Lecture Note on Options

Barry W. Ickes

Fall 2003

1. Introduction

Options are an interesting type of financial contract. Currency futures are an important way in which currency risks can be hedged. They offer the right, rather than the obligation, to purchase or sell at some given price (strike price). More generally, they offer flexibility to take an action or to get out of taking it. They are very useful in a variety of transactions. We discuss here their use as a hedge and as a speculative investment.

2. Options and Hedging

Suppose we need to deliver pounds, \$, in the future at a price of \$1.50. If the \$ appreciates it will cost me more to deliver on this pledge. In figure 2.1 we can see how my short position's profitability varies with the future price of the pound. Let us map the return to a short position on the pound. This is given in figure 2.1, where we have the dollar price of the pound on the horizontal axis. Clearly a short position loses money as the pound appreciates above the contract price of \$1.50. Appreciation means I will have to pay more to cover my position. When the pound is low I profit because I can buy pounds cheap. When the price rises I lose because I will have to pay a higher price to obtain pounds. This is very straightforward and simple. But I bear all the risk of the price fluctuation.

Hence I may want to hedge the risk. One way to do that is to purchase pounds in the forward market. This eliminates my currency risk. But it also means that I am obligated to purchase the pounds. But what if the pound is low in the future? Then I would have been better off not purchasing the forward contract. If I want to retain this flexibility I can purchase an option.

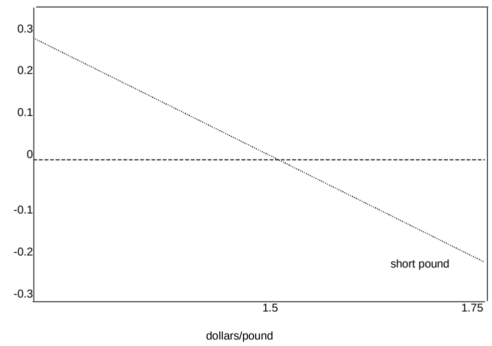

Figure 2.1: Return with no hedge

How do options work? These are rights, rather than obligations, to buy or sell. We can see that the payoff from a call or put option depends on whether it is exercised or not. Let S be the spot price of a currency and K the strike price. Then the payoff from a call option is

$$C = \max[0, S - K] \tag{2.1}$$

where C is the payoff from the call option. If $S < K$ then the option would not be exercised. Its value is zero, and the purchaser loses only the premium paid for the option. Of course, the seller of the option does not have such a limited liability. The seller of a call option faces unlimited liability as S could rise without limit.

We can similarly write the payoff from a put option as:

$$P = \max[0, K - S]. \tag{2.2}$$

For a put option, if $S > K$ the option will not be exercised – I can make more by selling in the spot market. Hence, its payoff is zero. If the spot price falls, however, then the put increases in value. Notice that for the put option again the purchaser cannot lose more than the option premium, but the seller's potential losses, though still large, cannot exceed the strike price, K, because the spot price cannot fall below zero.

One interesting use of a put option is as insurance against a severe decline in an asset price. I can insure my portfolio by purchasing an out of the money put option on that portfolio. Under normal circumstances I am just out the premium which is cheap given that the option is out of the money. But if the market declines the option limits the downside risk. The seller of such put options is providing catastrophe insurance. Taking a small premium against a very large, but unlikely, loss.

Now suppose I am a fund manager and I sell such out of the money options. I will gain small premium income each period. If I also have invested in an index, I will beat the market, until the catastrophe strikes. If my bosses look only at my returns I will be rewarded. What if the catastrophe strikes? The worst that can happen to me is that I am fired.[1] Hence, I have an incentive to take such risks (moral hazard). For this reason many contracts with fund managers are based on more than just returns; they also limit the type of investments that can be made.

In the mid-80's insights about put options gave rise to **portfolio insurance**. The idea was to insure a portfolio against declines using put options to set a floor. If the value of the portfolio fell the option would be exercised and losses limited. But put options were not widespread at the time,[2] so the portfolio insurers instead focused on using short sales to mimic the put option. This exploited the insight of Black-Scholes that a risk-free asset could be created by combining shares of stock with an option. So Leland decided to create an option by combining stock with cash.[3] To insure against declines in the S&P 500 the insurers calculated appropriate short sales of the futures index. As the value of the portfolio falls short sales increase and more of the portfolio is held in cash. As the portfolio value drops towards the strike price of the hypothetical put option the portfolio is all cash. A put option is replicated. A complicated program calculates the appropriate short sales. Eventually quite a large market is created.

This worked fine until the market jumped discontinuously. When the market jumped down the required short sales were huge, the prices dropped more, and the market crashed; the Dow falling 500 points on October 19, 1987. The key problem was the lack of the put option – which is a contract – and the substitute of short sales which require liquidity.[4] In normal times markets are liquid and

[1] The payoff to the manager of an investment fund is a call option. Clearly moral hazard is a problem here. Many funds limit the ability of managers to pursue such strategies.

[2] Most of the options that were traded at that time were call options.

[3] See the example in 3.1 on replication. In that example a call option is combined with holding a share of stock. The resulting portfolio has the same risk as cash. To replicate a put option you would hold the stock and cash. As the price of the stock goes down you hold more of your portfolio in cash. As the price of the stock reaches some floor your portfolio is all cash. Any further declines in the price leave your portfolio unaffected. You can see the relation of this to a put option by replacing the idea of a floor with the strike price.

[4] Alternatively one could point to the assumption in Black-Scholes that prices move contin-

there is no problem. But when everyone wants to sell at the same time markets do not function. Discontinuous price jumps implied very large short selling. But somebody has to purchase for a market to work. Lack of liquidity makes it impossible to make the transactions that are part of the insurance strategy.[5] This ought to have been a warning but it was not really heeded.

2.1. A Simple Example

Suppose I purchase a call option.

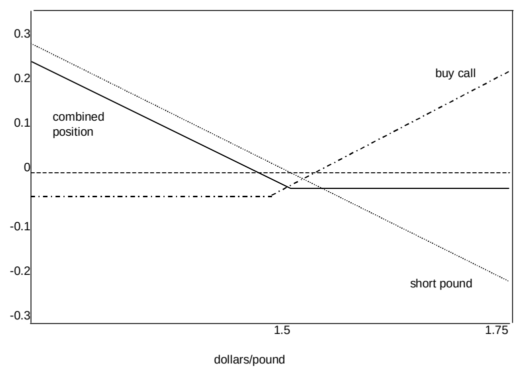

Figure 2.2: Using a Call to Hedge a Short Pound

I can buy an option today to purchase $ in the future at a price of \$1.5/£.[6] If the pound appreciates above this I will purchase it, else I will not. Let us assume that the call option costs \$0.05 per cost of the option.[7] Then if the pound

uously. This was violated at that date. October 19, 1987 was a Monday, and at the start of trading there were \$8 billion of futures waiting to be sold at Chicago. Prices opened 7% below the previous closing price, which immediately caused a new stream of short sale orders.

[5] If returns were distributed normally the size of the price jumps observed on this Black Monday would happen once a millenium. The S&P 500 index fell 20.5% on October 19, 1987. This was 20 standard deviations from the mean. That ought to occur once in the history of the universe if returns are normally distributed. But extreme price events like this are more common than would be predicted by the normal distribution.

[6] The actual (September 19, 2002) price of a November call option with a strike price of 1520 (\$1.520/pound) was 3.12 cents per pound, for a contract of 62,500 pounds. So the cost of purchasing the contract would be $62,500 * (\$0.0312) = \$1,950$ plus commission.

[7] How is an option priced? The basic theory was developed by Black, Scholes, and Merton.

appreciates my risk is reduced. The value of my position is given in figure 2.2. Notice the asymmetric nature of the hedge – the combined position – in figure 2.2. If the pound depreciates the buyer of the option lets it expire, and he earns a high profit. This is superior to purchasing pounds future. If the pound appreciates, the option holder strikes, and the downside loss is bounded.

We can similarly examine the position of someone who is long in pounds. This individual wants to hedge against the possibility of the pound depreciating. The purchase of a put option can help here. Basically, the agent obtains the right to sell pounds in some specified future period at a given price, say $1.50. If the pound depreciates I can sell the pounds at the strike price. If the pound appreciates, however, I just let the option expire and benefit from the appreciation. Hence, my combined position is as in figure 2.3. Again, notice how the combined position is asymmetric, allowing a large gain but putting a floor on my potential losses.

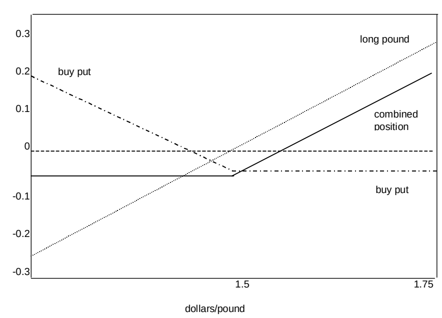

Figure 2.3: Using a Put to Hedge a Long Pound

3. Volatility and the Option Price

Examining the value of the hedge we have just examined demonstrates the role of volatility in determining the price of an option. Consider for example using

They first showed that a portfolio that consists of stock and options can eliminate risk. One offsets another. Then they noted that arbitrage implied that the return on this portfolio must equal the risk-free rate, else there would be profit opportunities. Then they noted that hedging risks of a volatile asset is worth more than for a less risky asset, hence the price must be higher for such ano option. This led to the key insight that it is the expected volatility, not the expected price that matters. See 3 below.

the call option to hedge a short pound. The value of this option is the insurance it gives you as the pound appreciates. The greater the volatility the larger loss that is averted by holding the combined position, hence it becomes more valuable in this case. Notice that volatility does not alter the expected value of the short position – this is equal to zero if appreciation and depreciation are equally likely. But volatility makes the call option, and thus this hedge, more valuable.

Suppose that the current price is \$1.50, and that with equal probability the pound next period will be \$1.40 or \$1.60. What is the expected value of the combined position? There are only two cases to consider:

- $\$ = \$1.40$: The option is not exercised, I buy pounds at \$1.40 and sell them at \$1.50, earning ten cents.

- $\$ = \$1.60$: The option is exercised, I buy at the strike price which equals the price I shorted the pound, so the payoff is zero.

Of course, we must also pay something for the option. Let this price be 5 cents. We can thus write the expected value of this hedge as follows:

$$EV(H) = \frac{1}{2}[1.5 - 1.4] + \frac{1}{2}[0] - .05 = 0 \tag{3.1}$$

where the first term in the brackets is the gain in the short position if the option is not exercised, and the second position is the gain if it is (this always equals zero because we assumed the current price and the strike price were the same for simplicity). Now suppose that the volatility of the pound is higher – that is, assume that next period it will either be \$1.30 or \$1.70. Notice that the expected value is unchanged under these circumstances, only the variance. What is the expected value now?

$$EV(H) = \frac{1}{2}[1.5 - 1.3] + \frac{1}{2}[0] - .05 = 0.05 \tag{3.2}$$

The net payoff is now larger, even though the expected value of the exchange rate has not changed. The reason why this occurs is the non-linear nature of the payoffs from the option. Notice from figures 2.2 and 2.3 how the combined payoff is nonlinear. If there is very little fluctuation in the exchange rate the expected value hardly changes. But as volatility increases the value of the option increases in one direction and is unchanged in the other. The average of this must then increase with volatility.

3.1. Volatility

Higher volatility thus makes the hedge more worthwhile. So the amount you would pay for such an option is increasing in volatility. Hence, the price must be positively related to volatility. Indeed, you can see from this simple example that if the only outcomes were \$1.30 or \$1.70 with equal probability then the option price could **not** be \$0.05. From 3.2 we see that the expected value is positive. For the seller of this call option, however, the price is too low – the expected value is negative. Sellers would demand a larger premium given these potential outcomes. Purchasers would bid up the price in this case. If the option cost 6 cents I would still be willing to purchase it. Indeed, given these odds, the price of the option would have to be ten cents. This simple example illustrates that once we know the volatility we can figure out the price of the option. It also suggests that if we know the market price of the option (along with the spot price, strike price, time to maturity and the rate of interest) then we can determine the market's estimate of volatility.[8]

Table 1: Value of a Portfolio of 2 Options and the Stock

	return on the option	return on the stock	total
if the stock price = 110	−(2)(\$10) + 2(\$5)	10	0
if the stock price = 90	2(\$5)	−10	0

One **very interesting** feature to note here is that we did not have to say anything about risk aversion (or expected value) to obtain the price of the option. The reason is that we can replicate an option with an alternative portfolio that eliminates risk. Consider the following simple example (see table 1). Suppose we have sold two call options on a stock that costs \$100 today and will either be \$110 or \$90 next period, with equal probability. If the price goes up the option is exercised and we lose \$10 per option; if it goes down we gain some premium income. So this is risky. But suppose we hold a share of the stock. Then if the price were to go up we gain \$10 on the stock itself. If the price goes down the options are not exercised but we lose on the stock. We lose \$10 in either case. But we forgot the premium income. If the options are priced at \$5 each, then we have \$10 in premium income. Now our portfolio of shorting the two options and long in the underlying stock earns a certain return (of zero). Holding this combined

[8] This last point is important. We know past volatility, but investors need to know future volatility. They can extrapolate from the past, but the past may not be prologue. But since we can observe the market price, we can use options pricing formulas to infer the market's estimate of volatility: **implied volatility**. So if we are looking at currency options, we can infer what the market expects the volatility to be of the underlying exchange rate. See 4.1.

portfolio eliminates the risk, which is why risk aversion does not matter for the price of the option.

This is called replication. Of course in practice the price of the assets change daily, so you will have to adjust the options you short in your portfolio each day to replicate.[9] But this technical problem is easily solved in practice. The key point is that throug replication you eliminate the risk, and this does not depend on the underlying price of the asset, just its volatility.

To see the role of volatility in the price of the option now suppose that the stock will be worth either \$120 or \$80. The expected value is the same, but now the option must be priced at \$10. The reason why volatility increases the option price is that it is more difficult to construct the replicating portfolio when the price is more volatile. A higher price for the option offsets the increased price volatility.

4. Straddles

A straddle is a combination of two options designed to speculate in volatility rather than in the price of the asset. An investor may think that volatility will be low in the future, and hence may want to bet on this. For example, a currency that has been very volatile may be stabilized in the future if a proper program is implemented. Hence, a bet on stability may be warranted.

This can be done by selling both a call option and a put option. This is called a short straddle. To see how it works, first look at the profit position for the seller of a call option. In figure 4.1 we map the profit position for the seller of a call option with a price of \$0.10. Notice that if the price of the pound is low the option will expire unused. The seller will then earn the option premium of ten cents. If the pound appreciates, however, the seller will have to deliver the pounds at \$1.50, and this will result in a loss. The seller of the call option breaks even at \$1.60.[10]

[9] Our simple example is like a fork in a tree where two branches emerge. Next day there is another fork in the tree. Over many days you have many forks. Tracing these out you get the actual price path that the asset follows. So how to price this? Note that just before the option expires you have only two forks. So you can value that option. Now work backwards to the preceding fork and continue to the original one. This is called backward induction. Then you simply sum all the paths that the price could follow, since you have valued the option in all the cases. It turns out that the formula for doing this is the Feynman-Kac formula from Quantum Electrodynamics. But that is really another story.

[10] Notice that the buyer of a call option also breaks even at \$1.60. Can you draw the corre-

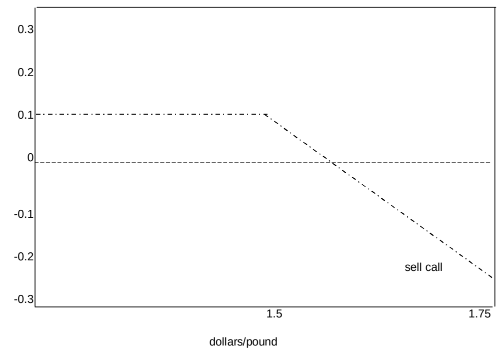

Figure 4.1: Profit Position for Seller of a Call Option

What about the seller of a put option? This is a promise to purchase pounds at \$1.50 in the future (recall that the purchaser of a put option receives the option to sell pounds at that rate). If the pound rises above \$1.50 in the future the option will not be exercised – why sell pounds at \$1.50 when you could sell them at a higher price! Hence, the seller of the put earns the premium of ten cents. But if the pound falls below \$1.50 the seller of the put loses. Hence, we have:

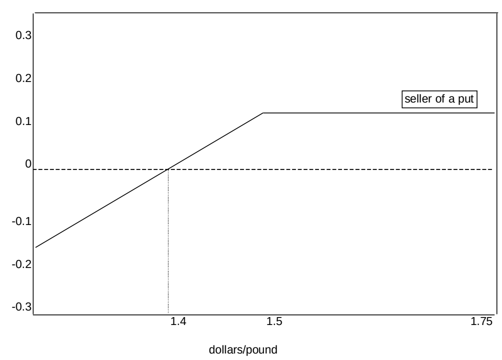

Figure 4.2: Profit Position of the Seller of a Put Option

Now we can put the two profiles together to consider a straddle. In this case the sold put option is in the money with a high price of the pound and the sold

sponding position for the buyer of the call option?

call option is in the money for the low value of the pound. Hence, if the pound remains between the points of \$1.30 and \$1.70 the short straddle will be in the money. This is evident in figure 4.3.

To see how the straddle works, notice that if the pound stays unchanged at \$1.50 neither the call option nor the put option is exercised. Then the short straddle will pay 20 cents, the sum of the two premiums. This is the maximum the straddle can earn. If the pound stays above \$1.3 the straddle breaks even, because the premium on the call option offsets the loss on the put option. Below \$1.3, however, the ten cents received on the call option is insufficient and the straddle loses. Now suppose that the pound rises to \$1.60. The call option just breaks even at that point, but with the premium from the sale of the put we are still in the money on the straddle. Once the pound rises above \$1.7, however, the straddle loses. So the short straddle is profitable if the pound is not too volatile. But if the pound fluctuates a lot the straddle loses.

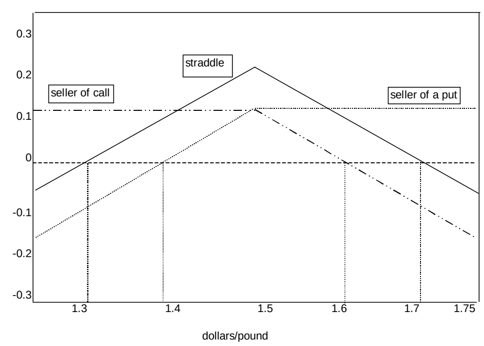

Figure 4.3: Profit Position on a Short Straddle

To bet on more volatility we would buy a call and put option and create a long straddle. Obviously, our position would be the mirror image of that in figure 4.3, and we would profit if the pound fluctuated a lot. An example of a long straddle with a \$.05 premium per option contract is given in figure 4.4. It is evident from figure 4.4 that this bet is in the money if the pound fluctuates a lot in the future. With a low pound the put pays off, and with a high pound the call option pays. But if the pound does not fluctuate, the holder of this straddle is out two premia – in this case \$.10 per contract.

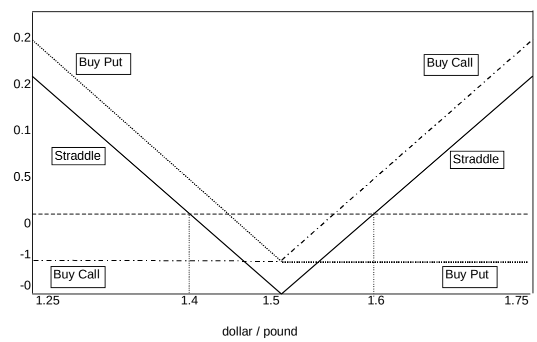

Figure 4.4: Profit Position on a Long Straddle

4.1. Implied and Historical Volatility

To make a volatility bet one has to compare the market's judgement of future volatility with one's own. How can we measure the market's estimate of future volatility? Fortunately, options theory provides the answer. Options theory explains how such contracts are priced. As we noted earlier options are more valuable the more volatile is the asset. The Black-Scholes formula says how an option should be priced given the interest rate, strike price, spot price and duration of the contract, plus an estimate of the future volatility of the price. Once you have the estimate of volatility you can figure out the price an option should be.

But you can also go backwards. Since the market price of the option is known, you can use the formula to obtain the **implied** volatility. Then this can be compared with the historical volatility.

Implied volatility tends to be a fairly poor predictor of actual volatility. Implied volatility explains relatively little of future volatility, but it tends to get the direction of changes correct. It is an open question of whether there are trading profits available.

5. Bad Volatility Bets

Two notorious cases of volatility bets gone wrong are Barings in 1995 and LTCM in 1998. These are two interesting cases in that they both involve the same type of bets, but the source of the ultimate problems were directly opposite. In the Barings case a "rouge" trader was able to bet huge amounts because his bosses were willing to believe that he could beat the system; in other words, that markets are not efficient. His record did not arouse suspicion until it was too late. In the LTCM, on the other hand, it was precisely the belief in market efficiency that got them into trouble. They believed so much in efficiency they were convinced that anomalies must be wiped out. That is how they got into trouble. It is worth a little discussion of each.

5.0.1. Barings

Perhaps the most notorious use of straddles was by Nick Leeson, who managed to bankrupt Barings in early 1995. At the time of its collapse, he had bet against volatility in the Japanese Stock Market, by creating a short straddle. He sold about 70,000 call and put options on the Nikkei 225 stock index, worth about $7 billion.[11] He was betting that volatility on this market would decrease. Meanwhile he was earning lots of premium income from the sales of the options, and this was helping his other speculations from being exposed. The problem is that volatility increased after the Kobe earthquake.

The strike prices of most of Leeson's straddle positions ranged from 18,500 to 20,000. He thus needed the Nikkei 225 to continue to trade in its pre-Kobe earthquake range of 19,000 - 20,000 if he was to make money on his option trades. The Kobe earthquake shattered Leeson's options strategy. On the day of the quake, January 17, the Nikkei 225 was at 19,350. It ended that week slightly lower at 18,950 so Leeson's straddle positions were starting to look shaky. The call options Leeson had sold were beginning to look worthless but the put options would become very valuable to their buyers if the Nikkei continued to decline. Leeson's losses on these puts were unlimited and totally dependent on the level of the Nikkei at expiry, while the profits on the calls were limited to the premium earned.

Some of Leeson's biggest bets happened at this point. He started going very long on the Nikkei. Either he thought the market over-reacted, or he wanted to shore up the Nikkei to protect the long position which arose from the option straddles. (Leeson did not hedge his option positions prior to the earthquake and his Nikkei 225 futures purchases after the quake cannot be construed as part of a belated hedging programme since he should have been selling rather than buying.)

When the Nikkei dropped 1000 points to 17,950 on Monday January 23, 1995, Leeson found himself showing losses on his two-day old long futures position and

[11] At the time, Barings reported capital was about $618 million.

facing unlimited damage from selling put options. There was no turning back. Leeson, tried single-handedly to reverse the negative post-Kobe sentiment that swamped the Japanese stock market. On 27 January, account '88888' showed a long position of 27,158 March 1995 contracts. Over the next three weeks, Leeson doubled this long position to reach a high on 22nd February of 55,206 March 1995 contracts and 5640 June 1995 contracts.

The large falls in Japanese equities, post-earthquake, also made the market more volatile. This did not help Leeson's short option position either - a seller of options wants volatility to decline so that the value of the options decrease. With volatility on the rise, Leeson's short options would have shown losses even if the Tokyo stock market had not plunged.

The puzzle in the Barings case is how Leeson was allowed to bet such large stakes. Basically, he had found an accounting system mistake and exploited it. He put all his losses into a hidden account and showed his gains in his trading account. His superiors thought he was a super trader and they kept giving him more collateral to trade with. He needed it for margin calls on deteriorating positions, while he claimed it was needed to fund his new clients' trades. By the time anyone found out he had bankrupted the firm. The puzzle is that they did not inquire more into "how" he could be beating the market so regularly. Aren't markets efficient?

5.0.2. Long Term Capital Management

Another important case was LTCM. This famous hedge fund was to profit from relative trading. The idea is to bet not on market levels but on pricing anomalies. The basic idea went something like this. Imagine two assets–say, Italian and German government bonds–whose prices usually move together. But Italian bonds pay higher interest. So someone who "shorts" German bonds–receives money now, in return for a promise to deliver those bonds at a later date–then invests the proceeds in Italian bonds, can earn money for nothing. Another noted play was to short 30 year Treasury Bills and purchase 29 year T- bills. The idea is that the new ones are more liquid so the 29 year T- bills bear a higher return. Notice that these differences are likely to be very small, so it is critical to make many and very large bets.[12] You earn small amounts on each transaction, but with enough

[12] LTCM had something like $7 billion before returning ($2.7 billion in early 1998 to maintain leverage) in capital and $1.25 trillion in outstanding positions (during 1998 capital fell to $2.5 billion as their trades turned bad). Of course, many of these positions were offsetting, so it

leverage you can really make something. Here LTCM's reputation was critical. This allowed for huge leverage and no transaction fees.

Of course, it's not that simple. The people who provide money now in return for future bonds are aware that if the prices of Italian and German bonds happen not to move in sync, you might not be able to deliver on your promise. So they will demand evidence that you have enough capital to make up any likely losses, plus extra compensation for the remaining risk. But if the required compensation and the capital you need to put up aren't too large, there may still be an opportunity for an exceptionally favorable trade-off between risk and return. LTCM was able to borrow at very low costs and paid very small fees due to their reputation and to the fact that major investment banks were also investors in the fund.

Actually it's still not that simple. Any opportunity that straightforward would probably have been snapped up already. What LTCM did, or at least claimed to do, was find less obvious opportunities along the same lines, by engaging in complicated transactions involving many assets. For example, suppose that historically, increases in the spread between the price of Italian as compared with German bonds were correlated with declines in the Milan stock market. Then the riskiness of the bet on the Italian-German interest differential could be reduced by taking out a side bet, shorting Italian stocks–and so on. In principle, at least, LTCM's computers — programmed by those Nobel laureates — allowed the firm to search for complex trading strategies that took advantage of even subtle market mispricings, providing high returns with very little risk.

Volatility trades were an important bet that LTCM was making in the summer of 1998. Volatility was expensive. Investors wanted to hold long straddles rather than short ones, as markets were very volatile. LTCM conjectured that markets were too volatile. Hence, they sold the long straddles and held the short straddles. They were betting on a return to historic volatility levels. But they turned out very wrong. Volatility increased after the Russian crisis and these bets turned out very bad. An unprecedented flight to quality in all markets widened spreads rather than shrinking them. Then two problem struck. First, LTCM was losing the same type of bet in many markets. Second, it was so large it could not unwind its positions.

The underlying principle guiding LTCM was market efficiency. If markets are efficient then the pricing anomalies that LTCM found would have to be eliminated. That is what they were betting on. Given enough arbitrage activity they were sure to win.

---

sounds more dramatic than it really was. But they did borrow extreme amounts.

The problem that LTCM ran into was a lack of liquidity. Notice that if you are betting on convergence in markets the greatest profit opportunities are when the spreads widen. If you believe that arbitrage will eventually work you can make very large gains. The problem is sufficient liquidity.[13] It may take a lot of time for the arbitrage to work. In the meantime your portfolio is deteriorating. You may give up or run out of assets before the arbitrage pays off.[14]

The problem is even worse if you are a manager of a fund. Then you are judged on your relative performance. If your portfolio starts to do poorly you may start to dump it. If you have investors they may only notice the deterioration of the portfolio, not the potential for larger gains. Hence, they may demand repayments. This hastens the liquidity crisis. And it weakens the forces of arbitrage.

A related problem occurs when everyone in a market starts to expect the same thing. Markets are liquid when there are buyers and sellers. Differences of belief spur markets. When everyone thinks the market will fall no one wants to buy. The market becomes very illiquid. In August 1998 investors around the world, to an unprecedented degree, fled to safety (mostly T-Bills). Now LTCM was betting that price anomalies would narrow. But the flight to quality killed them. A simple example is a T-bill corporate bond play – short the T-Bill and buy the corporate bond. If the latter does not default its higher yield should produce profits. But with the flight to quality the T-Bill becomes more expensive. The problem for LTCM was that this divergence was happening for all of its bets. Were these independent bets the odds would be astronomical against such an outcome. But these were not independent events. The same phenomenon – the flight to quality – was at work in all the bets.

---

[13]To see the problem first consider a "martingale bet." Bet one dollar on a coin toss (heads you win) and then keep doubling the bet after each loss. If you win quit. Notice that this strategy guarantees a one dollar win no matter how many tails come up in a row. If you lose the first four tosses you bet 16 dollars. If you win your net position is one dollar $= 16 - 8 - 4 - 2 - 1$. This is true even if there are 100 tails in a row. The problem, however, is that you may run out of money before you get to a heads. Notice that the amount of money you need to bet n times is $2^n - 1$, so that when $n = 10$ you would need 1023, but when $n = 20$ you need \$1,048,575.

[14]You may wonder how long it can take for arbitrage to work. Consider the case of Royal Dutch Shell. Royal Dutch Petroleum and Shell Transport are two legal entities which been allied since 1907, and split revenues 60-40 ever since. The value of Royal Dutch shares should thus equal 1.5 times Shell shares. It is, after all, the same company. Yet the shares do not follow this pattern. Suppose in mid 1983 you purchased Royal Dutch and sold Shell short when the latter was at a 10% premium (compared to the 60-40 split). This ought to result in pure profit, but six months later the premium was 25% and you might be wiped out. Or consider, that the premium was 30% in 1980 and did not reach zero till 1984. Puzzle indeed!

The basic problem was that the Russian crisis hurt many investment banks and hedge funds. Many had invested heavily in Russia. They were not stupid, they anticipated potential problems. Savvy investors assumed that a default on foreign denominated debt and a ruble collapse would be likely. So they purchased domestic debt with high interest rates and sold rubles forward. If the ruble collapsed they would benefit from the forward contracts. They also sold short Russian foreign debt. The assumption was that the Russian government would never default on domestic bonds without defaulting on foreign bonds. But the Russians surprised all the hedge funds and investment banks. They defaulted on domestic debt, not foreign debt, and they placed a moratorium on currency transactions of the domestic banks. The latter meant that the forward contracts could not be executed. All the hedges failed. Investment banks and hedge funds took very large hits. CSFB lost over $1 billion alone.

Investment banks used VAR models to measure their risk. Marked to market their risks increased, so they had to increase their capital or lower their risk. They chose to unwind some positions. Many of these banks and hedge funds had similar positions as LTCM, so as they tried to unwind their positions the amount of money bet on spreads narrowing declined. Relative investing is sort of like betting that a rubber band that is stretched out will return to its normal level. But if stretched too far the rubber band can collapse. That is what happened. Notice that to LTCM it seemed that the world was getting more irrational. The rubber band should not be getting more stretched out. But all the other players were leaving the bets because of their need to reduce risks. This left too little betting on mean reversion.

At this point the volatility bets that LTCM engaged in really crashed. They were selling volatility – mostly in terms of index options on the DAX and CAC-40. The historical long-term average for volatility for these indexes was roughly 15% per year, but the market was pricing it at 22%. So they sold. But volatility increased, and the value of their positions thus declined. They needed to reduce their positions, but this meant purchase the options on the market. But this only raised prices, further worsening volatility. Especially as there were few sellers. Then a peculiar feature of LTCM's dealings made everything worse. LTCM's counterparties were responsible for marking to market the long-dated options, as there were few quotes. Because LTCM had not paid haircuts, many counterparties now took advantage by marking to market at as high a rate as possible. Implied volatility now rose over 30% and towards 40% just due to LTCM's collateral problems.[15]

Although LTCM lost some large money on bets, most of it was lost on core arbitrage activity – about \$1.3 billion on volatility trading and another \$1.2 billion on fixed income arbitrage.

The Fed organized a bail-out for LTCM. Some consider this to be the worst policy – creating moral hazard. The Fed argues, however, that it did not lend to LTCM; rather it got LTCM's creditors to put more money into it. So it could be called a bail-in. LTCM had very complex trades, and it dominated many markets. To sell these positions quickly would result in a "fire sale." The losses that would ensue would erode the collateral on the debt owed to other investment banks. Hence, the investors had an interest in a bail-in.

The Fed helped coordinate the bail in because the free rider problem normally plagues such situations. After all, if 90% of the creditors put in some extra funds the last 10% may not matter. So why should I risk the bail in. But that reasoning applies to each investor. Hence, the Fed needed to coordinate such a bail in.

[15] Had LTCM defaulted the counterparties would take the options which were now priced at levels that implied huge volatility. They could have sold these options and certainly made a profit because volatility had never reached 40% before.