
A SURE PROFIT: THE ESSENCE OF ARBITRAGE

Arbitrage represents the holy grail of investing because it allows investors to invest no money, take no risk and walk away with sure profits. In other words, it is the ultimate money machine that investors hope to access.

In this chapter, we consider three types of arbitrage. The first is pure arbitrage, where, in fact, you risk nothing and earn more than the riskless rate. For pure arbitrage to be feasible, you need two assets with identical cashflows, different market values at the same point in time and a given point in time in the future at which the values have to converge. This type of arbitrage is most likely to occur in derivatives markets – options and futures- and in some parts of the bond market. The second is near arbitrage, where you have assets that have identical or almost identical cash flows, trading at different prices, but there is no guarantee that the prices will converge and there exist significant constraints on the investors forcing convergence. The third is speculative arbitrage, which may not really be arbitrage in the first place. Here, investors take advantage of what they see as mispriced and similar (though not identical) assets, buying the cheaper one and selling the more expensive one. If they are right, the difference should narrow over time, yielding profits. It is in this category that we consider hedge funds in their numerous forms. As we will see, the peril of this strategy is that the initial assessment of mispricing is usually based upon a view of the world that may or may not be justified.

Pure Arbitrage

The requirement that you have two assets with identical cashflows and different market prices makes pure arbitrage difficult to find in financial markets. First, identical assets are not common in the real world, especially if you are an equity investor. Second, assuming two identical assets exist, you have to wonder why financial markets would allow pricing differences to persist. If in addition, we add the constraint that there is a point in time where the market prices converge, it is not surprising that pure arbitrage is most likely to occur with derivative assets – options and futures and in fixed income markets, especially with default-free government bonds.

Futures Arbitrage

A futures contract is a contract to buy (and sell) a specified asset at a fixed price in a future time period. There are two parties to every futures contract - the seller of the contract, who agrees to deliver the asset at the specified time in the future, and the buyer of the contract, who agrees to pay a fixed price and take delivery of the asset. If the asset that

underlies the futures contract is traded and is not perishable, you can construct a pure arbitrage if the futures contract is mispriced. In this section, we will consider the potential for arbitrage first with storable commodities and then with financial assets and then look at whether such arbitrage is possible.

The Arbitrage Relationships

The basic arbitrage relationship can be derived fairly easily for futures contracts on any asset, by estimating the cashflows on two strategies that deliver the same end result – the ownership of the asset at a fixed price in the future. In the first strategy, you buy the futures contract, wait until the end of the contract period and buy the underlying asset at the futures price. In the second strategy, you borrow the money and buy the underlying asset today and store it for the period of the futures contract. In both strategies, you end up with the asset at the end of the period and are exposed to no price risk during the period – in the first, because you have locked in the futures price and in the second because you bought the asset at the start of the period. Consequently, you should expect the cost of setting up the two strategies to be exactly the same. Across different types of futures contracts, there are individual details that cause the final pricing relationship to vary – commodities have to be stored and create storage costs whereas stocks may pay a dividend while you are holding them.

a. Storable Commodities

The distinction between storable and perishable goods is that storable goods can be acquired today at the spot price and stored till the expiration of the futures contract, which is the practical equivalent of buying a futures contract and taking delivery at expiration. Since the two approaches provide the same result, in terms of having possession of the commodity at expiration, the futures contract, if priced right, should cost the same as a strategy of buying and storing the commodity. The two additional costs of the latter strategy are as follows.

(a) Since the commodity has to be acquired now, rather than at expiration, there is an added financing cost associated with borrowing the funds needed for the acquisition now.

$$\text{Added Interest Cost} = (\text{Spot price}) \left((1 + \text{Interest Rate})^{\text{Life of Futures contract}} - 1 \right)$$

(b) If there is a storage cost associated with storing the commodity until the expiration of the futures contract, this cost has to be reflected in the strategy as well. In addition, there may be a benefit to having physical ownership of the commodity. This benefit is called the convenience yield and will reduce the futures price. The net

storage cost is defined to be the difference between the total storage cost and the convenience yield.

If F is the futures contract price, S is the spot price, r is the annualized interest rate, t is the life of the futures contract and k is the net annual storage costs (as a percentage of the spot price) for the commodity, the two equivalent strategies and their costs can be written as follows.

Strategy 1: Buy the futures contract. Take delivery at expiration. Pay $\$F$.

Strategy 2: Borrow the spot price (S) of the commodity and buy the commodity. Pay the additional costs.

$$(a) \text{ Interest cost} = S((1+r)^t - 1)$$

$$(b) \text{ Cost of storage, net of convenience yield} = S k t$$

If the two strategies have the same costs,

$$\begin{aligned} F^* &= S((1+r)^t - 1) + S k t \\ &= S((1+r)^t + k t) \end{aligned}$$

This is the basic arbitrage relationship between futures and spot prices. Note that the futures price does not depend upon your expectations of what will happen to the spot price over time but on the spot price today. Any deviation from this arbitrage relationship should provide an opportunity for arbitrage, i.e., a strategy with no risk and no initial investment, and for positive profits. These arbitrage opportunities are described in Figure 11.1.

This arbitrage is based upon several assumptions. First, investors are assumed to borrow and lend at the same rate, which is the riskless rate. Second, when the futures contract is over priced, it is assumed that the seller of the futures contract (the arbitrageur) can sell short on the commodity and that he can recover, from the owner of the commodity, the storage costs that are saved as a consequence. To the extent that these assumptions are unrealistic, the bounds on prices within which arbitrage is not feasible expand. Assume, for instance, that the rate of borrowing is r_b and the rate of lending is r_a , and that short seller cannot recover any of the saved storage costs and has to pay a transactions cost of t_s . The futures price will then fall within a bound.

$$(S - t_s)(1 + r_a)^t < F^* < S((1 + r_b)^t + k t)$$

If the futures price falls outside this bound, there is a possibility of arbitrage and this is illustrated in Figure 11.2.

Figure 11.1: Storable Commodity Futures: Pricing and Arbitrage

$$F^* = S((1+r)^t + kt)$$

If $F > F^*$

If $F < F^*$

Time	Action	Cashflows	Action	Cashflows
Now:	1. Sell futures contract 2. Borrow spot price at riskfree r 3. Buy spot commodity	0 S -S	1. Buy futures contract 2. Sell short on commodity 3. Lend money at riskfree rate	0 S -S
At t:	1. Collect commodity; Pay storage cost. 2. Deliver on futures contract 3. Pay back loan	-Skt F -S(1+r) ^t	1. Collect on loan 2. Take delivery of futures contract 3. Return borrowed commodity; Collect storage costs	S(1+r) ^t -F +Skt
Net Cash Flow=		F-S((1+r)^t - kt) > 0		S((1+r)^t + kt) - F > 0

Key inputs:

F* = Theoretical futures price

F = Actual futures price

S = Spot price of commodity

Key assumptions

1. The investor can lend and borrow at the riskless rate.
2. There are no transactions costs associated with buying or selling short the commodity.
3. The short seller can collect all storage costs saved because of the short selling.

r = Riskless rate of interest (annualized)


t = Time to expiration on the futures contract

k = Annualized carrying cost, net of convenience yield (as % of spot price)

Figure 11.2: Storable Commodity Futures: Pricing and Arbitrage with Modified Assumptions

Modified Assumptions

1. Investor can borrow at r_b ($r_b > r$) and lend at r_a ($r_a < r$).
2. The transactions cost associated with selling short is t_s (where t_s is the dollar transactions cost).
3. The short seller does not collect any of the storage costs saved by the short selling.



Time	Action	Cashflows	Action	Cashflows
Now:	1. Sell futures contract 2. Borrow spot price at r_b 3. Buy spot commodity	0 S -S	1. Buy futures contract 2. Sell short on commodity 3. Lend money at r_a	0 S - t_s -(S - t_s)
At t:	1. Collect commodity from storage 2. Delivery on futures contract 3. Pay back loan	-Skt F -S(1+r _b) ^t	1. Collect on loan 2. Take delivery of futures contract 3. Return borrowed commodity; Collect storage costs	(S- t_s)(1+r _a) ^t -F 0
Net Cash Flow=		F-S((1+r_b)^t - kt) > 0		(S-t_s) (1+r_a)^t - F > 0

F_h = Upper limit for arbitrage bound on futures prices

F_l = Lower limit for arbitrage bound on futures prices

b. Stock Index Futures

Futures on stock indices have become an important and growing part of most financial markets. Today, you can buy or sell futures on the Dow Jones, the S&P 500, the NASDAQ and the Value Line indices, as well as many indices in other countries. An index future entitles the buyer to any appreciation in the index over and above the index futures price and the seller to any depreciation in the index from the same benchmark. To evaluate the arbitrage pricing of an index future, consider the following strategies.

Strategy 1: Sell short on the stocks in the index for the duration of the index futures contract. Invest the proceeds at the riskless rate. This strategy requires that the owners of the stocks that are sold short be compensated for the dividends they would have received on the stocks.

Strategy 2: Sell the index futures contract.

Both strategies require the same initial investment, have the same risk and should provide the same proceeds. Again, if S is the spot price of the index, F is the futures price, y is the annualized dividend yield on the stock and r is the riskless rate, the arbitrage relationship can be written as follows:

$$F^* = S (1 + r - y)^t$$

If the futures price deviates from this arbitrage price, there should be an opportunity from arbitrage. This is illustrated in Figure 11.3.

This arbitrage is also conditioned on several assumptions. First, we assume that investors can lend and borrow at the riskless rate. Second, we ignore transactions costs on both buying stock and selling short on stocks. Third, we assume that the dividends paid on the stocks in the index are known with certainty at the start of the period. If these assumptions are unrealistic, the index futures arbitrage will be feasible only if prices fall outside a band, the size of which will depend upon the seriousness of the violations in the assumptions.

Assume that investors can borrow money at r_b and lend money at r_a and that the transactions costs of buying stock is t_c and selling short is t_s . The band within which the futures price must stay can be written as:


$$(S - t_s)(1 + r_a - y) < F^* < (S + t_c)(1 + r_b - y)^t$$

The arbitrage that is possible if the futures price strays outside this band is illustrated in Figure 11.4.

In practice, one of the issues that you have to factor in is the seasonality of dividends since the dividends paid by stocks tend to be higher in some months than others. Figure

11.5 graphs out dividends paid as a percent of the S&P 500 index on U.S. stocks in 2000 by month of the year.

Figure 11.5: Dividend Yields by Month of Year- 2000



Thus, dividend yields seem to peak in February, May, August and November. An index future coming due in these months is much more likely to be affected by dividend yields especially as maturity draws closer.

Figure 11.3: Stock Index Futures: Pricing and Arbitrage

$$F^* = S (1+r-y)^t$$

If $F < F^*$

Transaction	Cashflows	Action	Cashflows
Buy futures contract	0	1. Buy futures contract	0
Short price of index at riskfree rate	S	2. Sell short stocks in the index	S
Borrow money	$-S$	3. Lend money at riskfree rate	$-S$
Lends on stocks	$S((1+y)^t-1)$	1. Collect on loan	$S(1+r)^t$
Buy futures contract	F	2. Take delivery of futures contract	$-F$
Return money	$-S(1+r)^t$	3. Return borrowed stocks; Pay foregone dividends	$-S((1+y)^t-1)$
	$F - S(1+r-y)^t > 0$		$S(1+r)^t - F > 0$


r = Riskless rate of interest (annualized)
 t = Time to expiration on the futures contract
 y = Dividend yield over lifetime of futures contract as % of current index level

and borrow at the riskless rate.
 No costs associated with buying or selling short stocks.
 Arbitrage with certainty.

Figure 11.4: Stock Index Futures: Pricing and Arbitrage with modified assumptions
Modified Assumptions

borrow (where $r_b > r$) and lend at r_a ($r_a < r$).

associated with selling short is t_s (where t_s is the dollar transactions cost) and the transactions cost associated with



	Cashflows	Action	Cashflows
Buy futures contract	0	1. Buy futures contract	0
Borrow price at r_b	$S + t_c$	2. Sell short stocks in the index	$S - t_s$
Lend in the index	$-S - t_c$	3. Lend money at r_a	$-(S - t_s)$
Borrow on stocks	$S((1+y)^t - 1)$	1. Collect on loan	$(S - t_s)(1 + r_a)^t$
Buy futures contract	F	2. Take delivery of futures contract	$-F$
Lend in	$-(S + t_c)(1 + r_b)^t$	3. Return borrowed stocks; Pay foregone dividends	$-S((1+y)^t - 1)$
	$F - (S + t_c) (1 + r_b - y)^t > 0$		$(S - t_s) (1 + r_a - y)^t - F > 0$

Upper bound on futures prices

F_L^* = Lower limit for arbitrage bound on futures prices

c. Treasury Bond Futures

The treasury bond futures traded on the Chicago Board of Trade require the delivery of any government bond with a maturity greater than fifteen years, with a no-call feature for at least the first fifteen years. Since bonds of different maturities and coupons will have different prices, the CBOT has a procedure for adjusting the price of the bond for its characteristics. The conversion factor itself is fairly simple to compute and is based upon the value of the bond on the first day of the delivery month, with the assumption that the interest rate for all maturities equals 8% per annum (with semi-annual compounding). For instance, you can compute the conversion factor for a 9% coupon bond with 18 years to maturity. Working in terms of a \$100 face value of the bond, the value of the bond can be written as follows, using the interest rate of 8%.

$$\text{PV of Bond} = \sum_{t=0.5}^{t=20} \frac{4.50}{(1.08)^t} + \frac{100}{(1.08)^{20}} = \$111.55$$

The conversion factor for this bond is 111.55. Generally, the conversion factor will increase as the coupon rate increases and with the maturity of the delivered bond.

This feature of treasury bond futures, i.e., that any one of a menu of treasury bonds can be delivered to fulfill the obligation on the bond, provides an advantage to the seller of the futures contract. Naturally, the cheapest bond on the menu, after adjusting for the conversion factor, will be delivered. This *delivery option* has to be priced into the futures contract. There is an additional option embedded in treasury bond futures contracts that arises from the fact that the T.Bond futures market closes at 2 p.m., whereas the bonds themselves continue trading until 4 p.m. The seller does not have to notify the clearing house until 8 p.m. about his intention to deliver. If bond prices decline after 2 p.m., the seller can notify the clearing house of his intention to deliver the cheapest bond that day. If not, the seller can wait for the next day. This option is called the *wild card option*.

The valuation of a treasury bond futures contract follows the same lines as the valuation of a stock index future, with the coupons of the treasury bond replacing the dividend yield of the stock index. The theoretical value of a futures contract should be –

$$F^* = (S - PVC)(1 + r)^t$$

where,

F* = Theoretical futures price for Treasury Bond futures contract

S = Spot price of Treasury bond

PVC = Present Value of coupons during life of futures contract

r = Riskfree interest rate corresponding to futures life

t = Life of the futures contract

If the futures price deviates from this theoretical price, there should be the opportunity for arbitrage. These arbitrage opportunities are illustrated in Figure 11.6.

This valuation ignores the two options described above - the option to deliver the cheapest-to-deliver bond and the option to have a wild card play. These give an advantage to the seller of the futures contract and should be priced into the futures contract. One way to build this into the valuation is to use the cheapest deliverable bond to calculate both the current spot price and the present value of the coupons. Once the futures price is estimated, it can be divided by the conversion factor to arrive at the standardized futures price.

d. Currency Futures

In a currency futures contract, you enter into a contract to buy a foreign currency at a price fixed today. To see how spot and futures currency prices are related, note that holding the foreign currency enables the investor to earn the risk-free interest rate (R_f) prevailing in that country while the domestic currency earn the domestic riskfree rate (R_d). Since investors can buy currency at spot rates and assuming that there are no restrictions on investing at the riskfree rate, we can derive the relationship between the spot and futures prices. *Interest rate parity* relates the differential between futures and spot prices to interest rates in the domestic and foreign market.

$$\frac{\text{Futures Price}_{d,f}}{\text{Spot Price}_{d,f}} = \frac{(1 + R_d)}{(1 + R_f)}$$

where $\text{Futures Price}_{d,f}$ is the number of units of the domestic currency that will be received for a unit of the foreign currency in a forward contract and $\text{Spot Price}_{d,f}$ is the number of units of the domestic currency that will be received for a unit of the same foreign currency in a spot contract. For instance, assume that the one-year interest rate in the United States is 5% and the one-year interest rate in Germany is 4%. Furthermore, assume that the spot exchange rate is \$0.65 per Deutsche Mark. The one-year futures price, based upon interest rate parity, should be as follows:

$$\frac{\text{Futures Price}_{d,f}}{\$ 0.65} = \frac{(1.05)}{(1.04)}$$

This results in a futures price of \$0.65625 per Deutsche Mark.

Why does this have to be the futures price? If the futures price were greater than \$0.65625, say \$0.67, an investor could take advantage of the mispricing by selling the futures contract, completely hedging against risk and ending up with a return greater than the riskfree rate. The actions the investor would need to take are summarized in Table 11.1, with the cash flows associated with each action in brackets next to the action.

Table 11.1: Arbitrage when currency futures contracts are mispriced

<i>Forward Rate Mispricing</i>	<i>Actions to take today</i>	<i>Actions at expiration of futures contract</i>
If futures price > \$0.65625 e.g. \$0.67	1. Sell a futures contract at \$0.67 per Deutsche Mark. (\$0.00) \$ 0.00 2. Borrow the spot price in the U.S. domestic markets @ 5%.	1. Collect on Deutsche Mark investment. (+1.04 DM) 1.04 DM 2. Convert into dollars at futures price. (-1.04 DM/

	(+\$0.65) + \$ 0.65 3. Convert the dollars into Deutsche Marks at spot price. (-\$0.65/+1 DM) - \$ 0.65/+ 1 DM 4. Invest Deutsche Marks in the German market @ 4%. (-1 DM) - 1 DM	+\$0.6968) -1.04 DM to + \$ 0.6968 3. Repay dollar borrowing with interest. (-\$0.6825) Profit = \$0.6968 - \$0.6825 = \$ 0.0143
If futures price < \$0.65625 e.g. \$0.64	1. Buy a futures price at \$0.64 per Deutsche Mark. (\$0.00) \$ 0.00 2. Borrow the spot rate in the German market @4%. (+1 DM) + 1 DM 3. Convert the Deutsche Marks into Dollars at spot rate. (-1 DM/+\$0.65) - 1 DM/ \$ 0.65 4. Invest dollars in the U.S. market @ 5%. (-\$0.65) - \$ 0.65	1. Collect on Dollar investment. (+\$0.6825) \$ 0.6825 2. Convert into dollars at futures price. (-\$0.6825/1.0664 DM) - \$ 0.6825 /+1.0664 DM 3. Repay DM borrowing with interest. (1.04 DM) Profit = 1.0664-1.04 = 0.0264 DM - 1.04 DM + 0.0264 DM

The first arbitrage of Table 11.1 results in a riskless profit of \$0.0143, with no initial investment. The process of arbitrage will push down futures price towards the equilibrium price.

If the futures price were lower than \$0.65625, the actions would be reversed, with the same final conclusion. Investors would be able to take no risk, invest no money and still end up with a positive cash flow at expiration. In the second arbitrage of Table 34.3, we lay out the actions that would lead to a riskless profit of .0264 DM.

Special Features of Futures Markets

There are two special features of futures markets that can make arbitrage tricky. The first is the existence of margins. While we assumed, when constructing the arbitrage, that buying and selling futures contracts would create no cashflows at the time of the transaction, you would have to put up a portion of the futures contract price (about 5-10%) as a margin in the real world. To compound the problem, this margin is recomputed every day based upon futures prices that day – this process is called marking to market - and you may be

required to come up with more margin if the price moves against you (down, if you are a buyer and up, if you are a seller). If this margin call is not met, your position can be liquidated and you may never get to see your arbitrage profits.

The second is that the futures exchanges generally impose ‘price movement limits’ on most futures contracts. For instance, the daily price movement limit on orange juice futures contract on the New York Board of Trade is 5 cents per pound or \$750 per contract (which covers 15,000 pounds). If the price of the contract drops or increases by the amount of the price limit, trading is generally suspended for the day, though the exchange reserves the discretion to reopen trading in the contract later in the day. The rationale for introducing price limits is to prevent panic buying and selling on an asset, based upon faulty information or rumors, and to prevent overreaction to real information. By allowing investors more time to react to extreme information, it is argued, the price reaction will be more rational and reasoned. In the process, though, you can create a disconnect between the spot markets, where no price limits exist, and futures markets, where they do.



Price Limits and Contract specifications:
Take a look at the price limits and contract specifications on widely traded futures contracts.

Feasibility of Arbitrage and Potential for Success

If futures arbitrage is so simple, you may ask, how in a reasonably efficient market would arbitrage opportunities even exist? In the commodity futures market, for instance, Garbade and Silber (1983) find little evidence of arbitrage opportunities and their findings are echoed in other studies. In the financial futures markets, there is evidence that indicates that arbitrage is indeed feasible but only to a sub-set of investors. Differences in transactions cost seem to explain most of the differences. Large institutional investors, with close to zero transactions costs and instantaneous access to both the underlying asset and futures markets may be able to find arbitrage opportunities, where you and I as individual investors would not. In addition, these investors are also more likely to meet the requirements for arbitrage – being able to borrow at rates close to the riskless rate and sell short on the underlying asset.

Note, though, that the returns are small¹ even to these large investors and that arbitrage will not be a reliable source of profits, unless you can establish a competitive advantage on one of three dimensions. First, you can try to establish a transactions cost

¹ A study of 835 index arbitrage trades on the S&P 500 futures contracts estimated that the average gross profit from such trades was only 0.30%.

advantage over other investors, which will be difficult to do since you are competing with other large institutional investors. Second, you may be able to develop an information advantage over other investors by having access to information earlier than others. Again, though much of the information is pricing information and is public. Third, you may find a quirk in the data or pricing of a particular futures contract before others learn about it. The arbitrage possibilities seem to be greater when futures contracts are first introduced on an asset, since investors take time to understand the details of futures pricing. For instance, it took investors a while to learn to incorporate the effect of uneven dividends into stock index futures and the wild card option into treasury bond futures. Presumably, investors who learnt faster than the market would have been able to take advantage of the mispricing of futures contracts in these early periods and earn excess returns.

Options Arbitrage

As derivative securities, options differ from futures in a very important respect. They represent rights rather than obligations – calls gives you the right to buy and puts gives you the right to sell. Consequently, a key feature of options is that the losses on an option position are limited to what you paid for the option, if you are a buyer. Since there is usually an underlying asset that is traded, you can, as with futures contracts, construct positions that essentially are riskfree by combining options with the underlying asset.

Exercise Arbitrage

The easiest arbitrage opportunities in the option market exist when options violate simple pricing bounds. No option, for instance, should sell for less than its exercise value. With a call option: Value of call > Value of Underlying Asset – Strike Price
With a put option: Value of put > Strike Price – Value of Underlying Asset
For instance, a call option with a strike price of \$ 30 on a stock that is currently trading at \$ 40 should never sell for less than \$ 10. If it did, you could make an immediate profit by buying the call for less than \$ 10 and exercising right away to make \$ 10.

In fact, you can tighten these bounds for call options, if you are willing to create a portfolio of the underlying asset and the option and hold it through the option's expiration. The bounds then become:

With a call option: Value of call > Value of Underlying Asset – Present value of Strike Price

With a put option: Value of put > Present value of Strike Price – Value of Underlying Asset

Too see why, consider the call option in the previous example. Assume that you have one year to expiration and that the riskless interest rate is 10%.

Present value of Strike Price = $\$ 30/1.10 = \27.27

Lower Bound on call value = $\$ 40 - \$27.27 = \$12.73$

The call has to trade for more than \$12.73. What would happen if it traded for less, say \$ 12? You would buy the call for \$ 12, sell short a share of stock for \$ 40 and invest the net proceeds of \$ 28 ($\$40 - 12$) at the riskless rate of 10%. Consider what happens a year from now:

If the stock price $>$ strike price (\$ 30): You first collect the proceeds from the riskless investment ($\$28(1.10) = \30.80), exercise the option (buy the share at \$ 30) and cover your short sale. You will then get to keep the difference of \$0.80.

If the stock price $<$ strike price (\$ 30): You collect the proceeds from the riskless investment (\$30.80), buy a share in the open market for the prevailing price then (which is less than \$30) and keep the difference.


In other words, you invest nothing today and are guaranteed a positive payoff in the future. You could construct a similar example with puts.

The arbitrage bounds work best for non-dividend paying stocks and for options that can be exercised only at expiration (European options). Most options in the real world can be exercised prior to expiration (American options) and are on stocks that pay dividends. Even with these options, though, you should not see short term options trading violating these bounds by large margins, partly because exercise is so rare even with listed American options and dividends tend to be small. As options become long term and dividends become larger and more uncertain, you may very well find options that violate these pricing bounds, but you may not be able to profit off them.

Replicating Portfolio

One of the key insights that Fischer Black and Myron Scholes had about options in the 1970s that revolutionized option pricing was that a portfolio composed of the underlying asset and the riskless asset could be constructed to have exactly the same cash flows as a call or put option. This portfolio is called the replicating portfolio. In fact, Black and Scholes used the arbitrage argument to derive their option pricing model by noting that since the replicating portfolio and the traded option have the same cash flows, they would have to sell at the same price.

To understand how replication works, let us consider a very simple model for stock prices where prices can jump to one of two points in each time period. This model, which is called a binomial model, allows us to model the replicating portfolio fairly easily. In the figure below, we have the binomial distribution of a stock, currently trading at \$ 50 for the next two time periods. Note that in two time periods, this stock can be trading for as much as \$ 100 or as little as \$ 25. Assume that the objective is to value a call with a strike price of \$ 50, which is expected to expire in two time periods:




Now assume that the interest rate is 11%. In addition, define

Δ = Number of shares in the replicating portfolio

B = Dollars of borrowing in replicating portfolio

The objective is to combine Δ shares of stock and B dollars of borrowing to replicate the cash flows from the call with a strike price of 50. Since we know the cashflows on the option with certainty at expiration, it is best to start with the last period and work back through the binomial tree.

Step 1: Start with the end nodes and work backwards. Note that the call option expires at $t=2$, and the gross payoff on the option will be the difference between the stock price and the exercise price, if the stock price $>$ exercise price, and zero, if the stock price $<$ exercise price.



The objective is to construct a portfolio of Δ shares of stock and B in borrowing at $t=1$, when the stock price is \$ 70, that will have the same cashflows at $t=2$ as the call option with a strike price of 50. Consider what the portfolio will generate in cash flows under each of the two stock price scenarios, after you pay back the borrowing with interest (11% per period) and set the cash flows equal to the cash flows you would have received on the call.


$$\text{If stock price} = \$ 100: \quad \text{Portfolio Value} = 100 \Delta - 1.11 B = 50$$

$$\text{If stock price} = \$ 50: \quad \text{Portfolio Value} = 50 \Delta - 1.11 B = 0$$

We can solve for both the number of shares of stock you will need to buy (Δ) and the amount you will need to borrow (\$ 45) at $t=1$. Thus, if the stock price is \$70 at $t=1$, borrowing \$45 and buying one share of the stock will give the same cash flows as buying the call. To prevent arbitrage, the value of the call at $t=1$, if the stock price is \$70, has to be equal to the cost (to you as an investor) of setting up the replicating position:


$$\text{Value of Call} = \text{Cost of Replicating Position} = 70\Delta - B = (70)(1) - 45 = 25$$

Considering the other leg of the binomial tree at $t=1$,



If the stock price is 35 at $t=1$, then the call is worth nothing.

Step 2: Now that we know how much the call will be worth at $t=1$ (\$25 if the stock price goes to \$ 70 and \$0 if it goes down to \$ 35), we can move backwards to the earlier time period and create a replicating portfolio that will provide the values that the option will provide.



In other words, borrowing \$22.5 and buying $5/7$ of a share today will provide the same cash flows as a call with a strike price of \$50. The value of the call therefore has to be the same as the cost of creating this position.

Value of Call = Cost of replicating position =

$$\left(\frac{5}{7}\right)(\text{Current Stock Price}) - 22.5 = \left(\frac{5}{7}\right)(50) - 22.5 = 13.21$$

Consider for the moment the possibilities for arbitrage if the call traded at less than \$13.21, say \$ 13.00. You would buy the call for \$13.00 and sell the replicating portfolio for \$13.21 and claim the difference of \$0.21. Since the cashflows on the two positions are identical, you would be exposed to no risk and make a certain profit. If the call trade for more than \$13.21, say \$13.50, you would buy the replicating portfolio, sell the call and claim the \$0.29 difference. Again, you would not have been exposed to any risk.

You could construct a similar example using puts. The replicating portfolio in that case would be created by selling short on the underlying stock and lending the money at the riskless rate. Again, if puts are priced at a value different from the replicating portfolio, you could capture the difference and be exposed to no risk.

What are the assumptions that underlie this arbitrage? The first is that both the traded asset and the option are traded and that you can trade simultaneously in both markets, thus locking in your profits. The second is that there are no (or at least very low transactions costs). If transactions costs are large, prices will have to move outside the band created by these costs for arbitrage to be feasible. The third is that you can borrow at the riskless rate and sell short, if necessary. If you cannot, arbitrage may no longer be feasible.

Arbitrage across options

When you have multiple options listed on the same asset, you may be able to take advantage of relative mispricing – how one option is priced relative to another - and lock in

riskless profits. We will look first at the pricing of calls relative to puts and then consider how options with different exercise prices and maturities should be priced, relative to each other.

Put-Call Parity

When you have a put and a call option with the same exercise price and the same maturity, you can create a riskless position by selling the call, buying the put and buying the underlying asset at the same time. To see why, consider selling a call and buying a put with exercise price K and expiration date t , and simultaneously buying the underlying asset at the current price S . The payoff from this position is riskless and always yields a cashflow of K at expiration t . To see this, assume that the stock price at expiration is S^* . The payoff on each of the positions in the portfolio can be written as follows:

Position	Payoffs at t if $S^* > K$	Payoffs at t if $S^* < K$
Sell call	$-(S^* - K)$	0
Buy put	0	$K - S^*$
Buy stock	S^*	S^*
<i>Total</i>	K	K

Since this position yields K with certainty, the cost of creating this position must be equal to the present value of K at the riskless rate ($K e^{-rt}$).

$$S + P - C = K e^{-rt}$$

$$C - P = S - K e^{-rt}$$

This relationship between put and call prices is called put call parity. If it is violated, you have arbitrage.

If $C - P > S - K e^{-rt}$, you would sell the call, buy the put and buy the stock. You would earn more than the riskless rate on a riskless investment.

If $C - P < S - K e^{-rt}$, you would buy the call, sell the put and sell short the stock. You would then invest the proceeds at the riskless rate and end up with a riskless profit at maturity.

Note that violations of put call parity create arbitrage opportunities only for options that can be exercised only at maturity (European options) and may not hold if options can be exercised early (American options).

Does put-call parity hold up in practice or are there arbitrage opportunities? One study² examined option pricing data from the Chicago Board of Options from 1977 to 1978 and found potential arbitrage opportunities in a few cases. However, the arbitrage opportunities were small and persisted only for short periods. Furthermore, the options examined were American options, where arbitrage may not be feasible even if put-call parity is violated. A more recent study by Kamara and Miller of options on the S&P 500 (which are European options) between 1986 and 1989 finds fewer violations of put-call parity and the deviations tend to be small, even when there are violations.

Mispricing across Strike Prices and Maturities

A spread is a combination of two or more options of the same type (call or put) on the same underlying asset. You can combine two options with the same maturity but different exercise prices (bull and bear spreads), two options with the same strike price but different maturities (calendar spreads), two options with different exercise prices and maturities (diagonal spreads) and more than two options (butterfly spreads). You may be able to use spreads to take advantage of relative mispricing of options on the same underlying stock.

Strike Prices: A call with a lower strike price should never sell for less than a call with a higher strike price, assuming that they both have the same maturity. If it did, you could buy the lower strike price call and sell the higher strike price call, and lock in a riskless profit. Similarly, a put with a lower strike price should never sell for more than a put with a higher strike price and the same maturity. If it did, you could buy the higher strike price put, sell the lower strike price put and make an arbitrage profit.

Maturity: A call (put) with a shorter time to expiration should never sell for more than a call (put) with the same strike price with a long time to expiration. If it did, you would buy the call (put) with the shorter maturity and sell the call (put) with the longer maturity (i.e, create a calendar spread) and lock in a profit today. When the first call expires, you will either exercise the second call (and have no cashflows) or sell it (and make a further profit).

Even a casual perusal of the option prices listed in the newspaper each day should make it clear that it is very unlikely that pricing violations that are this egregious will exist in a liquid options market.

² See Klemkosky, R.C. and B.G. Resnick, 1979, Put-Call Parity and Market Efficiency, *Journal of Finance*, v 34, pg 1141-1155.

Fixed Income Arbitrage

Fixed income securities lend themselves to arbitrage more easily than equity because they have finite lives and fixed cash flows. This is especially so, when you have default free bonds, where the fixed cash flows are also guaranteed. Consider one very simple example. You could replicate a 10-year treasury bond's cash flows by buying zero-coupon treasuries with expirations matching those of the coupon payment dates on the treasury bond. For instance, if you invest \$ 100 million in a ten-year treasury bond with an 8% coupon rate, you can expect to get cashflows of \$ 4 million every 6 months for the next 10 years and \$ 100 million at the end of the tenth year. You could have obtained exactly the same cashflows by buying zero-coupon treasuries with face values of \$ 4 million, expiring every 6 months for the next ten years, and an additional 10-year zero coupon bond with a face value of \$ 100 million. Since the cashflows are identical, you would expect the two positions to trade for the same price. If they do not trade at the same price, you would buy the cheaper position and sell the more expensive one, locking in the profit today and having no cashflow or risk exposure in the future.

With corporate bonds, you have the extra component of default risk. Since no two firms are exactly identical when it comes to default risk, you may be exposed to some risk if you are using corporate bonds issued by different entities. In fact, two bonds issued by the same entity may not be equivalent because of differences in how they are secured and structured. There are some arbitrageurs who argue that bond ratings are a good proxy for default risk, and that buying one AA rated bond and selling another should be riskless, but bond ratings are not perfect proxies for default risk. In fact, you see arbitrage attempted on a wide variety of securities with promised cashflows, such as mortgage backed bonds. While you can hedge away much of the cashflow risk, the nature of the cashflow claims will still leave you exposed to some risk. With mortgage backed bonds, for instance, the unpredictability of prepayments by homeowners has exposed many "riskless" positions to risk.

Is there any evidence that investors are able to find treasuries mispriced enough to generate arbitrage profits? Grinblatt and Longstaff, in an assessment of the treasury strips program – a program allowing investors to break up a treasury bond and sell its individual cash flows – note that there are potential arbitrage opportunities in these markets but find little evidence of trading driven by these opportunities. A study by Balbas and Lopez of the Spanish bond market may shed some light on this question. Examining default free and option free bonds in the Spanish market between 1994 and 1998, they conclude that there were arbitrage opportunities especially surrounding innovations in financial markets. We would extend their findings to argue that opportunities for arbitrage with fixed income

securities are probably greatest when new types of bonds are introduced – mortgage backed securities in the early 1980s, inflation-indexed treasuries in the late 1990s and the treasury strips program in the late 1980s. As investors become more informed about these bonds and how they should be priced, arbitrage opportunities seem to subside.

Determinants of Success

The nature of pure arbitrage – two identical assets that are priced differently – makes it likely that it will be short lived. In other words, in a market where investors are on the look out for riskless profits, it is very likely that small pricing differences will be exploited quickly, and in the process, disappear. Consequently, the first two requirements for success at pure arbitrage are access to real-time prices and instantaneous execution. It is also very likely that the pricing differences in pure arbitrage will be very small – often a few hundredths of a percent. To make pure arbitrage feasible, therefore, you can add two more conditions. The first is access to substantial debt at favorable interest rates, since it can magnify the small pricing differences. Note that many of the arbitrage positions require you to be able to borrow at the riskless rate. The second is economies of scale, with transactions amounting to millions of dollars rather than thousands. Institutions that are successful at pure arbitrage often are able to borrow several times their equity at the riskless rate to fund arbitrage transactions, using the guaranteed profits on the transaction as collateral.

With these requirements, it is not surprising that individual investors have generally not been able to succeed at pure arbitrage. Even among institutions, pure arbitrage is feasible only to a few, and even to those, it is a transient source of profits in two senses. First, you cannot count on the existence of pure arbitrage opportunities in the future, since it requires that markets repeat their errors over time. Second, the very fact that some institutions make profits from arbitrage attracts other institutions into the market, reducing the likelihood of future arbitrage profits. To succeed in the long term with arbitrage, you will need to be constantly on the lookout for new arbitrage opportunities.

Near Arbitrage

In near arbitrage, you either have two assets that are very similar but not identical, which are priced differently, or identical assets that are mispriced, but with no guaranteed price convergence. No matter how sophisticated your trading strategies may be in these scenarios, your positions will no longer be riskless.

Same Security, Multiple Markets

In today's global markets, there are a number of stocks that are listed on more than one market. If you can buy the same stock at one price in one market and simultaneously

sell it at a higher price in another market, you can lock in a riskless profit. As we will see in this section, things are seldom this simple.

Dual and Multiple Listings

Many large companies such as Royal Dutch, General Electric and Microsoft trade on multiple markets on different continents. Since there are time periods during the day when there is trading occurring on more than one market on the same stock, it is conceivable (though not likely) that you could buy the stock for one price in one market and sell the same stock at the same time for a different (and higher price) in another market. The stock will trade in different currencies, and for this to be a riskless transaction, the trades have to at precisely the same time and you have to eliminate any exchange rate risk by converting the foreign currency proceeds into the domestic currency instantaneously. Your trade profits will also have to cover the different bid-ask spreads in the two markets and transactions costs in each.

There are some exceptional cases, where the same stock trades in different markets in one country. Swaicki and Hric examine 84 Czech stocks that trade on the two Czech exchanges – the Prague Stock Exchange (PSE) and the Registration Places System (RMS)- and find that prices adjust slowly across the two markets, and that arbitrage opportunities exist (at least on paper) –the prices in the two markets differ by about 2%. These arbitrage opportunities seem to increase for less liquid stocks. While the authors consider transactions cost, they do not consider the price impact that trading itself would have on these stocks and whether the arbitrage profits would survive the trading.

Depository Receipts

Many Latin American and European companies have American Depository Receipts (ADRs) listed on the U.S. market. These depository receipts create a claim equivalent to the one you would have had if you had bought shares in the local market and should therefore trade at a price consistent with the local shares. What makes them different and potentially riskier than the stocks with dual listings is that ADRs are not always directly comparable to the common shares traded locally – one ADR on Telmex, the Mexican telecommunications company, is convertible into 20 Telmex shares. In addition, converting an ADR into local shares can be both costly and time consuming. In some cases, there can be differences in voting rights as well. In spite of these constraints, you would expect the price of an ADR to closely track the price of the



Most widely traded ADRs: Take a look at the 50 most widely traded ADRs on the U.S. market.

shares in the local market, albeit with a currency overlay, since ADRs are denominated in dollars.


In a study conducted in 2000 that looks at the link between ADRs and local shares, Kin, Szakmary and Mathur conclude that about 60 to 70% of the variation in ADR prices can be attributed to movements in the underlying share prices and that ADRs overreact to the U.S. market and under react to exchange rates and the underlying stock. However, they also conclude that investors cannot take advantage of the pricing errors in ADRs because convergence does not occur quickly or in predictable ways. With a longer time horizon and/or the capacity to convert ADRs into local shares, though, you should be able to take advantage of significant pricing differences.

Closed End Funds

In a conventional mutual fund, the number of shares increases and decreases as money comes in and leaves the fund, and each share is priced at net asset value – the market value of the securities of the fund divided by the number of shares. Closed end mutual funds differ from other mutual funds in one very important respect. They have a fixed number of shares that trade in the market like other publicly traded companies, and the market price can be different from the net asset value.

In both the United States and the United Kingdom, closed end mutual funds have shared a very strange characteristic. When they are created, the price is usually set at a premium on the net asset value per share. As closed end funds trade, though, the market price tends to drop below the net asset value and stay there. Figure 11.7 provides the distribution of price to net asset value for all closed end funds in the United States in early 2002.


Figure 11.7: Discounts/Premiums on Closed End Funds- June 2002



Note that almost 70% of the closed end funds trade at a discount to net asset value and that the median discount is about 5%.

So what, you might ask? Lots of firms trade at less than the estimated market value of their assets. That might be true, but closed end funds are unique for two reasons. First, the assets are all traded stocks and the market value is therefore known at any point in time and not an estimate. Second, liquidating a closed end fund's assets should not be difficult to do, since you are selling securities to the market. Thus, liquidation should neither be costly nor time consuming. Given these two conditions, you may wonder why you should not buy closed end funds that trade at a discount and either liquidate them yourself or hope that some one else will liquidate them. Alternatively, you may be able to push a closed-end fund to open-end and see prices converge on net asset value. Figure 11.8 reports on the performance of closed-end funds when they open end, based upon a study of 94 UK closed-end funds that open ended:

Figure 11.8: Relative Discount on Closed End Funds that Open End



Note that as you get closer to the open-ending date (day 0), the discount becomes smaller relative to the average closed-end fund. For instance, the discount goes from being on par with the discount on other funds to being about 10% lower than the typical closed-end fund.


So what is the catch? In practice, taking over a closed-end fund while paying less than net asset value for its shares seems to be very difficult to do for several reasons- some related to corporate governance and some related to market liquidity. The potential profit is also narrowed by the mispricing of illiquid assets in closed end fund portfolios (leading to an overstatement of the NAV) and tax liabilities from liquidating securities. There have been a few cases of closed end funds being liquidated, but they remain the exception. What about the strategy of buying discounted funds and hoping that the discount disappears? This strategy is clearly not riskless but it does offer some promise. In one of the first studies of this strategy in 1978, Thompson studied closed end funds from 1940 to 1975 and reported that you could earn an annualized excess return of 4% from buying discounted funds. A study in 1986 by Anderson reports excess returns from a strategy of buying closed end funds whose discounts had widened and selling funds whose discounts had narrowed – a contrarian strategy applied to closed end



Most discounted closed end funds: Take a look at the 50 closed end funds with the largest discounts.

funds. In 1995, Pontiff reported that closed end funds with a discount of 20% or higher earn about 6% more than other closed end funds. This, as well as studies in the UK, seem to indicate a strong mean reversion component to discounts at closed funds. Figure 11.9, which is from a study of the discounts on closed end funds in the UK, tracks relative discounts on the most discounted and least discounted funds over time:

Figure 11.9: Discounts on most discounted and least discounted funds over time



Source: Minio-Paluello (1998)

Note that the discounts on the most discounted funds decrease whereas the discounts on the least discounted funds increase, and the difference narrows over time.

Convertible Arbitrage

A convertible bond has two securities embedded in it – a conventional bond and a conversion option on the company's stock. When companies have convertible bonds or convertible preferred stock outstanding in conjunction with common stock, warrants, preferred stock and conventional bonds, it is entirely possible that you could find one of these securities mispriced relative to the other, and be able to construct a near-riskless strategy by combining two or more of the securities in a portfolio.

In the simplest form of this strategy, note that since the conversion option is a call option on the stock, you could construct a conversion option by combining the underlying stock and the treasury bond (a replicating portfolio). Adding a conventional bond to this

should create the equivalent of the convertible bond. Once you can do this, you can take advantage of differences between the pricing of the convertible bond and synthetic convertible bond and potentially make arbitrage profits. In the more complex forms, when you have warrants, convertible preferred and other options trading simultaneously on a firm, you could look for options that are mispriced relative to each other, and then buy the cheaper option and sell the more expensive one.

In practice, there are several possible impediments. First, many firms that issue convertible bonds do not have straight bonds outstanding, and you have to substitute in a straight bond issued by a company with similar default risk. Second, companies can force conversion of convertible bonds, which can wreak havoc on arbitrage positions. Third, convertible bonds have long maturities. Thus, there may be no convergence for long periods, and you have to be able to maintain the arbitrage position over these periods. Fourth, transactions costs and execution problems (associated with trading the different securities) may prevent arbitrage.

Determinants of Success

Studies that have looked at closed end funds, dual listed stocks and convertibles all seem to conclude that there are pockets of inefficiency that can be exploited to make money. However, there is residual risk in all of these strategies, arising sometimes from the fact that the assets are not perfectly identical (convertibles versus synthetic convertibles) or because there are no mechanisms for forcing the prices to converge (closed end funds).

So, what would you need to succeed with near arbitrage strategies? The first thing to note is that these strategies will not work for small investors or for very large investors. Small investors will be stymied both by transactions costs and execution problems. Very large investors will quickly drive discounts to parity and eliminate excess returns. If you decide to adopt these strategies, you need to refine and focus your strategies on those opportunities where convergence is most likely. For instance, if you decide to try to exploit the discounts of closed-end funds, you should focus on the closed end funds that are most discounted and concentrate especially on funds where there is the potential to bring pressure on management to open end the funds. You should also avoid funds with substantial illiquid or non-traded stocks in their portfolios, since the net asset values of these funds may be significantly overstated.

The Limits of Arbitrage

In a perfect world (at least for financial economists), any relative mispricing of assets attracts thousands of investors who borrow risklessly and take advantage of the

arbitrage. In the process, they drive it out of existence. In the real world, it is much more likely that any assets that are mispriced are not perfectly identical (thus introducing some risk into the mix) and that only a few large investors have the capacity to access low-cost debt and take advantage of arbitrage opportunities. It is entirely possible then that near arbitrage opportunities will be left unexploited because these large investors are unwilling to risk their capital in these investments. Vishny and Shleifer provide a fascinating twist on this argument. They note that the more mispriced assets become on a relative basis, the greater the risk to arbitrageurs that the mispricing will move against them. Hence, they argue that arbitrageurs will pull back from investing in the most mispriced assets, especially if there are thousands of other traders in the market who are pushing prices in the opposite direction.

Speculative Arbitrage

The word arbitrage is used much too loosely in investments and there are a large number of strategies that are characterized as arbitrage, but actually expose investors to significant risk. In fact, the strategies covered in this section would probably be better characterized as pseudo arbitrage strategies.

Paired Arbitrage

In classic arbitrage, you buy an asset at one price and sell an exactly identical asset at a different (and higher) price. In paired arbitrage, you buy one stock (say GM) and sell another stock that you view as very similar (say Ford), and argue that you are not that exposed to risk. Clearly, this strategy is not riskless since no two equities are exactly identical, and even if they were very similar, there may be no convergence in prices.

Let us consider first how you pair up stocks. The conventional practice among those who have used this strategy on Wall Street has been to look for two stocks whose prices have historically moved together – i.e., have high correlation over time. This often leads to two stocks in the same sector, such as GM and Ford. Once you have paired the stocks, you compute the spread between them and compare this spread to historic norms. If the spread is too wide, you buy the cheaper stock and short the more expensive stock. In many cases, the strategy is self-financing. For example, if Ford is trading at \$ 20 and GM is trading at \$ 40 and you believe that GM is overpriced relative to Ford, you would buy two shares of Ford and sell short one share of GM. If you are right, and the spread narrows between the shares, you will profit on your paired position.

Can such a simplistic strategy, based entirely upon past prices, make excess returns? In 1999, Gatev, Goetzmann and Rouwenhorst tested a variety of trading rules based upon pairs trading from 1982-1997, using the following process:

- Screening first for only stocks that traded every day, the authors found a matching partner for each stock by looking for the stock with the minimum squared deviation in normalized price series³. Intuitively, note that if two stocks move together all the time, the squared distance in returns should be zero. Once they had paired all the stocks, they studied the pairs with the smallest squared deviation separating them.
- With each pair, they tracked the normalized prices of each stock and took a position on the pair, if the difference exceeded the historical range by two standard deviations, buying the cheaper stock and selling the more expensive one.

Over the 15 year period, the pairs trading strategy did significantly better than a buy-and-hold strategy. Strategies of investing in the top 20 pairs earned an excess return of about 6% over a 6-month period, and while the returns drop off for the pairs below the top 20, you continue to earn excess returns. When the pairs are constructed by industry group (rather than just based upon historical prices), the excess returns persist but they are smaller. Controlling for the bid-ask spread in the strategy reduces the excess returns by about a fifth, but the returns are still significant.

While the overall trading strategy looks promising, there are two points worth emphasizing that should also act as cautionary notes about this strategy. The first is that the study quoted above found that the pairs trading strategy created negative returns in about one out of every six periods, and that the difference between pairs often widened before it narrowed. In other words, it is a risky investment strategy that also requires the capacity to trade instantaneously and at low cost. The second is a quote from a well known quantitative analyst, David Shaw, who bemoaned the fact that by the late 1990s, the pickings for quantitative strategies (like pairs trading) had become slim because so many investment banks were adopting the strategies. As the novelty has worn off, it seems unlikely that the pairs trading will generate the kinds of profits it generated during the 1980s.

Merger Arbitrage

As we noted in the last chapter, the stock price of a target company jumps on the announcement of a takeover. However, it trades at a discount usually to the price offered by the acquiring company. The difference between the post-announcement price and the offer

³ If you use absolute prices, a stock with a higher price will always look more volatile. You can normalize the prices around 1 and use these series.

price is called the arbitrage spread, and there are investors who try to profit off this spread in a strategy called merger or risk arbitrage. If the merger succeeds, the arbitrageur captures the arbitrage spreads, but if it fails, he or she could make a substantial loss. In a more sophisticated variant in stock mergers (where shares of the acquiring company are exchanged for shares in the target company), the arbitrageur will sell the acquiring firm's stock in addition to buying the target firm's stock.

To begin with, we should note that the term risk arbitrage is extremely misleading. It is clearly not arbitrage in the classic sense since there are no guaranteed profits and it is not quite clear why the prefix "risk" is attached to it. Notwithstanding this quarrel with terminology, we can examine whether risk arbitrage delivers the kinds of returns we often hear about anecdotally, and if it does, is it compensation for risk (that the merger may not go through) or is it an excess return? Mitchell and Pulvino (2000) use a sample of 4750 mergers and acquisitions to examine this question. They conclude that there are excess returns associated with buying target companies after acquisition announcements of about 9.25% annually, but that you lost about two thirds of these excess returns if you factor in transactions costs and the price impact that you have when you trade (especially on the less liquid companies).

While the overall strategy returns look attractive, Mitchell and Pulvino also point to one unappealing aspect of this strategy. The strategy earns moderate positive returns much of the time, but earns large negative returns when it fails. Thus, they argue that this strategy has payoffs that resemble those you would observe if you sell puts – when the market goes up, you keep the put premium but when it goes down, you lost much more. Does this make it a bad strategy? Not at all, but it points to the dangers of risk arbitrage when it is restricted to a few big-name takeover stocks (as it often is)- an investor who adopts this strategy is generally is one big failure away from going under. If you use leverage to do risk arbitrage, the dangers are multiplied.

Determinants of Success

The fact that we categorize the strategies in this section as speculative arbitrage is not meant to be a negative comment on the strategies. We believe that these are promising investment strategies that have a history of delivering excess returns but they are not riskfree. More ominously, it is easy for those who have done pure arbitrage to drift into near arbitrage and then into speculative arbitrage as they have funds to invest. In some cases, their success at pure or near arbitrage may bring in funds which require this shift. In doing so, however, there are two caveats that have to be kept in mind:

- The use of financial leverage has to be scaled to reflect the riskiness of the strategy. With pure arbitrage, you can borrow 100% of what you need to put the strategy into play. In futures arbitrage, for instance, you borrow 100% of the spot price and borrow the commodity. Since there is no risk, the leverage does not create any damage. As you move to near and speculative arbitrage, this leverage has to be reduced. How much it has to be reduced will depend upon both the degree of risk in the strategy and the speed with which you think prices will converge. The more risky a strategy and the less certain you are about convergence, the less debt you should take on.
- These strategies work best if you can operate without a market impact. As you get more funds to invest and your strategy becomes more visible to others, you run the risk of driving out the very mispricing that attracted you to the market in the first place.

In many ways, the rise and fall of Long Term Capital Management (see below) should stand as testimony to how even the brightest minds in investing can sometimes either miss or willfully ignore these realities. Long Term Capital Management's eventual undoing can be traced to many causes but the immediate cause was the number of speculative arbitrage positions they put in place – pairs trading, interest rate bets – with tremendous leverage.

The Fall of Long Term Capital

Investors considering arbitrage as their preferred investment philosophy should pay heed to the experiences of Long Term Capital Management (LTCM). The firm, which was founded in the early 1990s by ex-Salomon trader, John Merriweather, promised to bring together the best minds in finance to find and take advantage of arbitrage opportunities around the world. Delivering on the first part of the promise, Merriweather lured the best bond traders from Salomon and brought on board two Nobel prize winners – Myron Scholes and Bob Merton. In the first few years of its existence, the firm also lived up to the second part of the promise, earning extraordinary returns for the elite of Wall Street. In those years, LTCM was the envy of the rest of the street as it used low cost debt to lever up its capital and invest in pure and near arbitrage opportunities.

As the funds at their disposal got larger, the firm had to widen its search to include pseudo arbitrage investments. By itself, this would not have been fatal but the firm continued to use the same leverage on these riskier investments as it did on its safe investments. It bet on paired trades in Europe and decreasing spreads in country bond markets, arguing that the sheer number of investments in had in its portfolio would create

diversification – if it lost on one investment, it would gain on another. In 1997, the strategy unraveled as collapses in one market (Russia) spread into other markets as well. As the portfolio dropped in value, LTCM found itself facing the downside of its size and high leverage. Unable to unwind its large positions without affecting market prices and facing the pressures of lenders, LTCM faced certain bankruptcy. Fearing that it would bring down other investors in the market, the Federal Reserve engineered a bailout of the firm.

What are the lessons that we can learn from the fiasco? Besides the cynical one that it is good to have friends in high places, you could argue that the fall of LTCM teaches us that

- (a) Size can be a double-edged sword. While it gives you economies of scale in transactions costs and lowers the cost of funding, it also makes it more difficult for you to unwind positions that you have taken.
- (b) Leverage can make low-risk positions into high-risk investments, since small moves in the price can translate into large changes in equity
- (c) The most brilliant minds in the world and the best analytical tools cannot insulate you from the vagaries of the market.

For an excellent analysis of LTCM, see “When Genius Failed” by Roger Lowenstein

Long Short Strategies – Hedge Funds

In the last few years, hedge funds have become one of the fastest growing parts of the money management business. Largely unregulated, headed by outsized personalities like George Soros and Julian Robertson and seemingly delivering huge returns to their investors, hedge funds have become serious players in the money management game. At the outset of this section, we have to note that it is probably not quite accurate to categorize hedge funds as having any specific strategy, since you can have hedge funds specializing in almost every strategy we have listed in this book. You can have value and growth investing hedge funds, hedge funds that specialize in market timing, hedge funds that invest on information and hedge funds that do convertible arbitrage. The reason that we consider it in this chapter is because it lends itself particularly well to arbitrage strategies, which require that you buy some assets and sell short on others at the same time. In this section, we will take a closer look at hedge fund strategies and how well they really have performed.

Background and History

What makes a fund into a hedge fund? The common characteristic shared by all hedge funds is that they not only buy assets that they feel are undervalued, but that they simultaneously sell short on assets that they believe to be overvalued. Defined this way,

hedge funds have probably been around as long as stock markets have been in existence, though they have traditionally been accessible only to the very wealthy. In the last decade, however, hedge funds have taken a larger and larger market share of total investment funds. While the magnitude of the funds under hedge fund management is disputed, it is estimated that 500 to 600 billion dollars in assets were managed by about 6000 hedge funds in early 2002.

Performance

Are the storied returns to investing in hedge funds true? Are small investors who are often shut out from investing in hedge funds losing out because of this? To answer these questions, we need to look not at anecdotal evidence or the performance of the best hedge funds, but at all hedge funds. In a study of all offshore hedge funds from 1989 to 1995, Brown, Goetzmann and Ibbotson (1999) chronicled the returns in table 11.2.

Table 11.2: Offshore Hedge Fund Returns – 1989 to 1995


<i>Year</i>	<i>No of funds in sample</i>	<i>Arithmetic Average Return</i>	<i>Median Return</i>	<i>Return on S&P 500</i>	<i>Average Annual Fee (as % of money under management)</i>	<i>Average Incentive Fee (as % of excess returns)</i>
1988-89	78	18.08%	20.30%		1.74%	19.76%
1989-90	108	4.36%	3.80%		1.65%	19.52%
1990-91	142	17.13%	15.90%		1.79%	19.55%
1991-92	176	11.98%	10.70%		1.81%	19.34%
1992-93	265	24.59%	22.15%		1.62%	19.10%
1993-94	313	-1.60%	-2.00%		1.64%	18.75%
1994-95	399	18.32%	14.70%		1.55%	18.50%
Entire Period		13.26%		16.47%%		

^aReturns are net of fees

There are several interesting numbers in this table. First, the average hedge fund earned a lower return (13.26%) over the period than the S&P 500 (16.47%), but it also had a lower standard deviation in returns (9.07%) than the S & P 500 (16.32%). Thus, it seems to offer a better payoff to risk, if you divide the average return by the standard deviation – this is the commonly used Sharpe ratio for evaluating money managers. Second, these funds are much more expensive than traditional mutual funds, with much higher annual fees and annual incentive fees that take away one out of every five dollars of excess returns.

As noted earlier, hedge funds come in all flavors. A study in 1999 by Ackermann, McEnally and Ravenscraft looked at the returns on a variety of hedge funds, and their findings are summarized in figure 11.10 below:

Figure 11.10: Hedge Funds: Average Returns and Standard Deviations - 1989-1995



Source: Ackermann, McEnally and Ravenscraft

Note that event-driven and opportunistic funds (which look for arbitrage opportunities) provide the best Sharpe ratios and fund of funds (which invest across what they claim are the best mutual funds) offer the worst.

Liang examined 2016 hedge funds from 1990 to 1999. While his overall conclusions matched those of Brown et al., i.e. that these hedge funds earned a lower return than the S&P 500 (14.2% versus 18.8%), they were less risky and had higher Sharpe ratios (0.41 for the hedge funds versus 0.27 for the S&P 500), he also noted that there a large number of hedge funds die each year. Of the 2016 funds over the period for instance, only 1407 remained live at the end of the period.

In summary, what are we to make of these findings? First, the biggest advantage of hedge funds does not seem to lie in high returns but in lower risk. Every study that we have quoted finds that hedge funds under perform the market, especially over bull markets, but

that they are more efficient in their risk taking.⁴ Second, the high failure rate of hedge funds has to be factored into any investment strategy. An investor considering hedge funds should either hold several hedge funds or view a hedge fund as a supplemental investment. Third, the fee structure is tilted towards management and the expenses are larger for hedge funds. Since many hedge fund strategies are successful only on the small scale, it will be interesting to see what happens to returns as the business grows. In our view, success at attracting more money may ultimately prove fatal to this business.

Conclusion

For many practitioners, the promise of being able to invest no money, take no risk and still make a profit remains alluring. That is essentially what arbitrage allows you to do. In pure arbitrage, you buy a security at one price and sell an exactly identical security at a higher price, thus ensuring yourself a riskless profit. The markets where pure arbitrage is feasible tend to be the derivatives markets, since you can construct equivalent securities using the underlying asset and lending or borrowing. In futures markets, for instance, you can attain equivalent results by borrowing money, buying the underlying asset (storable commodity, stock index or bond) and storing it until the maturity of a futures contract, or buying the futures contract directly. In the options markets, you can replicate a call option by borrowing money and buying the underlying asset and a put option, by selling short on the underlying asset and lending at the riskless rate. Pure arbitrage may also be feasible with default free bonds. If opportunities exist for pure arbitrage, they are likely to be few and far between and available only to a subset of large institutional investors with very low transactions costs and the capacity to take on very high leverage at close to riskless rates.

Near arbitrage refers to scenarios where the two assets being bought and sold are either not exactly identical or where there is no point in time where prices have to converge. We considered three examples – equities on the same company that trade in different markets at different prices, closed end funds that trade at a discount on the net asset value of the securities in the fund and convertible bonds that trade at prices that are inconsistent with the prices of other securities – warrants and convertible preferred - issued by the company. In each, the mispricing may be obvious but the profits are not guaranteed because you cannot force convergence (liquidating the closed end fund or converting the ADR into local

⁴ A contrary viewpoint is offered by three hedge fund managers at AQR Capital Management in Chicago. They argue that the returns on many hedge funds are based upon self-assessments of value for illiquid securities. The resulting smoothing out of returns creates the illusion of low or no correlation with the market and low standard deviations.

shares) or because you cannot create exactly identical securities (convertible arbitrage). You may still be able to construct low risk strategies that earn high returns, but riskless profits are not feasible.

In the final section, we examine what we term speculative or pseudo arbitrage. In pairs trading, two stocks that tend to move together are paired up and traded when the price difference moves out of a historical range – you sell the more expensive stock and you buy the cheaper one. In merger arbitrage, you buy stocks of target firms after acquisitions are announced and hope to make the difference between the price post-announcement and the offer price. While both these strategies offer promising returns, neither is close to being riskless.

In closing, we look at the part of the money management business that is closest in philosophy to arbitrage– the hedge fund business. While they share the common characteristic of a long-short strategy, hedge funds come in all varieties. On average, they have returns that are lower than a strategy of buying and holding stocks, but have risk that is significantly lower. As hedge funds use this finding to bring in more money, they run the risk of being the victims of their own success, since it is not clear whether many hedge fund strategies will scale up – i.e. a convertible arbitrage strategy that works with \$ 100 million may not work with a billion.

Lessons for investors

To be a successful arbitrageur, you need to:

- Understand that near arbitrage is more likely than pure arbitrage: Pure arbitrage opportunities, if they exist, are most likely to be found in derivatives and government bond markets and will be very quickly exploited. Near arbitrage, where you have two almost-identical assets that are priced differently or where you have no forced convergence, should be more common.
- Have excellent execution capabilities and low execution costs: Arbitrage requires you to trade large quantities instantaneously in two or more markets.
- Have access to low cost debt: The pricing differences between two similar assets, even if they exist, are likely to be very small and can be made into substantial returns only by using leverage.

If you decide to move to pseudo arbitrage, you need to

- Keep leverage under control: As you move from pure to near arbitrage and from near to pseudo arbitrage, the risk in your strategy will increase and you should reduce the financial leverage in your strategy accordingly.
- Recognize that size is a double edged sword: As you get more funds to invest, you may be able to reduce your execution costs, but you will also have a much more difficult time getting in and out of your positions quickly and without a price impact.