

# Commodity forward curves: models and data

Craig Pirrong

Bauer College of Business

University of Houston



# Commodity Forward Curves

- A forward curve is a locus of points relating the forward price to the associated delivery date
- Forward curves can take many shapes
- Upward sloping
- Downward sloping
- Humped



# Elementary Theory: Cash and Carry

- The simple theory of forward curves is that for storable commodities, forward prices should exceed nearby prices plus the cost of carrying inventory (interest, warehousing charges, insurance) to the expiration date
- Holds for virtually all financial forwards
- Seldom holds for commodity forwards: why not?



# Commodity Forward Curves

- Full carry (“carrying charge market”)
- Contango: forward prices are increasing with time to expiry, but do not necessarily cover carrying charges
- Backwardation (“inverted market”): prices fall with time to expiration



# Why the Diversity of Forward Curve Shapes?

- We need an explanation of the diversity of forward curve shapes
- We can draw on the basic insight of our earlier analysis: calendar spreads are price signals that guide the *intertemporal* allocation of commodities
- Intertemporal allocation economics will depend on the type of commodity
- Therefore need to have a basic taxonomy of commodity types



# A Commodity Taxonomy

- The Most Basic Divide
  - Storables
  - Non-storables (or effectively non-storables)
- Costs of storage in fact lie along a continuum, and it is perhaps better to think of things that are very costly to store vs. those that aren't quite so costly



# Non-Storables

- Electricity
  - Very costly to store, although hydro introduces an element of storability into some markets
- Freight
- Weather (duh!)
- Livestock



# A Taxonomy of Storables

- Continuously Produced
  - Exhaustible (e.g., oil)
  - Non-exhaustible (e.g., CU, AL)
- Periodically Produced
  - Grains and oilseeds
  - Tree crops





# Storables Forward Curves

- Again, the forward curve is providing price signals on how to allocate resources over time
- In a well functioning market, forward prices will reflect—and guide—the optimal intertemporal allocation
- So, to understand the forward curve in a well-functioning storables market, you need to understand the economics of intertemporal resource allocation



# Intertemporal Allocation Basics

- Want to carry goods from where they are abundant to where they are scarce but . . .
- With intertemporal allocation, can only move in one direction: from the present to the future
- Can't move in the other direction: from the future to the present (though that would be nice)
- This constraint is essential in understanding forward curves



# Scenario 1

- Supplies of the commodity are currently abundant because demand is low
- Substantial excess capacity
- Already have a lot of inventory
- Demand is expected to rebound, so commodity is abundant today relative to what we expect in the future
- Under these circumstances, optimal to store commodity for future use



# Forward Curves in Scenario 1

- Forward curve should adjust to provide an incentive to store
- Futures prices at or near full carry provides this incentive



# Scenario 2

- Current supplies very tight
- Inventories low, and operating at nearly full capacity
- Commodity scarce today relative to what is expected in the future
- Would like to bring stuff from the future to today, but can't
- The best we can do is NOT to store
- May want to consume all inventories



# Forward Curves in Scenario 2

- Want to punish storage
- If spot price is above forward price net of carrying costs, anticipate losing money by storing: certainly lose money if you hedge inventory
- Therefore, backwardation in time of scarcity gives the right price signals



# Implication

- Markets should be in full carry when inventory is very high
- Markets should be in backwardation when inventory is low
- “Supply of Storage Curve”



# A Refinement

- Market may be in full carry for very short-dated deliveries, but below full carry for more distant dates
- EG, LME metals: one day prices may be at carry, cash-3 month spread at less than full carry
- This occurs when there are positive inventories, but there is some possibility of a stockout in the period prior to the expiry of the forward in question
- Departure from full carry/amount of backwardation depends on likelihood of stockout over period prior to expiry: higher the likelihood, bigger departure from full carry
- In a competitive market, backwardation between spot and 1 day forward occurs only during a stockout





# Additional Implications for Continuously Produced Commodities

- The foregoing analysis implies that the shape of the forward curve depends on scarcity
- Price volatility and correlations should also depend on scarcity
- Therefore, volatility and correlation should be related to the shape of the forward curve



# Implication: Volatilities

- Volatilities should be higher when the market is in backwardation than in contango
- Spot and forward prices should be equally volatile when the market is at full carry because storage connects the prices
- Spot prices should be more volatile than forward prices when the market is in backwardation because there is more flexibility to adjust quantity over a longer period of time, and this ability reduces the burden placed on price to respond to demand shocks
- Some empirical evidence supports this, but cannot explain all of the dynamics of volatility



# Implication: Correlations

- Correlations between spot and forward prices should be near 1.00 when market is at full carry, because storage connects these prices
- Since backwardation signals the possibility of a stockout that would break the direct link between spot and forward prices, correlations along the curve should be lower in backwardated markets
- Data support this



# Implication: Mean Reversion

- “Mean reversion” means that prices tend to rise (fall) when they are below (above) their long run mean values
- Prices tend to mean revert in this model for two reasons
  - Shocks are mean reverting
  - Equilibrium effects: capacity tends to enter when prices are high (thereby driving down prices) and exit when prices are low



# Empirical Analysis: Continuously Produced Commodities

- These theories difficult to test because the intertemporal linkage greatly complicates modeling
- I've done some empirical analysis to see how well the model works
- Verdict: mixed



# Basic Idea

- 1 day decision horizon: can choose amount of inventory to hold every day
- Two demand shocks, one persistent (integrated or very close), one more transitory (half-life in days)
- Maximizing agents choose carry-out conditional on demand shocks and carry-in
- 1 day forward prices solves PDE:



$$0 = F_t + \mu_z z_t F_z + \mu_y y_t F_y + .5\sigma_z^2 F_{zz} + .5\sigma_y^2 F_{yy} + \rho_{yz} \sigma_z \sigma_y F_{yz}$$

# Solution Methodology

- Make an initial guess for spot price function at each point in a grid of  $z$ ,  $y$ , and carry-in  $x$ . Use this as a boundary condition for the PDE, and solve the PDE
- At each point in the grid, choose carry-out so that spot price equals discounted forward price: if this carry-out is negative (an impossibility) set it to zero. Market can be in daily backwardation in this case
- Wash, rinse, repeat until spot price function converges



# Some Empirical Work

- Pirrong (2010) examines copper market data using a dynamic programming model calibrated to copper market data, and an extended Kalman filter
- Two demand shocks—one has very long persistence, another a short half life
- Can explain some of the dynamics of forward curves and inventory
- Does an OK job of explaining spot price dynamics, medium tenor forward prices, storage; poorer job explaining longer tenor prices, volatilities





# Extension: Stochastic Fundamental Volatility


- Weakest aspect of this model is inability to explain volatility behavior: although fundamental factors generate time varying volatility, fluctuations in volatility aren't big enough
- Introduce stochastic fundamental volatility
- EG, 1 demand shock  $z$ , the variance of which is stochastic, such as:



$$dV_t = \kappa(\theta - V_t)dt + \sigma_v V_t^{.5} dB_t$$


# Same Basic Routine

- Forward price solves a PDE. Make an initial guess for spot price as a function of variance, demand shock, and carry-in. Solve PDE. Choose carry-out to equate spot price and expected forward price.
- Only difference is that forward price PDE is now:


$$0 = F_t - \mu_z F_z + \kappa(\theta - V_t) F_V + .5V_t F_{zz} + .5\sigma_v^2 V_t F_{VV} + \rho_{zV} \sigma_v \sqrt{V_t} F_{zv}$$

# Benefits of This Model

- This model can generate much more realistic dynamics for implied volatilities
- Plausible that economic fundamentals do exhibit stochastic volatility
- Look at the time-varying volatility of the stock market (which reflects macroeconomic volatility)
- Commodity specific factors, e.g., hurricane risk in the US Gulf, military conflict or political risk in oil producing regions



# Preliminary Findings

- This model can provide much more realistic dynamics of commodity price volatility
- It can also explain some otherwise anomalous things, such as simultaneous rises in inventories and prices (which some argue are an indicator of speculative distortion)
- Greater uncertainty makes it more advantageous to hold inventory as a precaution against future shocks
- The only way to accumulate inventory is to encourage production and discourage consumption: higher prices do this
- Therefore, a positive volatility shock should be associated with higher prices *and* inventories



# Seasonal Commodities

- The previous analysis focuses on continuously produced commodities—they are easier to handle and understand
- Seasonal commodities (such as grains) are more challenging



# Forward Curves for Seasonals

- The basic economic insight about intertemporal allocation suggests that prices should be at full carry for delivery dates prior to the next harvest, but “new crop” forward prices are likely to be less than “old crop” prices
- Big slug of new supply due to harvest: should want to consume current inventory before that new supply is available
- Only if expected harvest is expected to be really small should it be optimal to carry inventory to the new crop year; only then should there be a carry between old crop and new crop prices



# Empirical Evidence

- Foregoing provides a pretty good rough characterization of seasonal forward curves
- But there are some anomalies: old crop-old crop prices sometimes fall below full carry
- Aggregation phenomenon that reflects spatial dispersion of stocks and costs of transportation?



# An Anomaly

- The basic story suggests that old crop and new crop prices should not be that highly correlated, and that old crop prices should not respond much to news about the coming harvest
- New crop prices driven by anticipated crop size, but old crop price should not be because little or no storage between old and new crop
- But: (a) old crop and new crop prices are highly correlated, and (b) old crop prices respond almost as much to government crop forecasts as new crop prices





# Non-Competitive Pricing: Manipulation

- The foregoing analysis is predicated on the assumption that the market is competitive
- But there is no guarantee that will be the case
- Indeed, futures and forward markets create the potential for the exercise of market power
- Market power can distort forward curves (and basis relations)



# Corners, Squeezes, Hugs

- A large trader can sometimes accumulate a futures/forward position that is larger than the supply of the commodity in the delivery market at the competitive price
- Additional supplies can be brought into the market, but only by distorting flows, and distortions are costly
- Upward sloping supply curve in the delivery market due to transformation costs



# Exploiting Transformation Costs/Frictions

- By demanding delivery of more than the competitive quantity, a large long can force the market up the supply curve, thereby driving up prices
- The large long can liquidate his remaining positions at this elevated price
- If his position is sufficiently larger than the competitive quantity in the delivery market, the profit on the contracts liquidated at the inflated price is larger than the loss he takes on the units delivered to him, making this profitable



# Burying the Corpse

- A corner works by demanding excessive deliveries
- The cornerer has to dispose of this stuff after the corner is over: it is said that this is the “corpse he has to bury”
- Lose money on burying the corpse



# Price Effects of a Corner

- Artificial demand for the commodity elevates the price of the manipulated contract (and the spot price) until the manipulation ends
- Anticipation of the corpse being dumped on the market depresses forward price for expiries later than cornered contract
- These effects mean that corners can cause backwardations
- Price of deliverable rises relative to prices of related commodities, or of the same commodity at other locations (i.e., basis effects)
- Spot price crashes when corpse is buried



# Quantity Effects of a Corner

- Excessive flows of the commodity to the delivery point before the corner ends, excessive flows away from the delivery point after it ends
- Atypical directions of commodity flows (“water flowing uphill”)
- Inventories in delivery market elevated
- Level of inventories high even though market is in backwardation (or less than full carry): completely contrary to what you’d expect if the market were competitive



# How Do You Know Somebody Wants to Corner?

- Cornerers make transactions that would otherwise be uneconomic
- They take delivery of the commodity at a higher price than they could obtain it in the markets where they actually want to consume it (example to follow)
- Hold inventories of the commodity even though spreads say that is very costly
- “Sharp pencil tests”



# Corner Examples

- Ferruzzi soybeans, 1989
- Sumitomo copper, 1995
- BP propane, 2004





# Implications for Trading

- Don't do it—it's illegal (in the US, certainly)
- Try not to get caught short in a corner
- Commercial intelligence is important: want to know about positions being accumulated in the market, because they might be used to run a corner
- Keep an eye out for anomalous pricing relations, and anomalous commodity flows



# Power Forward Curves

- Power prices exhibit behavior that is not well captured by standard reduced form models
- Structural approaches have some advantages: exploit the transparency of fundamentals
- Structural approaches (relatively) straightforward in fossil fuel dominated markets, more difficult in hydro dominated ones



# Structural Modeling Tools—Non-Storables

- In some respects, non-storables are much easier to deal with; dynamic programming considerations absent
- E.g., Pirrong-Jermakyan (2008)
- For some non-storables, structural models can actually be used for contingent claims valuation
- Market price of risk (market incompleteness)



# Some Challenges

- For markets with strong hydro presence, there are intertemporal linkages that require dynamic programming tools
- Even primarily fossil fuel driven markets are high dimensional
- Oh, if only the whole world were like Texas
- Still, the one eyed man is king in the land of the blind—even a crude structural model is likely to be better than feasible reduced form ones



# Basic Approach: Spot Price Modeling

- Identify major drivers of power prices
  - Load (weather)
  - Fuel prices
  - Outages
- Estimate relationship between power spot prices (e.g., hourly prices) and the major drivers
- Estimate the dynamics of the drivers
- Given driver dynamics and driver-price relationship, can determine price dynamics



# Relation Between Spot Prices and Fundamentals

- Power markets relatively transparent: good data on main drivers and prices (at least in markets with RT markets)
- Econometric approach: use data on prices, loads, fuel prices, and outages (if available) to fit a (nonlinear, perhaps non-parametric) relationship
- Bid stack approach: some markets report (on a delayed basis) the actual generator bids that gives the relationship between loads and prices



# Driver Dynamics

- Load (or temp):
  - Estimate mean loads as a function of time of day, day of year
  - Use econometric methods to estimate the dynamics of deviations between observed and mean loads (mean reversion)
- Fuel:
  - Challenge in multi-fuel markets: use fuel that is usually at the margin
  - Standard models (e.g., GBM)



# Forward Curves

- Given the price-drivers relationship and a specification of the dynamics of the drivers, possible to forecast spot power prices
- But remember, a forward price is NOT a price forecast/expectation
- Forward prices embed a risk adjustment
- Indeed, in power this adjustment is very large





# Market Price of Risk

- Power markets “incomplete”: power price is not an asset, and important drivers (notably load) are not traded either
- Can use relatively advanced quantitative methods to determine the risk premium—the “market price of risk”—associated with non-traded load
- This MPR is very large
- Big deviation between forward prices and expected spot prices



# Sources of the Risk Premium

- In some respects, large risk premium surprising given that the relevant risk is idiosyncratic (e.g., temperature shocks not correlated with asset prices)
- Probably reflects incomplete integration of power markets and broader financial markets
- As markets have become more integrated, premia have declined
- Keynes-type story: “spike-o-phobia”
- Power prices can spike, imposing big losses on shorts: they demand a big premium to sell forward



# Hydro Markets

- In the PNW of the US, and in northern Europe, hydro is a major source of generation
- Although power can't be stored, water can be
- Optimal use of water over time is like a storage problem—need to use dynamic programming methods
- Presence of hydro affects price dynamics: more autocorrelation, slower mean reversion

