest conditional on the spot price are shown for the pre ($\tau = \infty$) and post ($\tau = 0.08$) financialization model calibrations. The results resemble the data in two respects: open interest is lower pre-financialization, and open interest rises sharply with the spot price post-financialization, leveling out beyond a certain price. Although open interest is also somewhat positively related to the spot price in the pre-financialization model, the relationship is much weaker.

Figure 1.19: Household Futures Positions: 1-and-2-period Contracts



The plot shows household futures positions in 1-and-2-period contracts, and in 1-period contracts for the original model with only one actively traded contract. Positions are averages over the stationary distribution of commodity production states, conditional on inventory (given on the x-axis). Households always trade the 1-period contract more heavily. The spread between the positions is larger in low-inventory states, when a stockout is possible in the next period. When next-period stockout is very unlikely, the ratio of the positions converges to $f_{t,1}/(f_{t,1} + k)$. The household's 1-period contract position is larger in the original model, but the household has a much larger combined long position when trading multiple contracts.

Figure 1.20: Futures Risk Premium: 1-and-2-period Contracts



The plot shows the risk premium (expected excess return) on actively traded 1-and-2-period contracts, and for the original model with only one actively traded contract. Results are averages over the stationary distribution of commodity production states, conditional on inventory (given on the x-axis). The risk premium is lower with multiple contracts trading. This is because households effectively face reduced transaction costs when able to trade two contracts, so they trade more heavily. A larger combined long position in futures by households means the dealer has reduced long exposure to the commodity, since he takes the short side of household futures. This leads to a reduction in equilibrium risk premia to the long side of the contracts.

Chapter 2

The Price of Oil Risk

with Bryan R. Routledge

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.

2.1 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 2.1). 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. Buyuksahin et al. (2011) 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 Buyuksahin et al. (2011)). 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 heterogeneous 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 “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

¹See for example “The food bubble:How Wall Street starved millions and got away with it” by Frederick Kaufman, Harpers July 2010 <http://arpers.org/archive/2010/07/0083022>

²The 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)

³Often, 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 “convenience 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).

“commercial” or hedging demand was on the short side. This is documented in Buyuksahin et al. (2011) 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 2.2 plots the one-month holding period expected excess returns. At all maturities, you can see from Figure 2.2(a), excess return or risk premium is much higher in the post-2004 sample. In the time series, Figure 2.2(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 heterogeneous 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 needed 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) explana-

tion 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 Gørleanu 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.

2.2 Facts

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

$$r_{t+1}^n = \log \left(\frac{F_{t+1,n-1} - F_{t,n} + (F_{t,n}(\exp r_{t+1}^f))}{F_{t,n}} \right) \quad (2.1)$$

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

log-change in futures prices.

$$r_{t+1}^n - r_{t+1}^f \approx \log F_{t+1,n-1} - \log F_{t,n} \quad (2.2)$$

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 2.2 plots the realized returns for one-month return for a long position in an crude oil future at different maturities. Figure 2.2(a) realized excess return, which is a noisy proxy for risk premium, is much higher in the post-2004 sample. Figure 2.2(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 2.1 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 sloped forward curve is much less in the post-2004 period. Backwardation occurs 68% of the time pre-2004 and only 41% of the months 2004 and beyond (the full sample has a 60% frequency of backwardation). Despite this change, excess returns are higher post 2004. They are higher in both cases where we condition on the sign of the slope of the oil futures curve.⁵

⁴Data 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

⁵Casassus and Collin-Dufresne (2005) use the closely related fact that the negative slope is highly correlated to a high spot price

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 2.2 looks at the one-month excess returns on holding US Treasury bonds over the same time-frame and conditioning information as Table 2.1. 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).

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

And note that if $\phi_{t,n}$ is a constant, then b_n should be one. Table 2.3 confirms that the b_n , particularly at longer horizons is significantly less than one. Interestingly, the coefficient is also broadly decreasing in maturity suggesting more variation in the risk premium associated with the longer-dated oil futures contracts.

of oil.

⁶It 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 2.3, $k = 1$ and we also transform things by log.

There is certainly more work to do here, but we think we have established an interesting fact worth pursuing: 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.

2.3 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.

2.3.1 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

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

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.

$$m_{t+1} = \frac{\partial W_t / \partial x_{t+1}}{\partial W_t / \partial x_t} (\pi_{t+1})^{-1} = \beta \left(\frac{x_{t+1}}{x_t} \right)^{-1} \left(\frac{A_{t+1}}{A_t} \right)^\rho \left(\frac{W_{t+1}}{\mu_t(W_{t+1})} \right)^{\alpha-\rho} \quad (2.3)$$

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,

$$P_t = \frac{\partial W_t / \partial y_t}{\partial W_t / \partial x_t} = \frac{\gamma x_t}{(1 - \gamma) y_t}. \quad (2.4)$$

The pricing kernel 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

$$m_{t+1} = \beta (f_{t+1})^{-1} (A(f_{t+1}, g_{t+1}))^\rho \left(\frac{A(f_{t+1}, g_{t+1}) \hat{W}_{t+1}}{\mu_t (A(f_{t+1}, g_{t+1}) \hat{W}_{t+1})} \right)^{\alpha - \rho}. \quad (2.5)$$

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).

2.3.2 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 heterogeneity 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 $c_{1,t}^x + c_{2,t}^x = x_t$ and $c_{1,t}^y + c_{2,t}^y = y_t$. Agent one derives utility from consumption of the aggregated good $A_1(c_{1,t}^x, c_{1,t}^y) = (c_{1,t}^x)^{1-\gamma_1} (c_{1,t}^y)^{\gamma_1}$. The utility from the stochastic stream of this aggregated good has the same recursive form as above.

$$\begin{aligned} W_t &= W(c_{1,t}^x, c_{1,t}^y, W_{t+1}) = \left[(1-\beta)A_1(c_{1,t}^x, c_{1,t}^y)^{\rho_1} + \beta\mu_{1,t}(W_{t+1})^{\rho_1} \right]^{1/\rho_1}, \\ \mu_{1,t}(W_{t+1}) &= E_t \left[W_{t+1}^{\alpha_1} \right]^{1/\alpha_1}, \end{aligned} \quad (2.6)$$

Agent 2 has similar preference structure with $A_2(c_{2,t}^x, c_{2,t}^y) = (c_{2,t}^x)^{2-\gamma_2} (c_{2,t}^y)^{\gamma_2}$ and recursive preferences

$$\begin{aligned} V_t &= V(c_{2,t}^x, c_{2,t}^y, V_{t+1}) = \left[(1-\beta)A_2(c_{2,t}^x, c_{2,t}^y)^{\rho_2} + \beta\mu_{2,t}(V_{t+1})^{\rho_2} \right]^{1/\rho_2}, \\ \mu_{2,t}(V_{t+1}) &= E_t \left[V_{t+1}^{\alpha_2} \right]^{1/\alpha_2}. \end{aligned} \quad (2.7)$$

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 $\{c_{1,t}^x, c_{1,t}^y, c_{2,t}^x, c_{2,t}^y\}$ that maximizes the weighted average of date-0 utilities subject to the aggregate resource constraint which binds at each date and state:

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

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

⁷Differing 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)

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

$$J(x_t, y_t, V_t) = \max_{c_{1,t}^x, c_{1,t}^y, V_{t+1}} \left[(1 - \beta) A_1 (c_{1,t}^x, c_{1,t}^y)^{\rho_1} + \beta \mu_{1,t} (J(x_{t+1}, y_{t+1}, V_{t+1}))^{\rho_1} \right]^{1/\rho_1} \quad (2.8)$$

s.t. $V(x_t - c_{1,t}^x, y_t - c_{1,t}^y, V_{t+1}) \geq V_t$.

The optimal policy involves choosing agent one's date- t consumption, $c_{1,t}^x$, $c_{1,t}^y$, 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 $W_t = W(c_{1,t}^x, c_{1,t}^y, W_{t+1}) = J(x_t, y_t, V_t)$ and $V_t = V(x_t - c_{1,t}^x, y_t - c_{1,t}^y, V_{t+1})$. The first order and envelope conditions for the maximization problem with date- t -dependent Lagrange multiplier λ_t ⁸ are

$$W_x(c_{1,t}^x, c_{1,t}^y, W_{t+1}) = \lambda_t V_x(x_t - c_{1,t}^x, y_t - c_{1,t}^y, V_{t+1}) \quad (2.9)$$

$$W_y(c_{1,t}^x, c_{1,t}^y, W_{t+1}) = \lambda_t V_y(x_t - c_{1,t}^x, y_t - c_{1,t}^y, V_{t+1}) \quad (2.10)$$

$$W_V(c_{1,t}^x, c_{1,t}^y, W_{t+1}) = \lambda_t V_V(x_t - c_{1,t}^x, y_t - c_{1,t}^y, V_{t+1}) \quad (2.11)$$

$$J_V(x_t, y_t, V_t) = -\lambda_t, \quad (2.12)$$

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

$$\begin{aligned} m_{t+1} &= \beta \left(\frac{c_{1,t+1}^x}{c_{1,t}^x} \right)^{-1} \left(\frac{A_{1,t+1}}{A_{1,t}} \right)^{\rho_1} \left(\frac{W_{t+1}}{\mu_{1,t}(W_{t+1})} \right)^{\alpha_1 - \rho_1} \\ &= \beta \left(\frac{c_{2,t+1}^x}{c_{2,t}^x} \right)^{-1} \left(\frac{A_{2,t+1}}{A_{2,t}} \right)^{\rho_2} \left(\frac{V_{t+1}}{\mu_{2,t}(V_{t+1})} \right)^{\alpha_2 - \rho_2}. \end{aligned} \quad (2.13)$$

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.

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

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:

$$P_t = \frac{\gamma_1 c_{1,t}^x}{(1 - \gamma_1) c_{1,t}^y} = \frac{\gamma_2 c_{2,t}^x}{(1 - \gamma_2) c_{2,t}^y}. \quad (2.14)$$

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

$$\begin{aligned} \hat{c}_{1,t}^x &= \frac{c_{1,t}^x}{x_t} & , & & \hat{c}_{2,t}^x &= \frac{c_{2,t}^x}{x_t} = 1 - \hat{c}_{1,t}^x \\ \hat{c}_{1,t}^y &= \frac{c_{1,t}^y}{y_t} & , & & \hat{c}_{2,t}^y &= \frac{c_{2,t}^y}{y_t} = 1 - \hat{c}_{1,t}^y \end{aligned}$$

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 (2.13), and we can state the pricing kernel as

$$m_{t+1} = \beta \left(\frac{f_{t+1} \hat{c}_{1,t+1}^x}{\hat{c}_{1,t}^x} \right)^{-1} \left(\frac{A_1(f_{t+1} \hat{c}_{1,t+1}^x, g_{t+1} \hat{c}_{1,t+1}^y)}{A_1(\hat{c}_{1,t}^x, \hat{c}_{1,t}^y)} \right)^{\rho_1} \left(\frac{A_1(f_{t+1}, g_{t+1}) \hat{W}_{t+1}}{\mu_1 (A_1(f_{t+1}, g_{t+1}) \hat{W}_{t+1})} \right)^{\alpha_1 - \rho_1} \quad (2.15)$$

(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 expectations) and the future growth state s_{t+1} . Now, in the two agent case, the pricing kernel 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.

2.3.3 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 $C_{1,t}^x$

$$C_{1,t}^x = E_t \left[\sum_{\tau=t}^{\infty} m_{\tau} c_{1,\tau}^x \right]. \quad (2.16)$$

Note, by convention this is the “cum dividend” value including current consumption. To solve this, conjecture that $C_{1,t}^x = \frac{c_{1,t}^x W_t^{\rho_1}}{(1-\beta)A_{1,t}^{\rho_1}}$, and verify. Note it is easier here to use the kernel defined in equation (2.13)

$$\begin{aligned}
C_{1,t}^x &= c_{1,t}^x + \beta E_t \left[\left(\frac{c_{1,t+1}^x}{c_{1,t}^x} \right)^{-1} \left(\frac{A_{1,t+1}}{A_{1,t}} \right)^{\rho_1} \left(\frac{W_{t+1}}{\mu_{1,t}(W_{t+1})} \right)^{\alpha_1 - \rho_1} \frac{c_{1,t+1}^x W_{t+1}^{\rho_1}}{(1-\beta)A_{1,t+1}^{\rho_1}} \right] \\
&= c_{1,t}^x + \frac{\beta c_{1,t}^x \mu_{1,t}(W_{t+1})^{\rho_1 - \alpha_1} E_t [W_{t+1}^{\alpha_1}]}{(1-\beta)A_{1,t}^{\rho_1}} \\
&= \frac{[(1-\beta)A_{1,t}^{\rho_1} + \beta \mu_{1,t}(W_{t+1})^{\rho_1}]}{(1-\beta)A_{1,t}^{\rho_1}} \\
&= \frac{c_{1,t}^x W_t^{\rho_1}}{(1-\beta)A_{1,t}^{\rho_1}} \\
&= \left(\frac{\hat{c}_{1,t}^x \hat{W}_t^{\rho_1}}{(1-\beta)A_1(\hat{c}_{1,t}^x, \hat{c}_{1,t}^y)^{\rho_1}} \right) x_t.
\end{aligned} \tag{2.17}$$

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 $C_{1,t}^x/x_t$ 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.

$$C_{1,t}^y = E_t \left[\sum_{\tau=t}^{\infty} m_{\tau} P_{\tau} c_{1,\tau}^y \right].$$

From the spot price, (2.14), we can write $P_t c_{1,t}^y = \frac{\gamma_1 c_{1,t}^x}{1-\gamma_1}$, so

$$\begin{aligned}
C_{1,t}^y &= E_t \left[\sum_{\tau=t}^{\infty} m_{\tau} c_{1,\tau}^x \frac{\gamma_1}{1-\gamma_1} \right] \\
&= \frac{\gamma_1}{1-\gamma_1} C_{1,t}^x
\end{aligned}$$

Again, note the level of oil does not play a role and the ratio $C_{1,t}^y/x_t$ 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.

$$C_{1,t} = C_{1,t}^x + C_{1,t}^y = \frac{1}{1-\gamma_1} C_{1,t}^x \tag{2.18}$$

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

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

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,

$$C_t^x = C_{1,t}^x + C_{2,t}^x, \quad (2.19)$$

in the oil sector,

$$C_t^y = \frac{\gamma_1 C_{1,t}^x}{1-\gamma_1} + \frac{\gamma_2 C_{2,t}^x}{1-\gamma_2}, \quad (2.20)$$

and the overall wealth in the economy,

$$C_t = \frac{C_{1,t}^x}{1-\gamma_1} + \frac{C_{2,t}^x}{1-\gamma_2}. \quad (2.21)$$

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

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

(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

$$B_{t,n} = E_t[m_{t+1} B_{t+1,n-1}], \quad (2.22)$$

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 $r_{t+1}^n = -n^{-1} \log(B_{t,n})$ (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

$$\begin{aligned} 0 &= E_t[m_{t+1}(F_{t+1,n-1} - F_{t,n})] \\ F_{t,n} &= (B_{t,1})^{-1} E_t[m_{t+1}F_{t+1,n-1}], \\ \log F_{t,n} &= -\log B_{t,1} + \log E_t[\exp\{\log m_{t+1} + \log F_{t+1,n-1}\}], \end{aligned} \tag{2.23}$$

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

2.3.4 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_t^x is the value of a claim to the stream of numeraire good and C_t^y 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 $\phi_{1,t}^x$ shares numeraire and $\phi_{1,t}^y$ 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:

$$\begin{aligned} C_{1,t} &= \phi_{1,t}^x C_t^x + \phi_{1,t}^y C_t^y \\ C_{2,t} &= \phi_{2,t}^x C_t^x + \phi_{2,t}^y C_t^y \end{aligned}$$

Substitute in the definition of the aggregate value of the numeraire sector in equation (2.19) and oil sector in equation 2.20). 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.,

$$\frac{C_{1,t}^y}{C_{1,t}^x} = \frac{\gamma_1}{1 - \gamma_1} \quad , \quad \frac{C_{2,t}^y}{C_{2,t}^x} = \frac{\gamma_2}{1 - \gamma_2}$$

This all implies, for agent one:

$$C_{1,t}^x = \frac{(1 - \gamma_1) \left((-\gamma_2 C_t^x + (1 - \gamma_2) C_t^y) \right)}{\gamma_1 - \gamma_2} \tag{2.24}$$

and

$$\begin{aligned} \phi_{1,t}^x &= \frac{-\gamma_2}{\gamma_1 - \gamma_2} \\ \phi_{1,t}^y &= \frac{1 - \gamma_2}{\gamma_1 - \gamma_2}. \end{aligned} \tag{2.25}$$

(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.

2.4 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 2.4 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 2.5. 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 will 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.

⁹One 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.”

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 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 2.5 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 2.4. Section 2.4.1 shows how \hat{V} governs wealth and consumption sharing between agents. Section 2.4.2 demonstrates the importance of shifts in wealth for asset prices and risk premia, whereas Section 2.4.3 relates these effects to trade in financial markets. Section 2.4.4 discusses the model's 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.

2.4.1 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 2.3 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

¹⁰Although 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. $c_{2,t}^x = x_t$, $c_{2,t}^y = y_t$ 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.

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

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 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.

2.4.2 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 2.4 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 2.4 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 2.4.4. 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 2.5. 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 2.6 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 2.1 we see a negatively sloping oil futures curve implies higher excess returns. Table 2.6 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 2.1. 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 2.8 reproduces the Fama predictive regressions from Table 2.3. 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 2.7 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 2.2 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 2.8, 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.

2.4.3 Portfolios and open interest

As we saw in Section 2.3.3, 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

¹²When 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 $V \rightarrow 0$ or $V \rightarrow 1$.

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, a one-period bond, and “collateralized” 1 and 2 year oil futures contracts (labeled F_1 and F_2 , respectively). Figure 2.9 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 to 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 2.10, 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 2.10 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 ($V \rightarrow 1$), 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.

2.4.4 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 2.11 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 2.12 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 2.12 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 2.10, 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 2.12 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.

2.5 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.

2.6 Additional Discussion and Extensions

2.6.1 Risk Aversion with Recursive Preferences and Two Goods

With Epstein Zin preferences, attitudes toward risk are an amalgam of two parameters: α , associated with static risk-aversion (gambles over states of nature), and ρ , associated with intertemporal variation in utility. To this our model adds an extra layer: γ affects attitudes towards risk over goods. This section characterizes each agent's attitude towards risk generally and in each of the two goods, and relates this to the dynamics of wealth share (proxied by V).

Figure 2.13 shows the mean growth in each agent's aggregated consumption (A) in the left panel, and the standard deviation of aggregated consumption in the right panel. Results are shown conditional on V (essentially agent 2's wealth-share), taking expectations over growth states using their stationary distribution. Although mean and standard deviation are insufficient statistics for characterizing risk-aversion with recursive preferences, they provide some intuition. Based on his Epstein-Zin preference parameters, agent 2 is more tolerant of risk over states (α) and over time (ρ) than agent 1. The intuition from these parameters alone holds despite the complicating factor of heterogeneous goods aggregation (γ). Notice that agent 2's aggregated consumption (A_2) grows with a higher mean but also generally a higher variance than agent 1's. When agent 1 has most of the wealth, agent 2 exposes himself to more risk, in exchange for faster mean growth. When agent 2 is dominant, equilibrium prices adjust so that he is content with the mean and standard deviation of growth that is inherent to the economy (that of the aggregated endowments). Growth in aggregated consumption reflects growth in wealth, which also drifts toward agent 2 more quickly if his initial share is small (see for example Figure 2.11).

The existence of two goods adds an extra layer to risk-sharing. For each agent we see the mean and standard deviation of good X (numeraire) consumption growth in the left column of Figure 2.14, and mean and standard deviation of good Y (oil) consumption growth on the right. Results are again shown conditional on V , but unconditional of growth state. Good X results are similar to those for the aggregated good, as X consumption represents most of the value of total consumption. Good Y consumption is less valuable overall, but agent 2 values consumption of oil about twice as highly as agent 1. Although agent 2 is generally more risk tolerant than agent 1, the standard deviation of agent 2's oil consumption growth is lower than that of agent 1 in most states of the world. By this measure, agent 2 is less tolerant of oil risk than agent 1, even though he is more tolerant of aggregated risk. Therefore a preference for high oil consumption could be an important factor in the risk premium on oil related assets.

2.6.2 Sources of Time-variation in the Risk Premium

Our model is an attempt to understand excess returns, particularly on oil futures. What makes excess returns to a long position so high on average, and what makes them so variable? We also study time-variation in the risk premium, expected excess returns. The former concerns the difference between the return on a fully collateralized futures contract and the risk-free return, whereas the latter studies the difference in differences. The model has two interesting moving parts that drive variation: persistent states of nature s , which are exogenous, and an endogenous variable V , which maps to the wealth share of agent 2. To understand what drives variation in excess returns, we separate the effects of these two variables.

Figure 2.15 illustrates the contribution of the two state variables, growth state s and wealth share V , to the risk premium on 1-period oil futures. Results are shown conditional on current V_t (x-axis), but unconditional of current growth state s_t . We consider three scenarios: allowing only V to vary (by imposing $s_{t+1} = s_t$), allowing only s to vary (by imposing $V_{t+1} = V_t$), and allowing both V and s to vary (the equilibrium scenario). When variation in s_t is “shut down”, it is done in such a way as to eliminate the effects of *persistence* in the growth state, but period $t + 1$ consumption growth is drawn from the usual distribution, as if on an iid basis.¹³

¹³Of course it is possible to shut down growth also. The contribution of persistent growth states to the standard deviation of excess returns is reduced in this case, since growth realizations and the effects of persistence are positively correlated.

The top panels of Figure 2.15 show mean and standard deviation of excess returns. Variation in V alone generates mean excess returns at nearly equilibrium levels, whereas there are lower mean excess returns with only s variation. However V contributes less to the standard deviation of excess returns, much of which is due to s . Together V and s appear to amplify each other, generating slightly higher combined standard deviation than the sum of their independent contributions. The bottom panels illustrate time-variation in the risk premium, which is expected excess returns. Earlier we discussed a drift in wealth-share from agent 1 to agent 2, i.e., a positive drift in V . Because agent 2 is less risk-averse overall, this generates a negative drift in the risk premium, which is seen in the bottom left panel. There is no drift in risk premium due to changes in s . Interestingly s and V interact to produce a subtly different drift together, even though s generates no drift alone. The bottom right panel shows standard deviation in the risk premium, for which s is almost solely responsible. Therefore heterogeneity in preferences contributes only small short-term variability to the risk premium, but is chiefly responsible for trends. Notice that changes in V contribute more to standard deviation when current V is small. That is because the growth of V is also more volatile conditional on low current V . Presumably mechanisms that increase the volatility of V would also increase the volatility of the risk premium.

2.6.3 Improving the Calibration with Transient Growth Shocks

Ongoing work concerns the introduction of transient log-normal growth shocks into the model. Instead of driving growth directly via a Markov chain, consider a Markov switching model where

$$\log \left(\left[\begin{array}{c} \frac{x_{t+1}}{x_t}, \frac{y_{t+1}}{y_t} \end{array} \right] \right) \sim \mathcal{N}(\mu(s), \Omega(s)). \quad (2.26)$$

There is persistence in the mean and variance of growth that is Markovian, but growth realizations are normally distributed and transient. This offers several advantages. Because the solution technique for the original model can only manage a moderate state space for the Markov chain, there is a tension between generating realistic short-term variation, and generating persistence in growth. This tension would grow if we were to move from the current annual calibration to a quarterly calibration. Much of the trade in oil futures is in contracts with less than a year to delivery, so a quarterly calibration would be more realistic. The Markov switching model can deliver persistence in the mean and variance using only a few Markovian states, while generating short-term variation through transient normal shocks, solving this dilemma. Another advantage is that shock realizations will generate variation in V even if the Markov state s does not switch.

As discussed in the previous section, this could generate interesting variation in the risk premium that is specific to the multi-agent setting. Finally, the switching setup makes it easier to disentangle the effects of changing the means versus changing the variance-covariance of growth in the two goods.

Unfortunately the Markov switching setup introduces several difficulties. The optimal policy for promised utility V is now defined on a continuous state space, due to the normal shock. Therefore the policy involves solving for a function, rather than one point for each Markovian state. And taking expectations requires numerical integration, which is much slower than computing expectations over a small discrete state space.

At present we have solved a Markov switching model with multiple Markov states but only a scalar normal shock. That is, transient normal shocks to x and y are either perfectly positively correlated, or perfectly negatively correlated conditional on the Markov state realization s . Current experiments indicate that the optimal promised utility policy is nearly linear in the normal shock realizations, which simplifies implementation. However use of numerical integration makes futures prices extremely slow to compute via a naive recursive method. We are working towards more efficient computation of futures prices, and a 2-d normal shock process. We hope this will lead to a more plausible, yet computationally efficient, quarterly calibration.

2.7 References

- Buyuksahin, B. et al. (2011). *Fundamentals, trader activity and derivative pricing*. Tech. rep. American University.
- Hamilton, J.D. and J.C. Wu (2011). “Risk Premia in Crude Oil Futures Prices”. Working Paper.
- Hamilton, J.D. (2009). “Causes and Consequences of the Oil Shock of 2007-08”. In: *NBER Working Paper*.
- Wright, Brian D. (2011). “The Economics of Grain Price Volatility”. In: *Applied Economic Perspectives and Policy* 33 (1), pp. 32–58.
- Keynes, John M. (1936). *The general theory of employment, interest and money*. London: Macmillan and Co.
- Schwartz, Eduardo S. (July 1997). “The Stochastic Behavior of Commodity Prices: Implications for Valuation and Hedging”. In: *Journal of Finance* 52, pp. 923–973.
- Routledge, Bryan R., Duane J. Seppi, and Chester S. Spatt (June 2000). “Equilibrium Forward Curves for Commodities”. In: *Journal of Finance* 60, pp. 1297–1338.
- (2001). “The Spark Spread: An Equilibrium Model of Cross-Commodity Price Relationships in Electricity”. Carnegie Mellon University Working Paper.
- Campbell, John Y. and John H. Cochrane (Apr. 1999). “By Force of Habit: A Consumption-Based Explanation of Aggregate Stock Market Behavior”. In: *Journal of Political Economy* 107, pp. 205–251.
- Bansal, Ravi and Amir Yaron (2005). “Risks For the Long Run: A Potential Resolution of Asset Pricing Puzzles”. In: *Journal of Finance (forthcoming)*. Wharton Working Paper.
- Routledge, Bryan R. and Stanley E. Zin (Aug. 2010). “Generalized Disappointment Aversion and Asset Prices”. In: *Journal of Finance* 65.
- Cochrane, John H. and Monika Piazzesi (2005). “Bond Risk Premia”. English. In: *The American Economic Review* 95.1,
- Epstein, Larry G. and Stanley E. Zin (July 1989). “Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Framework”. In: *Econometrica* 57.4, pp. 937–969.
- Kreps, David M. and Evan L. Porteus (Jan. 1978). “Temporal Resolution of Uncertainty”. In: *Econometrica* 46, pp. 185–200.
- Kihlstrom, Richard E. and Leonard Jay Mirman (1974). “Risk aversion with many commodities”. In: *Journal of Economic Theory* 8.3, pp. 361–388.

- Backus, David K., Bryan R. Routledge, and Stanley E. Zin (Sept. 2009). “Who Holds Risky Assets?”
Carnegie Mellon Working Paper.
- (Sept. 2008). “Recursive Risk Sharing: Microfoundations for Representative-Agent Asset Pricing”. Carnegie Mellon Working Paper.
- Lucas, Robert E. and Nancy L. Stokey (1984). “Optimal Growth with Many Consumers”. In: *Journal of Economic Theory* 32.1, pp. 139–171.
- Kan, Rui (Aug. 1995). “Structure of Pareto Optima When Agents Have Stochastic Recursive Preferences”. In: *Journal of Economic Theory* 66.2, pp. 626–631.
- Anderson, Evan W. (2005). “The dynamics of risk-sensitive allocations”. In: *Journal of Economic Theory* 125.2, pp. 93–150.
- Colacito, Riccardo and Mariano Massimiliano Croce (Dec. 2011). “Risks for the Long-Run and the Real Exchange Rate”. In: *Journal of Political Economy* 119 (1), pp. 153–181.
- (2012). “International Asset Pricing with Recursive Preferences”. In: *Journal of Finance* forthcoming.
- Gørleanu, N. and S. Panageas (2010). “Young, old, conservative and bold: The implications of heterogeneity and finite lives for asset pricing”. In: *Unpublished Manuscript, Haas School of Business, University of California at Berkeley*.
- Carlson, Murray, Zeigham Khokher, and Sheridan Titman (2007). “Equilibrium Exhaustible Resource Price Dynamics”. In: *Journal of Finance* 62.4, pp. 1663–1703.
- Casassus, Jaime, Pierre Collin-Dufresne, and Bryan Routledge (Apr. 2007). “Equilibrium Commodity Prices with Irreversible Investment and Non-linear Technologies”. Hass Working Paper.
- Kogan, Leonid, Dmitry Livdan, and Amir Yaron (June 2009). “Futures Prices in a Production Economy with Investment Constraints”. In: *Journal of Finance* 64 (3). The Wharton School, University of Pennsylvania, Working Paper, 13451375.
- Casassus, Jaime and Pierre Collin-Dufresne (Oct. 2005). “Stochastic Convenience Yield Implied from Commodity Futures and Interest Rates”. In: *Journal of Finance* 60.5, pp. 2283–2331.
- Hamilton, James D. (2008). “Oil and the Macroeconomy”. In: *The New Palgrave Dictionary of Economics*. Ed. by Steven N. Durlauf and Lawrence E. Blume. Palgrave Macmillan.
- Fama, Eugene F. and Kenneth R. French (1987). “Commodity Futures Prices: Some Evidence on Forecast Power, Premiums and the Theory of Storage”. In: *Journal of Business* 60.1, pp. 55–73.
- Lucas Jr., Robert E. (1978). “Asset Prices in an Exchange Economy”. In: *Econometrica* 46, pp. 1426–14.

- Backus, David K., Bryan R. Routledge, and Stanley E. Zin (2005). “Exotic Preferences for Macroeconomists”. In: *NBER Macroeconomics Annual 2004*. Ed. by Mark Gertler and Kenneth Rogoff. Vol. 19. Cambridge, MA: MIT Press, pp. 319–391.
- Hansen, Lars Peter, John C. Heaton, and Nan Li (2008). “Consumption Strikes Back? Measuring Long-Run Risk”. In: *Journal of Political Economy* 116.2, pp. 260–302.
- Ready, Robert (2010). “Oil Prices and Long-Run Risk”. English. In: *SSRN eLibrary*.
- Yan, H. (2010). “Natural selection in financial markets: Does it work?” In: *Management Science* 54.11, p. 1935.
- Basak, Suleyman and Domenico Cuoco (1998). “An Equilibrium Model with Restricted Stock Market Participation”. English. In: *The Review of Financial Studies* 11.2,
- Basak, S. (2005). “Asset pricing with heterogeneous beliefs”. In: *Journal of Banking and Finance* 29.11, pp. 2849–2881.
- Dumas, B., A. Kurshev, and R. Uppal (2009). “Equilibrium Portfolio Strategies in the Presence of Sentiment Risk and Excess Volatility”. In: *The Journal of Finance* 64.2, pp. 579–629.
- Barro, Robert J. (2009). “Rare Disasters, Asset Prices, and Welfare Costs”. English. In: *The American Economic Review* 99.1,
- Hamilton, J.D. (2011). *Historical oil shocks*. Tech. rep. National Bureau of Economic Research.

2.8 Tables

Table 2.1: 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
Slope +	Horizon	Obs	Mean	St. Dev.	Obs	Mean	St. Dev.	Obs	Mean	St. Dev.
	1	95	-0.61	9.44	53	-0.38	9.16	42	-0.89	9.89
	3	95	-0.17	8.36	53	-0.29	7.65	42	-0.03	9.28
	6	95	-0.03	7.31	53	-0.24	6.17	42	0.23	8.60
	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	
Slope -	Horizon	Obs	Mean	St. Dev.	Obs	Mean	St. Dev.	Obs	Mean	St. Dev.
	1	145	1.52	9.27	114	1.26	9.34	31	2.49	9.12
	3	144	1.84	8.09	113	1.51	7.94	31	3.05	8.65
	6	144	1.80	6.89	113	1.37	6.46	31	3.36	8.21
	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 2.2: Monthly Excess returns on US Treasury Bonds

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

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

Table 2.3: **Predictive regression for crude oil**

Horizon (n)	a		b		R^2	nobs
	a.	t ($a = 0$)	b	t ($b = 1$)		
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:

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

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

Table 2.4: **Aggregate Consumption Growth Process**

$$\begin{aligned} \begin{bmatrix} s \in S \\ f(s) \\ g(s) \end{bmatrix} &= \begin{bmatrix} 1 & 2 & 3 & 4 \\ 0.99 & 1.03 & 1.05 & 0.93 \\ 0.90 & 1.04 & 1.06 & 1.07 \end{bmatrix} \\ \pi &= \begin{bmatrix} 0.80 & 0.10 & 0.05 & 0.05 \\ 0.05 & 0.85 & 0.05 & 0.05 \\ 0.05 & 0.18 & 0.72 & 0.05 \\ 0.05 & 0.05 & 0.63 & 0.27 \end{bmatrix} \\ \bar{\pi} &= [0.20 \quad 0.48 \quad 0.26 \quad 0.06] \end{aligned}$$

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 2.5: **Parameters**

<i>Parameter</i>	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 2.6: **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 2.7: **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 2.8: **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

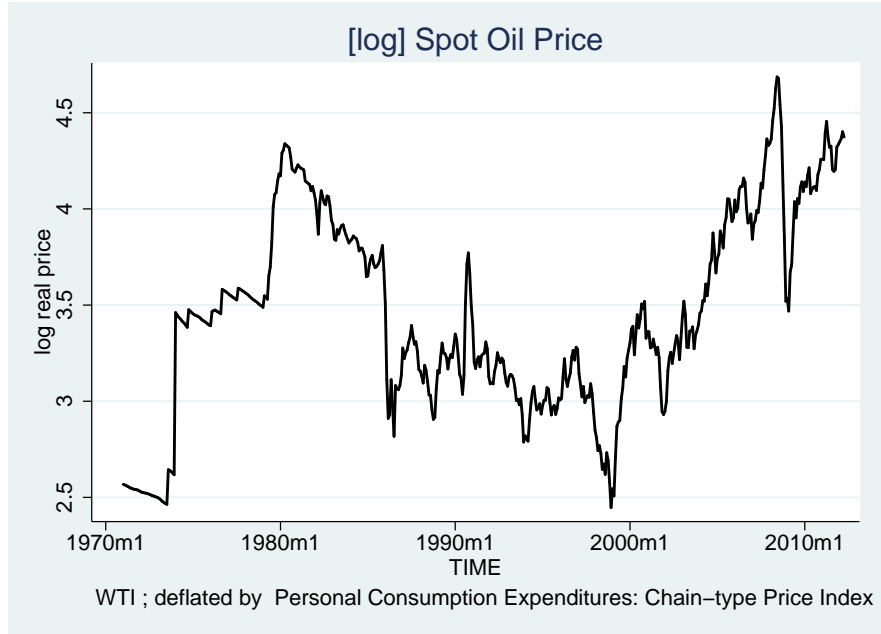
On simulated data from our model, regress:

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


2.9 Figures

Figure 2.1: 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 2.2: Term Structure of Holding Period Returns

a. By Horizon

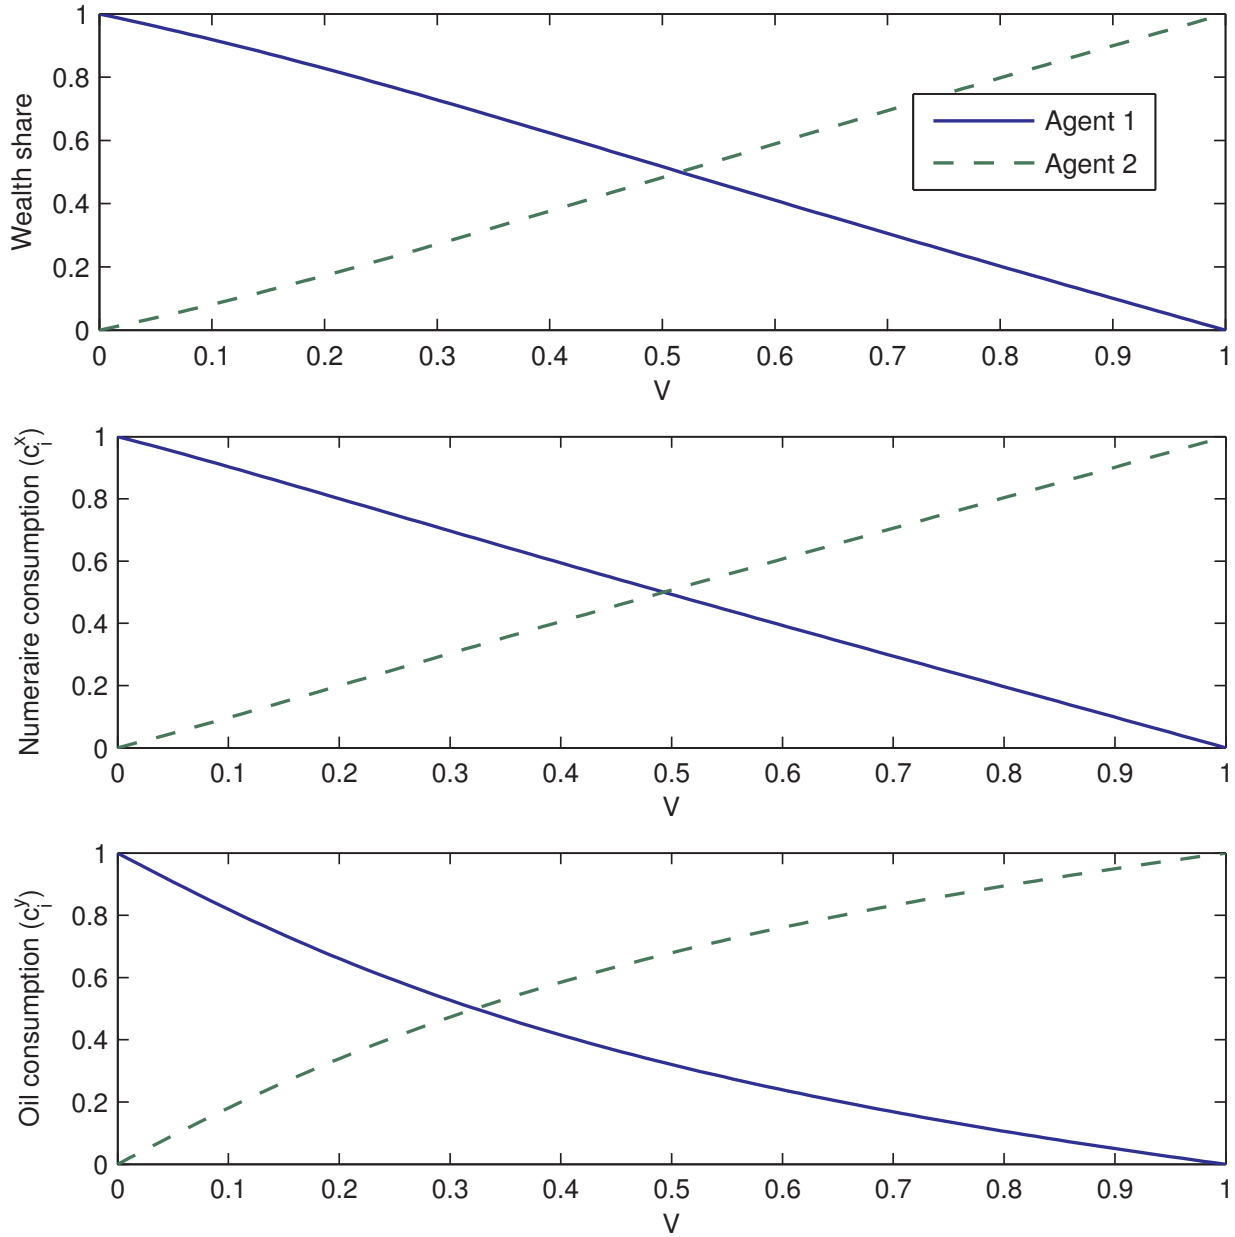


b. Time Series




One month expected holding excess returns on fully-collateralized 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 2.3: 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 2.4: 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 2.5: 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 2.6: 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 2.7: Zero-coupon bond rates



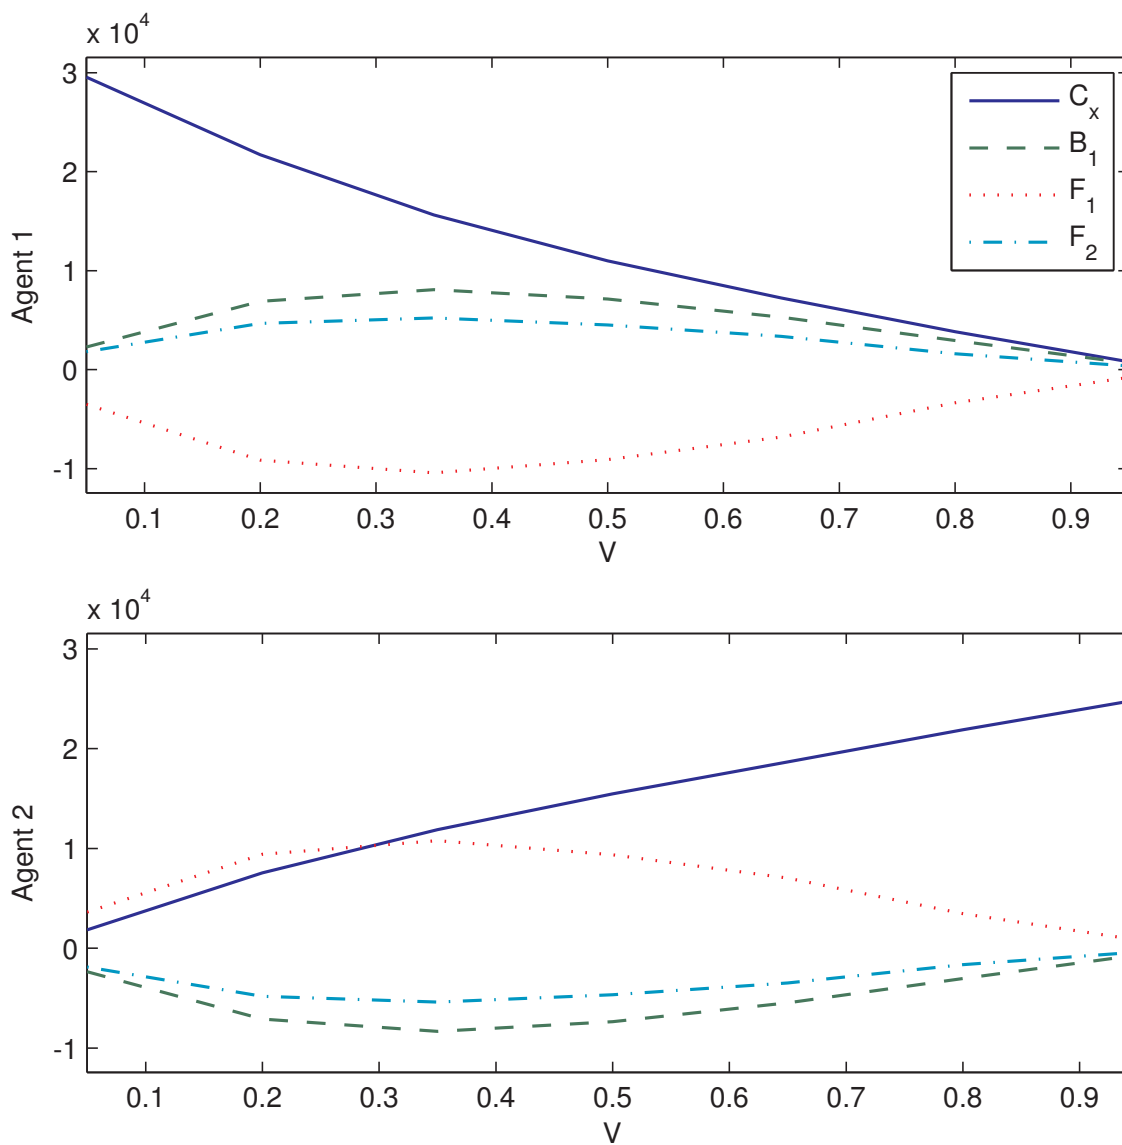
Annual interest rates on zero coupon bonds, averaged over the stationary distribution of growth states. Depending on V , the term structure may be downward sloping (low V) or upward sloping (high V). In addition, the risk-free rate (rate on the 1-year bond) is non-monotonic in V . The interaction of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

Figure 2.8: 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 2.9: 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 2.10: Open interest in oil futures and spot prices



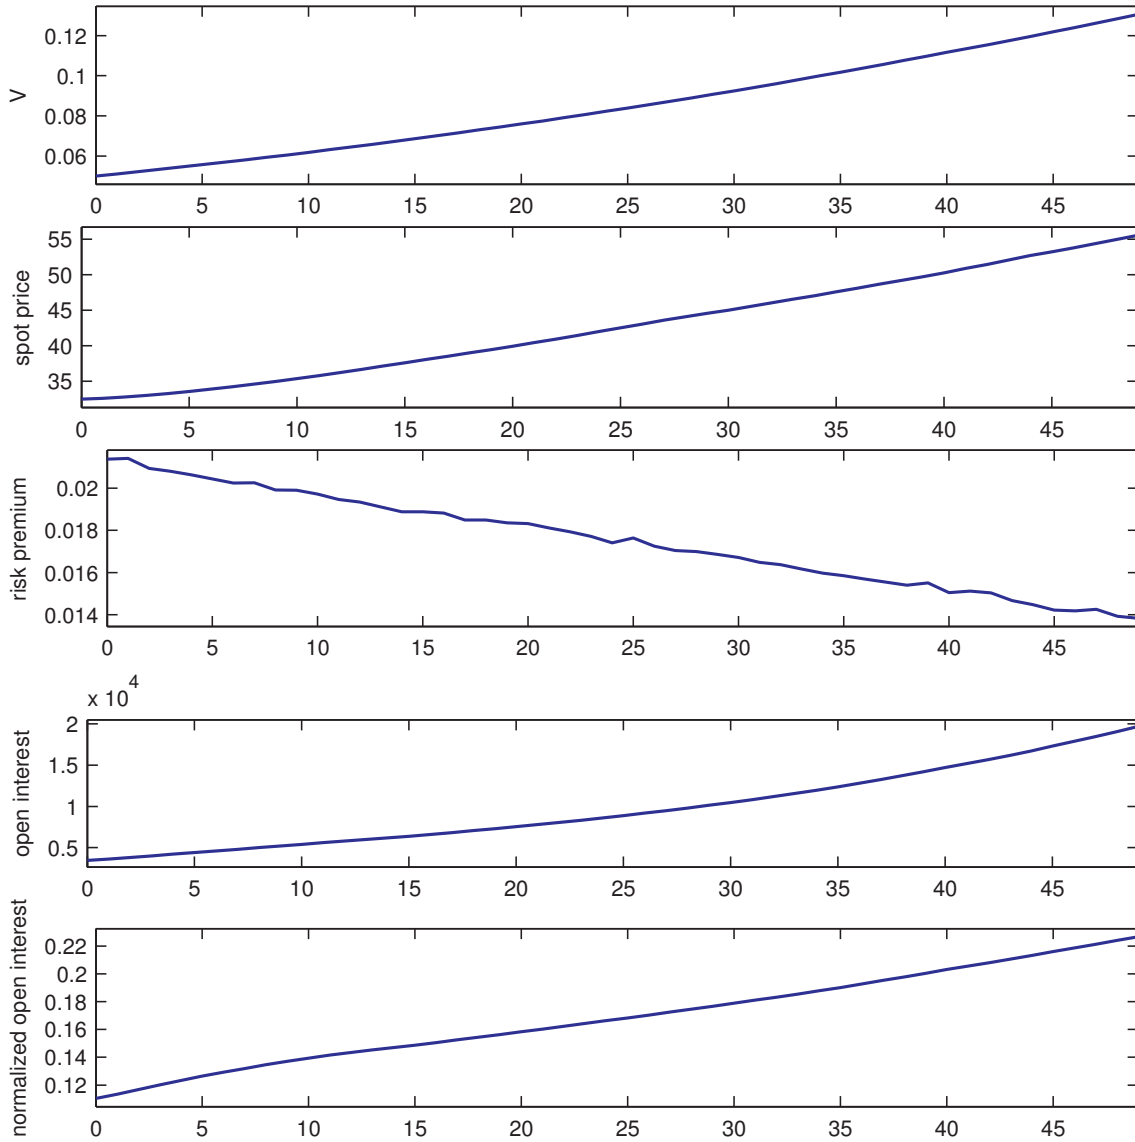
The top panel shows open interest in 1 and 2 year oil futures contracts, expressed as the numeraire value of the contracts. The bottom panel shows spot price of oil for various V . Results are averaged over the stationary distribution of growth states (s). However results conditional on a particular 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 2.11: 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 2.12: Evolution over time



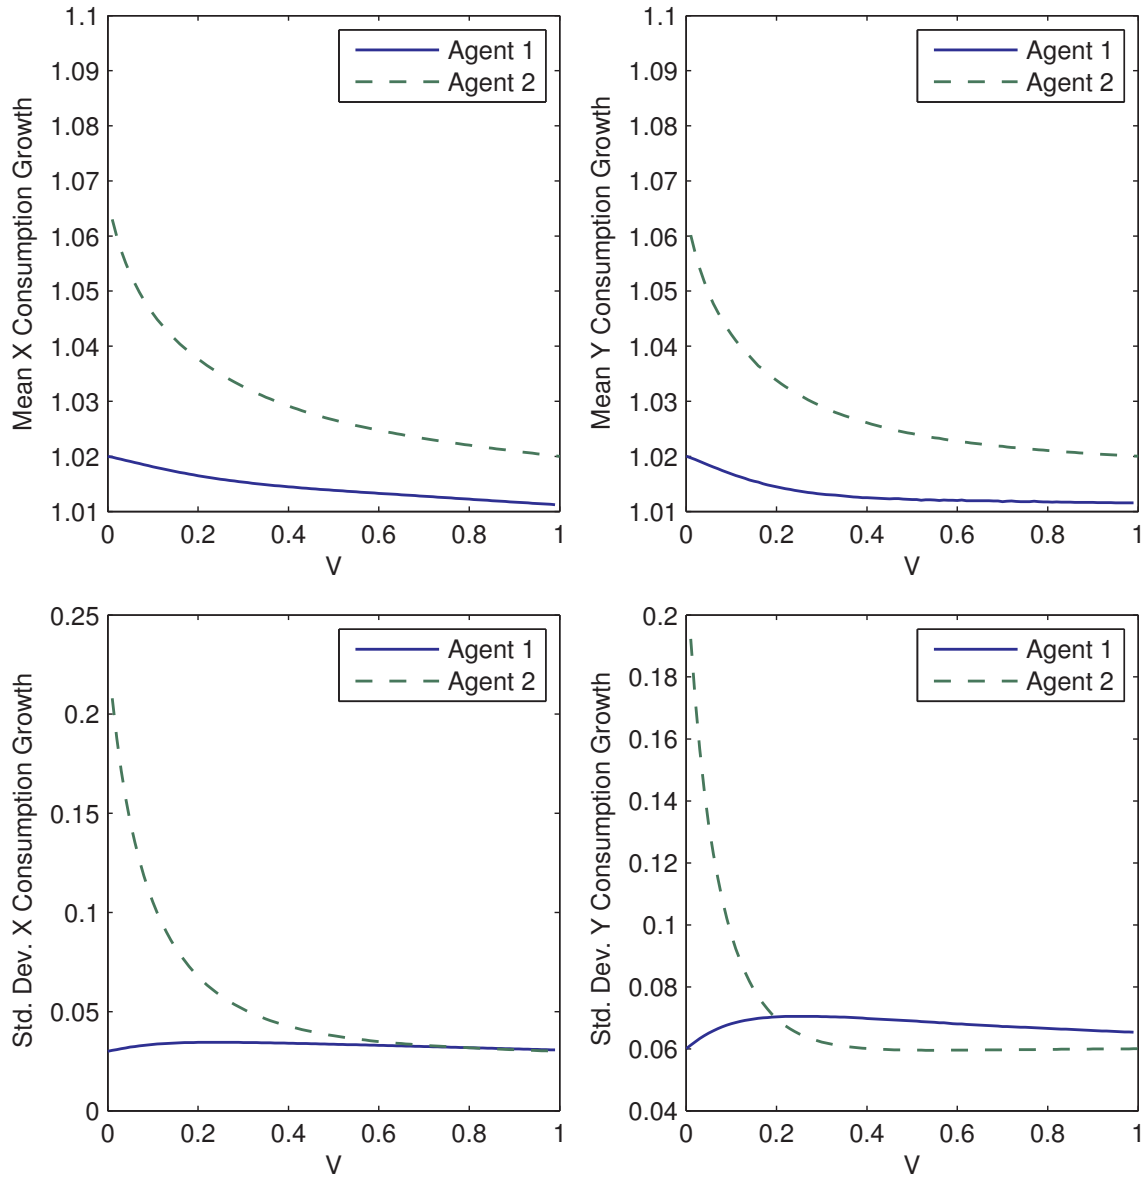
The plots illustrate the average path of the economy over a 50-year period, conditional initial $V_0 = 0.05$. The initial growth state is selected according to the stationary distribution. From top to bottom, the panels show V (indicative of agent 2's wealth share), the oil spot price, risk premium on 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.

Figure 2.13: Characterization of Risk Aversion: Aggregated Consumption Growth




Above we see the mean growth in each agent's aggregated consumption in the left panel, and the standard deviation of aggregated consumption in the right panel. Results are shown conditional on V (essentially agent 2's wealth-share), taking expectations over growth states using their stationary distribution. Based on his Epstein-Zin preference parameters, agent 2 is more tolerant of risk over states (α) and over time (ρ) than agent 1. Notice that agent 2's aggregated consumption (A_2) grows with a higher mean but also generally a higher variance than agent 1's. When agent 1 has most of the wealth, agent 2 willingly absorbs more than his share of aggregate risk, in exchange for faster mean growth. When agent 2 is dominant, equilibrium prices adjust so that he is content with the mean and standard deviation of growth that is inherent to the economy (that of the aggregated endowments).

Figure 2.14: Characterization of Risk Aversion: X and Y Consumption Growth



The existence of two goods adds an extra layer to risk-sharing. For each agent we see the mean and standard deviation of good X (numeraire) consumption growth in the left column, and mean and standard deviation of good Y (oil) consumption growth on the right. Results are shown conditional on V (essentially agent 2's wealth-share), taking expectations over growth states using their stationary distribution. Good X results are similar to those for the aggregated good (see Figure 2.13), as X consumption represents most of the value of total consumption. Good Y consumption is less valuable overall, but agent 2 values consumption of oil about twice as highly as agent 1. Although agent 2 is generally more risk tolerant than agent 1, the standard deviation of agent 2's oil consumption growth is lower than that of agent 1 in most states of the world. By this measure, agent 2 is less tolerant of oil risk than agent 1, even though he is more tolerant of aggregated risk.

Figure 2.15: What Drives the Risk Premium



The above plots illustrate the contribution of the two state variables, growth state s and wealth share V , to the risk premium on 1-period oil futures. Results are shown conditional on initial V (approximately agent 2's wealth-share). The top panels show mean and standard deviation of excess returns for three scenarios: allowing only V to vary, allowing only s to vary, and allowing both V and s to vary (the equilibrium scenario). The bottom panels illustrate time-variation in the risk premium, which is expected excess returns, under the same three scenarios.

Chapter 3

Disagreement, Financial Markets, and the Real Economy

I study consumption, asset prices, and portfolios in a production economy with two agents who disagree regarding expected stock returns. The economy is characterized by continual overconsumption by individuals, and periodic “consumption booms” at the aggregate level. An optimistic investor believes stocks offer high excess returns, whereas a pessimist perceives low or negative excess returns. Each investor is able to construct a portfolio that seems to offer higher returns than he could achieve without the presence of his “misinformed” counterpart. As a result, each believes that his portfolio can support a higher level of consumption. When the wealth distribution tilts toward the optimist, aggregate consumption rises to a level not seen in either agent’s homogeneous economy: a consumption boom. In a multi-sector version of the economy, controversy regarding one small firm is sufficient to cause significant movements in the price of a larger firm.

3.1 Introduction

This paper links disagreement among investors in financial markets to the real economy. I analyze an economy in which investors assign different expected returns to stocks because they disagree about the firms’ productivities. This disagreement spills into the money market, as different expected stock returns support

different risk-free rates. When prices are determined in equilibrium, the relative returns on investments are altered from those the agents would encounter in homogeneous economies of their respective types. This has two main consequences for the real economy. One is that the allocation of capital among firms may be distorted. Another is that the trade-off between consumption and saving is altered, possibly leading to overconsumption or overinvestment. These effects are seen in the portfolios of individual agents, and at an aggregate level.

To examine these issues, I present two examples of the canonical linear AK production economy extended to allow for two agent types. The agents are identical except in their perceptions of productivity, which govern expected stock returns. The first example focuses on questions related to consumption and saving, and assumes a single firm type (or sector). One agent is optimistic, expecting high firm productivity, whereas the other is pessimistic. A key result is that both agents “overconsume” in the heterogeneous economy relative to how they would consume in their respective homogeneous economies. The effects are strongest when an agent holds a small fraction of the economy’s total wealth. Suppose the pessimist expects stock returns roughly 2% per year lower than the optimist. When the pessimist is a small player in a market dominated by optimists, he consumes a fraction of his wealth almost 50 percent higher than he would consume in an economy dominated by pessimists. Likewise the optimist facing a market dominated by pessimists consumes over 40 percent more than in a purely optimistic economy. Overconsumption is more moderate when wealth is more evenly distributed, but agents always consume more in the heterogeneous economy than in their respective homogeneous economies.

Agent perceptions of risk premia explain these results, as they determine achievable returns on individual portfolios. In a standard representative agent economy, the equilibrium risk-free rate is such that no trade occurs in the money market. With different agent perceptions of expected stock returns, the equilibrium risk-free rate falls between the lower rate that would obtain in the homogeneous pessimist’s economy and the higher one of the homogeneous optimist’s economy. Hence the risk-free rate offered to the pessimist in the heterogeneous market seems attractive relative to the low expected stock returns he perceives. Consequently he lends in the money market. The optimist is happy to be his counterparty, as to him the risk-free rate appears low: in his estimation the stock offers high excess returns, so he takes a levered long position in the stock. The essential point is that each agent perceives his portfolio as obtaining higher (risk-adjusted) returns in the mixed economy than he could obtain in his homogeneous economy. In a sense, given he has the same

amount of capital to invest in either case, each agent feels richer when he has a “misinformed” counterparty to trade with. Under the assumption that agents have an elasticity of inter-temporal substitution less than 1, this leads them to consume a greater fraction of their wealths in the heterogeneous economy.¹

The interplay between wealth distribution, prices, and individual agent consumption leads to nontrivial dynamics in aggregate consumption that run counter to the intuition from homogeneous economies. When the firm is productive, the stock return is high, and wealth shifts toward the optimist. In isolation, the optimist would consume more of his wealth than the pessimist (approximately 27 % more in my example), so one might expect greater optimist wealth to result in higher aggregate consumption. However in the mixed economy a wealthier optimist implies a higher risk-free rate. This raises the pessimist’s consumption (as he invests in the money market), but lowers the optimist’s consumption (the leveraged stock position is less attractive). Whether high stock returns increase or decrease aggregate consumption depends upon how wealth is initially distributed. Aggregate consumption achieves its maximum when the optimist controls roughly 85 % of aggregate wealth. In this economic state both the optimist and the pessimist consume more than would be consumed in a homogeneous optimist economy. Such occurrences are called “consumption booms.”

To study the effects of disagreement upon capital allocation between firms, I introduce a two-firm parameterization of the model where agents disagree moderately about the productivity of one firm. Parameters are chosen such that the controversial firm is relatively small, with approximately 15 % market share in the initial period. However its market share may grow to over 25 % if stock returns favor the optimist, or sink to as little as 5 % if they favor the pessimist. Despite these shifts, aggregate consumption and saving behavior remains almost unchanged. Basic accounting implies that “spillover” from the controversy regarding the small firm significantly influences the price of the large firm. The altered capital allocation affects the bond market through changes in the variance of aggregate output. There is little diversification under the pessimist, which leads to greater variation in aggregate output and a lower risk-free rate. The optimist invests more in the small firm, which reduces variance in output and supports a higher risk-free rate. Therefore a shift in wealth toward the optimist produces two effects: greater investment in the small firm (and hence higher total investment in stocks), and a higher risk-free rate (which depresses investment in stocks). The

¹In the examples that follow I assume constant relative risk aversion $\gamma = 2$, which corresponds to $EIS = 1/2$. As one might expect, an $EIS > 1$ leads to “underconsumption” in heterogeneous economies, and the knife-edge case of log utility implies that each agent consumes the same as he would in his homogeneous economy.

result is a shift in allocation across sectors, but little change in total capital allocated to production.

Related Literature

Numerous papers investigate the impact of investor disagreement upon asset prices. Recent examples include Banerjee and Kremer (2010), David (2008), Dumas, Kurshev, and Uppal (2009), Gallmeyer and Hollifield (2008), Li (2007), and Yan (2008). Each of these papers incorporates interesting belief dynamics, frictions, or additional forms of heterogeneity for increased realism, or to address empirical asset pricing phenomena such as trading volume, price volatility, or the equity premium. However they restrict analysis to endowment economies in the style of Lucas (1978), in which investor disagreement impacts prices and portfolios but leaves aggregate consumption and the path of economic growth unchanged. The purpose of my paper is to address precisely those issues entwined with macroeconomic quantities not addressed in this extensive literature. My results have implications for the extension of heterogeneous beliefs asset pricing results to the production setting. For example, the price volatility produced by my model is far smaller than that which would obtain in an endowment economy with similar disagreement regarding the aggregate dividend process. In line with Jermann (1998), this suggests that some form of capital adjustment friction is required to extend disagreement-based explanations of price volatility to a production setting.

In contrast to the endowment setting, theoretical studies of production economies with investor disagreement are few. Detemple and Murthy (1994) is closest to this paper. They also investigate a frictionless production economy with investor disagreement, obtaining closed-form results for prices and portfolios in a continuous-time framework with one productive sector. However they restrict preferences to log utility, which leaves consumption unaffected by beliefs. More recently Branch and McGough (2011) investigate the effects of belief heterogeneity on the internal propagation of real business cycle models. Their work focuses the consumption and output dynamics induced by TFP shocks, as opposed to connections between markets, portfolios and consumption.

The paper is organized as follows. Section 2 describes the model. Section 3 solves the planner's problem, and characterizes prices and portfolios of the corresponding competitive equilibrium. Section 4 presents the single-firm version of the model. Section 5 considers a two-firm parameterization. Section 6 concludes.

3.2 The Model

I construct a discrete-time, discrete-state stochastic optimal growth model in the style of Brock and Mirman (1972), but extended to two agents and M firms. However I abstract from labor. There is a single good, which is both consumable and used as capital in production. The next subsections describe firms, households, and markets.

3.2.1 Firms

There are M competitive firms responsible for production. Each firm $i \in M$ has access to a simple linear technology that produces output according to $F_t^i = \tilde{A}_t^i K_{t-1}^i$, where F_t^i is output produced at the beginning of period t , K_{t-1}^i is capital input selected at the end of the previous period $t - 1$, and \tilde{A}_t^i is the stochastic productivity of technology i , which is observed at the beginning of period t . The input K_t^i is fully depreciated in the production process, and there are no capital adjustment costs or other frictions to concern the firm. Define the vector $\tilde{A}_t = [\tilde{A}_t^1, \tilde{A}_t^2, \dots, \tilde{A}_t^M]$. I assume \tilde{A}_t is discrete and i.i.d. each period with outcomes $A_t \in \mathbb{A} \subset \mathbb{R}^M$, $\forall t$, and denote the number of states $N = |\mathbb{A}| < \infty$. Note that although I assume \tilde{A} is i.i.d. over time, I allow productivity to be correlated across firms.

At time $t = 0$, each firm i raises capital by issuing a single, infinitely divisible share of stock. It maximizes the value of its stock by announcing a plan for optimal dividends,

$$\begin{aligned}
 p_0^{s,i} &= \max_{\{d_t^i\}_{t=1}^{\infty}} E_0 \left[\sum_{t=1}^{\infty} m_t d_t^i \right] \\
 \text{s.t. } d_t^i &= F_t^i - K_t^i, \\
 F_t^i &= \tilde{A}_t^i K_{t-1}^i,
 \end{aligned} \tag{3.1}$$

where $K_0^i = p_0^{s,i}$ is the value of capital raised in the IPO, and m_t is a stochastic discount factor, which I discuss in section 3.

3.2.2 Households

The economy is populated by an optimist and a pessimist, $j \in \{O, P\}$, who live forever. The agents have identical CRRA preferences over consumption of the numeraire,

$$u(c^j) = \frac{(c^j)^{1-\gamma}}{1-\gamma}, \quad u(c^j) = \log(c^j) \text{ if } \gamma = 1. \quad (3.2)$$

In the first period, agent O is endowed with a fraction $\theta_O \in [0, 1]$ of the initial stock of the good, and agent P receives the remainder, $\theta_P = 1 - \theta_O$. Through investment of his endowment, each agent maximizes his expected utility of lifetime consumption,

$$\begin{aligned} \max_{\{c_t^j\}_{t=0}^{\infty}} E_0^j \left[\sum_0^{\infty} \beta^t u(c_t^j) \right] \\ \text{s.t. } E_0^j \left[\sum_{t=0}^{\infty} m_t^j c_t^j \right] \leq \theta_j f_0, \end{aligned} \quad (3.3)$$

where $\beta \in (0, 1)$ is a constant reflecting impatience, and E_t^j denotes expectation with respect to agent j 's beliefs given the information set at time t . Although identical in most respects, the two agents disagree about the probability distribution of the firm productivities \tilde{A} . That is, although they agree on the state space (their measures are equivalent), they assess state probabilities according to $P_O[\tilde{A} = A]$, $A \in \mathbb{A}$ and $P_P[\tilde{A} = A]$, respectively. Therefore the two agents will not generally agree on the optimal way to invest a given amount of wealth. In fact, the technology actually behaves according to a third distribution, that of the econometrician, $P_E[\tilde{A} = A]$. Although it is reasonable to assume that rational agents would eventually learn the true distribution of \tilde{A} , O and P are irrational insofar as their beliefs are dogmatic.²

For analytical convenience I introduce a representative agent who is a composite of O and P. In each period, the representative agent solves

$$\bar{u}(c, \lambda) = \max_{c^O} u(c^O) + \lambda u(c - c^O) \quad (3.4)$$

where c is current aggregate consumption, λ is a weighting factor reflecting the current relative wealth of the two agents, and $c^P = c - c^O$. λ evolves stochastically over time according to the law of motion

$$\lambda_t = \tilde{z}_t \lambda_{t-1}, \quad (3.5)$$

²Even if agents are fully rational learners, different beliefs may persist for long periods given sufficiently different priors. This is particularly true if the productivity process were to be made more realistic. See for example Yan, 2008, p. 1946.

where \tilde{z}_t is a discrete random variable that is i.i.d. across periods (and hence, as with \tilde{A} , I will usually omit the time subscript). \tilde{z} represents disagreement between the agents regarding the probabilities of technological outcomes in the next period. It is a “one-period change of measure” from the beliefs of O to those of P , i.e., $\tilde{z} = \frac{\lambda_{t+1}}{\lambda_t} = \frac{P_P[\tilde{A}=A]}{P_O[\tilde{A}=A]}$, $\forall A \in \mathbb{A}$, $\forall t$. λ_0 is a constant chosen such that each agent’s budget constraint is satisfied, an issue that I address in a later section.³

The above formulation of the representative agent’s utility function and the accompanying law of motion for λ implicitly assume that we will optimize expected utility under the optimist’s measure. Henceforth and without loss of generality, I take expectations with respect to beliefs of the optimist O unless otherwise specified.⁴

Taking aggregate consumption and the weighting factor as given, the solution to Equation (3.4) is

$$c^O(c, \lambda) = \frac{c}{1 + \lambda^{1/\gamma}}, \quad (3.6)$$

from which it follows that the utility of the representative agent is

$$\bar{u}(c, \lambda) = \frac{c^{1-\gamma}}{1-\gamma} (1 + \lambda^{1/\gamma})^\gamma. \quad (3.7)$$

Note that $\bar{u}(c, \lambda)$ is homogeneous of degree $1 - \gamma$ in c .

3.2.3 Markets

Markets are dynamically complete, with as many linearly independent assets as states N , and trade is frictionless. In particular, I assume that there is a one period bond (or equivalently a money market), which pays one unit of the good regardless of technological outcome, and one share of stock in each firm. With the exception of the stocks, which are each in net supply one, all assets are assumed to be in zero net supply.

³For a practical discussion of change of measure in a discrete setting, see for example Shreve, 2004, p. 61. For additional discussion of the representative investor in a heterogeneous beliefs economy, see Basak (2005).

⁴The model could equivalently be solved under the pessimist’s measure by placing the stochastic weighting factor λ on the optimist’s utility $u(c^O)$, and defining $\tilde{z} = \frac{P_O[\tilde{A}=A]}{P_P[\tilde{A}=A]}$ instead. The resulting equilibrium capital allocation would be the same.

3.3 Equilibrium

Equilibrium is determined by two key state variables realized at the start of each period: the stochastic weighting factor λ_t , and aggregate output $f_t = \sum_{i=1}^M F_t^i$. The fact that we need only keep track of aggregate output and not sectoral allocations is due to the assumption of frictionless capital adjustment. I solve for equilibrium in two main steps: first I solve the planner's problem, then I construct the corresponding competitive equilibrium.

To solve the planner's problem, I determine optimal aggregate consumption c_t and capital allocation $K_t = [K_t^1, K_t^2, \dots, K_t^M]'$ taking the initial weighting factor λ_0 and aggregate stock of the good f_0 as given. Subsequently I construct a stochastic discount factor m_t . This is used to derive the wealth w_t^j of each agent as a function of λ_t and f_t , which is finally used to determine λ_0 consistent with the initial wealth allocations given by θ_O , θ_P and f_0 .

With the solution to the planner's problem in hand, the prices and portfolios forming the competitive equilibrium derive from the stochastic discount factor m_t and the wealth functions w_t^j . I discuss stock prices and the risk-free rate, and explain the solution procedure for portfolios.

All proofs are in the appendix.

3.3.1 Aggregate Saving and Capital Allocation

I select K_t as the choice variable, denoting aggregate saving $k_t = \sum_{i=1}^M K_t^i$ and aggregate consumption $c_t = f_t - k_t$. The objective is to solve for an optimal policy function $K_t = K(f_t, \lambda_t)$ that gives capital allocation as a function of the state variables. Formulated as a sequence problem, the value function is defined as

$$\begin{aligned} \bar{v}(f_0, \lambda_0) &= \max_{\{K_t\}_{t=0}^{\infty}} E_0 \left[\sum_{t=0}^{\infty} \beta^t \bar{u}(f_t - k_t, \lambda_t) \right] \\ \text{s.t. } \lambda_t &= \tilde{z}_t \lambda_{t-1}, \\ f_t &= \tilde{A}_t K_{t-1}. \end{aligned} \tag{3.8}$$

Typically we also desire nonnegative and bounded capital allocations $K_t^i \in [0, 1)$, $k_t \in (0, 1)$. I do not enforce this formally, but it results naturally from the form of the utility function and judicious selection of parameters.

Due to the infinite horizon setting, the solution can be given in terms of the state variables f and λ without explicit reference to time. Under technical conditions the value function can be formulated recursively as

$$\bar{v}(f, \lambda) = \max_K \bar{u}(f - k, \lambda) + \beta E[\bar{v}(\tilde{A}K, \tilde{z}\lambda)], \quad (3.9)$$

The right hand side has first order conditions

$$\left(\frac{1 + \lambda^{1/\gamma}}{f - k} \right)^\gamma = \beta E[\bar{v}'(\tilde{A}K, \lambda\tilde{z})\tilde{A}^i], \quad i \in \{1, \dots, M\}. \quad (3.10)$$

The envelope theorem implies $\bar{v}'(f, \lambda) = \bar{u}'(f - k(f, \lambda), \lambda)$. In combination with first order conditions, this yields the Euler equations,

$$\left(\frac{1 + \lambda^{1/\gamma}}{f - k(f, \lambda)} \right)^\gamma = \beta E \left[\left(\frac{1 + (\tilde{z}\lambda)^{1/\gamma}}{\tilde{A}K(f, \lambda) - k(\tilde{A}K(f, \lambda), \tilde{z}\lambda)} \right)^\gamma \tilde{A}^i \right], \quad i \in \{1, \dots, M\}. \quad (3.11)$$

The functional form of the solution is characterized as follows.

Proposition 1. *Any function $K(f, \lambda)$ satisfying Equation (3.11) is homogeneous of degree one in f , and has the form $K(f, \lambda) = fB(\lambda)$, where $B : \mathbb{R}_+ \rightarrow \mathbb{R}^M$ satisfies*

$$\left(\frac{1 + \lambda^{1/\gamma}}{1 - b(\lambda)} \right)^\gamma = \beta E \left[\left(\frac{1 + (\tilde{z}\lambda)^{1/\gamma}}{1 - (\tilde{A}B(\lambda))b(\tilde{z}\lambda)} \right)^\gamma \tilde{A}^i \right], \quad i \in \{1, \dots, M\} \quad (3.12)$$

for $b(\lambda) = \sum_{i=1}^M B^i(\lambda)$.

Unfortunately a closed form solution to Equation (3.12) seems unlikely. However an accurate numerical approximation to the solution can be computed without difficulty. Details are available upon request.

3.3.2 State Prices

The representative agent values consumption in a given time and state according to a stochastic discount factor m_t reflecting his marginal utility of consumption in that time and state. Specifically, if t is the current period, then

$$m_{t+1} = \beta \frac{\bar{u}'(c_{t+1}, \lambda_{t+1})}{\bar{u}'(c_t, \lambda_t)} = \beta \frac{u'(c_{t+1}^O)}{u'(c_t^O)} = \beta \tilde{z} \frac{u'(c_{t+1}^P)}{u'(c_t^P)}. \quad (3.13)$$

The equivalence of the representative agent's SDF and those of P and O is implied by the envelope theorem, and is easily verified algebraically. By way of example, let a_{t+1} be an asset that offers an arbitrary stochastic

payoff in the next period only, and recall that \tilde{z} is a change of measure from beliefs of P to those of O. Therefore the price of a_{t+1} is

$$p_a = E_t^O \left[\beta \frac{u'(c_{t+1}^O)}{u'(c_t^O)} a_{t+1} \right] = E_t^O \left[\beta \tilde{z} \frac{u'(c_{t+1}^P)}{u'(c_t^P)} a_{t+1} \right] = E_t^P \left[\beta \frac{u'(c_{t+1}^P)}{u'(c_t^P)} a_{t+1} \right]. \quad (3.14)$$

Whether agent P values a_{t+1} under his beliefs and according to his marginal utility of consumption, or O does the same using his beliefs and marginal utility, the value of a_{t+1} is the same. So O and P disagree about state probabilities, but agree upon state prices. An important implication is that, when the firm maximizes the value of its shares, the choice of measure and SDF is irrelevant. Finally, homogeneity of \bar{u} in c and of K in f imply that we can express

$$m_{t+1} = m(\lambda_t) = \beta \left(\frac{1 + \lambda_t^{1/\gamma}}{1 - b(\lambda_t)} \right)^{-\gamma} \left(\frac{1 + (\tilde{z}\lambda_t)^{1/\gamma}}{(1 - b(\tilde{z}\lambda_t))(\tilde{A}B(\lambda))} \right)^\gamma, \quad (3.15)$$

i.e. the SDF does not depend upon the current period level of aggregate output.

3.3.3 Wealth and Budget Constraints

I define an agent's wealth as the discounted value of his contingent claims for present and future consumption. Here I develop an expression for the wealth of the optimist, $w_O(f, \lambda)$, as a function of the aggregate capital stock and weighting factor. With $w_O(f, \lambda)$ in hand it is easy to solve for λ_0 such that each agent's budget constraint is satisfied, and also to determine their optimal portfolios. O's wealth (calling the current period $t = 0$) is defined as

$$w_O = E_0 \left[\sum_{t=0}^{\infty} m_t c_t^O \right]. \quad (3.16)$$

Using the laws of motion for f and λ , wealth can be expressed recursively as

$$\begin{aligned} w_O(f, \lambda) &= c_O(f, \lambda) + \beta E \left[\frac{c_O(\tilde{A}B(\lambda)f, \tilde{z}\lambda)^{-\gamma}}{c_O(f, \lambda)^{-\gamma}} w_O(\tilde{A}B(\lambda)f, \tilde{z}\lambda) \right] \\ &= \frac{f(1 - b(\lambda))}{1 + \lambda^{1/\gamma}} + \beta E \left[\left(\frac{\tilde{A}B(\lambda)(1 - b(\tilde{z}\lambda))}{1 + (\tilde{z}\lambda)^{1/\gamma}} \right)^{-\gamma} \left(\frac{1 - b(\lambda)}{1 + \lambda^{1/\gamma}} \right)^\gamma w_O(\tilde{A}B(\lambda)f, \tilde{z}\lambda) \right] \end{aligned} \quad (3.17)$$

A function of the form $w_O(f, \lambda) = D(\lambda)f$ will satisfy the equation, where $D(\lambda)$ satisfies

$$D(\lambda) = \frac{1 - b(\lambda)}{1 + \lambda^{1/\gamma}} + \beta E \left[\left(\frac{1 - b(\tilde{z}\lambda)}{1 + (\tilde{z}\lambda)^{1/\gamma}} \right)^{-\gamma} \left(\frac{1 - b(\lambda)}{1 + \lambda^{1/\gamma}} \right)^\gamma D(\tilde{z}\lambda)(\tilde{A}B(\lambda))^{1-\gamma} \right]. \quad (3.18)$$

The function $D(\lambda)$ gives the optimist's fraction of aggregate wealth. With one additional assumption it is possible to establish the existence of a unique solution to Equation (3.18).

Assumption 1. Let $\hat{\beta} = \sup_{\lambda} \beta E \left[\left(\frac{1-b(\bar{z}\lambda)}{1+(\bar{z}\lambda)^{1/\gamma}} \right)^{-\gamma} \left(\frac{1-b(\lambda)}{1+\lambda^{1/\gamma}} \right)^{\gamma} (\tilde{A}B(\lambda))^{1-\gamma} \right]$. Assume model parameters s.t. $\hat{\beta} < 1$.

Although satisfaction of Assumption 1 cannot generally be verified analytically, I verify it computationally for parameters used in numerical results later in the paper.

Proposition 2. Define a space of continuous, bounded functions $\mathcal{D}(\Lambda)$, $\Lambda \equiv \mathbb{R}_+$, $g \in \mathcal{D}(\Lambda)$ s.t. $g : \Lambda \rightarrow [0, 1]$, with the sup norm $\|g\| = \sup_{\lambda \in \Lambda} |g(\lambda)|$. Let \mathcal{T}_D be the mapping given by Equation (3.18),

$$[\mathcal{T}_D g](\lambda) = \frac{1-b(\lambda)}{1+\lambda^{1/\gamma}} + \beta E \left[\left(\frac{1-b(\bar{z}\lambda)}{1+(\bar{z}\lambda)^{1/\gamma}} \right)^{-\gamma} \left(\frac{1-b(\lambda)}{1+\lambda^{1/\gamma}} \right)^{\gamma} g(\bar{z}\lambda) (\tilde{A}B(\lambda))^{1-\gamma} \right]. \quad (3.19)$$

Under Assumption 1, there is a unique solution $D(\lambda) \in \mathcal{D}(\Lambda)$ to Equation (3.18), and $\forall g \in \mathcal{D}(\Lambda)$, $\lim_{N \rightarrow \infty} [\mathcal{T}_D^N g](\lambda) \rightarrow D(\lambda)$.

The optimist's wealth w_O is not simple to characterize analytically for general parameters. The exception is log utility ($\gamma = 1$), in which case $w_O(f, \lambda) = \frac{f}{1+\lambda}$, i.e., the agent's share of wealth is identical to his share of consumption. As with K , I compute w_O numerically for general parameters satisfying Assumption 1.

In order to complete the equilibrium solution the initial weighting factor λ_0 must be chosen such that each agent's budget constraint is satisfied. That is, λ_0 is chosen to solve

$$w_O(f, \lambda_0) = \theta_O f_0. \quad (3.20)$$

Since the pessimist holds any endowment and consumption claims not held by the optimist, satisfaction of O's budget constraint also implies satisfaction of P's.

3.3.4 Asset Prices

Using the optimal capital allocation that solves the planners problem, I derive prices supporting a competitive equilibrium. From the stochastic discount factor, the price of the one-period bond is

$$\begin{aligned} p_b(\lambda) &= E[m(\lambda)] \\ &= \beta \left(\frac{1+\lambda^{1/\gamma}}{1-b(\lambda)} \right)^{-\gamma} E \left[\left(\frac{1+(\bar{z}\lambda)^{1/\gamma}}{(1-b(\bar{z}\lambda))(\tilde{A}B(\lambda))} \right)^{\gamma} \right], \end{aligned} \quad (3.21)$$

and the gross risk-free rate is simply the reciprocal of the bond price,

$$r(\lambda) + 1 = \frac{1}{p_b(\lambda)}. \quad (3.22)$$

Following logic equivalent to that of Cox, Ingersoll Jr, and Ross, 1985, p. 382, the ex-dividend price of the stock must be equal to the capital held by the firm, i.e. for firm i ,

$$p_s^i(f, \lambda) = K^i(f, \lambda). \quad (3.23)$$

The basic argument behind this result relies on the assumptions of free entry and frictionless capital adjustment. If $p_s^i > K^i$, then the firm's owners could profitably sell their shares and invest K^i in a new firm employing the same technology (and hence capable of generating the same dividend stream as the original firm). If $p_s^i < K^i$, then the firm can liquidate its capital and pay K^i to its owners as a dividend, who could then invest K^i in a new firm employing the same technology (where by nature of the IPO the value of the firm would be $p_s^i = K^i$). The fact that the firm's owners desire K^i to be invested in technology i follows by the fact that they are (as a group) the representative agent, and it is precisely this allocation that supports the optimal aggregate dividend stream, $c(f, \lambda) = \sum_{i=1}^M d^i = f - \sum_{i=1}^M K^i(f, \lambda)$.

The fact that price is equal to capital is significant for stock returns. Gross returns are

$$\frac{d_{t+1}^i + p_{t+1}^i}{p_t^i} = \frac{F_{t+1}^i}{K_t^i} = \tilde{A}_{t+1}^i, \quad (3.24)$$

that is, the gross returns on a firm's stock are equal to the productivity of the firm's technology. It follows that gross returns are i.i.d., and that disagreement regarding technology processes could equally be characterized as disagreement regarding the distribution of stock returns.

3.3.5 Portfolios

In equilibrium, each agent follows a portfolio rule that implements his optimal wealth process, which in turn supports his optimal consumption plan. Having solved for the wealth of the optimist as a function of the state variables, I reverse engineer his optimal portfolio rule to satisfy this relationship.

Let $R(f, \lambda)$ be an $N \times N$ matrix of gross asset payoffs, such that $R_{j,i}(f, \lambda)$ is the next-period payoff of asset i in state j given the current state is described by (f, λ) .⁵ Let $\phi(f, \lambda)$ be an $N \times 1$ vector implementing the portfolio rule, and finally $\Omega(f, \lambda)$ an $N \times 1$ vector describing target wealth outcomes in each possible next-

⁵Of course it is possible to have any number of assets $\geq N$ s.t. the rank of R is N , but I assume for simplicity that the portfolio rule uses the minimal number of assets.

period state. Then

$$\begin{aligned} R(f, \lambda)\phi(f, \lambda) &= \Omega(f, \lambda) \\ \Rightarrow \phi(f, \lambda) &= R^{-1}(f, \lambda) \Omega(f, \lambda) \end{aligned} \tag{3.25}$$

Given the optimist O's portfolio, that of the pessimist is recoverable by market clearing.

In particular, consider the case $N = 2$, $M = 1$ with only the stock and the bond as assets.⁶ Then the optimal portfolio satisfies

$$\phi(f, \lambda) = \begin{bmatrix} A_1 k(f, \lambda) & 1 \\ A_2 k(f, \lambda) & 1 \end{bmatrix}^{-1} \begin{bmatrix} w_O(A_1 k(f, \lambda), Z_1 \lambda) \\ w_O(A_2 k(f, \lambda), Z_2 \lambda) \end{bmatrix} \tag{3.26}$$

3.4 An Economy with One Firm

To provide a simple example that offers some intuition regarding the behavior of the economy, I continue with the one firm, two state formulation introduced in the previous section. This has the advantage of restricting portfolios to the two assets of primary interest, the stock and the one-period bond, as additional assets needn't be introduced in order to complete markets.

Parameter	Value	Description
\mathbb{A}	{H=1.14, L=1.02}	Production technology states
P_P	{0.4, 0.6}	Beliefs, agent P (pessimist)
P_O	{0.6, 0.4}	Beliefs, agent O (optimist)
P_E	{0.5, 0.5}	Beliefs, econometrician
β	0.99	Impatience
f_0	1	Initial stock of the good
θ_O	0.5	Initial wealth share, agent O
γ	2	Relative risk aversion

Table 3.1: Baseline parameter values.

I calibrate the model's parameters to match some aspects of the US economy. There are three probability

⁶Analogous to $A = \{A_1, A_2\}$, let $Z = \{Z_1, Z_2\}$ be the state space of \tilde{z} .

measures: those of the pessimist P, the optimist O, and the econometrician E, where the latter is the measure by which the economy actually evolves. As E plays no active role in the economy (he holds no assets), his beliefs have not been involved in the derivation of equilibrium, but they determine the dynamics of the state variables. They also provide a point of reference. The econometrician knows that each state in \mathbb{A} is equally likely. The optimist believes the high state is more likely than the low state by 0.1, whereas the pessimist believes the low state is more likely, also by 0.1. This amounts to a substantial level of disagreement between the agents, but their errors in beliefs are symmetric about the econometrician's. By selecting equal initial wealth shares $\theta_O = \theta_P = 1/2$ I make a wealth-weighted average of the population's beliefs equal to the econometrician's, i.e., on average the agents are correct.

Table 3.1 gives the full set of parameters. Although the agents are dogmatic in their divergent beliefs, they agree on the levels of productivity corresponding to the two technology states. I choose the high (H) and low (L) states of \mathbb{A} such that, in a homogeneous beliefs economy with agents believing as the econometrician, the mean and variance of output growth match the sample mean and variance of annual US real GDP growth from 1930 through 2009, as recorded by the BEA.⁷ Mean growth is around 3.41%, with standard deviation 5.07%. To match these, the technology has mean return $\mu_A = 7.74\%$ and standard deviation $\sigma_A = 7.52\%$. The agents have identical, moderate risk aversion $\gamma = 2$. The level of risk aversion was chosen to highlight non-monotonicity in savings that is less apparent (but still present) with higher risk aversion.

Obviously a more precise parameter estimation is possible, taking the levels of disagreement and risk aversion as free parameters and attempting to match additional empirical moments. My objective is a simple example that highlights the basic implications of the model. Alternative parameters were explored in previous drafts of this paper; the results presented here are representative of those the model produces. One less innocuous assumption is $\gamma > 1$ and consequently $EIS < 1$, which determines whether the main result of disagreement is overconsumption or overinvestment. Selecting $\gamma < 1$ would, roughly speaking, flip the consumption figures about the horizontal axis. Since it is more common to assume $\gamma > 1$ I do so for purposes of example.

⁷See <http://www.bea.gov/national/index.htm>. I use the data for chained 2005 dollars.

3.4.1 Evolution of the Optimist's Consumption Share

After developing the model it is clear that the state variable of interest is λ , the stochastic weighting factor, so I present its dynamics now. The level of aggregate output f acts simply as a linear scaling factor on consumption, wealth, and stock prices, and does not directly impact state-prices or the interest rate. I discuss dynamics of f - which are influenced by those of λ - in a later section.

Productivity (A_1)	Econometrician's Probability	z_1	λ_1	ω_1
L	1/2	3/2	3/2	0.45
H	1/2	2/3	2/3	0.55

Table 3.2: Period $t = 1$ values of the pareto weight λ_1 and optimist's consumption share ω_1 given realized productivity under the baseline parameters under the simplifying assumption $\lambda_0 = 1$ ($\omega_0 = 0.5$). Low productivity increases the weight on the pessimist's utility and decreases the optimist's consumption share. The opposite occurs given high productivity. Under the baseline parameters, agent disagreement is symmetric about the econometrician's beliefs, which assign each state a 50% probability. It follows that λ is as likely to increase as decrease, and the magnitude of an increase is the reciprocal magnitude of the decrease. Both agents will survive indefinitely in this economy.

There is only one exogenous process driving this model: the productivity of the firm, \tilde{A} . The realization of \tilde{A} will determine that of change of measure \tilde{z} , which will in turn drive λ and most variables of economic interest. For example, consider the optimist's share of aggregate consumption, $\omega \in [0, 1]$,

$$\omega(\lambda) = \frac{1}{1 + \lambda^{1/\gamma}}. \quad (3.27)$$

If productivity is low (L), which the pessimist finds more likely than the optimist, then the realization z^L of \tilde{z} is greater than one. This increases λ , the Pareto weight on the pessimist's utility, which implies a decrease in the optimist's consumption share ω . The *value* of \tilde{z} in each state (high (H) or low (L) productivity) is determined by the relative beliefs of the optimist and the pessimist, but the *probability* that high or low productivity occurs is determined under the econometrician's beliefs, which reflect the true behavior of the production technology. Assuming for simplicity that $\lambda_0 = 1$, Table 3.2 provides a concrete example under the baseline parameters of what occurs in the next period $t = 1$, given realized productivity.




Figure 3.1: Expectation and density of the optimist’s consumption share (ω_t) over time. Initial value $\omega_0 \approx 0.56$ is chosen to satisfy individual budget constraints. At longer horizons ω_t is more likely to be far from its initial position. Although ω_t has a discrete distribution, for purposes of visualization the right panel shows a continuous approximation to the probability mass function.

I characterize economic behavior using ω as the state variable. The optimist’s share of aggregate consumption, ω , is bounded between 0 and 1, and has a more natural interpretation than λ , which is unbounded above. There is also a monotonic mapping between the two variables. An alternative would be to use the optimist’s wealth share $D(\lambda)$ as the state variable, but it is not available in closed form, which makes D a more complicated and less intuitive option.

Figure 3.1 shows the evolution of the optimist’s consumption over time, in expectation and in distribution. The baseline parameters assign each agent equal wealth in the initial period. In contrast to the simplifying assumption in Table 3.2, this usually *does not* imply an equal consumption share $\omega_0 = 0.5$, or equivalently $\lambda_0 = 1$. To satisfy individual budget constraints under baseline parameters a value of $\omega_0 \approx 0.56$ is required, i.e., the optimist initially consumes more of his wealth than the pessimist. This is reflected in the plots. Consumption share is initially tilted slightly towards the optimist, and this is still visible in the distribution of ω_{25} , at the 25 year horizon. Over time the impact of ω_0 washes out, and $E[\omega_t]$ goes to 0.5. As the economy evolves, successive productivity realizations allow ω_t to move further away from its initial value, as shown in the right panel. At long horizons there is a significant probability that the economy is dominated by the pessimist ($\omega_t \approx 0$) or by the optimist ($\omega_t \approx 1$). Although ω_t can come arbitrarily close to 0 or 1, it

never reaches the absorbing states; both agents will survive indefinitely. In fact, as time increases without bound, each value in the domain of ω is reached with probability 1. Therefore it is reasonable to consider economic behavior over the entire domain of ω , while keeping in mind that states near ω_0 are more likely at short horizons. The next few sections explain the economic impact of the shifting consumption share. For a more detailed discussion of the asymptotic behavior of ω (and correspondingly λ), see the appendix.

3.4.2 Consumption

Most prior studies of disagreement take place in endowment economies, where aggregate consumption is exogenous. If the optimist and pessimist were situated in an endowment economy with disagreement about the growth of c_t , then ω_t would retain its interpretation, with the optimist consuming $c_t\omega_t$ and the pessimist the remainder $c_t(1 - \omega_t)$.⁸ If c_t experienced a positive exogenous shock, ω_{t+1} would increase versus ω_t , and optimist would consume more, both as a fraction of the total and in absolute terms. Likewise the pessimist's consumption must decrease in the same way, such that total consumption equals exogenous c_t ; the consumption "decisions" of the two agents are entwined. The obvious distinction in the production setting is that each agent determines - in some respects independently of the other - how much of his wealth he will consume. The sum of these consumption choices determines aggregate consumption endogenously. The optimist's consumption share ω_t remains sufficient as a state variable (with f_t), but it is no longer sufficient to describe the optimist's consumption *decision*.

The right panel of Figure 3.2 shows how much of his wealth each agent chooses to consume, as a function of state variable ω . At the leftmost extreme of the plot, ω approaches 0, the optimist is relatively impoverished, and the pessimist dominates the economy. The pessimist behaves as if the economy is a homogeneous one populated by pessimists.⁹ He consumes little, as he expects slow economic growth. By contrast the comparatively poor optimist consumes a large portion of his wealth, roughly 80% more than the pessimist, and roughly 40 % more than he would consume if he were the only agent type in the economy. These results have nothing to do with each agent's absolute wealth; they obtain whether aggregate output f is high or low.

⁸So the state variables in the endowment economy would be c_t and ω_t , with c_t replacing f_t .

⁹In a homogeneous economy the agent consumes a constant fraction of his wealth, and his wealth is equal to aggregate output f .




Figure 3.2: Aggregate consumption and individual consumption as a fraction of wealth.

As we move rightwards on the horizontal axis, the optimist's consumption share increases, as does his fraction of total wealth. But the fraction of his wealth *that he consumes* decreases sharply, whereas the pessimist consumes more of his wealth as his economic influence wanes. As ω approaches 1, the optimist follows his homogeneous economy policy, whereas the pessimist's consumption is sharply elevated, by roughly 50 % relative to wealth.

These results are remarkable in three respects. First, both agents always over-consume in the heterogeneous economy, although the magnitude of overconsumption declines to nothing as an agent becomes dominant. Second, when the pessimist has little economic influence ($\omega \rightarrow 1$), he consumes a substantially higher fraction of his wealth than the optimist does. This is in sharp contrast to the homogeneous economy analysis, in which, *ceteris paribus*, the pessimist consumes significantly less than the optimist. Finally, for a wide interval on ω , both agents over-consume substantially versus their homogeneous levels, and in the rightmost fifth of the plot both agents consume more than the optimist's homogeneous consumption level.

These facts are apparent in results for aggregate consumption, in the left panel of Figure 3.2. Consistent with individual consumption results, the leftmost extreme of the plot, where $\omega = 0$, corresponds to the pessimist's homogeneous economy consumption rule. Likewise the rightmost extreme with $\omega = 1$ matches the homogeneous optimist economy. But rather than displaying a weighted average characteristic, aggregate consumption is always higher than the weighted average of the homogeneous economy levels, because both

agents are overconsuming. When the average investor is quite optimistic, around $\omega \approx 0.8$, a consumption boom occurs; both optimistic and pessimistic investors over-consume sufficiently that aggregate consumption is driven to levels beyond the range of either the pessimist's or the optimist's homogeneous economy level. In contrast, Detemple and Murthy (1994) find that aggregate consumption is a weighted average of the homogeneous economy levels, and individual agents do not alter their consumption rules from their homogeneous economy plans. These results are due to the assumption of log utility, which I relax.

Casting back to the dynamic properties of the model illustrated by Figure 3.1, the plausible range of values for ω_t becomes quite broad after a century has passed. It is likely that boom and bust cycles will occur over the course of decades, even though productivity is i.i.d.

3.4.3 Asset Prices, Portfolios, and Perceptions




Figure 3.3: Stock price and risk-free rate. The stock price assumes $f = 1$.

The consumption behavior of the previous section becomes more intuitive once competitive equilibrium prices and portfolios are examined. The risk-free rate is the main determinant of portfolios, whereas the price of the stock is important in relation to economic growth. Each of these is shown in Figure 3.3. The risk-free rate r , in the right panel, exhibits a weighted average characteristic. Intuitively the rate is higher when the average investor is optimistic (at the right of the plot) than when he is pessimistic (at the left),

and the difference of the two extremes - over 200 basis points - is rather large. As a point of reference, the risk-free rate that would occur under the econometrician's beliefs (which are the midway between the optimist and the pessimist) is found in the center of the range, at the dotted line marked 'E'.

The stock price (identical to the firm's capital stock) is found in the left panel, assuming for simplicity that aggregate output $f = 1$. Equivalently the plot may be interpreted as the stock price relative to aggregate output (price/GDP). In the single sector economy the shape of the stock price curve is simply a mirror image of aggregate consumption, as $p_s = f - c$. Although aggregate consumption varies substantially in response to shifts in consumption share, by roughly 25 % from peak to trough, this corresponds to relatively minor stock price fluctuations, of only 1% between extremes. This result is very different from what obtains in a simple endowment economy with heterogeneous beliefs, where comparable levels of disagreement would generate price fluctuations orders of magnitude larger.¹⁰ As discussed in Jermann (1998) and elsewhere, the frictionless linear production model with CRRA utility does a poor job of explaining asset prices. Although difference of opinion induces time-variation in a model that would otherwise have a constant risk-free rate and price to GDP ratio, the risk-free rate is too high, and price fluctuations are far too small. The model's performance in this regard would likely be improved by introducing capital adjustment costs or other frictions, and calibrating with a view toward asset pricing.

From the perspective of investors in this model, price volatility is irrelevant: it affects how returns are decomposed into price appreciation and dividend yields, but returns remain i.i.d. What is essential to the investor's decision making is his perception of expected stock returns relative to the risk-free rate. In the top left panel, Figure 3.4 shows each agent's perception of excess returns of the stock. The econometrician's reading of the true excess return falls midway between the two agents. Although investors agree upon the risk-free rate, the optimist perceives much higher excess returns than the pessimist. When the pessimist is dominant ($\omega \approx 0$), the risk-free rate is low. However the pessimist also expects low stock returns, so his perceived excess return approaches his homogeneous economy level. Accordingly he holds the stock and does not lend significantly. (Portfolio weights on bonds and stocks are shown in the left and right panels of Figure 3.4, respectively.) However the optimist facing the same risk-free rate perceives very high excess returns. His wealth is trivial, but he borrows heavily relative to his wealth to take a levered long position in the stock. Perceptions of the Sharpe Ratio, in the top right panel, also correspond closely to excess returns,

¹⁰For examples of price volatility due to heterogeneous beliefs in an endowment economy, see for example Li (2007) and Dumas, Kurshev, and Uppal (2009).

and support each agent's portfolio strategy.




Figure 3.4: Excess returns and the Sharpe ratio as perceived by each agent, and his resulting investment decisions.

Consistent with an increasing risk-free rate, a wealthier optimist implies a lower excess return. All investors agree upon the direction of change, but they disagree regarding levels. The pessimist's perception of excess returns turns negative once the optimist becomes a significant market force. The pessimist sells the stock short, using the proceeds to lend ever more to the optimist, who immediately turns his borrowed money around and invests most of it back in the stock! There is a good deal of trade, but the net effect on investment in the firm is rather slight, consistent with Figure 3.3. Both agents pursue flawed portfolio strategies corresponding to their flawed beliefs. However the pessimist's strategy seems particularly bad, as

he often shorts the stock and lends the proceeds even though the stock does, in reality, offer positive excess returns (e.g., for $\omega \approx 0.5$). Nevertheless, the pessimist is *not* expected to be driven from the market. The resolution to this apparent contradiction is two-fold. First, his lending is not entirely financed through short sales of the stock (he does have positive net wealth), so the expected return on the pessimist's portfolio remains positive under the econometrician's measure. Second, the pessimist consumes at a significantly lower rate than the optimist, except when the optimist dominates the economy. As we see at the rightmost extreme of the plot, excess returns are in fact negative under the true measure when the optimist dominates the economy. Therefore in situations where the pessimist consumes more than the optimist, his portfolio strategy of shorting the stock and buying bonds is expected to pay off.




Figure 3.5: Perceived and actual expected returns on portfolios. In the left panel, the horizontal dotted line marked 'P' indicates the return the pessimist would expect to receive on his portfolio if he were the only agent in the economy (in which case he holds only stock), and likewise the dotted line marked 'O' for the optimist. In the right panel, the dotted horizontal line marked 'E' is the econometrician's homogeneous economy expected portfolio return, which correspond to the true expected return on the stock.

Figure 3.5 offers a visual summary of the preceding logic, depicting perceived (left panel) and actual (right panel) expected portfolio returns for each agent. Perceived expected portfolio returns are strikingly similar to the plot of individual consumption behavior shown in Figure 3.2. The horizontal dotted line marked 'P' indicates the return the pessimist would expect to receive on his portfolio if he were the only agent in the economy (in which case he holds only stock). As the optimist becomes wealthier, interest rates rise, and the

pessimist begins to buy bonds. When the optimist has more than 40 % of the consumption share, the pessimist thinks that his bond-buying strategy delivers higher expected returns than the underlying technology of the economy (the stock) could generate. For large ω the pessimist views bonds as so undervalued relative to the stock that he is actually more ‘optimistic’ about expected portfolio returns than the optimist is! This region roughly corresponds to that in which the pessimist consumes more than the optimist. The optimist’s perceptions follow roughly a mirror image of the pessimist’s: he expects high returns on his levered stock portfolio when the pessimist is dominant. For values greater than $\omega \approx 0.7$, both agents are in effect very optimistic about expected portfolio returns. It is in these situations that consumption booms are observed at the aggregate level.¹¹

The actual expected portfolio returns - computed under the econometrician’s measure - are of course different. The pessimist’s strategy of shorting the stock to buy bonds does not work so well as he thinks. However, his true expected returns are always positive, and when the optimist is dominant the pessimist does actually earn higher expected returns than the optimist. Similarly the optimist’s portfolio does offer high expected returns when the pessimist is dominant, although they are not as sensational as he hopes. Recall that these plots do not directly indicate how much of his wealth each agent invests, so higher expected portfolio returns do not imply an expected increase in wealth or consumption share.

3.4.4 Long-run Impact of Overconsumption

The stock price - equal to the capital stock of the firm - determines the expected rate of economic growth. Whereas the bond market shapes how agents split aggregate output, the stock price influences how much output will be available to split. We observed in Figure 3.3 that the stock price is typically depressed in the heterogeneous economy, a consequence of overconsumption. However the impact of disagreement upon the stock price is small in percentage terms.

¹¹Because agents are risk averse, expected returns are obviously not a sufficient statistic for the desirability of portfolios, although they offer a reasonable summary. For example, when the optimist has a small amount of wealth, the pessimist accepts a lower expected portfolio return than he could achieve by holding the stock, because he trades some expected return for the certainty offered by the bond. For this reason the agents’ consumption decisions do not exactly correspond to expected portfolio returns.




Figure 3.6: Expected aggregate output over time, and distribution of aggregate output after $T = 100$ years.

Figure 3.6 assesses the long-run impact of overconsumption on economic growth. In the left panel, the time-path of expected aggregate output in the heterogeneous economy is plotted versus that of a homogeneous economy populated by agents with correct beliefs - those of the econometrician. Recall that the optimist and pessimist split wealth equally at first, that their beliefs average to those of the econometrician, and that they are both expected to survive indefinitely. The slower economic growth in the heterogeneous economy reflects the effects of difference of opinion per se, rather than a mean error in beliefs. Since all investors would choose the econometrician's path of aggregate output if only they had correct beliefs, the lower rate of economic growth may be viewed as a "cost" of disagreement. Although the magnitude of underinvestment is small in each period, compounding amplifies the effect. After a century, aggregate output is expected to be over 30 % higher in the econometrician's economy. The right panel shows that the result carries over to the distribution of aggregate output at the 100 year horizon: based on a Monte Carlo approximation, the econometrician's output stochastically dominates heterogeneous economy output.

3.5 An Economy with Two Firms

This section presents a model parameterization with two firms, each with access to its own production technology. Whereas the previous section focused on the trade-off between consumption and saving, the

objective here is to study the effects of disagreement upon capital allocation *between* firms. Rather than repeating the discussion from the one-firm economy, I focus on a few elements novel to the two-firm setting. The main result is that moderate disagreement regarding one relatively small firm can cause large swings in capital allocations across firms. This occurs without significantly altering consumption/saving decisions at either the individual or the aggregate level.

Agent	Firm 1: $P[H], P[L]$	Firm 2: $P[H], P[L]$
Econometrician	0.55, 0.45	0.45, 0.55
Pessimist	0.55, 0.45	0.42, 0.58
Optimist	0.55, 0.45	0.48, 0.52




Figure 3.7: Beliefs and the evolution of the optimist’s consumption share. The table at the left shows beliefs regarding the productivity of each firm. Everyone agrees regarding Firm 1’s productivity, but there is moderate disagreement regarding Firm 2. The right panel shows a continuous approximation to the probability mass function of ω_t at various horizons.

I maintain all assumptions regarding risk-aversion and initial conditions from the previous section, as given in Table 3.1, and alter only the specification of productive technologies and agent beliefs. There are two firms, called 1 and 2, each with a access to a single technology. For simplicity the technologies have the same two possible levels of productivity, $\tilde{A}^i \in \{1.2 (H), 0.9 (L)\}$, $i \in \{1, 2\}$, and I further assume that productivity outcomes are independent across firms. However the technologies differ in the probability that a high or a low productivity state occurs, as detailed in Figure 3.7. There is complete agreement among agents and the econometrician regarding Firm 1. In addition, all agents perceive Firm 1 as more productive than Firm 2; the productivity of Firm 1 is more likely to be high. However there is disagreement regarding how disadvantaged Firm 2 is. The optimist believes Firm 2 is highly productive with probability 0.48,

whereas the pessimist thinks the probability is 0.42. The true probability, known to the econometrician, is 0.45. Because Firm 2 is less productive it will receive less investment in equilibrium than Firm 1, but the benefits of diversification lead to positive investment in each firm.¹² Although I refer to the operators of each technology as “Firms” they are open to interpretation. One reasonable interpretation is that Firm 1 is the agglomeration of most industry, whereas Firm 2 represents a single sector about which there is controversy. Alternatively one could think of the firms as countries, with one more developed than the other.

With the introduction of a second binomial technology process there are now four states of nature ($N = 4$), with states defined over tuples $A \in \mathbb{A} = \{H, L\}^2$. The variables \bar{z} , λ , and ω remain scalars. The right panel of Figure 3.7 summarizes the dynamics of ω in the 2 Firm economy. Because disagreement between the optimist and the pessimist is less acute in this parameterization, ω_t is more likely to remain close to its origin over long horizons than it was in the previous section.

The expanded state space requires four linearly independent assets to dynamically complete markets. In addition to three natural assets - stock in each firm and one-period bonds - I add an Arrow-Debreu security that pays one unit of the good when both firms experience low productivity. It may be thought of as a sort of “disaster insurance”, and has price

$$p_d(\lambda) = E_{\lambda, f}[m(f, \lambda)1_d], \quad (3.28)$$

where 1_d is an indicator for the disaster state $\{L, L\}$. In practice agents are nearly able to implement their optimal consumption plans without trading disaster insurance - at most it accounts for slightly over 1 % of either agent’s portfolio - and so I do not discuss it in subsequent analysis.

The focus of this section is capital allocation, which is presented in Figure 3.8. The top left panel shows aggregate saving k , the fraction of aggregate output allocated to production rather than consumption. As in the single-firm economy, aggregate saving is highest when the pessimist dominates the economy ($\omega \approx 0$) and lower when the optimist is ascendant ($\omega \approx 1$). However disagreement is more modest in this economy, and the savings curve is now monotonic, unlike the single-firm setting with more severe disagreement. To examine fractional changes in capital allocation, I also show results relative to a fictitious homogeneous economy populated by econometricians, who hold correct beliefs that are an equally weighted average of the

¹²This parameterization is not calibrated to match GDP growth, but expected growth in aggregate output remains reasonable, at 3.25 % per year. To induce positive equilibrium investment in Firm 2 despite its lower productivity it is necessary to increase the variability of the technology somewhat. Therefore the standard deviation of GDP growth is increased to 12.5 %.

optimist and the pessimist. The econometrician's level of saving is the dotted line marked 'E: k', and saving relative to that baseline is shown in the bottom left panel of Figure 3.8. The range of values for aggregate saving k is very narrow, diverging less than 0.1 % from the econometrician's baseline. By extension there is very little variation in aggregate consumption. This theme extends to individual consumption, where agents over-consume as in the previous section, but by a maximum of 8 basis points relative to wealth (not shown).




Figure 3.8: Absolute and relative changes to aggregate and sectoral capital allocation. The dotted lines marked 'E' indicate allocations that would occur in a fictitious homogeneous economy populated by econometricians. The bottom panels are relative to the econometrician's baseline.

This apparent tranquility belies the sensitivity of sectoral allocations to changes in the wealth distribution. The right panels of Figure 3.8 show capital allocated to each firm, in absolute terms (top) and relative to the econometrician’s baseline allocation (bottom). Results in the top panel may also be interpreted as the price of stock in each firm, and I proceed with that interpretation. Horizontal dotted lines give prices in the econometrician’s baseline economy. The higher valuation of Firm 1 follows from its higher productivity. Although there is no disagreement regarding Firm 1, “spillover” from disagreement regarding Firm 2 causes investment in the large firm to vary by up to 12 % versus the econometrician’s baseline, as illustrated in the bottom right panel. Since aggregate saving remains approximately constant, a decrease in Firm 1’s price implies the same *absolute* increase in Firm 2’s price. This leads to far larger *relative* changes in the small firm’s price. As shown in the bottom right panel, Firm 2 sees variation in excess of 60 % relative to the econometrician’s baseline over the domain of optimist consumption share. The caveat is that this change would occur only slowly, as Figure 3.7 shows that extreme values of ω (near 0 or 1) are unlikely even after 100 years. Annual fractional changes in Firm 2’s price, relative to the size of the economy, will only be a few percentage points.

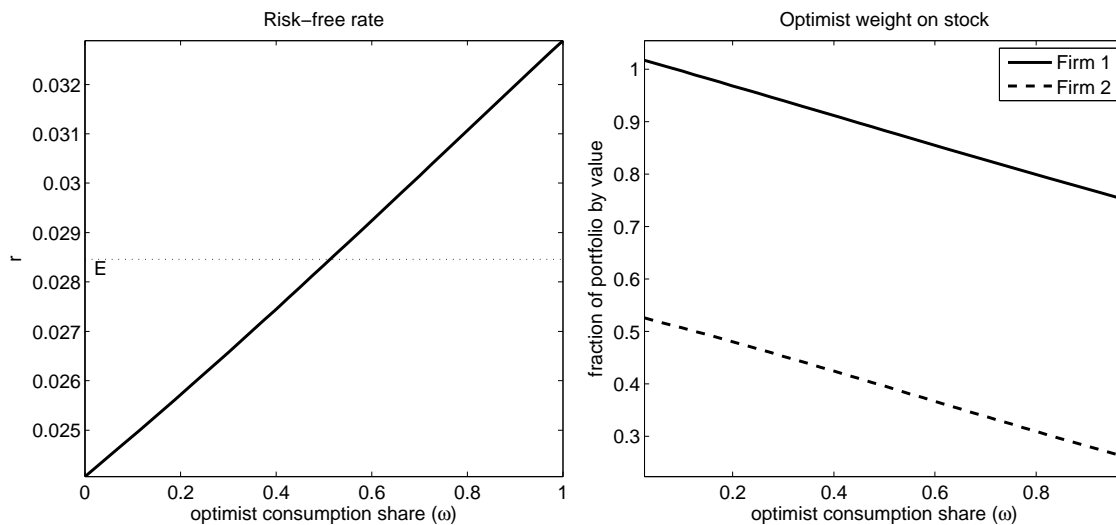


Figure 3.9: Risk-free rate and the optimist’s stock portfolio.

Finally I highlight some aspects of the optimist’s portfolio problem. Figure 3.9 shows the risk-free rate in the left panel, and the optimist’s portfolio weights on the two stocks in the right panel. The risk-free rate is roughly 90 basis points lower when the pessimist is dominant than when the optimist is dominant. This is a result of diversification. Because the optimist perceives Firm 2 as only moderately less productive

than Firm 1, he invests more in Firm 2 than the pessimist would. The result is a reduction in perceived variability of output, which supports a higher risk-free rate.¹³ However, from the standpoint of an optimist choosing his portfolio, the risk-free rate determines the attractiveness of stocks *in general*, as it affects excess returns. The result is that as the optimist's consumption share declines, excess returns increase, and he invests more relative to his wealth in *both* stocks, by borrowing from the pessimist. Individually, the optimist is only directly concerned with the relative weight on Firm 2 in his portfolio, and not with Firm 2's market valuation. The changes in stock prices seen in Figure 3.8 are mainly due to shifting wealth from the optimist to the pessimist, rather than to changes in the relative weight on each stock in each agent's portfolio.

3.6 Conclusion

Aggregate consumption may be higher in an economy with optimistic and pessimistic agents than in an economy with only optimistic agents, even though optimists prefer higher consumption than pessimists in homogeneous economies. The result stems from disagreement regarding excess stock returns. The optimist perceives the stock as offering a high excess return, which is driven by the pessimist's willingness to lend at a low risk-free rate. Enticed by "cheap credit", the optimist borrows heavily to finance high current consumption and a large long position in the stock. In contrast, the pessimist expects stock returns to be lower than the risk-free rate, so a strategy of shorting the stock and buying bonds appears very profitable. Each agent believes his portfolio strategy is capable of supporting an elevated level of consumption. As a consequence of overconsumption, economic growth is depressed. I extend my analysis to a multi-sector economy, where controversy regarding a small sector spills over to a large, uncontroversial sector. The altered capital allocation affects the bond market, with implications for agents' individual portfolio strategies.

¹³It is not the case that shifting capital to Firm 2 increases the expected GDP growth. In fact, since all agents agree Firm 1 is the most productive, there is agreement that expected growth is *lower* when following the optimist's policy. The disagreement is over whether too much expected growth is sacrificed to reduce variance.

3.7 References

- Banerjee, S. and I. Kremer (2010). “Disagreement and learning: Dynamic patterns of trade”. In: *The Journal of Finance* 65.4, pp. 1269–1302.
- David, A. (2008). “Heterogeneous Beliefs, Speculation, and the Equity Premium”. In: *The Journal of Finance* 63.1, pp. 41–83.
- Dumas, B., A. Kurshev, and R. Uppal (2009). “Equilibrium Portfolio Strategies in the Presence of Sentiment Risk and Excess Volatility”. In: *The Journal of Finance* 64.2, pp. 579–629.
- Gallmeyer, M.F. and B. Hollifield (2008). “An Examination of Heterogeneous Beliefs with a Short-Sale Constraint in a Dynamic Economy”. In: *Review of Finance* 12.2, pp. 323–364.
- Li, T. (2007). “Heterogeneous beliefs, asset prices, and volatility in a pure exchange economy”. In: *Journal of Economic Dynamics and Control* 31.5, pp. 1697–1727.
- Yan, H. (2008). “Natural Selection in Financial Markets: Does It Work?” In: *Management Science* 54.11, p. 1935.
- Lucas, R.E. (1978). “Asset prices in an exchange economy”. In: *Econometrica: Journal of the Econometric Society*, pp. 1429–1445.
- Jermann, U.J. (1998). “Asset pricing in production economies”. In: *Journal of Monetary Economics* 41.2, pp. 257–275.
- Detemple, J. and S. Murthy (1994). “Intertemporal asset pricing with heterogeneous beliefs”. In: *Journal of Economic Theory* 62.2, pp. 294–320.
- Branch, W.A. and B. McGough (2011). “Business cycle amplification with heterogeneous expectations”. In: *Economic Theory*, pp. 1–27.
- Brock, W.A. and L.J. Mirman (1972). “Optimal economic growth and uncertainty: The discounted case”. In: *Journal of Economic Theory* 4.3, pp. 479–513.
- Shreve, S.E. (2004). *Stochastic Calculus for Finance: The binomial asset pricing model*. Springer Verlag.
- Basak, S. (2005). “Asset pricing with heterogeneous beliefs”. In: *Journal of Banking and Finance* 29.11, pp. 2849–2881.
- Cox, J.C., J.E. Ingersoll Jr, and S.A. Ross (1985). “An intertemporal general equilibrium model of asset prices”. In: *Econometrica: Journal of the Econometric Society* 53.2, pp. 363–384.
- Cox, D.R. and H.D. Miller (1980). *The theory of stochastic processes*. Chapman and Hall.
- Hoel, P.G., S.C. Port, and C.J. Stone (1972). *Introduction to stochastic processes*. Houghton Mifflin Boston.

3.A Proofs

Proposition 1. Any function $K(f, \lambda)$ satisfying Equation (3.11) is homogeneous of degree one in f , and has the form $K(f, \lambda) = fB(\lambda)$, where $B : \mathbb{R}_+ \rightarrow \mathbb{R}^M$ satisfies

$$\left(\frac{1 + \lambda^{1/\gamma}}{1 - b(\lambda)} \right)^\gamma = \beta E \left[\left(\frac{1 + (\tilde{z}\lambda)^{1/\gamma}}{1 - (\tilde{A}B(\lambda))b(\tilde{z}\lambda)} \right)^\gamma \tilde{A}^i \right], \quad i \in \{1, \dots, M\} \quad (3.29)$$

for $b(\lambda) = \sum_{i=1}^M B^i(\lambda)$.

Proof. The proof is by contradiction. The proposition only concerns homogeneity w.r.t. f . Since λ evolves independently of f the following applies for general paths of λ_t .

Suppose $\{K_t^*\}$ solves the sequence problem for some f_0^* and λ_0 , and let $c_t^* = f_t^* - k_t^*$. Now consider $\hat{f}_0 = \alpha f_0^*$ for constant $\alpha > 0$. Suppose the optimal policy for \hat{f}_0 and λ_0 is $\{\hat{K}_t\}$, and that $\{\alpha K_t^*\}$ is not optimal. Clearly the policy αK_t^* defines a feasible consumption plan αc_t^* . Then optimality of $\{\hat{K}_t\}$, $\hat{c}_t = \hat{f}_t - \hat{k}_t$, and homogeneity of \bar{u} in c imply

$$\begin{aligned} E_0 \left[\sum_{t=0}^{\infty} \beta^t \bar{u}(\hat{c}_t, \lambda_t) \right] &> E_0 \left[\sum_{t=0}^{\infty} \beta^t \bar{u}(\alpha c_t^*, \lambda_t) \right] \\ E_0 \left[\sum_{t=0}^{\infty} \beta^t \bar{u}(\hat{c}_t, \lambda_t) \right] &> \alpha^{1-\gamma} E_0 \left[\sum_{t=0}^{\infty} \beta^t \bar{u}(c_t^*, \lambda_t) \right] \\ E_0 \left[\sum_{t=0}^{\infty} \beta^t \bar{u}\left(\frac{\hat{c}_t}{\alpha}, \lambda_t\right) \right] &> E_0 \left[\sum_{t=0}^{\infty} \beta^t \bar{u}(c_t^*, \lambda_t) \right] \end{aligned} \quad (3.30)$$

But this contradicts the assumption that $\{K_t^*\}$ is optimal for f_t^* , since $\{\frac{\hat{K}_t}{\alpha}\}$ is feasible and supports the superior consumption plan $\{\frac{\hat{c}_t}{\alpha}\}$. Therefore it must be that if $\{K_t^*\}$ solves the problem for f_0^* , then $\{\alpha K_t^*\}$ solves it for αf_0^* . Noting that the recursive formulation of the policy $K(f, \lambda)$ must also solve the sequence problem yields the decomposition $K(f, \lambda) = fB(\lambda)$. Equation (3.12) derives readily from Equation (3.11). \square

Assumption 1. Let $\hat{\beta} = \sup_{\lambda} \beta E \left[\left(\frac{1 - b(\tilde{z}\lambda)}{1 + (\tilde{z}\lambda)^{1/\gamma}} \right)^{-\gamma} \left(\frac{1 - b(\lambda)}{1 + \lambda^{1/\gamma}} \right)^\gamma (\tilde{A}B(\lambda))^{1-\gamma} \right]$. Assume model parameters s.t. $\hat{\beta} < 1$.

Proposition 2. Define a space of continuous, bounded functions $\mathcal{D}(\Lambda)$, $\Lambda \equiv \mathbb{R}_+$, $g \in \mathcal{D}(\Lambda)$ s.t. $g : \Lambda \rightarrow [0, 1]$, with the sup norm $\|g\| = \sup_{\lambda \in \Lambda} |g(\lambda)|$. Let \mathcal{T}_D be the mapping given by Equation (3.18),

$$[\mathcal{T}_D g](\lambda) = \frac{1 - b(\lambda)}{1 + \lambda^{1/\gamma}} + \beta E \left[\left(\frac{1 - b(\tilde{z}\lambda)}{1 + (\tilde{z}\lambda)^{1/\gamma}} \right)^{-\gamma} \left(\frac{1 - b(\lambda)}{1 + \lambda^{1/\gamma}} \right)^\gamma g(\tilde{z}\lambda) (\tilde{A}B(\lambda))^{1-\gamma} \right]. \quad (3.31)$$

Under Assumption 1, there is a unique solution $D(\lambda) \in \mathcal{D}(\Lambda)$ to Equation (3.18), and $\forall g \in \mathcal{D}(\Lambda)$, $\lim_{N \rightarrow \infty} [\mathcal{T}_D^N g](\lambda) \rightarrow D(\lambda)$.

Proof. The result follows from the contraction mapping theorem. I first apply Blackwell's sufficient conditions to demonstrate that \mathcal{T}_D is a contraction. Let $g, h \in \mathcal{D}(\Lambda)$ be s.t. $g(\lambda) \leq h(\lambda)$, $\forall \lambda \in \Lambda$. Then \mathcal{T}_D is monotonic, i.e.,

$$[\mathcal{T}_D g](\lambda) \leq [\mathcal{T}_D h](\lambda), \forall \lambda \quad (3.32)$$

since for all realizations z, A of \tilde{z}, \tilde{A} we have

$$\left(\frac{1 - b(z\lambda)}{1 + (z\lambda)^{1/\gamma}} \right)^{-\gamma} g(z\lambda)(AB(\lambda))^{1-\gamma} \leq \left(\frac{1 - b(z\lambda)}{1 + (z\lambda)^{1/\gamma}} \right)^{-\gamma} h(z\lambda)(AB(\lambda))^{1-\gamma}, \quad (3.33)$$

all other terms on either side of the inequality being deterministic and identical. To demonstrate that \mathcal{T}_D has a discounting property, let $g \in \mathcal{D}(\Lambda)$ and $a \geq 0$, a constant. Then

$$\begin{aligned} [\mathcal{T}_D g + a](\lambda) &= \frac{1 - b(\lambda)}{1 + \lambda^{1/\gamma}} + \beta E \left[\left(\frac{1 - b(\tilde{z}\lambda)}{1 + (\tilde{z}\lambda)^{1/\gamma}} \right)^{-\gamma} \left(\frac{1 - b(\lambda)}{1 + \lambda^{1/\gamma}} \right)^\gamma (g(\tilde{z}\lambda) + a)(\tilde{A}B(\lambda))^{1-\gamma} \right] \\ &\leq [\mathcal{T}_D g](\lambda) + \hat{\beta} a. \end{aligned} \quad (3.34)$$

Therefore \mathcal{T}_D is a contraction, and the proposition follows by application of the contraction mapping theorem. \square

3.B Asymptotic Behavior of Optimist Consumption Share (ω)

Survival analysis for the case where agents have identical preferences and impatience but *different magnitudes* of error in their beliefs is presented by Yan (2008). For the case where errors in beliefs are of *identical magnitude* but where agents nonetheless disagree, e.g., the case of an equally erroneous optimist and pessimist, Yan states that neither agent's consumption share converges to zero. However he does not formally characterize its asymptotic behavior. I do so below for a simple representative case.

The analysis is also novel in that the setting is a discrete-time, discrete state economy. I use the baseline parameters listed in Table 3.1, although I relax the assumption that $\theta^O = \theta^P = 0.5$. Instead I assume for convenience (but w.l.o.g. for asymptotic analysis) that $\lambda_0 = 1$. I focus purely on the asymptotic behavior of the stochastic weighting factor λ and the consequences for consumption share $\omega(\lambda)$. Thus the analysis

here is relevant to either an endowment economy or a production economy: the behavior of λ is determined by disagreement regarding the states of nature, and is independent of whether those states correspond to different realizations of TFP or directly to different levels of aggregate output. The issue is how the pie is split, not its size or origin.

Suppose $\tilde{z}_1 = 3/2$, i.e., a low state L is realized. Then $\lambda_1 = 3/2$, i.e., P's consumption share increases. If $\tilde{z}_2 = 3/2$, another L realization, then $\lambda_2 = (3/2)^2$. A subsequent high state H would bring $\tilde{z}_3 = 2/3$ and $\lambda_3 = 3/2$, moving consumption share back towards O. It should be clear by now that, due to the symmetry of the problem, the state space of λ is

$$\Lambda = \dots, \left(\frac{3}{2}\right)^{-2}, \left(\frac{3}{2}\right)^{-1}, \left(\frac{3}{2}\right)^0, \left(\frac{3}{2}\right)^1, \left(\frac{3}{2}\right)^2, \dots \quad (3.35)$$

We can relabel to the states in Λ according to the exponent i as in $\left(\frac{3}{2}\right)^i$. Further, since realizations of \tilde{z} are determined under the Econometrician's measure, it is equally likely that λ will increase by a factor of $3/2$ ($i_t = i_{t-1} + 1$) or decrease by a factor of $2/3$ ($i_t = i_{t-1} - 1$). We can model λ by looking at the process for i , which is simply a symmetric random walk on the integers, a well known example of a Markov chain. It has the following properties ¹⁴:

1. It is recurrent, specifically null-recurrent.
2. Consequently any current state $i_t \in \mathbb{Z}$ is revisited with probability 1 as $t \rightarrow \infty$, so it cannot be that $\lambda \rightarrow 0$ or $\lambda \rightarrow \infty$.
3. Null-recurrent Markov chains have no stationary distribution.

Despite the absence of a stationary distribution, we can say something useful about the asymptotic distribution of λ . For $i_0 = 0$ the probability distribution of i_t is approximated by

$$P_t(i) = \begin{cases} \frac{2}{\sqrt{2\pi t}} e^{-i^2/2t}, & t \bmod 2 = i \bmod 2 \\ 0, & t \bmod 2 \neq i \bmod 2 \end{cases} \quad (3.36)$$

¹⁴See for example Cox and Miller (1980), Hoel, Port, and Stone (1972).

for values $|i|$ much smaller than t . This formula gives us the probability that i_t takes a value close to the origin for large t . Taking limits, we can see that for any finite distance d from the origin the probability $|i| < d$ is

$$P[|i| < d] = \lim_{t \rightarrow \infty} 2 \sum_{i=0}^d P_t(i) = 0, \quad (3.37)$$

since $P_t(i)$ takes on vanishingly small values near the origin for large t and there are only finitely many states within d of the origin. Since this is a symmetric random walk, it is equally likely that i will be either very negative or very positive at long time horizons. It follows that λ_t is asymptotically very nearly 0 or very nearly infinite, each with a roughly 50% chance, but it converges to neither. Consequently consumption share will belong almost entirely to one agent or the other asymptotically, but neither agent is extinguished.