



Smart Investor Insight into fundamental analysis

“The most important factor to look at when assessing a possible investment is how much profit it makes.”

Part 2.1

Understanding profits

This is the first of the advanced series of Smart Investor notes. We reviewed the basic concepts related to JSE-investing in the Basics Series. In this series of advanced notes we'll cover issues to help you develop your analytical insight into investing in shares.

This first note considers an often used but seldom examined concept in finance: profits. They are the point of all investing. So it's a good place to start.

There are three different levels of profitability, and each tells you different things:

Gross profit is useful in assessing a company's pricing power and its ability to source inputs cheaply. It is simply the total revenue minus the costs of goods sold. A lower gross profit margin – which is the difference between the revenue and cost of goods sold measured as a percentage – is an indication that

input costs are rising or prices are being forced down, or both. For example, retailers are facing higher input costs because of growing labour and energy costs. Mining companies, on the other hand, face major pricing pressure when international mineral prices fall, pushing down their profit margins.

Operating profit is a good measure of a company's ability to manage its fixed costs – the rent, salaries of staff and other costs that don't depend on sales. When this measure rises, it indicates that a company is getting a handle on its costs and overheads. Operating profit is less dependent on market conditions than the gross margin and is a more accurate measure of how efficient management is.

Net profit is the bottom line. It tells you how well a company also manages non-operating costs such as debt, financing and depreciation of assets. These extra costs reflect less on management's ability to operate a profitable business than the decisions of shareholders on how the company should be financed.

Introducing profit measures

What is it? Gross profit, operating profit and net profit and their associated margins are used to measure the performance of a company.

Why are they useful? Profit margins are a key way of assessing how well the management of a company is doing its job. It

is helpful to compare margins to peers as well as to examine margins over time. A company with shrinking margins is usually one facing increasing competition and therefore pressure on the prices it can charge, or facing increasing costs. This can reflect poor strategy or inefficient operations.

Gross profit

This is the simplest calculation as it reflects total sales less the cost of sales. "Cost of sales" is the direct variable cost of a company's sales activities. It's easiest to think of it in terms of retail or wholesale businesses where cost of sales is the cost of stock. It excludes overheads and other non-sales related costs.

Gross profit = revenue – cost of goods sold

Simply put, if a widget costs you R100 to produce or buy and you sell it for R200, your gross profit is R100 per widget.

The gross profit margin in that example is 50%, which is calculated using the basic equation:

Gross profit margin = gross profit/net sales x 100

This figure is particularly useful when making comparisons between companies in the same industry as it represents a fair basis for comparison, and reveals quickly who has higher prices and the better ability to source stock cheaply.

Operating profit

Operating profit is a measure of a company's operating efficiency as a whole, including all the expenses it incurs in its daily business activities.

Operating profit can also be referred to as EBIT (earnings before interest and tax) or sometimes EBITDA (which also excludes depreciation and amortisation).

Operating profit is calculated in this way:

Operating profit = gross profit – operating expenses

This produces a figure reflecting the income that is left after all costs associated with running the company are deducted. This typically includes the same costs as for gross margin, and adds overheads such as marketing, administration expenses and

rent, but excludes interest and tax. For example, in the case of a retailer, its operating profit includes the cost of stock as well as the cost of salaries, rental, head office administration, and so on.

The operating profit margin represents the operating profit as a percentage of sales and measures a company's operating efficiency.

The formula is simply:

Operating profit margin = operating profit/net sales x 100

Reducing the profit to a percentage of sales produces a more accurate means to compare the relative performance of different companies.

Snapshot

It is important to note that different industries have differing margins. Retailers tend to have low margins and focus on high

volumes to generate revenue. High-margin businesses tend to be found where competition is restricted.

Trading insight

Profits are everything to shareholders! When a company surprises with an earnings announcement or trading update the share price is likely to respond. Understanding how to

quickly analyse the company's margins, and whether a change in profits is sustainable or a once-off, is critical to knowing whether there is a short-term trading opportunity available.

Net profit

Net profit is the bottom line profit that is attributable to shareholders as it reflects the income left after all expenses are paid, including interest and tax.

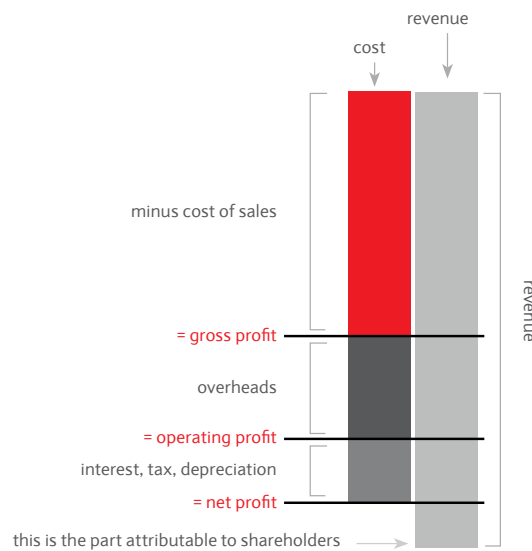
The net profit margin, by contrast, shows how much of each sales rand shows up as net income after all expenses.

This is calculated as:

$$\text{Net profit margin} = \text{net profit} / \text{net sales}$$

Therefore, if a company's net profit margin is 5% it means that it generates R5 of net profit for every R100 of revenue generated.

Here's a summary of how to think about different "levels" of profit:



Which kind of company should you invest in?

The most important factor to look at when assessing a possible investment is how much profit it makes. The profitability ratios discussed here provide a view into the profit levels: for example, gross profit can be expected to be similar for two companies in the same sector, but the further you drill down

to the operating profit and net profit figures, the better the picture of how efficiently each company is producing or selling goods. The higher the margins, the better management is at sweating more profit out of each rand consumed in the sales or production process.

Some examples

The table below shows the operating profit margins of selected food and clothing retailers. You can see that clothing retailers enjoy bigger margins than food retailers.

It also looks like the management teams at Shoprite and Mr Price have been doing excellent work to widen their profit margins, while Pick n Pay's shrinking margins reflect the group's struggling fortunes of the past few years.

For most companies this is a valuable measurement but it is particularly useful for companies that operate with high turnover on low margins, where management is challenged with trying to increase the percentage profit on every product produced or sold.

When comparing similar companies, for example Pick n Pay and Spar, one can quickly tell which is the more efficient by looking at the operating profit margin.

Operating profit margins (%)						
Company	Ave. growth (%)	2013	2012	2011	2010	2009
Shoprite Holdings	4.2%	5,78	5,52	5,41	5,03	4,91
Pick n Pay Stores	-21.1%	1,44	2,29	2,77	3,35	3,83
The Spar Group	-2.2%	3,48	3,53	3,69	3,78	3,80
Woolworths Holdings	10.4%	9,80	9,33	8,21	6,99	6,63
The Foschini Group	1.3%	8,10	9,15	8,08	5,99	8,72
Mr Price Group	12.1%	15,10	14,37	13,03	10,49	9,63

 **Part 2.2 – Next we discuss the key measure in fundamental analysis: the price:earnings ratio.**



Smart Investor Insight into fundamental analysis

“The PE ratio is based on historic earnings, but share prices are all about what the company is expected to deliver in the future.”

Part 2.2

The PE Ratio

Share prices are not useful in comparing companies because the actual price of a share is arbitrary. For example, the fact that Anglo American is trading (at the time of writing) at R214/share and Capitec Bank is trading at R554/share does not mean that Capitec is worth twice as much as Anglo. Share price levels are largely a function of the number of shares in issue and can be changed through share consolidations, share splits, buy-backs and new share issues. In this note we introduce one common way to compare shares – the price:earnings ratio.

Price:earnings ratio

The PE ratio is a ratio of the share price to the latest full year's profits of the company per share. In South Africa, headline earnings per share (HEPS) are used, but this may vary from country to country. The formula is:

$$\frac{\text{Share price (cents)}}{\text{HEPS (cents)}}$$

If Company A is trading at 100c/share and earned 10c/share in its last financial year, the PE ratio would be $100c/10c = 10$. The share can be said to be trading on a multiple of 10 times earnings.

But what if Company B is also trading at 100c/share but its headline earnings were 20c/share in its last financial year?

The PE ratio would $100/20 = 5$. Immediately you can tell that, even though their share prices are the same, Company A is actually twice as expensive as Company B.

If you are interested in investing in either share, the next step in your analysis would be to try to find out why this is so.

The PE ratio is based on historic earnings, but share prices are all about what the company is expected to deliver in the future. So if there is a difference in PE ratios, it must be because the market is expecting different future performance from the two shares.

Snapshot

What is it? Share price divided by historic earnings.

Why are they useful? It is a quick way to determine if a company's share is cheap or expensive relative to its peers and the general market. It also provides a guide to the market expectations of a company's future earnings – a high PE means the market expects strong earnings growth in the future.

Why is it not useful? Just going on past earnings is dangerous. A share price is based on what the market believes the company's future earnings will be (see forward PE below), so a low PE ratio might simply mean the market is expecting future earnings to be poor.

Live market example

In the JSE's food retailers and wholesalers category, Pick n Pay is on a PE of 32.9 (at close of trade on May 4 2015) and its competitor, The Spar Group, is on a PE of 24.2.

Clearly the market is expecting substantial growth in earnings from Pick n Pay and less so for Spar. This places more pressure on Pick n Pay to produce high earnings, in line with its high PE ratio.

It is impossible to provide a general guide as to what levels a PE ratio is considered cheap or expensive. For example, in a bull market when companies across most industries are growing earnings at a healthy rate, a high PE ratio in itself should not be a deterrent to an investor as a company with good historic earnings is likely to meet the market expectations of higher future earnings.

Very generally though, a PE of around 24 times (like Spar's) is not considered cheap but a PE above about 25 times is usually considered expensive. But it is important that the retailers meet these expectations.

If a future earnings announcement disappoints the market, you can expect share prices to move down in response, bringing the PE ratios lower.

Retailers are currently priced rather high compared with the rest of the market. As a comparison, the average PE ratio of the entire All Share Index on May 4 was 19.2, which is above its historic average of around 14.

Varieties of PE ratio: historic vs trailing PEs

PE ratios are based on historic earnings, usually taking the company's latest set of published full-year earnings. However, if a company has since published a set of interim results covering the first six months of its financial year, the historic PE will be out of date.

Analysts then use trailing 12 months' earnings, taking the headline earnings from the latest interim results and adding the headline earnings from the final six months of its full-year results. That is known as a trailing PE ratio.

Forward PE

As mentioned, a share price reflects what the market believes about the future earnings of a company, rather than what the company has actually earned in the past. The forward PE ratio attempts to determine what the company's future earnings will be, so the formula would be:

However, forecasts are just a best guess as to what the company is going to earn and there is always plenty of room for error. Usually when the media talk about forward PEs, they are using the consensus forecasts of professional financial analysts who publish their earnings expectations.

$$\frac{\text{Share price (cents)}}{\text{HEPS (cents)}}$$

About PE ratios

PE ratios are a great way to compare companies and the prices they are trading at. It is a metric that is often used by “value” investors – Warren Buffett is an oft-cited example – who assume that the lower the PE, the greater the value in the company relative to the share price.

One way to think of the PE is as the number of years it will take for the company to earn its own share price. So a PE of 10 means the company will take 10 years to generate the profits to cover the price you are paying for its shares at its current earnings level.

This is clearer when looking at the earnings yield which is just the inverse of the PE ratio (“inverse” just means to divide 1 by the PE). In our example of Company A, the earnings yield is $1/10$ which is 10%.

This shows the “yield” to the investor from the company’s profits, i.e. the value that the investor derives each year from holding the stock. This is not the cash you actually receive, but is the theoretical value that is attributable to you as the investor.

Final thought

A PE ratio is a useful indicator, but certainly not something that can be used alone. It is important that you do research to understand why the PE ratio is low. It may be because the

company is genuinely cheap, or it may be that profits are going to be under severe pressure in future. You want to be buying the companies that are genuinely cheap.

 **Part 2.3 – We introduce the PEG ratio which provides a better way to see the PE in the context of a company’s growth prospects.**



Smart Investor Basic concepts used in fundamental analysis

“Prospective earnings may differ markedly from historic earnings, and it is future earnings that a shareholder should be focused on.”

Part 2.3

The PEG Ratio

The PEG (price:earnings/growth) ratio takes the use of the price:earnings ratio – discussed in our previous Smart Investor note – a step further. The PEG ratio can help determine whether a high PE ratio simply means a company’s share is expensive or whether it reflects strong growth prospects. It provides another level of insight to compare companies in determining which are cheap or expensive.

As we discussed, by understanding a company’s PE and growth prospects, one can anticipate how earnings announcements and other news will affect share prices when trading those stocks.

Introducing PEG ratios

What is it? Price:earnings ratio divided by company's earnings growth. It represents the ratio of the PE to the expected future earnings growth rate.

Why are they useful? It is a more complete measure than the straight PE ratio of whether a company's share price is cheap or expensive, as earnings growth is taken into account.

Why is it risky? The PEG ratio relies heavily on earnings forecasts, which have a wide margin of error. Circumstances in markets can change quickly. For example, a 5% depreciation in the rand can have a huge impact on the future earnings of an exporter, or a rapid rise in the gold price fundamentally affects the performance of gold mining companies.

PEG ratio

The PEG ratio is a company's PE ratio divided by its percentage growth in earnings for a specified period. By adding the extra step of dividing by growth, we reveal whether or not a high or low PE ratio is justified. In South Africa, headline earnings per share (HEPS) are used.

The formula is:

$$\text{PEG} = \frac{\text{Share price (cents)}}{\text{HEPS (cents)}}$$

Example: If a company's PE ratio (price divided by headline earnings per share) is 10 and its headline earnings have grown by an average of 10% a year for the past five years, the PEG ratio would be $10/10 = 1$ (call this example Company A). The PEG ratio is useful in assessing whether a high PE ratio is warranted for a company. For example, if Company B's PE is 25 times and its five-year average earnings growth is 10% a year, its PEG would be 2,5 (25/105).

Immediately one can tell that the high PE is not justified compared to company A. But, there could be an explanation: that the forward growth scenario is very positive for B. Prospective earnings may differ markedly from historic earnings, and it is future earnings that a shareholder should be focused on. To determine that, one would have to project what a company's earnings will be in the foreseeable future (say three years). Alternatively, analysts make earnings predictions for most of the JSE's highly traded companies and their consensus forecasts can be used. These can be accessed through some financial services news websites. Bear in mind though that there is a high margin for error in forecasting.

Assuming Company B's projected earnings growth averages 20% a year, the forward PEG ratio would be 1,25. Immediately one can tell that the high PE is based on expectations of good earnings growth in the future. While the PEG of 1,25 indicates it is still on the expensive side, it is not too demanding relative to expected earnings.

Snapshot

There are variations to the basic formula for a PEG ratio:
 PEG = PE/earnings growth (taking all figures from historic earnings)
 Current PEG = Current PE/historic earnings growth
 Forward PEG = Forward PE/forward earnings growth

Live market example

To take the example from our previous article on PE ratios further, in the JSE's food retailers' and wholesalers' category, Pick n Pay is on a PE of 32.9 (at close of trade on May 4 2015) and its competitor, The Spar Group, is on a PE of 24.2.

Just taking those numbers at face value with Pick n Pay's PE nearly 50% higher than Spar's, the market appears to be expecting substantial growth in earnings from Pick n Pay and less so for Spar. Measuring their forward PEG ratios would indicate if this was so. Projected earnings are drawn from I-Net Bridge's consensus forecasts.

	Last 12 months' HEPS	Projected HEPS: next financial year	Projected HEPS: two years	Projected HEPS: three years
Spar	801,2c (2014)	887,7c (2015)	993,5 (2016)	1089,0 (2017)
Pick n Pay	174,7c (2014)	222,5c (2015)	281,9 (2016)	N/A

Because their financial years differ, it makes it a bit more complicated to compare the two, so we have made two calculations: the forward PEG for the next set of results,

and the forward PEG taking the compound annual growth rate (CAGR) over two years. (See addendum on next page: Calculating CAGR.)

Company	PE	Projected earnings growth for next year	Forward PEG (1 year ahead)	Projected earnings growth: two years	Forward PEG (2 years ahead)
Spar	24,2	10.8%	2,24	24.0%	1,01
Pick n Pay	32,7	27.3%	1,20	61.4%	0,53

Immediately you can see that it is in fact Pick n Pay which is the cheaper, relative to expected earnings growth, even though its PE ratio is twice that of Spar's. That's because Pick n Pay's growth rate is projected at 27.3% for the next year while Spar's is 10.8%.

So the high PE looks much more justified now. The one year PEG ratio for Pick n Pay is 1.2 and for Spar it is 2.2.

While the PE measure made Pick n Pay look very expensive, the PEG measure makes it look cheap relative to Spar.

Of course, the pressure is on Pick n Pay to produce those high earnings that the market expects – should it not deliver, you can expect serious downward pressure on the share price to bring the PE back in line with appropriate earnings expectations.

DY% **Part 2.4 – We introduce the PEG ratio which provides a better way to see the PE in the context of a company's growth prospects.**



Smart Investor Basic concepts used in fundamental analysis

“Some companies have an explicit dividend policy, usually expressed as a percentage of profits that will be paid to shareholders.

Part 2.4

The dividend yield

Cash never lies, goes the old adage. That’s true in all sorts of fundamental analysis. From a shareholder perspective, there is no more important form of cash than the dividends you receive from shares you own.

Not all companies pay dividends, sometimes for good reason. But, depending on what your objectives are, the dividend yield of a company is always an important component to consider. (You may want to review note 4 in our Basics series on the sources of returns for shares to get the background).

The formula:

The dividend yield is the amount of dividends paid out expressed as a percentage of the share price:

$$\frac{\text{dividend paid (cents)}}{\text{share price (cents)}} \times 100$$

For example, if two companies both pay annual dividends of 100c/share, but Company A’s share price is R20 and Company B’s is R40, Company A has a dividend yield of 5% and Company B has 2.5%. Assuming all other factors are equal, an investor looking to supplement his or her income would likely prefer Company A’s stock.

Snapshot

What is it? The profits a company returns to shareholders as a percentage of the share price.

Why are they useful? The dividend yield is a very important component of a shareholder's returns and a key ratio used in fundamental analysis. It allows for comparisons between different shares based on the hard cash paid out to shareholders.

What do you need to know? Some companies pay dividends, others do not, but opt to reinvest profits to grow the company. There are two components to returns you get from owning a share – the dividend and the capital appreciation of the value of the share itself. These two elements are in tension with each other. The more a company pays out in dividends, the less cash it will have to fund growth, so the lower the capital appreciation in the share price is likely to be.

Historic dividend yields

Dividend yields quoted in the media are historic dividend yields. Companies pay out dividends once or twice a year, and the dividend yield reflects the dividends paid in the past 12 months. Sometimes analysts forecast what dividends a company is expected to pay and calculate a yield based on this expectation. This represents the forward dividend yield.

One weakness of the historic dividend yield is its backwards-looking nature. Like many other indicators based on historic information, it is risky because the company's prospects may change. A share price reflects what the market believes about the future prospects of a company. Some companies may have a high dividend yield because of their last dividend payment, but will not be able to pay such a dividend in future.

Dividend payers versus non-payers

The decision on whether or not to pay a dividend rests with the board and the essential question directors should ask is: would our company be able to generate better returns from the cash than shareholders could earn if they gave the money back to them?

If the company is in a high growth phase, shareholders will be better off if the company uses its profits to fund future growth. This provides a better return than shareholders can get on the cash. This relationship is clear on days that a company goes "ex dividend". This is the set date when the dividend is declared and buyers of the share are no longer entitled to the dividend. If a company goes ex dividend on 11 October, you would receive the dividend payout if you buy the share on 10 October, but not if you buy it on 11 October. On 11 October the share price will fall by the amount of the dividend.

There is also an important tax consideration. When a company pays a dividend the shareholder has to pay dividend tax at a rate of 15% on the payment (in South Africa the company retains this and pays SARS directly. The dividend tax replaced the secondary tax on companies in April 2012). However, if the company does not distribute its profits and instead uses them to fund growth, there is no dividend tax payable. So it is more tax efficient for companies not to pay out dividends – shareholders get their returns in the form of capital growth. Of course, capital gains tax is then payable on this growth, but this is at least deferred until you actually sell the shares.

Some companies have an explicit dividend policy, usually expressed as a percentage of profits that will be paid to shareholders. For example, a company may have a policy of paying 50% of after-tax earnings per share to shareholders. This would then balance the need to fund its own growth with the desire of shareholders to receive some cash income.

Few companies would pay out 100% of their profits, although there are some that do. Property companies usually do because they have no need to retain profits. But a dividend policy is seldom a commitment to always pay a dividend; if the company is not profitable, it will most likely not pay anything. However, there are exceptions to this: sometimes a company will still pay a dividend out of its cash reserves if it believes the cash is surplus to its needs.

Companies sometimes also pay out a special dividend, which is exceptional and should not be treated as the norm. This sometimes happens when a company gets a large, exceptional cash inflow such as when it disposes of an asset. For example, Old Mutual paid out a special dividend of £1bn in 2012 after it sold its Nordic subsidiary.

Special dividends are meant to be seen as outside of the normal dividend policy and should not be included in calculations of the dividend yield.

Snapshot

A dividend yield is an important indicator, allowing for comparisons between companies and to understand whether the company will be able to generate income, if that is what your objective as an investor is. It should be looked at in

conjunction with the company's dividend policy. And, as with all investing, one has to assess what the prospects of the company are in order to understand whether it will reliably pay dividends in future.

Growth versus value

A dividend yield, like a PE ratio (reviewed in note 2), is used as an indicator of "value" versus "growth" companies. A high dividend yield can indicate a company that is relatively cheap compared with other companies. On the other hand, a company that has a low or zero dividend yield usually indicates a growth company that is expected to grow earnings in the future.

When undertaking fundamental analysis a key theoretical assumption is that shareholders will receive cash from shares at some point in the future. What we are really doing is estimating how much that cash will be and what the right price is to pay for the right to receive that cash. Owning a share really just means having a right to a portion of the future profits of the company. So even when a company pays no dividends at all, we are implicitly assuming that one day it will, and that's why we own the share. A growth company pays no dividends but we are, at least in theory, counting on the company paying out a much larger dividend in future.

One indicator that can tell us whether a company should pay out its dividends or keep them to fund growth is return on equity (see Smart Investor note 6). When return on equity is high it may be preferable for the company to hold on to profits, increasing the amount of equity on its balance sheet, and thereby ensuring higher future profits.

Some investors depend on income from their share portfolio. Retirees, for instance, may want to hold reliable, high dividend yield stocks to provide for their basic needs. For this reason, high dividend yield stocks are sometimes called income stocks. Other forms of yield instruments are bonds which pay out a fixed amount of interest.

For example, Vodacom has over the past few years switched from being a growth company to a value company. It currently (28 May 2015) has a dividend yield of just about 5.6%, which is above average. It pays out most of its profits now, having reached a plateau in its growth rate.

Trading insight

There are two dividend-related issues that can lead to short term share price movements. First are dividend surprises, when a company that is expected to pay a dividend doesn't, or else a company unexpectedly announces a dividend. These are seen

as signals about the health of the company. Second, when a company goes ex-dividend, there is often a disproportionately large fall in the share price. Some traders exploit this by trading around the ex-dividend date.

Examples of dividend yields

We have selected stocks of interest. The low dividend payouts by Naspers and Aspen can be justified as they are growth stocks while Barclays Africa Group and Vodacom have high dividend yields, implying less potential for growth. Massmart's "middle of the road" dividend yield is perhaps reflective of a

company with growth plans but believes it can still return a decent amount of cash to investors, and currently (28 May 2015) has a dividend yield of just about 5.6%, which is above average. It pays out most of its profits now, having reached a plateau in its growth rate.

Five dividend yields on the JSE (large non-property companies)

Company	Dividend yield
Vodacom	5.66%
Barclays Africa Group	5.14%
Massmart Holdings	2.47%
Naspers	0.23%
Aspen Pharmacare	0%



Part 2.5 – In our next note we discuss net asset value.



Smart Investor Basic concepts used in fundamental analysis

“To make sense of the relationship between the share price and the NAV, we often look at the price-to-book ratio, which expresses the share price as a multiple of the NAV/share.”

Part 2.5

Net asset value and price: book ratio

Net asset value (NAV) is one measure of the value of a company. It simply considers its assets net of any liabilities. It is the value that is attributable to shareholders, rather than lenders or suppliers who are owed money. In theory, it is the value that could be derived by liquidating the company, although if this were actually done there would probably be tax liabilities generated that would absorb some of the NAV.

Snapshot

What is it? The net asset value (or book value) of a company is the value shareholders can really count as their own. It is the amount left over once all a company's liabilities are deducted from its assets.

Why are they useful? The NAV is a bottom-line value that does not depend on forecasts of future cash flows and so is often a more reliable indicator of value. A company should trade at least at its after-tax NAV, or else shareholders should demand that it be liquidated.

Why is it risky? It does not take cash flows into account, and while the company's assets may be substantial, it is not always easy to unlock the value for shareholders.

To determine the net asset value, we have to consult the financial statements of a company, in particular the balance sheet, which in most annual reports is called the “statement of financial position”. The NAV is the company’s assets less the company’s liabilities.

For example, consider the following abridged balance sheet (which is from Tiger Brands’ interim report for March 2015).

ASSETS	Rm
Property, plant and equipment	5 680
Goodwill	2 384
Intangible assets	1 060
Investments	3 448

Current assets	11 054
Total assets	24 928

LIABILITIES	Rm
Total equity	14 145
Non-current liabilities	1 313
Current liabilities	9 469
Total equity & liabilities	24 928

To determine the NAV, we deduct liabilities from assets. So:

NAV	=	Total assets – total liabilities
	=	R24 928 – (R1 313 + R9 469)
	=	R14 145 (in Rm)

Snapshot

Variations of NAV are used for certain types of companies. When valuing mines, for example, in-situ NAV can be used – the mine’s measured, indicated and inferred mineral resources, less its liabilities. Some analysts value mines on a weighted

average or multiple valuations using different methods (typically in-situ NAV plus discounted cash flow).

You can see that this is equal to the “total equity” line of the balance sheet, and usually you can determine NAV by looking at the total equity. Sometimes, however, equity includes elements that are not attributable to ordinary shareholders, such as preference share funding. That should be properly considered as a liability. On the other hand, sometimes current and non-current liabilities include shareholder loans. It would then often be better included with equity.

NAV is often expressed in terms of NAV/share. To get this number you divide the NAV by the number of shares in issue (usually in the annual report). At the time of this balance sheet, Tiger Brands had 192 069 868 shares in issue on a fully diluted

basis. That means that the NAV/share is R73.64. The Tiger Brands share price at time of writing (28 May) is R290.85.

To make sense of the relationship between the share price and the NAV, we often look at the price-to-book ratio, which expresses the share price as a multiple of the NAV/share. So in this case, Tiger Brands’ price-to-book ratio is $290/73,64 = 3,95$.

This is a fairly high price-to-book ratio which shows the market values the company at about four times the value of its net assets. So shareholders are expecting it to be very profitable and earn high returns out of its assets.

TNAV

There is one further important measure, the tangible net asset value (TNAV). “Tangible” assets are physical assets you can see and feel, as distinct from intangible assets such as intellectual property and goodwill.

“Goodwill” is an accounting fiction that is used to account for premiums paid when buying another company. For example, were a company to buy Tiger Brands at its current share price and then consolidate Tiger Brands’ balance sheet onto its own, there would be a significant difference between the price it paid and the net assets it would acquire (indeed, it would pay four times the value of the assets). Accountants solve this dilemma with the fictitious asset they call goodwill.

You can see on Tiger Brands’ balance sheet that it has R2 384m-worth of goodwill, which represents the premium it has paid for companies it has acquired over and above the assets of

those companies. But were the company to be liquidated, this fictitious asset would be worth nothing. Also, Tiger Brands has R1 060m-worth of other intangible assets, which may consist of brands, software, patents and so on. We usually don’t think of these as “real” assets either. We are generally cautious about intangible assets because there is a lot of management discretion involved in how they are valued and there have been plenty of examples in the corporate history of shareholders being misled about the values of companies through inflated intangible assets (the Corpcapital and Regal Treasury Bank scandals a decade ago are examples).

We therefore often calculate the TNAV by excluding the intangibles from the total assets. When this is done for the Tiger Brands case, TNAV comes out at R55,71/share. Quite a bit less than the NAV.

Trading insight

Sometimes a share price goes out of whack with its component assets. For instance, the value of Anglo American depends substantially on its exposures to Kumba Iron Ore and Anglo Platinum. The share price of the holding company should exhibit some relationship with the share prices of its subsidiaries. When this relationship breaks down it is often a

signal of a trading opportunity. Other companies have assets in the form of large inventories like commodities of copper. If the value of the inventory changes, the share price should too. Traders can arbitrage these two values.

Why it's useful

NAV is particularly useful when it comes to investment holding companies, property companies and unit trusts, the value of which really do consist of their assets.

For investment holding companies, which consist of a variety of interests in other companies, for example, NAV is a useful tool to determine whether they are trading above or below the sum total of the value of the underlying investments. In theory, the holding company should equal the sum of the parts, less any capital gains tax and dividend tax, and the costs of the holding company management layer.

As is clear, a company's share price is often lower or higher than its NAV. This does not automatically mean it is under- or

overvalued; share prices are a factor of historical information plus expectations of future earnings.

If a company is trading above its NAV per share, it indicates that investors believe it has good growth prospects so are prepared to pay more for it. In other words, they believe the forward NAV will be higher.

Investment holding companies usually trade below their NAV per share – but this does not necessarily mean they offer value. Rather, the discount reflects the added layer of costs involved in an investment holding company structure (auditors, a board and executive committee with subcommittees and administrative costs) and the tax implications of liquidating it.

EV/Ebitda ratio

Known as the “enterprise multiple” or the “Ebitda multiple”, this ratio is a comparison of enterprise value (EV) before interest, taxes, depreciation and amortisation. It values a company, including debt and other liabilities, to the actual cash earnings excluding non-cash expenses. The enterprise value is the same as the net asset value, therefore the formula is:

$$\text{EV/Ebitda ratio} = \frac{\text{NAV}}{\text{Ebitda}}$$

The Ebitda multiple is a better measure of a company's value than the PE ratio because it is not affected by changes in a company's capital structure – for example,

when a company raises equity finance to pay off debt, it usually results in lower earnings per share.

Another benefit is that the Ebitda ratio can be used to compare companies with different capital structures. Finally, the Ebitda ratio removes the effect of non-cash expenses such as depreciation and amortisation and makes cash flows the key component.



Part 2.6 – In our next note we look at a key measure of profitability, return on equity (RoE).



Smart Investor Building insight into trading shares

“A low RoE suggests that shareholders would be better off if the company gave the cash back to them, while a high RoE suggests shareholders would be better off giving the company more cash.”

Part 2.6

Return on Equity (RoE)

This is the first in a series of notes on ratios that reflect the profitability of a company, beginning with return on equity.

To determine RoE, simply divide the company’s headline earnings by shareholders’ equity. That is the practice in South Africa, but in other countries profit after tax is used as the profit measure. The headline earnings measure removes some once-off income statement lines, so it attempts to capture sustainable earnings before dividends are paid out. You can find it in any company’s reported financials as well as the net equity figure. Shareholders’ equity, found in the balance sheet,

excludes preference shares and is usually used as an annual average – the average of the opening and closing balances on the balance sheet.

The formula is:

$$\text{RoE} = \frac{\text{headline earnings}}{\text{shareholders' equity}}$$

Snapshot

What is it? RoE represents the profits the company is able to generate as a percentage of the shareholder funds that it holds.

What does it do? Companies have a choice on how to structure their balance sheets. If they hold shareholders' funds – often

called equity – that money should be working for shareholders. RoE shows how hard that money is working for shareholders. It is one measure of the efficiency of the balance sheet, showing that it has the right balance of debt and equity.

Why is it useful?

For one, it shows how well a company is using money that belongs to shareholders. A low RoE suggests that shareholders would be better off if the company gave the cash back to them, while a high RoE suggests shareholders would be better off giving the company more cash.

RoE is also a useful input to guide forecasts of a company's growth. If the amount of shareholder capital a company holds is the main determinant of its ability to grow, RoE indicates how profits may grow in future. For example, if a company has an RoE of 10% and its shareholder funds double from year 1 to year 2, we would expect the profits to double as well, provided the RoE stays the same.

Note, however, that the usefulness of RoE depends on what industry you are looking at as well as the specific company. Some industries are not very capital intensive and so have higher RoEs, for example, services companies which have very

little capital equipment they need to finance. Other companies are very capital intensive so have low RoEs, such as mining or telecoms companies. Capital-intensive companies often have less competition, because having a large balance sheet provides an effective barrier to entry while high RoE companies often face more competition over time. For that reason, RoE is most useful when comparing companies within an industry or when judging the performance of a company over time.

In some companies the amount of shareholder capital is a more critical determinant of growth than in others. For example, banks are legally required to hold a certain amount of capital for every loan that they make. So having more equity increases the amount of lending they do. RoE therefore becomes important in determining whether the bank should be increasing or reducing its lending. A bank with a high RoE should be increasing shareholder funds in order to expand its lending potential.

Snapshot

RoE can be a useful tool when assessing companies. It is particularly useful to study trends in a company's performance – an improving RoE shows a company that is improving

the efficiency of its balance sheet and its profitability. It is also a useful tool to start to predict the growth trends in a company.

The Du Pont formula

Financial analysts often break RoE down into three different parts in order to analyse what causes changes in RoE over time. Remember that the basic formula (above) has profits at the top and shareholder funds at the bottom. By factoring out that equation, one can tease out the different things that affect RoE. The basic three-term Du Pont formula is:

$$\text{RoE} = \frac{\text{net income}}{\text{sales}} \times \frac{\text{sales}}{\text{average total assets}} \times \frac{\text{average total assets}}{\text{average shareholder equity}}$$

You can "cancel out" the same numerators and denominators in the Du Pont formula to end up back at the basic RoE equation again. But teasing out these three terms in the equation gives us three well-known ratios:

- Net profit margin is the profit ("net income") as a percentage of sales
- Sales over total assets represents the total asset turnover
- Total assets over shareholder equity represents the financial leverage ratio and indicates how indebted the company is

The Du Pont formula helps identify what is driving RoE in a company. High profit margins, a high turnover relative to the assets on the balance sheet, and a high leverage ratio translate into high RoEs.

If the net margin increases, it means every sale brings in more money, resulting in a higher overall RoE. Similarly, if the asset turnover increases, it means the company generates more sales for every unit of assets owned, again resulting in a higher overall RoE.

Market example

We've calculated the Du Pont formula for Mr Price, the JSE-listed retailing company, for its 2014 financial year (all these figures come from the annual reports on the Mr Price website). The company had equity of R3 992m at the end of 2014 and R3 309m at the start, so the average for the year was R3 651. It had headline earnings of R1 888 at the end of the year. This means RoE was a very healthy 52%.

To make sense of the key drivers of this figure, we can break down the RoE into the Du Pont formula (all Rm): Now we can

$$\frac{R1\ 888}{R15\ 892} \times \frac{R12\ 122}{R5\ 731} \times \frac{R5\ 731}{R3\ 651}$$

From this formula we can determine that Mr Price's profit margin is 12%, its asset turnover is 2,1 times and its leverage ratio is 157%. But if we look at the same formula for Mr Price for the year before, this is what we get:

$$\frac{R1\ 537}{R13\ 720} \times \frac{R13\ 720}{R4\ 597} \times \frac{R4\ 597}{R3\ 049}$$

Problem areas

As with most financial ratios, there are variations in ratio definitions which may cause problems.

Numerator (the profit measure at the top): Which profit number should be used, before or after taxes? How should it be adjusted for non-recurring items? Should operating profit figures be used (earnings before interest and taxes, or earnings before interest, taxes, depreciation and amortisation) or earnings attributable to shareholders?

Denominator (the measure of equity at the bottom): A balance sheet figure is for a specific point in time (the year-end or interim period) while the numerator is for a full reporting

period. While the most common usage is for an annual average, some use equity at the beginning of the year or the figure at year-end. In certain companies there may be difficulties with knowing which figures to include in equity, such as when the company is financed with shareholder loans (which could be either debt or equity, depending on the context).

Increased debt will make a positive contribution to a company's RoE only if the matching RoE of that debt exceeds the interest rate on the debt.

see the profit margin is 11.2%, the asset turnover is 3,0 times, and the financial ratio is 150%.

This gives a total RoE of 50%, so RoE improved between 2013 and 2014. And we can see that it improved because profit margins increased (from 11.2% to 12%) but its asset turnover fell from 3 to 2.1. This means it is holding more stock relative to the sales it is making in 2014 compared to 2013. Its leverage ratio was also slightly higher in 2014, which helped RoE.

So this exercise tells us some important things. Mr Price did well to increase its sales margins, but did less well in the efficiency of its stock.

From here, to forecast Mr Price's earnings we'd want to know what was going to happen to margins in the year ahead, whether it would turn over its stock faster and whether it would increase debt. Getting answers from management on those questions would allow us to forecast what growth is going to be and therefore what future profits will be.

Be careful of using the list of computer-generated RoE figures provided for companies on some financial websites – if the company is running at a loss and if the shareholders' equity itself becomes negative, the RoE figure will still read as a positive because you're dividing a negative figure with another negative figure, but such a company would be technically insolvent!



Part 2.7 – We examine cash flows, the numbers that never lie.



Smart Investor Basic concepts used in fundamental analysis

“Assessing the relationship between profits and cash flows is an important measure to verify what a company is reporting in its financial statements.”

Part 2.7

Cash flow – the king of financial analysis

Understanding the implications of a business’s cash flow on its sustainability is an invaluable insight that can help you better assess the performance of a company.

A company’s accounts can get murky very quickly. Accounting works on the principle of accrual, which means costs and income are recognised at the time they are incurred, rather than when the cash is actually received or spent. Investors like to say “cash never lies” because the actual receipt of cash is when you know for certain that money has been earned.

Aggressive managers can find ways to manipulate the accrual system to show profits that will never actually lead to real cash.

For instance, a company may report higher sales and profit by virtue of a sales campaign that offers liberal credit terms, although this may lead to liquidity problems or bad debts down the line that might not be immediately apparent.

For this reason, we explore the different types of cash flow that will be reflected in a company’s financial statements under the Statement of Cash Flows section. These different measures produce a cumulative picture of cash in and cash out.

Snapshot

What is it? Cash flow is an indicator of the in and outflow of hard currency in a business.

Why is it useful? Cash really doesn't lie. Assessing the relationship between profits and cash flows is an important measure to verify what a company is reporting in its financial statements.

Why is it risky? Cash in the bank is a metric that can be checked fairly easily, although it is equally important to examine the sources of cash flow. A company that has negative cash flows from operations must be either selling assets or raising finance from somewhere to sustain itself.

Operating cash flow

The basic source of cash for a company is its sales, and the basic use of cash is payments to suppliers. These cash flows make up the operating cash flow of a firm. Cash generated from operations differs from the turnover or revenue figure in that it strips out the effect of sales on credit or credit received from suppliers.

The cash flow statement shows you this figure, but is related to the income statement and balance sheets by adjusting revenue for changes in current debtors and creditors. Any income received from the revelation of investments and other non-operational items also need to be removed from the calculation to get an accurate operating cash flow picture.

Shoprite Operating cash flow	June 2014 (Rm)
Plus increase in payables	3 658
Interest received	252
Interest paid	-342
Dividends received	30
Dividends paid	-1 868
Income tax paid	-937
Total operating cash flow	5 720

For example, consider the cash flow statement of Shoprite above. It is a high-volume cash retailer, so its cash flows are quite considerable. The cash generated from operations is slightly higher than the operating profit figure (R5.708bn compared to R5.72bn). Looking at the reconciliation of the two shows why this is the case. Non-cash costs include depreciation and amortisation which are costs, but are not cash costs, so these can be added back to the operating profit. But the company saw an almost R2bn increase in

inventories, probably because it opened new stores during the period. It also had more receivables, which is the cash that is owed by its customers, so although it books this as income, it doesn't get the cash. What really helped matters, though, is that Shoprite squeezed suppliers with a R3.6bn increase in its payables – the money it owes suppliers. This is good for cash flows, but can be unsustainable if suppliers start pushing back.

Investing cash flow

These items are an important indicator of a company's investment in new factories and machines, subsidiary companies, any financial instruments like bonds and

shares as well as any proceeds from the sale of any such assets. Using the example of Shoprite again, its cash flow from investing activities was:

Shoprite Operating cash flow	June 2014 (Rm)
Investment in property plant and equipment to expand operations	3 658
Property, plant and equipment to maintain existing operations	252
Proceeds from the sale of property, plant and equipment-	342
Other investing proceeds	30
Investment in associates	-1 868
Acquisition of subsidiaries and operations	-937
Total cash flow utilised by investing activities	5 720

Here you can see that investing cash flow was strongly negative at -R4.2bn largely because Shoprite has been investing in expanding its operations, such as opening new

stores, at a cost of almost R3bn. It sold some assets to generate a relatively small amount of R126m of investing cash inflow. Everything else was negative.

Cash flow from financing activities

The last item reflecting monetary inflows is cash received from shareholders or lenders that a company receives in return for issuing new shares in the first case and through

issuing bonds or obtaining loans in the latter. In the case of Shoprite, cash flow from financing activities was:

Shoprite Operating cash flow	June 2014 (Rm)
Proceeds from ordinary shares issued	0
Proceeds from convertible bonds issued	224
Increase in bank loan	229
Total cash flow from financing activities	453

So here we can see that Shoprite borrowed R453m during the year to finance its business, through a combination of issuing bonds and bank loans. By looking at the cash flows broken down into operating, investing and financing, we can see that

Shoprite is using the cash flows from its operations (largely generated by squeezing suppliers) to fund its investment in new stores. Some borrowing has also added a bit, while it has also paid out a fair bit in dividends.

Interest, dividends and tax

An area of cash flow accounting that causes some confusion is the treatment of interest and dividend payments. The reason for this is that it is not always clear where these inflows are accounted for within the framework of the three types of cash flow explained above.

In the case of Shoprite, it specifies its interest, tax and dividend payments as part of operating cash flow, shown above.

There is greater clarity for financial institutions in accounting for this as such inflows should be treated as operating cash flows, given the nature of the business model. But in other companies it is less clear and accountants are usually free to decide for themselves.

The sum of all these calculations is reflected in a company's financial statements as the net movement in cash and cash equivalents. In Shoprite's case that was a positive R2bn. That is a positive sign for investors and indicates high quality earnings in a business that is able to fund its growth out of operating revenue.

The games execs play

South African corporate history is littered with examples of companies that have booked big profits, allowing execs to pay themselves big bonuses. But when the cash flows are

examined it becomes clear they have been selling assets or simply handing out credit to boost sales figures. Those types of games can only be played for so long.

1:1 Part 2.8 – Discounted cash flow modelling.



1:1

Smart Investor Basic concepts used in fundamental analysis

“The accrual principle in accounting can lead to a company’s cash flow from operations differing from its earnings.”

Part 2.8

Cash flow ratios

In our last note we discussed cash flow statements and how to understand them. In this note we consider some analytical tools that can be applied to cash flow statements in order to tease out important information about companies’ performances.

The cash ratio: operating cash flow/operating profit ratio

This ratio is the starting point in assessing the quality of earnings of a company. It is a quick way to determine just how much of a company’s operating profit comes from non-cash items.

Operating cash flow is one way to assess the quality of earnings of a company. Alarms should go off if it’s booking profits but not receiving any cash. Cash flow is also important

to understand the company’s liquidity position. You can assess whether it is able to meet all its obligations when due, particularly payments to lenders. Cash flow also provides an important measure of returns to shareholders: the free cash flow. We discuss how to use each of these below.

The ratio expresses the total “cash flow from operations” (which can be found in the cash flow statement) and the total revenue as a percentage:

$$\frac{\text{Cash generated from operations} \times 100}{\text{operating profit}}$$

Snapshot

What is it? Cash flow ratios provide powerful financial analysis tools that help determine how healthy a company is, particularly regarding earnings quality.

Why is it useful? Cash flow is important to understand a company's liquidity position as it will help you assess whether it is able to meet all its obligations and provides a critically

What to look out for? The accrual principle in accounting can lead to a company's cash flow from operations differing from its earnings. Clever accounting can therefore make a company's profits look really good, but in reality they are accounting oddities rather than real cash earnings.

When calculating this ratio, we use cash generated from operations excluding any interest and dividend payments (we discussed this issue in our previous note).

In our last note we discussed Shoprite as an example and showed it had total operating cash flow of R5.72bn while its operating profit was R5.708bn. This implies a cash ratio of slightly more than 100% because it gets more cash than its profits – 100.2%, to be exact.

Companies with a ratio higher than 100% show healthy cash-generating ability.

For cash retailers such as Shoprite, this ratio tends to be above the 100% mark if they manage their creditors 'correctly' by paying them as late as possible. As we saw in the previous note, Shoprite generated a lot of cash by paying its suppliers later than the previous year.

Trading insight

Ultimately the value of any financial asset depends on the cash flows it delivers to the owners of that financial asse

If the ratio is below 100%, check to see why. Usually it's because the company has been selling more on credit. Then, check if the company is able to efficiently obtain the cash for such sales. In other cases it is because it is revaluing assets,

an issue common in investment companies. Sometimes it is legitimate, but valuations upwards are always rather convenient for management. Be sceptical.

Danger alert

The cash ratio can pick up serious problems in reported results. One celebrated example was Corpcapital, which in 2001 reported an operating profit of R387m but operating cash flows of R205m – so its ratio was 53%. This raised alarms and a study of the cash flow statement made it clear that significant income was being booked for adjustments to the value of investments the company held.

In effect it was booking profit by revaluing its investments – a non-cash form of income. One in particular was highly suspicious – an offshore entity that had not been audited and had suddenly been valued upwards by R500m. The profits that year allowed directors to pay themselves R14m in bonuses. Two years later that value had to be written off. Shareholders, needless to say, were not pleased.

Cash flow coverage ratio

This ratio is of most interest to lenders to a company, like its banks and bond holders, who want to assess the company's ability to meet its financing costs every year.

Remaining with the Shoprite example, in 2014 cash generated from operations was R5.72bn and total debt was R4.7bn. That gives a cash flow coverage ratio of 122%. This is a very strong balance sheet as Shoprite can effectively pay back all its debt and then some with one year of operating cash flow. However, Shoprite has a large amount of creditors in the form of suppliers, and it owes them R16.3bn. If you include that as debt, the ratio falls to 27%. That is less comforting and indicates that creditors

cash flows to the total debt, expressed as a percentage:

$$\frac{\text{Cash generated from operations} \times 100}{\text{total debt}}$$

should be more cautious about Shoprite's liquidity position. If all its suppliers started demanding faster payment, Shoprite could find it challenging to do so.

The cash flow coverage ratio is also useful to shareholders because companies with a low ratio could be expected to be cash conservative and probably limit investment in growth or pay out cash to shareholders.

Free cash flow

Free cash flow is not a ratio, but can be used in a variety of different analytical approaches to price a company. It is important to investors as the figure indicates the amount of cash available to distribute to shareholders because it is the money remaining after a company has settled all its obligations, including financing costs and the costs of investing in plant and equipment to continue operating at current levels.

It is calculated by deducting investing cash flow from operating cash flow, except where the investing cash

flow is in new business areas. It is what the owners of the business can call theirs – it is the cash that is surplus to the needs of the business to keep its current operations going. If it chooses to use some of that cash to fund its growth, shareholders should be pleased provided that the growth will be in ways that will add to profits and thereby increase the cash flows shareholders get in the future.

Ultimately the value of any financial asset depends on the cash flows it delivers to the owners of that financial asset.



Part 2.9 – Next we consider an important way of valuing companies by discounting their free cash flows.



1:1

Smart Investor Basic concepts used in fundamental analysis

“Forecast periods can be of varying lengths, but it must be remembered that cash flow estimates become less accurate the further into the future you project.”

Part 2.9

Discounted cash flow valuations

In this note we discuss one of the most important methods of valuing shares: discounted cash flow. It is a key way to figure out what a share is “worth” so that you can compare it to the share price to look for value.

If you know you are going to get R1000 in one year’s time, what is the current value of that cash flow? Obviously if you could have R1000 now rather than in a year’s time, you’d take it now, so something must compensate you for the wait.

The idea behind cash flow discounting is that there should be a discount applied to future amounts to get a present value. So for instance, if you used the discount rate of 10%, the present value would be R900 to get the cash flow of R1000 in a year. We would say that you’d pay R900 now to receive R1000 in a year, which implies a discount rate of 10%.

One way to think of it is as an interest rate. If you could get an interest rate of 11%, you’d deposit R900 in a bank now in order for it to be worth R1000 in one year. You are “buying” a future value of R1000.

This is the basic idea behind the valuation of all shares and other financial assets. All financial assets are basically pieces of paper that give the holder a right to a future stream of cash flows (although now those pieces of paper are actually just electronic records). A bond gives you the right to interest payments plus the return of the capital amount at the end of the bond term. A share gives you the right to a share of the profits of a company for as long

as the company exists (and to its assets if it ever winds up). Financial assets are distinct from other assets such as gold, jewellery and artworks that don't necessarily generate any financial returns.

When undertaking cash flow analysis there are two big questions that need to be answered:

1. How much are the cash flows that we will receive in future?

2. What discount rate is reasonable to apply to those cash flows to determine the present value?

There is no end to the discussion about ways to estimate these two key measures, and we're only going to give a brief overview. Then we'll explain the way financial analysts calculate present value once these estimates have been determined.

Snapshot

What is it? Cash flow discounting is one of the key ways to calculate the value of any financial asset. The principle behind it is one of the most important ideas in financial theory – that the value of any asset must be the present value of its future cash flows.

Why is it useful? A discounted cash flow valuation typically requires more effort than most relative valuation techniques, but it provides a good picture of the key drivers in the share price: expected growth in earnings, capital efficiency, cost of equity and debt, and expected duration of growth phases.

Estimating future cash flows

These estimates are especially important, and rather complicated, but help to provide some clarity on a company's expected future profit. With the insights you've gained in this series of notes, you are well-positioned to start estimating how a company's profits will grow in future. The key issues to think about are:

Is the company in a growing market with increasing product sales? Does it have the ability to increase prices and margins? Is it investing in its own growth by expanding production and does it earn a good return on equity (RoE)? Also, is equity growing?

A company's historic growth rate can give some indication of future growth trends. This is part art, part science, but an understanding of the profit drivers in a company and the market it operates in allows investors to narrow down the range of their projections.

Certain websites carry the consensus forecasts of analysts who calculate expected earnings per share. These are a good guide to forecasts.

Our focus is on estimating "free cash flow" which we discuss in our previous note, number 2.8. That is the cash that is surplus to the company's needs to maintain current operations.

Estimating the discount rate

Determining the discount to apply is essential to calculating the future value of a company's cash flow.

Two factors will determine the extent of that discount. Firstly, the prevailing interest rate because the return from investing in a company's shares has to be higher than a risk-free deposit in a savings account.

The second factor to consider is the risk inherent in investing in a particular company. Companies with highly volatile earnings and big risks to future profits should compensate you for those risks, which can be determined to some extent by looking at

historic share price movements. If a share price has been very volatile we assume the company is very risky.

Bearing these factors in mind then, an investor can apply a higher discount on future cash flows when both the interest rate and risks associated with a company are higher.

The discount rate indicates the expected return from an investment. So, once you have estimates of cash flow and the discount rate, you can calculate a present value for a financial asset. If this present value is above the share price, you buy, and conversely if it is below the share price, you sell.

Trading insight

A key advantage of this type of valuation is that they are less likely to be affected by earnings manipulation or aggressive accounting as they are based on cash flows, which are hard to

manipulate. This valuation is useful for comparing companies with diverse capital structures.

The discounted cash flow (DCF) calculations

Here's the formula to use to make these calculations. It looks complicated, but it contains a simple idea!

Here is an explanation of what the various abbreviations represent:

Free cash flow (FCF) – Cash generated by the business available for distribution to all providers of capital (debt and equity).

Discount rate (r) – The rate used to discount projected FCFs and terminal value to their present values.

Growth rate (g) – The rate at which free cash flows are projected to grow into the future.

Essentially, what this formula says is that the present value

is the sum of each year's free cash flow, discounted by the discount rate:

Here are the two basic ideas:

1. All that matters when owning a financial asset is the cash flows that ownership gives you access to. The tricky bit is figuring out just what those cash flows will be. In the case of equity (rather than debt), future cash flows are less certain.
2. Before making an investment, you should ensure that the yield from that investment compensates you for the risk. The more certain you are about the cash flows, the less you need to demand to make the investment.

The DCF formula determines the total enterprise value. From this number we have to deduct total debt, the part of the enterprise value that is "owned" by the debt providers. Whatever is left after that is owned by the shareholders. That number can be divided by the number of shares in issue to obtain a value per share. That is the key number to compare to the market price to tell if it is worth buying.

Some more insight into making estimates

It is important in the first step to determine the forecast time period and expected number of phases. For example, you might identify two stages of growth: a high growth stage and thereafter a low growth phase; or you may use a three-stage growth phase with three periods: high, moderate and low growth.

Forecast periods can be of varying lengths, but it must be remembered that cash-flow estimates become less accurate the further into the future you project. So an estimated terminal value is as good as trying to make yearly estimates.

The quality on the inputs in the valuation will determine the quality of the output, and unfortunately the process is not a simple case of inputting the numbers into the formula because subjective projections and some assumptions are required.

The discount rate that results from the calculation is the rate at which the free cash flow in each of the periods is discounted to a present value. This can also be expressed as a discount factor, which is a single figure by which each cash flow is multiplied to arrive at a present value.

The good

You cannot know whether a company's share is undervalued or overvalued until you have analysed it, and a DCF valuation method is one of the most important methods of estimating a company's intrinsic value. This method of valuation can be checked against other valuation techniques, such as relative valuations derived from such measures as the price-to-earnings, price-to-book value multiples, and other valuation methods we have discussed in previous Smart Investor notes.

While a DCF valuation typically requires considerably more effort than most relative valuation techniques (like price to book or price to earnings ratios), it does provide a good picture of the key drivers in the share price: expected growth in earnings; capital efficiency; balance sheet capital structure; cost of equity and debt; and expected duration of growth phases.

The good

DCF is only as good as the estimates that are made. The further into the future one is trying to estimate, the less accurate the results are going to be. Present values are also highly sensitive

to the discount rates that are estimated. The vast majority of analysts' estimates turn out to be quite wrong.

The bad

DCF is only as good as the estimates that are made. The further into the future one is trying to estimate, the less accurate the results are going to be. Present values are also highly sensitive

to the discount rates that are estimated. The vast majority of analysts' estimates turn out to be quite wrong.

The bottom line

DCF is always a worthwhile exercise even if it is quite hard work to do properly. At the very least it reveals what the market currently believes about future earnings, which can help you

make judgment calls over whether the share price represents value or not.

 **Part 2.10 – Next we discuss important issues in thinking about a company's risk level.**



Smart Investor Basic concepts used in fundamental analysis

“In common usage, “risk” refers more to the fear that we could lose some or all of our investment, not how a share price gyrates around the average return.”

Part 2.10

Profiting from risk (a)

We have focused mainly on fundamental analysis so far in this series, looking at how to calculate expected returns. The other crucial factor in any investment decision is risk and there is invariably a trade-off between risk and returns. Any investment offering wondrous returns probably carries much higher risk than an investment that promises low returns.

This is the first of a two-part series focused on understanding risk, measuring it and learning how to profit from it. We start with analysis of equity market risk, and the all-important measure called Beta.

Snapshot

What is it? The probability that you are actually going to get the returns you expect.

A useful rule of thumb: The higher the expected returns, the higher the risk.

Remember: This trade-off between risk and return is a key feature of investments and has been implicit in much of the discussion in the previous Smart Investor notes.

Risk and investment

The use of statistics to analyse financial decisions was introduced for the first time by economist Harry Markowitz in the 1950s. He took the concept of volatility from statistics and used it to explain how investment decisions can be optimised by treating volatility the same as risk.

Markowitz showed that it is not just returns that matter, but the volatility of returns. Investors should aim to maximise returns but minimise volatility, and portfolios can be constructed to do this efficiently. That insight earned him a Nobel prize decades later.

A share that has big swings up or down of, say, 10% each day is a highly volatile one. A share that has swings of less than 1% is

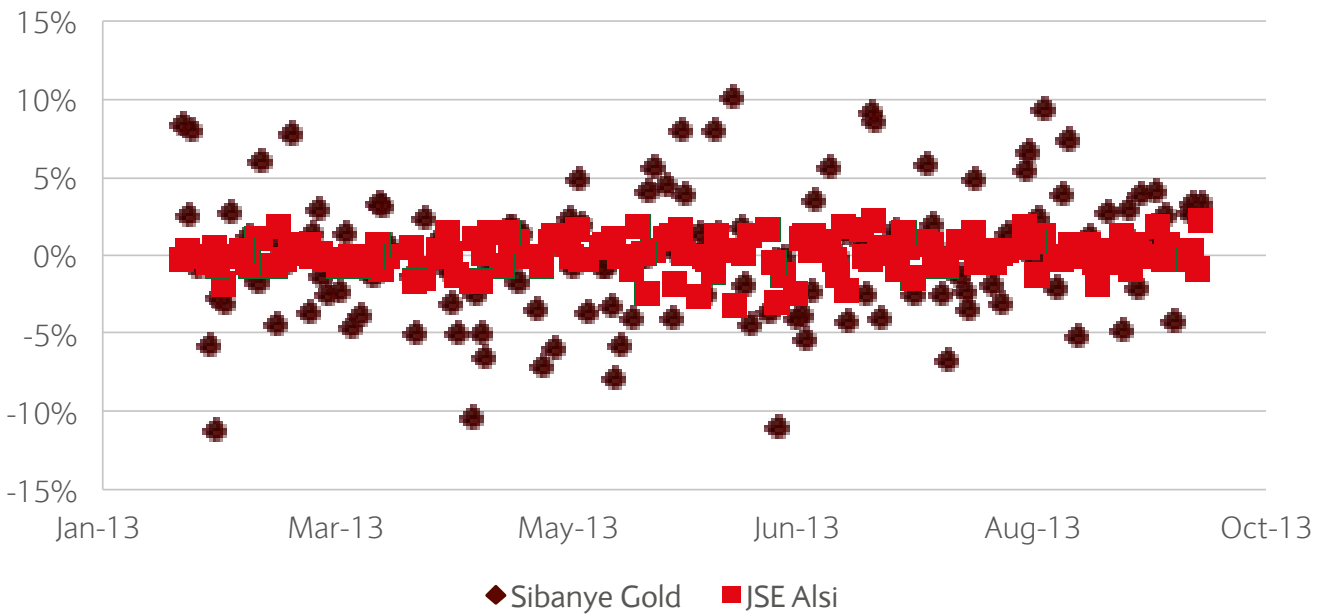
far less volatile. This seems to capture something about what we mean when we think of risk. A highly volatile share tends to keep an investor a tad stressed.

A clear example of this is reflected in the chart below showing the daily movement in the share price of Sibanye Gold versus the All Share Index. Each red dot shows the movement in the closing price of the share. In the period covered, there a numerous days with a swing of nearly 10% both up and down, while there are three days with a swing down of 10% or more and one day with a swing up of 10%. It is clear that Sibanye Gold is more volatile than the All Share Index (green dots).

Risks: low and high

A bond is a low-risk investment – it pays a fixed amount of interest that is usually quite low. At the other end of the scale are high-risk equity investments such as mineral exploration

companies that could make a fortune if they strike oil but are worthless if they don't. This range of risks can be multiplied or lessened through the use of derivatives.



Daily returns of Sibanye Gold (blue) versus the All Share Index (red)

In common usage, “risk” refers more to the fear that we could lose some or all of our investment, not how a share price gyrates around the average return. Certainly we don’t consider the possibility that returns will be too high as a risk factor. But statistical volatility is a measure of dispersion around the mean – whether it is up or down. That is why there is a trade-off between risk and return – volatile returns tend to average out at a higher rate over time than non-volatile, stable returns.

Fundamental analysis has many different way of thinking about financial risks. These include how a company manages risks it faces, such as credit risk, exchange rate risk, commodity price risks, interest rate risk, liquidity risk and so on. Of course, there are also non-financial risks such as the risk of a factory burning down, and part of analysing a company is checking to ensure that it can manage those risks through insurance and other tools. In the next Insights article on risk we will take up some of these issues.

Equity market risk

Markowitz's idea was to create "efficient portfolios" by combining different assets to minimise variance (a kind of volatility) at the same time as maximising return. "Efficient" meant you couldn't get any more return without sacrificing some variance, or get less variance without sacrificing some return. This relies explicitly on the concept of a trade-off between risk and return. The problem with Markowitz's approach is that even once you understand the volatility of an individual asset, you'd need to work out how volatile all the possible asset combinations would be. Each combination is a possible portfolio and each portfolio has a unique profile stemming from the covariance of the assets within it.

Computers at the time couldn't handle the calculations. In the 1960s, one of Markowitz's protégé's, William Sharpe, developed what we now call the Capital Asset Pricing Model. It drew on Markowitz's idea of volatility but produced a simplified approach to managing it.

Sharpe began with a basic assumption: the market itself is the most efficient portfolio – every single available asset in a single portfolio. The JSE's All Share Index (Alsi), would be taken as the market, as it includes all eligible listed shares. The question then is how an individual asset correlates with the volatility of the Alsi – is it more or less volatile?

Beta

The underlying theory in using Beta stems from the efficient markets hypothesis, which says share prices really do reflect the actual risk in that stock. While that is debatable, the way we like to think of Beta is as a reflection of more fundamental risks in a company.

We know from previous notes that a share price is merely the present value of future cash flows of that company. So changes in the share price (volatility) should reflect changing views over what the future cash flows will actually be. A utility company

This question is answered by a measure called Beta, which also comes from statistics. Its mathematical formula is:

$$\beta = \frac{\text{cov}(r_a, r_m)}{\sigma_m^2}$$

It may look complicated but the idea is simple: Beta is the covariance of the returns from a particular asset with the returns of the market. $[\text{cov}(r_a, r_m)]$

It measures how closely volatility in the returns track each other – divided by the variance of the market as a whole (σ_m^2).

The number comes out as 1 if the asset is just as volatile as the market. A Beta of 1.1 means the asset is 10% more volatile than the market, while a Beta of 0.8 means it is 20% less volatile than the market.

The important idea of Beta is that it measures the risk that cannot be diversified away without sacrificing returns. A high Beta stock therefore must be one that offers higher returns than the market and a lower Beta stock must offer returns lower than the market. It is important to remember that the market itself is still volatile. All Beta is saying is how a particular investment should perform relative to the market as a whole. It is only a relative, not an absolute, measure of risk.

By looking at the Beta we can then get an idea of just how much return we should expect to receive for holding a particular stock. That is the essence of Sharpe's Capital Asset Pricing Model.

that has a monopoly and whose customers have little choice but to buy its output will have a very predictable cash flow stream and therefore a low Beta. Pharmaceuticals and mining exploration companies, though, depend on big discoveries in order to generate profits, so can be highly volatile and therefore have a high Beta. High Beta stocks are also those which have high financial leverage or high fixed costs relative to variable costs, and so struggle to absorb changing market conditions.

Control risk with Beta

What is important is that Beta is a useful way to control risk in a share portfolio. If we want to lower risk then we add low Beta stocks, or add high Beta stocks to increase risk. Low Beta

stocks are better if you are expecting a recession, whereas high Beta stocks are better if you are expecting a market rally.

Risk in Beta

We measure Beta by looking at historic return figures and applying the statistical tools to them. It does not look forward. Immediately this raises problems with some stocks – for example, if a company is a takeover target, its share price tends

to shoot up to the level at which the offer is made, or collapse if the deal does not happen. In such a case, past volatility has little bearing on future volatility.

Finding Beta

It is not hard to calculate Beta using the formulas above and Excel, but Beta is often provided by various financial data

websites, such as the Reuters website where it is available for free.



Part 2.11 – In our next note we discuss risk and the Sharpe ratio.



Smart Investor Basic concepts used in fundamental analysis

“Good investors understand the trade-off between risk and return and aim to maximise it.”

Part 2.11

Risk and the Sharpe ratio (b)

In Smart Investor 10, we learnt to see risk as an essential component in any share’s expected returns. The important point is that we should expect to be compensated with higher returns for taking on extra risk.

Snapshot

What is the Sharpe ratio? It tells us how much risk an investor took on for the return earned. Why it’s useful: The Sharpe ratio

is often used as a way to assess professional investors as it reveals how well they manage risk.

Managing risk

The previous note showed how risk and return stand in opposition to one another. Risk is determined from Beta, a statistical measure of the relative volatility of an asset compared to the market as a whole, or volatility which measures the risk of a share in absolute terms. Good investors understand the trade-off between risk and return and aim to maximise it.

In order to assess how investors are doing, we use a tool called the Sharpe ratio, named after the economist who invested

The Sharpe ratio

The Sharpe ratio tells us how much return is generated for a given unit of risk. As we discussed in the last note, a “unit of risk” is the standard deviation of the returns of an investment. Standard deviation is the distribution of returns around the average return. So if the share delivers an average daily return of 0.5%, but often varies significantly from that, then the standard deviation is higher than a share

The Sharpe ratio’s formula is:

$$\text{Sharpe ratio} = \frac{r_i - r_f}{\sigma_i}$$

(which is the rate of interest we could earn for taking no risk at all – i.e. putting it in a bank deposit or a government retail bond), divided by the standard deviation of the). We won’t go into how to calculate standard deviation in this note, but it is easily researched online and can be calculated using an Excel spreadsheet. It is usually expressed in percentage points, reflecting how much the returns vary from the mean. So the ratio tells us how much in excess returns an investment has generated for each percentage point of standard deviation. The higher the ratio, the better the investment.

Looking at the returns per unit of risk enables us to compare investments that have very different risk profiles. A portfolio of small-cap shares, for instance, will have a higher standard

Examples

You can, however, also use the Sharpe ratio to look at the returns of individual stocks. We’ve done that for a few popular South African shares in the table below.

the Capital Asset Pricing Model in the 1960s (he called it the “reward-to-variability ratio”). The ratio is often used as a way to assess professional investors as it reveals how well they manage risk. It is often used to assess unit trust fund managers, for example. The ratio tells us how much return the investor earned for the amount of risk taken on. Minimising risk while maximising returns is an important objective. If two fund managers achieve the same returns, you’d want to be with the one that took on less risk.

that consistently delivers 0.5%. A share with a high variance (variance is just standard deviation squared) is considered higher risk than one with low variance. Once we have the standard deviation of the share we then consider the returns and how much excess returns are earned per standard deviation.

deviation than a portfolio of government bonds, but the returns per unit of standard deviation may still be higher, making the small-cap portfolio a better investment.

The Sharpe ratio is particularly useful for assessing portfolios, because it can draw out the skill of the portfolio manager in combining investments so as to lower the overall volatility of the portfolio. The ultimate aim is an “efficient portfolio” – one that could not be any less risky without giving up some returns, or deliver no higher returns without adding some risk. We can compare managers by looking at the Sharpe ratio, especially when the two managers have the same investment mandate and available investment instruments.

Share	Average annual return	Standard deviation	Sharpe ratio
Naspers	37.83%	36.53%	0.87
MTN	15.26%	13.49%	0.68
FirstRand	10.34%	25.71%	0.17
BHP Billiton4	.95%	17.56%	-0.06
Standard Bank3	.43%	15.44%	-0.17
Anglo American	-8.74%	21.30%	-0.7

(Based on annual returns. Source: Intellidex)

So the Sharpe ratio tells us that, for example, Naspers earned 0.87 percentage points of return for each standard deviation.

Naspers is therefore the most efficient stock on the list in terms of its risk/return ratio.

Variations to the Sharpe ratio

William Sharpe won the 1990 Nobel Prize for the Capital Asset Pricing Model. In 1994 he went further: he decided we should rather be looking at standard deviation of returns over the investor’s benchmark. So if a fund manager is being measured against a benchmark’s returns, his fund’s risk should also be measured relative to the risk in the benchmark. Most managers, however, still use the original Sharpe ratio. Another useful variation is called Roy’s ratio. This defines the top

part of the equation (numerator) as the returns above a given minimum return, rather than returns above the risk-free rate. This is useful if you have a particular return objective such as being able to fund your retirement. Roy’s ratio allows you to see how efficient the manager is being at using your portfolio to generate excess returns: a bonus to the minimum returns you need for your investment goal.

Problems with the Sharpe ratio

Like much of financial analysis, it rests entirely on historic data. What a share has done in the past is not necessarily what it is going to do in the future. Past performance may, however, give us some useful insight into the nature of the share. One with a low standard deviation usually reflects a company with reliable and predictable earnings. But it is always important to think about how reliable the past performance is as a guide to the future. A company that makes major changes, such as an

acquisition, will probably have a different risk profile in future. There are also problems with the use of standard deviation as a measure of risk. The technique makes some assumptions about how returns are distributed, what is usually called the “normal” distribution, which assumes, among other things, that returns are symmetrical; i.e. that returns deviate from the mean upwards just as often and by just as much as they do downwards. But we know in fact that share prices tend to fall faster than they rise.

Final thought

We have mentioned just some of the variations to the Sharpe ratio here – there are others that have been developed in an attempt to compensate for weaknesses in the Sharpe ratio. However, the Sharpe ratio remains one of the most commonly used methods in finance today to assess

performance. While it would be worthwhile using it on your own investments, it is definitely something that should be used to examine the performance of any professionally managed investment portfolio such as a unit trust, especially to compare it with competitors.

This is the last of the series of advanced notes in the Smart Investor education series. We hope you’ve enjoyed them!