Lecture Notes Part 3
Economics of Sustainability
K Foster, CCNY, Spring 2016

Table of Contents

Basics of Oil

Many sustainability students would consider themselves opposed to fossil fuels. Nevertheless it is important to understand your opponent.

I can heartily recommend Jim Hamilton's papers ("Historical Oil Shocks" 2011; "Oil Prices, Exhaustible Resources, and Economic Growth" 2012) as well as the book, *Oil 101*, by Morgan Downey, which is a great non-technical but highly informative read.

There is a myth that oil is made of dinosaurs – please discard this belief, if you want to be taken seriously! Oil **is** from fossilized creatures, but not the charismatic dinosaurs, rather tiny plankton and algae from old seabeds. Most oil is not from ancient fossils but relatively more recent (ie since the dinosaurs went extinct) less than 60m years. That organic material was buried under mud and sank downward, becoming kerogen (sometimes called source rock). As you know, the temperature of the earth gets hotter as you go deeper so the buried material, pushed downward, was cooked. There is a "window" in which oil can be formed – deeper forms natural gas. Much or most then evaporated up through the porous and permeable rock – oil and gas deposits are only found underneath cap rock, often shale or salt, an impermeable shell that prevents these volatile gases and liquids from bubbling up to the surface. Oil is often discussed as being in pools but it is actually in the pores of rock – which must be sufficiently porous (enough holes in it) and permeable (whether the holes are connected).

People noticed that there were springs with 'funny' smells or even tar pits, but this was a curiosity, not important to the economy, until Col. Edwin Drake struck oil in Pennsylvania in 1859. That began our modern era of petroleum fuels. While total US oil production increased steadily, this was a result of new exploration – existing fields were often quickly drained – many individual states hit "peak oil" and declined thereafter but total national production gained as new locations were found and new technologies allowed more effective drilling and extraction.

While oil drilling is sometimes celebrated as the free market in action, it was originally monopolized by Standard Oil (that built Rock Center), then much of the development of oil fields in Texas and Oklahoma was shepherded by strong government policy. Oil fields can be thought of as like a lake of water: a pipe in one place, pumping the liquid out, can drain away the liquid that other property-holders might believe is theirs. Property rights to the underlying oil were not clearly defined. Further, pumping too quickly (as from too many wells, all competing to suck up the oil first) could strand a large fraction of the oil. This is a tragedy of commons, of the type we discussed earlier – same as overfishing. The Texas Railroad Commission (see Hamilton 2011 "Historical Oil Shocks") was formed to regulate and control the extraction, acting as a cartel with federal government support.

While the Suez Crisis of the 1950s left Europe without oil and encouraged more shipments from the Western hemisphere, there was not much of a world market for oil before the 1970s. A confluence of events in the early 1970s, including the ending of the Bretton Woods (gold-based) international payments systems, Nixon's wage and price controls, a peak in US oil production, and finally the OPEC embargo for the 1973 Yom Kippur War, led to the first modern oil crisis. Five years later the revolution in Iran led to another price spike. In the early 1990s there was another war in the Middle East that again spiked the oil price.

These are important for their relation to the US economy, "All but one of the 11 postwar recessions were associated with an increase in the price of oil, the single exception being the recession of 1960. Likewise, all but one of the 12 oil price episodes listed in Table 1 were accompanied by U.S. recessions, the single exception being the 2003 oil price increase associated with the Venezuelan unrest and second Persian Gulf War." (Hamilton 2011)

These graphs from Hamilton (2011) shows the historical oil price:

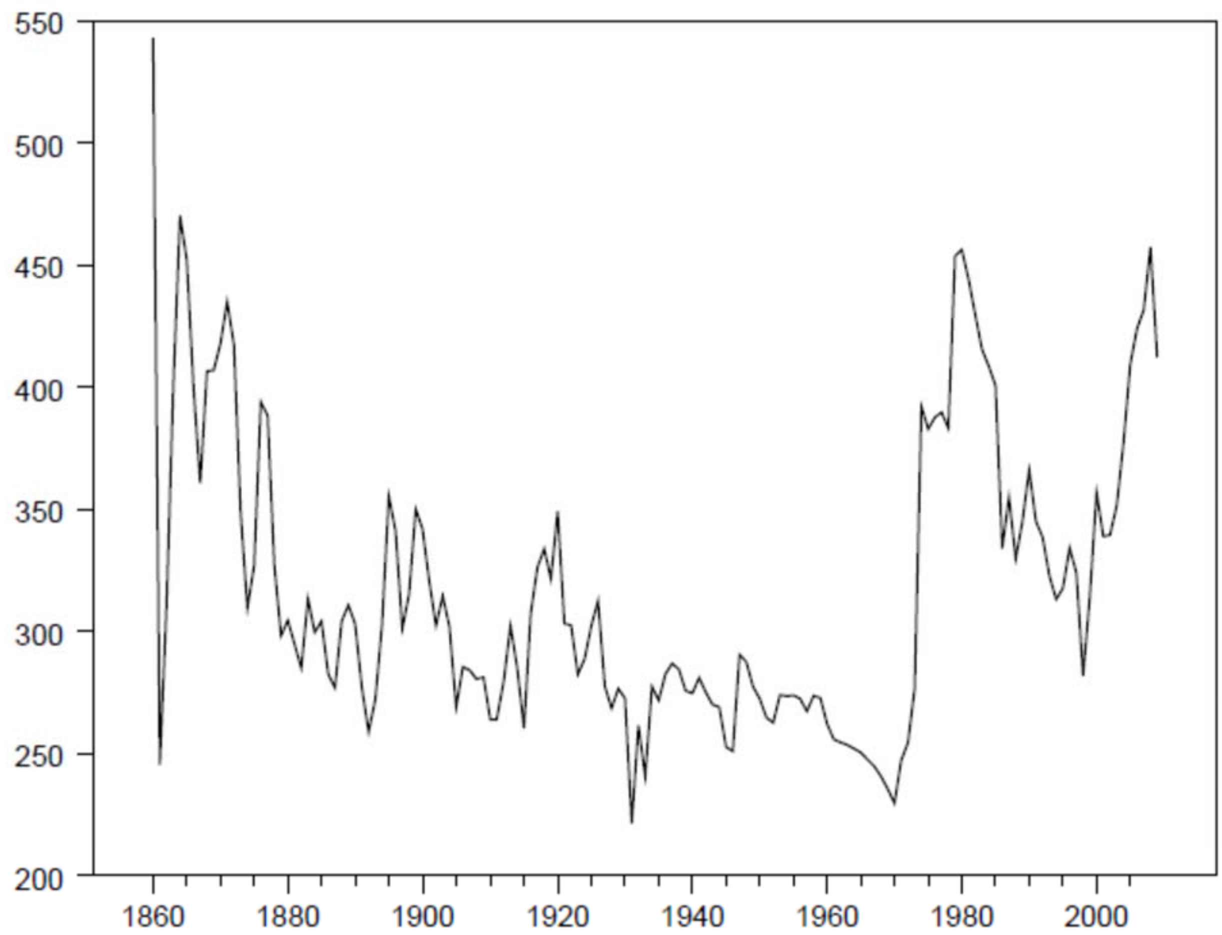

Figure 2. One hundred times the natural logarithm of the real price of oil, 1861-2009, in 2009 U.S. dollars. Data source: *Statistical Review of World Energy 2010,* BP; Jenkins (1985, Table 18); and *Historical Statistics of the United States,* Table E 135-166, Consumer Prices Indexes (BLS), All Items, 1800 to 1970, as detailed in footnote 1.

Figure 14. Price of oil in 2009 dollars, 1973:M1-2010:M10. Price of West Texas Intermediate deflated by CPI.

One common question is about "Peak Oil" – ever since M King Hubbert proposed an estimation in 1956 that production could be modeled as a logistic distribution curve; in 1956 he predicted the US peak of production. Estimates of the global peak are more difficult however, particularly in the face of expanding technology. One problem is that the data is so limited: although publicly traded oil firms such as ExxonMobil must publish their best estimates, most of the world's oil is controlled by national governments (NOCs are National Oil Companies) that treat even basic information as a state secret. Even statistics for the KSA's Ghawar field (the world's largest) produce more heat than light. All the oil majors distinguish between "proven reserves" that very likely could be extracted with current technologies at current prices, and "probable reserves" that are more uncertain. Although these estimates involve a degree of uncertainty, in the US the SEC has jurisdiction over publicly-traded companies. Nevertheless we can be certain that the world will eventually slow down the release of CO_2 into the atmosphere, the question is whether this is done by deliberate policy in response to climate change or by a shortage of oil to burn. (See Hamilton 2012 "Oil Prices, Exhaustible Resources, and Economic Growth.")

While people work very hard to transform oil into a homogenous commodity, it does not start that way – every field is different, sometimes dramatically different. There are global standards such as WTI (West Texas Intermediate) or Brent Blend.

Crude oil is differentiated by a number of factors including density – how heavy the liquid is. The American Petroleum Institute created API density, ranging from zero to 100. Water is at 10°; lower numbers are heavier (the stuff used to pave the roads), up to 100° which is about 60% less dense than water (some could even be lighter than 100). Grades such as WTI, Brent, and even the so-called Arab Light are all intermediate density (in the thirties of API density). Venezuelan oil is heavy with API in the twenties (about 90% of the density of water). Oil sands produce oil about as dense as water. Of course much crude oil from the ground has large amounts of water – usually several times more water than oil is brought to the surface, so the oil must be de-watered.

Crudes also vary by sulfur content – "sweet" refers to oil with low sulfur, and "sour" has high levels. Sulfur is a pollutant and also corrodes equipment so it reduces the value of the oil. There are many other characteristics – oil is as varied as the life that produced it. Companies work hard to transform idiosyncratic substances into a commodity.

Refining is a very complicated process where plant operators look at the grades available (at various prices and delivery times) and figure out which outputs to make. Some refineries have a wide array of technologies to produce many different types. But the heart of refining is simply heating the crude in a tall tower and letting the vapors rise to different condensing trays – the lighter outputs go to the top and the heavier products barely rise. Light products are gases like methane and propane; then gasoline; then kerosene, diesel and heating oil, motor oil, down to grease/wax, bitumen (what roads are paved with), and coke (a solid burned like coal). Since gasoline is more valuable than many of the heavier products, those can be "cracked" into gasoline with heat and pressure. The "refinery gain" shows that they produce a greater volume of output than the heavy dense input going in.

These various types of products are then blended together to suit the market. You might be familiar with octane ratings for a car – higher octane fuels generally have less energy density but can be better compressed without igniting (knocking) so more can be injected, so the engine can be more powerful. Lead is a cheap way of boosting octane.

Oil exploration uses a variety of tools including "thumpers" where sound waves are bounced off deep rock formations.

As for getting the oil out of the ground, you can explore that in more detail by looking at BP's Deepwater Horizon disaster. The National Commission's Report to the President, "Deep Water: The Gulf Oil Disaster and the Future of Offshore Drilling," is here, http://docs.lib.noaa.gov/noaa_documents/NOAA_related_docs/oil_spills/DWH_report-to-president.pdf. Chapter 4 gives most of the detail of the drilling process and what went wrong; really though the entire report is worth a careful read.

Hotelling on Resource Extraction

Hotelling result on resource extraction: for an exhaustible resource, the price ought to grow at a rate equivalent to the market rate of interest, so if p is the price of this resource and r is the rate of interest then $p_t = p_0 e^{rt}$, the price will grow exponentially. Why?

Arbitrage between risk-free investment (getting r) and keeping resource in the ground. Keeping resource in the ground returns $\frac{p_t}{p_0} = 1 + \frac{\Delta p}{p_0}$, the percent increase of its price. Note that if extraction becomes more difficult (diminishing returns) then more investment is required to get the same rate of return so this will eventually become unprofitable, even when there is still some resource available.

Sadly, while the theory is elegant it does not explain markets for things like oil. It might be a better guide for natural resource managers of forests, though.

Prisoner's Dilemma and Cartels

In the past there have been instances when OPEC was able to successfully (from its perspective) raise the price of oil and increase the revenues of its members. Why don't they still do that? To understand their problem, it is useful to consider the "Prisoner's Dilemma" – which seems like a completely different topic at first.

Consider two accused robbers. The police don't have enough evidence to get them on anything more than minor charges (each would get 1 year in jail) but they try to get each prisoner to confess and accuse the other. The police go to prisoner A and tell him that he can get a reduced sentence (just 6 months in minimum security) if he gives them evidence to convict prisoner B (who will get 20 years). They got to prisoner B and make him the same offer. If both confess, each will get 15 years.

What is the likely outcome? Both prisoners are likely to confess. Why? Draw a table of their choices and outcomes.

	A silent	*A confess*
B silent	*A: 1, B: 1*	*A: ½, B: 20*
B confess	*A: 20, B: ½*	*A: 15, B: 15*

The key insight is that, no matter what the other guy does, prisoner A is better off if he confesses. If B stays quiet then A reduces his prison time from 1 year to 6 months; if B confesses then A reduces his prison time from 20 years to 15 years. Same for prisoner B.

You might at first think this requires that the prisoners be in separate cells but this is not required – they can meet ahead of time and strategize, it won't change the outcome. Of course they would lie to each other but they should each realize that they are being lied to.

The key is that, although they would like to both stay silent, they cannot trust the other player to achieve this result (even though it would be optimal for them).

How is this relevant to the behavior of cartels? A cartel has the same basic pattern of choices. If there are 2 players (companies or nations) then each has the choice: restrict production or produce a lot. Restricting production raises prices and profits. But restricting production means not selling and so not getting some revenue – better if the other player restricts production.

Social Entrepreneurship

Social Entrepreneurs identify resources where people only see problems. They view the villagers as the solution, not the passive beneficiary. They begin with the assumption of competence and unleash resources in the communities they are serving. - David Bornstein

Please read Chapters 1, 2, 16, 21 of David Bornstein's book, *How to Change the World: Social Entrepreneurs and the Power of New Ideas*. It's a classic, which served the purpose of bringing wide attention to what was (back in 2007) an emerging field. We'll talk about it in class. Please bring ideas for your own ventures.

Some background on corporate types (note: this is not legal advice)

Used to be for-profit vs non-profit [typically 501(c)]
Although now more middle ground, whether for-profit with strong CSR or triple-bottom-line; or B-corp that explicitly blends purposes

Directors of for-profit companies have a fiduciary responsibility to maximize shareholder value – usually measured in money. Courts give broad discretion to a company to determine exactly what/how maximizes, but this responsibility is always in the background. That's a reason why so much CSR talk from corporations circles back to the allegation that good CSR maximizes (long-run) profits.

So CSR has a strange moral status – it has altruistic stated motivation but redounds to selfish ends.

Setting up a non-profit is not difficult but requires patience with forms – it is a tax-exempt organization and therefore the IRS is interested

In NY State, this is a 4-page document plus a $75 fee
http://www.dos.ny.gov/forms/corporations/1511-f-l.pdf
with instructions here http://www.dos.ny.gov/forms/corporations/1511-f-l_instructions.pdf
mainly the issue is to figure out "Type A" or B or C or D

Now granted the IRS rules are a bit more complicated
http://www.irs.gov/pub/irs-pdf/p557.pdf
but a lot of that is irrelevant to your likely purpose

Benefit corporations (B-corp) have a stated mission in addition to profits – they're for-profit companies but with a stated social mission as well; they must be explicit about what these missions are and annual reports by the company should state how they're doing (see https://www.bcorporation.net/) eg Etsy, Patagonia, King Arthur Flour, Warby Parker, Greyston Bakery …

Unsurprisingly these are mostly consumer-facing companies where the marketing edge makes sense even from a relatively narrow profit-maximizing view

These types of mission-driven organizations have a few distinct audiences:

- Customers, who have ethical desires to support the mission
- Employees, whose ethical desires propel them to work for a particular organization
- Investors, who support particular missions
- Regulators or other government bodies

Other goals can be seen as part of these – eg broader concerns about reputation are ultimately about customer's views or employees or somebody

Stepping back, the basic idea is that every organization wants to be seen as being good – but also has incentives to backslide or cheapen this "good". If cheap talk can sufficiently sway profits, then every organization will talk cheaply and eventually the value will be eroded (eg products labelled "natural"; greenwashing). So organization wants a commitment mechanism, some way to be held accountable. This might be an external source (Angie's List? Rating agency?) or certification (LEED). Or some explicit easily-measured target (5% of sales donated).

General question: what are consumers purchasing? Shifting definitions of organic; fair trade; carbon neutral; no child labor?

But! Recent news on conflict minerals, "No party has suggested that the conflict minerals rule is related to preventing consumer deception. In the district court the Commission admitted that it was not." "By compelling an issuer to confess blood on its hands, the statute interferes with that exercise of the freedom of speech under the First Amendment." [US Court of Appeals, DC, Apr 14, 2014, Assoc of Manuf vs SEC] see http://www.bloombergview.com/articles/2014-04-14/the-first-amendment-lets-companies-keep-quiet-about-blood-diamonds

Fracking

Basic Framework: distinguish costs and benefits at variety of scales and internal/external incidence

Background: fracking depends on two related technologies: horizontal drilling where one vertical well can branch out along seams of source rock (so cheaper to access deeper rock) and fracturing where water under high pressure is pumped into the wells to open cracks for the gas to better flow; it is mixed with sand to keep these fractures open. (Recall ideal source rock needs lots of holes that are connected to one another; fracking opens connections.) Neither of these had a revolutionary 'aha' moment, rather gradual improvements in technique, usually acquired through experience, made them more useful and able to exploit resources that had been previously unobtainable.

Shale gas fracking implies modest increases in local employment and income (Weber 2012, Brown 2014).

There is an increase in wealth to local landowners or whoever owns the property rights; there are separate issues of estimating the increase in income/consumption expenditure after wealth increases. The wealth increases might not incur to local residents.

From basic theory, distinguish between tradables and non-tradables; an increase in local income means a rise in demand for both. The price effect on most tradables is likely small – and gains go to those other regions. Non-tradables have price increases depending on elasticity of supply (non-tradables include property, housing, services). Transportation services are important and transport infrastructure.

The extent to which people migrate (in and out) makes it more difficult to determine which populations are affected (Gamper Rabindran & Timmins 2011).

Can cause "Dutch Disease" as local workers become expensive so local exporters suffer. Mixed evidence for natural "resource curse" where more resources actually hurt long-term growth – this might depend on quality of governance.

Income/jobs from gas can be separated as initial (well-drilling) and production. Likely that drilling jobs are specialized (so income benefit might be external to local area) and transitory.

Klaiber and Gopalakrishnan (2012) find that properties close to drilling are negatively affected possibly via water safety concerns or other neighborhood effects. They estimate a hedonic regression on property characterisics including agricultural land – which is easily developed for drilling.

How close is too close?

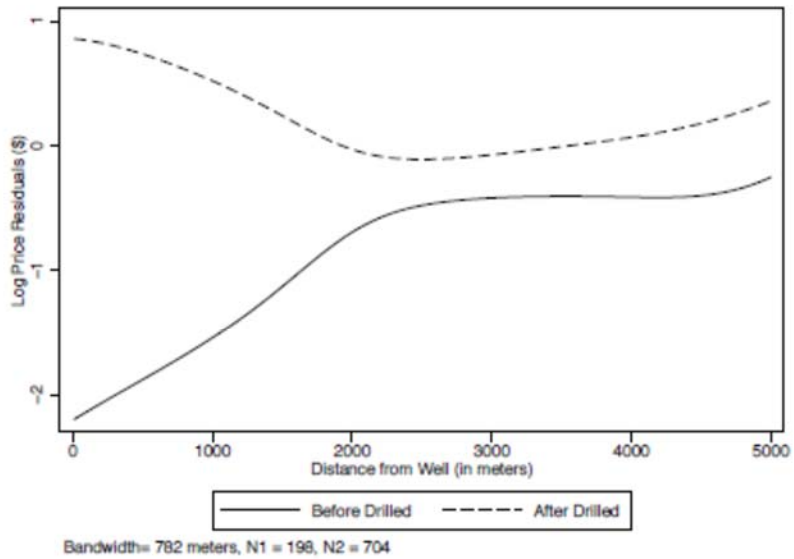

Figure 1: Sales Price Gradient of From Local Polynomial Regressions on Distance from Current/Future Well

Muehlenbachs, Spiller, and Timmons (2014) find positive effects on property values unless it has well water. Their graph (labeled Figure 1 in above) shows the price differences by distance, whether before the well was drilled (when there were negative impacts on the price) or after (when there were positive impacts). Their estimation suggests that local impacts of about a mile (2000m).

One square mile property (640 acres) can get a signing bonus of about $1.5m (Hefley et al 2011). Then royalty payments if gas is struck. Much of the variation in cost depends on how many separate parcels of property are involved.

One of the chief externalities is from the water used since in many areas of fracking there are water shortages. Recycling the water is difficult since the source rocks contain pollutants such as metals and the fracking liquid has additional chemicals added. There are reports of widespread contamination of drinking water wells from methane, ethane and propane, when gas wells are drilled nearby.

There are also air quality externalities both from trucks (often diesel particulates) and from the wells themselves (emissions of methane and polycyclic aromatic hydrocarbons).

However the net air pollution effects are uncertain: an area that switches from coal-fired electrical generation to natural gas could see overall air pollution reduced while particular sites near the wells would increase. (See Holladay and LaRiviere 2013)

Water pollution can be partly avoided by private provision such as bottled water (eg Graff Zivin, Neidell, Schlenker 2011) although at increased cost. Avoiding air pollution is more difficult.

There are also externalities on local roads as heavy trucks mean more deterioration; some wells pay explicit road bond fees.

Rising income and property values may have benefits to communities through induced expenditures and via increased tax revenue.

The exact characterization of how increased population/income affects locals is difficult to determine. If many public goods are not at capacity then increase government revenue without an increase in expenditure can be good for most residents. (What is the production function of government services?) This question arises in other contexts, e.g. how do immigrants affect local economies.

Other externalities on local industries such as tourism if there had been a draw to "unspoiled nature".

There appears to be external harm to infants born near fracking areas (Hill 2013), however the existence of these effects is debated particularly since the mechanism is unclear (since it is not mediated by well/public water supply). Hill compares mothers living near gas wells that were being drilled with mothers living near gas wells that were not yet drilled (so otherwise comparable), finding increased incidence of low birth weight and lower APGAR scores, although no change in premature birth. With reasonable assumptions about public health costs of LBW, could be $250m of additional costs from drilling, which would be allocated as less than 10% of the private cost of drilling a well.

How do we find overall benefits of shale gas drilling? The increased supply resulted in lower prices paid by consumers, which is an increase in consumer surplus.

What is effect of price changes on fracking activity? Need to consider marginal vs average cost difference – once a well has been drilled, production is cheap so most of the margin is new exploration and drilling (so change in rig count leads change in production).

Bibliography

Brown, J P (2014). Production of Natural Gas from Shale in Local Economies: A Resource Blessing or Curse? Economic Review of Federal Reserve Bank of Kansas City.

Gamper-Rabindran, S., & Timmins, C. (2011). Hazardous waste cleanup, neighborhood gentrification, and environmental justice: Evidence from restricted access census block data. The American Economic Review, 101(3), 620-624.

Gopalakrishnan, S., & Klaiber, H. A. (in press). Is the shale energy boom a bust for nearby residents? evidence from housing values in Pennsylvania. American Journal of Agricultural Economics, aat065.

Graff Zivin, J., Neidell, M., & Schlenker, W. (2011). Water quality violations and avoidance behavior: Evidence from bottled water consumption. American Economic Review: Papers and Proceedings, 101, 448-453.

Hefley, W, S Seydor, et al. (2011). The Economic Impact of the Value Chain of a Marcellus Shale Well. University of Pittsburgh Katz Graduate School of Business.

Hill, E L (2013). Shale Gas Development and Infant Health: Evidence from Pennsylvania. Working Paper, Cornell University.

Holladay, J S, & J LaRiviere (2013). The Effect of Abundant Natural Gas on Air Pollution from Electricity Production. Working Paper, Univ of Tennessee.

Klaiber, H A, & S Gopalakrishnan (2012). The Impact of Shale Exploration on Housing Values in Pennsylvania. Presented at Agricultural and Applied Economics Association meeting.

Marchand, J. (2012). Local labor market impacts of energy boom-bust-boom in western Canada. Journal of Urban Economics, 71(1), 165–174.

Muehlenbachs, L, E Spiller & C Timmins (2014). The Housing Market Impacts of Shale Gas Development. National Bureau of Economic Research, Working Paper 19796.

Weber, J. G. (2012). The effects of a natural gas boom on employment and income in Colorado, Texas, and Wyoming. Energy Economics, 34(5), pp. 1580–1588.

Financial Markets

Financial markets for commodities such as oil can reveal important information about expected future prices. These also help understand impacts of price changes throughout the global economy.

These notes are based on John C Hull, *Options, Futures, and Other Derivatives* and Frederic S. Mishkin, *The Economics of Money, Banking, and Financial Markets*.

To understand the place and function of commodity markets, we need to start with a bit of perspective on financial markets overall.

Financial Markets and Securities

Financial markets are intentionally creating commodities, where these commodities are little pieces of paper representing legal claims to a particular commodity or cash flow (securities such as stocks and bonds). Then there are other derivative securities, with a value that depends on (is derived from) the value of another security (which could in turn be a derivative). These securities can be for current transactions (spot markets) or transactions happening in the future (forward contracts). Some securities are contingent, only paying out if some event occurs.

Financial Markets trade money. If you only think of money as currency then trading money seems odd. But money is much broader. Financial markets allow me to trade money **through time** – if I have money now, I can turn it into money later; six months from now or 30 years from now. Depositors have too much money now; those who take loans don't have enough money now but will have money later. Financial markets trade money **from different countries**; euro to dollars to renminbi or many other transactions. Financial markets also trade money **in different states of the world** (throughout the multiverse), depending if different events occur – contingent claims. Insurance pays only if some event occurs (if a person dies, the life insurance company pays money). Stocks pay only if the company makes a profit. Options pay only if a stock or other security has a value in some range.

Financial markets promote the efficient use of scarce capital by ensuring that firms with the most productive possibilities get investments. The broader the scope of individuals and firms that can get loans, the more likely it is that loans will be based on the merit of the project to be funded. Of course stupid ideas can still get funding, nobody is perfect! But this is the ideal. The problem is avoiding "dead capital" – saving hidden away, in case I need it in the future, but where nobody else can access it to invest (like gold buried in the garden).

Before there were organized markets, loans were made in individual transactions (peer-to-peer). But markets bring several advantages: standardized structures bring lower transactions costs (much of this through the legal accretion of case law) and a sharing of risks over more participants. The standardization brings enormous savings in solving some of the essential problems of giving money to someone else, whose incentives and information are quite different from my own.

Consider the problems that rich people would have without capital markets (in history, this is what was done). They want to make interest on their savings by loaning it out. But they would have to evaluate each borrower on their own and then monitor each borrower. Each lender would have to pay a lawyer to write up their loan contracts. In a large market there are economies of scale as well as liquidity services – if lenders get together then if one needs liquidity then she can get it.

This also makes possible risk sharing: if there are many lenders then losses can be spread over more of them. Eg: if there are 100 people each making a loan of $1000 then if 3% of loans are bad (no repayment) then any single lender has the possibility of a large loss. But if these 100 people get together and form a bank that makes 100 loans of $1000 each, then if 3% are bad then the bank loses $3000, so each person loses just $30 out of their $1000. Assuming they are risk averse, this outcome is much better! This may also give diversification, since an investor can avoid putting all of her wealth into one large project but instead put a little bit into many projects in different industries, to get a diversified portfolio.

So finance combines with the limited-liability company to spread risks.

Financial intermediaries also have a wider selection of tools to deal with information asymmetries (where one side of a deal knows more than the other), adverse selection (asymmetric information before the transactions), and moral hazard (asymmetric information after the transaction). Large financial institutions can screen out bad credit risks more effectively and monitor better.

Commodity Markets

People have been trading commodities for centuries; there were contracts involving future transactions in rice in Japan in the 1600s. Agricultural commodities are most common but there are contacts for many raw materials such as metals and energy, weather, real estate, and financial commodities such as foreign exchange, interest rates, credit ratings/defaults, and stock indexes (other indexes), even volatility on these. There are quantos, which are markets for things that don't really exist.

Both producers and consumers want the widest possible markets: buyers want to find the lowest possible price and sellers want to find the highest possible price. A central commodity market makes it easy to shop around. In NYC there are many markets where both buyers and sellers cluster such as the diamond district or fashion district. Commodity exchanges provide buyers and sellers access to a large number of counterparties. They also have lots of people working in the middle.

The ability to buy and sell in the future helps firms on both sides. A well or mine can sell its output in advance, so that it will not be exposed to risks of its output falling in price. A manufacturer, that uses some input, can buy this input in advance and lock in the price, so that it will not be exposed to the risk of that input rising in price. For example, an airline might buy its fuel in advance so that it knows, when selling tickets to customers, that it will make a particular profit that is not put at risk by fluctuations in the price of jet fuel.

Most financial markets have hedgers (participants using the markets to reduce risk), speculators (participants looking to make bets) and arbitrageurs (market makers). (Details below.)

Commodities are traded in carefully defined and structured contracts. Crude oil contracts can be traded on CME (Chicago Mercantile Exchange) for delivery in any month for the next five year, and every six months for a decade more. Each contract is for 1000 barrels (a "mini" is 500 barrels), with minimum movements of 1 cent per barrel (so the contract can move up or down by $10). That's a lot of oil! More than most backyard swimming pools. Average price in 2011 was $94.88 per barrel so a single contract would be worth $94,880.

Useful webinar from the CME on the Fundamentals of Energy Trading,
http://www.cmegroup.com/education/interactive/webinars-archived/fundamentals-of-energy-trading.html
Sundaram has a great explainer here
http://people.stern.nyu.edu/rsundara/papers/RangarajanSundaramFinal.pdf

Futures and Options

Many commodities and financial instruments can be either exchanged now or a contract can be arranged for a future transaction.

Definitions:
Derivative: value depends on another variable; examples options (calls, puts, swaps, etc)
Forward contract: agree to buy/sell a particular asset at given price and date/time.
Spot contract: agree to buy/sell a particular asset at given price NOW.
Forward vs **Future** contracts: futures are traded on an exchange in regulated sizes

Valuation:
$\mathbf{S_t}$ is the value of some asset at time t (so S_0 is value at time zero, the beginning; S_T is time at time T, often expiration date)
$\mathbf{F_t}$ is the value of some forward price at time t (the price at time t, of some asset to be delivered at date T>t; sometimes for clarity denoted F(t,T) the value at date t, of an asset to be delivered at T.

For now we just take these prices as given: some trader or exchange tells us what the price is. (This is equivalent to saying that the financial markets are perfectly competitive so that our position will not affect the market price – we don't have market power.) Later we will ask what the prices, theoretically, ought to be. But first we have to understand the details of how portfolios can be constructed when the prices are just taken as given.

Long on forward: agree to buy at future date
Short on forward: agree to sell at future date

Exchanges vs **Over-the-Counter** (**OTC**): BIS reports that OTC derivatives have notional value of about $640 trillion while exchange-traded contracts were about $25 trillion. Commodities are smaller, just under $3 trillion, with gold about $0.5 trillion and metals $0.1 trillion. (http://www.bis.org/statistics/derstats.htm)

The particular price at which the option or forward trade will take place is **Exercise** Price or **Strike** Price (**K**).

The particular date by which the option or forward must be exercised is **Expiration Date** or **Expiry** or **Maturity**.

Payoff to Long position on a forward = $S_T - K$
Payoff to short position on a forward = - payoff to long position = $K - S_T$

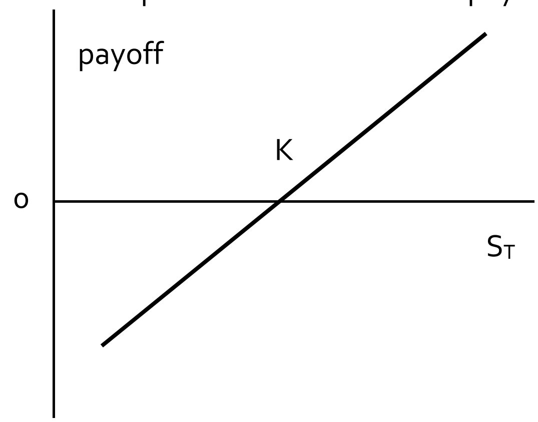

Spot and Forward Prices

How can we develop a relationship between the current price of an asset (spot) and its future value? First we have to think about how these two prices are set. Clearly there is a relationship between them, but what?

Consider if the spot price were 100 and the forward price, for delivery in a year, were 110. Would you rather buy it now for 100 or spend a little more to lock in the price?

There are a couple things that we immediately notice we're missing. First, since we're comparing money in two different times, we need to worry about the relative values of these dollars – i.e. the interest rate, which gives the price of next year's dollars. We also need to know something about how/if the value of the underlying asset changes – if we're buying ripe tomatoes then they'll go bad long before a year is up; if we're buying oil then we have to store it somewhere; if we're buying stock shares they pay dividends.

Interest rate: assume the rate is given as "**r**" and that we're working in continuous time so the present value of each dollar, paid in a year's time, is $\mathbf{e^{-rT}}$, where T=1 so it is $\mathbf{e^{-r}}$.

In the example above, where spot is 100 and forward is 110, if the interest rate is low then we could borrow money today to buy at spot, sell it at the forward price, make $10 per transaction and if the $100 borrowed costs, say, $3 or $4, then that's a nice profit from the arbitrage. On the other hand, if the interest rate were very high then the opposite transaction would be more worthwhile. If I have $100 I could put it in the bank and get more than $110 after a year. Sell short at the spot rate (100) and buy forward at 110 to lock in the price at which I return the underlying asset. (Like Hotelling result for resources.) The difference (how much more I earn from interest over the 110 forward price) is arbitrage profit. In either case, the arbitrage trades work to change demand and supply to bring the prices back into line.

We might be confused because we might think that the forward price is a predictor of the price that will be set at that future date. But it's not – the spot price is a predictor. Why? Again, we consider what actions might be taken by a smart financial trader. Suppose that it is known that, on Friday, the price of some asset will jump from 50 to 75. Clearly, someone who holds the asset on Friday will get a huge return on their money. So what is likely to be the demand for that asset on Thursday? Wednesday? Tuesday? Today? The argument gets more complicated if the asset is difficult to store or if it changes value when held. (Under some circumstances, the futures price can be an unbiased estimate of the expected future spot price, but we still worry about the interest rate.) But the core arbitrage argument is clear.

These examples assume that the value of the asset does not change much over the time period. So we differentiate between an investment asset and a consumption asset. This tells if large numbers of market participants will be able to arbitrage (as outlined above) or whether large numbers will be eating what they buy. (I wouldn't sell a forward for a pint of Ben & Jerry's because I'd eat it and wouldn't have anything to deliver at the end of the contract!)

In some way we can think of putting money in the bank as buying money forward: if I put $100 in the bank and get (without risk) some return, r, so that after a time of T, I get $100e^{rT}$. This is like buying forward $100e^{rT}$ at a price of $100. Any other forward contract can be thought of as delivering in some different units of measure – but still, in the end, I should get the same rate of return. Whether I buy forward 1 gallon of crude oil, or some equivalent number of liters, doesn't matter. Similarly it doesn't matter whether I buy forward dollars or yen or euro or hog bellies or gasoline or S&P index contracts….

All of these arbitrage arguments get us to our first equation: $\mathbf{F_0 = S_0e^{rT}}$, today's forward price is the future value of today's spot price. This is strictly true for investment assets in markets where arbitrageurs can borrow and lend at the same riskless rate, there are no transactions costs or other taxes, and there are enough (potential) arbitrageurs. You can think of it as just offering one more way to invest – you could get the riskless rate on the money or buy an asset that would (again, risklessly) provide some payment in a year.

Of course some stocks have a known income or known dividend yield, so we can modify the equation to take account of these complications. Other assets have storage costs (negative known income) or convenience yields. The convenience yield is defined as the amount that we observe that market participants are willing to forfeit in order to have the actual physical asset rather than a futures contract.

If we generalize about the "cost of carrying" some asset forward, whether that is the interest rate to finance it, or the interest rate less the income actually earned, or the interest rate less the foreign interest rate, or interest rate plus storage cost, denote the "cost of carry" as c so that for investment assets,

$$F_0 = S_0 e^{cT},$$

while for consumption assets, where y is the convenience yield,

$$F_0 = S_0 e^{(c-y)T}.$$

Also there are many contracts that offer interest held in different currencies – again the same arbitrage arguments should hold. If I can risklessly get some r interest rate in US dollars then I should be able to lock in an equivalent rate in euro or yen or any other major currency. If I have a unit of foreign currency (FX) then I can either buy dollars at S_0 and invest in the US to get S_0e^{rT} at the end of T time, or I could invest the FX at the foreign rate to get $e^{r_{FX}T}$ and buy forward at F_0 to end up with $F_0 e^{r_{FX}T}$. Set these two end possibilities equal, $S_0e^{rT} = F_0 e^{r_{FX}T}$ or $F_0 = S_0 e^{(r-r_{FX})T}$.

Next we move to valuing these forward contracts. The forward price is $\mathbf{F_0}$ but the value of the futures contract (agreeing to buy at that forward price) is **f**. This sounds confusing but it is the simple result of the distinction between the value of a contract and its notional price – for example you could buy insurance that will pay $25,000 if you die – but you don't pay $25,000 for it! Of course a forward contract is not probabilistic – the whole point is that there is an ironclad agreement to trade at F_0.

Suppose for some asset the spot price is 100 and forward price is 110. If I enter into a forward contract that sets a strike price (denote it as K) of 110 then the value of this contract, f, is exactly zero. Tomorrow is a new day so the prices will change (but not K – that's written into the contract) and $\mathbf{f = (F_0 - K)e^{-rT}}$.

You might be asking why anyone would enter into a contract where the value of it is zero. This is what arbitrage means – that although everyone is trying to make money, on net the prices must give no arbitrage profit. A significant fraction of the parties buying and selling are hedging: they're not looking for arbitrage profit but rather to lock in some price. (Also, leverage.) Later on, as the forward price changes away from the strike price, the value of the forward contract will change – that's the point.

Finally we discuss the relation between the current futures price (F_0) and the expected future spot price ($E(S_T)$). An investor might have some expected future spot price that is different from the market, so she could put the necessary cash in the bank today (cost F_0e^{-rT}) and expect to get S_T. However this expected rate in the future should be discounted by the investor's required rate of return given the risk (systematic or non-systematic) that she is taking on. But speculators would enter the market until $F_0 = E(S_T)e^{(r-k)T}$. If the asset risk is uncorrelated with the stock market, then r=k and $F_0 = E(S_T)$. When the futures price is below the expected future spot price (so k>r) this is called "normal backwardation"; when the futures price is above the expected future spot price (so k<r) then it is "contango". (There are a number of linguistic theories about where that word comes from.)

Options

Call Option: the right (but not the obligation) to **BUY** a particular asset on or by a particular date at a particular price

Put Option: the right (but not the obligation) to **SELL** a particular asset on or by a particular date at a particular price

The asset, from which the option value is derived, is the **underlying asset** or underlier.

American vs **European** vs **Asian** options: American options can be exercised at any date up to the expiry; European options can only be exercised on the date; Asian options are exercised on the date but payoff depends on average price.

Puts and calls can be bought when they are "**in-the-money**," "**out-of-the-money**," or "**at-the-money**" (ATM). In the money means that the option would have value today given the current trading price; out of the money means it would have zero value if the expiration date were to be right now; at the money means that the strike price is just equal to the current price of the underlying asset. (so, for a call, $S_T>K$ is in the money, $S_T=K$ is at the money, $S_T<K$ is out of the money; for a put $S_T<K$ is in the money, $S_T>K$ is out of the money, $S_T=K$ is at the money).

If you think of these as being like insurance, then buying insurance against a flood, when the sun is shining, is buying out of the money. If you wait until the storm is washing up to your home, then you're buying in the money. (Some people buy a AAA membership from their broken-down car on the side of the road to get free towing.)

payoff to **European call** = $\max\{0, S_T - K\}$

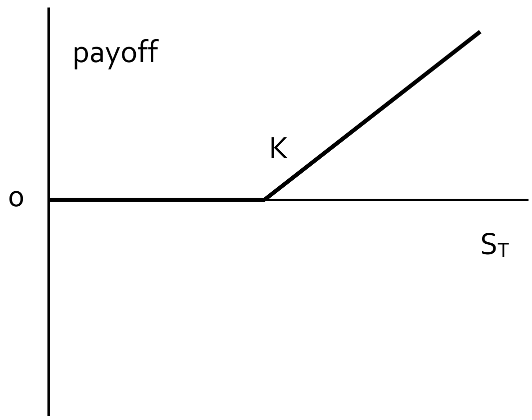

payoff to **European put** = max{o, K - S_T }

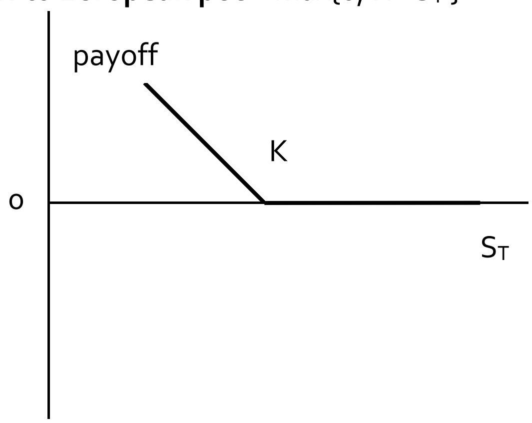

Positions **Closed Out** not usually delivered
As delivery date nears, futures price should **converge** to spot price
Futures contracts have **daily settlement** (so cash flows)

- factors affecting option prices: strike, time, volatility, interest rate ("the Greeks")
- swaps are based on interest rates, where parties swap rates, usually fixed for floating (often pegged to LIBOR, which had recent scandals that you might have read about)

Types of Trading

Traders are generally classified as:

- hedgers – reduce risk of other positions
- speculators – bet on market movements
- arbitrageurs – make multiple positions and profit from the gaps

Hedge

A hedge is basically locking in cash flows at an early date before the asset changes hands. Consider a position, S_t, that will have value S_T at some future date, T. If that asset is hedged then a forward is sold at time T and bought at time t, so that the net asset position is $S_T - F_T + F_t$. We expect that, by the time of expiration, the spot and futures price will be equal (or else there would be arbitrage opportunities) so we expect that, by date T, $S_T - F_T = 0$. So the hedge is exchanging a volatile price (S_t becoming S_T) for a known price, F_t.

Short Hedge: own an asset and short a forward to sell at a pre-specified price.

examples: gold mines might sell the gold, that's still in the ground, at pre-determined prices at some date in the future to "lock in" a profit; farmer can sell the crop forward; exporter with short-term receivables might pre-sell (sell forward) to lock in a profit rate. Hedge can be considered by comparing the money lost on the asset position with the money gained from the offsetting hedge.

Consider an insurer selling annuities in Japan that doesn't want the business affected by FX fluctuations so it could sell forward contracts for 3, 6, 9, 12 months (based on expected sales over the year). If these revenues are to be invested in, say, US Treasury securities, then these securities can be bought forward as well. For instance, if the insurer above is getting ¥100,000,000 in 3 months. If the spot rate is 112¥/$ then this is worth $892,857.14. If the rate increases to 122¥/$ then this is worth only $819,672.13. The movement of ¥10 in the FX rate meant a loss of $73,185.01 on the asset position. If the forward price is also 112¥/$ then selling ¥112 forward (getting one dollar delivered in 3 months) would mean that, if the yen increased to 122 per dollar then the short forward position would mean that the company could sell ¥112 for $1 and still have ¥10 left over to buy dollars (0.0819 worth). This 8-penny gain is small compared to the $73,185.01 loss – but the company could sell more than ¥112. How many 112¥/$ contracts? 73,185.01/0.0819 = 892,857.14 worth (which is exactly the number we discovered earlier).

This might seem like the long way to go about it but it is worth showing the basic method: one position loses a certain amount; it can be hedged if I can find some other position that would gain that same amount. Most companies use hedges with much more complicated structures, but the basic idea remains: construct two offsetting positions so that, as one loses the other gains (and vice versa).

Long Hedge: will buy an asset in the future and buy a forward to but at a pre-specified price.

examples: manufacturers that use mining products (gold, copper, etc) or plastics can buy in advance and lock-in their costs. Southwest Airlines made huge profits, compared with some of their competitors, when they bought fuel forward for the first half of 2005 before the oil price rose so drastically. Their competitors had to pay higher prices while Southwest reaped the profits. (Their competitors, of course, noted that had fuel prices fallen, then Southwest would have been paying extra for unnecessary insurance.)

Can work an example in reverse (as above): how to hedge a short position today with a long forward.

Plenty of individuals hedge, even though they might not realize it. Property owners choosing mortgages must choose between fixed-rate (where the interest rate paid is constant for the life of the loan) or variable or varieties in-between (sometimes a rate is fixed for a few years at first and then varies more often). A business, that employs a person at a fixed salary even though the employee's productivity might vary, is, in some way, hedging (eg contracts for a star sports player). Most insurance companies pass along risks through re-insurance, which are then shared among a wide net of different financial companies.

Why do so many companies hedge? Wouldn't their investors want exposure to certain risks? For instance investors might buy shares in both ExxonMobil (that does well when oil prices rise) as well as GM (which does worse as oil prices rise). If both companies hedge their positions, then that risk-diversification is lost. An investor would have to buy shares in the counter-parties. Insurers take a hit from hurricanes (like Katrina) but many pass along the risks as they hedge their positions (there are catastrophe bonds that are linked to occurrences of natural disasters).

But the reality is that many companies forecast their earnings and their share price falls when they don't meet expectations; it's difficult to communicate the many sources of risk that might be faced by a global company with revenues in many different currencies and costs paid for many different goods. Even internally, a company might want to sort out whether a particular division made money by luck (a favorable FX move) or skill (even after hedging they still out-performed). A hedge means that the company can set its benchmarks and make profits only in its particular areas of comparative advantage. Return to the example of the gold mine: by selling forward they commit that they will make profit if they are efficient at extracting gold; they will lose money if they are not efficient at that. Random fluctuations in gold prices will not drive their results; their profits instead come only from their own efficiency. Competitive pressures can also be important and so every industry (even every firm) must make decisions based on their own particular needs. Finally, these hedges might allow a company to spread the risk more broadly to willing investors. An individual company might hedge in order to pass the risk on to the global financial markets. Instead of a small number of companies losing a lot of money, a large number of investors around the world can each lose a small amount.

Most hedging is not perfect – the real world is messier. **Basis** = $S_t - F_t$. If the asset that is held and the futures contract on the market are the same then at expiration the Basis should be zero. So if the position is opened at time t and closed out at time T, then we would like $S_T=F_T$. Define $b_t = S_t - F_t$ and $b_T = S_T - F_T$. Then if the company has assets S_t, it could chose not to hedge, in which case it would have S_T at the end of the period. If it hedges then it would still get the return of S_T at the end but would then accrue profits to the forward positions, $F_t - F_T$, so the net position would be *[this is the same formula as at the beginning, just with a slightly different interpretation]*

$$S_T + F_t - F_T = F_t + b_T.$$

If it is a perfect hedge then the basis is zero at time T and the value is known at time t; if the basis is not zero then there is residual risk – **basis risk**. This is generally common when the asset position is not one of the standard contracts traded on exchanges.

Commodities markets give many examples. If I am a local heating oil company then I can hedge some of my risk by buying heating oil futures on NYMEX but these are for delivery in New York harbor. I diversify much of the risk of oil price changes but still very local events (e.g. any supply disruption between NY harbor and my customer's oil tank) can impact my results. Crack spreads are similar. Oil producers might hedge using a common blend (Brent or WTI) even though their own oil pumped has different characteristics.

If we consider a distinction between S_t (the asset held) and S_t* (the asset that is traded on a market) then the position will be $S_T – F_T + F_t$. If we add and subtract S_T* then we can rearrange that to get $F_t + (S_T* - F_T) + (S_T – S_T*)$ – the first term in parentheses is the basis risk between the "ideal" asset and its forward price; the second term in parentheses is the basis from the difference in assets. For instance, a bank or insurance company might want to hedge positions, where customers are guaranteed some rate of return, using mixtures of Treasuries and private debt.

In oil markets there are common products to manage the basis risk such as the crack spread. As the EIA explains, "One type of crack spread contract bundles the purchase of three crude oil futures (30,000 barrels) with the sale a month later of two unleaded gasoline futures (20,000 barrels) and one heating oil future (10,000 barrels). The 3-2-1 ratio approximates the real-world ratio of refinery output—2 barrels of unleaded gasoline and 1 barrel of heating oil from 3 barrels of crude oil." (http://www.eia.gov/oiaf/servicerpt/derivative/chapter3.html)

There is the further complication of when the forward should mature. Often traders do not want a forward that expires at the same time – they are planning on closing out the position for cash and don't want to be bothered with actual delivery! So they choose a contract that matures as short a time afterward as is possible. (After, since they don't want to accidentally take delivery.) Since short-term markets have the greatest liquidity, someone hedging a large position might use a series of short contracts (again this does not deliver complete hedging).

Cross Hedging is used if there is no contract traded forward that is exactly what the firm desires. High-paid professionals exert a great deal of ingenuity to figure out how to hedge various positions that their companies enter. The **hedge ratio, h,** is the number of forwards that must be bought per unit of the asset. If the asset and forward are the same thing, then the hedge ratio is one. But generally it will be different.

A **stock index** is often used as a hedge. Since they weight by market capitalization, however, a hedge can gradually erode as weights change slightly. A stock bubble can lead to distortions.

A stock or portfolio can be described with a **β (Beta)** – it measures how sensitive the stock or portfolio is relative to movements in the whole market; in the simplest case it can be found by regression of the excess returns on the stock (over the risk-free rate) upon the excess market returns (again, over the risk-free rate). A stock with a high beta will track the market closely; a stock with a low or near-zero beta will be uncorrelated with the market. In general, $h* = \beta$.

A hedge can change the beta of a portfolio as well, so a portfolio might be incompletely hedged in order to take on more risk or shed risk (as the portfolio manager desires). If the original portfolio has β, and desires to change to β', then instead of taking a hedge ratio $h = \beta$, should take either $h = (\beta - \beta')$ {if $\beta > \beta'$} or $h = (\beta' - \beta)$ {if $\beta < \beta'$}.

A fully-hedged portfolio in the stock market will grow at the risk-free rate. (Since a fully-hedged portfolio is riskless, this makes sense – two riskless assets should have the same return.) Why hedge, then? This allows the company to earn returns entirely from its ability to pick stocks, for example: a company with a meticulously-chosen portfolio that is fully hedged against an index will earn the risk-free rate plus the differential return accruing to its stock-picking skill. Many other companies are just hedging their exposure to other asset baskets and want to minimize their exposure to market risks.

Note that one person's hedge is sometimes another person's speculation. Hedge funds were originally set up to take positions that were well hedged (thus the name) but gradually moved into assets where the basis risk got larger and larger, until they were essentially speculating.