How options *should* be valued

I liked our present method of valuing options pretty well right up until the moment when it was finally time to put some of my own hard-earned money at risk by selling one. I could still understand the logic of the mathematical argument, but somehow the idea of selling an uncovered call, or even an incompletely covered call, myself, just seemed crazy to me. There are ideas we can advocate in public without looking foolish, and then there are ideas we would really act on in private, even if nobody else could see what we were doing. Somehow it suddenly seemed to me, then, that Black-Scholes, like so much of the rest of modern finance theory, actually belonged in the first category but not the second.

What I want to talk about here is what I would really do instead of just using Black-Scholes in its present form, and why. The question that interests me is a practical one about how you or I, as prudent people, ought to actually account for the actual options that we might actually buy or sell some day soon. I'm going to confess to you how I, personally, would do things, and try to rationally justify my own choices, which I admit is not as good as giving you a scientific theory that spits out a precise numerical value as the absolute, eternal, observer-independent, peer-reviewed truth. But, after all, my subject is a mere accounting rule, and Aristotle tells us that we ought to treat each subject with the degree of precision it deserves, so perhaps there's even something principled about my refusal to claim to be doing hard, objective science.

What is it, intuitively, that seems so wrong about Black-Scholes? Any reasonably cautious person who sees the way options are promoted to small investors on CNBC can't help but be bothered by the fact that our present theory of option valuation presents the risks of buying options and the risks associated with selling them as being exactly the same, even though one activity involves only a limited possible loss, while the potential losses from the other are unlimited.

This is probably no accident. There's a very interesting interview with Paul Samuelson on the website of the American Finance Association (they're the publishers of the *Journal of Finance*) in which he gives some of the historical background needed to understand why we now look at options the way we do. It's hard not to smile when you learn why European and American options were first given those names (Samuelson seems to have been a bit of a patriot) but reading between the lines of his story about his great discovery concerning the tax treatment of put options is even more amusing.

Everyone knows that commodity producers and other operating businesses often have to hedge their bets to avoid the risk of bankruptcy as a result of sudden price-movements. Before Samuelson came along, they were mostly doing this by shorting (or going long) the commodity itself. Going long a commodity presents no special problems, but selling short is not tax efficient. Why? Because a short-seller, having borrowed some amount of a commodity, immediately sells it, and doesn't own it again until the day he closes out the trade, when he buys enough to repay the original loan. Since he only owns the actual asset for a short time, any profit he makes by owning it is subject to capital gains tax at the higher short-term rate, even if the short position has been on his books for years.

Samuelson, at a time when options were a relatively marginal, somewhat shady business mostly aimed at small speculators, realized that hedging the same position by buying a put option could allow the commodity producer or consumer to pay taxes at a lower rate, because the put option, an asset rather than a liability like the obligation to repay a loan, would have been on his books the whole time *as* an asset. So it seemed like there might be a lot of money to be made, for somebody, by persuading agricultural businesses and miners and airlines that they ought to change their hedging strategy from one built around short-selling to one built around buying puts. (Many people actually hedge with futures contracts, but the way we account for those makes the potential liabilities associated with them fairly obvious, so they're not really part of this story.) There were, however, two problems. Buyers would be reluctant to own assets they couldn't attach a definite value to in some defensible way. What number were they supposed to put on their balance sheet? And the issuers of options had to know how much of the underlying commodity they needed to be long or short of to ensure that their own capital was not at risk from changes in its price.

Conservatism was a basic principle of accounting in those days, as it still is, outside of the domain of the Basel Accords, even now. A genuinely conservative accountant might have argued that selling an option was a risky move, requiring rather a lot of cash or the underlying commodity to be kept on hand at all times just in case. The same tedious person could have insisted that the buyer, if he was really going to be conservative, should record the value of an out-of-the-money option as zero, reasoning that it is not the job of accountants to speculate about the possibility or probability of prices going up or down at some time in the future and impute values to assets on the basis of what they *might* be worth if the universe would only cooperate. But if the issuer had to set aside a lot of capital to cover the risks associated with selling an option, while the buyer had no choice but to carry it on his books as essentially worthless, it is hard to see how there could be a very large market for the things. The seller would have to demand a higher price than he does now, while the buyer, if he did not want to book an immediate loss as soon as he bought his hedge, would have to insist on a lower one.

The great and immediate social utility of Black-Scholes was that it allowed the employees of firms on both sides to get around the accounting principle of conservatism and meet in the middle, by creating the impression that the values of options (unlike any other asset, good, or service) could be scientifically determined by an impressive-looking equation directly linked to the grand theory that markets are always perfectly efficient. (No need for conservatism, because no possibility of error.) And given certain assumptions, the model makes perfect sense. The problem is that 'given certain assumptions' is only an innocuous phrase as long as you're not really doing anything very crucial, as long as you're just cranking out papers so your graduate students can get jobs…

How, exactly, is Black-Scholes linked to the theory that markets are perfectly efficient? It is an implication of that theory, at least as its proponents understand it, that prices in all markets must always follow a random walk, or to be more precise, a geometric Browninan motion with constant drift and volatility. This, in turn, implies that daily price movements, stated as percentages, will always be normally distributed, so the distribution assumed in calculating the variances for Black-Scholes values is a normal one. However, as Nassim Taleb has pointed out at length in several publications, the normal distribution is rather poorly named, as it is, in the universe of possible distributions, actually a somewhat *ab*normal one, in that it has very thin tails. To use the normal distribution as a model is to assert that your statistical universe includes relatively few unusual events or snowballing trends. But many actual statistical universes aren't like that. (Imagine a very diverse ecosystem where you seldom see the same kind of organism twice, except for a few kinds of trees and birds that are very common, or some cascading critical-mass phenomenon like nuclear fission, in which events are either very small or else arbitrarily large.)

My purpose, here, is not to reiterate and readjudicate the whole Black Swan/Fat Tails discussion. I simply want to ask what strikes me as the first, naïve, simple-minded question a person would think of as soon as it occurred to him that the normal distribution might not be the right one to use all the time and in all situations. (What if markets aren't always perfectly efficient, what if some markets sometimes trend strongly, and occasionally crash? A crash that occurs over multiple trading days is a strong-form inefficiency because each day influences the next, and daily price moves during crashes can be almost arbitrarily large.) What would happen to the Black-Scholes value of an option when the distribution of daily price movements follows, say, a power law, or conforms to some other common type of fat-tailed distribution?

Asked in this precise form, the question doesn't really have an answer. It is one of the explicit assumptions of the model that daily price movements are normally distributed, so we're subjecting it to a test that it was never meant to survive. It would be mathematically nonsensical to take the variance of some unspecified fat-tailed distribution and plug it into the cumulative distribution function for the normal distribution.

The derivation of the equation itself via Ito's Lemma assumes that daily price changes in the underlying commodity are a stochastic process of a particular type. The arbitrage pricing argument used to justify the value of the risk neutral portfolio assumes that daily price moves are normally distributed and that the distribution is known. In a world where we don't know whether the distribution of daily price-moves is fat-tailed or thin-tailed, or even whether there is a stable stationary distribution in the first place, we can't know what mix of call options and short stock positions it should consist of. But if nobody can know what the risk-neutral portfolio is, what use is it to know that it ought to have the same yield as a riskless asset? Without the very unrealistic assumptions that the distribution is knowable, and stationary, and known to be stationary (but how could anyone ever know something like that, something which involves predicting future price moves in markets in great detail?) the portion of the argument that is driven by arbitrage is invalidated. When we don't know which member of an infinite population of urns balls are being drawn from, and opinions about this among participants in the market differ, one man's arbitrage opportunity is another's elaborate form of suicide.

None of this constitutes a defense of Black Scholes – basically I'm saying that the model is so artificial that it just can't accommodate the complexities of any real situation – but it does make it hard to say what it would predict if prices were not normally distributed.

You just can't use it then, it doesn't predict anything. To make an actual rival model you'd also have to assume some particular distribution or family of possible distributions. Really, what you want is some model in which balls are being drawn from a family of urns, and you don't know which urn is being used at the moment. The problem is that there's no principled way to pick a particular finite family of urns, or pattern of switching between them. The real world is so complicated that the population of urns needed to model it keeps evolving and changing. So a really satisfactory formal model would be impossible to make. If you made an unsatisfactory one, the risk is that you'd fall in love with it. The very attempt would be perilous.

What we can do is something less ambitious. We can recast the question as one about what would happen, given normally distributed daily price moves, as the variance of the distribution tends towards infinity. This isn't quite the same as using a fat-tailed distribution, but it might give you some rough idea of how wrong you could be, as a result of using Black Scholes, in a worst-case scenario. After all, what I really want to know is a little more than I do now about how badly this existing model might misguide me. A rough-and-ready sketch of a worst-case scenario is better than nothing. An infinite variance of the kind we find in fat-tailed-distributions can be read as the distribution saying to you 'look, don't ever think you can safely predict that I'll stay within any particular limits'. Which, if you think about it, is pretty much the way the world works anyway; Murphy's Law tells us that anything that can go wrong eventually does. So why not just do a Murphy's Law stress-test on Black Scholes, why not try to plug an infinite variance into the existing equation, even if that's not quite as good as coming up with a whole new model, just to see what you get?

Here's the Black-Scholes equation itself, straight from Wikipedia[1]:

$$C\ (S, t) = SN(d1) - Ke^{-r(T-t)}N(d2)$$

Where $d1 = \ln(S/K) + (r + \sigma^2/2)(T - t)/\sigma\sqrt{T - t}$

And $\quad d2 = d1 - \sigma\sqrt{T - t}$

I should define all the terms and carefully explain the meaning of each part, but really that would only waste the reader's time, because there are plenty of other better places he can get that information, so let me make a long story short.

As sigma approaches infinity, d1 also approaches infinity. [It approximates $((\sigma^2/2)/\sigma)$.] $N(\infty) = 1$.

What about d2? It is $d1 - \sigma(T\text{-}t)$. That is, it's:

$[\ln(S/K) + (r + \sigma^2/2)(T - t)/\sigma\sqrt{T - t}] - \sigma\sqrt{T\text{-}t}$

---

[1] The original paper containing the pricing model is F. Black, M. Scholes, (1973). "The Pricing of Options and Corporate Liabilities". *Journal of Political Economy* **81** (3): 637–654. Wikipedia's version is used here for the convenience of readers who may want to quickly check the definitions of the various terms, or have the derivation and current use of the equation itself explained, or who want a list of other sources.

To give the second term the same denominator as the first, we can multiply it by $\sigma\sqrt{T-t}/\sigma\sqrt{T-t}$, since that's just 1. Then it becomes:

$[\ln(S/K) + (r + \sigma^2/2)(T - t) - \sigma^2(T-t)]/\sigma\sqrt{T-t}$

which is just:

$[\ln(S/K) + (r + \sigma^2/2 - \sigma^2)(T-t)]/\sigma\sqrt{T-t}$

or in other words:

$[\ln(S/K) + (r - \sigma^2/2)(T-t)]/\sigma\sqrt{T-t}$

When $\sigma^2/2 >> r$ and $\ln(S/K)$, this approximates $-(\sigma^2/2)(T-t)]/\sigma\sqrt{T-t}$, which should approach $-\infty$ as σ approaches ∞. $N(-\infty)$ the fraction of a normally distributed quantity that has a value less than $-\infty$, is zero. So as sigma approaches infinity, the Black Scholes value approaches:

$$C(S, t) = (1)S - (0)Ke^{-r(T-t)}$$

Which is the same thing as:

$$C(S,t) = S$$

That is, with an infinite variance, the Black-Scholes value of a call option is the full value of the underlying stock or commodity.

The issuer of a call can still hedge his delta in this almost-fat-tailed world, can still arrange things so his profit is unaffected by price movements in the underlying commodity, it's just that to do it he always has to own the full amount of the underlying commodity represented by the option. If the call is exercised, he just turns it over, he doesn't need to care whether its price has suddenly gone to infinity (as a result of something like Israel bombing Iran or Iran bombing Israel, or a sudden unseasonable frost) or how illiquid the market is.

The infinite-variance delta of any call should always be zero, because the changing value of the underlying commodity doesn't influence the amount of that commodity needed to maintain your hedge in the worst of all possible worlds, the world where countries sometimes suddenly go to war, and there is occasional freakish weather. When variance is infinite, the issuer of a call option can only ever be fully hedged by owning the full amount of the commodity which the buyer has a call on. The issuer of a put option can only ever be fully hedged by being short the same number of units, though he can be sure of avoiding bankruptcy by at least having cash or bonds on hand that are worth as much as the number of units that can be put to him at the strike price.

How should we interpret these results, from a practical point of view? If the value of any call option is equal to the full value of the underlying commodity, while their present prices are generally much less, should we all mortgage our houses to go out and buy them? No, of course not, that would be reckless. It's only when *selling* options that we need to be more cautious than Black-Scholes suggests, need to regard our liability, if the options are not covered, as potentially unlimited. Fat tails make it much easier for us to lose everything we have at some one particular point in time, to hit the absorbing boundary of bankruptcy during a crash or a price spike and get stuck there, but they don't necessarily make it more likely that we will routinely win over long intervals. The difference is that you only have to go bankrupt once to be bankrupt forever, so to avoid it you have to completely avoid extreme events, even very rare ones, while to keep on making a profit you have to do well most of the time, which means you care a lot about what *typical* events are like.

(An employee, as opposed to an owner, can always get a job somewhere else, so it's often very hard to get them to care much about extreme events, typically they'd rather hope for the best and try for a high bonus this year. This gives reason to doubt that firms without strong proprietors can survive indefinitely as dealers or traders of options, swaps, and other risky derivatives. Won't their employees always prefer to gamble by assuming that price movements are normally distributed? Doesn't it take some monomaniacal son-of-a-gun to keep insisting that positions must be fully hedged, year after year, while rivals do better by constantly risking bankruptcy? AIG only blew up when Greenberg left, Fanny and Freddy never had private owners at all, Prince Bin Talal isn't managing Citicorp himself, and Bank of America apparently belongs to America. The idea that some half-baked statistical value-at-risk methodology, in the hands of rather junior risk-management staff, is an adequate institutional substitute for a genuine owner with his awful thunderbolt, is ludicrous.)

One of Taleb's most interesting points about fat-tailed distributions is an epistemological one. He points out that it can be hard or impossible, on the basis of a finite amount of evidence, to tell whether a distribution is fat-tailed or thin-tailed, because even a fat-tailed distribution can behave just like a thin-tailed one right up until the moment it finally throws a black swan at you. So you really can't tell, a lot of the time, which world you're in, a fat-tailed world of infinite variance or a thin-tailed world where things are normally distributed. If going bankrupt even once is something you'd prefer to avoid at all costs, the safe thing for you to do is probably to assume, when you are *buying* options, that the distribution is thin-tailed, because if you are wrong, the only consequence is that things will be better than you supposed. On the other hand, if you ever are so rash as to *sell* someone else an option, to actually originate an option, you probably had better assume that the distribution of daily price-movements is fat-tailed, and the variance of the distribution effectively infinite, because in that case you don't only care about what *usually* happens. You also care about the once-every-decade or once-every-generation event that could put you out of business, and force you to get a job working for someone else.

So I, personally, would adopt Taleb's 'avoid the fourth quadrant' maxim – avoid the realm in which problems of unknown magnitude can occur with unknown frequency - as the basis for a prudential accounting rule, which would cause me to put any naked call options I sold, or identical call options I bought, on my own private mental balance sheet at very different values.

The defender of our present way of valuing options would probably say that this rule is patently ridiculous, because we can fully hedge a call we've sold by buying a call on the same underlying commodity. So how can we possibly carry them on our balance sheets at different values? But of course, I fully agree that the two paired positions, provided the underlying commodity, maturity, and strike are exactly the same, present little net risk, although execution problems can still cause one side of the trade to fail while the other goes through. So it isn't quite as good as really having the commodity itself, but it's good enough. Even if the strikes are different, the risk is mostly limited to the difference between them. There should be an accounting rule allowing the carrying of these sorts of perfectly paired positions, which protect the balance sheet in almost every possible world, at a very small capital charge.

It's just that nothing about this special case should be imagined to *generalize* to options that aren't part of such a perfectly offsetting dyad. In those cases, there may be possible worlds in which the naked option you've sold ends up costing you more money than you have. You need completely different accounting rules for the two different cases, covered and naked calls, because the chance that each will ruin you is completely different.

If you go ahead and get rid of half of the hedged position, if you sell the call you bought, this suggests that you'll have to increase the liability associated with the one you issued, resulting in a charge to owner's equity. And so? Aren't you actually in a worse, riskier position, then? Aren't you in greater danger of going bankrupt during a crisis, precisely the time when it's most important to have a strong balance sheet? Isn't that exactly what 'less owner's equity' is supposed to mean? Why should you pretend otherwise, if avoiding bankruptcy is your first priority? You need the capital to cover the call even in a crisis if you want to stay solvent during the unanticipated emergency that will inevitably occur, sooner or later.

(There's always an unanticipated emergency. The past is full of them. Why should the future be different? And that is exactly when you are most likely to go bankrupt, if you're going to at all, so planning for unexpected disasters is one of the keys to survival over the long term. Anyway, I'm just telling you what I would do, you're welcome to go on taking your chances. It'll probably even work out well for you, for a while, and it certainly makes options cheaper for *me*.)

Why wouldn't you just go out and buy another call, instead, to avoid the charge to capital? Well, if you could, you probably would, and you certainly should, unless you're deliberately speculating on the price of the underlying commodity. But if you can't, you won't, and if you don't, you actually will have a new potential problem of unlimited magnitude, we're just arguing about whether or not you should admit it to yourself.

But wait a second, why *am* I treating bankruptcy as a special sort of event which has to be avoided 'at all costs'? Surely this is just a normal business risk, like the risk of making less money than you would like to? Why take such costly precautions just to avoid it? Economists, historically, have had a hard time coming up with an intuitively satisfying story about why people treat going bankrupt as such an awful, costly event. A business that's about to go bankrupt isn't worth anything anyway. So why should crossing a line from one sort of worthlessness to another sort of worthlessness be perceived as so costly? Really, by releasing assets for more productive uses, bankruptcy should actually add value to the economy. Perhaps there are fire-sale costs associated with liquidation….

But to anyone whose capitalist enterprise has ever gone bankrupt, or been close to that point, it is obvious that this explanation does not capture very much about what people are actually worried about in that sort of situation. To know whether the possibility of an occasional unlimited loss should actually make us behave very differently, we would have to be able to say what it is they *do* worry about, and whether those worries are ever justified. But remember, the reason we're having this whole discussion in the first place is that firms that produce and consume commodities are willing to spend money to hedge against extreme price movements. Presumably, since hedging always costs something (there's no such thing as a money machine, a free lunch, and the intermediaries have to be paid some fee if they're to stay in existence, so when hedging we should always expect to make a net loss, just as we do when buying insurance or gambling in a casino) this means that they are willing to pay quite a lot to avoid the risk of bankruptcy resulting from, say, a collapse in the price of corn or a spike in the price of jet fuel.

If bankruptcy is not really costly, the people that run these firms are simply insane, and the whole derivatives market is nothing but a den of vice designed to take advantage of that insanity. But mass insanity should never be the first explanation we reach for when seeking to explain the behavior of people we otherwise treat as rational profit-maximizers. It seems more reasonable to suppose that they have an economically rational reason for taking on this continuing expense, and try to figure out what it is.

What would certainly be inconsistent would be to believe, for example, both the Modigliani-Miller argument that how leveraged a company is shouldn't matter to its shareholders, *and* that deep liquid derivatives markets should exist because hedging is economically rational. The first belief is grounded in the assumption that the costs of bankruptcy are negligible, while the second really only makes sense if they are fairly large. It seems difficult to consistently believe both that it is rational to hedge against the risk of bankruptcy in a costly way and that bankruptcy is costless. Any really serious theoretical discussion of hedging depends on there being some story in the deep background about what the costs of bankruptcy actually are.

(Geometric-mean arguments ignore the fact that the person who sells you a hedge must find a way to cover his own costs, and so will always charge you more to compensate him for the induced variability of his earnings than you can possibly recover by smoothing your earnings yourself.)

I think I have a fairly good one. The truth about the costs of bankruptcy, as I presently understand it, is somewhat awkward because it's impossible to quantify, and requires us to think in a slightly philosophical way. But the story is one I'd like to tell you anyway. Though it may seem like a digression, I can't really finish explaining why I wouldn't use Black-Scholes all the time until I've explained what I think those costs are. Dealing with these two things in isolation from each other would leave the reader still grappling with the same kind of fragmented and dismembered world-picture he could get by reading something peer-reviewed. (A paper that has to pass peer review should, if possible, avoid any discussion of larger ramifications or controversial background assumptions, as that would just give the reviewers free ammunition. Which is fine if all those things are in perfectly good order already, as they never are in the social sciences.) But I think I have a little more than that to offer you, I think I have an idea that ties a lot of things together, and, in passing, tells us something about the logic behind our whole way of life. So I've decided to gamble on your curiosity, on your wanting to hear the whole story. To give it to you, I will have to explain the concept of Knightian uncertainty.

Many people cite Frank Knight, but from what they say, it is sometimes unclear whether their acquaintance with his ideas is first- or second-hand. Often it seems that the person has read a paper by some later author about Knightian uncertainty, but never actually read *Risk, Uncertainty, and Profit* itself. This is, perhaps, understandable, since the book was first published in 1911 and contains very little mathematics, but from a practical point of view it's unfortunate, because the later chapters are full of important and useful ideas.

At the heart of Knight's argument, as you might guess from his book's title, are the distinction he makes between *risk* and *uncertainty*, and the connection he suggests between his kind of uncertainty and *economic profits*.

Modern readers, steeped in the naïve psychological Bayesianism that is the finance theorist's philosophical interpretation of probability talk, tend to want to reinterpret his argument as an attempt to make a distinction between two possible states of mind, in one of which our reasoning can be expected to follow the standard Savage probability axioms, while in the other, it may seem reasonable to us to violate them. But Knight's argument, which predates this whole way of looking at probability by decades, is actually an epistemological, Humean one, about what can possibly be known by observing the frequencies of different types of events in the world.

Certain sorts of events, he suggests – the outcomes of rolling dice, the frequency of different sorts of automobile accidents, the age at which people die – can reasonably be attributed numerical probabilities on the basis of past observations, because they belong to large classes of fairly similar events in which more or less the same sort of causes are operating in more or less the same ways over and over. Other events, however, are effectively unique.

When Bill Gates first encountered the idea of an 'operating system' for a computer, he didn't have the option of looking at all the cases, in previous history, where someone had invented digital computers and then come up with operating systems for them, because there weren't any. This was the first time the human race had ever done any of these things, so assembling a population of past cases and counting the outcomes simply wasn't an option. Instead, Mr. Gates had to follow his hunch about the likely outcome. Famously, his hunch about DOS was right, and the guy who sold it to him was dead wrong. But who knew, at the time? Or consider George Soros's idea of *reflexivity*, the idea that our own understanding of economic phenomena affects the way they unfold. This implies that as soon as we have a genuinely new economic theory, we face a genuinely new economic situation (the one in which people know the new theory) which no human being has ever encountered before. Thus, one effect of Sorosian reflexivity is to continually create a species of Knightian uncertainty. It is events like Gates' purchase of DOS, or the introduction of Black-Scholes itself, whose possible outcomes in terms of profit and loss for particular players can't reasonably be given numerical probabilities in advance, on the basis of statistics derived from a population of similar past events, because they are entirely new and unique, that Knight calls 'uncertain', as opposed to 'risky'.

How do we form hunches about uncertain things? Some of us, Knight points out, are demonstrably much better at it than others. A guy like Steve Jobs, to use a contemporary example, seems to have an aptitude for determining what's plausible on the basis of general considerations about what people are like and how the world works that's something like the skill a good novelist has. Again and again, he has encountered situations no human has ever faced before, guessed how they would come out, and been right. Like a good poet or a good science fiction writer, the successful capitalist has the ability to strain analogies to their limits and still get useful results. There were no statistics about whether or not people would buy iPods until well after a lot of money had already been spent on design, factories, and marketing. Jobs had to be right, despite a lack of exactly comparable cases to extrapolate from, and in fact he was, which is one reason we all know his name.

The implausible economic phenomenon that Knight himself was actually trying to explain with his idea of objective Humean uncertainty was the existence of *economic profit*. Our usual models of the way economies work seem to imply that nobody should ever really make a net profit, once the costs of labor, materials, taxes, etc., have been taken out of their revenues. If some new business were to become profitable, other entrepreneurs would start competing companies, devoting more labor and capital to the profitable line of business, until there were so many firms doing the same thing that the profit was wiped out. At equilibrium – the condition most of our economic models assume – there can, apparently, be no economic profit at all. This fact was behind Karl Marx's belief that a capitalist system can only survive by continually expanding and finding new people to exploit; and within the formal science of mathematical economics, the basic approach still hasn't changed in any way that solves this rather fundamental problem.

(This is why I am sometimes tempted to worry that Samuelsonian economists, though they might know something about the efficient allocation of resources in a static, unchanging, equilibrium world like Pharaonic Egypt, don't really understand capitalism at all. If the world were the way all their models assume it is, the profits that drive the whole capitalist system could not possibly exist. When you think you've proven that something that actually exists can't possibly exist, it becomes somewhat doubtful that you're the person who really understands why it is the way it is.)

Knight's largely forgotten solution to this paradox was to suggest that all technological change necessarily creates Knightian uncertainty, creates lots of new situations nobody has ever seen before, and that all real economic profit in a modern economy derives from the resulting series of protracted disequilibria. We don't know how these new situations are going to evolve, because we've never seen a wind-powered sawmill or a steam engine or a railroad or a telegraph or a telephone or an automobile or an airplane or a computer or an iPad before, but some of us have much better hunches about these sorts of things than most people. It's those individual people with good hunches who earn economic profits, by making unusually good guesses about how things are likely to work out. The reason the profits don't quickly get arbitraged away is that it takes people who aren't quite as creative a long time to realize that the real entrepreneur is not crazy and begin to imitate him, by which time he's already moved on to some other uncertain project on the basis of some new keen hunch. You become Bill Gates or Steve Jobs or Warren Buffet by doing this over and over and being right almost every time. People seldom get rich by just correctly measuring a risk; the real source of economic profit in a technologically dynamic capitalist economy is individual intuitions about objective uncertainties.

In an economy with no new technology, Knight argued, the social role of entrepreneur would eventually go away; ancient Egypt didn't have them because it didn't need them, and if we froze our own technology at 1911 levels for a thousand years, we, too, would eventually be able to centrally plan our economy without any loss of efficiency. Modern economic growth – a phenomenon of the last few centuries, as medieval economies usually only had annual growth of tiny fractions of a percent - is the consequence of modern science, the new technology it continually produces, and a society set up to continually discover the uses of that new technology by letting people play their private hunches about what they might be, and succeed or fail on their own merits. (Of course, modern science and technology are, reciprocally, themselves the results of having a certain kind of social set-up; the two things are mutually reinforcing, and had to both come into existence, slowly, asynchronously, and by fits and starts, during the same few centuries. No human being thought of the whole system; it just evolved.)

This Knightian explanation for the existence of economic profits in a capitalist society is satisfying in a number of different ways. It helps explain the historical incidence of capitalism as a social form, why the invention of the rudder and the compass and gunpowder and the printing press and the telescope and the clock and the multitude of other early modern inventions, small and large, had such a transformative effect on our whole society, and why that rapidly transforming society eventually became so good at producing and integrating yet more inventions. It helps explain why planned economies are never really able to catch up, or keep up, technologically. (Central planners are not poets, they always hate uncertainty.) It attaches some economic logic to our historically atypical modern tolerance for unusual points of view, and our strange patience with what Locke's pupil Shaftesbury called 'raillery', the mocking of received notions. The reason it's worth talking about here is that it also provides us with a perfectly good explanation for our shared perception of bankruptcy as a costly and disastrous event.

At the moment your company goes bankrupt, it is, as far as anyone else is concerned, worth nothing, so nothing is lost. But you would never have bothered with all the effort involved in starting it if you didn't have a hunch that some new line of business was likely to be more profitable than most might suppose. If you're a fool, the hunch was probably wrong, and the world is doing you a favor by making you stop wasting your time. But if you're Steve Jobs, and you go broke because you can't afford the rent on the garage where you're building the first Apple computer, and so have to go work at your uncle's auto dealership, and forever give up your dream of becoming a computer mogul, you're losing rather a lot, even if you're the only one who knows it.

Once a company has been in existence for a generation or two, once it becomes General Motors or Microsoft as they are today, the original Knightian intuition behind its success may have played itself out, and not been replaced by anything nearly as good. It's a lot easier to gamble everything on wild hunches, anyway, when you have almost nothing to lose. Those companies could go bankrupt, now, if their share prices were to happen to fall to zero, without causing any further loss to shareholders beyond the ones they'd already suffered in the fall. And, in fact, this is probably one of the reasons why we have the institution of bankruptcy, to clear mature, rent-seeking, uninventive businesses out of the way, so we can move the next step forward. But bankruptcy for Microsoft in 1981 would have been a personal tragedy for Gates no matter what the share price happened to be. It is precisely the fact that nobody else can see an opportunity that makes it worth gambling everything on in the first place. If the opportunity is real, if you really are a genius, but you lose by bad luck, or by having failed to hedge some crucial cost, then though only you know it, you've lost a lot more than the accounts would ever show. And so has your society. In fact, there isn't any other situation in which you can ever lose as much, because it's only when everybody else is wrong that you stand to make abnormally high returns.

This, presumably, is why people spend money to hedge their bets on things like the price of corn or the price of jet fuel – because that particular price over that particular period isn't what their Knightian hunch or variant perception is about, and they don't want their grand vision of the future disrupted by exogenous accidents before the value that only they can see is realized in the eyes of the rest of the world. Bankruptcy is either costly or costless, in terms of lost opportunities, depending on whether the hunch is right or wrong. Fools, if they only knew it, should actively seek bankruptcy, because it will release them from their self-imposed Sysiphean task, but the wise man should fear it, because he has a lot more to lose than it seems on paper.

So if you happen to be wise, and not a fool, then you really should try very hard to avoid ever going bankrupt, and that means that selling options, in the sense of actually issuing a call or a put, is something you should only ever do when you are, for some reason, prepared to continually hold the full amount of cash or the underlying commodity (or calls on the underlying commodity) which you would need to provide the buyer in a worst-case scenario, and yet not averse to giving up that cash or stock of the commodity, or exercising your calls, at an earlier date should the buyer ask you to. If you are a fool, on the other hand, you and your money will be soon parted anyway if you aren't very cautious indeed, so there is no point in advising you on how to trade options with it, the best thing you could do is put it away somewhere safe.

Buyers of options should probably value them as if daily or weekly price movements are normally distributed, because that's a conservative accounting assumption for a *buyer* to make; but in my opinion prudent sellers will follow the accounting principle of conservatism and act as if the variance is arbitrarily large, because it sometimes is, and always might be. This implies that they should only ever issue covered calls, and short enough stock, or reserve enough cash, to fully cover the potential loss from all outstanding puts.

Of course, if this approach to valuing options were to be widely adopted, options and related derivative markets, as they are today, could hardly exist. Buyers and sellers would want very different prices, and fewer of them would end up meeting in the middle. Dealers would try much harder to match up their clients rather than taking on risk themselves, since the first way of doing things would tie up a lot less of their capital. (Unfortunately, that means they would probably go on trying to encourage their *clients* to sell naked calls, so they wouldn't have to themselves, and that means that it would be in their interest to go on publicly professing strong belief in Black-Scholes…)

But much of the options-related activity in contemporary markets is nonsense, mere casino gambling, anyway, because much of it is done by people who aren't actually hedging any sort of bankruptcy risk in the first place. Which is not to say nobody ever makes money trading options. In my experience, traders make money, trading options, if and only if they would actually be just as good at trading the underlying asset, and as in any casino, the house always takes a cut.

It's probably too late to tell Wall Street how options *should* be valued, because they're already institutionally committed to doing it incorrectly. For those firms whose demise would not actually interrupt any very brilliant and unprecedented entrepreneurial plan, the hidden costs of this sort of recklessness, to shareholders, are likely to be small anyway, though there may be systemic risks to consider. But if you're not in the banking business, somebody should probably tell you that our present method of valuing options is not something you ought to take without a grain of salt, and that you should never sell an uncovered call or put.

A good thing about my very cautious approach to options, as far as I'm concerned, is that it hasn't ruined me yet. Not everyone can make the same boast. But we don't really seem to learn, somehow in this kind of science failed experiments aren't seen as falsifying theories, the theorist gets a do-over or some kind of free pass. (As a supposed expert in the philosophy of science, I, of all people, ought to know why this is okay, if it is… But I don't seem to be aware of knowing any such thing.) Our collective punishment for bailing out LTCM and pretending that the whole incident never happened is that we now have to see Options Action every weekday at 5:51 on Channel Zero. Well, *Mundus Vult Decipi*; you ought to know better, anyway, than to invest in things you see advertised on TV, whether the person selling them is G. Gordon Liddy, or some guy with a fancy equation.