The Financialization of Storable Commodities

Steven D. Baker *

sdbaker@cmu.edu

Carnegie Mellon University

Pittsburgh, PA

November 30, 2012

Abstract

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.

*Thanks to Bryan Routledge for numerous helpful discussions, and to workshop participants at Carnegie Mellon and the Federal Reserve Board for comments. Additional thanks to Emilio Osambela and Burton Hollifield. For the most recent version of this paper, please see http://www.andrew.cmu.edu/user/sdbaker/papers/bakerOilFinancialization.pdf

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, 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

[1]For 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 et al. [2009].

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

[2] See for example the GSCI fact sheet at

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

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 frequencies, 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, 1996]. Its appeal is simplicity coupled with an ability to produce autocorrelated spot prices and occasional, dramatic price "spikes" characteristic of the data. Routledge et al. [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

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

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 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 et al. [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 et al. [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 et al. [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.

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.

2.1 Markets

Before proceding 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) $\varphi_t(s_{t+1} - f_t) - \tau f_t \varphi_t^2$.[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

[4]The 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.

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].

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 et al. [2009] and Casassus et al. [2009], whereas Carlson et al. [2007] takes account of the fact that oil is an exhaustible resource.

[5] Abstracting 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.

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

[7] This 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 $f(K_t, L_t) = a_t K_t^{\phi}$. (Assume N firms, each endowed with $L = 1/N$ labor, and let $N \to \infty$.) 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.

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

$$\begin{aligned}
&\max_{\{I_t,\phi_t\}_{t=0}^{\infty}} \sum_{t=1}^{\infty}(1+r)^{-t}\left(E_0[p_t^p] - \frac{\theta}{2}Var_0[p_t^p]\right), \\
&\text{s.t. } p_{t+1}^p = s_{t+1}a_{t+1}(1+I_t^{1/2}) - (1+r)I_t + \phi_t(s_{t+1} - f_t), \\
&\qquad I_t \geq 0,\ \forall t
\end{aligned} \tag{1}$$

for I_{-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

$$Var_t[p_{t+1}^p] = \underbrace{Var_t[s_{t+1}a_{t+1}]}_{\sigma^2_{sa,t}}(1+I_t^{1/2})^2 + 2\underbrace{Cov_t[s_{t+1}a_{t+1}, s_{t+1}]}_{\sigma_{sas,t}}(1+I_t^{1/2})\phi_t + \underbrace{Var_t[s_{t+1}]}_{\sigma^2_{s,t}}\phi_t^2 \tag{2}$$

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

$$\begin{aligned}
\phi_t &= \frac{E_t[s_{t+1}] - f_t - \theta\sigma_{sas,t}(1+I^{1/2})}{\theta\sigma^2_{s,t}} \\
I_t &= \left(\frac{\sigma^2_{s,t}E_t[s_{t+1}a_{t+1}] - \sigma_{sas,t}(E_t[s_{t+1}] - f_t) - \theta(\sigma^2_{s,t}\sigma^2_{sa,t} - \sigma^2_{sas,t})}{\theta\sigma^2_{s,t}\sigma^2_{sa,t} + \theta\sigma^2_{sas,t} + 2\sigma^2_{s,t}(1+r)}\right)^2
\end{aligned} \tag{3}$$

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.

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

$$\begin{aligned} &\max_{\{Q_t,\psi_t\}_{t=0}^{\infty}} \sum_{t=1}^{\infty}(1+r)^{-t}\left(E_0[p_t^d] - \frac{\rho}{2}Var_0[p_t^d]\right), \\ &\text{s.t. } p_{t+1}^d = s_{t+1}Q_t - (1+r)(s_t+k)Q_t + \psi_t(s_{t+1} - f_t), \\ &\quad Q_t \geq 0, \ \forall t. \end{aligned} \tag{4}$$

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

$$\begin{aligned} Q_t &= \frac{E_t[s_{t+1}] - (1+r)(s_t+k)}{\rho\sigma_{s,t}^2} - \psi_t, \\ \psi_t &= \frac{E_t[s_{t+1}] - f_t}{\rho\sigma_{s,t}^2} - Q_t. \end{aligned} \tag{5}$$

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

$$\begin{aligned} Q_t &= 0 \\ \psi_t &= \frac{E_t[s_{t+1}] - f_t}{\rho\sigma_{s,t}^2}. \end{aligned} \tag{6}$$

[8] An 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 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 (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.

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

$$c_t = A(c_{x,t}, c_{y,t}) = \left[(1-\gamma)c_{x,t}^{\eta} + \gamma c_{y,t}^{\eta}\right]^{1/\eta}, \tag{7}$$

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

$$c_t = A(c_{x,t}, c_{y,t}) = c_{x,t}^{1-\gamma} c_{y,t}^{\gamma}. \tag{8}$$

Utility is time additively separable, given in each period by

$$u(c_t) = log(c_t). \tag{9}$$

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

$$\hat{x}_t = 1 + \varphi_{t-1}(s_t - f_{t-1}) - \tau f_{t-1}\varphi_{t-1}^2. \tag{10}$$

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

$$\begin{aligned} \max_{\{c_{x,t},c_{y,t},\varphi_t\}_{t=0}^{\infty}} (1-\beta)E_0 \sum_{t=0}^{\infty} \beta^t \log(A(c_{x,t}, c_{y,t})) \\ s.t.\ c_{x,t} + s_t c_{y,t} \leq \hat{x}_t, \\ \hat{x}_t = 1 + \varphi_{t-1}(s_t - f_{t-1}) - \tau f_{t-1}\varphi_{t-1}^2 \end{aligned} \tag{11}$$

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

$$\begin{aligned} c_{y,t} &= \hat{x}_t s_t^{-1} \left[1 + \left(\frac{\gamma}{(1-\gamma)s_t^{\eta}}\right)^{\frac{1}{\eta-1}}\right]^{-1} \\ c_{x,t} &= \hat{x}_t - s_t c_{y,t} \\ 0 &= E_t\left[\frac{s_{t+1} - f_t - 2\tau f_t \varphi_t}{1 + \varphi_t(s_{t+1} - f_t) - \tau f_t \varphi_t^2}\right]. \end{aligned} \tag{12}$$

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.

3 Equilibrium

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

$$y_t = a_t(1 + I_{t-1}^{1/2}). \tag{13}$$

In each period, commodity market clearing requires

$$y_t - c_{y,t} = Q_t - Q_{t-1} = \Delta Q_t \tag{14}$$

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

$$(1 + \varphi_{t-1}(s_t - f_{t-1}) - \tau f_{t-1}\varphi_{t-1}^2)s_t^{-1}\left(1 + \left(\frac{\gamma}{(1-\gamma)s_t^{\eta}}\right)^{\frac{1}{\eta-1}}\right)^{-1} = y_t - \Delta Q_t, \tag{15}$$

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,

$$\phi_t + \varphi_t + \psi_t = 0. \tag{16}$$

Writing $\psi_t = -(\phi_t + \varphi_t)$ and substituting into dealer FOC Equation (5), the futures price is

$$f_t = E_t[s_{t+1}] - \rho\sigma_{s,t}^2\,(Q_t - (\phi_t + \varphi_t)) \tag{17}$$

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

$$Q_t = J(Q_{t-1}, f_{t-1}, \varphi_{t-1}, a_t), \tag{18}$$

spot price function

$$s_t = S(Q_{t-1}, f_{t-1}, \varphi_{t-1}, a_t), \tag{19}$$

and household portfolio function

$$\varphi_t = \Phi(f_t, s_{t+1}). \tag{20}$$

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

$$\begin{aligned} Q_t &= J(Q_{t-1}, a_{t-1}, a_t) \\ s_t &= S(Q_{t-1}, a_{t-1}, a_t) \\ \varphi_t &= \Phi(Q_t, a_t) \end{aligned} \tag{21}$$

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. Cafiero 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.

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. 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 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 et al. [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 prices for a rare, but essential, commodity. These features are seen in Figure 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.

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 is[9]

$$\frac{(1+r)f_{t,n} + f_{t+1,n-1} - f_{t,n}}{f_{t,n}} - 1 = \frac{f_{t+1,n-1}}{f_{t,n}} - 1 + r, \tag{22}$$

and so the excess holding return is

$$\frac{f_{t+1,n-1}}{f_{t,n}} - 1. \tag{23}$$

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

$$\frac{E_t[s_{t+1}] - f_t}{f_t} = \frac{\rho\sigma^2_{s,t}\,(Q_t + \psi_t)}{f_t} \tag{24}$$

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

[9]When 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, $E_t\left[\log(f_{t+1,n-1})\right] - \log(f_{t,n})$.

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.

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

$$\begin{aligned} -\frac{E_t[s_{t+1}] - f_t - \theta\sigma_{sas,t}(1+I^{1/2})}{\theta\sigma^2_{s,t}} - \varphi_t &= \frac{E_t[s_{t+1}] - f_t}{\rho\sigma^2_{s,t}} - Q_t \\ \Rightarrow \frac{E_t[s_{t+1}] - f_t}{f_t} &= \frac{\rho\theta}{f_t(\rho+\theta)}\left(\sigma_{sas,t}(1+I^{1/2}) + \sigma^2_{s,t}(Q_t - \varphi_t)\right). \end{aligned} \tag{25}$$

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

$$\frac{E_t[s_{t+1}] - f_t}{f_t} = \frac{\rho\theta}{f_t(\rho+\theta)}\left(\sigma^2_{s,t}Q_t + \sigma_{sas,t}(1+I^{1/2})\right). \tag{26}$$

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 (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

$$\begin{aligned} &\max_{Q_t,\psi_t} E_t[\hat{p}^d_{t+1}] - \frac{\rho}{2}Var_t[\hat{p}^d_{t+1}], \\ &\text{s.t. } \hat{p}^d_{t+1} = p^d_t + \psi_{t,n}(f_{t+1,n-1} - f_{t,n}), \\ &p^d_{t+1} = s_{t+1}Q_t - (1+r)(s_t+k)Q_t + \psi_{t,1}(s_{t+1} - f_{t,1}). \end{aligned} \tag{27}$$

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

$$f_{t,n} = E_t[f_{t+1,n-1}] - \rho(Q_t + \psi_{t,1})Cov(s_{t+1}, f_{t+1,n-1}) \tag{28}$$

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.

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 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 2, US production actually declines by roughly 25 % over the

[10] www.eia.gov

[11] GDP data is from the BEA.

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 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.

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 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 et al. [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 2.

Detailed data on futures positions is not freely available, so I rely upon references for summary statistics. Acharya et al. [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 finan-

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

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

[14]Quarterly 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 futher 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.

cial 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 adjusting only one parameter.

Parameter values for the baseline calibration are given in Table 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.

5 Results

Given the simplicity of the model, the baseline calibration performs fairly well overall. Table 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.

[15]Data 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/)

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 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 autocorrellated (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 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 7. The ability of the storage model to match unconditional standard deviation over the term structure was emphasized in Routledge et al. [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 8, with a slope similar to that in the data.

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 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.

5.2.1 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 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 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 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.

[16]The 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.

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

5.2.2 Households

The effects of household entry upon the commercial sector are almost directly offset by the effects upon households themselves. Figure 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 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.

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 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

[17] For 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.

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 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.

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.

References

V. Acharya, L. Lochstoer, and T. Ramadorai. Limits to arbitrage and hedging: Evidence from commodity markets. Technical report, NYU, 2012.

D. M. Arseneau and S. Leduc. Commodity price movements in a general equilibrium model of storage. Technical report, Federal Reserve Board, 2012.

S. D. Baker and B. R. Routledge. The price of oil risk. Technical report, Carnegie Mellon University, 2012.

H. . Bessembinder, A. Carrion, K. Venkataraman, and L. Tuttle. Predatory or Sunshine Trading? Evidence from Crude Oil ETF Rolls. *SSRN eLibrary*, 2012. doi: 10.2139/ssrn.2026802.

C. Brunetti and D. Reiffen. Commodity index trading and hedging costs. *Finance and Economics Discussion Series*, 2011.

B. Buyuksahin, M. Haigh, J. Harris, J. Overdahl, and M. Robe. Fundamentals, trader activity and derivative pricing. Technical report, American University, 2011.

R. Caballero, E. Farhi, and P. Gourinchas. Financial crash, commodity prices and global imbalances. Technical report, National Bureau of Economic Research, 2008.

C. Cafiero, E. Bobenrieth H, J. Bobenrieth H, and B. Wright. The empirical relevance of the competitive storage model. *Journal of Econometrics*, 162(1):44–54, 2011.

M. Carlson, Z. Khokher, and S. Titman. Equilibrium exhaustible resource price dynamics. *The Journal of Finance*, 62(4):1663–1703, 2007.

J. Casassus, P. Collin-Dufresne, and B. Routledge. Equilibrium commodity prices with irreversible investment and non-linear technology. Technical report, National Bureau of Economic Research, 2009.

CFTC. Staff report on commodity swap dealers and index traders with commission recommendations. Technical report, Commodity Futures Trading Commission, September 2008.

A. Deaton and G. Laroque. On the behaviour of commodity prices. *Review of Economic Studies*, 59:1–23, January 1992.

A. Deaton and G. Laroque. Competitive storage and commodity price dynamics. *Journal of Political Economy*, 104:896–923, October 1996.

M. Floden. A note on the accuracy of markov-chain approximations to highly persistent ar (1) processes. *Economics Letters*, 99(3):516–520, 2008.

G. Gorton, F. Hayashi, and K. Rouwenhorst. The fundamentals of commodity futures returns. *Review of Finance*, 2012.

J. Hamilton and J. Wu. Risk premia in crude oil futures prices. Technical report, UCSD, 2011.

J. D. Hamilton. Oil and the macroeconomy. In S. N. Durlauf and L. E. Blume, editors, *The New Palgrave Dictionary of Economics*. Palgrave Macmillan, 2008.

O. Hart. On the optimality of equilibrium when the market structure is incomplete. *Journal of Economic Theory*, 11(3):418–443, 1975.

D. Hirshleifer. Hedging pressure and futures price movements in a general equilibrium model. *Econometrica: Journal of the Econometric Society*, pages 411–428, 1990.

S. Irwin and D. Sanders. Index funds, financialization, and commodity futures markets. *Applied Economic Perspectives and Policy*, 33(1):1–31, 2011.

S. Irwin, D. Sanders, R. Merrin, et al. Devil or angel? the role of speculation in the recent commodity price boom (and bust). *Journal of Agricultural and Applied Economics*, 41(2):377–391, 2009.

L. Kilian and R. Vigfusson. Do oil prices help forecast us real gdp? the role of nonlinearities and asymmetries. Technical report, University of Michigan, 2012.

L. Kogan, D. Livdan, and A. Yaron. Oil futures prices in a production economy with investment constraints. *The Journal of Finance*, 64(3):1345–1375, 2009.

X. N. Pan. Investor Beliefs and State Price Densities in the Crude Oil Market. *SSRN eLibrary*, 2011. doi: 10.2139/ssrn.1988449.

R. Ready. Oil prices and long-run risk. Technical report, University of Rochester, 2012.

B. R. Routledge, D. J. Seppi, and C. S. Spatt. Equilibrium forward curves for commodities. *Journal of Finance*, 60:1297–1338, June 2000.

K. J. Singleton. Investor Flows and the 2008 Boom/Bust in Oil Prices. *SSRN eLibrary*, 2011. doi: 10.2139/ssrn.1793449.

H. Stoll and R. Whaley. Commodity index investing and commodity futures prices. *Journal of Applied Finance*, 20(1):7–46, 2010.

K. Tang and W. Xiong. Index investment and financialization of commodities. Technical report, Princeton University, 2011.

G. Tauchen. Finite state markov-chain approximations to univariate and vector autoregressions. *Economics letters*, 20(2):177–181, 1986.

J. Williams and B. Wright. *Storage and Commodity Markets*. Cambridge University Press, 1991.

B. Wright and J. Williams. The welfare effects of the introduction of storage. *The Quarterly Journal of Economics*, 99(1):169–192, 1984.

Table 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 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 et al. [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 financialiation (decreased transaction costs, given by τ). The baseline calibration is desiged 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.

Figure 1: Real Spot Prices (WTI Oil)

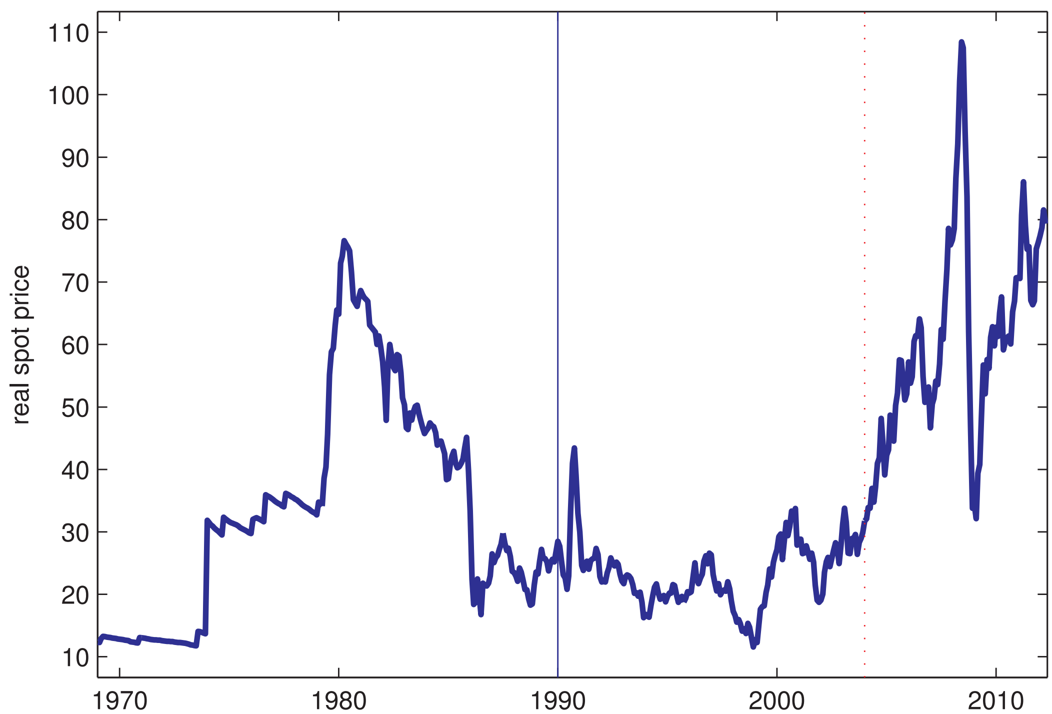

The plot above shows real spot prices for WTI crude oil over time. Data is from the St. Louis Fed. I focus on the period from 1990-2012, for which futures prices are also available (marked with the solid vertical line). I split the sample into 1990-2003 and 2004-2012 (marked with the dashed vertical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

Figure 2: Oil Consumption and Production

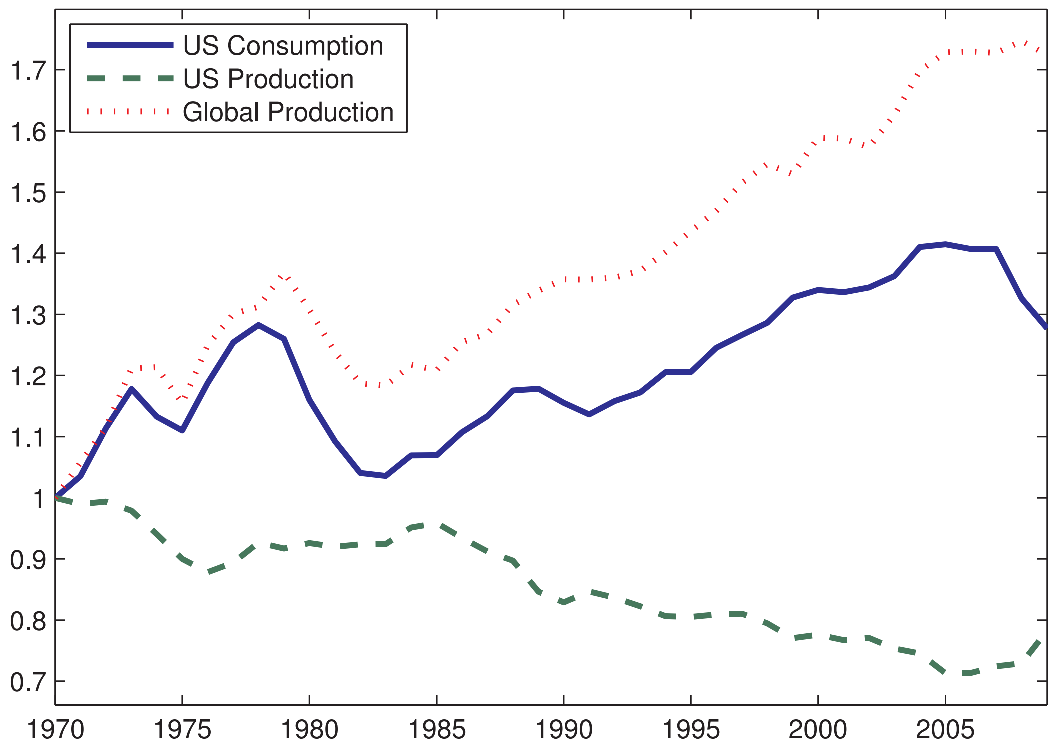

Annual US oil consumption is shown alongside worldwide and US production. I normalize each series by its initial value. Because oil is a globally traded commodity, the path of US consumption follows global production, not US production.

Figure 3: Detrended Oil Consumption and Production

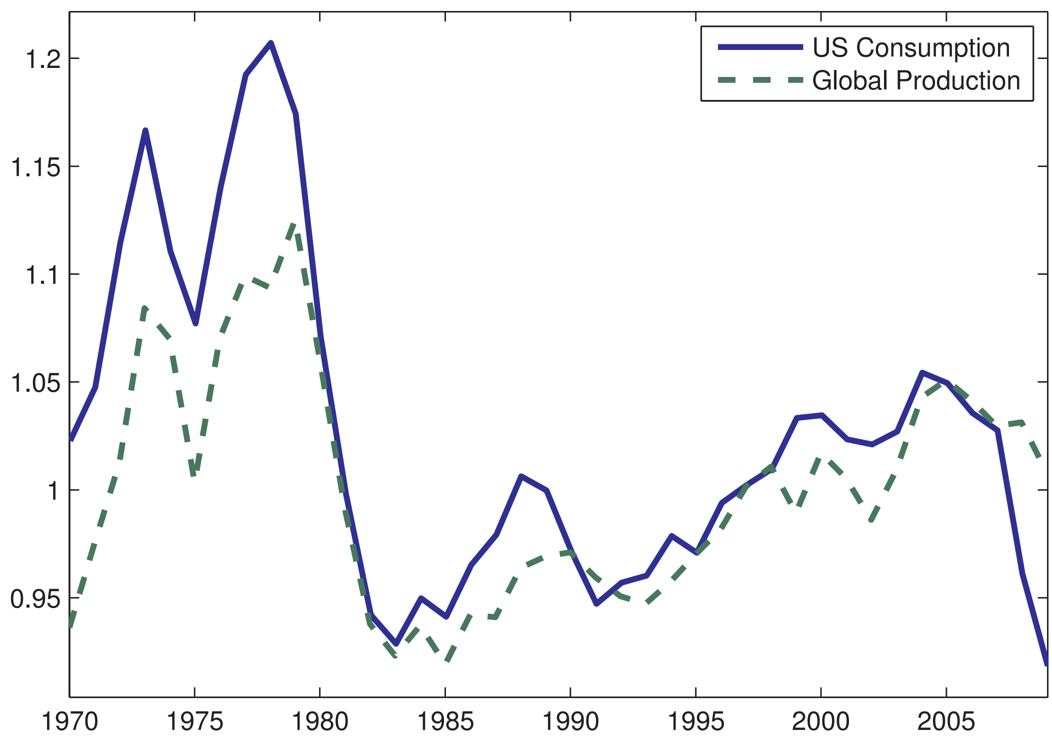

The figure shows detrended annual US consumption and global production. For each series I estimate a linear time trend via OLS, then divide each series by the trend. The two detrended series are very correlated (Pearson's coefficient of 0.85).

Figure 4: Inventory Policy

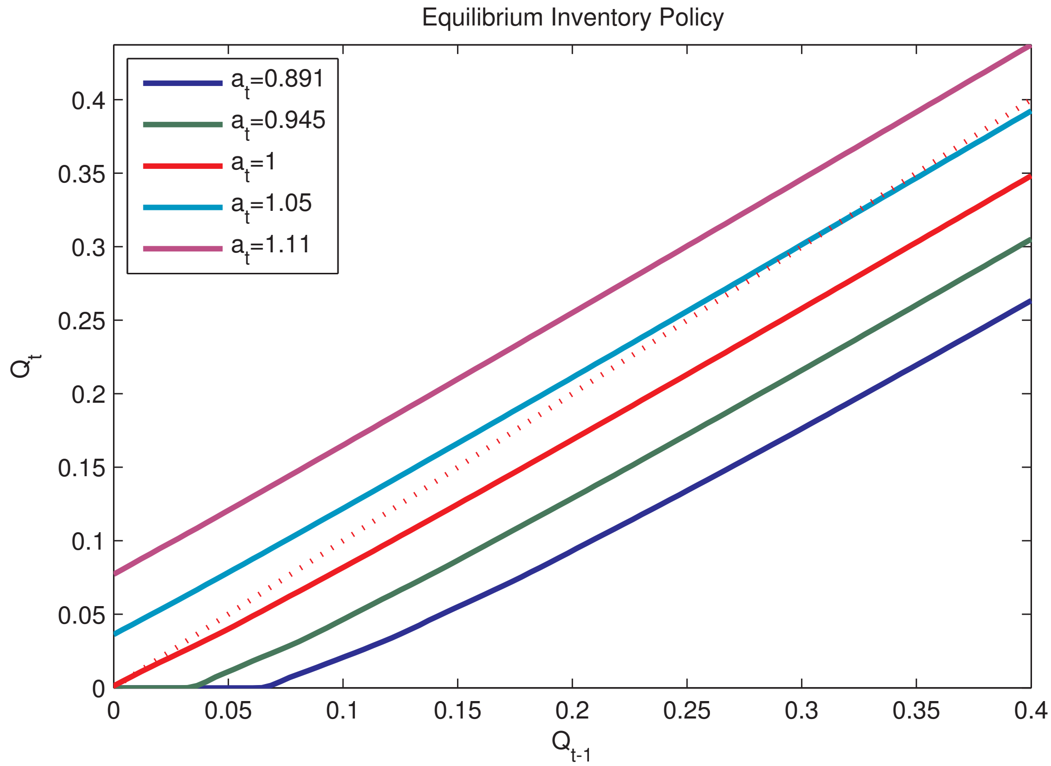

The plot shows optimal inventory policy for the baseline model. The dealer's end of period inventory Q_t depends on inventory carried over from the previous period (Q_{t-1}, given on the x-axis), and the productivity realization (a_t, shown by the different curves). Inventory policy is also affected 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 5: State-contingent Spot Prices

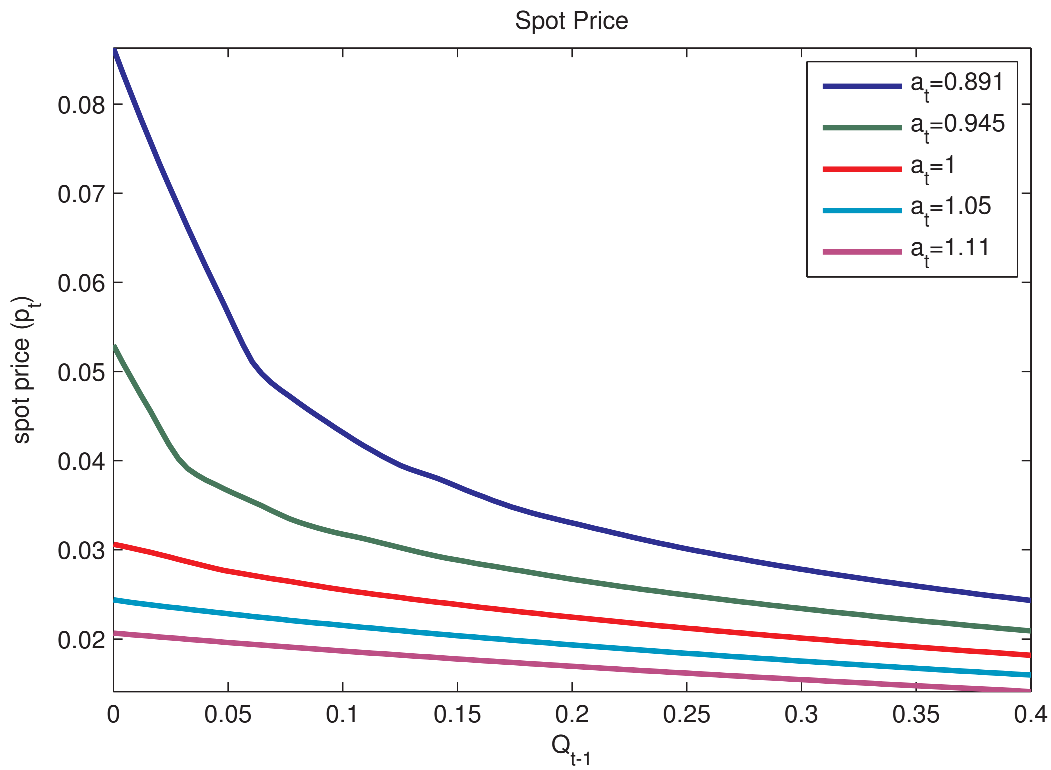

For the baseline calibration, spot prices are shown for entering inventories (Q_{t-1}, given on the x-axis), and each possible productivity realization (a_t, shown by the different curves). If there's plenty of inventory, enough will be sold in response to poor productivity that the price impact of the 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 would of course mean more profits for him! The trade-off is governed by the cost of storage, risk aversion, and opportunities in the futures market.

Figure 6: Mean Futures Curve: 1990-2003

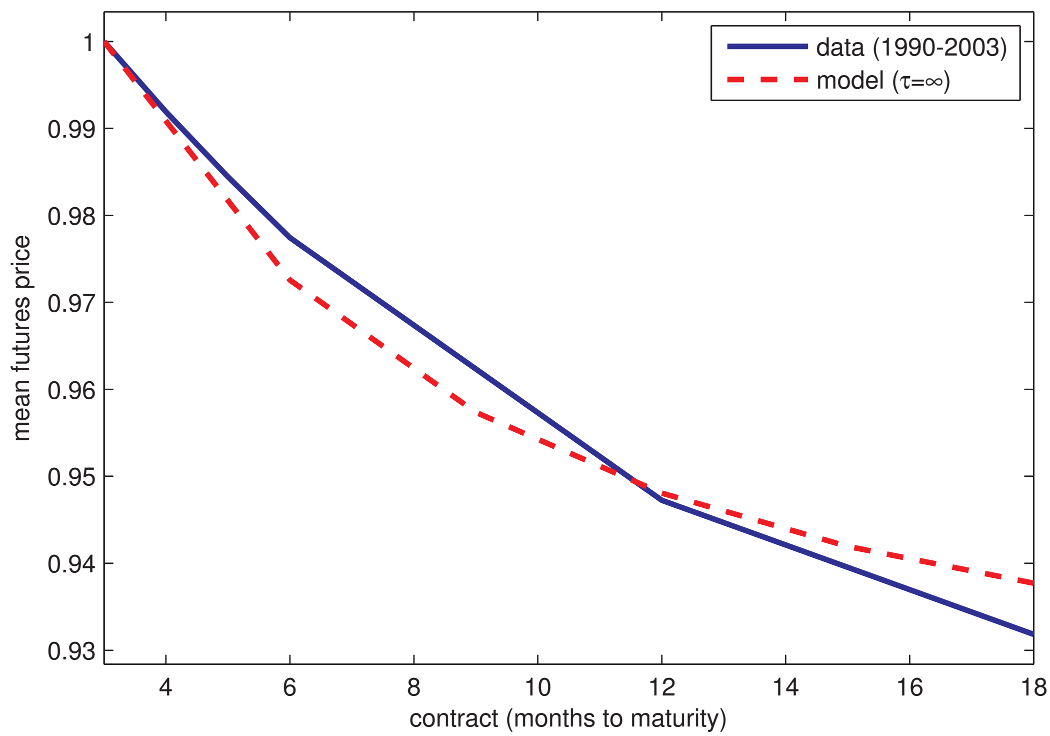

The plot shows the mean futures curve before financialization, in the data (1990-2003) and the model ($\tau = \infty$). Each curve is normalized by the price of its front contract. The model provides a good approximation to the shape of the mean futures curve: the relationship between futures of different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

Figure 7: Term Structure of Futures Volatility: 1990-2003

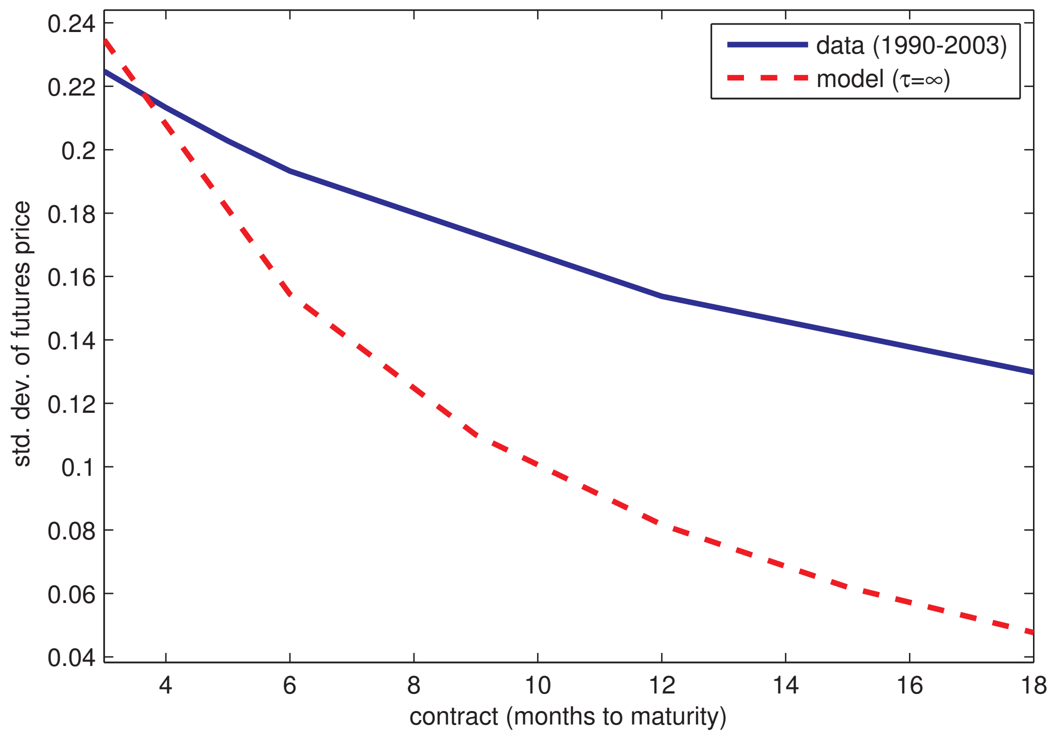

The plot shows the unconditional quarterly standard deviation of futures contracts before financialization, in the data (1990-2003) and the model ($\tau = \infty$). Standard deviations are expressed as percentages of the mean price of the front contract. The model matches the standard deviation of the 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 8: Expected Excess Returns on Futures: 1990-2003

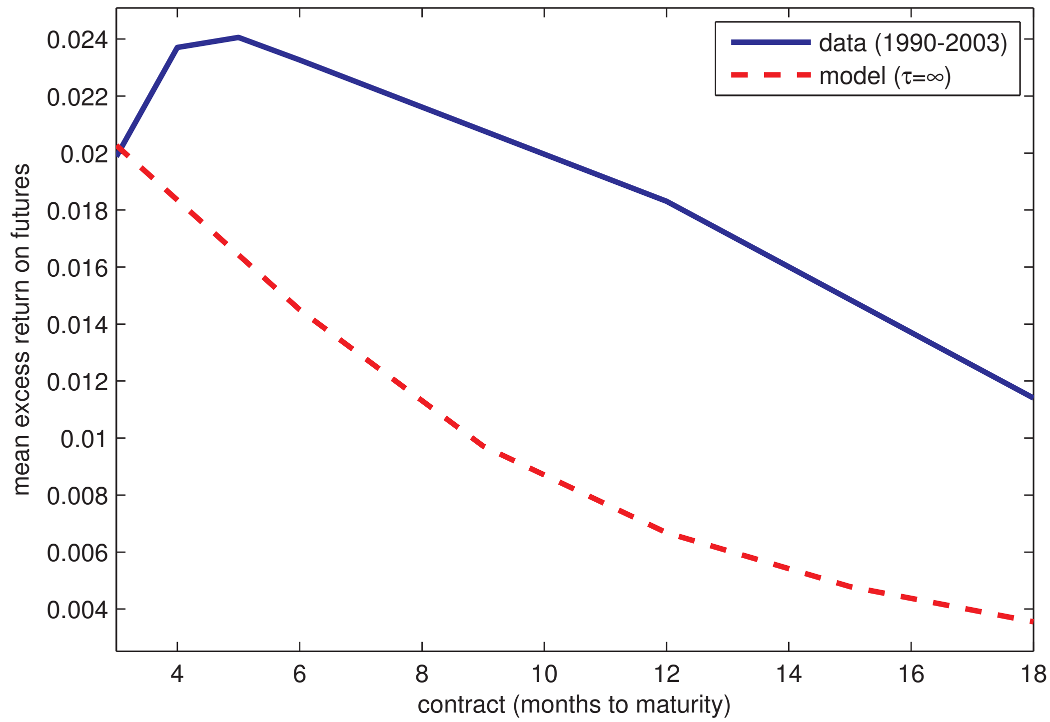

The plot shows the unconditional expected excess quarterly holding returns (the risk premium) on futures contracts before financialization, in the data (1990-2003) and the model ($\tau = \infty$). The model does well at matching the risk premium of the front (3-month) contract, and approximates the slope 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 9: Dealer Net Futures Position

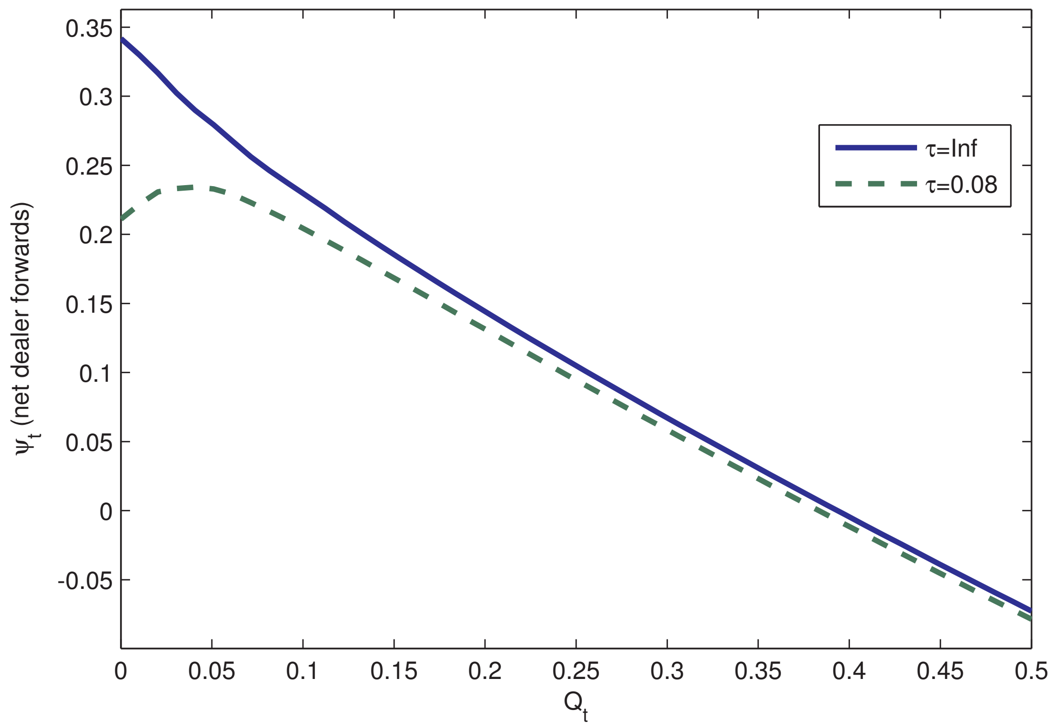

Above we see the dealer's net position in the futures market before ($\tau = \infty$) and after ($\tau = 0.08$) financialization, plotted against inventory (Q_t) on the x-axis. Before financialization, the dealer is the only counterparty to the producer, who has a lot of exposure to oil spot price fluctuations. 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 10: Distribution of Inventory (Q)

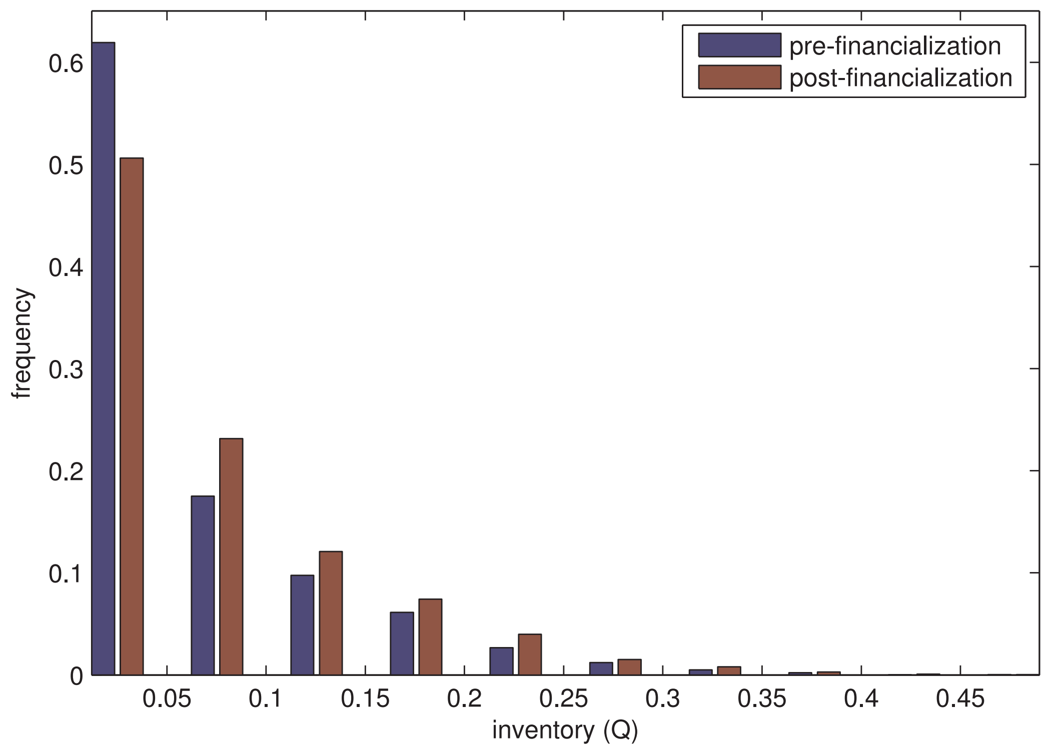

The histogram shows the unconditional distribution of inventory (Q) before and after financialization. Financialization reduces the frequency of very low inventory or stock-out conditions.

Figure 11: Financialization and Investment

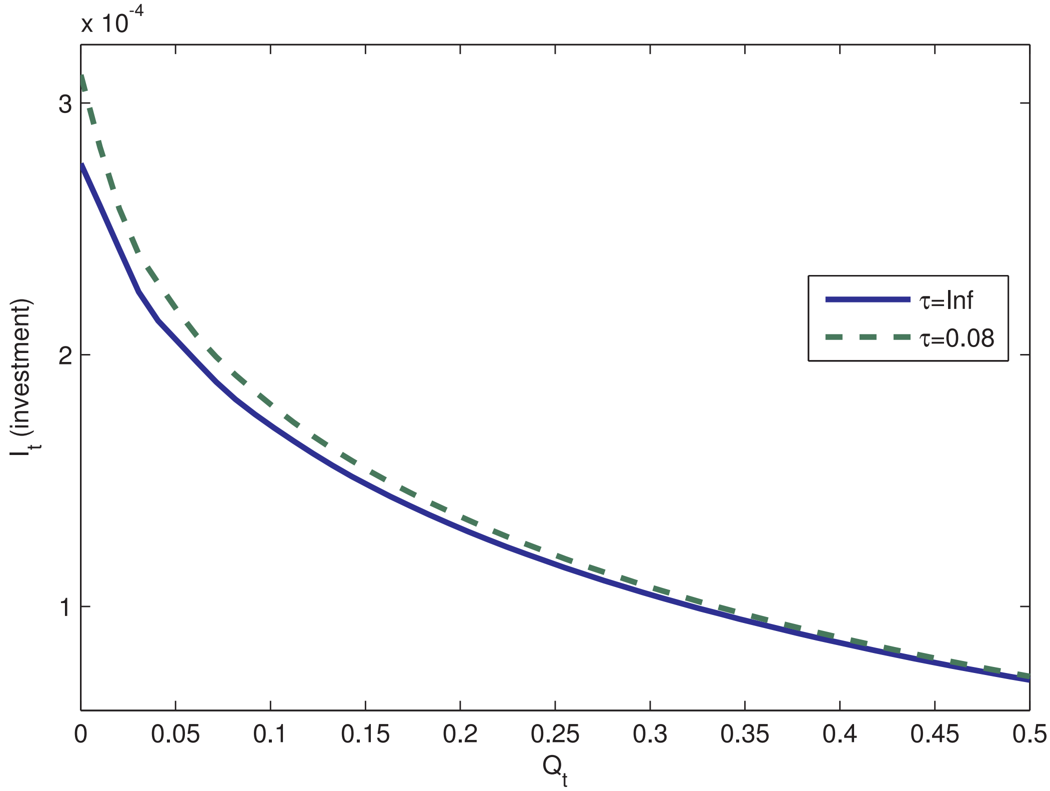

Financialization increases investment by the commodity producing firm. Taken in isolation, this should increase supply of the commodity relative to the numeraire, and so decrease the spot price.

Figure 12: Distribution of Spot Prices

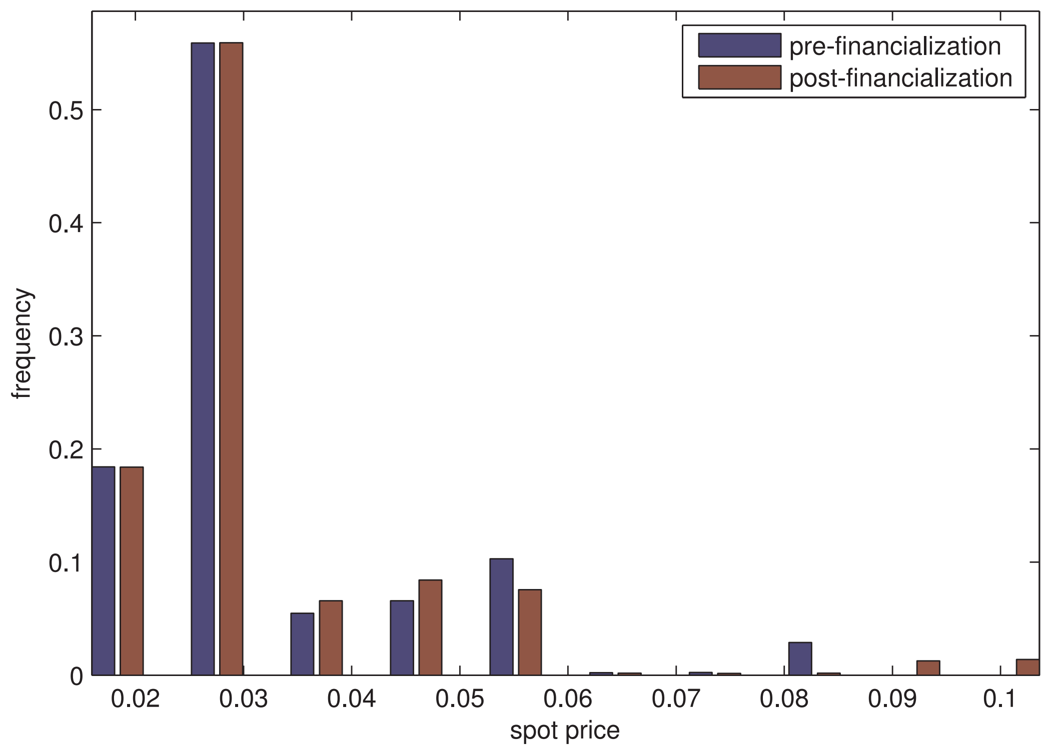

The histogram shows the unconditional distribution of the spot price (s) before and after financialization. The results reflect the offsetting effects of financialization on inventory dynamics and the household budget. Low-inventory states are less likely after financialization, which tends to reduce 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 13: State-contingent Spot Prices and Financialization

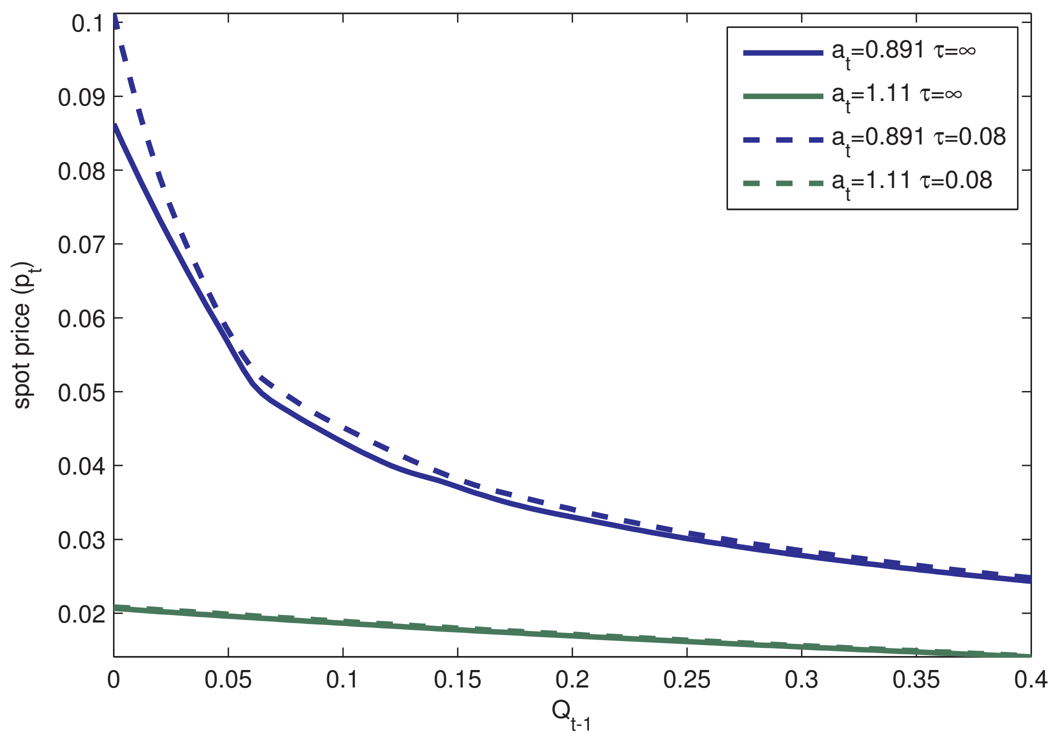

State-contingent spot prices are shown before and after financialization, for the most extreme productivity realizations. Although spot prices in high-productivity states are little changed, prices are elevated in low-productivity states after financialization. Stockouts lead to more extreme price spikes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

Figure 14: Expected Household Utility

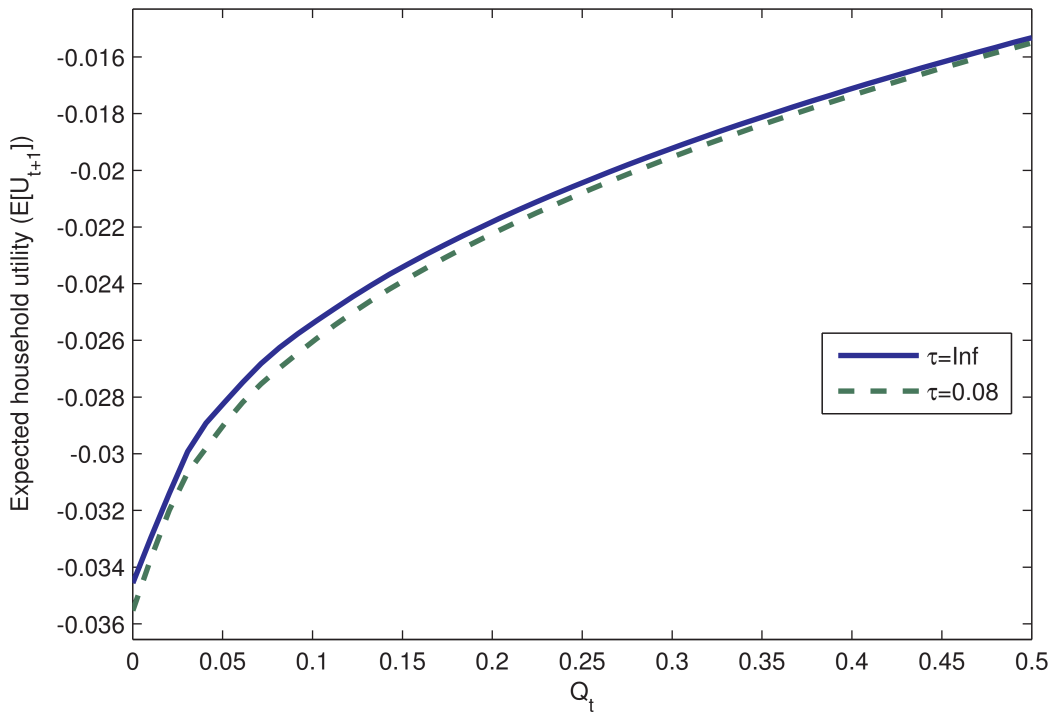

The plot shows expected next-period household utility conditional on the level of inventory. On a conditional basis, financialization reduces household utility. Although futures allow households to hedge commodity price risk, they also make spot prices more volatile. The net affect is a reduction in 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.