GOOD STOCKS CHEAP

Value Investing with Confidence for a Lifetime of Stock Market Outperformance by Kenneth Jeffrey Marshall

Book Summary

Chapter-by-chapter summary of the key takeaways derived from the book.

The book is available for purchase from Amazon **HERE**

Cowley, Alistair
alistair.cowley@growthwithvalue.com
www.growthwithvalue.com

Good Stocks Cheap – Value Investing with Confidence for a Lifetime of Stock Market Outperformance

Table of Contents

PART I – FOUNDATIONS

Chapter 1 – The Quiet Outperformer

- Value investing performs better over time than other strategies.
- Value investing's efficacy is underreported, and will remain so.

Chapter 2 – Why Stocks?

Since I was born in 1967, the S&P 500 Index has returned just over 10 percent per year including dividends.

It's easy to see why listed stocks outperform bonds. Bonds represent a promise to pay someone a fixed amount plus interest. The money that can come out of a bond investment is thereby limited. It's limited to the principal, or the amount originally loaned, plus interest. And the loaned amount – the principal – doesn't go up at all. That's why there's a ceiling on how much a bond can return. Its why bonds are called a fixed income security.

Stocks have no ceiling. Business don't have a set limit on how well they can do grow. There's no guarantee that they will continue to do well and grow, but they can. That's why stocks perform better over time. They're nothing less than claims on ever-growing earnings. They offer ceiling-free upside.

My family moved to Southern California in 1979. That summer, a Balboa Bar cost $1.25. Thirty-seven years later in 2016, I bought one for $3.75. That comes out to an annual average price increase just above 3 percent. Even if I had deposited my 1979 $1.25 in an interest-bearing savings account, the 2016 balance would have bought less than one Balboa Bar. This is due to taxes. Earned interest is taxable income. If I paid an average of 50 percent of my income in taxes every year – which summing California and federal rates is roughly right – I would have had to earn twice the rate of inflation just to stay even. But my bank's average interest rate was much less than that.

Commodities also fall short as a long-term investment. Over decades, the prices of crops and metals do nothing more than track inflation. And when one takes into account carrying costs – the fees involved in moving and warehousing copper, for example – returns can actually lag inflation.

Unlisted stocks – ownership interests in privately held companies – also perform modestly average. The history of unlisted stock performance suffers from selection bias. The successes are trumpeted, but the far more numerous failures are muted.

We often adopt the perspective that we're buying entire businesses. Our approach is therefore the same whether we're buying 100 shares of a company or 100 percent of a company. To analyse an investment, we start by analysing a business.

This approach is called fundamental analysis. It's looking at company's financial statements, strategic positioning, and other firm-specific factors. It has little use for stock price or charts or macroeconomics. Our primary data source in fundamental analysis is annual reports, in particular the consolidated financial statements.

We prefer to gather our data direct from financial reports rather than online data services. This is for two reasons. First, annual reports are – on the margin – more reliable. Second, online data sources summarize narrative portions of annual reports. This requires them to interpret. And this can lead to

distortions. A proxy statement contains information on executive compensation, fees paid to the members of the board of directors, and potential conflicts of interest.

Another data source is investor relations departments. I have found investor relations people to range from helpful to unresponsive. Lesser ones do less. But an imperfect reply is useful information in itself. How investor relations answer questions is part of the answer.

- Value investors focus on listed stocks because they return best over time.
- Stocks are business ownership interests.
- Intelligent stock investing begins with fundamental analysis.
- The most reliable, free data source for fundamental analysis is annual reports.

Chapter 3 – Price and Value Are Different

Value investing has at its core that price and value are different. Price is what something can be purchased or sold for at a given time. Price fluctuates. Value, by contrast, is what something is worth, it fluctuates less.

Value investors concede that over time, the average price of a stock mirrors its actual worth. But at a single point in time, it may not. And those single points are enough to shatter the efficient market hypothesis. The moments when price is far away from value is when the value investor acts.

I wind up with two intrinsic values that bookend a range. Working within ranges turns out to be sufficient. If a good company is priced well below the lower end of my range, it's inexpensive. I buy it. If its priced wildly above the upper end of my range, it's expensive. I may sell it. My inability to pin down an exact intrinsic value doesn't emerge as a limitation.

- Some companies are superior.
- Superior companies are worth buying at some price, but not at any price.
- Price and value are commonly conflated.
- Investors who distinguish between price and value enjoy a rare advantage.

Chapter 4 – Measuring Performance

Calculating one's historic performance correctly is useful. It provides a baseline, as well as a template for tracking future returns. Plus, it reveals whether or not there is any room for improvement.

I measure my own performance in three steps. First, I measure my percentage return for each year. Second, I calculate my yearly average rate. Third, I compare that average to a standard. There are two kinds of standards: relative and absolute. The relative standard is an index, like the S&P 500 Total Return. The other standard is absolute. It's a fixed percentage, like 10 percent.

When we correctly measure our investment performance, we identify it. Yes, calculating one's average return is exacting and bothersome. But it's no sideline. It's a prerequisite. It's a must if we're ever to manage it.

- Measure the percentage return for each individual year.
- Calculate the geometric average.
- Compare the average to a standard.

PART II – THE VALUE INVESTING MODEL

Chapter 5 – Understanding the Business

The value investing model begins with a fundamental question: do I understand it? More specifically, do I understand the business? To understand a business is to be able to describe it in a single, unambiguous sentence. I call it an understanding statement. It's written in simple language that anyone can understand.

There's a process used to formulate sentences like this. It involves defining the business along six parameters.

The first parameter is products. Are the company's products goods or services? It is also useful to distinguish between commodity and differentiated products. Think plainly. It would not be wrong to define Unilever's products as fast-moving consumer goods. But it's clearer to say soap and food.

The second parameter is customers. Customers buy a company's products. Are the customers consumers or organisations? The more specific the customer definition, the better. Customers are different from users. I use Yahoo. But I've never paid Yahoo a nickel. The customers are the ones that pay the company.

A third parameter is industry.

A fourth parameter is form. Form is the way a business is structured, both legally and operationally. When talking about form in the operational context, we could say the company is a franchisor. Or it could be a franchisee. A company could be a multilevel marketer. Or it could be vertically integrated, operating along many nodes of its industry value chain.

The fifth parameter is geography. Geography describes where the business's customers, operations, and headquarters are located, with the emphasis placed on whatever is most enlightening. Wal-Mart's most relevant geography concerns its stores. They're located mostly in North America.

The sixth parameter is status. This catchall category can be used for a company's prominence, age, transformation, or whatever else merits singling out. Other times the status parameter can be used to define a company's size.

Product understanding can be a trap when writing an understanding statement. It can masquerade as the whole picture. It's one of four traps to look for. A second trap is marketing messages. Also called slogans, they're designed to motivate customers. They shouldn't be plugged unthinkingly into an understanding statement. The third trap is mission statements. Mission statements are marketing messages pointed inward. They highlight companies' goals and strengths, and are written to inspire employees, partners, and investors. Communications like this are important. But they're not ready-made modules of understanding statements. The fourth trap is aspirations. Sometimes a company wants to have a characteristic that it presently does not. An understanding statement should focus on what the business is about and not what it could be.

Sometimes companies are composed of very different units. Each unit is a distinct company with its own products, customers, industry, form, geography, and status. When I encounter this, I do one of two things. If one unit constitutes the bulk of a company's activities, I define the six parameters for just that one. If the units are more equally weighted, I define all six separately for each.

Our circles of competence naturally evolve. Expertise we have in some fields atrophies, and in other fields it grows. If we truly can understand a business, we're right to look elsewhere. But if it just stretches us – if it invites us to learn something new – we should take advantage of that chance. But just because one has to read up on a new subject doesn't mean that an idea is beyond reach. And even if it does turn out to be beyond reach, that doesn't mean that something constructive didn't happen. Proficiency accumulates.

Define the business along six parameters:

1. Products
2. Customers
3. Industry
4. Form
5. Geography
6. Status

Don't be misled by:

1. Product familiarity
2. Marketing messages
3. Mission statements
4. Aspirations

Chapter 6 – Accounting is a Language

The second step in the model asks another fundamental question: Is it good? This question breaks into three sub questions. First, has the business been historically good? If so, is it likely to remain good in the future? And third, is it shareholder-friendly?

The Income Statement:

The income statement first reports revenue, then cost of goods sold, then it states operating expenses. Operating expenses are different from cost of goods sold as they are not triggered by revenue. Finally, the income statement measures income.

The Cash Flow Statement:

Cash inflow is cash coming into a business. Cash outflow is cash going out of a business. Cash inflow minus cash outflow equals net cash flow. The cash flow statement sorts cash flows into three categories. Each category has its own cash inflows, cash outflows, and net cash flows.

The first category is cash flow from operations. It's the cash flow that results from the business doing what it does.

The second type of cash flow on a cash flow statement is cash flow from investments. It includes the purchase and sale of equipment that will last for multiple years. Such purchases are called capital expenditures. They appear on a cash flow statement as a cash outflow from investments called purchase of property, plant, and equipment.

The third category of cash flow is cash flow from financing. A company borrowing money from a bank is an example of cash inflows from financing. A company paying the bank back, or buying back its own shares, are examples of cash outflows from financing.

The main difference between the cash flow statement and income statement is that the cash flow statement uses the cash basis of accounting, and the income statement uses the accrual basis of accounting. Both the cash flow statement and the income statement look at a business over a period of time.

The Balance Sheet:

The balance sheet looks at a business at a single point in time. The balance sheet shows what a business owns, what a business owes, and the difference between the two. What a business owns is called assets. What a business owes is called liabilities. And the difference between the two is called equity.

Assets are things that a business controls, finds valuable, and bought. There are two kinds of assets: current assets and noncurrent assets. Any asset that could be used within a year is a current asset. Noncurrent assets generally take more than a year to use.

Generally, a business is said to make money if it shows positive income on its income statement. If this income is retained – that is, not paid out as dividends – then equity increases by that amount. Income that is kept in the business is called retained earnings.

Chapter 7 – Capital Employed

One key number is capital employed. Capital employed is a company's required financial base. It's the resources that have to be there in order for the company to maintain its level of operations. I measure capital employed by starting with total assets and then subtracting excess cash and non-interest-bearing liabilities. Note that the type of financing doesn't matter. It could be debt, or it could be equity. If the business borrowed $1,000, capital employed is $1,000. If the business raised the $1,000 by selling new stock in itself, capital employed is still $1,000. Capital employed is blind to the type of financing.

Excess Cash

Spotting excess cash is difficult. How much cash does a business really need? Different business are different. How much cash a firm needs to have depends on its business model, how fast customers pay, and how fast vendors need to be paid. Not that there's a perfect solution. But absent any specific guidance, one approach is to calculate two versions of capital employed: one with all cash subtracted, and one with no cash subtracted.

Debt and Non-interest-bearing Liabilities

My formula for capital employed also calls for the subtraction of non-interest-bearing current liabilities. These include accounts payable, deferred income, and accrued expenses. Why the non-interest-bearing bit? To keep financial debt in the calculation of capital employed. When a company borrows money, it often does so for terms longer than a year. On the balance sheet, such loans are categorised as noncurrent liabilities, sometimes called long-term liabilities. As such, this debt would naturally wind up in the calculation of capital employed. But as these loans come due, they slip into the current liabilities section.

Goodwill

The astute investor makes a judgment in deciding whether or not to deduct goodwill. To me, the key issue is how recurring the increases in goodwill are. The more frequently a company makes acquisitions, the less deducting goodwill makes sense. I often leave it in.

Leases

Accounting sees two kinds of leases. One is called an operating lease. Lease expenses are recognised on the income statement when they're due. Cash outflows from operations are recognised on the cash flow statement when they're paid. But that's it. Nothing shows up on the balance sheet. For this reason, an operating lease is said to be off balance sheet.

The second kind of lease accounting sees is called a capital lease. The company reports as if it had purchased the equipment and had borrowed money to do so. It's on balance sheet. It sees the lease as a debt every bit as real as money owed to a lender. A capital lease drags leased equipment onto the balance sheet. It increases assets and liabilities by an equivalent amount. This forces leased equipment into any measurement of capital employed, since it's now part of total assets. If substantially all the risks and rewards of ownership go to the lessee, then it's a capital lease. A capital lease is often referred to a finance lease.

The challenge comes when accounting standards see an operating lease, but the investor smells a capital lease. The solution is to adjust capital employed upward by capitalising the operating lease.

Present value is important in the capitalisation of operating leases because it determines how much of a liability, and corresponding asset, is added to the balance sheet. When an operating lease is capitalised, all three financial statements change. Specifically, six things happen. Each of them makes sense when one remembers that the point is to pretend that equipment that was leased was instead bought with borrowed money.

First, the lease expense is added back on the income statement. Second, the lease payments are added back on the cash flow statement. Third, a liability is added to the balance sheet, as is an equivalent asset. Fourth, interest is expensed on the income statement. Fifth, interest payments are outflowed on the cash flow statement. Sixth, the lease asset is depreciated on the income statement.

The amount of the liability added to the balance sheet, and the corresponding asset, is the present value of the future expenses under the lease. Calculating it require two things: the future interest expenses and a discount rate.

Future lease expenses appear in a table following the financial statements. The table lays out the lease expenses scheduled for each of the next five years, plus a sixth amount equal to the sum of all lease expenses scheduled from year six through the year in which the last lease expires. One must make a rough guess to allocate this sixth amount among the last set of years.

The discount rate is the percentage the company pays to borrow money. The financial report often reveals the company's most recent cost of debt, expressed as an interest rate. This is the percentage to use. The higher the discount rate, the smaller the present value.

Whatever the asset, operating leases can mask the full scope of a company's liabilities.

Invested Capital

Capital employed is only one way to conceptualise a company's financial base. Another is invested capital. By some formulas, it's equivalent to capital employed. But the thinking behind it is different. Invested capital is the total amount of money committed to a company by parties that expect a financial return. Shareholders expect a financial return. So do bondholders. But vendors – like landlords that just want to get paid – don't. they're not expecting interest, dividends, appreciation,

or anything else besides their rent. They wouldn't think of themselves as investors in the company in the way that shareholders and noteholders would. The invested capital notion of a company's required financial base suggests different math. It suggests adding up all the capital that has been contributed by parties that expect a financial return. Equity plus financial debt, basically. But this additive method may lack a benefit of the subtractive method used to calculate capital employed. Starting with total assets and deducting only those things that clearly don't belong reduces the chance of underestimating a company's required financial base.

- Capital employed is one measure of a company's required financial base.
- Capital employed is total assets minus excess cash minis non-interest-bearing current liabilities, and possibly minus goodwill.
- Total assets can include capitalised operating leases.
- Capitalising operating leases causes all three financial statements to change.
- Considering a company's financial base is more important than measuring it perfectly.

Chapter 8 – Operating Income

A second key number is operating income. At first glance, net income seems like a perfectly good measure of the profitability of a business. It is, after all, the bottom line. So why focus on operating income? Because it makes comparisons easier. To see this, see what operating income ignores. For one, it ignores non-operating income. It also ignores interest expense. Interest expense captures how a company was financed, not just how it operates. Operating income also ignores taxes. Tax expense reflects where a business is located. In other words, operating income measures the profitability of a company without regard for the capital structure or tax regime in which it operates. It lets one compare the same enterprise at different points in time. It also lets one compare different businesses at the same point in time. It's the best measure of income for looking at companies that have different tax rates or debt situations but are otherwise similar.

This is not to deny that a company has to pay interest, or that it has to pay taxes. Nor is it to suggest that interest or tax expense don't have an impact on shareholder returns. They do. But its more instructive to begin by looking at what the heart of an enterprise produces.

When we start calculating metrics, we'll consider operating income and capital employed together. Adjustments shaping one must therefore shape the other. If operating leases were capitalised in the calculation of capital employed, some changes need to be made to the published operating income figure.

First, the lease expense needs to be added back. Second, a new depreciation expense needs to be subtracted. Recall that the capitalisation of operating leases compels one to put new interest and depreciation expenses on the income statement. They're swapped in for the lease expense that's taken out. In fact, a common assumption is that the sum of new interest and depreciation expenses equals the old lease expense. That is, net income doesn't change. Interest income isn't directly relevant to operating income. But interest is the key when estimating depreciation, which is captured in operating income. So, it has to be quantified. When adjustments were made to capital employed, an interest rate was identified. That same rate can be used to gauge interest expense. Find the amount that was added to both sides of the balance sheet as the capitalised operating lease present value. Effectively, that's the size of the loan used to buy the equipment. Multiply that number by the interest rate. The result is the interest expense. Subtract this number from the old lease expense for the same year. The result is depreciation expense.

Operating income works for me because its compatible with the broad range of tax and debt options available to the companies that I see. The other thing that accounts for my focus on operating income is my long investment horizon. I hold stocks for decades. As far as I'm concerned, companies I own have years to get their taxes and capital structure in order.

- Operating income is the best measure of the core earning capacity of a business.
- Adjustments that shaped capital employed must also shape operating income.
- If operating leases were capitalised, the calculation of operating income requires adding back an old lease expense and subtracting a new depreciation expense.

Chapter 9 – Free Cash Flow

A third key number is free cash flow. It measure the amount of cash a company threw off just by operating. I calculate free cash flow by subtracting operating expenses from cash flow from operations. Recall that a capital expenditure – also called capex – is the purchase of a non-current asset. There are two types: maintenance capex and growth capex. Maintenance capex is the purchase of equipment that will take over from equipment that is wearing out. It's replacement. It's an outlay a company makes just to maintain its level of operations. Growth capex is different. It's an outlay a company makes to expand operations. It results in a larger enterprise.

To calculate free cash flow well, only maintenance capex should be deducted from operating cash flow. After all, it's the cash-generating power of the present operations that free cash flow aims to assess. It would be wrong to underestimate the performance of a business by subtracting all capex. The problem is that the cash flow statement doesn't divide capex into maintenance and growth. A popular workaround is to accept depreciation – the accounting measure of how much accounting equipment wore out – as a proxy for maintenance capex. This has merit since deployed equipment depreciates. But I dislike this approach, for two reasons. One is inflation. depreciation is a portion of old cost. It's a fraction of what was paid for a piece of equipment. That same piece of equipment might cost more now. Two is quality. Pieces of equipment of equipment wearing out might be of different quality than their replacement.

A workaround better than relying on depreciation is to search for some breakout of maintenance capex in an annual report. Another is to listen to one during webcast earnings calls. When both of these approaches fail me, I pretend that all capex is maintenance capex. Knowing that I'm painting an overly dark picture of the enterprise.

Any adjustments that shaped capital employed and operating income should also shape free cash flow. So, if one capitalised operating leases, one must make two adjustments. First, lease payments should be added back to cash flow from operations. Assume that lease payments equal the lease expense that was added back to published operating income. Second, interest payments on the pretend loan should be subtracted from cash flow from operations. The amount to subtract equals the interest expense that was estimated in the calculation of operating income.

Some companies that report under IFRS put interest payments not in cash flow from operations, but rather in cash flow from financing. If one capitalised operating leases of such a company, one should not only subtract the new interest payments, but also move existing interest payments from the finance section to the operations section. This formula for free cash flow – cash flow from operations minus capex, adjusted for the capitalisation of operating leases – yields a specific measure of free cash flow. It's called levered free cash flow. Levered means indebted. It indicates the inclusion of interest payments.

Unlevered free cash flow is different. It excludes interest payments. To calculate it, first add interest payments back to levered free cash flow. Include both interest payments created when operating leases were capitalised, and any interest payments that existed beforehand. Second, increase tax payments. Interest is a tax-deductable expense. So erasing interest expense boost earnings before taxes, compelling the company to pay more income tax.

It's simple to calculate how much tax payments should go up. Start by finding the company's income tax rate. Then, subtract that percentage from 100 percent. If the tax rate was 30 percent, then the number is 70 percent. Finally, multiply that number by the interest payments that were added back. That's how much tax payments should increase.

I prefer levered free cash flow because it gives me such a different look at a company than does operating income. Operating income is an accrual accounting figure that ignores taxes and interest, while levered free cash flow is a cash accounting figure that captures taxes and interest. Together, these two perspectives help me to see everything important related to how much money came out of the business.

Revenue is recognised on the income statement before the associated inflows from operations happen on the cash flow statement. This is particularly true with growing businesses. More and more cash is outlaid as companies build. With shrinking businesses, the reverse can be true. Operating income can be lower. Cash from previously large order bases pours in as purchases slow. But most companies worthy of an investor's attention do more growing than shrinking.

Occasionally a company's business model creates a negative cash cycle. This happy situation has a company getting paid by customers before it pays vendors. An example is a magazine publisher that requires payments for subscriptions. Unlevered free cash flow can exceed operating income even though the business is growing.

But besides those that are paid in advance and those that are in decline, businesses normally see revenue exceed cash inflows from operations. This will become important when we start measuring past performance.

- Free cash flow equals cash flow from operations minus maintenance capex.
- If maintenance capex isn't distinguishable from growth capex, subtracting all capex is a conservative – though imperfect – approach.
- If operating leases were capitalised, the calculation of free cash flow requires adding back an old lease payment and subtracting a new interest payment.
- Levered free cash flow captures interest payments, while unlevered free cash flow does not.
- Most businesses have a normal cash cycle.

Chapter 10 – Book Values and Shares

A fourth key number is book value. When applied to a whole company, book value means equity.

Noncontrolling interests – redeemable or not – are normal. They're the difference between shareholders' equity and the larger total shareholders' equity. Two things make equity go up. One is retained earnings. Another is issuance of new stock for cash. Correspondingly, two things make equity go down. One is the payment of dividends. Another is share repurchases. This means that any metric based on book value is influenced by repurchases. This isn't automatically good or bad. It's just something to be aware of.

A fifth key number is tangible book value. Its book value minus intangible assets. Intangible assets are any assets that aren't physical. They include goodwill, trademarks, and patents. I find tangible book value useful for companies that overpay for acquisitions. If book value is much higher than tangible book value, its often because goodwill is enormous. This directs my attention to a company's acquisition history. If it recently bought another public company, I analyse the acquired company right before the acquisition. Would I have paid that price? If not, I stop my analysis. I don't care what the purported synergies were. Overpriced acquisitions are the beginning of a story that ends poorly. I'm not in.

A sixth final key number is shares. Many useful metrics are expressed on a per-share basis. They capture the dilutive effects of an increase in the number of shares, or the concentrating effect of a decrease in the number of shares.

Shares outstanding, sometimes called basic shares, are the number of shares that people, investment funds, and other entities hold. Fully diluted shares is a bigger number. Its shares outstanding plus the number of shares that could become outstanding if other securities issued by the company were exercised or converted. Shares authorised is bigger still. It's the number of shares the company is allowed to issue according to its charter or bylaws.

Which to choose? Many pick shares outstanding. But when considering the purchase of a stock, I prefer fully diluted shares. This is because in a successful company, anything that can become a share will become a share. The fully diluted shares measure isn't perfect. For instance, it doesn't take into account the cash the company could receive in the event of conversions. Exercising stock options, for example, often entails the payment of a strike price by the holder. This would lead to a cash inflow from financing for the company, and an uptick in the company's cash balance. But unless one knows the strike price, capturing this effect is difficult.

Here's another one. During the course of a year, the number of fully diluted shares fluctuates. Secondary offerings happen, options get issued, and the number of shares generally bops around. To capture this variability, annual reports often report the weighted average number of fully diluted shares.

- Book value equals equity excluding noncontrolling interests.
- Tangible book value equals book value less all nonphysical assets.
- Fully diluted shares is the most meaningful of the three share counts.

Chapter 11 – Past Performance

Extracting key numbers from financial statements lets one calculate key performance metrics. There are seven. They make clear whether or not a business has been historically good.

Return on Capital Employed (ROCE)

The first performance metric is return on capital employed, or ROCE. ROCE is a ratio. It's expressed as a percentage. The numerator is operating income. The denominator is capital employed. The purpose of ROCE is to show how much money a business made relative to the amount of capital it needed. The numerator, operating income, describes a period of time. This is because it comes from the income statement. The denominator, capital employed, describes a point in time. It's calculated from the balance sheet.

Capital employed can change a lot over 12 months. There are three choices. One is the beginning of the year. This means calculating capital employed from the prior year's ending balance sheet. Another is the end of the year, which means using the current years ending balance sheet. A third is to use an average, by taking the arithmetic mean of the beginning and the end.

One method is to calculate two versions of capital employed: one at the beginning of the year, and one at the end of the year. This yields a range of possible values. So, for a single period I wind up with four measures of ROCE. Each is defined by my treatment of capital employed: beginning of the period with cash, beginning of the period without cash, end of the period with cash, and end of the period without cash.

Free Cash Flow Return on Capital Employed (FCFROCE)

A second performance metric is free cash flow return on capital employed, FCFROCE. It's levered free cash flow divided by capital employed. Like ROCE, FCFROCE requires picking a date on which to measure capital employed. Using a range again makes sense.

In a growing business, FCFROCE will often be lower than ROCE. This is because levered free cash flow tends to lag operating income. It captures interest, tax, and the normal cash cycle. Operating income doesn't.

Change in Operating Income per Fully Diluted Share (ΔOI/FDS)

A third performance metric is growth in operating income per fully diluted share. Calculating ΔOI/FDS is easy. Divide operating income by fully diluted shares for the first year. Then do the same for the second. Next, subtract the first from the second, and divide the result by the first. A percentage results.

Change in Free Cash Flow per Fully Diluted Share (ΔFCF/FDS)

The fourth performance metric is growth in free cash flow per fully diluted share, ΔFCF/FDS. Its computed by dividing levered free cash flow by the number of fully diluted shares for the first year, then doing the same for the second. Then, subtract the first from the second, and divide the result by the first.

Change in Book Value per Fully Diluted Share (ΔBV/FDS)

The fifth performance metric is growth in book value per fully diluted share, ΔBV/FDS. Divide book value by the fully diluted shares for the first year, then doing the same for the second. Then, subtract the first from the second, and divide the result by the first. This metric shows the increase in worth over time from a strict accounting standpoint. But it has limited utility by itself. This is for two reasons. One, it fails to capture dividends. Dividends come directly out of book value. So meaningfully using ΔBV/FDS requires knowing about any dividends paid.

While book value book value return and dividend growth are both related to an investor's return, they're different. It would be wrong to regard the dollar of one as equivalent to a dollar of the other. This is in part because dividends are usually taxed, while growth in book value usually isn't. So, it would be overly simplistic to lump the two together in a single operating metric.

Another reason that ΔBV/FDS has limited utility by itself is because it's influenced by share repurchases. Buybacks can make ΔBV/FDS negative even if a company has positive ΔOI/FDS and pays dividends.

Change in Tangible Book Value per Fully Diluted Share (ΔTBV/FDS)

The sixth performance metric is growth in tangible book value per fully diluted share, ΔTBV/FDS. It's calculated by dividing tangible book value by the number of fully diluted shares for the first year, then doing the same for the second year. Then subtract the first from the second and divide the result by the first. This metric has merit, but it too doesn't capture dividends and is influenced by buybacks. So again, it's of limited use by itself.

Liabilities-to-equity Ratio

The seventh performance metric is liabilities-to-equity ratio. It sizes the prominence of obligations on a company's balance sheet. It's computed by dividing total liabilities by book value. I maniacally expand the numerator to include all liabilities: bonds, bank loans, account payable, accrued expenses, unearned revenue – anything owed, interest-bearing or not. Why? Because companies can crumble under the weight of too many obligations of any time.

While I prefer the liabilities-to-equity ratio, the alternatives can work well too. Liabilities amplify results. If things go well for a highly leveraged company, they go really well. The economic benefits of success gush in the direction of the shareholders because the debtors and vendors have limited upside. But if things go poorly for a highly leveraged company, they go quite poorly. Operating income falters and cash flow slows to a trickle, but the liabilities burden lets up not one iota. Lenders demand their interest and principal. Vendors want their invoices paid. Insolvency can result. It's this technical possibility of disaster that tilts value investors towards business with manageable obligations.

A high liabilities-to-equity ratio may be palatable in two cases. One is when a very credit worthy company has a high-return use for cash borrowed in a time of low interest rates. Another is when equity has been reduced through dividends or – under certain circumstances, defined later – buybacks. Occasionally these two conditions combine. A strong company borrows money at a low interest rate to buy back its stock at prices well below value. Debt soars and book value condenses, creating a grotesquely high liabilities-to-equity ratio. But the situation may nonetheless be attractive to the astute investor.

Normalised Results

It would be misleading to judge a company based on one bad year. At the same time, it would be foolish to pretend that it didn't happen. What one really wants is to make judgments on normalised results. One way to normalise results is to look at several years together. Consider ROCE. If a business hasn't fundamentally changed over the last five years, one could calculate average ROCE by adding up the operating income of all five years and dividing that number by the sum of each year's capital employed. Average FCFROCE can be calculated the same way.

I like normalising by taking multiyear averages. But some prefer doing so by tweaking individual years. They remove atypical events. By this method a one-time litigation expense, for example, would get added back. I dislike this approach for three reasons. First, only future periods can determine what's truly irregular. Was that litigation expense really a one-off? If it happens again, no. Second, to the extent that there are real irregularities, management-level knowledge is required to identify them. Annual reports, press releases, and news items are detailed enough to fool outside investors into thinking that they can spot them too. But they can't. Third, every year has irregularities of some sort. Once one starts tweaking, there's no end to the justifiable tweakings. For these reasons I prefer to regard historic operating performance as just that: historic. I don't try to clean it up and pretend that its sanitised version will repeat.

With the liabilities-to-equity ratio, normalisation doesn't make sense. What matters is the most recent version of that figure.

Performance Metric Benchmarks

When we calculate the seven performance metrics, we wind up with actual numbers. Consider ROCE. It may be 9 percent, or 19 percent. What's good, and what isn't? The answer can be found in the stock market. Since the year of my birth – 1967 – the stock market has returned 10.2 percent per year on average.

Picture buying all the shares of the company and delisting the company from the stock exchange. No more stock quotes. As the years rolled on, what kind of return could one expect? One would get all of the company's net income. Over time, that figure would accurately reflect – via depreciation – the cost of replacing equipment wearing out, or maintenance capex. As the years rolled on, this relationship between net income and capital employed would come to govern one's return. If that figure isn't at least 10.2 percent annually, the company isn't worth buying.

Of course, ROCE doesn't use net income as its numerator, but rather the larger operating income. This excludes tax and interest expense. So, 10.2 percent is a too low a benchmark for ROCE. Instead I use 15 percent. That's my rough estimate of the average annual S&P 500 total return if none of its component companies had to pay interest or taxes.

What about FCFROCE? It can be benchmarked against a stock market return for the same reasons. But 10.2 percent is too high. Because of the normal cash cycle. I use 8 percent.

Growth metrics also need benchmarks. As an absolute baseline, consider inflation. as noted, U.S. inflation has average around 3 percent in recent decades. So, if we're looking at a growth rate of under three precent, we're probably looking at a business that's shrinking. That's because financial statement figures are nominal. They're not adjusted for inflation. So, in real terms, the business isn't growing at all.

The liability-to-equity ratio is much harder. As a general rule I like it to be no higher than 2. But if all of the other performance metrics are great and the company is using cheap borrowed money to buy back its under-priced stock, I can be comfortable with a ratio as high as 7.

There are really only about a dozen figures that need to be pulled from the financial statements for each year. From the income statement, there's operating income, and possibly interest expense. From the cash flow statement, there's cash flow from operations, and capital expenditures. From the balance sheet, there's cash, goodwill, intangible assets, total assets, accounts payable, deferred income, accrued expenses, total liabilities, and shareholder's equity. The only additional figure needed are the number of fully diluted shares and maybe the effective tax rate.

Sometimes a performance metric is negative. Average FCFROCE or ΔOI/FDS may be minus. Other times, a performance metric may be simply underwhelming. ROCE may be 7 percent, or ΔFCF/FDS may be 1 percent with no other compensating growth metric. Avoid making excuses for unremarkable companies. The value investing model is designed to flag them. If a business looks bad, consider it no further. There are others more worthy of attention.

A business's historic performance can be measure with seven metrics:

1. Return on capital employed (ROCE).
2. Free cash flow return on capital employed (FCFROCE).

3. Growth in operating income per fully diluted share (ΔOI/FDS).
4. Growth in free cash flow per fully diluted share (ΔFCF/FDS).
5. Growth in book value per fully diluted share (ΔBV/FDS).
6. Growth in tangible book value per fully diluted share (ΔTBV/FDS).
7. Liabilities-to-equity ratio.

Chapter 12 – Future Performance

Strategic analysis is required to analyse future performance. Specifically, four qualitative tools help to assess a firm's prospects. Two resources become useful when undertaking strategic analysis. First is investor relations. Questions with factual answers are the nest sort to ask of investor relations departments. The other resource was people in the target market. Explain to them that you are researching for an investment and that you know nothing about the business or industry and would be grateful for any insights.

Breadth analysis

First is breadth analysis. It asks two questions. One, is the company's consumer base broad, and unlikely to consolidate? And two, is the company's supplier base broad, and unlikely to consolidate? The business isn't good unless the answer to both these questions is yes.

I define a broad customer base as one where no single customer accounts for over a tenth of revenue. Similarly, I define a broad supplier base as one where no single supplier accounts for over a tenth of cost of goods sold or operating expenses. There's one exception, however. A good company may concentrate purchases on a limited number of suppliers. This can yield bulk discounts and service concessions not available otherwise. But it can also cause one supplier to emerge as disproportionately important.

Analysing customer breadth is straightforward. Annual reports often contain a clear statement like "no single customer accounted for over 10 percent of our revenue.: the absence of such a statement suggest the absence of this happy condition.

Analysing supplier breadth is harder. It's often best just to think about what a business needs to buy. Must inventory be a certain brand, or will any do? Do key components come from a country where one family controls commerce, or are they available globally? Are qualified workers scarce and unionised, or plentiful and easily trained?

Forces analysis

A second qualitative tool is forces analysis.

Bargaining Power of Customers

The first force is the bargaining power of customers. If a company's clients demand lower prices or faster delivery times, would they get them? Three things can lower customer bargaining power:

1. One is the number of customers. The more there are, the lower their power.
2. Another is the improbability of backward integration. When a customer backward integrates, it starts doing something previously done by the company. If it can't – because it would need a competence far away from its core activity, for example – its bargaining power is weaker.
3. A third is high switching costs. If a customer would incur large expenses by switching to one of the company's competitors – because its operations are tightly integrated into those of the

company, for example – customer bargaining power is weaker. Example, boaters are unlikely to backward integrate by making their own life preservers and outboard motors, of course. But they have low switching costs, and there are plenty of other places to shop for accessories. So, the bargaining power of the company's customers is strong.

Bargaining Power of Suppliers

The second force is the bargaining power of suppliers. If the suppliers hiked prices or demanded faster payment, would the company give in? Three things that lower a supplier bargaining power:

1. One is the number of suppliers. More means weaker. If breadth analysis revealed a sustainably broad supplier base, supplier bargaining power is probably weak.
2. Another is the improbability of forward integration. When a vendor forward integrates, it starts doing something previously done by the company. The less conceivable this scenario, the weaker suppliers' bargaining power.
3. A third is low switching costs. The less expensive it is for the company to change vendors, the weaker supplier bargaining power is.

Threat of Substitutes

The third force is the threat of substitutes. Substitutes take three forms. One is direct substitutes, or similar products offered by existing competitors. Two is doing without, a concept from economics. Three is wholly different products that perform the same basic function. Two things can keep the threat of substitutes low. One is if the substitutes offer lower value the do the company's offerings. Another is if the company's customers would incur high switching costs by adopting substitutes.

Threat of New Entrants

The fourth force is threat of new entrants. New entrants come in the form of both start-ups and new divisions of existing companies. A company may have attributes that may keep potential new entrants on the sidelines. One is economies of scale. The company's production volume may result in lower per-unit costs than those of a newcomer. Another is high customer switching costs, which make a new player's offerings expensive to try. Last is hard-to-get permits if the company is in a regulated industry. Just one of these three attributes can dull the threat of new entrants.

To effectively use forces analysis, one must make decisions. One must call each force as either strong or weak. Medium isn't useful. If all four forces acting on a company are weak, further consideration of the investment is warranted. If two or more a strong, the opposite is true. One should stop. When in doubt, quit the analysis and find another company to look at.

Moat Identification

A third qualitative tool is what I call moat identification. A moat is a barrier that protects a business from competition. There are six potential sources.

Government

The first is government. Governments grant special rights to some companies.

Network

A second source of moat is network. Network is an accumulation of users or customers. It constitutes a moat if it yields a product benefit derived from the other users of a product.

Cost

A third source of moat is cost. Sometimes a company has a low-cost structure that enables it to produce a product for less than competitors. A low-cost structure may come from a proprietary manufacturing technology, an inexpensive operating method, or some quirk of history.

Brand

A fourth source of moat is brand. Some brands are so strong that customers rarely consider substitutes. The Maggi name is so accepted in India that many people there don't say "instant noodle." The say "Maggi." That's strong evidence of a brand formidable enough to keep many would-be competitors away.

So how do far-away investors know if there's a brand-based moat? By talking to people in the target market. They avoid saying the brand name and see what gets mentioned. They unaided recall questions like "What's your favourite instant noodle?"

Switching Costs

A fifth source of moat is switching costs. High switching costs have a powerful retention effect on customers.

Ingrainedness

A sixth source of moat is what I call ingrainedness. It's like brand, but it applies to the channel instead of the end user. A company enjoys ingrainedness when it's so integral to a value chain that it's difficult to picture an industry without it. Like brand, ingrainedness can be hard for someone outside of the channel to see. So again, astute investors talk to people in the channel. They ask open-ended questions to see if the name of a favoured firm comes up.

Market Growth Assessment

A fourth qualitative tool is market growth assessment. This is a straightforward view of whether the company's market is growing or not. It's important because a business with an expanding market – everything else being equal – has a brighter future than one that doesn't.

Quantitative Forecasting

As the four tools make clear, qualitative forecasting is my preferred way to assess a company's prospects. But many prefer quantitative forecasting. This involves using spreadsheets to calculate expected revenues, margins, interest rates and the like. It's the dummying up of future period financial statements. I don't like this for two reasons.

First, it's hard. Making predictions about prices, costs and market share is just plain difficult. It often devolves in nothing more than guesswork cloaked in mathematics.

Second, it can be self-deceiving. Those dummied up statements look awfully official. They create the risk that one will mistake precision for accuracy. The truth is that one can't know how tomorrow's financials will look. But if one owns parts of businesses with both good operating histories and solid strategic positionings, it's likely that the future will play out well.

Assess a company's prospects with four qualitative tools:

1. Breadth Analysis
2. Forces Analysis
3. Moat Identification
4. Market Growth Assessment

Forces analysis involves assessing each of the following four forces as either strong or weak:

1. Bargaining power of customers
2. Bargaining power of suppliers
3. Threat of substitutes
4. Threat of new entrants

Moats, which are rare, come from one of six sources:

1. Government
2. Cost
3. Brand
4. Network
5. Switching costs
6. Ingrainedness

Chapter 13 – Shareholder-Friendliness

The final checkpoint is to see if the company is shareholder-friendly. Gauging shareholder-friendliness is a qualitative tool based on four indicators.

Compensation and Ownership

The first indicator is compensation and ownership. In the ideal situation, executives earn a reasonable salary. Board members are paid normal fees for their part-time jobs. And they all use their own cash to purchase shares in the company at market prices such that their holdings constitute the bulk of their net worth's.

Stock option plans are commonplace. These schemes allow for shares to be purchased at the low price of an earlier year, and then sold right away. This lets officers and directors participate in the upside of equity ownership without exposure to the corresponding downside. Such automatic gain arrangements hardly represent an equal yoking with investors. Any suggestion that they align incentives is laughable.

A useful shortcut is to search for the summary compensation table. This concise presentation lays plain the total paid to each executive for each of the last three years. The far-right column lumps together each officer's salary, bonus, and stock-based consideration.

I focus on the total amount that went to the highest paid employees. I don't care how much of it was salary, how much of it was bonus, and how much of it was stock. I gauge this number with floating guidelines. My current standard is if it's over $30,000,000, or over 5 percent of the lower of free cash flow or net income, I'm not interested. What revenue was doesn't matter.

Also disclosed are fees payed to members of the board. Search the proxy statement for director compensation. I look for the amount that went to the highest paid outside director. Outside means not also a full-time company employee. This is often the non-executive chair or the lead independent director. Currently I like to see the highest paid outside director get total annual consideration of around $250,000. I stop my analysis if its more than twice that.

These standards – $30,000,000 for an executive, $500,000 for a director – are very American. Both numbers would be absurdly high in most other countries. So, fees must be considered in context.

Officer compensation can just be a sliver of a company's overall expenses. Whether a chief executive makes $5,000,000 or $10,000,000 has little direct mathematical impact on a company with $1,000,000,000 in income. At the earnings per share level, the difference is undetectable. So why do we care? We care because compensation offers a rare window into the decision-making process at the highest levels of a company. Executives and directors that approve outsized pay are prone to other aggrandisements. They're more likely to consent to overprices acquisitions, for example.

Paying too much for management may be bearable by itself but paying too much for acquisitions can kill companies. That's why I favour teams that are properly paid. This preference causes me to miss some opportunities. But it keeps me clear of more catastrophise.

Significant, outright holdings by executives and directors of stock in the company they steward is a positive sign. Consider only owners that also serve as directors and executives of the company. Investment firms – which are on the chart of they own over a certain percentage of the stock – should be ignored. They're not the ones running the company. Outside the united states, search for the annual report for the term shareholder or ownership.

Material insider buying at market prices is also a promising indicator. The exercise of an option doesn't point toward shareholder-friendliness by itself. It's the open market purchases that are the positives.

Inside stock sales are never positive, but may not be negative either. Directors may need cash to pay for a child's college tuition. Executives may want to pay off mortgages. Founders that own most of their companies may sell a bit each month as part of personal financial plans. These are all legitimate reasons to sell. Of course, if several insiders dump most of their shares at once, that's bad. But routine disposals are so conclusive.

Related-party Transactions

A second indicator of shareholder-friendliness is related-party transactions. A party is related if it has dealings with the company that could pose a conflict of interest. A vendor to the company that's owned by a company board member would be an example of such a party. They're fine if negotiated at arm's length.

As a general rule, the number and size of such dealings, as well as the opacity in disclosure, are inversely correlated with shareholder-friendliness. If there's just one or two small related-party transactions and the detail offered on them is extensive, they're less likely to hurt the outside investor.

Search the proxy statement for the term related or relationships. What comes p will commonly involve leasing, consulting, or hiring.

One blunt test for fairness is to see how much of a company's total lease expense when to the insider entity. H&M's annual report says that rent expenses in 2015 totalled 20,554,000,000 Swedish kronor. Some quick maths reveals that the chairman's entity got only 2 percent of that. So, supporting an insider's real estate holdings clearly isn't a primary function of H&M. that finding, plus the address disclosure makes this setup seem less troubling.

Another type of related-party transaction involves consulting. Sometimes a director will be paid to provide special advisory services to the company. If few details are given, that by itself is a negative. But sometimes the best thing to do is to file a fact away, and judge it later in the context of other, less ambiguous shareholder-friendliness indicators.

A third type of potentially questionable dealings involves hiring. Sometimes a company job goes to a member of an insider's family. Officerships may pass from one generation to the next. I personally view management dynasties as questionable, but not disqualifying. If they're unhealthy, that fact tends to show up elsewhere. It may drive ROCE down, or executive compensation up. Bad managers cause bad outcomes, whatever families they come from.

Virtually all related-party disclosures come with some assurance of fairness. "This was purchased in the ordinary course of business at market rates," for example. Such phrases can be unconvincing. After all, they're thin on details. We don't know the size of the stores, the number of hours worked, the IQ of the offspring. But that's okay. We don't need to know everything. We just want to avoid blatant plunderage. We want a basic sense that there aren't undeserved bonuses masquerading as something else.

We needn't consider transactions between a company and its subsidiaries. They're normal, and don't threaten to pose a conflict of interest in the way that transactions with insider-owned entities do.

Share Repurchases

A third indicator of shareholder-friendliness is share repurchases. Many investors view companies buying their own stock back in the open market as positive. They may be. But only if the shares are bought at a discount to value. Purchases at other-than-discount prices aren't useful to the outside investor. They're as destructive as overprices acquisitions. Alas, management may be motivated to repurchase even without a discount. This is because of incentive compensation plans. Some bonuses kick in when earnings per share rise. By using free cash flow to purchase stock on the open market, executives can juice earnings per share even if total income doesn't increase.

Stock repurchases show up on the balance sheet. Cash decreases on the asset side, and equity decreases by an equal amount. Repurchases also register on the cash flow statement. They're a cash outflow from financing. But they're absent from the income statement. No expense is recognised. In other words, accounting has no way to call out overpayments at the time that they're made. This is why buybacks also have the short-term effects of turbocharging ROCE. The denominator drops, but the numerator doesn't. buybacks are disclosed in annual reports. Search for the term repurchase.

A commonly cited justification for buybacks is that they counterbalance the issuance of options under a stock options plan. By repurchasing shares while new options are granted, the number of fully diluted shares can be made to stay roughly the same. I hereby nominate this as one of the worst rationales in the history of finance. Using company money to transfer shares from outsiders to insiders has a name. Theft. The only good reason to buy back shares is because they're inexpensive relative to their worth.

Dividends

The fourth indicator is dividends. A problem with dividends is that they may signal that a company has no growth opportunities. I consider dividends most shareholder-friendly when they're paid by fully priced companies in mature industries with long histories of strong free cash flow. Dividends represent a use of cash for which there are no alternatives. Shares could be repurchased, growth capex could be increased, and other businesses could be bought. A company's dividends practices must therefore be weighed in the context of varying customs, expectations, and regulations around the world.

One dicey indicator is enough to compel the astute investor to drop an idea. Shareholder-friendliness is not a four-indicator average. It's defined by the worst. After all, gauging shareholder-friendliness is essentially an attempt to answer a basic question: *Do I want to be in business with these people?*

Gauging shareholder-friendliness is a qualitative process based on four indicators:

1. Compensation and ownership
2. Related-party transactions
3. Share repurchases
4. Dividends

Chapter 14 – Inexpensiveness

The third step in the value investing model is to ask a final fundamental question: Is it inexpensive? Even if we understand a business and find it to be good, we can still make a bad investment by paying too much. We deploy four price metrics.

Times Free Cash Flow (MCAP/FCF)

The first is times free cash flow. It equals a company's market capitalisation divided by its levered free cash flow, MCAP/FCF. The numerator, market capitalisation, is often shortened to market cap. It's the number of shares outstanding times the current price per share.

Notice the consistency between numerator and denominator. Market cap is the price for the equity only. Levered free cash flow is the cash thrown off by the business after debtholders have been paid their interest. Hence levered free cash flow goes to the equity holders.

Theoretically, market cap is what it would cost to buy all of a company's outstanding shares. But it's actually an underestimate. That's because the current share price reflects only what some shareholders were – moments ago – willing to take for their stock. Most are holding out for more.

Enterprise Value to Operating Income (EV/OI)

A second price metric is enterprise value to operating income, EV/OI. The denominator, operating income, was the subject of chapter 8. It's revenue, minus cost of goods sold, minus operating expenses. Note that it's not net of interest and tax expenses. The numerator, enterprise value, is the theoretical takeover price. It's what one would fork over to buy the entire company, not just the outstanding shares. It equals market cap plus the market price of all the company's preferred equity, non-controlling interest, and debt and minus cash.

Noncontrolling interest is portions of subsidiaries not owned by the company. Sometimes it's called minority interest because it reflects ownership by a minority partner. The number to add is on the balance sheet of the most recent report. It's in the equity section.

Three methods are used when recording a subsidiary on the financial statements:

1. If the company owns less than 20 percent of the investee, the cost method is used. The investment is carried on the balance sheet as an asset at cost, and that's it.
2. If the company owns between 20 percent and 50 percent of the investee, the equity method is used. The investment is initially carried on the balance sheet at cost. Then the investee has net income, that net income is multiplied by the percent of the investee that the company owns. This proportionate share of investee net income is then added to the company's income

statement. It's called something like earnings in affiliate. It flows through to the balance sheet as an addition to the carrying value of the investment.

3. If the company owns over half over half of the investee, the consolidation method is used. This makes the investee a consolidated subsidiary. The investees revenues and expenses are consolidated – mashed in – with the company's revenue and expenses on the income statement. Further down there's a line called something like earnings attributable to noncontrolling interest. That's the amount of the investee's earnings that belong to its other owners. It's the proportionate share of earnings that aren't the company's. it's subtracted. The consolidated method also calls for all of the investee's assets and liabilities to be included on the company's balance sheet. Noncontrolling interest – which appears either in the liabilities or equity section, as noted earlier – corrects for the proportion of the investee that the company doesn't own. On the income statement, earnings attributable to noncontrolling interest is below operating earnings. Operating earnings therefore reflect 100 percent ownership of the investee, which is fiction.

EV/OI is a particularly sage metric. Like capital employed, it's blind to type of financing. This contrasts with market cap-based valuation metrics that can be fooled by capital structures. Consider market cap to operating income, MCAP/OI. A company with a lot of debt could look inexpensive just because the numerator prices only the common stock. It doesn't capture the direct cost of buying out any bondholders, noteholders, or bank creditors. But enterprise value is never blindsided by capital structure. It prices everything.

Price to Book Value (MCAP/BV)

A third price metric is price to book, MCAP/BV. It equals market cap divided by book value.

Price to Tangible Book Value (MCAP/TBV)

A final metric is price to tangible book value, MCAP/TBV. It's MCAP/BV with intangible assets removed from the denominator.

Expensiveness Benchmarks

I like MCAP/FCF to be no higher than 8, and EV/OI to be no higher than 7. I prefer both MCAP/BV and MCAP/TBV to be no higher than 3.

Valuation

Note how late valuation comes in the value investing model. It follows two time-consuming steps. This contrasts to a more common approach to investment analysis that starts with valuation. Valuation comes late for two reasons. First, ignorance of valuation fosters objectivity in the earlier steps. It keeps one from being from being unproductively swayed by apparent cheapness or priciness. The second reason valuation comes late in the model is because the analyses that precede it are durable. An understanding of a business, and a finding that its good, both have shelf life. They're valid for a while. They can be called into service whenever a stock price takes a sharp turn.

The price metrics often afford us the unsatisfying insight that a company – one that we understand, and one that is good – isn't inexpensive. This can be frustrating since we'd like to buy it. What do we do? We wait. We wait for years, if necessary.

Inexpensiveness can be detected with four different price metrics:

1. Times free cash flow (MCAP/FCF)
2. Enterprise value to operating income (EV/OI)

3. Price to book (MCAP/BV)
4. Price to tangible book value (MCAP/TBV)

Waiting for understandable, good companies to become inexpensive is made easier by recognising:

- The difference between action and progress
- The option value of cash

Chapter 15 – Price Drives Risk

Value investors buy for two reasons. First, it increases returns. The second reason is that buying inexpensively lowers risk. Value investors see the risk of investing in a stock differently than do most in the money management industry. We define it as the chance of losing money. But the industry defines it as the average daily change in the price of that stock over the last month. The industry calls this volatility.

Value investing's take on risk upsets another apple cart as well. It's the risk-return trade-off. This principle holds that the way to achieve higher potential returns is to take on greater risk. By this logic, the road to 15 percent performance is paved with angel investments, and misunderstood commodity funds.

Value investing doesn't just reject the risk-return trade-off. It flips it over. Instead of insisting that outperformance can happen in the face of lower risk, it says that outperformance happens because of lower risk. This perspective gives rise to one of value investing's better -known concepts: the margin of safety. It's the big discount that astute investors require to buy shares. It's the space we allow ourselves to be wrong. It's why we don't wait for a stock to get a little inexpensive. We wait for it to get undeniably inexpensive. Then if we've botched part of our analysis, we're still unlikely to lose significant money. Surprisingly for most, but intuitively for us, covering downsides this way has the happy effect of generating uncommon upsides.

Chapter 16 – Misjudgement and Misaction

Know what to do, do it, and don't do anything else. The first of these three steps, know what to do, is the one that calls for the lower three steps. We know what to do after we make sure that we understand the business, that it's good, and that it's inexpensive. Know what to do pulls from finance and strategy. But the next two supra-steps – dot it, and don't do anything else – are different. They pull entirely from psychology. To do it is to decisively take action when it's time to do so. It means, for example, to buy when an understood, good business is under-priced. The third supra-step, don't do anything else, is even harder. Since value is the approach that performs best over the long term, diverting cash to other strategies is likely to dilute overall performance. This means no over-priced glamour stocks, no misunderstood mutual funds, and no angel investments in technology start-ups.

Cognitive Biases

Knowing how our instincts can lead us astray is the key to keeping from doing so. This is why the astute investor rejects judgments born of cognitive biases. I focus on 18 of them.

Affinity

First is affinity. Affinity is liking. This bias urges us to buy a stock because we are fond of something associated with it. Affinity can also work in reverse. We may prematurely stop considering an investment because we dislike something associated with it.

Reciprocity

Second is reciprocity. Reciprocity is the tendency to treat others as they have treated us. Professional money managers are often courted by public company investor relations departments. Gracious treatment gently nudges them toward buying, and away from scrutinising. Reciprocity is different from affinity. While affinity doesn't require specific action by the investee toward the investor, reciprocity does.

Anchoring

Third is anchoring. Anchoring is benchmarking against an insignificant baseline. It can manifest itself as buying just because price drops from a high level. Anchoring also works on the sell side.

Authority

Fourth is authority. The authority bias inclines us towards investment ideas endorsed by those we admire. It's following. Correspondingly, the authority bias can motivate us to ignore ideas backed by those we consider beneath us.

Availability

Fifth is availability. The availability bias emphasises information that's foremost in our minds, regardless of its overall relevance. We're predisposed to place undue wait on memories that surface with force. This can cause us to overlook other meaningful facts. Investing well requires more balanced recollections.

Cleverness

Sixth is what I call cleverness. The cleverness bias inclines us toward ideas that make us feel smart. Flipped over, the cleverness bias pushes us away from investments that seem too simple.

Incomprehensibility

Seventh is what I call incomprehensibility. The incomprehensibility bias causes us to want an investment more the less we understand it. It has us conflating unintelligibility with specialness. Inverted, the incomprehensibility bias causes us to reject opportunities that present plainly. The incomprehensibility bias is different from the cleverness bias. The cleverness bias thrives on ideas that are truly complicated, while the incomprehensibility bias just needs ideas to be expressed in a complicated way.

Consensus

Eight is consensus. Consensus is doing what everyone else is doing. It tilts us toward popular investments. This creates the risk of overpaying.

Peculiarity

Ninth is what I call peculiarity. Peculiarity inclines us toward investments that are unique. The opposite of consensus, it afflict contrarians and those particularly fond of their own ideas. It's similar to the cleverness bias except that it hungers for opportunities that are merely different, not necessarily complicated. It keeps us from learning from others.

Intermixing

Tenth is what I call intermixing. It causes us to favour holdings that seem offbeat. It predisposes us to securities that would stand out in our portfolio. Note that intermixing is distinct from diversification. Diversification is filling a portfolio with different kinds of assets so that not all of them are ever down in price at the same time. It may be an acceptable consideration. But intermixing is picking different assets just for kicks.

Consistency

Eleventh is consistency. Also known as commitment and consistency, it encourages us to actin concert with our actions. It creates an unwillingness to change our mind, particularly when doing so would require us to reverse a position others know us to have taken. Consistency is sticking with a stock merely because one already owns it. If the fundamentals of a company have truly deteriorated, not selling in the name of consistency is a mistake.

Confirmation

Twelfth is confirmation. The confirmation bias inclines us toward opportunities that support our pre-existing views. Correspondingly, the confirmation bias repels us from ideas that run counter to out convictions. It causes us to reject good opportunities because they fall outside of our conception of a trend, an industry, or market.

Hope

Thirteen is hope. As much as disposition is a bias, hope urges us to believe in an investment. It has us curating available information to make the positive stand out. Correspondingly, it has us supressing any unflattering data.

Lossophobia

The fourteenth is what I call lossophobia. It's more formally known as fear of loss. We feel it when we have the urge to place a sell order as the price of a holding drops, even though he fundamentals of the company remain solid.

Scarcity

Fifteenth is scarcity. The scarcity bias predisposes us toward investments that we perceive to be in short supply. The rarer it seems, the more we want it. Promotors of initial public offerings leverage this bias when they tout an IPO as oversubscribed.

Hotness

The sixteenth is what I call hotness. It's the belief that one is on a roll. More formally called the hot hand fallacy, it surfaces after a string of good investment outcomes. It wells up as the sense of being on a winning streak. Operating inversely, hotness causes us to withhold investments after a few consecutive bad results. We may come to believe that we are on a losing streak.

Miscontrast

The seventeenth is what I call miscontrast. This bias causes us to favour investments that aren't good in an absolute sense, but merely better than others we're seeing at a particular time. It creates the risk that we'll tie up capital in a marginal investment, leaving us without the money necessary to participate in a truly great one when it surfaces.

Windfallapathy

The eighteenth is what I call windfallapathy. People who receive a lot of cash unexpectedly – a big inheritance, for example, or proceeds from the sale of a family business – may develop a certain carelessness toward investing.

Several cognitive biases can act simultaneously. This can strengthen their detrimental impact.

Lack of Restraint

There are two sorts of lack of restraints. The first is impetuosity. It's acting on an urge. It's doing something without thinking about it. It's hard to fight, because it feels like the instincts that

serve us so well in sports, personal relationships, and other aspects of life. Beating back on impetuosity hinges on recognising it as fleeting. It feels strong when it wells up, but it resides just as sharply. But choices made under its influence last. They can result in permanent loss of capital, a result that's anything but fleeting.

The second lack of restraint is weakness. It's passively capitulating to passion after deliberation. It's doing something fully aware that it's foolish. The is the sort of lack of restraint seen during the financial crisis.

Cialdini studied sales situations. He found that goods sales professionals use a combination of six methods to inspire customers to buy: affinity, authority, commitment and consistency, consensus, reciprocity, and scarcity. Value investors were quick to respond to his findings. They saw the six methods as reasons why financial actions that shouldn't be taken are taken.

Once one accepts psychology as a force, it can be harnessed to advantage. I try to put it to work for me. Before placing a buy or sell order, I tick through the list of 20. I consider the 18 cognitive biases and the two forms of lack of restraint. The process doesn't take long, and it occasions just enough introspection to keep me from doing something foolish. I may sill make a mistake, but it's less likely to be of the psychological sort.

Just knowing the dangers of an activity does a great deal to inoculate us against them.

Investor misjudgement is caused by some combination of 18 cognitive biases:

1. Affinity
2. Reciprocity
3. Anchoring
4. Authority
5. Availability
6. Cleverness
7. Incomprehensibility
8. Consensus
9. Peculiarity
10. Intermixing
11. Consistency
12. Confirmation
13. Hope
14. Lossophobia
15. Scarcity
16. Hotness
17. Miscontrast
18. Windfallapathy

Investor misaction is caused by one of two forms of lack of restraint:

1. Impetuosity
2. Weakness

PART III – MAINTENANCE

Chapter 17 – Portfolios and Selling

Buying

When an understood, good company gets inexpensive, we buy its stock. But how much? My rule is simple. Provided that I have enough uninvested cash, I put 10 percent of the portfolio in it. If I'm not comfortable putting at least 10 percent of the portfolio into an equity, I don't want the equity. If my conviction is lower, I don't buy less, I buy none.

A strong conviction is important in part because right after a buy, the price of a stock is almost certain to drop. That's the corollary to another near-certainty: that the price paid for a stock is unlikely to be a low. Rock-bottoms don't send out invitations. So, knowing when one will happen is impossible. The astute investor counts on missing them.

Correspondingly, I prefer not to put more than a tenth of the portfolio into a single equity. This reduces the chance that I'll lack the cash necessary to take advantage of other opportunities that emerge.

Selling and Taxes

There are two problems with selling. The first is taxes. Just how much this eats into long-term returns is best illustrated by example. Picture two portfolios. Each starts with inly cash, buys only non-dividend-paying stocks, and liquidates after 30 years. Assume that any stock sales are subject to a total long-term capital gains tax rate of 30 percent.

Portfolio one uses all its cash to buy stock on the first day. It appreciates 15 percent before taxes every year. It doesn't sell anything until the liquidation date, at which point it immediately pays any taxes due.

Portfolio two also uses all its cash to buy stock on the first day. It too appreciates 15 percent before taxes every year. But it churns its holdings annually. At the end of every year, it sells everything, and uses all the after-tax proceeds to instantly buy different stocks. When it liquidates after 30 years, it too promptly pays any taxes due.

Portfolio one would end the 30-year period with more money. But what's striking is just how much more. It would wind up with over twice as much cash.

Selling and Alternatives

The second problem with selling is alternatives. Companies that are understandable and good don't go on sale every day. They're hard to find. So, absent an acute cash requirement, each stock sales mandates a hunt for the next opportunity.

Reasons to Sell

Selling does make sense in some instances. I see four:

The first is when price flies past value. If EV/OI is over 25, and there are no mitigating facts, I find it hard to justify holding.

The second instance is when a company that originally registers as good turns out not to be. The cognitive bias of consistency can make it hard to see such instances, we may want to hold just to validate our buys. But analysis really can be wrong, and contexts really can change. Selling in such

situations keeps a snag from ripping into both a realised loss and a missed chance to redeploy cash into a better opportunity.

The third instance is when one is bought out. Public companies sometimes get acquired.

The fourth instance is when cash is needed to make an investment that's clearly better than one already held. The problem with this is that fresh ideas often glow with a special promise. They're new. The hope bias gets a prime shot at causing mischief. As such, I get extremely suspicious of my reasoning when I think that I'm spotting such a circumstance.

When a sale is wise, its justification is distinct. It's an overpricing, an analytical error, a contextual change, a buyout, or a better opportunity. Absent that clarity I hold.

Dividends

Even without active selling, an equity portfolio can generate cash through dividends. Dividends can become sizable. This fact gets lost in the commonly quoted metric of dividend yield. Recall that dividend yield equals annual dividends divided by current stock price. But to an owner, current only counts in the numerator.

Concentration

I choose to concentrate because I've observed over time that good, focused stock portfolios outperform diversified stock portfolios. This is because diversified portfolios are more like an index. Of course, a bad focused equity portfolio can certainly lag a diversified stock portfolio. Concentration isn't enough to assure outperformance.

Cash

While I don't diversify within my equity portfolio, I do diversify outside of it. I always keep enough cash on hand to cover expenses for a few years. Sequestering cash enables me to confidently ride the wild swings guaranteed to come with a concentrated portfolio. It's what lets me take the long view.

Many governments insure bank deposits. It's wise to stay well under the insured limit. Opening up accounts at several different banks is not hard. Sequestered cash is best held in the same currency as one's expenses.

- Conviction prepares one for the likely price drop that follows a stock buy.
- Selling stocks can make sense when price flies past value, when a company thought to be good turns out not to be, in buyouts, or when a clearly better opportunity emerges.
- The problems with selling are taxes and alternatives.
- Questionable reasons for selling include rebalancing, memorialising success, and industry compensation.
- Equity portfolios can generate cash without active selling through buyouts and dividends.
- Good, focused equity portfolios outperform diversified equity portfolios over the long term.
- Cash sequestered for ordinary expenses un government-insured accounts makes equity portfolio price gyrations less troubling.

Chapter 18 – Endurance

To achieve long-term after-tax outperformance, one must hold onto stocks in good companies for a long time. It's easier to do this with a portfolio that's consistent with one's moral posture. It's therefore useful to define one's ethical disposition early. I've seen smart people do it in one of four ways.

The first is amorality. It's not viewing investments through an ethical lens. It regards money management as an activity that sits outside of moral consideration. Put differently, it sees the ethical imperative of investing as growing wealth.

The second is what I call moral failure abstention. It's not investing in companies with certain unsavoury characteristics. It's a list of don'ts. It often entails dismissing firms that make certain products. Cigarettes and handguns are two common current examples.

The third is what I call moral success affirmation. It involves only investing in companies with certain desirable characteristics. It's a list of do's. Again, this is commonly related to a firm's products. A popular current example is renewable energy.

Fourth is what I call moral failure activism. It involves buying stock in companies with undesirable characteristics for the purpose of pushing for change as a shareholder.

There are four different moral postures toward investing:

1. Amorality
2. Moral failure abstention
3. Moral success affirmation
4. Moral failure activism

Chapter 19 – Generating Ideas

The value investing model needs to be fed stock ideas. These ideas can come from different sources. I look at seven:

Bad News

The first source is bad news. Stories about companies often emphasise an extreme element of an event. These extremes get amplified in headlines. Headlines drive human reaction, sometimes too far. This can cause stock price swings deeper then would a more sober take on the facts.

Spin-offs

A second source of ideas is spin-offs. A spin-off is the public listing of a company that was previously part of another listed company. Often, some of the old parent's shareholders are institutional investors. When their spin-off shares start trading, they may sell them automatically. This is because the new stock doesn't meet their formal investment criteria, such as minimum market cap. This forced selling in shares can depress prices in companies that, if they're both understood and good, are worth owning.

Regulatory Filings

A third source is regulatory filings. Many governments require large investors to periodically report their holdings. These filings are public. One can compare reports between periods to see which stocks talented professionals bought.

Reorganisations

The fourth idea source is reorganisations. It could be a merger, a big change in capital structure, or the sale of a major division.

Small Cap Stocks

A fifth source is small capitalisation stocks, also called small caps.

Stock Screeners

A sixth source is stock screeners.

Serendipity

Seventh is serendipity. Serendipity is the mental preparedness to receive tips from everyday life.

Once one feeds an idea into a model, it's best to forget where it came from. When we forget that we're looking at a company because it's a spin-off, it's owned by a hero, or its stock price plunged, we keep a whole raft of cognitive biases at bay.

Promising sources of investment ideas include:

1. Bad news
2. Spin-offs
3. Regulatory filings
4. Reorganisations
5. Small-caps
6. Stock screeners
7. Serendipity

Chapter 20 – Differences Among Value Investors

Dimensions on which bona fide value investors can differ include:

1. Asset class
2. Holding period
3. Activism
4. Diversification
5. Quality
6. Leverage
7. Complexity
8. Shorting

Chapter 21 – Preservation

- Capital preservation is a value investors priority because of the mathematics of realised loses.
- The risk-return trade-off blinds most asset managers to the primacy of capital preservation.
- Most actively managed equity funds fail to beat basic market indexes over time.
- The bottom layer of the value investing model is the part most likely to change.
- Thinking in percents encourages habits that work hard over a lifetime.
- Value investing has benefits beyond remunerativeness.