



City Research Online

City, University of London Institutional Repository

Citation: Hayley, S. (2014). Hindsight Effects in Dollar-Weighted Returns. *Journal of Financial and Quantitative Analysis*, 49(1), pp. 249-269. doi: 10.1017/s0022109014000155

This is the accepted version of the paper.

This version of the publication may differ from the final published version.

Permanent repository link: <https://openaccess.city.ac.uk/id/eprint/6226/>

Link to published version: <https://doi.org/10.1017/s0022109014000155>

Copyright: City Research Online aims to make research outputs of City, University of London available to a wider audience. Copyright and Moral Rights remain with the author(s) and/or copyright holders. URLs from City Research Online may be freely distributed and linked to.

Reuse: Copies of full items can be used for personal research or study, educational, or not-for-profit purposes without prior permission or charge. Provided that the authors, title and full bibliographic details are credited, a hyperlink and/or URL is given for the original metadata page and the content is not changed in any way.

Measuring Investors' Historical Returns: Hindsight Bias In Dollar- Weighted Returns

Simon Hayley
Cass Business School,
106 Bunhill Row, London EC1Y 8TZ, UK

E-mail: simon.hayley.1@city.ac.uk

Tel +44 20 7040 0230

First version: 26 October 2010

This version: 12 March 2012

I am grateful to Ilia Dichev, Stewart Hodges, Aneel Keswani, Ian Marsh and participants at the Midwest Finance Association 2012 annual meeting for useful comments. The usual disclaimer applies.

Measuring Investors' Historical Returns: Hindsight Bias In Dollar-Weighted Returns

Abstract

A growing number of studies use dollar-weighted returns as evidence that consistently bad timing substantially reduces investor returns, and that consequently the equity risk premium must be considerably lower than previously thought. These studies measure the impact of bad timing as the difference between the geometric mean return (corresponding to a buy-and-hold strategy) and the dollar-weighted return. However, the present paper demonstrates that this differential combines two distinct effects: The correlation of investor cashflows with (i) future asset returns, and (ii) past asset returns. Both correlations tend to alter the dollar-weighted return, but only the first affects investors' expected wealth. The second generates a hindsight bias.

This paper also derives a method which separates these two effects. The results show that the great majority of the return differential for mainstream US equities has been due to hindsight bias, and very little due to bad investor timing. Dollar-weighted returns have been low because aggregate investment flows reflect past returns rather than future returns, and these low returns should not lead us to adopt correspondingly low estimates of the risk premium. The decomposition method which is derived here also has many applications in other fields where dollar-weighted returns are used, such as project finance and investment management.

JEL Classification: G11, G12, G15

Keywords: risk premium, dollar-weighted return, investment, equity

Measuring Investors' Historical Returns: Hindsight Bias In Dollar-Weighted Returns

Few figures are of such central importance in finance as the equity risk premium, yet estimates vary widely. In particular, a growing number of studies argue that investors time their investments so badly that on average they earn returns which are significantly below the buy-and-hold return on the corresponding market index. They conclude from this that the risk premium must be substantially lower than had previously been thought.

This emerging literature stems from an influential paper by Dichev (2007), who argues that the impact of bad timing on aggregate investor returns can be deduced using a simple and elegant method. The geometric mean (GM) of monthly market returns gives the return that would be earned if investors followed a strict buy-and-hold strategy, immediately re-investing any dividends. By contrast, the dollar-weighted (DW) return takes account of the net cashflows paid or received by the average investor ahead of the terminal period, such as share issues, dividend payments or share buybacks. The difference between these two rates is then used as a measure of the effect the timing of these cashflows has on investor returns.

Using this method Dichev concludes that poor timing has led to a substantial reduction in investor returns: A 1.3% per annum reduction for equities traded on NYSE and AMEX exchanges (1926 to 2002) and a 5.3% reduction for NASDAQ stocks (1973 to 2002), as shown in Table I. This would imply that the equity risk premium earned by investors (and firms' cost of capital) must be considerably lower than previously estimated.

Table I
Investor Timing Effects Identified by Previous Studies

The table shows the annualized Geometric Mean (GM) and Dollar-Weighted (DW) returns derived by previous studies for the markets and periods shown. A positive differential (final column) is interpreted as the reduction in the return received by investors as a result of bad timing. Distributions are defined as net cash distributions by firms to investors – a negative distribution represents an additional net investment (eg. as investors buy new share issues). The correlation coefficients in columns 4 and 5 are calculated on mean returns over the previous/subsequent three years (Dichev and Dichev/Yu), and one year (Clare/Motson).

Market	Period	Authors	Correlation of distributions with:		GM	DW	GM - DW
			Past returns	Future returns			
NYSE/AMEX	1926-2002	Dichev	-0.26	0.51	9.9%	8.6%	1.3%
NYSE/AMEX	1926-1964	Dichev	-0.41	0.54	9.6%	8.0%	1.6%
NYSE/AMEX	1965-2002	Dichev	0.09	0.44	10.1%	9.4%	0.7%
NYSE/AMEX	1926-1951	Keswani/Stolin			7.5%	5.8%	1.8%
NYSE/AMEX	1951-1977	Keswani/Stolin			9.5%	9.7%	-0.2%
NYSE/AMEX	1977-2002	Keswani/Stolin			12.6%	12.9%	-0.3%
NASDAQ	1973-2002	Dichev	-0.57	0.28	9.6%	4.3%	5.3%
NASDAQ	1973-2006	Keswani/Stolin			10.4%	7.5%	2.9%
19 International stock exchanges	1973-2004	Dichev	-0.24	0.16			1.5%
19 International stock exchanges	1973-2004	Keswani/Stolin					0.7%
UK mutual funds (all flows)	1992-2009	Clare/Motson	-0.18	-0.02	6.5%	5.7%	0.8%
UK mutual funds (retail flows)	1992-2009	Clare/Motson	-0.37	0.09	6.5%	5.4%	1.2%
UK mutual funds (instn'l flows)	1992-2009	Clare/Motson	-0.03	-0.08	6.5%	6.2%	0.3%
US mutual funds (all)	1991-2004	Friesen/Sapp					1.6%
Hedge funds (7190 funds)	1980-2006	Dichev/Yu	-0.22	0.04	10.0%	6.4%	3.6%

Keswani and Stolin (2009) challenge the robustness of these results, noting that the differential for NYSE/AMEX stocks is sensitive to the exact start and end dates chosen. Dichev finds that the differential falls to 0.7% per annum in the second half of the period, but Keswani and Stolin find that it disappears entirely if the time series is split at different points. They also find that the differential for NASDAQ stocks shrinks substantially when four years' subsequent

data are included, and that the differential recorded for international stock exchanges was influenced by a dramatic increase in the proportion of stocks included in these indexes.

However, studies using the same method have found differentials in other markets which are similar to those reported by Dichev. Friesen and Sapp (2007) find a differential of 1.6% per annum for US mutual funds, Clare and Motson (2010) a differential of 0.8% for UK funds, and Dichev and Yu (2011) a differential of 3.6% for hedge funds. A consensus has thus emerged that the aggregate effects of bad investor timing have been substantial.

These studies all use the difference between GM and DW returns to measure the impact of bad investment timing. However, we demonstrate below that this differential combines two distinct effects: the correlation of investor cashflows with (i) future asset returns, and (ii) past asset returns. Both correlations tend to alter the DW return, but only the first affects investors' expected wealth. The second generates a hindsight bias.

Section II derives a method for quantifying and removing the effects of this hindsight bias. The results show that for mainstream US equities (those traded on NYSE and AMEX) the great majority of the differential between DW and GM returns has been due to the hindsight bias, and very little due to bad investor timing. DW returns are low because aggregate investment flows reflect past returns, rather than future returns. The effect of bad timing of investment in NASDAQ stocks is also much smaller than was initially calculated. Furthermore, Table I shows that distributions in other markets are generally much more strongly correlated with previous

returns than with future returns. This suggests that the return differentials recorded for these markets are also likely to be largely due to hindsight bias.

This contribution of the present paper is: (i) it helps to resolve the current debate about the equity risk premium by showing that low DW returns do not imply correspondingly low risk premia; (ii) it derives a new method which can be used to separate genuine bad timing from hindsight bias in any context in which DW returns are used. This method is likely to find applications in many other fields, since the DW return is still commonly used (as the internal rate of return, IRR) in project finance and investment management.

The structure of this paper is as follows: Section I demonstrates the hindsight bias in DW returns which can be mistaken for the effect of bad investor timing. Section II sets out a method for decomposing the GM-DW return differential into the bias component and the genuine effect of investor timing. Subsequent sections apply this decomposition to data for NYSE/AMEX stocks (sections III and IV) and NASDAQ stocks (section V). Conclusions are drawn in the final section.

I. Identifying the Bias

The source of the hindsight bias can be illustrated with reference to the simple game illustrated below, in which the player faces a gain/loss of 10% in each of two rounds (Figure 1a). If we assume that the outturns in each round have a probability of exactly 50%, then the expected terminal wealth equals the initial stake.




Figure 1a




Figure 1b

The player may instead be able to quit the game following a win in the first round (Figure 1b). The expected terminal wealth for this truncated tree is still 1.0, but by quitting following the initial win the player can bias other performance measures.

For example, quitting early allows the player to claim a higher expected win rate. The full tree structure shows a 50% success rate. For example if we are trying to show “heads” in a fair coin toss then we have: HH, HT, TH, TT, giving success rates of 100%, 50%, 50%, 0%. But if the player quits if the first outcome is a head, then the tree shrinks to H, H, TH, TT, and the success rates shift to 100%, 100%, 50%, 0%, giving an impressive overall average of 62.5%. Quitting whilst ahead, and thus preserving a 100% winning record, biases this performance measure in a way which appears to suggest that the player has the ability to forecast the coin. The opposite incentive – to keep gambling when behind – can also be found. One simple example from outside the realm of finance is the child who agrees to toss a coin to settle an issue but, having lost, demands “best of three”.

Quitting whilst ahead also biases the expected Internal Rate of Return (IRR, the term which is generally used for the dollar-weighted return in investment management – the two terms

are synonymous). Table II shows that the full tree gives an average IRR of close to zero (fractionally negative due to the arithmetic/geometric mean inequality). This rises to 2.4% if the player quits after a win in the first round, since quitting locks in the early gains and gives the same IRR as if another win was guaranteed in the second round. This quit-whilst-ahead bias is similar to the familiar problem of the re-investment assumption used in calculating the yield to maturity on bonds.

Table II
IRRs of Illustrative Two-Round Game

The table shows the payoffs and associated internal rates of return of the two games shown in Figures 1a and 1b. An initial investment of 1 unit is assumed. The average IRR is the simple unweighted average of the IRRs calculated for the four scenarios.

Period	Lose-lose	Lose-win	Win-lose	Win-win	Avg.IRR
(a) Game Played over two periods					
0	-1	-1	-1	-1	
1	0	0	0	0	
2	0.81	0.99	0.99	1.21	
IRR	-10.0%	-0.5%	-0.5%	10.0%	-0.25%
(b) Player quits if ahead after round one					
0	-1	-1	-1	-1	
1	0	0	1.1	1.1	
2	0.81	0.99			
IRR	-10.0%	-0.5%	10.0%	10.0%	2.4%

Phalippou (2008) notes that private equity managers can in this way boost their recorded IRRs by altering the time horizon of their investments - returning cash to investors rapidly for successful projects and extending the life of poorly-performing projects. We show below that IRRs can also be biased up when the time horizon is fixed. Ingersoll et al. (2007) show that investment managers can manipulate conventional performance measures by reducing risk exposure following a good performance and increasing exposure after a poor performance. The

underlying strategy is to quit whilst ahead, but gamble more following poor outturns. They show that measures such as the Sharpe ratio and Jensen's alpha can be biased by this means, although they do not cover the IRR in their analysis.

Individual investors have no corresponding incentive to bias the IRR recorded for their own savings, but the typical pattern of investor cashflows tends to introduce this bias accidentally. To demonstrate this we need to examine the reasons for this bias more formally. Dichev derives net distributions from data for market returns and market capitalization using the clean surplus identity identified by Peasnell (1982). If in any period the market capitalization K_t is less than would have been suggested by applying the monthly rate of return r_t to the previous capitalization, then the differential must represent a distribution d_t :

$$d_t = K_{t-1}(1 + r_t) - K_t \quad (1)$$

If we regard the market capitalization K_t as the aggregate portfolio value across all investors, then when we discount at the internal rate of return (r_{dw}), the present value of future cashflows and the final liquidation value by definition sum to the value of the initial investment:

$$K_0 = \sum_{t=1}^T \frac{d_t}{(1 + r_{dw})^t} + \frac{K_T}{(1 + r_{dw})^T} \quad (2)$$

As set out in Dichev and Yu (2011), substituting equation 1 into equation 2 eliminates the distributions and shows that the IRR can be considered to be a dollar-weighted average of the individual monthly returns. Specifically, the relative weight that this DW return puts on the market return in any month (r_t) is determined by the NPV of the assets that the investor holds in this market at the start of this period (discounted at the DW return):

$$r_{dw} \sum_{t=1}^T \frac{K_{t-1}}{(1+r_{dw})^t} = \sum_{t=1}^T \left(\frac{K_{t-1}}{(1+r_{dw})^t} \times r_t \right) \quad (3)$$

This formula shows how distributions and injections of additional funds re-weight the monthly returns r_t . The resulting shifts in the DW return may represent either a genuine effect on expected investor wealth or a retrospective bias in the calculation. To illustrate these different effects we assume in Table III that we are calculating the overall DW return over an investment horizon of ten successive periods. We initially assume that there are no further cashflows after the initial investment, and that returns are IID. Our ex ante expectation would then be that the ten returns will be given equal weight (the first scenario shown in the table). We would expect the portfolio value to increase over time, but at a rate equal to the DW return, implying that the expected NPV of this portfolio would be equal in each period.

After the event we are likely to see some variation in these NPVs, due to volatility in r_t . But, to keep the illustration simple, we assume in Table III that this effect is small. If we invest a further amount, equal to the current portfolio value, after period 9, then the weight given to the period 10 return will be increased from 1/10 to 2/11. All earlier periods now have 1/11 weight, keeping the weights summing to unity. This is the second scenario illustrated in Table III.

Table III
Illustrative Effects of Net Distributions on Return Weights

The table shows the expected weights given in the DW return calculation to the returns in each period of a ten-period investment horizon, given the cash injections/distributions shown in the first column. Four different cashflow profiles are considered. For simplicity returns are assumed IID, with low volatility.

	1	2	3	4	5	6	7	8	9	10
(1) No injections/withdrawals	1/10	1/10	1/10	1/10	1/10	1/10	1/10	1/10	1/10	1/10
(2) Injection ($=K_t$) after period 9	1/11	1/11	1/11	1/11	1/11	1/11	1/11	1/11	1/11	2/11
(3) Injection ($=K_t$) after period 1	1/19	2/19	2/19	2/19	2/19	2/19	2/19	2/19	2/19	2/19
(4) Distrib'n ($=K_t/2$) after period 1	2/11	1/11	1/11	1/11	1/11	1/11	1/11	1/11	1/11	1/11

If instead we had made a corresponding injection after period 1, then the weight on subsequent periods would have been raised only to 2/19. An injection or withdrawal cannot have a substantial effect on the weights given to a large number of subsequent periods since this would raise the overall sum of the NPVs across all periods, with limited impact on the relative weights. But this injection has a substantial impact on the first period's weight, which falls from 1/10 to 1/19 (scenario 3). Indeed, if we had instead distributed half the portfolio after period 1, halving the value of the remaining portfolio, then we would expect the first period return to be given twice the weight of each subsequent return (scenario 4).

Thus injections/distributions can affect the weights given to previous returns just as much as the weights given to future returns. For example, comparing scenarios 2 and 4 shows that we can just as easily boost the expected weight given to r_1 (by distributing half the portfolio after period 1) as the weight given to r_{10} (by doubling the size of the portfolio after period 9). In addition, comparing scenarios 2 and 3 shows that the effect of a given distribution/injection

depends on its timing within the overall investment horizon – a result confirmed in Appendix A using simulated data.

We can re-arrange equation 3 further to show the deviation of periodic returns from the DW return. Periodic returns r_t will be either above or below r_{dw} , but the weighted sum of these differentials must be zero:

$$\sum_{t=1}^T \left(\frac{K_{t-1}}{(1+r_{dw})^t} (r_t - r_{dw}) \right) = 0 \quad (4)$$

This gives us a convenient form in which to consider the effect on the DW return of a distribution d (expressed as a percentage of portfolio value) at the end of period m :

$$\sum_{t=1}^m \left(\frac{K_{t-1}}{(1+r_{dw})^t} (r_t - r_{dw}) \right) + (1-d) \sum_{t=m+1}^T \left(\frac{K_{t-1}^*}{(1+r_{dw})^t} (r_t - r_{dw}) \right) = 0 \quad (5)$$

The distribution reduces the weight given to future returns in calculating the DW return, by reducing the future portfolio values to a fraction $(1-d)$ of what they otherwise would have been (K_t^*). A negative distribution (a further investment, for example as the result of a share issue) correspondingly increases future portfolio values. In the extreme, an investor could liquidate the entire portfolio ($d=1$). The DW return would then be calculated just on the returns up to period m , giving no weight to subsequent market returns.

Equation 5 shows that the two types of correlation that affect the GM-DW differential act in very different ways. A negative correlation between distributions and future returns would tend to boost the DW return, with negative distributions (injections) raising the start-of-period portfolio value ahead of periods of above-average returns, and positive distributions lowering it

ahead of weaker returns. This would represent good investor timing. Unfortunately this correlation is generally positive (see Table I), with investors tending to reduce their exposures ahead of periods of above-average returns and increase them ahead of poor returns.

The correlation of distributions with previous returns can also affect the DW return by retrospectively altering the relative weight given to earlier returns. This correlation tends to be negative (eg. with above-average¹ returns tending to be followed by injections of new funds). This will reduce the expected DW return by increasing the relative weight given to subsequent returns and correspondingly reducing the weight given to these earlier strong returns.

The arithmetic appears similar for the backward-looking and forward-looking correlations, but these effects are very different. The forward-looking correlation works by altering investors' portfolio size ahead of unusually strong/weak returns. Thus the change in the weight given to these returns in the DW return calculation corresponds to a change in investors' exposure to these returns. By contrast, the correlation of distributions with past returns does not affect the portfolio value until after the relevant returns have already taken place - the relative weight given to these returns in the DW return calculation is adjusted retrospectively. A forward-looking correlation between distributions and future returns represents good/bad timing, and clearly affects investor welfare. The backward-looking correlation does not.

There is also an important distinction to be made in the information content of these different effects. Ingersoll et al. (2007) state an important principle: That a manipulation-proof

¹ The correlation coefficient is calculated using the arithmetic mean rather than the DW, but these two measures will of course be highly correlated.

performance measure must not reward information-free trading. The correlation of distributions with future returns clearly depends on trades which have a high information content, since they forecast future returns. The correlation of distributions with prior returns instead affects the DW return by means of trades which have a very low information content - all that is required is that the investor is sometimes able to judge that returns to date have been unusually high or low compared with likely returns in future. This is a much easier task.²

An investment manager with no forecasting ability can boost his recorded DW return by making large distributions following a lucky period of strong returns, thus giving less weight to subsequent returns and correspondingly more weight to the returns already recorded. This is a form of the quit-whilst-ahead bias discussed above. Conversely a negative distribution could be used to increase the relative weight given to future returns after disappointing returns to date.

For illustration we can consider a situation in which periodic returns are drawn from a distribution with a fixed mean μ . An investor with a negative forward-looking correlation will tend to invest more (negative distribution) ahead of periods where $r_t > \mu$. This investor will achieve higher returns over time because her forecasting ability means that her ex ante conditional expectation (conditioned on these forecasts) is greater than μ .

An investor with no forecasting ability is still able to boost his expected DW return by retrospectively re-weighting the returns in previous periods. This can boost the expected DW

² It may be difficult to judge whether the return to date in any specific period differs significantly from the long-term mean, but investors can be opportunistic: if there is ever a period when the return to date has reached levels which are clearly different from any plausible estimate of the long-term mean then investors can use net distributions at this point to bias the DW return (for example, the cumulative return drops to well below zero in the early years of our NYSE/AMEX dataset – see Table V). By contrast, forecasting future returns is always likely to be difficult.

return, but not in any way which is likely to help meet his underlying investment objectives since his ex ante expected return in each period is still μ . The evidence in Table I suggests that distributions tend to show a significant negative correlation with past returns, biasing the DW return downwards.

Monte Carlo simulations confirm that this bias can spuriously affect the DW return. Friesen and Sapp (2007) show the results of simulations where returns are NIID. The ex ante expected return each period is identical, so any weighting of these ex ante returns must give the same average, regardless of the relative weights used. Volatility in ex post returns will drag the geometric mean below this arithmetic mean, but the simulations show that when distributions are correlated with previous returns the DW return is significantly lower than the geometric mean. By construction, there is no correlation between distributions and future returns, so this reduction must be due to ex post re-weighting of past returns. The simulations presented in Appendix A of the present paper confirm this result.

Hayley (2010) shows that the same hindsight bias is also responsible for the superior returns claimed for value averaging (a formula investment strategy which requires investors to make regular periodic investments to keep their portfolio growing at a pre-specified target rate). This strategy builds in a strong correlation of periodic investments with prior returns, since a smaller (larger) additional investment is required following strong (weak) market returns, thus giving relatively less (more) weight to future returns, which are likely to be lower (higher). This gives rise to an IRR which is greater than the geometric mean even in simulated random walk data where the ex ante expected return in each period is constant by construction.

The following section sets out a method for decomposing the observed GM-DW return differential into the effects of the backward-looking correlation (hindsight bias) and forward-looking correlation (bad timing) effects.

II. A Method For Decomposing The Effects Of Distributions

Equation 5 allows us to calculate the effect of a distribution in period m on the expected DW return, but it cannot in itself distinguish the effects of the forward-looking and backward-looking correlations in historic time series. The period weights sum to unity so, for example, increasing the weight given to above-average returns after period m would automatically reduce the weight given to earlier below-average returns. Thus we cannot directly distinguish between the two effects discussed above.

At first sight this seems to be an insoluble problem. However, a fund manager would require very limited information to use the quit-whilst-ahead strategy to deliberately bias the DW return - all that is required is an estimate of μ . Thus we should be able to identify retrospectively the expected impact of such a strategy on the DW return conditional on a similar assumption for future returns. Specifically, if we assume μ is constant, then we can evaluate the retrospective bias that each distribution has on the expected DW return. Repeating the process for each successive distribution will give us the total bias.

We start by assuming that the expected return in each period is equal to the geometric mean recorded for the whole investment horizon (this prevents the AM/GM differential from biasing our results) and that the distribution each month is zero. We will relax each of these

assumptions later. On these assumptions, the DW return over the investment horizon as a whole will initially be equal to the GM. We then substitute in the historical value for the return in the first period (r_1), recalculate the DW return for the entire series, and record the amount by which this is different from our initial DW return estimate. Next we substitute in the historical distribution for that period, recalculate the DW return again and note how much this has changed from our previous estimate (which was based on r_1 and assumed values for all other data). This sequence reflects the assumption in equation 1 that distributions are made at the end of each month, after the return for the month is known.

Substituting in the actual distribution data can be interpreted as replicating the process by which a cynical investment manager would bias the DW return. Each month, once the monthly return is known, he decides on the net distribution. Having no short-term forecasting ability, the manager assumes that all future returns will be equal to μ . If returns to date are significantly different from μ , then distributions/injections can immediately increase the expected DW return by increasing the weight given to previous good returns or reducing the weight given to previous bad returns.

These distribution decisions would have low information content in the sense that they are not predicated on a forecast of short-term asset returns. They require only an estimate of μ (and as we will see below, the results are not very sensitive to the accuracy of this estimate). The sum of the changes in the estimated DW return as a result of each of these net distributions thus gives us the cumulative effect of these low information trades.

If the distributions turn out to anticipate future returns, then this will be recorded when these subsequent returns are substituted into the calculation. The effect of successive return data on the estimated DW return is likely to be noisy, but if there is no relationship between these returns and earlier distributions then these effects will tend to cancel out over time. The cumulative effect on the estimated DW return will be positive only to the extent that previous net distributions resulted in relatively large start-of-period portfolio values (with correspondingly large NPVs) for periods when the returns were high, and relatively low NPVs ahead of periods when returns were low. This captures the effect of the good/bad timing of previous net distributions. As discussed above, this is a genuine economic effect rather than a bias, and will only come about if previous distributions contained information about future returns.

Our decomposition starts with the DW return equal to the GM, but by recalculating the DW return after each new piece of data is substituted in, we gradually move to the historic DW return. We consider separately (i) the aggregate effect on the DW return of including the distribution data (this gives the bias effect resulting from re-weighing past returns), and (ii) the aggregate effect on the DW return of the monthly return data (this gives the timing effects, reflecting any information in the distributions about future returns). These two components sum to give the total GM-DW return differential.

Table IV
Impact of Distributions on the DW Return (NYSE/AMEX)

This shows for early years of the NYSE/AMEX dataset the effect each year's return (timing effect) and distribution (bias) have on the expected DW return for the whole investment horizon (1926-2002). Future returns are assumed constant at 9.87% per annum. For clarity the table shows annual data, but underlying calculations use monthly data, so the precise effects depend on the timing of these distributions within these years.

	Annual return	Annualised return to date	Net distribution	Timing effect	Bias
1926	9.6%	9.6%	-3.1%	0.00%	0.00%
1927	33.3%	20.9%	-3.2%	0.25%	-0.01%
1928	39.0%	26.6%	0.1%	0.31%	-0.01%
1929	-14.6%	14.7%	-9.4%	-0.33%	-0.06%
1930	-28.8%	4.3%	-5.3%	-0.57%	0.00%
1931	-44.4%	-6.1%	6.6%	-0.89%	-0.04%
1932	-8.5%	-6.4%	7.4%	-0.24%	-0.07%
1933	9.9%	.	0%	.	.
1934	9.9%	.	0%	.	.
.
2002	9.9%	.	0%	.	.

Table IV illustrates this process, with historical data having been substituted over our starting assumptions up to 1932. The large negative distribution in 1929 had two effects. In aggregate investors added new cash equivalent to 9.4% of their existing portfolios. This subsequently turned out to be very bad timing: It increased portfolio values and so boosted the effect on the DW return of the subsequent negative returns. However, even before any further return data were included, the cash injection in 1929 had an immediate impact on the expected DW return by increasing the weight given to future returns (assumed equal to the 9.9% overall GM) and reducing the weight given to returns up to 1929, which were then well above this average. This re-weighting resulted in the immediate -0.06% bias effect shown for 1929.

Conversely, the large cash distributions in 1931 and 1932 reduced the weight given to future returns and boosted the weight given to returns to date, which by then were far below average. This gave rise to an immediate hindsight bias which again reduced the DW return.

This decomposition reflects the distinction between the forward-looking correlation of distributions with future returns and the backward-looking correlation with past returns. Future returns are assumed constant, so by construction the decomposed effects are calculated on the basis that past return and distribution data bear no relationship to future returns. But as we substitute in the actual distribution for each successive period, the impact on the estimated DW return reflects any relationship between this distribution and previous returns. Correspondingly, the effect on the estimated DW return as we substitute in each new piece of return data reflects any relationship with previous distributions. Thus substituting in the distribution data captures any effects of the backward-looking correlation with previous returns (the hindsight bias), whilst substituting in the return data captures the forward-looking correlation of previous distributions with the current r_t (the effect of good or bad investor timing).

III. Decomposing The Effects On the DW Return For NYSE/AMEX Stocks

We use the same dataset as Dichev (NYSE and AMEX stocks January 1926 to December 2002), and the same method to infer net distributions from the capitalization and return figures. We then step through the entire dataset adding first the monthly return, then the monthly distribution, calculating the DW return after each piece of data is added. The GM is 9.87% and the DW 8.61% (1926 to 2002). However, we find that the overall -1.26% per annum differential decomposes

into an annualized -0.21% from adding the return data and -0.95% from adding the distribution data (see Table V). This shows that the large majority of the GM-DW differential is due to hindsight bias, with only a limited effect from bad investor timing.³

Table V
Decomposition Of Timing And Bias Effects (NYSE/AMEX)

Cumulative effect on the DW return of substituting in (i) return data (timing effect), (ii) distributions (bias). By construction, for monthly data DW return = GM return + timing effect + bias, although this summation is only approximate for the annualized returns shown here.

	GM return	Timing effect	Bias	DW return
Jan. 1926 - Dec. 2002	9.87%	-0.21%	-0.95%	8.61%
Jan. 1926 - Dec. 2006	10.22%	-0.25%	-0.98%	8.88%

This method allows us to identify the impact of each new data point as we step through the data. Figure 2 presents the annualized return to date and the annual net distribution as a proportion of the implied portfolio value at the time. Figure 3 shows the corresponding incremental effects on the DW return resulting from adding successive data for returns (the investor timing effect) and distributions (the hindsight bias). The timing effect is noisy, but small in aggregate, whereas the hindsight bias is consistently negative in the early part of the period.

³ Some papers have also found bad investor timing using data on equity issues (e.g. Ritter, 1991, Ritter and Loughran, 1995), although others question these results (e.g. Brav and Gompers, 1997 and Schultz, 2003). The present paper does not revisit this well-established debate: It shows instead that whatever their statistical significance, the economic significance of these bad timing effects is far smaller than is suggested by the studies detailed in Table I.

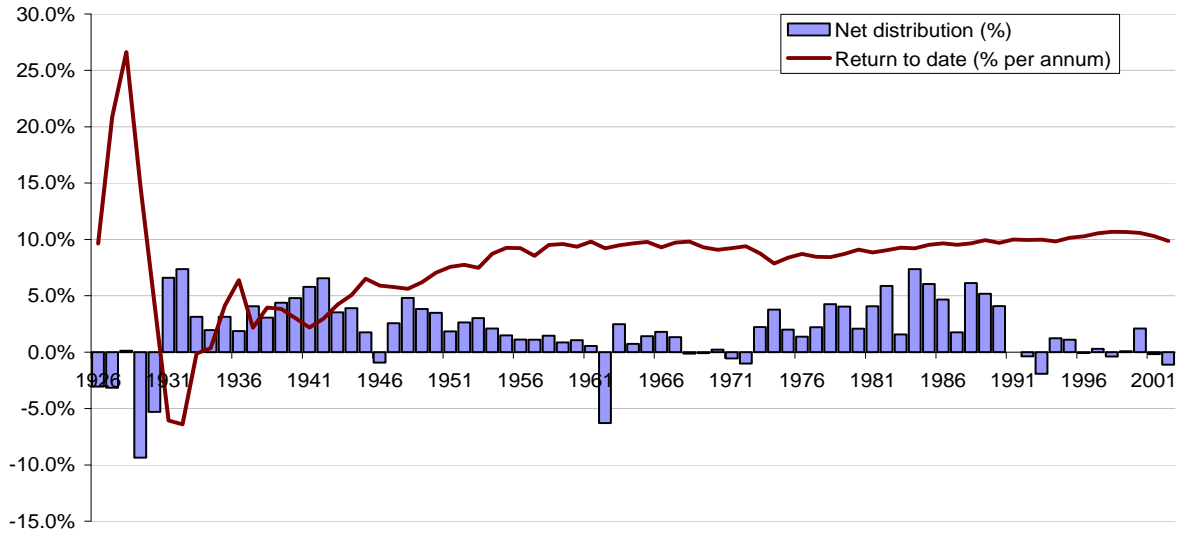


Figure 2: Returns to date and net distributions (NYSE/AMEX stocks). The line shows the annualized cumulative return to date (%) from the start of the dataset in January 1926. The bars show net distributions as a percentage of the implied total market capitalization before the distribution ($K_{t-1}(1+r_t)$). A positive distribution is a return of cash to investors, a negative distribution is a net investment.




Figure 3: Timing and bias effects in DW returns (NYSE/AMEX stocks). This shows the change in the expected DW return for the whole period (January 1926 - December 2002) resulting from substituting in (i) the monthly market return ('timing effect', the volatile bold line), and then (ii) the monthly aggregate net distribution ('bias', the thinner, more stable line). The DW return is calculated on the initial assumption that future returns are equal to the geometric mean (9.87% per annum) and that future distributions are zero. The underlying calculations use monthly data, but annual effects are shown here for clarity.

It is also reassuring that Figure 3 shows no massive monthly jumps. This suggests that the DW return calculation has found the same underlying root to the IRR polynomial in every calculation, with only modest changes due to the new data. If it ever jumped between multiple solutions then we might expect to see a far larger monthly move.

The large cash distributions in 1931 and 1932 reduced the weight given to future returns and boosted the weight given to returns to date, which by then were far below average. As we saw above, this gave rise to an immediate hindsight bias which again reduced the DW return. The incremental bias effect remained negative in the late 1930s and early 1940s as consistently large distributions were made whilst the return to date was below 5%. Distributions in later years caused relatively little bias since by this stage the return to date had inevitably converged towards the overall average. Consistent with this, extending the dataset makes little difference, with the return and distribution effects shifting only from -0.21% and -0.95% respectively (1926 to 2002) to -0.25% and -0.98% (1926 to 2006).

The method derived here for decomposing the GM – DW return differential clearly has applications in other fields where DW returns are used, notably project finance and investment management. Good timing by investment manager should clearly be separated from the effect of hindsight bias (whether deliberate or accidental). For this purpose the method above should be used to calculate a hindsight-corrected DW return, which can be derived by adding the timing effect to the GM return or, equivalently, subtracting the hindsight bias from the DW return:

$$R_H (\text{hindsight-corrected DW return}) = GM \text{ return} + \text{timing effect} = DW \text{ return} - \text{bias} \quad (6)$$

IV. Sensitivity Analysis

In this section we examine the robustness of our results as we relax the initial assumptions made above: (i) that future returns are equal to the GM of 9.87% per annum; (ii) that future distributions are zero. Our assumption that distributions are made at the end of each month makes minimal difference: when we assume instead that they are made at the start of each month, the decomposed effects differ by less than 0.01%.

Table VI sets out the timing and bias effects derived using a wide range of assumptions for future returns. The timing effect calculated on such counterfactual assumptions is relatively uninformative: Assuming returns which are well below the historical mean naturally leads to a more positive timing effect as returns subsequently tend to be higher than this (high assumed returns lead to correspondingly negative return surprises). Our interest is instead in the bias effect. The table covers a huge range of assumed average returns, but our key finding is robust, since for any plausible figure in the middle part of this range the bias effect is clearly substantial and negative, and accounts for a large part of the -1.3% historical GM-DW differential.

Table VI
Decomposition on Alternative Return Assumptions

The table shows the cumulative effect of monthly returns (timing effect) and distributions (bias) on the expected DW return for NYSE/AMEX stocks 1926-2002. Future returns are initially assumed constant at the levels shown in the first column. By construction, for monthly returns the DW return = assumed GM + timing effect + bias, although this summation is only approximate for the annualized returns shown here.

Assumed GM	Timing effect	Bias	DW return
5%	3.49%	-0.04%	8.61%
6%	2.76%	-0.28%	8.61%
7%	2.01%	-0.49%	8.61%
8%	1.24%	-0.67%	8.61%
9%	0.47%	-0.83%	8.61%
10%	-0.31%	-0.97%	8.61%
11%	-1.10%	-1.09%	8.61%
12%	-1.89%	-1.19%	8.61%
13%	-2.67%	-1.28%	8.61%
14%	-3.46%	-1.37%	8.61%
15%	-4.24%	-1.44%	8.61%

Moreover, we do not need to interpret these assumptions as reflecting investor expectations. Considering how a cynical investment manager could attempt to bias the DW return helped lead us to the decomposition set out above, but this should be seen as just an analogy. As discussed in section II, the key requirement is that our assumed future returns are constant, with no relationship with past returns or distributions. This ensures that all forward-looking correlation of distributions with future returns is captured in the “timing effect” column as the subsequent return data is substituted in, whilst all the backward-looking correlation between distributions and previous returns is captured in the “bias” column. This holds regardless of whether the return assumption actually reflects investor expectations. The key advantage of setting the assumed future returns equal to the historical geometric mean is that this removes the

effect of consistent return surprises in either direction, leaving only the pure effect of the timing of investment flows compared to periods of above/below average return.

We also investigate the effect of changing our initial assumption for future distributions (set to zero for all periods in the decompositions above). Setting each instead to 0.082% of market capitalization (giving an average distribution equal to that in the historic sample) substantially alters the decomposition, with the aggregate bias effect increasing to -1.15% and the timing effect almost vanishing (-0.04%). The reason can be seen in Figure 4, which shows that returns on NYSE/AMEX stocks (cumulated over 10 year periods to reduce short-term noise) trended upwards over our sample period of January 1926 to December 2002. Given this trend, any early distribution would appear to be bad timing compared to our initial assumption that future distributions were zero.




Figure 4: Long-term equity returns. The chart shows the annualized 10 year geometric mean market return starting in the period shown. A linear time trend has been added, fitted to NYSE/AMEX returns Jan. 1926-Dec. 2002. The longer time series shows the S&P index return (source: Shiller). Data after Dec. 2002 is shown as dotted lines.

However we should not accept the effects of this trend at face value. First, bad investment timing is generally interpreted as a short-term cyclical effect as investors chase returns during booms. Apparent bad timing caused by this very long term trend is a very different effect. As most investors have horizons which are substantially shorter than this 77 year time series, they had no realistic option to time their investments better. Moreover it is entirely implausible to suppose that this upward trend will continue in future – this would imply that expected equity returns are currently over 15% per annum and will continue to rise by almost 1% per decade.

Figure 4 also shows the similar S&P index returns going back to 1871 and forward to 2010. This shows that the 1926 to 2002 period was almost unique in showing such a sustained uptrend. Almost any other period of equal length would have given us very different results. Thus if we are to obtain results that can be plausibly applied to the future (e.g. in estimating the expected risk premium) we need to strip out the effects of this trend before decomposing the residual into the timing and bias effects. This can be achieved either by de-trending the return series or de-meaning the distribution. For robustness we do both, on a variety of different assumptions, both individually and combined.

Table VII presents the results using four alternative treatments of the distribution data. We have already seen the first two, which use unadjusted historical distribution data. The third and fourth variants de-mean the distribution data by subtracting a percentage of portfolio value such that (a) the average distribution is zero (b) the average percentage of capitalisation which is distributed is zero. The last four variants repeat the first four, but with return data from which a log-linear trend has been extracted (thus keeping the GM return unchanged).

Table VII
Decomposition With Alternative Corrections For Trend In Returns (% pa.)

The table shows the cumulative effect of monthly return and distribution data on the expected DW return for NYSE/AMEX stocks 1926-2002 using a range of measures to correct for the long-term uptrend in returns and the positive mean distribution. The first two use raw distribution data with an assumption that future distributions are (1) zero, (2) set to the sample average as a percentage of implied market capitalization. These decompositions are repeated for distribution data which has been de-meanned by subtracting a percentage of market capitalization such that (3) the average distribution is zero, (4) the average percentage of market capitalization which is distributed is zero. The last four variants repeat the first four, but with return data from which a log-linear trend has been removed.

Return Data	Distribution data and starting assumption	Timing effect	Bias	DW return
1. Unadjusted r_t	Unadjusted d_t (future d_t initially set at zero)	-0.21%	-0.95%	8.61%
2. Unadjusted r_t	Unadjusted d_t (future d_t (%) initially set at sample avg.)	-0.04%	-1.15%	8.61%
3. Unadjusted r_t	De-meanned d_t (future d_t^* initially set at zero)	-0.12%	-0.53%	9.17%
4. Unadjusted r_t	De-meanned d_t (future d_t^* (%) initially set at zero)	-0.06%	-0.29%	9.50%
5. Detrended r_t	Unadjusted d_t (future d_t initially set at zero)	-0.03%	-0.49%	9.31%
6. Detrended r_t	Unadjusted d_t (future d_t (%) initially set at sample avg.)	0.04%	-0.58%	9.31%
7. Detrended r_t	De-meanned d_t (future d_t^* initially set at zero)	-0.06%	-0.26%	9.51%
8. Detrended r_t	De-meanned d_t (future d_t^* (%) initially set at zero)	-0.06%	-0.36%	9.43%

The first two variants use unadjusted historical data, thus ending up with the historical DW return of 8.61%. By contrast the other variants (3 to 8) adjust the historic data to remove the effect of the long-term trend in returns. These new variants all give substantially higher final DW returns but, reassuringly, these lie within a limited range 9.17% to 9.51%. Thus all these methods suggest that this long-term uptrend accounted for a substantial part of the raw GM-DW differential, and only around half resulted from shorter-term effects. Thus even before decomposing the residual into bias and timing effects it is clear that these two together have much less effect once we strip out the long-term uptrend in returns. Moreover, decomposing the remaining differential shows that the timing effect is very small in all cases, ranging from -0.12% to +0.04%.

Thus after adjusting for (i) the unsustainable uptrend in returns, and (ii) hindsight bias in the DW return, we find that bad timing actually had only a very small impact on the return received by investors. Thus, in contrast to the claims made elsewhere, bad investor timing does not justify reducing our estimates of the equity risk premium.

V. Decomposing The Return Differential For NASDAQ Stocks

NASDAQ stocks show a much larger differential than NYSE/AMEX stocks, with a GM of 9.6%, but a DW return of only 4.2% (January 1973 to December 2002). When we decompose this differential using the method set out above, we find that the large majority (-4.0%) is due to bad investor timing, with only -1.0% due to hindsight bias (see Table VIII).

Table VIII
Decomposition of Timing and Bias Effects (NASDAQ)

Cumulative effect of monthly returns (timing effect) and distributions (bias) on the expected DW return for NASDAQ stocks. By construction, for monthly returns DW return = GM return + timing effect + bias, although this summation is only approximate for the annualized returns shown here.

	GM return	Timing effect	Bias	DW return
Jan. 1973 - Dec. 2002	9.6%	-4.0%	-1.0%	4.2%
Jan. 1973 - Dec. 2006	10.4%	-1.8%	-0.9%	7.5%

The main effect comes from investors' terrible timing during the dotcom boom. Additional funds equivalent to 8.6% of market capitalization were invested in 1999 and 13.7% in 2000, just ahead of the crash (see Figure 5).

Naturally, the recorded effect of bad timing rises as we increase our assumption for future returns, but it is reassuring that the degree of hindsight bias remains fairly stable (see Table IX). Moreover, the decomposition shows very little sensitivity to the assumed level of future distributions (consistent with the absence of any long-term trend in returns). Thus our estimates of the bias are robust to shifts in both these assumptions.

Table IX
NASDAQ Return Decomposition: Alternative Return Assumptions

The table shows how cumulative timing and bias effects vary as we alter our assumption for future returns. Returns are initially assumed constant at the levels shown in the first column before historical returns are substituted in. For monthly returns DW return = assumed GM return + timing effect + bias, but this summation is only approximate for the annualized returns shown. Coverage: NASDAQ stocks 1973-2002.

Assumed GM	Timing effect	Bias	DW return
5%	0.59%	-1.31%	4.25%
6%	-0.44%	-1.23%	4.25%
7%	-1.45%	-1.15%	4.25%
8%	-2.44%	-1.08%	4.25%
9%	-3.40%	-1.02%	4.25%
10%	-4.35%	-0.96%	4.25%
11%	-5.28%	-0.90%	4.25%
12%	-6.19%	-0.85%	4.25%
13%	-7.08%	-0.80%	4.25%
14%	-7.95%	-0.75%	4.25%
15%	-8.80%	-0.71%	4.25%

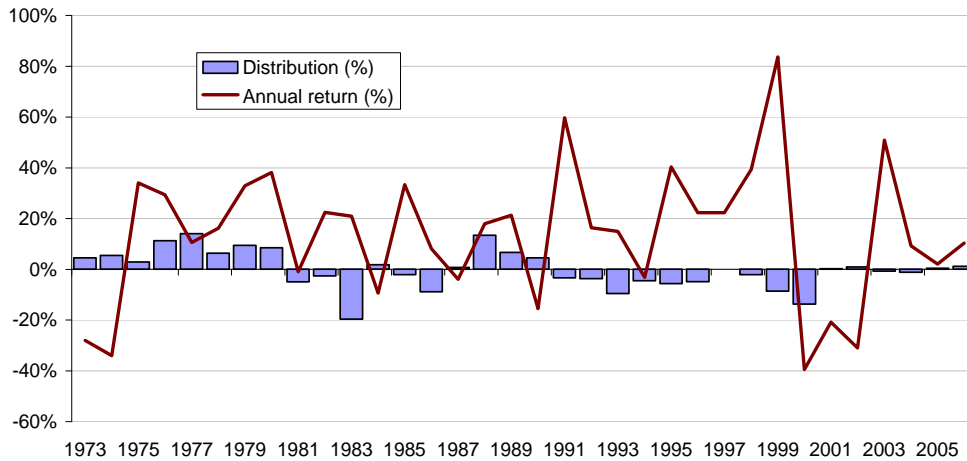


Figure 5: Distributions and annual returns (NASDAQ stocks). The line shows the annual return on NASDAQ stocks (Jan. 1973 – Dec. 2002). The bars show annual net distributions as a percentage of the market capitalization ahead of the distribution ($K_{t-1}(1+r_t)$). A positive distribution is a return of cash to investors, a negative distribution a net investment.

However, as Keswani and Stolin (2008) showed, when the dataset is extended to 2006 the GM-DW return differential shrinks markedly. The decomposition shows that the cumulative timing effect changes from -4.0% to only -1.8%. One reason for this is the positive returns seen after 2002, but the timing effect would have diminished even if the additional data were unexceptional. As we saw in Section I (and is confirmed by the simulations in Appendix A), for any given relationship between distributions and subsequent returns, timing effects are far more powerful for distributions close to the end of the investment horizon, since they then have a large effect on the weights given to subsequent returns in the DW return calculation. The same pattern of distributions and subsequent returns would tend to have much less impact further from the end of the horizon, since the distributions would then affect start-of-month portfolio values over a

larger number of subsequent periods, implying less impact on their relative weights in the DW return calculation.⁴

The historical data 1973 to 2006 still shows a -1.8% effect from bad investor timing, but we face two problems in assuming that bad timing will continue to have such an effect in future. First, this effect stems from what should be seen as a single massive event – the dotcom crash – so we must question whether this is statistically significant. Second, we must expect the measured bad timing effect to shrink further as more data is added, pushing the 2000 to 2002 period further away from the end of the investment horizon.

VI. Conclusion

A growing number of papers use the difference between the dollar-weighted return and the geometric mean return as a measure of the effect of bad investment timing. They generally find that poor timing has reduced the return actually received by investors to well below the buy-and-hold return on the assets concerned. As a result they conclude that our estimates of the equity risk premium need to be revised down substantially. Given the central role that this figure plays in finance, this would have profound implications.

However, the present paper finds that the DW return is affected by the correlation of net investor cashflows with both future asset returns and previous asset returns. The first effect

⁴ Using artificial data for these extra years (returns set equal to the GM to date (9.62%) and distributions set to zero) reduces the aggregate timing effect to only -2.5% (from -4.0% for 1973-2002). This confirms that the majority of the reduction is inherent in the extension of the dataset, rather than being due to the specific data added.

clearly affects investors' expected wealth, but the second merely gives rise to a hindsight bias in the DW return.

This paper derives a method which allows us to separate these two effects. This shows that bad investment timing accounts for very little of the overall differential between the GM and DW returns for mainstream US equities (those traded on the NYSE and AMEX exchanges). The great majority is just hindsight bias. Thus low DW returns should not lead us to adopt correspondingly low figures for the equity risk premium.

There are likely to be applications in other fields. Wherever DW returns (IRRs) are quoted this technique can be used to separate hindsight bias from any genuine timing effect. This is likely to be useful in project finance and investment management, where IRRs are still routinely used as summary performance measures. The decomposition described above can be used to derive hindsight-corrected IRRs.

More specifically, future research could investigate the degree to which funds have benefited from hindsight bias. This could have come about if funds tend to choose between alternative cashflow options by comparing the projected IRRs, or if funds which benefit from this bias by luck tend to have higher survival rates. As we saw above, the effects of this bias can be substantial even for broad equity market indices. They could be much larger for individual funds which are likely to have considerably greater volatility in both their returns and their cashflows.

REFERENCES

- Brav, Alon and Paul A. Gompers.** 1997, Myth or reality? The long-run underperformance of initial public offerings: Evidence from venture and non-venture capital-backed companies, *Journal of Finance* 52: 1791-1821.
- Clare, Andrew and Nick Motson.** 2010. "Do UK investors buy at the top and sell at the bottom?" *Cass Business School Working Paper*.
- Dichev, Ilia D.** 2007. "What are Stock Investors' Actual Historical Returns? Evidence from Dollar-Weighted Returns." *American Economic Review*, 97(1): 386-401.
- Dichev, Ilia D and Gwen Yu.** 2011. "Higher Risk, Lower Returns: What Hedge Fund Investors Really Earn." *Journal of Financial Economics* 100(2): 248-263.
- Friesen, Geoffrey C. and Travis R.A. Sapp.** 2007. "Mutual Fund Flows and Investor Returns: An Empirical Examination of Fund Investor Timing Ability." *Journal of Banking and Finance* 31: 2796-2816.
- Gompers, Paul A, and Josh Lerner.** 2003, The really long-run performance of initial public offerings: The pre-NASDAQ evidence. *Journal of Finance* 58(4): 1355–1392.
- Hayley, Simon.** 2010. "Value Averaging and the Automated Bias of Performance Measures." *Cass Business School working paper*. Available at SSRN: <http://ssrn.com/abstract=1606347>
- Ingersoll, Jonathan, Matthew Spiegel, William Goetzmann and Ivo Welch.** 2007. "Portfolio Performance Manipulation and Manipulation-proof Performance Measures." *Review of Financial Studies* 20(5):1503-1546.
- Keswani, Aneel and David Stolin.** 2008. "Dollar-Weighted Returns to Stock Investors: A New Look at the Evidence." *Finance Research Letters*, 5(4).
- Loughran, Tim, and Jay R. Ritter.** 1995. "The New Issues Puzzle." *Journal of Finance*, 50(1): 23-51.
- Peasnell, Ken V.** 1982. "Some Formal Connections Between Economic Values And Yields And Accounting Numbers." *Journal of Business Finance & Accounting* 9(3): 361–381.
- Phalippou, Ludovic.** 2008. "The Hazards of Using IRR to Measure Performance: The Case of Private Equity". *Journal of Performance Measurement*.
- Schultz, Paul.** 2003. "Pseudo Market Timing and the Long-Run Underperformance of IPOs". *Journal of Finance*, 63(2): 483-517.
- Shiller, Robert J.** 2005. *Irrational Exuberance*. Princeton University Press, 2nd edition, and www.irrationalexuberance.com.
- Zweig, Jason.** "Funds that really make money for their investors." *Money Magazine*, 1st April 1997.

Appendix: Simulation Evidence

This section uses Monte-Carlo simulations to show that the correlation of distributions with prior returns can shift the average DW return away from the geometric mean even when the ex ante expected return in each period is identical by construction. We also confirm that the size of this bias depends on the timing of the cashflows within the investment horizon.

Returns are generated over an investment horizon comprising ten periods. These returns are NIID, and for convenience the mean is set so that the GM return over the ten periods averages zero. We then consider the impact that a single net distribution after each of periods 1-9 has on the DW return (all assets are assumed to be liquidated in the tenth period). Net distributions are either (a) negatively correlated with the previous return (as investors chase returns by investing more following strong returns), or (b) positively correlated with the return in the following period (bad investor timing). Note that these are the signs of the correlations generally found in the empirical studies shown in Table I.

Figure 6 shows that each of these correlations pulls the average DW return below the GM return (the opposite correlations – not shown – have a positive effect). The forward-looking correlation reduces the conditional expected return in the period following the distribution (conditioned on the amount which remains invested). By contrast, where the distributions are only correlated with past returns the ex ante expected return in each period is identical by

construction, so we should regard the shift in the DW return as a hindsight bias produced by the retrospective shifts in the weights given to past returns.

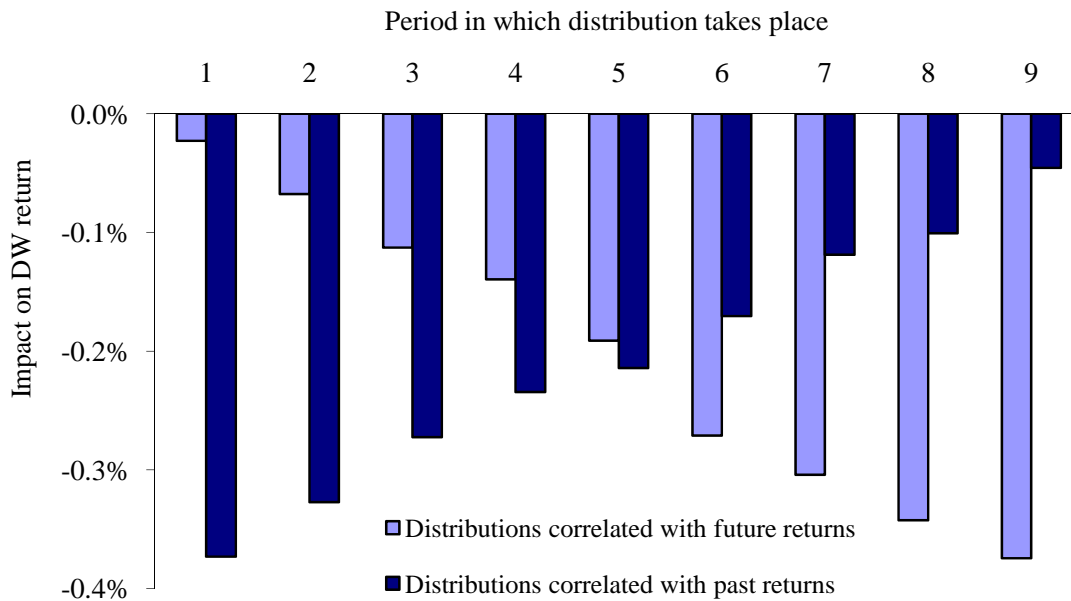


Figure 6: Impact on DW returns of correlation between distributions and returns. Each simulated path is of ten periods, with returns in each period NIID with standard deviation 20% and geometric mean zero. A single distribution is included for each path in the period shown. This distribution is set as a percentage of portfolio value which is (i) the previous period's percentage return multiplied by -1, or (ii) identical to the following period's percentage return. Net distributions in all other periods are zero for each path. The impact of these cashflows in pulling the DW return below the GM return is calculated for each of 5000 simulated return paths for each of the distribution patterns shown. The GM is, of course, unaffected by these cashflows, so the differential is interpreted as the negative impact of these cashflows on the DW return.

The size of these effects depends on when each distribution comes within the investment horizon, but the average size of the effect across all periods is roughly the same for the forward-looking and backward-looking correlations. This confirms the underlying symmetry apparent in

equation 5: That a distribution can in principle affect the DW return just as effectively by re-weighting either past returns or future returns.

A distribution which is correlated with returns in the coming period is most effective in period 9, since it can then have a substantial effect on the NPV of the portfolio value at the start of period 10, and hence on the relative weight given to this return in the DW return calculation. By contrast, a similar distribution after period 1 alters the portfolio value in all future periods, and so has little effect on their relative weights. But such early distributions strongly affect the relative weight given to previous returns. Thus a forward-looking correlation has more impact near the end of the investment horizon and the backward-looking correlation has more impact near the beginning.

Decomposing the effects on the DW return for NYSE/AMEX stocks (Section III) shows that the major impact comes from a backward-looking correlation of distributions early in the investment period with prior returns. The same pattern of distributions and returns would have had less impact on the DW return if our sample had started earlier.

For NASDAQ stocks we found instead that the major impact is the bad timing of the large net investments made at the height of the dotcom boom – ahead of the subsequent bust. But, again, the fact that these flows took place very near the end of the investment horizon gives them the maximum effect on the DW return. We found that the effect is reduced as more recent data is added, pushing these large distributions away from the end of the horizon, and we should expect further reductions as subsequent data is added.