Rational Numbers

At the tail-end of the 19th century, a small-town Carolinas chemist named Caleb Bradham hit upon an exciting new formulation. A marvellous tonic that served as both a refreshing pick-me-up and a remedy for indigestion, an affliction more formally known as *dyspepsia*.

And thus in 1893 Pepsi Cola was born. But Pepsi's early years were not problem free. The hit from WWI sugar rationing was so severe that it brought the company first to its knees and then to the feet of its bigger rival Coca-Cola. But in one of fizzy-pop's more epochal (and perhaps apocryphal) moments, Coke dismissed Pepsi's capitulation, declining their offer to buy the embattled company. Coke presumably deemed the $50,000 asking price too high - but in doing so unwittingly set the scene for decades of bitter competition. Said with the benefit of hindsight, it is hard to imagine Pepsi setting any price that would have been truly unreasonable in this instance[1]. With 2020 drawn to a close, Pepsi now matches Coke for soft drinks sales, earns the same again from snacks, and sits just shy of its $200+ billion market cap.

We can't really blame Coke though. How on earth should they have valued the then bankrupt Pepsi brand back in 1933? But as investors today can we avoid repeating their mistake? How much is PepsiCo really worth and is its reappraised/market-derived price-tag anywhere close to being correct?

Well, we can start by noting that the 1930s were not unique here. Sadly, we don't have a century of market data to scrutinize, but Bloomberg obliges us with four decades worth - from which it's apparent that 1980's investors also struggled with valuations. In 40 years, PepsiCo's market cap has grown 100-fold, whilst the S&P500 has risen 'just' 29x. So relatively speaking, something clearly went wrong. But 1980 was a long time ago. A simpler time, before the arrival of Bloomberg and the World Wide Web (or for that matter Pepsi Max). With better tools and more information, do today's sophisticated traders now have Pepsi sussed? Can there really be that much extra fizz left in this aged hundred-bagger? As it happens recent performance has been far more pedestrian (during 2020, PepsiCo shares appreciated 8.5% vs. the S&P500's 16.3%). Should we now expect boilerplate returns from here on out?

These questions are important and of direct relevance not just to Pepsi, but to any high growth/hard to value stock. And before you scoff at the idea of a 127-year-old consumer staples giant posing as a growth stock, reflect on the fact that over those 40 years PepsiCo's dividend per share compounded at a rate of 11% pa. The S&P's rate was half this. If not a sprinter, then think of Pepsi as a world champion distance runner. Company-sourced data takes us back even further, to 1955 from whence Pepsi's dividend has grown at over 10% pa (see figure 1). That's 65 years of double-digit growth.

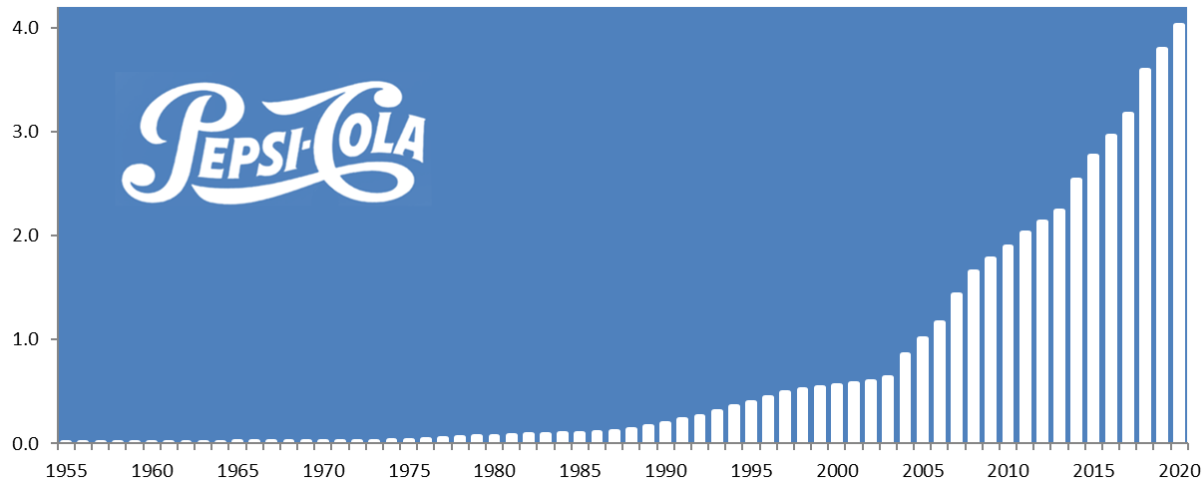

Figure 1: Pepsi's dividend per share chart (in USD and adjusted for share splits), which has compounded at a rate of 11% pa in the 40 years since 1980 and at a rate of 10% since 1955 (as far back as we can get clean data).

[1]For some numerical context, consider the following. For Coke to have received a nominal 8% p.a. market return on its investment (and assuming Pepsi's end-2020 market cap holds) they could have paid anything up to 5,000 times Pepsi's actual asking price in 1933.

So, to begin answering the above questions, let me propose the following: If, as we think plausible, Pepsi's historic rate of growth is maintained into the future, then today's already mammoth market cap will turn out to have *dramatically undervalued* the company's true worth.

Seeking elaboration, an informed sceptic might well query how we could make this excessive - perhaps even irrational claim. Surely in a relatively efficient market such as the US's, the price of a big liquid share like Pepsi's cannot be too far wrong? By what counter logic can we reason it to be so mispriced?

First, let's refresh ourselves on the fundamentals. Theory contends that pretty much every cash generating asset has a 'fair' price. A price at which the purchaser should then over time expect to receive a fair, market rate of return (which in principle varies given the level of 'risk' taken - but we'll leave that for another day). In fact, for a company destined to grow its dividends at a constant rate into perpetuity, this expected return is directly quantifiable via a popular formula known as the dividend discount model or DDM. This calculates the return (r) as next year's paid out cashflow - the dividend (d) in the case of stocks, the coupon for bonds - divided by the asset's price (P), plus the aforementioned constant rate of growth (g). i.e., r=d/P+g or alternatively arranged for price as P=d/(r-g). In theory, in an efficient system, participants collectively set a price such that everyone receives the same market rate (often known as 'R' or the discount rate) regardless of what is held. If not, then traders should rush to the cheaper assets and correspondingly flee from expensive stocks. Hence a higher rate of growth is compensated for by a steeper price, which normalises the return as defined by the formula. This is both accepted wisdom across the industry and foundational for many investment approaches.

But there is a universally acknowledged catch to the DDM. Dividends rarely (i.e. never) grow at a constant rate forever, and even if they did no one actually knows what it might be. Hence, as we'll all - 1980 Pepsi investors included - acknowledge from immediate experience, every asset *does not* deliver the same rate of return. Not even close. For another, deliberately starker example, try telling Groupon's 2011 IPO investors that they 'should' have done as well as Google's. At the time of writing Groupon's investors have lost 90% of their money, whilst over the same period Alphabet's have almost quintupled theirs.

As it happens, this hitch doesn't actually disqualify the theory (the price is only deemed correct as far as the facts were knowable at the time, so unpredictable events can permissibly drive disparity) but it certainly limits its application in practice. Nevertheless, for want of a general alternative, and with a few tweaks to assumptions here and there, the DDM is still very widely used. Notably, restated for cash, it forms the basis for most discount cash flow (DCF) terminal value calculations. The terminal value being the 'second stage' chunk of net present value (NPV) left from year five (or whenever your explicit predictions have dried up) onwards and usually emerges from a DDM-styled equation upon application of some 'equilibrium' growth rate. We can illustrate this with PepsiCo: five years of 11% growth applied to its c.$7bn of free cash flow, followed by a generic 3% perpetuity growth terminal value, all discounted back at 8% gets you (perhaps coincidentally) very close to the current market price. As the years project out, the discount rate's power law relation squeezes the present value of distant flows into irrelevance, meaning the terminal value is calculable as a finite entity. In this case the terminal value makes up 80% of the NPV, but oddly is arrived at almost by default. An obvious initial concern here is that in the (likely) event that a company doesn't last forever, the model can often overstate.

But the theory has another, more conceptual flaw that can lead the DDM to fundamentally *underestimate* value. This is more nuanced and only troubles a tiny minority of target companies - but it's central to the LT valuation approach, as it's these companies we're looking for. That is businesses whose expected long-term rate of growth is higher than the market's fair rate of return. If this is true into perpetuity, then looking back at the above r=d/P+g equation tells us that there is *no price* at which such a growth rate can be compensated for. In fact, the fair price to pay for such an asset becomes infinite. This phenomenon is sometimes known as the Growth Stock Paradox after David Durand's remarkable 1957 Journal of Finance paper which applied the famous and functionally analogous St. Petersburg gambling paradox to the problem.

The St. Petersburg Paradox as laid out by Bernoulli in the eighteenth century, debates the appropriate wager for a coin toss game whereby the payout doubles for every successive tail that's thrown. As the theoretical number of headless flips trends out to eternity (even with their decreasing likelihood) the series' expected value actually sums to infinity (i.e., P=1x1/2+2x1/4+...etc). By analogue to our growth stocks, the game's probability halving as the payout doubles has the same effect as a company's growth rate equalling the discount rate. The paradox in both cases stands that although a participant is in practice guaranteed to win/earn a finite amount, their theoretically fair entry fee is infinite. Many investors have referenced to this specific problem of valuing growth stocks, including John Burr Williams (widely credited with deriving DCFs in the first place) and Benjamin Graham, who included a whole chapter on them in his 1962 (fourth and final author-updated) edition of Security Analysis.

At this point though, most sensible practitioners will pause to remind us that this isn't really a problem, as such perpetual growth stocks don't actually exist in real life. (As the saying goes, *'anyone who believes exponential growth can go on forever in a finite world is either a madman or an economist*'). Some moderation is clearly called for - but how, and how much? Growth rates do not need to be that high before we run into trouble. If the market's fair rate of return is, say 8-9% nominal[2] then any double-digit perpetuity growth rate will blow up the equation. So, should we dial back the growth, or rationalise our time horizon? Arbitrarily constraining growth forecasts seems peremptory, but if left to run - even over finite periods - high rates eventually stop being credible. $1,000 compounded at 10% pa over 250 years will match the current gross domestic product of the US economy. Give it another 15 years and it will exceed today's entire global output. So, whilst bounding the problem resolves the infinity paradox, it can still yield unrealistic figures[3]. What company could possibly grow at such a rate for so long?

But we do not need to stretch into centuries. There is a multi-decade middle ground here too; of companies that can grow at unusually high rates for unnaturally long periods of time. Long enough to exhaust standard valuation models, but not so long as to test rational limits. Companies that compound over time and get much bigger - orders of magnitude so - but don't quite outgrow the planet. They are rare, but as PepsiCo has demonstrated, they do exist, and it is exactly this type of company that we are aiming for in our portfolios. Beyond Pepsi's superlative six-and-a-half decade run we can also count (again allowing for the availability of data) Unilever having grown its dividend at 10% p.a. for the past 30 years, Brown Forman at 10% for 35 years, Hershey at 10% for 40 years, Nintendo at 11% for 30 years and Disney at 14% for 40 years. If Pepsi and co continue to grow at similar paces over the coming decades, then the traditional two-stage discount formulation doesn't stand a chance.

The second, DDM-derived terminal value step will not work to uniformly assess all assets, least of all our above growth stocks. Growth rates as high as theirs will crash the calculation, whilst artificially lowering them in the name of conservatism seems at odds with empirical reality. Indeed, accommodating this latter option usually requires some desultory compromise between scenarios that could be starkly binary. A company facing terminal disruption is probably worth very little, whereas one set to grow at high rates for decades hence will be worth far more than even a bullish mainstream model might imply. Recall our earlier example of this divergence; when Groupon listed back in 2011 it traded on a very similar c.6x EV/sales multiple to Google, yet Groupon's subsequent growth was negative, Google's stratospheric. Standard valuation approaches demanding modest sub-R terminal growth rates would not have helped investors in either case. Put another way, conventional models only work for conventional companies.

Of course, as a quick fix the first stage of the discount model can be extended past the typical five or so years. But how far... and can we possibly have enough visibility to make any sensible estimates? Will Pepsi's historic double-digit dividend growth rate persist for another 40 years? And if so, then why not 50 or 60 years? Going back to our PepsiCo DCF example calculation from before - if we stretch the high growth period to say 40 years, PepsiCo's 'intrinsic' value jumps by a massive 4.6x. Push this out to 60 and you double again the theoretical NPV to over 9x the current market cap. How can we possibly know which if any of these assumptions are correct? So, yes, it is possible to coax growth stock valuations (and very high ones at that) from the first stage of the model by adapting it to more historically consistent ranges. However, to home in on anything solid requires pinning down the growth period's duration with impractical precision. Even then, we are still reliant on the procrustean terminal value for an unpalatably large part of the sum - it speaks for 40% of the NPV even in the 60-year growth example.

Can we do any better? Now that we have at least an idea of the relevant (multi-decade) time horizons, what other techniques do we have that might yield firmer figures? Here is another take on the problem. Picture our long-term 1980's Pepsi investors, 10 years into their journey. It's now 1990 and much has changed - Microsoft Windows, mobile phones and New Coke have all debuted - and throughout, Pepsi's owners have benefited from a decade of double-digit dividend growth. What's to come? Well, realise it or not, they've still got social media, smartphones and Pepsi Blue (Google it!) to look forward to - plus at least another three decades of double-digit dividend growth. Now picture Mr Market entering the scene as a 1990 Pepsi co-investor. And for the sake of argument let's credit him with the ability to see clearly ahead, but only for five years. As such we might assume he will construct his valuation models accordingly. i.e., by discounting back the five years of predictable growth before applying a generic second-stage DDM-based terminal value. The interesting result of this is that as Pepsi's growth rate continues to maintain its high

[2]No one really knows the true value for R (if it even exists) and this remains a fascinating, unsolved academic debate, with gilt yields, CAPMs, factors, and WACCs all employed in attempts to theoretically derive it. But, empirically at least, the US market's return has averaged at a 'Siegel number' of roughly 6.5% real over the past 200 years.

[3]If we assume inflation of say c.2%, then even converting to real figures you still arrive at a big (though arguably more feasible) number - as your $1,000 compounded at a real 8% for the same period would leave you with a cool $720 billion to spend. Put a 10x multiple on this cash flow and you have the equivalent of three-and-a-half Apples on your hands. Several companies have matched this longevity, few this size.

rate[4], then sentiment aside, there will be no reason at any point for Mr Market to change Pepsi's valuation multiple for at least the next 25 years. Five years from now, Mr Market will still see five years of high growth ahead, as he will five years on from that, and so on. Hence in this idealised scenario, the 25-year annualised return for a 1990's Pepsi investor should simply equal the starting dividend yield plus its growth rate as dictated by our r=d/P+g DDM formula.

How does this relate to the growth paradox? Well, under the above conditions, the initial price you pay can actually rise indefinitely without seriously damaging subsequent holding period returns. As the opening price rises towards infinity, the dividend yield will tend to zero. *But* once the yield has ceased to be an important contributor, and assuming the valuation multiple holds in line with the consistent rate of growth, then increasing the asking price will have no further impact on expected returns (which should now match the dividend's growth). This brings us exactly back to our paradox. This time however, conceptualising a finite holding period removes the requirement for genuinely perpetual growth; provided it satisfies myopic Mr Market's two-stage models for as long as you plan to invest.

As always, exercise caution here. Not only might you and others be wrong with future growth projections, but Mr Market is notoriously manic and shifts in sentiment do occur. This in fact more or less describes what actually happened in the above real-world sequence of events. From 1990 to 2015 Pepsi's dividend valuation took a meaningful hit (in part thanks to the 2008 financial crisis) with the yield jumping from 1.5% to 2.7%, annulling its reinvested return contribution. However, the dividend continued to compound at 11% pa - as did the total returns to investors.

Clearly the higher the price you pay, the more crushing any multiple compression could be, and I would advise thinking of high prices as heightening this particular risk. Hence, I advocate building in some, even quite severe, multiple compression into any fixed holding period model as a measure of warranted conservatism. Doing so is also the only way to resolve this version of the growth paradox into a system capable of yielding finite figures. In previous notes (see for example 2016's *Re-evaluating Valuations or 2019's Valuation Challenge*) I have explored this line of attack - tempering high, multi-decade holding period growth rates with gradual multiple compression (for example down to a market-average figure) headwinds. The advantage here is that, conceptually at least, you can sidestep the impossible precision of timing any supposed growth fade beyond your holding period. And whilst this too is an input dependent procedure, it has so far yielded usefully consistent, and encouragingly (in our eyes) bullish results. Applied to PepsiCo (assuming 11% nominal growth, a PE decay down to a market multiple of say 22x over a 20-year holding period, discounted at a nominal 8%) suggests 100% of 'upside' from the current market cap, or a warranted PE of closer to 50x.

Note that mathematically, this fixed holding period methodology mirrors that of a two-stage DCF (assuming an equivalent high growth period: 20 years in the prevailing example) with a presupposed market multiple used to obtain the terminal value. The comparison is not exact for the above referenced models as these incorporate other important details concerning the treatment of dividend reinvestment (complicated by the falling multiple, though iteratively resolvable) and the discount rate. More pertinently though, there is a conceptual difference to their construction. To reiterate the above, a fixed holding period alleviates the requirement to prescriptively bound a company's growth - you just need to anticipate (and enforce) the duration of your investment, which whilst the longer the better, will usually be dictated by personal or professional circumstances. Hence, although at the end of the holding period the stock is pessimistically assumed sold at a predicted market multiple, there is no presumption against ongoing underlying growth. Indeed, if high growth rates do persist, then this notional mean reversion could prove temporary. Multiple compression therefore provides a (relatively) conservative, paradox-resolving lower estimate for the valuation and attempts to account for the practical difficulty of timing a sale at other people's prices. This is a subtle point but means that whilst the model's construction can be roughly aped with a two-stage DCF, it is arrived at (and as necessary adjusted) via a different thought process. Thus, in the absence of satisfactorily objective measures, you can instead probe the stock's worth to *you*, given your holding period, your assumptions about long term growth, and your understanding of the approaches taken by others.

Here our sensible sceptic might caution us that this is all getting dangerously close to bubble talk. A few years after the dotcom-crash, Warren Buffett (in reference to Durand's paper) wisely counselled that, "*The idea of projecting out extremely high growth rates for very long periods of time has caused investors to lose very, very large sums of*

[4]Of course, Pepsi's dividend did not increase at a metronomic 11% pa every year, and five-year CAGRs over this period range from as high as 15% to as little as 5%, geometrically averaging out to 11% pa from 1990-2020. For simplicity I have assumed Mr Market sees a smoothed compounding progression, but fully acknowledge that real world deviation from this may cause short-term boosts/dips to be felt in his ongoing valuation.

money". Likewise, other contemporary observers explicitly linked this growth-stock thinking to what was widely seen as temporary turn-of-the-millennium madness (e.g., see Székely & Richards' 2004 article *The St. Petersburg Paradox and the Crash of High-Tech Stocks in 2000* for an interesting example). Both were and are right to urge caution, but don't forget the exceptions that prove the rule. Not every dotcom growth stock was overvalued[5]. As demonstrated, long-lasting high growers do exist. You just have to be *very* careful when identifying them - and keep your valuations rational. PepsiCo, with 65+ years of profitable double-digit growth under its belt, is not some loss-making fancy. It is an exceptionally special case. And yet caution is still required. Note for instance that some of Pepsi's historic dividend growth can be attributed to an expanding payout rate, which may not be sustainable into the future. The rate of growth has also (and perhaps accordingly) abated somewhat over the last few years and it is up to each investor to decide if this amounts to temporary slack or forewarning of a more permanent attenuation.

If we are prepared however to presume continued growth, then can we rationalise our bullish calculations into something more tangible than a table of CAGRs? Why should these companies be so valuable (and so hard to value)? Think in strategic terms. How much would you have to spend to recreate Pepsi today? $200 billion? Or would it really even be possible to replicate the recognition and reputation that comes from 127 years of brand investment - or its commanding positions in salty snacks and iconic soft drinks... *regardless* of how much you spent[6]? It is plainly a step too far to conclude that Pepsi has infinite value. But it's a step in the right direction. So, semantics perhaps, but let's instead think of Pepsi in colloquial terms as an 'invaluable', possibly even 'priceless' asset.

Similarly, how much would it cost to recreate Unilever's reach to a third of the planet's population? or Jack Daniels' 'money can't buy' rock-and-roll heritage? or Nintendo and Disney's classic cast of characters? or Cadbury's seven decades built, 65% chocolate market share in India - and Hershey's corresponding dominance in the US? How much would Hershey need to spend to replace Cadbury in the UK and India, or Mondelez to unseat Hershey in America? Can it be done? As it happens, both have tried, and both retreated. Instead, Hershey offered to buy Cadbury in 2009 and Cadbury's eventual owner Mondelez tried to buy Hershey seven years later.

Happily, strategic acquisitions like these can provide some additional insight. not least as the deals always take place at valuations *above* the undisturbed and supposedly correct market prices (typically rationalised as control premia). We use such data points to gauge our own analysis, even if we do not necessarily expect, or even want, our companies to be targeted. For example, it's interesting to know that Coke (having belatedly learnt their lesson?) thought it fair to pay 10x sales - almost three times Pepsi's current market multiple - for Glaceau's Vitaminwater. However, there's ample room for error here. Companies can always overpay (though in the above example Coke was able to triple the brand's sales in as many years) and anyway, available examples will never provide perfect comparators. On top of this, their cherry-picking invites confirmation bias, meaning this comparative argument is at best suitable as an adjunct to other valuation approaches.

But in the end, none of these models are doing anything truly fundamentally different; all discounted dividend or cash flow calculations have the same theoretical foundations, just as they all eventually condense into a single valuation multiple. But by highlighting alternatives to the standard two-stage approach presented above, we can at least allow the very specific considerations for these rare long duration growth stocks to be more logically reflected - and perhaps that is all we can realistically ask of a model. We will never actually know the 'correct' price to pay and I am doubtful that there even really is one. Value is too subjective a quantity. But the right model will do a better job of framing the problem and in turn, this better directs us towards the right questions to ask. The ultimate practical goal then is to get a sense for how a stock might perform *if* our input assumptions hold. And you need to make sure you have the appropriate model to allow these assumptions to be fully and consistently implemented.

---

[5]With the benefit of 15 years of hindsight, there is an interesting if slightly mischievous observation to be made here. Returning to our above Google example, Székely & Richards' article (and its 2005 follow up) reference freshly listed Google Inc. as a clear example of the risks of applying high perpetuity growth rates to common stocks, trading as it did at the time for a dizzying $198 per share and a PE of over 70x. They write that, *"For a quantitative measurement of the extent to which Google's common stock may have experienced irrational exuberance since its initial public offering, we would recommend the virtues of Durand's reformulation of the St. Petersburg paradox. ... 19 financial analysts currently estimate an average one-year growth-rate of 35.6% for Google's revenue. [Although] estimated for 2006 only, we find ourselves in a state of utter disbelief [and] in discordance with the 'price targets' offered by financial analysts, we hereby make the public prediction that Google's stock price will fall below $50 by December 31, 2010."* Székely & Richards' scepticism is understandable and would have been correctly applied to many other tech favourites of the time. However, Google genuinely was in possession of a defensible moat and in its case the hype was justified. In the subsequent decade-and-a-half, the company's revenues compounded at 30% p.a. and last month the shares hit an all time high of $1,800, valuing renamed Alphabet at over a trillion dollars.

[6]This 'irreplaceable assets' argument is an important underpin for the newly vogue (if weakly defined) 'quality at any price' school of investing.

So, do we finally have a definitive resolution to the paradox? Not really. As Durand warned over 60 years ago, "*With growth stocks, the uncritical use of conventional discount formulas is particularly likely to be hazardous; for, as we have seen, growth stocks represent the ultimate in investments of long duration. Likewise, they seem to represent the ultimate in difficulty of evaluation. The very fact that the Petersburg Problem has not yielded a unique and generally acceptable solution to more than 200 years of attack by some of the world's great intellects suggests, indeed, that the growth-stock problem offers no great hope of a satisfactory solution.*" The paradox has many putative 'solutions', but in essence each requires the problem to be reframed and hence rationally bounded in some way. Consequently, any realistic growth stock valuation attempt rests on some such compromise. The problem is that these necessary choices are often arbitrary yet dramatically affect the result. But whilst precise prices might elude us, it seems clear that genuinely long-duration high-growth stocks like Pepsi's both exist and are worth considerably more than reckonable via traditional techniques. Hence, to resolve approximate valuations far above the conventionally correct market levels seems to me to be both plausible and entirely rational.

James Bullock, Portfolio Manager
Lindsell Train Ltd

Sources: Lindsell Train & Bloomberg end-December 2020. PepsiCo 2020, Durand 1957 *Growth Stocks and the Petersburg Paradox*, Graham & Dodd 1962 *Security Analysis* 4th Ed., Pendergrast 1993 *For God, Country, and Coca-Cola*, Buffett & Munger *2004 Berkshire Hathaway Annual Meeting*, Székely & Richards 2004 *The St. Petersburg Paradox and the Crash of High-Tech Stocks in 2000*, Székely & Richards *2005 Remain Steadfast with the St. Petersburg Paradox to Quantify Irrational Exuberance*.

Lindsell Train Limited
66 Buckingham Gate
London SW1E 6AU
ENGLAND

Tel. 020 7808 1210
Fax. 020 7808 1229
www.LindsellTrain.com
Info@lindselltrain.com