

The Financialization of Storable Commodities

Steven D. Baker *

April 5, 2019

Abstract

I solve a dynamic equilibrium model of commodity spot and futures prices, incorporating an active futures market, heterogeneous risk-averse participants, and storage. When calibrated to data from the crude oil market, the model implies that financialization reduces the futures risk premium, and increases correlation between futures open interest and the spot price level. However there is no long-run increase in the mean spot price, and speculative storage generally attenuates financialization's effect on spot price volatility. Therefore financialization's effect on spot price dynamics through storage arbitrage is likely modest, even if futures positions and risk premia are substantially altered.

*McIntire School of Commerce, University of Virginia; sdb7e@virginia.edu. Thanks to Bryan Routledge, Burton Hollifield, Emilio Osambela, Antonio Gargano, Christian Heyerdahl-Larsen, and seminar participants at Carnegie Mellon University, the Federal Reserve Board, University of Calgary, University of Houston, London School of Economics, Boston College, University of Virginia, University of Rochester, the CFTC, and Finance Down Under, and the EFA for comments. A version of this paper was part of my dissertation at the Tepper School of Business, Carnegie Mellon University. Support from a William Larimer Mellon Fellowship is gratefully acknowledged.

1 Introduction

Since at least Keynes [1930], the dynamics of commodity futures markets have been informally characterized as the outcome of trade between commodity producers, who wish to hedge against an uncertain spot price, and speculative dealers. These dealers have no fundamental use for the commodity, but accept some risk for expected return, while simultaneously arbitraging away certain gains by carrying inventory to replicate the futures payoff. Characteristics of spot and futures markets – such as the shape of the futures curve or spot price volatility – vary over time with the state of production, the level of inventory, and the futures positions of market participants.

Surprisingly, the essential elements of this economy have never been formally studied together within a dynamic equilibrium model suitable for calibration to data. Existing theoretical studies can be coarsely divided into two categories. The first focuses on the dynamic storage mechanism, but does not incorporate an active futures market, in which the dealer trades futures with heterogeneous counterparties such as commodity producers or consumers. Thus a critical arbitrage relationship is present, in which futures and spot markets are linked through inventory, but meaningful trade in futures is absent. Examples in this category include work by Deaton and Laroque [1992, 1996], focusing on dynamic storage as a mechanism for introducing persistence into spot prices, and Routledge et al. [2000], studying futures and spot prices in a market populated by homogeneous dealers. A critical common ingredient in these models is costly and non-negative inventory, introducing the potential for stockouts, which is important for spot-price dynamics and the futures-spot relationship.

The second coarse category in the literature models active futures markets with heterogeneous participants such as dealers, producers, and consumers, but abstracts from dynamic storage. Thus futures markets are economically important, in facilitating risk-sharing for example, but a critical linkage between futures and spot markets is absent or greatly simplified. Examples in this category include work by Hirshleifer [1988, 1990], studying futures risk premia and hedging pressure with

trading costs. Hirshleifer [1989] is exceptional in combining an active futures market with storage, but under important simplifying assumptions, that storage is costless, or that time ends after two periods. By contrast, models in the first category feature costly storage and are generally stationary, making them more suitable for calibration to data.

This paper unites the two strands of literature. I formulate, solve, and compare to data a dynamic equilibrium model of commodity prices, incorporating an active futures market, heterogeneous risk-averse participants, and storage. Because the model assumes an infinite time horizon and is stationary, I am able to calibrate it to empirical moments, and to assess in magnitude the effect of key parameters and modeling assumptions.

The need for such a study has become more apparent in the last decade, as recent empirical works such as Buyuksahin et al. [2011], Irwin et al. [2009], and Tang and Xiong [2012] have documented a range of changes in spot and futures markets accompanying what is commonly called the financialization of commodities markets. Open interest in commodity futures contracts has risen along with the proliferation of commodity investment funds, with for example, crude oil futures open interest approximately tripling from 2004 to 2008, and remaining high but volatile through early 2015. New financial products, such as commodity exchange traded funds (ETFs), have made it practical even for ordinary household consumers to trade commodity futures. Therefore financialization relates to a key modeling assumption identified in Hirshleifer [1990]: whether transaction costs prevent commodity consumers from participating in the futures market. I subject consumers to such a transaction cost, and since I model the commodity as a consumption good, I refer to consumers as households.

I first calibrate the model to moments of prices and quantities prior to financialization. To assess the magnitudes and mechanisms of financialization's effects, I then reduce the cost to households of trading in the futures market, which is initially dominated by producers and dealers. Financialization in the model reproduces several changes observed since 2004 in futures market data,

including a change in the sign of expected excess quarterly returns on the three-month futures contract, which declines from 1.9% to -0.4% per quarter empirically, and from 2.8% to -0.8% per quarter in the model. In both the model and the data, there is a roughly 50% reduction in the frequency of a downward sloping futures curve (backwardation) following financialization.

The model also implies that futures open interest will become more correlated with the spot price, as observed in the data, because new futures market participants (households) choose to take a long position that is largest on average when the spot price is high. This is the model analog of the increasing relationship between commodity index fund investment and commodity spot prices, highlighted in Masters [2008], which drew political and regulatory scrutiny of financialization. Storage arbitrage is the linchpin of one argument for a causal linkage: if new entrants, demanding a long position, drive up the futures price, dealers can sell futures and buy physical inventory, causing the current spot price to rise until arbitrage profits are eliminated. However, this inventory will eventually be sold off, which will reduce the spot price in the future, leaving the mean effect on the spot price ambiguous. Furthermore, dealers anticipating the behavior of households (or commodity index funds) may hold more inventory in anticipation of future price swings. Therefore static analysis offers little guidance as to the long-run equilibrium effect, necessitating a dynamic inventory model.

Despite increased correlation between futures open interest and spot price levels, financialization causes no increase in the mean spot price in the dynamic model, even though mean open interest in futures increases. And while financialization does increase spot price volatility, the assumption that new entrants are commodity consuming households is central to the mechanism.

Two main channels connect increased futures trade to spot market dynamics: amplification via household hedging, and smoothing via dealer inventory. As in the data, households choose a long position in futures once transaction costs are lowered, to hedge their consumption risk. 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 hedging effect will be asymmetrical, and the mean spot price will increase. However a second 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 hedging 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 price volatility.

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.¹ 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 U.S. GDP growth [Kilian and Vigfusson, 2013]. 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 the possibility that financialization increases the spot price level or volatility. Tang and Xiong [2012] document increasing correlation between oil futures and non-energy commodity futures concurrent with increased index investment. The oil futures market is also one of the most liquid, with extensive trade in contracts up to three years from delivery and listings up to nine years from delivery (against four years for agricultural commodities such as

¹See for example the GSCI fact sheet at

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

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 five parameters used in calibration.

The model relates to the canonical commodity storage framework developed in Williams and Wright [1991], which is analyzed empirically in Deaton and Laroque [1992, 1996], and presented with several extensions in Pirrong [2011]. The competitive storage mechanism produces auto-correlated spot prices and occasional, dramatic price spikes characteristic of the data, even if the underlying production process is i.i.d. and normally distributed. 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. However these models abstract from actively traded futures contracts, focusing on either a competitive risk-neutral dealer or a representative household. I adopt the canonical model's infinite horizon, but combine it with elements of Hirshleifer [1989], modeling producers, dealers, and consumers as separate, risk-averse agents. Heterogeneous preferences and technological endowments motivate trade in spot and futures markets, and generate a time-varying risk premium.

Other recent empirical papers also focus on oil while analyzing the financialization of commodities. Singleton [2014] 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. 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 [2014] estimate a time-varying risk premium on oil futures using a vector autoregression (VAR) that incorporates the position of index traders. The spot price and the position of index traders is exogenous in their analysis, whereas I endogenize the spot price and the trading decisions of new entrants.

Although my main analysis is of entry into futures markets by household consumers, an extension considers entry by investors who do not consume the underlying commodity, but who add futures to their existing portfolio to diversify and earn a risk premium. Futures trade by portfolio investors can only indirectly affect inventory, and has no effect on consumer response to prices. Therefore entry by portfolio investors arguably has less potential to distort spot prices than entry by households. For example, entry by portfolio investors slightly decreases spot price volatility if they take a long position similar to commodity index funds, because dealers react by carrying more inventory on average; the mean spot price is approximately unchanged. An important caveat, studied in Goldstein and Yang [2016] and Sockin and Xiong [2015], is that portfolio investors might be privately informed regarding fundamentals, so that futures prices convey useful but noisy signals that influence production and consumption decisions. Their models complement mine, by incorporating asymmetric information and focusing on the information aggregation channel, while assuming a one-period spot market and abstracting from inventory.

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. Vercaemmen and Doroudian [2014] model financialization as portfolio diversification using agricultural futures, under the assumption that futures are replicated using inventory. I endogenize inventory, model portfolio choice for all market participants, and discuss hedging effects that offset inventory effects. Basak and Pavlova [2016] model financialization

as increased indexing to commodity futures by fund managers, and study changes across futures markets for different commodities. They consider implications for spot markets without dynamic inventory. Tang and Zhu [2016] study the increased use of commodities as collateral in China, using a two-period model with multiple countries and capital controls. Irwin and Sanders [2011] summarize additional literature on commodity financialization. No previous study of financialization combines an active futures market with a calibrated dynamic storage model of spot prices.

Several recent papers analyze structural models of oil markets without explicitly modeling financialization. In a frictionless DSGE model, Baker and Routledge [2015] show that changes in open interest and risk premia on oil futures arise endogenously as a result of heterogeneous risk-aversion. 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 two-period versions of the storage model with active futures markets, and find empirical support for the models' predictions. Acharya et al. [2013] 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. 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.

2 Model

I model a dynamic, stochastic, infinite-horizon economy with two goods: a numeraire general consumption good, and a commodity. The model is set in partial equilibrium: it is closed with respect

to the commodity, but not with respect to the numeraire, in keeping with the dynamic storage literature. 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 the futures market. All information is public, and time subscripts indicate the period in which a variable becomes measurable with respect to public information.

One interpretation of the model is that the household represents a nation participating competitively in an international market for the commodity, in which it trades with dealers and producers operating globally. Section 6 considers several extensions to the basic model, such as domestic or foreign demand shocks, the presence of portfolio investors who trade futures but do not trade the commodity, and full or partial household ownership of producer equity.

2.1 Producer

The producer is endowed with y_t units of the commodity in each period t , where y_t follows a finite-state Markov chain. He sells the commodity on the spot market at price s_t , and has mean-variance preferences over profits p_t^p , with risk-aversion coefficient $\theta > 0$. The producer maximizes expected utility by selecting position ϕ_t in the one-period futures contract with price f_t :

$$\begin{aligned} \max_{\phi_t} E_t[p_{t+1}^p] - \frac{\theta}{2} Var_t[p_{t+1}^p], \\ \text{s.t. } p_{t+1}^p = s_{t+1}y_{t+1} + \phi_t(s_{t+1} - f_t). \end{aligned} \tag{1}$$

Because the cost of entering a contract is zero, futures cannot be used to accumulate wealth over time. In combination with the price-taking assumption this makes the producer behave myopically, avoiding possible issues with dynamically inconsistent portfolio choice under mean-variance preferences, as discussed for example in Basak and Chabakauri [2010]. Producer risk-aversion is a simple proxy for corporate risk management motives, including managerial risk-aversion, avoiding financial distress, minimizing taxes or external financing costs, and others; see for example Tufano [1996] and Haushalter [2000] for empirical analysis of corporate hedging motives in com-

modity markets. Theoretical work on commodity markets often models producer hedging motives with mean-variance preferences, as for example in Hirshleifer [1988], Bessembinder and Lemmon [2002], Acharya et al. [2013], and Tang and Zhu [2016].

Decompose the variance of profits into

$$\text{Var}_t[p_{t+1}^p] = \underbrace{\text{Var}_t[s_{t+1}y_{t+1}]}_{\sigma_{sy,t}^2} + 2 \underbrace{\text{Cov}_t[s_{t+1}y_{t+1}, s_{t+1}]}_{\sigma_{sys,t}} \phi_t + \underbrace{\text{Var}_t[s_{t+1}]}_{\sigma_{s,t}^2} \phi_t^2. \quad (2)$$

The producer's optimal futures position is

$$\phi_t = \frac{E_t[s_{t+1}] - f_t - \theta \sigma_{sys,t}}{\theta \sigma_{s,t}^2}. \quad (3)$$

The expected excess payoff to a long position in the futures contract, $E_t[s_{t+1}] - f_t$, is the risk premium. The producer's problem captures elementary motives for trade in the futures market, in which hedging, captured by the term $\theta \sigma_{sys,t}$, and speculation, captured by the risk premium $E_t[s_{t+1}] - f_t$, may be complementary or offsetting motives that vary over time in magnitude or sign. For example, if $\sigma_{sys,t} > 0$, then the producer faces primarily "price risk," and profits are high if the spot price is high. If $\sigma_{sys,t} < 0$ then the producer faces primarily "quantity risk," and profits are low if the spot price is high. The selection of parameters supporting empirically plausible behavior is addressed in Section 4.

Although the risk premium and hedging motive are functions of endogenous prices, output dynamics are exogenous, which omits potentially interesting interplay between trade in futures and production decisions. Other papers study production dynamics, but abstract from active futures markets with endogenous trading and prices. In the most relevant case of oil, Kogan et al. [2009] and Casassus et al. [2009] study production models with irreversible investment, whereas Carlson et al. [2007] and David [2014] model oil as an exhaustible resource, and Hitzemann [2016] decomposes the reaction of a productive oil sector to long-run versus short-run macroeconomic shocks.

2.2 Dealer

The dealer is an intermediary who neither produces nor consumes the commodity. However he may trade futures, and he may also purchase the commodity on the spot market and store it for future sale. The dealer finances spot market purchases at interest rate r , and pays a numeraire fee of k per unit of commodity stored.²

The dealer has mean-variance preferences with risk-aversion parameter ρ over profits p_{t+1}^d , which he maximizes by choosing inventory q_t and trading ψ_t futures contracts:

$$\begin{aligned} \max_{\{q_t, \psi_t\}} E_t[p_{t+1}^d] - \frac{\rho}{2} \text{Var}_t[p_{t+1}^d], \\ \text{s.t. } p_{t+1}^d = s_{t+1}q_t - (1+r)(s_t+k)q_t + \psi_t(s_{t+1} - f_t), \\ q_t \geq 0. \end{aligned} \quad (4)$$

Inventory is marked-to-market each period to determine realized profits. In conjunction with price-taking, this implies that the dealer will behave myopically, even though aggregate inventory decisions can have a persistent effect on prices in equilibrium.

Inventory cannot be negative. When the constraint is nonbinding ($q_t > 0$), first order conditions imply

$$q_t = \frac{E_t[s_{t+1}] - (1+r)(s_t+k)}{\rho\sigma_{s,t}^2} - \psi_t, \quad (5)$$

$$\psi_t = \frac{E_t[s_{t+1}] - f_t}{\rho\sigma_{s,t}^2} - q_t. \quad (6)$$

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

²An 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 unreported extension of the model I implement both costs; for purposes of calibrating to oil the difference between the two is small.

solution to the dealer's problem is

$$q_t = 0, \quad (7)$$

$$\psi_t = \frac{E_t[s_{t+1}] - f_t}{\rho\sigma_{s,t}^2}. \quad (8)$$

In the canonical model of Deaton and Laroque [1992], dealers are called “speculators.” They are risk-neutral, so they hold inventory if in expectation its value will appreciate net of costs. To clear the futures market, the equilibrium risk premium must be zero, and there is no modeling of active trade in futures. My model nests risk neutral dealers as a special case, letting $\rho \rightarrow 0$.

I focus on the case of a risk-averse dealer ($\rho > 0$). With risk-aversion and an active futures market, some elements of dealer behavior may be viewed as hedging. The zero-inventory case highlights the speculative aspect of the dealer's behavior with an active futures market: when $q_t = 0$, his futures position reflects only the direction and magnitude of the risk premium adjusted for risk-aversion. However when $q_t > 0$ as in Equation (6), 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. Therefore the dealer's decisions reflect a combination of hedging and speculative motives. If the futures risk premium is sufficiently high, the dealer may choose to be long the physical commodity and the futures contract simultaneously.

2.3 Household

The household consumes the commodity and the numeraire. Specifically, the household enjoys utility over composite consumption with constant elasticity of substitution (CES) aggregator

$$c_t = \left[(1 - \gamma)c_{x,t}^\eta + \gamma c_{y,t}^\eta \right]^{1/\eta}, \quad (9)$$

with subscript x denoting the numeraire and y the commodity.

The household receives an endowment of one unit of the numeraire in each period. In addition he chooses position ω_t in the front futures contract, subject to a transaction cost with parameter $\tau \geq 0$. The transaction cost is paid at settlement, it is dissipative, and it is proportional to the absolute face value of the position. The household's endowment net of gains or losses from futures is

$$\hat{x}_t = 1 + \omega_{t-1}(s_t - f_{t-1}) - \tau f_{t-1} |\omega_{t-1}|. \quad (10)$$

Optimal consumption of the commodity $c_{y,t}$ and the numeraire $c_{x,t}$ are

$$c_{y,t} = c_y(\hat{x}_t, s_t) = \hat{x}_t s_t^{-1} \left[1 + \left(\frac{\gamma}{(1-\gamma)s_t^\eta} \right)^{\frac{1}{\eta-1}} \right]^{-1} \quad (11)$$

$$c_{x,t} = c_x(\hat{x}_t, s_t) = \hat{x}_t - s_t c_{y,t}. \quad (12)$$

Aggregated consumption is homogeneous in the net endowment,

$$c_t = \hat{x}_t \left[(1-\gamma) \left(\frac{c_{x,t}}{\hat{x}_t} \right)^\eta + \gamma \left(\frac{c_{y,t}}{\hat{x}_t} \right)^\eta \right]^{1/\eta} = \hat{x}_t \hat{c}_t, \quad (13)$$

where \hat{c}_t is a function of preference parameters and the spot price s_t .

The household has mean-variance preferences over aggregated consumption, choosing ω_t in each period to maximize utility:

$$\begin{aligned} \max_{\omega_t} E_t[\hat{x}_{t+1} \hat{c}_{t+1}] - \frac{\theta}{2} \text{Var}_t[\hat{x}_{t+1} \hat{c}_{t+1}] \\ \text{s.t. } \hat{x}_{t+1} = 1 + \omega_t(s_{t+1} - f_t) - \tau f_t |\omega_t|. \end{aligned} \quad (14)$$

The household shares its risk aversion parameter θ with the producer.

The household's optimal futures position satisfies

$$\begin{aligned} \omega_t > 0 \text{ and } \omega_t = \frac{E_t[(s_{t+1} - (1+\tau)f_t)\hat{c}_{t+1}] - \theta \text{Cov}_t[(s_{t+1} - (1+\tau)f_t)\hat{c}_{t+1}, \hat{c}_{t+1}]}{\theta \text{Var}_t[(s_{t+1} - (1+\tau)f_t)\hat{c}_{t+1}]}, \text{ or} \\ \omega_t < 0 \text{ and } \omega_t = \frac{E_t[(s_{t+1} - (1-\tau)f_t)\hat{c}_{t+1}] - \theta \text{Cov}_t[(s_{t+1} - (1-\tau)f_t)\hat{c}_{t+1}, \hat{c}_{t+1}]}{\theta \text{Var}_t[(s_{t+1} - (1-\tau)f_t)\hat{c}_{t+1}]}, \text{ or} \\ \omega_t = 0 \text{ otherwise.} \end{aligned} \quad (15)$$

To the price-taking household, the term \hat{c}_{t+1} represents exogenous consumption risk driven by uncertainty regarding the commodity spot price. Futures positions display the familiar combination of speculation and hedging motives. Households wish to increase expected consumption by earning the futures risk premium net of transaction costs, as seen by the expectation term in the numerators. But they also wish to hedge aggregated consumption risk by choosing futures ω_t such that the covariance of their net endowment \hat{x}_{t+1} with \hat{c}_{t+1} is low. This motive is captured by the covariance term in the numerator of the expressions in Equation (42). Because consumers are naturally “short” the commodity, the hedging motive favors a long position in futures, but given a sufficiently small or negative risk premium households will choose zero or negative futures.

3 Equilibrium

Prices are determined by market clearing. In each period, commodity goods and futures market clearing requires

$$y_t = c_{y,t} + q_t - q_{t-1}, \quad (16)$$

$$0 = \phi_t + \omega_t + \psi_t. \quad (17)$$

We can now define equilibrium.

Definition 1 (Equilibrium). *Equilibrium is a sequence of state-contingent prices and policies $\{s_t, f_t, q_t, \phi_t, \psi_t, \omega_t, c_{x,t}, c_{y,t}\}$ such that each agent’s policy solves his maximization problem, and commodity spot and futures markets clear $\forall y_t, t > 0$, for f_0, q_0, ω_0 , and y_0 given.*

Appendix A describes the numerical solution for equilibrium. I represent equilibrium in terms of a state vector $z_t = (q_{t-1}, y_{t-1}, y_t)$, equilibrium aggregate inventory function

$$q_t = Q(q_{t-1}, y_{t-1}, y_t), \quad (18)$$

futures price function

$$f_t = F(q_t, y_t), \quad (19)$$

and household futures position function

$$\omega_t = \Omega(q_t, y_t). \quad (20)$$

The state vector z_t is sufficient assuming that the initial futures price and position are consistent with the equilibrium functions, i.e., $f_0 = F(q_0, y_0)$ and $\omega_0 = \Omega(q_0, y_0)$.

Given the state vector and equilibrium functions above, household commodity consumption must clear the goods market per Equation (16). Therefore the spot price s_t satisfies

$$(1 + \omega_{t-1}(s_t - f_{t-1}) - \tau f_{t-1} |\omega_{t-1}|) \left(s_t + \left(\frac{\gamma}{(1 - \gamma)s_t} \right)^{\frac{1}{\eta-1}} \right)^{-1} = y_t + q_{t-1} - q_t. \quad (21)$$

Household numeraire consumption follows from Equation (12). The producer and dealer futures positions are given by Equation (3) and Equation (6), respectively.

Although equilibrium is not expressible in closed form, some intuition regarding qualitative model characteristics is gleaned from equilibrium conditions and some figures. The choice of parameter values used to generate the figures, listed in Table 1, is discussed in Section 4; the focus here is on general model characteristics that are insensitive to parameter values.

Inventory at the start of period t , q_{t-1} , is the key endogenous state variable. The choice of exiting inventory q_t reflects some fundamental shock, here productivity y_t . For the dynamic storage model to be of interest, there must be at least one high productivity state in which the dealer accumulates inventory, $q_t > q_{t-1}$, given low or zero entering inventory q_{t-1} . If productivity remains high, the dealer continues to accumulate inventory up to a limiting endogenous maximum, above which level storage costs exceed potential gains to accumulating further inventory. The basic motive for holding inventory is that there is at least one low productivity state in which the dealer can sell off part or all of his stocks at a profit. A succession of low productivity states leads to a “stockout,”

$q_t = 0$. Intermediate productivity states may involve either accumulation or sale of inventory depending on whether entering inventory q_{t-1} is low or high. The left panel of Figure 1 illustrates this basic mechanism given five states of productivity, with entering inventory q_{t-1} on the x-axis and exiting inventory q_t on the y-axis.

Deaton and Laroque [1992, 1996] find that storage can partially account for autocorrelation, heteroskedasticity, and skewness of commodity spot prices when such characteristics are absent or less pronounced in productivity. A period of high productivity, low spot price, and inventory accumulation lowers the spot price in the future, as inventory is sold off once productivity falls. This reduces the variance and increases the autocorrelation of the spot price. The spot price variance is conditionally higher when inventory is low or zero, as any drop in productivity must be fully absorbed as reduced consumption. Particularly if demand for the commodity is quite inelastic, the spot price spikes in periods of low productivity and low inventory. The relationship between inventory, productivity and the spot price is illustrated in the right panel of Figure 1.

This paper extends our understanding of inventory dynamics and price risk to a setting with an active futures market in which that risk is priced. In addition it analyzes the effects of changing participation in the futures market, including any changes to the underlying inventory and spot price processes. As a benchmark, I impose a parameter restriction such that only two agents, the producer and the dealer, participate in the futures market. Relative to this benchmark, I study the effects of household entry into the futures market, an event that I define as “financialization.” The following lemma and definition formalize this concept. All proofs are in Appendix B.

Lemma 1. *Assume equilibrium spot and futures prices are positive and bounded. There exists $\tau > 0$ such that households will always prefer not to trade futures: $\omega_t = 0, \forall z_t$.*

Definition 2 (Financialization). *Suppose the household transaction cost parameter $\tau = \bar{\tau}$ is sufficiently large that $\omega_t = 0, \forall z_t$. Financialization is an unanticipated reduction in costs to $\tau = \underline{\tau} < \bar{\tau}$ such that $|\omega_t| > 0$ for one or more states z_t with positive probability of occurrence.*

Without loss of generality let $\bar{\tau} \rightarrow \infty$, and call this case pre-financialization, whereas any $\tau \leq \bar{\tau}$ is post-financialization. While this definition is somewhat stark in assuming no household access to futures markets prior to financialization, it accords well with how financialization is often described: as the entry of a new type of investor into the commodity futures market, leading to a reduction in market segmentation. The primary means by which these new investors participate in practise is commodity index products, as emphasized by Masters [2008], for example. Such products implement relatively unsophisticated investment strategies designed to track a benchmark commodity index at low cost, such as by means of a long position in near-to-delivery futures contracts. They are categorized as passive investments, because they do not involve active selection from a complete menu of commodity-linked securities in an effort to exceed benchmark returns. Although minimalist, a reduction in τ captures the key idea of financial innovation that allows a new class of investor to more cheaply access the commodity market via a uniform product.

By contrast Basak and Pavlova [2016] model financialization as increased performance benchmarking to commodity indices, by introducing institutional investors who incorporate the commodity index directly into their utility function. By modeling futures on many commodities, they illustrate spillovers from indexed to non-indexed commodities. However markets are complete and frictionless both before and after financialization, and no investor follows a passive indexing strategy. Rather, financialization may be viewed as altering the objectives of sophisticated investors, by introducing a new passive benchmark. Another study more oriented towards sophisticated investors is Sockin and Xiong [2015], who study financial traders of commodity futures who introduce both useful information and noise into the futures market. While sophisticated or informed financial traders could presumably trade commodity futures without the aid of low-cost index products, they may nevertheless use such products to partially implement their strategies. Appendix C discusses the composition index investors, and argues that a large portion of such investors are either households, or intermediaries pursuing simple investment strategies on their behalf.

In the pre-financialization economy, the following proposition and corollary summarize futures market positions, how they relate to fundamental risk, and how that risk is priced.

Proposition 1. *Suppose the producer shorts the futures contract before financialization, $\phi_t < 0$. Then the dealer is long futures, the risk premium is positive, and the producer faces predominantly price risk: $\psi_t > 0$, $E_t[s_{t+1}] - f_t > 0$, and $\sigma_{sys,t} > 0$.*

Corollary 1. *Suppose the producer faces predominantly price risk before financialization, $\sigma_{sys,t} > 0$. Then the risk premium is always positive: $E_t[s_{t+1}] - f_t > 0$.*

In the canonical setting of Keynes [1930, p. 142], where producers hedge by shorting futures and the risk is born by speculators (here dealers) rather than consumers, Proposition 1 implies “normal backwardation” in the sense of a positive risk premium. Corollary 1 shows that there is a positive risk premium when producers face predominantly price risk rather than quantity risk. This setting aligns quite well with Keynes’ verbal description of the forward risk premium, which he argues will be positive whether the forward curve is in contango (upward sloping) or backwardation (downward sloping).³ A negative risk premium is possible pre-financialization, but it requires that producers face primarily quantity risk, and so choose a long futures position. Because the predominance of price versus quantity risk is crucial for hedging motives, Section 6.1 presents an extended model that introduces alternative shocks, with implications for who bears more price or quantity risk. Hirshleifer [1989] provides related theoretical analysis.

After financialization the connection between the producer’s hedging motive and the risk premium is loosened. From Equation (6), the futures risk premium satisfies

$$E_t[s_{t+1}] - f_t = \rho\sigma_{s,t}^2(\psi_t + q_t). \quad (22)$$

³“But the existence of a contango does not mean that a producer can hedge himself without paying the usual insurance against price changes. On the contrary, the additional element of uncertainty introduced by the existence of stocks and the additional supply of risk-bearing which they require mean that he must pay more than usual. In other words, the quoted forward price, though above the present spot price, must fall below the anticipated future spot price by at least the amount of the normal backwardation. . . .” [Keynes, 1930, p. 144].

With household participation in futures market, the dealer's futures position $\psi_t = -(\omega_t + \phi_t)$ may be negative on net even if producers take a short position. The equilibrium risk premium then turns negative when $\psi_t + q_t < 0$, i.e., when dealers carry negative total exposure to the commodity. This suggests that a negative risk premium is more likely in the post-financialization economy, if the household's hedging motive is strong. Presuming producers face predominantly price risk, they will hedge more (sell more futures) given a reduced risk premium. Therefore open interest, the total number of outstanding futures contracts $|\phi_t| + |\omega_t|$, will increase.

To study the effects of financialization on the futures curve, I price multi-period contracts under the assumption that only the dealer participates in these markets. This is similar to the approach taken by Routledge et al. [2000], where the dealer is the only agent in the model.⁴ The period t price of a futures contract for delivery in period $t + n$ is

$$f_{t,n} = E_t[f_{t+1,n-1}] - \rho(q_t + \psi_{t,1})Cov(s_{t+1}, f_{t+1,n-1}), \quad (23)$$

with $f_t = f_{t,1}$ the one-period contract, and $f_{t,0} = s_t$ the deliverable contract. I measure the slope of the futures curve as the difference between the 2-period and 1-period futures contracts, $f_{t,2} - f_{t,1}$. The futures curve is in backwardation when the slope is negative, whereas it is in contango when the slope is positive.⁵

Financialization will affect the spot market also. To provide some intuition, consider the case where $\eta \rightarrow 1$, such that the household's composite consumption is Cobb-Douglas. In this simplified case, the spot price in period t satisfies

$$s_t = \frac{\gamma \hat{x}_t}{c_{y,t}}. \quad (24)$$

⁴Section 6.2 relaxes this assumption by allowing the household to also trade the 2-period contract, and argues that the household will prefer to trade only the 1-period contract in any case.

⁵Backwardation is sometimes defined as the the expected spot price at time of expiration less the futures price. Defining backwardation based on the slope facilitates comparison with the data because the slope is observable, whereas the expected future spot price is not.

In the equation above, shocks to the numerator are sometimes described as demand shocks, whereas shocks to the denominator are supply shocks. An important distinction between the model in Section 2 and other studies of financialization is that both the numerator (demand) and the denominator (supply) are endogenously affected by financialization through the futures market. In Basak and Pavlova [2016], for example, both demand and supply are exogenous in the terminal period in which futures market participants consume, so the spot price is also exogenous.⁶

As a result of financialization, the variance of demand, the numerator in Equation (24), increases. Taken independently, this increases the variance of the spot price. But also as a result of financialization, the variance of supply, the denominator in Equation (24), decreases. Taken independently, this typically decreases the variance of the spot price.

Supply is affected through the dealer's storage decision, since $c_{y,t} = y_t + q_{t-1} - q_t$. Dealers typically increase inventory in response to financialization. With inventory in particular it is important to distinguish between immediate and long-term effects. An increase in current inventory will, all else equal, increase the current spot price, and decrease the expected future spot price. However once the increased level of inventory is established, continuing to carry that level of inventory neither adds nor subtracts from the current supply, but provides greater insurance against potential production shortfalls in the future. Therefore increased *average* inventory decreases supply volatility and spot price volatility.

Demand is affected by the household's futures position, since $\hat{x}_t = 1 + \omega_{t-1}(s_t - f_{t-1}) - \tau f_{t-1} |\omega_{t-1}|$.

⁶In an extension, Basak and Pavlova [2016] introduce hand-to-mouth consumers of the commodity in interim periods t and $t + 1$, prior to the terminal period, and allow for one-time storage across these two periods. Inventory traded in the interim periods is a redundant asset that does not affect the evolution of terminal spot or futures prices, so futures market participants are unaffected by the introduction of these markets. In the interim periods, the overall effect of financialization on the spot price level is ambiguous: the period t spot price increases, whereas the period $t + 1$ spot price decreases. The authors emphasize the period t increase, but note that an infinite horizon model with dynamic storage would be necessary to determine the unconditional impact of financialization on the spot price. Such analysis is one contribution of this paper.

Pre-financialization, $\omega_{t-1} = 0$, hence $\hat{x}_t = 1$ and demand is constant. Post-financialization, \hat{x}_t is stochastic and positively correlated with the spot price if households choose to hedge ($\omega_{t-1} > 0$). Increased demand volatility increases spot price volatility.

The above discussion captures key forces at work in the model, but at some risk of oversimplification. In addition to operating in opposite directions, the effects described above do not operate independently. All moments of the spot price may be impacted by financialization, including the mean spot price.

Since the model features several counterbalancing forces, a calibrated example is required to assess the direction and magnitude of financialization's effects.

4 Data and calibration

I calibrate the model to crude oil, mapping to the spot price of West Texas Intermediate (WTI) crude oil for Cushing, Oklahoma delivery, and the associated NYMEX futures contracts. I calibrate to a quarterly frequency. I match the 3-month futures contract to the front (one-period) model contract, the 6-month contract as the second (two-period) model contract, etc. The slope of the futures curve is the 6-month price less the 3-month price. Estimates of autocorrelation use quarter-on-quarter prices. I define excess returns as the log-difference in price for a quarterly holding period. Using the model notation, for the n 'th futures contract this is

$$\log(f_{t+1,n-1}) - \log(f_{t,n}), \quad (25)$$

which is approximately equal to the continuously compounded excess quarterly return on a fully collateralized contract. Henceforth I refer to the expected excess return on a given futures contract as the risk premium for that contract.⁷

⁷This is a unitless measure of the risk premium that facilitates comparison between the model and the data, as distinct from the expected payoff to the futures contract discussed earlier, which is expressed in units of the numeraire.

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 parameter choices. Additional details regarding data sources empirical target moments are in Appendix C. The production process y_t is the exogenous stochastic process that drives the baseline calibration of the model. Extensions in Section 6.1 consider alternatives such as stochastic demand. I calibrate the production process y_t to roughly match fractional variation in global production relative to trend, using a discrete approximation to an AR(1) process, with five states. 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 , used to model financialization. Summary statistics for data on prices and quantities are shown in the left column of Table 2. Moments for prices are split into 1990-2003 and 2004-2015 samples, for comparison with pre-and-post-financialization model moments, to be discussed later.

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 1990-2015, roughly 0.15%.⁸ This leaves five free parameters in addition to the transaction cost τ . I choose model parameters to approximately match the data during the pre-financialization period (1990-2003). For purposes of the calibration and to match the definition of financialization vis-a-vis the model, I assume that no households participated in futures markets during 1990-2003, setting $\tau = \infty$ to reflect no household trade in futures. Storage costs are $k = 0.001$ per unit per quarter, around 3% of the average unit price of oil, which affects the standard deviation and autocorrelation of consumption. Household goods aggregation parameters γ and η are used to match the average model value of oil consumption to oil/GDP, and the standard deviation of spot and futures prices. I use producer (θ) and dealer (ρ) risk aversion parameters to match producer hedging as a fraction of output, around 25% short position, and the futures risk premium, around 2% per quarter to the long position. In broad strokes, the

⁸Data 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/>)

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 household's risk aversion plays no role in the pre-financialization model; I assume the household is as risk averse as the producer.

For the post-financialization model, the objective is not to match post-financialization (2004-2015) data, but rather to assess what changes one would expect to observe in crude oil markets due to financialization alone. The 2004-2015 period includes events other than financialization of potentially greater importance for crude oil markets, including the 2008 financial crisis and the development of tight oil in the US. However it is necessary to calibrate the magnitude of financialization that is modeled. I consider three values of the transaction cost τ such that the average face value of household futures is around 20%.

5 Results

The following section first assesses the ability of the model to fit the pre-financialization data. Subsequently I reduce transaction cost τ while leaving the other parameters unchanged, and discuss the model's post-financialization behavior alongside 2004-2015 empirical observations.

5.1 Before financialization

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 ($\tau = \infty$). The model is able to match basic statistics of spot and futures prices. Spot price autocorrelation is 0.67, with futures correlation higher at around 0.72, each of which is comparable to the data. The standard deviation of the spot price is too high in the model, at 41% in the model to 23% in the data, but at 24% the

futures standard deviation is close to the data. The risk premium on futures is around 2.8% versus 1.9% empirically. In the model producers hedge around 28% of their production. The futures curve in the model is not in backwardation as frequently as in the data: 44% of the time, versus 70% in the data.

Regarding oil consumption, at 3.7% of the household's endowment its relative value is comparable to the data. Consumption dynamics depend heavily on the exogenous production process. However inventory changes cause consumption to become more autocorrelated than production, and also reduce the standard deviation of consumption relative to production. The conditional standard deviation of oil consumption is close to the data, but the unconditional standard deviation is low relative to the data.

The model also approximates some futures moments over the term structure. The left column of Figure 2 shows normalized mean futures prices over the term structure in the top row, illustrating a similar downward slope in the model and the data. The middle-left panel shows standard deviation declining with maturity, more steeply in the model than in the data. 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 to have reasonable production and consumption dynamics. In addition, the model implies a downward sloping term structure for the risk premium, shown in the bottom left panel of Figure 2, with a slope similar to that in the data. Overall this seems a good performance for a simple model.

5.2 The effects of financialization

I model financialization by reducing the household's transaction cost parameter, τ . Summary statistics for decreasing τ are given from left to right alongside the baseline calibration in Table 2. As a percentage of consumption, the mean face value of household futures ranges from 16% of con-

sumption for $\tau = 0.15$ to 25% for $\tau = 0.08$. Analysis focuses on the middle case, with $\tau = 0.1$ and household futures at 22% of consumption of average.

In broad strokes, the model suggests that financialization increases the autocorrelation and reduces the standard deviation of consumption, reflecting increased inventory. Yet spot prices become less autocorrelated, and the standard deviations of spot and futures prices increase. These apparently contradictory results for consumption versus prices reflect the offsetting effects of inventory and household hedging. The futures risk premium also decreases, and backwardation becomes about half as frequent. These two changes are also observed in the split sample moments from the data. However the mean spot price in 2009 USD is \$28 from 1990-2003 versus \$77 from 2004-2015, whereas in the model the spot price is essentially unchanged by financialization. Preliminary conclusions are that financialization could substantially alter the risk premium and the shape of the futures curve, but should have little impact on the mean level of the spot price.

5.2.1 Futures positions

One feature of financialization in the data is an overall increase in open interest. Relative to annual US crude oil consumption, open interest rises from an average of 24% before 2004 to 71% after 2004, as reported in Table 2. In the model, open interest rises somewhat less than in the data, from 28% before financialization to 61% after financialization, with $\tau = 0.1$.

To understand how financialization effects trade in futures in the model, I present some results conditional on the level of inventory. Figure 3 shows positions and open interest in the 3-month futures contract in the pre-financialization model ($\tau = \infty$, solid line) and the post-financialization model ($\tau = 0.1$, dashed line). Each agent's futures exposure varies with the level of inventory (q_t , x-axis). The productivity state (y_t) also affects futures positions, but the shape of the curves is similar for each value of y_t ; the plots show expectations using the stationary distribution of y_t . Households (top right) take a large positive position after financialization, hedging their consumption risk.

However producers (top left) also hedge more after financialization, because the risk premium to the long side of the futures contract is reduced. Because producers and households have offsetting hedging motives, the dealer's net futures position is reduced when households trade futures.

Open interest, the total number of outstanding futures contracts, increases after financialization, especially when inventory is low. Households and producers both hedge most when inventory is low, because that is when spot prices are most volatile. In fact households will not hedge at all when inventory is sufficiently high, because transaction costs outweigh the benefits. A byproduct of this is that the net futures position of dealers is non-monotonically related to inventory after financialization, i.e., even though dealers could offset decreased futures exposure by increasing inventory, it is not generally optimal for them to do so. This is in contrast to Vercammen and Doroudian [2014], in which inventory is assumed to be perfectly hedged by shorting futures.

5.2.2 Inventory

The dealer has access to two investment opportunities, stored oil and a one-period futures contract on oil. Although inventory cannot be negative, the dealer will usually choose to hold positive inventory, such that futures and inventory are perfect substitutes at the margin. Since financialization alters the dealer's futures position, one would expect financialization to alter his inventory also.

For any state z_t , the dealer always chooses weakly more inventory q_t after financialization than before. As shown in a histogram for inventory in Figure 4 (top left), higher inventory states are more probable after financialization, and there are fewer stockouts. Unconditionally, average inventory in the model rises from 5.2 days supply before financialization to 6.9 days supply after financialization, with $\tau = 0.1$. In the data, average days of inventory rises from 2.9 days before 2004 to 3.7 days after 2004, a percentage increase comparable to the model. While anecdotal evidence also supports an increase in speculative inventory after financialization, the figures reported in the table assume that crude oil stocks in the Midwestern PADD 2 region, containing the

Cushing, Oklahoma delivery point, are the relevant proxy for speculative inventory in the model. Appendix C.1 discusses alternatives, which do not robustly support increased inventory after financialization.

To understand the mechanism by which inventory increases in the model, it is helpful to think of the dealer's interaction with households and producers after financialization as occurring sequentially. 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 equivalence of futures and inventory requires that he does both.

5.2.3 The spot price

The effect of financialization on the spot price can be thought of as the net of changes to dealer behavior and household behavior. The equilibrium effect of increased inventory is to smooth spot prices. Given sufficiently inelastic demand for crude oil, increased inventory could also lower the mean spot price. However the smoothing effect due to increased inventory is undone by the effect of hedging on household income.

The bottom panel of Figure 4 shows spot prices versus inventory and productivity, before (solid lines) and after financialization (dashed lines). When oil production is good (bottom curves), 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). However when productivity is low (top curves), 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. The net of the inventory and hedging effects is illustrated in the top right panel of Figure 4, 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.

To illustrate the effect of financialization on the joint dynamics of inventory, the spot price, and open interest, Figure 5 shows how the pre-financialization and post-financialization models respond to the same series of oil production realizations. Inventory is initially set to zero, and oil production is low. To illustrate the process of inventory accumulation and sale, there are three consecutive periods (1-3) of high oil production, followed by four periods (4-7) of low oil production. Inventory accumulates more rapidly in the post-financialization model, and takes one period longer to reach a stockout than in the pre-financialization model. This additional inventory provides a larger buffer against low production, leading to a lower post-financialization spot price in simulated periods 5 and 6. However by period 7 inventory is exhausted even in the post-financialization model. The spot price in period 7 is highest in the post-financialization model, because the household receives a large positive payoff from his futures position, which acts like a positive demand shock. Finally, open interest is more responsive to the spot price post-financialization, as is discussed in more detail in Section 5.2.5.

The increased spot price volatility after financialization is driven by payments from the dealer to the household in low productivity states. If the payments were significantly positive on average (i.e., across all states), this would imply a de facto increase in household income, suggesting the mechanism would not extend to a general equilibrium setting. Table 3 reports summary statistics for transfers to the household from gross gains or losses in the futures market,

$$\omega_{t-1}(s_t - f_{t-1}), \tag{26}$$

and net gains or losses,

$$\omega_{t-1}(s_t - f_{t-1}) - \tau f_{t-1} |\omega_{t-1}|. \quad (27)$$

On average the transfers are approximately zero. As a percentage of the household's quarterly endowment, mean transfers are 0.03% gross of transaction costs, and -0.07% net of transaction costs. Therefore the hedging effect is actually accompanied by a small net reduction in average household income, stemming from dissipative transaction costs. Although transfers can be large in some states, ranging from -3.75% to 6.29% net of fees, the standard deviation of transfers, 0.963% net of fees, is not implausibly large given the volatility of crude oil spot prices in the data, which is over 20% quarterly.

5.2.4 The term structure of futures moments

The baseline calibration does a reasonably good job of matching the unconditional term structure of futures prices (mean and standard deviation) and excess returns, shown in the left column of Figure 2. In addition, summary statistics in Table 2 suggest that financialization can partially explain changes in moments of the heavily traded 3-month contract after financialization, such as decreased backwardation, reduced risk premium, and some increase in standard deviation of futures prices. A natural question is whether these results extend to contracts with more than three months to delivery.

The right column of Figure 2 shows unconditional mean futures prices, standard deviation of futures prices, and expected quarterly excess returns for futures contracts with increasing time to delivery. The solid line gives sample averages from the data (2004-2015), the dashed line shows expectations from the model ($\tau = 0.1$). Post-financialization, the frequency of backwardation decreases in the model (from 44% to 21%), but when backwardation occurs the slope is steeply negative, such that the mean futures curve is still downward sloping. The model also implies that standard deviations shift upward, whereas risk premiums shift downward but increase in slope.

In contrast the data shows a hump-shaped mean futures curve, and the standard deviation shifts upward but also takes on a flat term structure. The futures risk premium shifts downward and is upward sloping, as opposed to the humped model-implied term structure. In general the model-implied effects of financialization are more consistent with observed changes for the near futures contracts than for long-dated contracts. Ready [2014] notes that metals and agricultural futures also subject to financialization did not exhibit the same changes to the mean futures curve or the term structure of volatility seen in crude oil futures data, and suggests that a change in oil supply dynamics could account for the difference.

5.2.5 Correlation between futures open interest and spot prices

Figure 6 plots the path of open interest (number of outstanding futures contracts, solid line) and real WTI spot prices (dashed line) from 1990 through 2015. After 2004 the spot price rises rapidly, then experiences a rapid drop and rebound following the financial crisis in 2008, and another drop in late 2014. Although open interest appears to move independently of the spot price prior to 2004, it rises in tandem with the spot price after 2004, and tracks the volatile price movements after the financial crisis.

My model is unique in that joint open interest and spot price dynamics are endogenous and non-trivial before and after financialization. Other models feature only one investor before financialization, so open interest is zero or undefined, or they model index investors as an exogenous shock after financialization, such that open interest after financialization does not reflect endogenous portfolio choice by market participants. As previously discussed, financialization ought not to cause a sustained increase in the spot price, but in the model it does generate increased comovement of open interest and the spot price. Figure 7 plots open interest in crude oil futures contracts (y-axis) against the spot price of crude oil (x-axis), in the data (top row) and the model (bottom row). The left column shows pre-financialization data (1990-2003) and model ($\tau = \infty$); the right

column shows post-financialization data (2004-2015) and model ($\tau = 0.1$). Prior to financialization the relationship between levels of open interest and prices was flat and not statistically significant (top left, $R^2 = 0.012$, $p = 0.164$). After financialization there is a positive and statistically significant relationship (top right, $R^2 = 0.364$, $p = 0.000$). The post-financialization model (bottom right) also shows a stronger positive relationship between open interest and the spot price relative to the pre-financialization model (bottom left), where the relationship is mostly flat.

In the model, the increased comovement of open interest and the spot price results from the increased hedging that takes place when inventory decreases, which coincides with higher conditional spot price volatility and level. As illustrated in Figure 3, households and producers hedge extensively in these high volatility states after financialization. This also leads high open interest to coincide with high spot prices, despite the fact that financialization does not increase the mean spot price.

6 Extensions

Two important modeling assumptions simplify the analysis in the preceding sections: the model is driven by a single exogenous shock to commodity sector productivity, and only the 1-period (3-month) futures contract is actively traded. Sections 6.1 and 6.2 relax each of these assumptions, respectively, and argue that such changes to the model leave the main results intact.

Section 6.3 contrasts household entry into futures markets with an alternative form of financialization, where the new entrants are “portfolio investors” who do not consume the commodity. Without consumer hedging effects, financialization decreases spot price volatility due to increased inventory smoothing.

Finally Section 6.4 studies producer equity issuance. Illiquid endowments are intrinsic to agent identity in the baseline model. If the producer is able to sell shares of his endowment to the

household without restriction prior to financialization, then financialization has no effect, as the household already owns a claim to commodity output that satiates its hedging demand. However the producer ceases to be a producer as such, since he retains no claim to his endowment. Motivated by the large fraction of oil production that is owned by national oil companies and is not publicly traded, I constrain the producer to sell only a fraction of his shares. In this scenario the effects of financialization are similar to the baseline model. Simultaneous trading of equity and futures reveals that financialization increases the correlation of equity and futures returns, similar to results in Basak and Pavlova [2016].

6.1 Alternative shocks

Appendix C rationalizes global oil production as the random process driving the model. Of course, additional sources of randomness affect oil prices. For example, 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

$$\hat{x}_t = x_t + \omega_{t-1}(s_t - f_{t-1}) - \tau f_{t-1} |\omega_{t-1}|. \quad (28)$$

The model in Section 2 is a special case of the above with $x_t = 1$ for all t .

Given the calibration of the household as a representative U.S. consumer, a third source of risk might be foreign demand. A simple way to model this without explicitly introducing a fourth agent is to add a third Markov process \hat{y}_t , such that goods market clearing now requires

$$\hat{y}_t y_t = c_{y,t} + q_t - q_{t-1}. \quad (29)$$

Consider non-U.S. countries that are collectively self sufficient on average ($E[\hat{y}_t] = 1$), but may have net imports or exports from time to time, such that \hat{y}_t represents foreign net demand. Fluctu-

ations in \hat{y}_t will reduce or increase the amount of oil available for U.S. consumption. The model in the main text is a special case with $\hat{y}_t = 1$.

The three shocks have intuitive meaning as production (y_t), domestic demand (x_t) and foreign demand (\hat{y}_t) shocks. They also facilitate a decomposition of quantity versus price risk in the model. In the default calibration driven by y_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. The household, meanwhile, faces oil price risk, but the value of his endowment is constant by assumption. Alternatively, if we consider a model driven only by shocks to x_t , the direct benefit to the household of a high endowment outcome is offset by a high spot price; the household faces offsetting quantity and price effects. The producer, however, is subject only to price risk. Finally we might consider a model driven only by shocks to \hat{y}_t . The foreign demand shock affects neither producer nor household 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 y_t in the main calibration to each of the shocks in turn, while the other processes are held constant at unity. This keeps the amount of intrinsic risk in the oil market roughly the same, but shifts the exposure to quantity and price risk. Table 4 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. The model results are quite similar across shocks, but differences in risk exposure manifest in the futures market. Moving from column y_t to column \hat{y}_t , we see that producers short more futures when they face only price effects (with \hat{y}_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 (y_t) or foreign demand (\hat{y}_t) shocks. However the higher risk premium to the long side of the futures contract under stochastic \hat{y}_t leads households to go long more futures after financialization. Therefore financialization generates a

larger drop in the risk premium and frequency of backwardation than in the baseline model.

When the model is driven by domestic demand shocks (x_t), the producer again faces only price risk, so he shorts futures heavily. Although the household can earn a risk premium by going long the futures contract, he would now hedge by shorting the futures contract, because the spot price is low when the household's endowment is low.⁹ Whether the household takes a long or a short position after financialization depends on the calibration, and especially on the household's risk aversion. Table 4 shows that the household shorts in this calibration, causing the risk premium to rise rather than fall after financialization. Whether producers and household consumers have offsetting hedging motives depends on the source of risk that drives spot prices.

6.2 Active trade in multi-period futures contracts

This section relaxes the assumption that only the one-period (three-month) futures contract is actively traded. The previous focus on the three-month contract is justified 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 three-months or less, and a further 31% of open interest was in contracts of three 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 three-months or less, and 33 % for three to 12-months. Therefore I investigate whether the results change if households may trade both one and two-period (three and six-month) futures. I find that households choose not to trade the two-period contract if transaction costs for that contract are at least as large as for the one-period contract.

⁹Financial assets such as equity derivatives or bonds might be more appropriate for hedging numeraire risk than oil futures, but these assets are not modeled. The case $x_t = 1$ might be thought of as fully hedging numeraire risk using assets that are outside of the model.

Futures positions of the dealer and household are $\psi_{t,n}$ and $\omega_{t,n}$, respectively, with $n \in \{1, 2\}$. The producer trades only the one-period contract, with position $\phi_{t,1}$.¹⁰ Period $t + 1$ profits for the dealer and producer, and net income for the household, are

$$\begin{aligned} p_{t+1}^d &= q_t(s_{t+1} - (1+r)(s_t + k)) + \psi_{t,1}(s_{t+1} - f_{t,1}) + \psi_{t,2}(f_{t+1,1} - f_{t,2}), \\ p_{t+1}^p &= s_{t+1}a_{t+1} + \phi_{t,1}(s_{t+1} - f_{t,1}), \\ \hat{x}_{t+1} &= 1 + \omega_{t,1}(s_{t+1} - f_{t,1}) - \tau f_{t,1}|\omega_{t,1}| + \omega_{t,2}(f_{t+1,1} - f_{t,2}) - \tau f_{t,2}|\omega_{t,2}|. \end{aligned} \quad (30)$$

Effectively the household is able to buy two oil futures “funds,” one with a rolling position in the one-period contract, the other with a rolling position in the two-period contract, with quarterly rolls. I assume the two funds have identical cost parameter τ . Agents retain the objectives described in Section 2; the only modification is that the household and the dealer now have access to an additional tradable asset. However, it turns out that they will choose not to trade the new asset.

Earlier we saw the familiar result that when $q_t > 0$, $s_{t+1} - (1+r)(s_t + k) = s_{t+1} - f_{t,1}$: buying and storing inventory is equivalent to buying futures. With two actively traded futures, a similar result holds for the two-period contract. Suppose that in period t we are in a high-inventory state such that $q_{t+1} > 0$ for all y_{t+1} , i.e., if there is sufficient inventory that it will never be liquidated entirely in period $t + 1$. Under these circumstances the front two futures contracts can be replicated via storage,

$$\begin{aligned} f_{t,2} &= (1+r)(f_{t,1} + k) = (1+r)((1+r)(s_t + k) + k), \\ f_{t+1,1} &= (1+r)(s_{t+1} + k). \end{aligned} \quad (31)$$

In high-inventory states the front two futures contracts are equivalent to the dealer. It turns out the two contracts are not equivalent to households, because of transaction costs.

Proposition 2. *Suppose the economy has high-inventory, such that $q_t > 0$ and $q_{t+1} > 0$ for all y_{t+1} , and assume $\tau > 0$. Then the household prefers not to trade the two-period futures contract: $\omega_{t,2} = 0$.*

¹⁰In certain states the one and two-period contracts are redundant assets. In such states open interest is undefined if the producer is also able to trade the two-period contract without transaction costs.

Of course the household might choose to trade the two-period contract in low-inventory states. In general this possibility is difficult to eliminate. However for the main calibration described in Section 4 with transaction cost $\tau = 0.1$ for both contracts, I verify numerically that the household will choose not to trade the two-period contract even for low-inventory states.

6.3 Portfolio Investors

Previous analysis assumes that new participants in the futures markets are commodity consuming households. Another possibility is that new participants are “portfolio investors,” who buy futures to earn a risk premium, or to hedge risk in their existing portfolios. The important distinction is that portfolio investors are assumed not to consume the underlying commodity, which has implications for financialization’s effect on the spot market particularly.

Suppose a representative portfolio investor has exogenous income (net returns) v_t from the rest of his portfolio. He chooses futures ζ_t to maximize mean-variance utility by adding commodity futures to his portfolio, subject to transaction cost τ_ζ , and the same risk-aversion θ as the producer. Defining investment income inclusive futures \hat{z}_t , the portfolio investor solves

$$\begin{aligned} \max_{\zeta_t} E_t[\hat{z}_{t+1}] - \frac{\theta}{2} Var_t[\hat{z}_{t+1}], \\ \text{s.t. } \hat{z}_{t+1} = v_{t+1} + \zeta_t(s_{t+1} - f_t) - \tau_\zeta f_t |\zeta_t|. \end{aligned} \quad (32)$$

The solution has three cases, depending on the sign of ζ_t :

$$\begin{aligned} \zeta_t > 0 \text{ and } \zeta_t = \frac{E_t[(s_{t+1} - (1 + \tau_\zeta)f_t)] - \theta Cov_t[s_{t+1}, v_{t+1}]}{\theta Var_t[s_{t+1}]}, \text{ or} \\ \zeta_t < 0 \text{ and } \zeta_t = \frac{E_t[(s_{t+1} - (1 - \tau_\zeta)f_t)] - \theta Cov_t[s_{t+1}, v_{t+1}]}{\theta Var_t[s_{t+1}]}, \text{ or} \\ \zeta_t = 0 \text{ otherwise.} \end{aligned} \quad (33)$$

Equilibrium is altered to include the portfolio investor’s position in the futures market clearing

condition,

$$0 = \phi_t + \omega_t + \psi_t + \zeta_t. \quad (34)$$

A preliminary step is to characterize the portfolio investor's non-futures income, v_t . I assume the standard deviation of the portfolio investor's non-futures income is 15% annualized ($\sigma_v = 0.075$), comparable to returns on the S&P 500 index. Büyüksahin et al. [2010] estimate time-varying correlation between commodity and equity index returns in the range of -0.5 to 0.5, with an average near zero. For simplicity I assume that non-futures income has constant correlation with spot prices $\rho_{s,v}$, such that $Cov_t[s_{t+1}, v_{t+1}] = \rho_{s,v}\sigma_v\sigma_{s,t}$, but solve the model for $\rho_{s,v} \in \{-0.2, 0, 0.2\}$. Remaining parameter values follow Table 1.

Similar to earlier analysis with households, I assume that portfolio investors initially face such high transaction costs that they do not participate in the futures market, i.e., $\tau_\zeta = \infty$ and $\zeta_t = 0$. Financialization occurs when τ_ζ unexpectedly falls to a level where portfolio investors choose to participate. In this alternative form of financialization, households are assumed never to trade futures, i.e., $\tau = \infty$ both before and after financialization, so $\omega_t = 0, \forall t$.

By comparing the entry of portfolio investors to the entry of households, the effects of financialization on the spot market can be decomposed. Dealers may offset futures sold to portfolio investors with increased inventory, but realized futures returns to portfolio investors will not directly effect consumer demand in the spot market. So entry by portfolio investors isolates the inventory smoothing effect. Appropriate comparison to equilibrium with participation by households shows the incremental consumer hedging effect.

The effects of financialization with portfolio investors are summarized in the right set of columns in Table 5. Assume initially that futures payoffs are uncorrelated with portfolio investor income from other sources, $\rho_{s,v} = 0$. When the portfolio investor's transaction cost is reduced to $\tau_\zeta = 0.5\%$, his average futures holdings call for delivery of 20.9% of aggregate oil consumption. This is comparable to the household's mean futures position of 21.7% with a much higher transaction cost

$\tau = 10\%$, reflecting a strong hedging motive for households that is absent for portfolio investors with $\rho_{s,v} = 0$. As in Section 5, the equilibrium impact of portfolio investor participation may be evaluated relative to the calibration with only dealer and producer participation in the futures market. In the futures market, entry by the portfolio investor reduces the futures risk premium and the frequency of backwardation. These effects are smaller in magnitude than with household participation, but identical in sign. In the physical market, the mean spot price is essentially unchanged by portfolio investor futures trade. However mean inventory increases, from 5.2 days supply to 6.3 days supply after entry of the portfolio investor. In a simple storage model, higher average inventory implies less volatile and more autocorrelated spot prices: exactly the result seen in Table 5, where the standard deviation of spot prices declines by 2% and autocorrelation increases by about 4%. In contrast, when households enter the futures market, inventory rises to 6.9 days supply on average, but spot prices become *more* volatile, by about 7%, and *less* autocorrelated, by about 9%. This reflects the previously discussed household hedging effect, which overcomes the effects of increased inventory. Since portfolio investors do not consume the commodity, no equivalent effect is present when they enter the futures market, and the inventory smoothing effect prevails.

Since households will trade futures heavily even with $\tau = 10\%$, due to the hedging motive, Table 5 also provides summary statistics for $\tau_\zeta = 10\%$ and a range of values for $\rho_{s,v} \in \{-0.2, 0, 0.2\}$. For $\rho_{s,v} = 0$ the portfolio investor's mean futures position is only 2.01%, but this increases to 32.4% for $\rho_{s,v} = -0.2$, as a long position hedges portfolio risk. For positive $\rho_{s,v} = 0.2$, mean futures holdings are -1.61%, as a short position hedges portfolio risk. The portfolio investor's decision only affects the spot price indirectly, by influencing the dealer's inventory decision, so the economic intuition is unchanged regardless of what motivates the portfolio investor's decision: positive portfolio investor futures increase inventory, decrease spot volatility and increase spot autocorrelation, whereas negative portfolio investor futures do the reverse. However this analysis assumes correlation $\rho_{s,v}$ is constant. Given the sensitivity of portfolio investor futures to the choice of correlation, it is possible that realistic time-varying cross-market correlation could indirectly

drive swings in inventory large enough to increase spot market volatility, even if inventory does not decrease on average. Related analysis is presented in Basak and Pavlova [2016], who formalize cross-market relationships but abstract from dynamic inventory management. A comprehensive analysis of dynamic inventory and trade across markets is a topic for future research.

6.4 Producer equity

In Basak and Pavlova [2016], a stock market is modeled as a claim to the output of all commodity sectors and a numeraire sector, which pays off in terminal period T . Stock in each commodity sector is implicitly defined as a claim to the terminal output of each sector, and a futures contract maturing in the terminal period trades for each commodity.¹¹ Along similar lines, I allow the commodity producer to sell claims to his next-period endowment, which trades alongside the futures contract. I refer to these claims as equity shares, since the term stock often refers to inventory in connection with commodities. In the baseline model, producers are endowed with 100% of equity at the beginning of each period, and may not sell it.

Suppose that the producer may sell shares in his firm to the household at price e_t . The trade settles in period $t + 1$ at the agreed upon price, such that the net payoff is $s_{t+1}y_{t+1} - e_t$.¹² Let π_t^p be the fractional equity share retained by the producer. The producer solves

$$\begin{aligned} \max_{\phi_t, \pi_t^p} E_t[p_{t+1}^p] - \frac{\theta}{2} \text{Var}_t[p_{t+1}^p], \\ \text{s.t. } p_{t+1}^p = \pi_t^p (s_{t+1}y_{t+1} - e_t) + e_t + \phi_t (s_{t+1} - f_t). \end{aligned} \tag{35}$$

In the baseline model, $\pi_t^p = 1$, and $\phi_t = \frac{E_t[s_{t+1}] - f_t - \theta \sigma_{\text{sys},t}}{\theta \sigma_{s,t}^2}$. I first solve the extended model when

¹¹In Basak and Pavlova [2016], individual sector stocks would be redundant assets, hence they are not individually traded. Futures contracts for each of K commodities complete the market.

¹²An important modeling assumption is that agents do not make intertemporal consumption-saving decisions that would greatly complicate determination of equilibrium. The treatment of settlement as occurring in period $t + 1$ at a price agreed in period t maintains this assumption from the baseline model.

the producer may choose any value π_t^p . Later I constrain the producer to retain a fraction of equity $\pi_t^p \geq \bar{\pi}^p$.

Optimal producer equity shares are

$$\pi_t^p = \frac{E_t[s_{t+1}y_{t+1}] - e_t - \theta\phi_t\sigma_{sys,t}}{\theta\sigma_{sy,t}^2}, \quad (36)$$

and optimal producer futures holdings are

$$\phi_t = \frac{E_t[s_{t+1}] - f_t - \theta\pi_t^p\sigma_{sys,t}}{\theta\sigma_{s,t}^2}. \quad (37)$$

With unrestricted sale of shares, producer equity and futures holdings reduce to

$$\pi_t^p = \frac{E_t[s_{t+1}y_{t+1}] - e_t - \frac{\sigma_{sys,t}}{\sigma_{s,t}^2}(E_t[s_{t+1}] - f_t)}{\theta\sigma_{sy,t}^2 \left(1 - \frac{\sigma_{sys,t}^2}{\sigma_{sy,t}^2\sigma_{s,t}^2}\right)} \quad \text{and} \quad (38)$$

$$\phi_t = \frac{E_t[s_{t+1}] - f_t - \frac{\sigma_{sys,t}}{\sigma_{sy,t}^2}(E_t[s_{t+1}y_{t+1}] - e_t)}{\theta\sigma_{s,t}^2 \left(1 - \frac{\sigma_{sys,t}^2}{\sigma_{sy,t}^2\sigma_{s,t}^2}\right)}, \quad (39)$$

respectively, where $\frac{\sigma_{sys,t}}{\sigma_{sy,t}\sigma_{s,t}}$ is the conditional correlation of returns on equity and futures.

Incorporating trade in producer equity, the household solves

$$\begin{aligned} \max_{\omega_t, \pi_t^h} E_t[\hat{x}_{t+1}\hat{c}_{t+1}] - \frac{\theta}{2} Var_t[\hat{x}_{t+1}\hat{c}_{t+1}] \\ \text{s.t. } \hat{x}_{t+1} = 1 + \omega_t(s_{t+1} - f_t) - \tau f_t|\omega_t| + \pi_t^h(s_{t+1}y_{t+1} - e_t), \end{aligned} \quad (40)$$

in which the household's net endowment \hat{x}_{t+1} now includes gains or losses from futures and equity, and π_t^h is the household's equity share. The definition of \hat{c}_{t+1} , which is optimal aggregated consumption scaled by net endowment, is unchanged from the baseline model.

Optimal household equity shares are

$$\pi_t^h = \frac{E_t[(s_{t+1}y_{t+1} - e_t)\hat{c}_{t+1}] - \theta Cov_t[(s_{t+1}y_{t+1} - e_t)\hat{c}_{t+1}, (1 + \omega_t(s_{t+1} - f_t) - \tau f_t|\omega_t|)\hat{c}_{t+1}]}{\theta Var_t[(s_{t+1}y_{t+1} - e_t)\hat{c}_{t+1}]}, \quad (41)$$

and the household's optimal futures position is

$$\begin{aligned} \omega_t > 0 \text{ and } \omega_t &= \frac{E_t[(s_{t+1} - (1 + \tau)f_t)\hat{c}_{t+1}] - \theta Cov_t[(s_{t+1} - (1 + \tau)f_t)\hat{c}_{t+1}, (1 + \pi_t^h(s_{t+1}y_{t+1} - e_t))\hat{c}_{t+1}]}{\theta Var_t[(s_{t+1} - (1 + \tau)f_t)\hat{c}_{t+1}]}, \text{ or} \\ \omega_t < 0 \text{ and } \omega_t &= \frac{E_t[(s_{t+1} - (1 - \tau)f_t)\hat{c}_{t+1}] - \theta Cov_t[(s_{t+1} - (1 - \tau)f_t)\hat{c}_{t+1}, (1 + \pi_t^h(s_{t+1}y_{t+1} - e_t))\hat{c}_{t+1}]}{\theta Var_t[(s_{t+1} - (1 - \tau)f_t)\hat{c}_{t+1}]}, \text{ or} \\ \omega_t &= 0 \text{ otherwise.} \end{aligned} \tag{42}$$

The definition of equilibrium follows Section 3, with the additional requirements that the producer and household each optimize over equity shares, and e_t clears the equity market in each period:

$$\pi_t^p + \pi_t^h = 1. \tag{43}$$

The numerical model solution adopts the usual parameter values, in Table 1. Table 6 reports summary statistics for the baseline model without ($\tau = \infty$) and with ($\tau = 0.1$) financialization in the two left columns, and for the extended model with equity issuance in the remaining columns.

The assumption of an illiquid producer endowment is important to the effects of financialization in the baseline model, and to the Keynesian conception of hedging demand in commodity futures markets generally. If the producer is able to sell equity to the household before financialization, then he will sell all of his equity in the model, and will even sell some equity short on average, around 30%.¹³ In an important sense, the producer loses his identity, since he is no longer positively exposed to the commodity. As a result of his short equity position, the producer now hedges with a long position in the futures market, and the equilibrium futures risk premium turns negative even without financialization.

Because producer equity already allows the household to hedge commodity risk without paying

¹³Household demand for equity drives the equity risk premium negative, so the producer earns a risk premium on short sales. Parameters do not attempt to match the equity risk premium to data.

a transaction cost, a reduced impact from financialization is to be expected. But the difference is stark: financialization has *no* impact with transaction cost parameter $\tau = 0.1$, for example, as the household optimally chooses a large equity position and zero futures, $\omega_t = 0, \forall t$. In fact $\omega_t = 0$ satisfies the household's optimality conditions to a close numerical approximation even when $\tau = 0$. While difficult to prove analytically for general parameters, the optimality of $\omega_t = 0$ with unrestricted equity trading does not appear sensitive to parameter values, provided that commodity output y_t is the only source of risk in the model.¹⁴

While this example shows that financialization should be a non-event if households already have good access to alternative hedging instruments, perhaps it is not so realistic that producers issue equity without restriction. In reality much of the world's oil production is owned by national oil companies (NOCs), such that only a fraction of oil producer shares trade publicly. Tordo [2011] estimate that NOCs control 75% of world production and 90% of reserves, whereas Victor et al. [2011] estimate that NOCs own 61% of production and 73% of reserves. While some NOCs such as Equinor (formerly Statoil) and Pemex have publicly listed shares, their governments, Norway and Mexico respectively, retain a majority ownership stake. Other NOCs, such as Saudi Aramco, do not have publicly listed shares at all. The preceding analysis suggests NOCs ought to sell more shares to oil consumers; but perhaps unmodeled political and agency costs may prevent them from doing so.

Accordingly I constrain producers to retain equity shares $\pi_t^p \geq \bar{\pi}^p$. In the numerical example I set $\bar{\pi}^p = 0.75$, so producers can only sell a 25% stake. The constraint always binds the producer, so the equilibrium equity price e_t clears the household market, satisfying Equation (41). Table 6 reports summary statistics for the constrained model in the last four columns, beginning without financialization ($\tau = \infty$), and then for declining futures transaction cost τ . The constrained model remains similar to the baseline calibration. Unlike the unconstrained model, the producer shorts futures to

¹⁴It is also probable that the household views futures and equity as redundant assets in equilibrium when neither asset is subject to a transaction tax. In this case, equilibrium may not be unique.

hedge, and the futures risk premium is positive prior to financialization.¹⁵ Increasing financialization (decreasing τ) leads to increasing household futures and open interest, and decreasing futures risk premium. The main difference from the baseline analysis is that the household buys fewer futures on average for the same value of τ , a natural consequence of partial hedging via equity.

With equity trading alongside futures in the model, we can examine financialization's effect on the correlation of futures and equity returns. In theoretical work, Basak and Pavlova [2016] show that financialization increases equity-commodity correlation, and cite empirical evidence of increased correlation in Tang and Xiong [2012]. In Table 6 with constrained equity sales, financialization also causes futures returns become more correlated with equity returns, increasing from 89.8% without financialization to 98.4% when $\tau = 0.1$. The similar results suggest that increased cross-market correlation is a robust feature of financialization, whether due to decreased market segmentation, as in this paper, or due to increased performance benchmarking to commodities, as in Basak and Pavlova [2016].

7 Conclusion

I construct a model of storable commodities with producers, dealers, households, and an active futures market. I use the model to study how the financialization of commodities, occurring in isolation, impacts spot and futures prices. When calibrated to crude oil markets, the model implies that financialization has essentially no impact on the mean spot price, and reduces the frequency of low-inventory states and stockouts. But when stockouts occur, spot prices soar even higher than before financialization, leading to higher spot and futures price volatility. In addition the futures risk premium decreases, and futures open interest becomes more correlated with the spot price.

¹⁵A reasonable question is whether NOCs hedge. At least some, such as Pemex, are known for active hedging programs. See for example Parraga and Gibbons [2014].

These findings are robust to several changes in the model specification, such as driving fluctuations in U.S. oil consumption with foreign demand shocks as opposed to production shocks, or allowing households to trade futures contracts of differing time-to-expiration. However the finding that financialization increases spot price volatility is sensitive to the identity of new entrants. If new entrants are portfolio investors who only trade futures, financialization has little impact on spot prices, but implications for futures price and open interest dynamics remain.

The model suggests theoretical hedging and storage mechanisms through which financialization could impact spot and futures markets, providing insight and perhaps a laboratory for policy makers considering regulation. However I do not formally test these mechanisms against alternatives. Similar changes in moments may arise through different mechanisms, for example through changes in the distribution of wealth among oil consumers as in Baker and Routledge [2015], or through investment to develop increasingly difficult to extract oil deposits as in David [2014]. Whatever the effects of financialization, they were surely felt alongside other changes occurring in commodity markets. Future work should consider the interaction of these mechanisms to assess their joint effects.

A Numerical solution procedure

The state vector $z_t = (q_{t-1}, y_{t-1}, y_t)$ has one continuous variable, entering inventory q_{t-1} . I approximate q_{t-1} using a grid of 96 points, with more points allocated to smaller values of q_{t-1} . Solution functions use shape-preserving piecewise cubic interpolation over q_{t-1} conditional on the discrete production state(s).

I solve the model by policy iteration.

1. Initialize functions for the futures price $F(q_{t-1}, y_{t-1})$, household futures position $\Omega(q_{t-1}, y_{t-1})$, and dealer's inventory $Q(q_{t-1}, y_{t-1}, y_t)$.
2. For each gridded state z_t , update the functions by finding q_t , f_{t-1} , and ω_{t-1} such that the commodity spot market and the 1-period futures market clears and the household's FOC is satisfied.
3. Evaluate convergence. If the criterion is not satisfied, go to step (2).

In step (1) I use a simple linear guess for Q , the no-arbitrage relation with spot prices for F , and constant $\Omega = 0$ to solve the pre-financialization calibration. I use this solution to initialize subsequent calibrations with progressively lower values of τ . In step (2), functions from the previous iteration in combination with first order conditions furnish the necessary prices and quantities to evaluate market clearing and optimality. Step (2) is easily parallelized. In step (3) the convergence criterion is a scaled sum of the l^2 norms of the difference of functions in successive iterations. The solution procedure is implemented in Matlab.

B Proofs

Lemma 1. *Assume equilibrium spot and futures prices are positive and bounded. There exists $\tau > 0$ such that households will always prefer not to trade futures: $\omega_t = 0, \forall z_t$.*

Proof. This follows from inspection of the household's net endowment \hat{x}_t . The household can guarantee $\hat{x}_t = 1$ in all states by selecting $\omega_t = 0$, whereas for sufficiently large τ , $\hat{x}_t < 1$ for any bounded spot and futures prices and $\omega_t \neq 0$. \square

The following proposition and corollary assume that an equilibrium exists, and that $\sigma_{s,t}^2$ is positive and finite.

Proposition 1. *Suppose the producer shorts the futures contract before financialization, $\phi_t < 0$. Then the dealer is long futures, the risk premium is positive, and the producer faces predominantly price risk: $\psi_t > 0$, $E_t[s_{t+1}] - f_t > 0$, and $\sigma_{sys,t} > 0$.*

Proof. Since $\omega_t = 0$ before financialization, $\psi_t = -\phi_t > 0$ follows from market clearing. Since $q_t \geq 0$, $E[s_{t+1}] - f_t > 0$ follows from $\psi_t > 0$ and the dealer's first order condition Equation (6). The final result $\sigma_{sys,t} > 0$ follows from the producer's first order condition Equation (3). \square

Corollary 1. *Suppose the producer faces predominantly price risk before financialization, $\sigma_{sys,t} > 0$. Then the risk premium is always positive: $E_t[s_{t+1}] - f_t > 0$.*

Proof. Suppose to the contrary that $E[s_{t+1}] - f_t \leq 0$. Then from Equation (6) $\psi \leq 0$, and from Equation (3) $\phi < 0$. But then the futures market does not clear. Therefore any equilibrium with $\sigma_{sys,t} > 0$ must have $E_t[s_{t+1}] - f_t > 0$. \square

Proposition 2. *Suppose the economy has high-inventory, such that $q_t > 0$ and $q_{t+1} > 0$ for all y_{t+1} , and assume $\tau > 0$. Then the household prefers not to trade the two-period futures contract: $\omega_{t,2} = 0$.*

Proof. Suppose instead that the household chose futures position $\omega_{t,2} > 0$. In a high-inventory state, this position has payoff $\omega_{t,2}(1+r)(s_{t+1} - (1+\tau)f_{t,1} - \tau k)$. But the household could instead choose a position in the 1-period contract $\omega_{t,1} = \omega_{t,2}(1+r)$, which has a strictly higher payoff $\omega_{t,2}(1+r)(s_{t+1} - (1+\tau)f_{t,1})$ for all y_{t+1} . Therefore the household will not choose $\omega_{t,2} > 0$. The argument for $\omega_{t,2} < 0$ is symmetrical. It follows that the household chooses $\omega_{t,2} = 0$. \square

C Data sources and calibration details

I use end-of-month spot and futures prices from quandl.com, available from January 1990 through April 2015 for the nearest 18 continuous futures contracts. Where relevant, spot prices are converted from real to nominal using the Personal Consumption Expenditures index from the U.S. Bureau of Economic Analysis, from which I also obtain U.S. GDP data. Monthly open interest data aggregated over all crude oil futures contracts is from the U.S. Commodity Futures Trading Commission (CFTC)¹⁶, from January 1990 through April 2015. Data on crude oil quantities is from the Energy Information Administration (EIA).¹⁷ Because monthly data on quantities is limited, I use annual data for the period 1973-2014. To convert from barrels to constant dollars only for the purpose of comparing the value of U.S. oil consumption to U.S. GDP, I use the annual average real spot price from the St. Louis Fed's FRED database for the period 1973-2014. Summary statistics for data on prices and quantities are shown in the left column of Table 2. Moments for prices are split into 1990-2003 and 2004-2015 samples, for comparison with pre-and-post-financialization model moments, to be discussed later.

Dynamic storage models in previous papers such as Deaton and Laroque [1992] and Routledge et al. [2000] do not explicitly model or calibrate to consumption and production dynamics. Doing so imposes additional restrictions, and necessitates some discussion of how the model is mapped

¹⁶<http://www.cftc.gov/MarketReports/CommitmentsofTraders/HistoricalCompressed/index.htm>

¹⁷<http://www.eia.gov/petroleum/data.cfm#crude>

to the data. Within the model, the crude oil spot price reflects crude oil consumption, which in turn reflects production smoothed by inventory changes. I focus on production and consumption in the calibration, with inventory as the residual. Given the use of U.S. spot prices, I fit the model to U.S. consumption data also. However the choice of U.S. or global production data to fit the production process is less obvious, because oil consumed in the U.S. is sourced globally. Mechanically, production and consumption will be positively correlated within the model. The left panel of Figure C1 suggests positive correlation between global production and U.S. consumption in the data, whereas U.S. production appears inversely related to U.S. consumption. The right panel shows that global production and U.S. consumption are positively correlated (Pearson's coefficient 0.76) after removal of separate exponential time trends; accordingly I use global production data.

Since I abstract from growth in the model, I calibrate the production process y_t to roughly match fractional variation in global production relative to trend. I use a discrete approximation to an AR(1) process to model persistent production tractably and with a small number of parameters. 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. To compare with the data I estimate AR(1) parameters from global production via the Yule-Walker method, and adjust the parameters to a quarterly frequency. Adjusted values are in the bottom left panel of Table 2. The 5-state Markov process cannot match the high persistence of production (autocorrelation of 0.97) unless the probability of transitioning state is near zero on a quarterly basis. I choose autocorrelation of 0.6 and conditional standard deviation to 0.035 as inputs to Tauchen's algorithm, which matches the unconditional standard deviation of production and allows a range of production outcomes comparable to those illustrated in the right panel of Figure C1. 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.

As a broad measure of trade in futures, Table 2 reports open interest as a fraction of quarterly crude oil consumption. This equates open interest across a variety of maturities in the data with

open interest for the three-month contract in the model. Section 6.2 discusses the distribution of open interest by maturity in the data, and the extension of the model to multiple tradable contracts.

Although the CFTC provides data on the futures positions for several categories of market participant, they do not cleanly isolate producers or consumers as categories. Even if they did, producers and consumers also participate indirectly through intermediaries, such as swaps dealers, hedge funds, and index funds. Regarding producers, Acharya et al. [2013] 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, who are the primary U.S. consumers of crude oil in the form of gasoline, I assume that broad participation in the futures market was impossible prior to the proliferation of managed funds that hold crude oil futures, often as part of a broader commodity futures basket. The first commodity fund targeting retail investors was the Oppenheimer Real Asset Fund, established in 1996 with the purpose of pursuing investments linked to the GSCI index.¹⁸ However interest in commodity index funds increased among investors after 2004, as return characteristics documented in Gorton and Rouwenhorst [2006] became better known. The first commodity index ETF was the DB Commodity Index Tracking fund, established January 2006.¹⁹ The oil-only ETF USO began trading in April, 2006. Breaking the sample at the beginning of 2004 is consistent with Baker and Routledge [2015], and similar to Hamilton and Wu [2014], who split the sample at the beginning of 2005.

For the 2004-2015 period, I use figures from Stoll and Whaley [2010] and CFTC [2008]. According to 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 those such as pension funds that should, in principle, have some objectives in common

¹⁸See prospectus at <http://www.sec.gov/Archives/edgar/data/1018862/0001018862-97-000003.txt>

¹⁹See <http://www.sec.gov/Archives/edgar/data/1328237/000119312506118678/d424b3.htm>

with households. I attribute 25% of open interest to households, around 350,000 contracts at the start of 2008. This implies that households hedge around 20% of quarterly consumption, assuming that all of their hedging is done via indexed investment implemented using futures. This is, of course, a very rough estimate.

ETFs in particular proliferated during the financialization period, and warrant some additional discussion, as sophisticated investors such as hedge funds also use ETFs to partially implement their strategies. Estimates of commodity ETF ownership shares are available from industry research articles. Mercado [2017] reports institutional ownership of commodity exchange traded products (ETFs, ETNs, etc.) ranging from 25% to 40% from 2005 to 2016 (their Figure 24). Their definition of institutional owners is broad, and those investors not classified as institutional are considered retail investors. Therefore the majority of commodity ETP shares are thought to be owned by retail investors. Of institutional investors in ETFs (of all types), most are investment advisers (55% of institutional investor share) and private banks / wealth managers (25% of institutional investor share), with hedge funds only about 3% of institutional investor share, as of 2016. The objectives of investment advisers and wealth managers — asset allocation rather than speculation — align better with retail investors than hedge funds. From [Mercado, 2017, p. 17]: “In our experience, retail investors, investment advisers, and private banks tend to behave more like investors, that is they focus on the right exposure ... Brokers and hedge funds, on the other hand, are liquidity seekers and behave like traders ... At the end of 2016, almost all of the top 15 ETP issuers by institutional assets had at least 90% in the hands of asset allocators such as retail, investment advisers, and private banks.” Although informal, this assessment suggests that only a small fraction of ETF owners should be characterized as speculators.

C.1 Crude oil inventory

The model is not directly calibrated to match data on inventories. In earlier work such as Deaton and Laroque [1992, 1996] and Routledge et al. [2000], storage parameters are calibrated to properties of prices, especially autocorrelation, rather than directly to data on inventories.²⁰ In this paper the approach is similar, since storage affects the autocorrelation of spot prices and the slope of the futures curve through the autocorrelation of household consumption, for example.

However, one implication of financialization in the model is an increase in mean inventory. Therefore a brief discussion of relevant data on inventories is in order. In selecting and comparing data to model, there are two main questions: which inventory data best proxies for speculative inventory in the model, and how should inventory data be transformed to be consistent with the stationary detrended production process that drives the model?

We examine three time-series for inventory, shown in the left panel of Figure C2: total petroleum stocks in the OECD and the US, and crude oil stocks at tank farms and pipelines in the Midwestern US (PADD 2). All data is from the EIA. Following Kilian and Murphy [2014], we take OECD stocks as a proxy for world inventory. We show all series from 1984, when OECD data becomes available, and normalize all starting values to 1. OECD and US inventory increase by 10-20% over the post-financialization period starting in 2004, although there is no clear break beginning in that period. OECD and especially total US inventory may not be a good proxy for speculative inventory. For example, the Strategic Petroleum Reserve (SPR) constitutes roughly one third of total US petroleum inventory over the period, and some fraction of private inventory is also held for operational purposes. In some instances accumulation and release from the SPR may act to offset

²⁰To my knowledge, no paper calibrates or estimates the canonical storage model using crude oil inventory data and prices simultaneously. The most closely related paper is Dvir and Rogoff [2014], in which variants of the canonical model, with and without stochastic growth trends, are used to motivate differently signed cointegrating vectors for oil production, stocks, real prices, and global income. However they do not fit the model itself to data.

price shocks similar to speculative inventory, but they also occur for unrelated political reasons.²¹

As a narrower proxy for speculative inventory, Figure C2 also shows tank and pipeline inventory in the Midwestern PADD 2 region containing the Cushing, Oklahoma delivery point for NYMEX WTI futures contracts. PADD 2 inventory rises rapidly starting around the 2008 financial crisis. Storage capacity also increased rapidly during this period, in response to high utilization rates.²² While clearly relevant to the WTI futures market, which saw frequent contango characteristic of high inventory during this period, high Midwestern inventory also reflected transportation bottlenecks from rapid regional tight-oil development.

To render data comparable to model moments, the right panel of Figure C2 shows days supply of inventory. Since OECD stocks proxy for world inventory I calculate days supply from world crude oil production, whereas for US inventory data I calculate from US product supplied. Since oil supply increased significantly over the sample period, days supply looks quite different from raw inventory, with an overall decline in world days supply, and a temporary decline in US days supply prior to financialization, followed by a rebound thereafter. Inventory does not exhibit a simple cointegrating relationship with production: Kilian and Murphy [2014] find no evidence of cointegration between inventory and production, whereas Dvir and Rogoff [2014] find a time-varying cointegrating vector for oil production, inventory, real spot price, and global income. Therefore it is not obvious that a decline in days supply in the data should map to a decline in days supply in the model, where inventory and production are cointegrated by assumption.

The series do provide a sense of scale: mean days supply in the model calibrations range from around five days without financialization to around seven days with financialization, roughly comparable to the days of US supply in PADD 2 tanks alone, which rise from around 3 days before financialization to around 5 days afterwards. By contrast total US days supply is much larger than

²¹See <https://www.eia.gov/todayinenergy/detail.php?id=24072> for an example of recent legislative action concerning the SPR, and for a summary of historical SPR releases.

²²See <https://www.eia.gov/todayinenergy/detail.php?id=26772>

in the model, ranging from roughly 80 to 100 days. Stockouts also occur with non-negligible probability in the model, but are not observed in the data.

Overall there is some evidence of increased inventory levels following financialization, especially in the region of the main delivery point for US futures contracts, but there is little evidence of a change in global regime. To provide a rough estimate of inventory relevant to futures market behavior in the pre-and-post-financialization periods, Table 2 reports average days of supply for the PADD 2 region for 1990-2003 and 2004-2015.

References

- V. V. Acharya, L. A. Lochstoer, and T. Ramadorai. Limits to arbitrage and hedging: Evidence from commodity markets. *Journal of Financial Economics*, 109(2):441–465, 2013.
- 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, 2015.
- S. Basak and G. Chabakauri. Dynamic mean-variance asset allocation. *Review of Financial Studies*, 2010.
- S. Basak and A. Pavlova. A model of financialization of commodities. *The Journal of Finance*, 71(4):1511–1556, 2016.
- H. Bessembinder and M. L. Lemmon. Equilibrium pricing and optimal hedging in electricity forward markets. *the Journal of Finance*, 57(3):1347–1382, 2002.
- 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. Büyüksahin, M. S. Haigh, and M. A. Robe. Commodities and equities: Ever a” market of one”? *The Journal of Alternative Investments*, 12(3):76, 2010.
- B. Buyuksahin, M. Haigh, J. Harris, J. Overdahl, and M. Robe. Fundamentals, trader activity and derivative pricing. Technical report, American University, 2011.
- R. J. Caballero, E. Farhi, and P.-O. Gourinchas. Financial crash, commodity prices, and global imbalances. *Brookings Papers on Economic Activity*, 2008(2):1–55, 2008.

- C. Cafiero, E. Bobenrieth, J. Bobenrieth, 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. David. Exploration activity, long run decisions, and the risk premium in energy futures. Technical report, University of Calgary, 2014.
- 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.
- E. Dvir and K. Rogoff. Demand effects and speculation in oil markets: Theory and evidence. *Journal of International Money and Finance*, 42:113–128, 2014.
- M. Floden. A note on the accuracy of markov-chain approximations to highly persistent ar (1) processes. *Economics Letters*, 99(3):516–520, 2008.
- I. Goldstein and L. Yang. Commodity financialization: Risk sharing and price discovery in commodity futures markets. Technical report, University of Pennsylvania, 2016.
- G. Gorton and K. G. Rouwenhorst. Facts and fantasies about commodity futures. *Financial Analysts Journal*, 62(2):47–68, 2006.

- G. Gorton, F. Hayashi, and K. Rouwenhorst. The fundamentals of commodity futures returns. *Review of Finance*, 2012.
- 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.
- J. D. Hamilton and J. C. Wu. Risk premia in crude oil futures prices. *Journal of International Money and Finance*, 42:9–37, 2014.
- G. D. Haushalter. Financing policy, basis risk, and corporate hedging: Evidence from oil and gas producers. *The Journal of Finance*, 55(1):107–152, 2000.
- D. Hirshleifer. Residual risk, trading costs, and commodity futures risk premia. *Review of Financial Studies*, 1(2):173–193, 1988.
- D. Hirshleifer. Futures trading, storage, and the division of risk: a multiperiod analysis. *The Economic Journal*, 99(397):700–719, 1989.
- D. Hirshleifer. Hedging pressure and futures price movements in a general equilibrium model. *Econometrica: Journal of the Econometric Society*, pages 411–428, 1990.
- S. Hitzemann. Macroeconomic fluctuations, oil supply shocks, and equilibrium oil futures prices. Technical report, Ohio State University, 2016.
- 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.
- J. M. Keynes. *A Treatise on Money: The Applied Theory of Money, Vol. 2*. London: Macmillan, 1930.

- L. Kilian and D. P. Murphy. The role of inventories and speculative trading in the global market for crude oil. *Journal of Applied Econometrics*, 29(3):454–478, 2014.
- L. Kilian and R. J. Vigfusson. Do oil prices help forecast us real gdp? the role of nonlinearities and asymmetries. *Journal of Business & Economic Statistics*, 31(1):78–93, 2013.
- 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.
- M. Masters. Testimony before the committee on homeland security and governmental affairs. Technical report, US Senate, 2008.
- S. Mercado. The 2016 guide to institutional etf ownership. Technical report, Deutsche Bank Markets Research, April 2017.
- X. N. Pan. Investor Beliefs and State Price Densities in the Crude Oil Market. *SSRN eLibrary*, 2011. doi: 10.2139/ssrn.1988449.
- M. Parraga and R. Gibbons. Mexico said to begin 2015 oil hedge. *Reuters Market News*, August 2014.
- C. Pirrong. *Commodity Price Dynamics: A Structural Approach*. Cambridge University Press, 2011.
- R. Ready. Oil prices and long-run risk. Technical report, University of Rochester, 2012.
- R. C. Ready. Oil consumption, economic growth, and oil futures: A fundamental alternative to financialization. Technical report, University of Rochester, 2014.
- 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. *Management Science*, 60(2): 300–318, 2014.

- M. Sockin and W. Xiong. Informational frictions and commodity markets. *The Journal of Finance*, 2015.
- 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 the financialization of commodities. *Financial Analysts Journal*, 68(5):54–74, 2012.
- K. Tang and H. Zhu. Commodities as collateral. *The Review of Financial Studies*, 29(8):2110–2160, 2016.
- G. Tauchen. Finite state markov-chain approximations to univariate and vector autoregressions. *Economics letters*, 20(2):177–181, 1986.
- S. Tordo. *National oil companies and value creation*. The World Bank, 2011.
- P. Tufano. Who manages risk? an empirical examination of risk management practices in the gold mining industry. *The Journal of Finance*, 51(4):1097–1137, 1996.
- J. Vercaemmen and A. Doroudian. Portfolio speculation and commodity price volatility in a stochastic storage model. *American Journal of Agricultural Economics*, 2014.
- D. G. Victor, D. R. Hults, and M. C. Thurber. *Oil and governance: state-owned enterprises and the world energy supply*. Cambridge University Press, 2011.
- J. Williams and B. Wright. *Storage and Commodity Markets*. Cambridge University Press, 1991.

Table 1: **Parameters**

Parameter	Value	Purpose
γ	0.06	oil preference
η	-15	elasticity of substitution over goods
k	0.001	nominal storage cost (per unit)
r	0.15	risk-free rate (%)
ρ	60	dealer risk aversion
θ	40	producer and household risk aversion

The above parameter values are used in all numerical examples throughout the paper. The storage cost and risk-free rate are per quarter.

Table 2: **Summary statistics**

Moment	Data		Model			
	1990-2003	2004-2015	$\tau = \infty$	$\tau = 0.15$	$\tau = 0.10$	$\tau = 0.08$
Mean Spot (2009 USD)	28.3	76.9	28.3	28.2	28.2	28.3
Autocorrelation Spot	0.630	0.653	0.669	0.583	0.576	0.578
Std. Dev. Spot (% of Mean)	22.5	25.6	41.1	46.8	47.8	48.2
Autocorrelation Futures	0.721	0.718	0.732	0.709	0.719	0.728
Std. Dev. Futures (% of Mean)	20.6	24.7	23.9	28.8	30.0	30.6
Risk Premium Futures (%)	1.89	-0.412	2.81	-0.120	-0.799	-1.02
Std. Dev. Return Futures (%)	16.4	19.1	22.7	23.5	23.3	23.1
Percent Backwardation	69.6	34.8	43.8	21.8	21.4	21.2
Mean Open Interest (% of Consumption)	24.0	71.4	27.5	52.1	60.9	66.5
Mean Household Fut. (% of Consumption)	0.000	20.0	0.000	15.9	21.7	25.5
Mean Inventory (Days Supply)	2.88	3.70	5.19	6.26	6.92	7.33
Mean Producer Fut. (% of Production)		-25.0	-27.6	-36.3	-39.4	-41.3
Mean Oil Expenditure (% of GDP)		3.25	3.70	3.69	3.69	3.70
Uncond. Std. Dev. Oil Production		0.0441	0.0474	0.0474	0.0474	0.0475
Cond. Std. Dev. Oil Production		0.0113	0.0381	0.0382	0.0381	0.0381
Autocorrelation Oil Production		0.966	0.595	0.594	0.594	0.595
Uncond. Std. Dev. Oil Consumption		0.0623	0.0345	0.0332	0.0325	0.0323
Cond. Std. Dev. Oil Consumption		0.0170	0.0240	0.0201	0.0189	0.0184
Autocorrelation Oil Consumption		0.961	0.718	0.796	0.813	0.822

The table shows summary statistics from data and the model. Statistics for quantities use annual data from 1973-2014, converted to quarterly units, except for inventory, for which starts in 1990. Appendix C.1 describes inventory data in more detail. Price and open interest statistics use monthly data from January 1990 through April 2015, with quarterly holding periods. Price statistics are split between pre (1990-2003) and post (2004-2015) financialization periods. Standard deviations are given relative to the mean price of the asset during the respective period. Data entries for futures positions are approximations - see Appendix C for details. 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 attempts to match 1990-2003 asset prices and full-sample quantities; parameter values are in Table 1. The model numeraire is mapped to 2009 USD such that the mean spot price for the pre-financialization model ($\tau = \infty$) matches the 1990-2003 mean in the data. Model moments are computed by Monte Carlo simulation over 1,000,000 periods.

Table 3: **Transfers from dealers to households**

Moment	Gross	Net
Mean (%)	0.0333	-0.0731
Std. Dev. (%)	0.957	0.963
Min (%)	-3.21	-3.75
Max (%)	6.71	6.29

Summary statistics for transfers from dealers to households via the futures market are shown gross and net of transaction costs paid by the household. All statistics are percentages of the household's endowment, computed for parameters in Table 1 with $\tau = 0.1$. Moments are computed by Monte Carlo simulation over 1,000,000 periods.

Table 4: **Summary statistics: alternative shocks**

Moment	Data		Model: y_t		Model: \hat{y}_t		Model: x_t	
	'90-'03	'04-'15	$\tau = \infty$	$\tau = 0.15$	$\tau = \infty$	$\tau = 0.15$	$\tau = \infty$	$\tau = 0.15$
Mean Spot (2009 USD)	28.3	76.9	28.3	28.3	28.4	28.5	27.9	27.4
Autocorrelation Spot	0.630	0.653	0.669	0.583	0.659	0.538	0.678	0.696
Std. Dev. Spot (% of Mean)	22.5	25.6	41.1	47.0	42.3	52.6	40.0	41.5
Autocorrelation Futures	0.721	0.718	0.732	0.710	0.720	0.695	0.730	0.685
Std. Dev. Futures (% of Mean)	20.6	24.7	23.9	28.9	23.4	29.3	23.8	25.9
Risk Premium Futures (%)	1.89	-0.412	2.82	-0.0552	4.41	0.636	3.60	5.11
Std. Dev. Return Futures (%)	16.4	19.1	22.7	23.5	24.0	25.9	22.8	23.7
Percent Backwardation	69.6	34.8	43.7	21.9	48.9	23.1	47.4	53.2
Mean Household Fut. (% of Cons.)	0.000	20.0	0.000	16.0	0.000	21.6	0.000	-33.4
Mean Producer Fut. (% of Prod.)		-25.0	-27.5	-36.2	-37.1	-49.4	-37.0	-14.8
Mean Oil Expenditure (% of GDP)		3.25	3.70	3.69	3.71	3.72	3.66	3.58
Uncond. Std. Dev. Oil Production		0.0441	0.0474	0.0474	0.000	0.000	0.000	0.000
Cond. Std. Dev. Oil Production		0.0113	0.0381	0.0381	0.000	0.000	0.000	0.000
Autocorrelation Oil Production		0.966	0.594	0.595	0.000	0.000	0.000	0.000
Uncond. Std. Dev. Oil Consumption		0.0623	0.0345	0.0333	0.0357	0.0343	0.0239	0.0210
Cond. Std. Dev. Oil Consumption		0.0170	0.0240	0.0202	0.0252	0.0208	0.0218	0.0192
Autocorrelation Oil Consumption		0.961	0.719	0.796	0.709	0.795	0.415	0.400

The table shows summary statistics for the model with different sources of randomness: oil production (y_t), foreign demand (\hat{y}_t), and domestic demand (x_t) shocks. The stochastic process for y_t in the baseline calibration is assigned to each process in turn, while the remaining processes are held constant at unity. For each source of randomness the table shows pre-financialization ($\tau = \infty$) and post-financialization ($\tau = 0.15$) results. Other parameter values are per Table 1. Data statistics are repeated from Table 2. Model moments are computed by Monte Carlo simulation over 1,000,000 periods.

Table 5: **Summary statistics with portfolio investor**

τ	∞	0.1	∞	∞	∞	∞
τ_{ζ}	∞	∞	0.005	0.1	0.1	0.1
$\rho_{s,v}$	N/A	N/A	0	-0.2	0	0.2
Mean Household Fut. (%)	0.000	21.8	0.000	0.000	0.000	0.000
Mean Port. Inv. Fut. (%)	0.000	0.000	21.0	32.6	2.00	-1.60
Mean Producer Fut. (%)	-27.6	-39.3	-38.8	-45.6	-28.7	-26.6
Mean Open Interest (%)	27.5	60.9	38.7	45.4	28.7	26.6
Mean Inventory (Days Supply)	5.17	6.93	6.30	6.91	5.22	5.17
Mean Spot (2009 USD)	28.3	28.1	28.1	28.1	28.2	28.2
Autocorrelation Spot	0.670	0.578	0.707	0.727	0.669	0.668
Std. Dev. Spot (% of Mean)	41.2	48.1	39.1	38.2	41.1	41.2
Autocorrelation Futures	0.732	0.720	0.752	0.763	0.729	0.728
Std. Dev. Futures (% of Mean)	23.9	30.1	26.0	27.5	24.5	23.9
Risk Premium Futures (%)	2.89	-0.761	0.716	-0.473	2.50	2.98
Std. Dev. Return Futures (%)	22.7	23.3	20.3	19.4	22.6	22.8
Percent Backwardation	44.0	21.5	34.0	25.6	42.9	43.8


The table compares summary statistics for the model before financialization ($\tau = \tau_{\zeta} = \infty$), after financialization due only to household entry ($\tau = 0.1$, $\tau_{\zeta} = \infty$), and after financialization due only to portfolio investor entry ($\tau = \infty$, $\tau_{\zeta} \leq 0.1$). Values for parameters varied across experiments are given in the top section of the table. The standard deviation of the portfolio investor's non-futures income is $\sigma_v = 7.5\%$. All remaining parameter values follow Table 1. Model moments are computed by Monte Carlo simulation over 1,000,000 periods.

Table 6: Summary statistics with producer equity

τ	∞	0.1	$\infty/0.1$	∞	0.25	0.1	0.08
$\bar{\pi}^p$	1	1	$-\infty$	0.75	0.75	0.75	0.75
Mean Household Shares (%)	0.000	0.000	129.	25.0	25.0	25.0	25.0
Mean Producer Shares (%)	100	100	-29.5	75.0	75.0	75.0	75.0
Mean Household Futures (%)	0.000	21.8	0.000	0.000	4.26	12.1	14.7
Mean Producer Futures (%)	-27.6	-39.4	16.4	-18.6	-21.1	-25.3	-26.6
Mean Open Interest (%)	27.5	61.1	16.5	18.7	25.4	37.3	41.2
Corr. Equity and Fut. Ret. (%)	N/A	N/A	99.4	89.8	93.0	98.4	98.9
Mean Inventory (Days Supply)	5.18	6.90	12.9	6.40	6.55	7.47	7.84
Mean Spot (2009 USD)	28.3	28.2	27.9	28.0	27.9	28.1	28.1
Autocorrelation Spot	0.669	0.576	0.640	0.677	0.650	0.584	0.583
Std. Dev. Spot (% of Mean)	41.2	48.1	42.6	40.8	42.4	47.1	47.3
Autocorrelation Futures	0.732	0.719	0.783	0.750	0.741	0.727	0.732
Std. Dev. Futures (% of Mean)	23.9	30.1	30.1	25.0	26.1	29.7	30.2
Risk Premium Futures (%)	2.88	-0.757	-2.18	1.37	0.790	-0.863	-1.14
Std. Dev. Return Futures (%)	22.7	23.4	18.8	21.3	21.7	22.7	22.5
Percent Backwardation	44.0	21.5	14.3	36.5	27.8	20.7	20.4

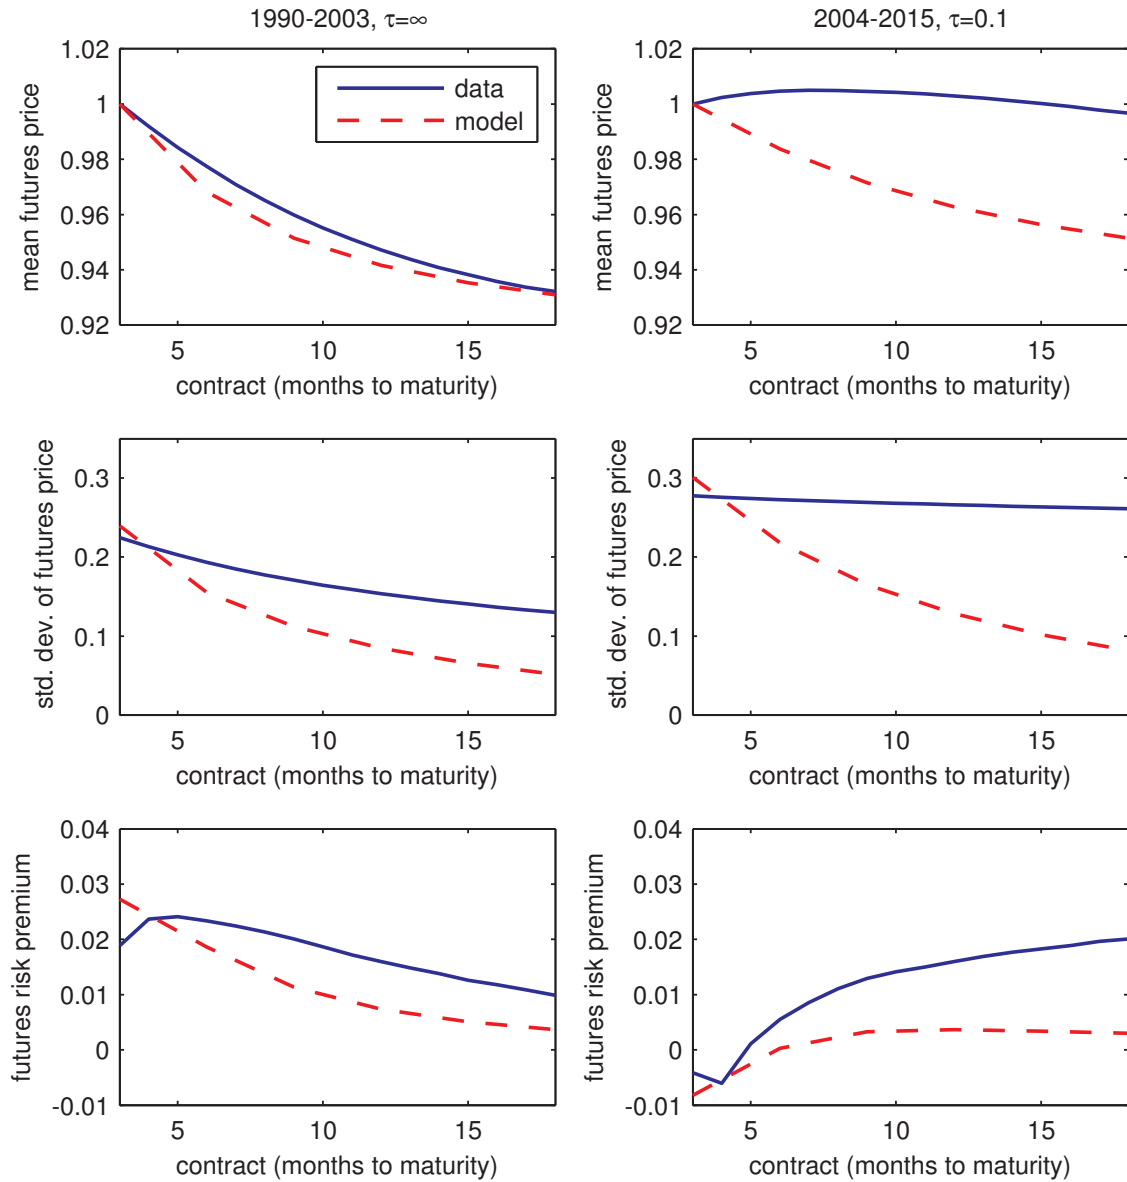
The table reports summary statistics illustrating the effect of producer equity issuance. The first two result columns are for the baseline model without and with financialization. In the third column the producer may issue equity without restriction, and results are the same with or without financialization, as the household will choose to hold no futures after unrestricted equity trading. In the last set of columns the producer faces a binding constraint to retain 75% of equity. The first column of the last set is without financialization, the remaining columns are for increasing levels of financialization (decreasing τ). Values for parameters varied across experiments are given in the top section of the table. All remaining parameter values follow Table 1. Model moments are computed by Monte Carlo simulation over 1,000,000 periods.

Figure 1: State-contingent inventory and spot price



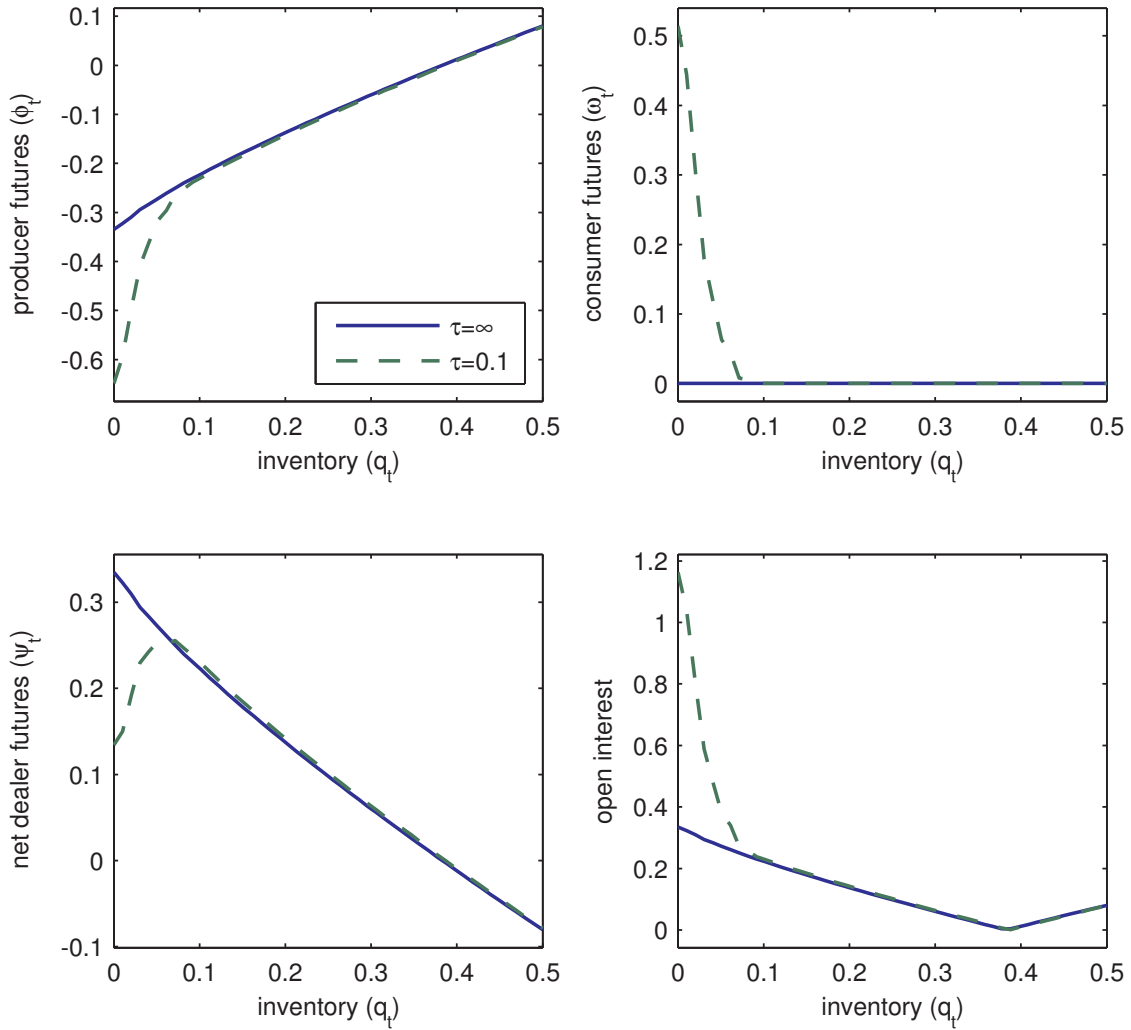
Numerical results are from the baseline model with parameter values per Table 1 and $\tau = \infty$. The left plot shows optimal inventory policy. The dealer's end of period inventory q_t depends on inventory carried over from the previous period (q_{t-1} , given on the x-axis), and the production realization (y_t , shown by the different curves). The 45° line shown in dashed red corresponds to unchanged inventory, with the region above that line indicating inventory accumulation and the one below indicating sell-off. The right panel shows the spot price conditional on entering inventory (q_{t-1} , x-axis) and realized production (y_t , different curves).

Figure 2: Term structure of futures moments



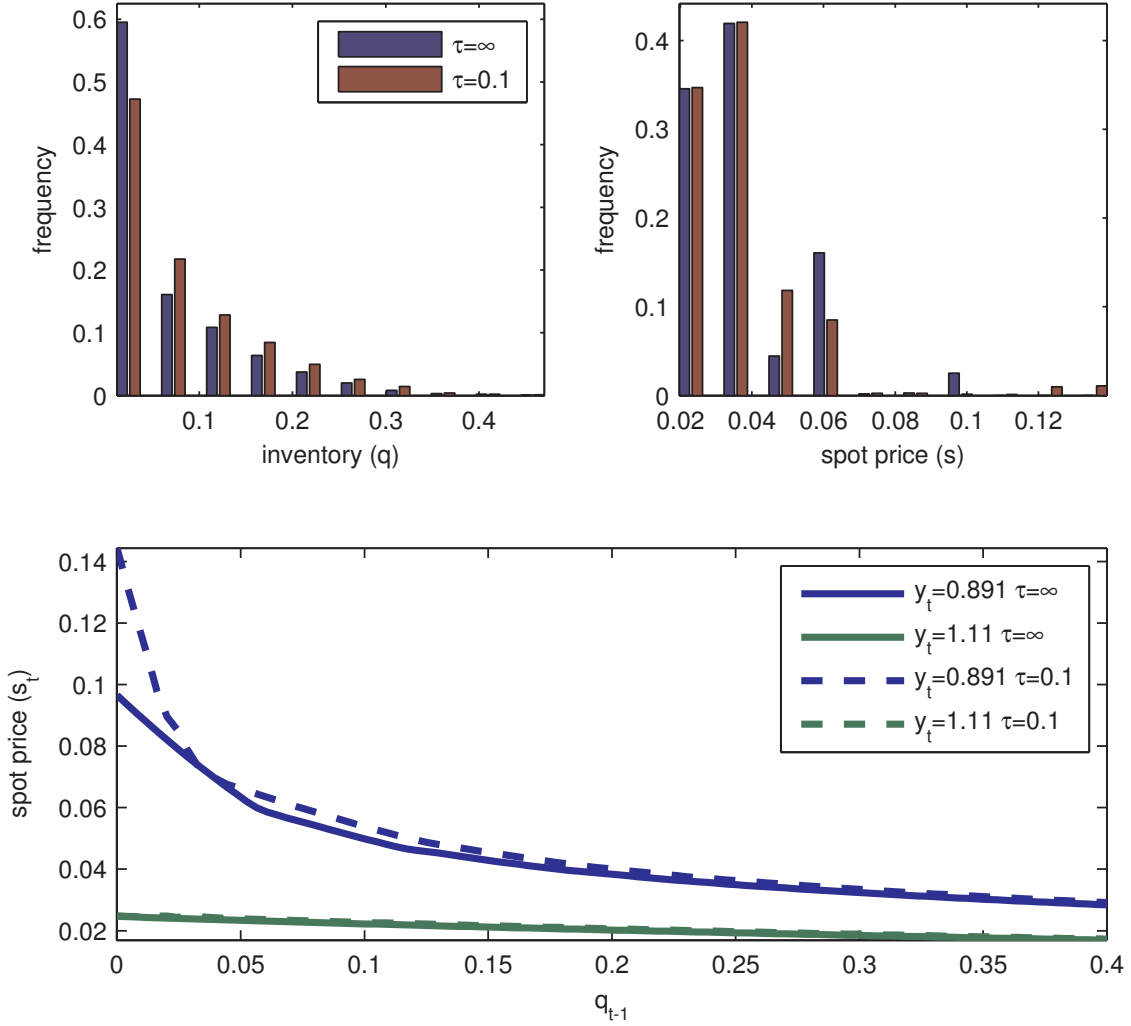
The figure shows unconditional moments of futures contracts in the data (blue solid line) and in the model (red dashed line). The number of months to delivery varies on the x-axis. The left column compares the pre-financialization data and model (1990-2003, $\tau = \infty$), whereas the right column shows post-financialization data and model (2004-2015, $\tau = 0.1$). Rows show the unconditional mean futures prices, quarterly standard deviation of prices, and expected excess returns with quarterly rolls.

Figure 3: Futures positions and open interest



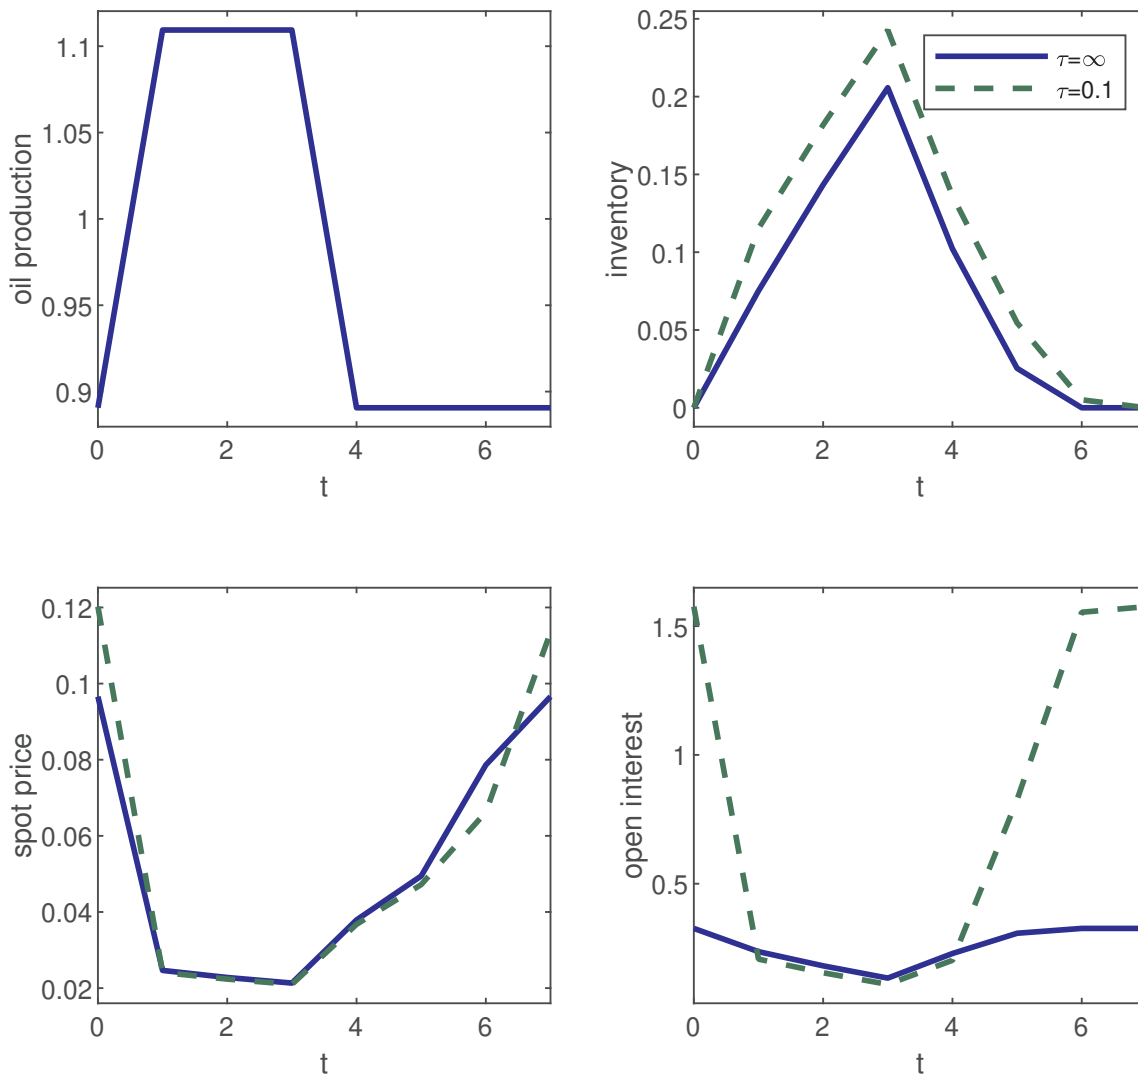
Positions and open interest in the 3-month futures contract are plotted for the pre-financialization model ($\tau = \infty$, solid line) and the post-financialization model ($\tau = 0.1$, dashed line). Open interest is the total number of contracts outstanding: $|\phi_t| + |\omega_t|$. One futures contract contract calls for delivery of the mean aggregate quarterly supply of oil. Results are conditional on the level of inventory (q_t , x-axis) but unconditional of production (y_t); expectations are taken using the stationary distribution of production.

Figure 4: Inventory and spot prices before and after financialization



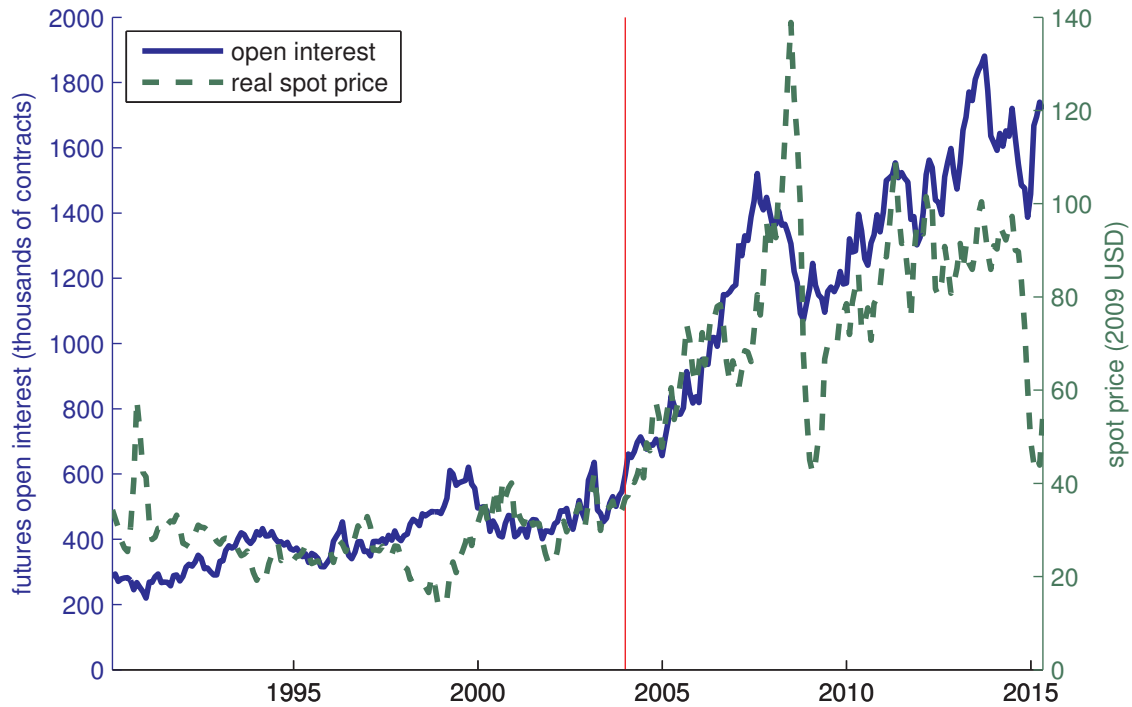
The top row above shows the unconditional distribution of inventory (q , left panel) and spot prices (s , right panel) in the model. The pre-financialization ($\tau = \infty$) distribution is the left pairwise columns, in blue, the post-financialization distribution ($\tau = 0.1$) is the right pairwise columns, in red. Expected inventory is higher after financialization. The bottom plot shows spot prices conditional on inventory for the lowest production state (top two curves) or the highest production state (overlapping bottom curves), in the pre-financialization ($\tau = \infty$) and post-financialization ($\tau = 0.1$) model.

Figure 5: Example model path



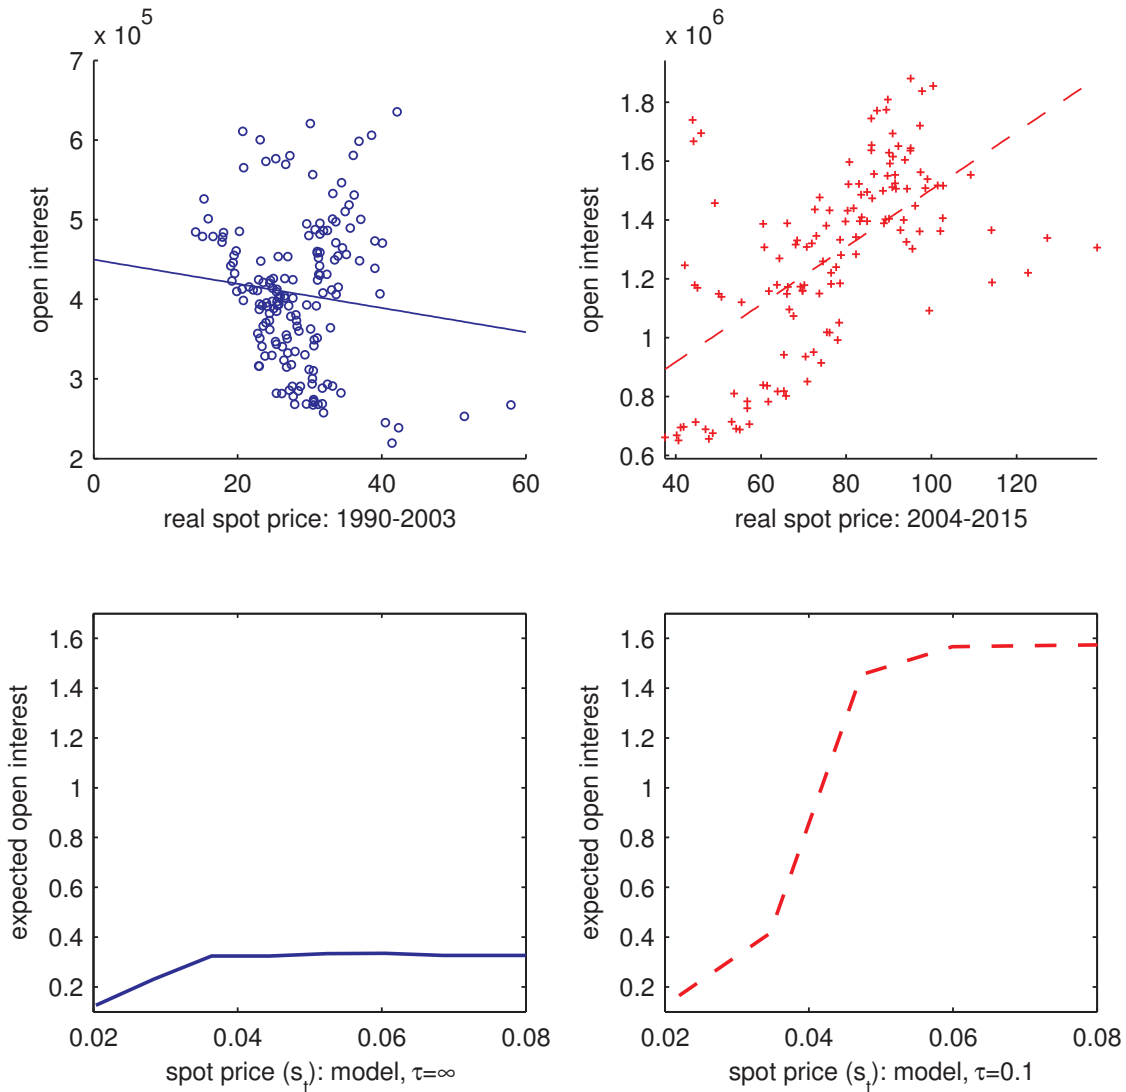
The figure compares the response of the pre-financialization ($\tau = \infty$, solid blue line) and post-financialization ($\tau = 0.1$, dashed green line) models to the same series of shocks. Production realizations, identical for both models, are in the top left panel. Inventory, which begins at zero for both models, is in the top right panel. The spot price and open interest are in the bottom left and bottom right panels, respectively.

Figure 6: WTI crude oil real spot price and futures open interest



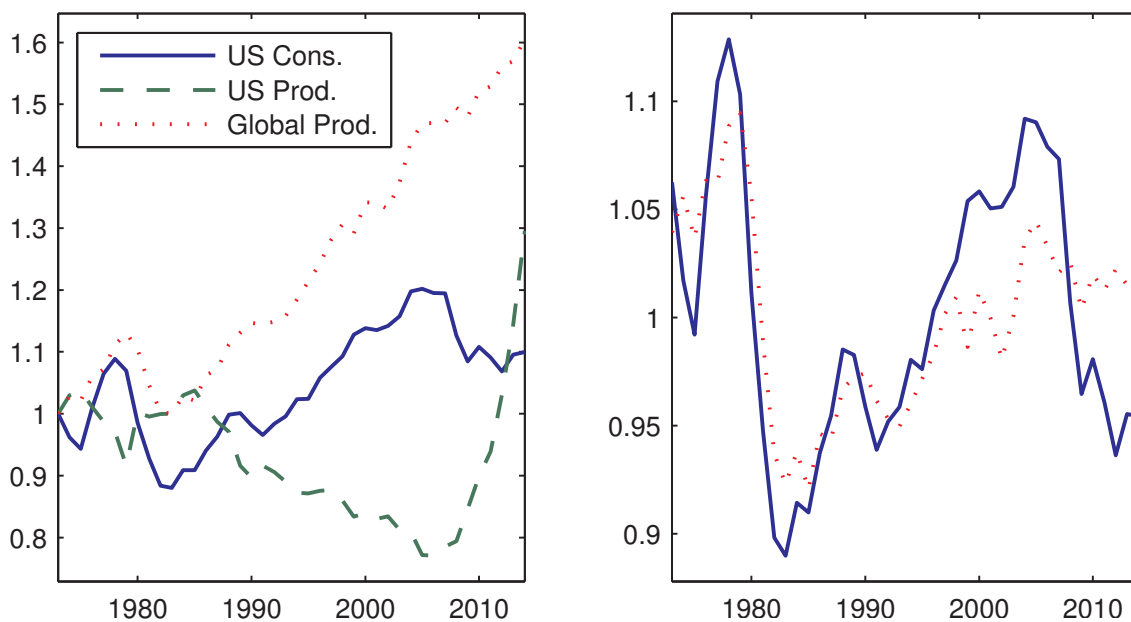
The plot shows the WTI crude oil real spot price (right axis, NYMEX data from Quandl) and futures open interest measured in thousands of contracts (left axis, from the CFTC) from Jan. 1990 through April 2015. Nominal prices are converted to 2009 USD using the PCE index from the BEA. The vertical red line is January 1 2004, the date used to split the sample into pre-financialization and post-financialization periods.

Figure 7: Correlation of spot price and open interest



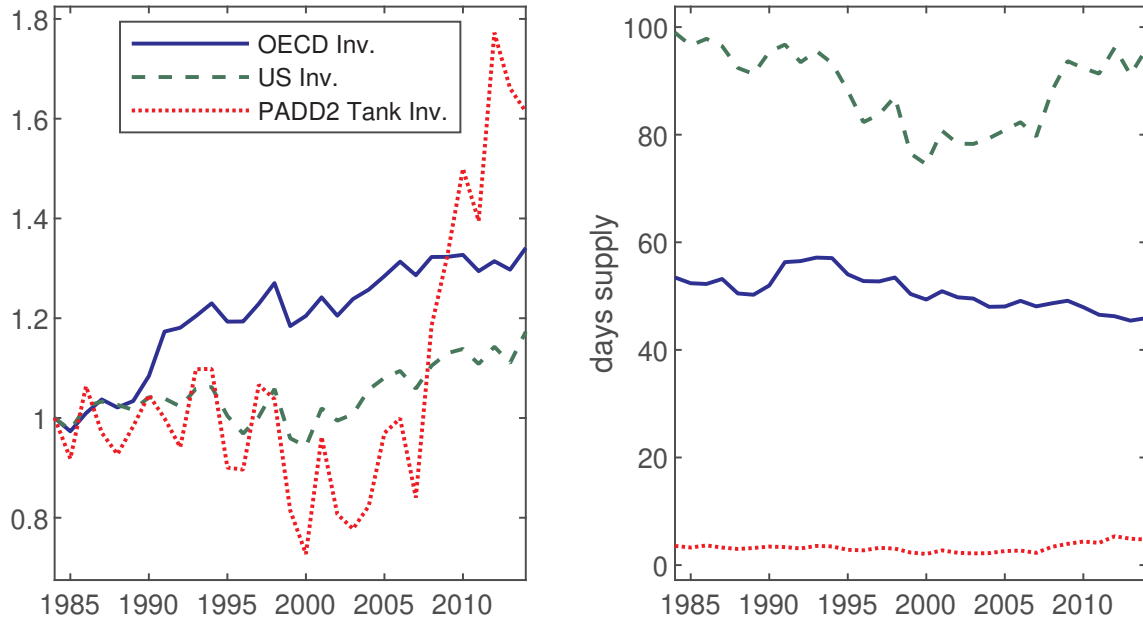
The figure illustrates the relationship between the spot price of crude oil (x-axis) and open interest in crude oil futures contracts (y-axis), in the data (top row) and the model (bottom row). The left column shows pre-financialization data (1990-2003) and model ($\tau = \infty$), the right column shows post-financialization data (2004-2015) and model ($\tau = 0.1$). Trend lines in for the data are fitted using simple OLS regression. Prior to financialization the relationship between open interest and prices is not statistically significant (top left, $R^2 = 0.012$, $p = 0.164$). After financialization there is a positive and statistically significant relationship (top right, $R^2 = 0.364$, $p = 0.000$). Axes for data are scaled to reflect increases in open interest and the spot price after 2004. Results in the bottom row show expected open interest (number of contracts) conditional on the spot price in the model, where one contract calls for delivery of the mean aggregate quarterly supply of oil.

Figure C1: Crude oil consumption and production



In the left panel, annual U.S. oil consumption is shown alongside global and U.S. production. Each series is normalized by its initial value. The right panel illustrates the positive correlation of detrended annual U.S. consumption and global production. Detrended U.S. consumption and global production have Pearson's correlation coefficient of 0.76. Data is from the Energy Information Administration.

Figure C2: Crude oil inventory



The left panel shows inventory, in log of thousands of barrels, for all OECD countries (total petroleum stocks), the US (total petroleum stocks), and at tank farms and pipelines in the Midwestern US region (PADD 2). OECD stocks are used to approximate world stocks, and PADD 2 stocks approximate ready inventory near the delivery point for NYMEX WTI crude oil futures at Cushing, Oklahoma. The right panel shows days supply, relative to world product supplied for the OECD, and US product supplied for the US and midwestern US. Data is from the Energy Information Administration.