INTRODUCTION TO CONVENIENCE YIELDS

Prepared for the 2000 Investment Conference

Andrew Smith
of the
Applications of Financial Theory Working Party

Abstract

This paper seeks to explain the concept of convenience yields. The idea was originally developed in the context of commodity futures pricing. In this paper, we apply the concept to more common institutional investments such as equities, cash and bonds. We show how pricing effects due to tax, liquidity, stock lending costs, audit costs and clientele issues, can all be considered in a convenience yield framework.

PART I: MOTIVATION

A Simple Example: Eclipse Sunglasses

On 11 August, 1999, a narrow band of the European peninsula experienced a solar eclipse. In the preceding months, it was possible to buy special sunglasses. Using these sunglasses, it was safe to stare at the sun without risking damage to the retina. Before the eclipse, these sunglasses were for sale at prices around 80p. After the event, the same sunglasses changed hands for around 10p.

It was also possible to place orders for sunglasses with retailers who were temporarily out of stock. A derivative practitioner might refer to these as "forward transactions." For example, I could agree on 9 August to purchase a pair of sunglasses on 10 August. In this case, the forward price would be close to the spot price, that is 80p. However, if the delivery date were after the eclipse, on, say, 12 August, the agreed forward price would have been closer to 10p. This is the case even though the forward agreement was entered before the eclipse. In other words, the market anticipated the fall in the price of sunglasses. A commodities trader might describe the sunglasses market on 9 August as in *backwardation*, that is, the price for forward delivery was lower than the spot price.

Why then would anybody wish to buy sunglasses on 9 August? As an investment decision, this would have been a recipe for disaster - not only did the price fall catastrophically August 11, but this was already anticipated by the market. The reason, of course, for buying the sunglasses had nothing to do with investment. Instead, people bought the sunglasses to stare safely at the sun. The objective of the purchase was not the financial return, but instead the non-financial *convenience* of having access to a physical pair of sunglasses during the eclipse. The 70p backwardation in the sunglasses market puts a financial value on this convenience effect. The fact that market prices fell on 11 August, and that this fall was accurately predicted by the market, does not imply any sort of market failure.

Efficient Markets and Price Protection

The existence of convenience yields has important implications for investment management and institutional investment.

Financial theory provides us with the concept of market efficiency, where asset prices in the market already reflect available information. This creates difficulties for stock-pickers, because all their research appears futile - by the time they produce their buy/sell recommendations, the price has already moved. But market efficiency also has many benefits, especially the concept of *price protection*. Price protection arises because in an efficient market, it is just as difficult to pick bad stocks as it is to pick good stocks. Even given unlimited resources of stupidity, it is difficult for a manager to under-perform consistently. The dumb investment managers are *price protected* by the efficient market mechanism. This observation results in the rather naive

suggestion that investment manager skill is irrelevant - performance is determined by luck alone.

However, an investor who relied on price protection in the sunglasses market could come badly unstuck. This is because much of the price of the sunglasses related to non-investment benefits, which are of no value to a pure investor. If we can find evidence of convenience yield effects on financial markets, this may give us some insight into the failure of price protection for investment managers, and hence some signals as to where investment managers could create value for their clients.

Rational and Irrational Explanations

There is of course another explanation for the failure of price protection in any market. Price protection could fail if the market was inefficient and if investors were acting irrationally. So why do we not accept that, to outperform, a manager must simply seek to exploit the irrationality of others?

The answer is that irrationality is hard to predict. The concept of rational behaviour is clearly defined, but irrational behaviour encompasses everything else. The assertion that investors are irrational does not of itself help us to exploit that irrationality. To make an analogy, suppose I ask for directions to Down Hall from Bishops Stortford station. Someone might advise to head South-East in a straight line, and this is a useful instruction because there is only one straight line heading South-East. If, instead, someone advises me to head South-East in a wiggly line, this is much less helpful.

In the same way, it is easy to construct vague rhetoric about inefficiency but much harder to demonstrate the validity of a specific bias. The literature on behavioural finance seems particularly prone to this pitfall, in that deviations in either direction from optimality, as viewed with perfect hindsight, tend to be treated as evidence of cognitive biases. So, for example, investors who failed to mortgage their houses in 1975 to invest in the stock market, showed "myopic loss aversion". On the other hand, investors who *did* mortgage their houses in 1990 to invest in the stock market were "overconfident". In essence, behavioural financiers extend the definition of irrationality so broadly that any action other than buying the winning lottery ticket is irrational. This may give them a wide field for study, but is useless for making tactical investment calls.

A more robust way to proceed is to consider whether there are any failures in price protection which we expect to see even if other investors were rational. Such failures are more likely to persist over time than any arising from irrational biases. We do not exclude the possibility of price distortions caused by irrational behaviour. Instead, we merely assert that distortions consistent with investor rationality are likely to be the easiest to identify.

PART II: CAUSES OF CONVENIENCE YIELDS

The Classical Paradigm

We first consider the classical one period mean-variance paradigm. This has a number of assumptions, all of which are unrealistic to some degree. The assumptions include:

- all agents agree on the underlying probability law
- agents rationally optimise risk against return, using single period variance to measure risk
- there are no limits of portfolio holdings or on short sales
- a given investor provides the same cash flows to anyone who holds it.

The major consequence of these assumptions is that each investors' optimal portfolio should lie on the so-called *capital market line* that efficiently balances risk against return. For any three investors, one will hold a mixture of the portfolios of the other two.

A further consequence of the model is that market prices are consistent with the capital asset pricing model, consisting of a base ("zero beta") return minus or plus a risk premium determined by covariance with the market as a whole. In this environment, price protection holds.

We now consider whether any of these assumptions can be relaxed without sacrificing rationality, and if so, what the effects would be on capital market pricing.

Consensus

Let us consider first what happens if we abandon the notion of consensus, and suppose that all investors have different beliefs.

The difficulty here is that we have no natural alternative hypothesis. There are many ways in which investors' views could be mistaken and these could all lead to different market pricing biases. Any theory based on diversity of opinion is unlikely to have predictive power, unless we have a systematic way of identifying the diversity.

We can wriggle out of this by saying that diversity of investor opinion is not, in any case, consistent with rationality. In many cases, it is relatively easy to reverse engineer market consensus views from observed market prices. Numerous studies have shown these market forecasts to be more robust than individual investors views or mechanical trading rules. A rational investor could not persist in the view that his own forecasts were superior to the market.

Use of Single Period Variance as a Measure of Risk

Numerous commentators have criticised the use of single period variance as a measure of risk. However, empirical evidence of the risk measure underlying actual investment decisions is hard to find. This is largely because different risk measures, for example those based on utility functions, or on probabilities of missing a target, seem to give portfolio recommendations which are indistinguishable from mean-variance efficient portfolios. In other words, although it might be that actual decision rules take a very different algebraic form to mean-variance optimisation, it appears that the variance assumption is not crucial to the analysis. In particular, the concept of rewards based on systematic risk, and the capital market line, extend readily to other utility measures. Price protection is likely to work in much the same way whatever utility functions investors are assumed to possess.

Limits on Short Sales

Restrictions on short sales, or on borrowing, provide a greater potential for convenience yield effects. We can replicate the optimisation for the classical case, but disallowing short positions. In this case, we find that

- Investors may hold a variety of efficient portfolios, which are not necessarily linear combinations of each other
- For any given investor, the assets actually held will trade at an efficient price given the correlation with that investor's portfolio. However, the market aggregate portfolio does not come into the calculation
- In addition, each investor may avoid some asset classes entirely. For these asset classes, the market price will exceed what the investor would see as the fair value by reference to the risks in the optimal portfolio.

Even the simple device of allowing for restrictions on short sales has provided some dramatic improvements in the ability of our model to explain real world features, in particular the diversity of actual portfolios held by investors.

The diversity of investors is sometimes also described as a *clientele effect*, so that some assets, rather than being held by the whole market, are held only by a subset of investors – the so-called clientele. In this case, investors outside the clientele would see the asset as being overpriced. To this extent, price protection has failed, and we have convenience yield effects creeping in.

Differential Cash Flows

The most important class of reasons for convenience yield effects arise from cash flow asymmetries, where different investors receive different cash flows from the same investments. The effect in each case is that the investment is attractive to a particular clientele, most of whom will be those who derive the most generous cash flow from the asset. Other investors outside the clientele would do well to avoid the asset entirely. The differential cash flow effects are sufficiently important that we give them a special section on their own.

PART III: CASH FLOW ASYMMETRIES

Tax – Induced Asymmetry

The impact of tax provides many opportunities for different parties to own different (net) cash flows from the same investment. There is scope for failure of price protection for investors who choose tax inefficient strategies (for example, realising capital gains too often inside taxed funds, or holding inside gross funds investments which are efficient for tax-paying investors).

There are a number of well documented pricing effects arising from tax asymmetries. Perhaps the most familiar to a UK audience would be the coupon effects which were evident in the gilt market prior to the introduction of capital gains tax in the mid 1990's. Before this change, the low coupon bonds were held mainly by net investors, while gross investors held the higher coupon bonds. This is a classic example of a clientele effect.

In this classic tax example, we could also observe the inequalities for assets that were not held. For example, low coupon bonds typically had a lower gross redemption yield than high coupon bonds; these bonds seemed expensive to gross investors. In the same way, high coupon bonds seemed expensive to net investors.

On some occasions, these tax effects can be large. For example, the removal in 1989 of certain tax subsidies available to home buyers is widely held to be responsible for the subsequent fall in property values. This led to a failure of price protection; a wise investor who did not benefit from the tax subsidy anyway should have deferred purchase until after the tax change.

Commodity Access Premiums

More complex convenience effects can arise because of the significant market impact costs of acquiring cash or inventory suddenly. If an organisation has cash flow problems, it can turn to other routes (for example, raising equity or debt finance or selling assets) to raise the cash, but these routes can be expensive, especially if approached in haste. The costs include market impact costs and also the fees from various financial intermediaries who arrange the deals. In the same way, a factory who suddenly identifies a need for oil, may have to pay high transportation costs to arrange delivery at short notice. We could refer to these costs as *inconvenience costs*.

These costs can be avoided (or at least, reduced) by holding a small inventory of assets which might be needed in a hurry. This could be worth doing *even if* there were a penalty for holding physical assets, for example, lower returns as expressed by the convenience yield (sometimes also called a liquidity premium). Provided this convenience yield is not greater than the inconvenience costs which would otherwise be incurred, it can still be rational to hold the inventory. This may explain why cash interest rates are often significantly lower than short bond yields.

A particular example of these costs arises in the gilt repo markets. There may be temporary squeezes on some stocks, because for example of a particular stock being cheapest to deliver on a futures contract. In this case, the gilt is described as being "special on repo", so that instead of swapping the gilt return for LIBOR, investors may only be able to achieve (for example) LIBOR minus 1%. However, there are also penalties, for those with short futures positions, to *not* holding this particular stock, because delivering an alternative stock would involve giving a free lunch to the other side of the futures contract. As a result, significant convenience yields can arise of those stocks. Any investor without short futures positions would be unwise to buy or hold these particular stocks.

Audit and Compliance Costs

Another example of convenience yield could arise because of audit or compliance costs. For example, suppose an insurer wants to hold long bonds in order to match certain liabilities. There is a choice between using government bonds or corporate bonds. The corporate bond market is less liquid and also less well researched. As a result, investments in corporate bonds can carry significant audit or compliance costs, because of the need to collect and manage credit risk information, and sometimes because of legal costs in establishing valid title to the investments. Other aspects of compliance are more costly if recent market prices are not readily available. Regulators may single out insurance companies who invest in unconventional assets for special (and costly) attention. These inconveniences all carry costs, which are likely to feature implicitly in the prices of different bonds.

One would expect these costs to be reflected in the yield spread between corporate and government bonds. Organisations with a lighter than average compliance regime (for example, private funds, offshore funds or pension funds) may constitute a natural clientele for corporate bonds.

Another aspect of audit costs arises in the use of assets as collateral. Lenders wish to ensure adequate security for a loan, which usually involves ensuring that the assets actually exist, and may involve charges on those assets to prevent their subsequent sale. This is relatively easy to implement in the case of assets such as equities, bonds or real estate. Synthetic alternatives to these assets (for example, cash + futures positions) may be much harder to use as collateral. As a result, some entities will exhibit a distinct preference for physical assets over their synthetic alternatives, which can create convenience yield effects on those assets. For the same reason, entities which do not wish to borrow against physical assets may prefer to lend their assets to borrowers, and swap back via derivatives into a synthetic alternative.

Index Entry and Exit Effects

Audit costs can go as far as to create a moral hazard. For example, investors have become increasingly suspicious of managed equity funds, because the charging structure may be unclear. Various charges may be debited to the investment funds (for example, bid offer spreads and commissions). These commissions may benefit other sister companies of the fund management organisation. Investors are right to be

suspicious of such behaviour, but the standard fund management products offer little protection against such practices. This is especially true where markets are illiquid, unquoted or not widely analysed.

On the other hand, newer products are often linked to a particular index, for example the FTSE 100, which is calculated independently of any particular fund manager. The charging structure is more transparent, which reduces the impact of moral hazard on the end user investor. As a result, there is a clientele effect where many investors prefer to invest only in managers whose structures are simple enough to monitor effectively.

As consumers flock to index products, the larger stocks are bid up by the financial institutions seeking to hedge FTSE 100-linked liabilities. This creates a number of strange effects, such as the FTSE 100 effect where the price of a company tends to rise immediately after joining the index, and the small companies effect where less well known or researched companies may achieve higher returns.

While these effects are sometimes claimed as evidence of market inefficiency, these supposed inefficiencies are hard to exploit once the audit costs are taken into account. For example, the investor in small companies funds faces the risk of exaggerated charges being debited to the fund because of less effective audit controls on such companies. The main consequence for investor behaviour is that, if those audit costs are to be incurred anyway (for example, because of a commitment to active management), then the larger companies should be avoid, to prevent the audit costs being effectively paid twice. To put this differently, an investor in large companies should first consider index or tracker products which minimise the discretion available to fund managers.

Non-Cash Flow Costs

Finally, convenience effects could arise because of investments held for reasons other than pure investment return, for example, sentimental reasons. In addition to the dividend returns, some shareholders may get a warm emotional feeling, or even social acceptance in particular sections of society, from the ownership of the shares. This warm feeling can be priced into the share price, rendering the shares unattractive to less sentimental investors. Probably the most notorious example is football club shares, whose performance has been awful if the warm feeling element is not taken into account. Other examples include transport companies or utility suppliers who offer discounted products, to their shareholders. Recent developments in ethical investments may have created similar effects.

PART IV: COMPUTING CONVENIENCE YIELDS

Convenience Yields on Commodities

The concept of a convenience yield first arose in the context of commodity pricing. For example, we might find that a commodity had the following prices:

Spot price: $100
Forward price in one year: $104
Market one-year interest rate: 6%

The *convenience yield* is defined by:

forward price * (1 + convenience yield) = spot price * (1 + interest rate)

which in this case gives a convenience yield around 2%. The convenience yield is a measure of the opportunity cost of holding inventory. For example, suppose I need some of this commodity in one year. I could either buy it now and store the commodity for a year, or save up know at buy a year forward. The latter course of action saves me about $2. This helps to explain why companies often try to avoid building up large inventories - it is generally cheaper to order just in time. On the other hand, sometimes the convenience yield is worth paying in order to guarantee fast access to a physical inventory, especially if the need for the product is difficult to predict.

For storable commodities, such as oil or metals, we would usually expect the convenience yield to be positive. Otherwise, if convenience yields were negative (and storage/insurance costs negligible) there would be a possible arbitrage buying physical now and selling forward.

These have important implications for any investors considering speculation in commodity markets. Even if you think it might be sensible for investors to accept commodity risk, it is *definitely* daft to speculate with a buy-and-hold strategy based on physical storage. A forward purchase will always work out cheaper, because then the investor does not pay the convenience yield. So here we have another failure of price protection.

Convenience Yields on Financial Assets

Physical commodities do not usually provide an income (although there are exceptions here for commodities sometimes seen as investments; for example the owner of a beautiful painting could perhaps charge visitors to look at it). On the other hand, financial investments usually do produce income, which complicates matters considerably. (Once again, there are exceptions; for example, one could consider equities held only between dividend dates).

As a result, we have to redefine the notion of convenience yields for financial assets. Our working definition is:

forward price * (1 + expected running yield) * (1 + convenience yield)
= spot price * (1 + interest rate)

This definition is messy, because the running yield should relate to the *expected* income, which is inevitably an estimate. Different analysts might reach different answers. This must be distinguished from the "implied" running yield used in derivatives markets, which are computed by working backwards from observed forward prices. We would therefore have (ignoring second order terms)

implied running yield = expected running yield + convenience yield

In the case of physical commodities, this plainly reduces to our previous formulation on setting expected running yield = 0.

We would not expect financial assets to carry large convenience yields - in this way, financial assets differ from physical commodities. This is because the main reason for holding financial assets is to obtain the financial return. Unlike sunglasses, the non-financial reasons to hold financial assets are likely to have only a marginal effect.

In the Absence of Derivatives

Where there is no comparable forward price, convenience yields are far more difficult to estimate. In principle, we could look for departures of actual returns from those implied by theory. Unfortunately, the general level of volatility in asset prices makes such risk premiums very difficult to pin down with confidence.

An alternative approach is to identify the natural clientele for a particular asset, and to identify the size of the cash flow benefit which is special to that clientele. This could give an upper bound on the size of the convenience yield.

Consequences for Corporate Valuation

The existence of convenience yields, if material, could have a large impact on the valuation of corporate projects. The existence of clientele effects goes some way to undermine belief in a universal fair value metric which is common to all enterprises.

The most difficult area following from convenience yields is the matter of tax. As this is the subject of a separate paper at this conference, we only cover some actuarial aspects briefly here.

One contentious question is whether index linked gilts have a natural clientele in the hand of higher rate tax payers. It could be that pension funds fall outside that clientele. Does this point undermine the argument that pension funds should discount their liabilities at index linked gilt yields? If so, what should be used instead?

The same issue arises for life insurers. In a fund taxed on income, it may be that even when liabilities would be theoretically matched with gilts, equities are a more tax efficient investment on the basis of the franking of dividend income. In this case, is it unfair to value liabilities at net gilt yields?

Conclusions

Modern finance theory has produced a number of general findings which while approximately right, remain contentious within the financial community. The question of market efficiency, and price protection is sensitive to various assumptions underlying equilibrium models, and these assumptions may not always be valid. This paper has investigated how some of these assumptions might break down, and how the resulting price effects could at first sight be mistaken for evidence of market malfunction.

Acknowledgements

I am grateful to Jon Exley and Shyam Mehta for their comments on earlier drafts of this paper. However, all errors remain mine alone.