THE TIME DIVERSIFICATION PUZZLE: *why trustees should care*

ROBERT J BIANCHI F Fin, Senior Lecturer, Griffith Business School, Griffith University
MICHAEL E DREW SF Fin, Professor of Finance, Griffith Business School, Griffith University
ADAM N WALK F Fin, Postdoctoral Research Fellow, Griffith Business School, Griffith University[1]

For 50 years, the time diversification debate has sought to understand the essential relationship between risk and investment horizon with little resolution. The answer seems to depend in part on how one views risk. This paper seeks to show that while the time diversification puzzle remains unsolved, the debate itself provides timely food for thought for trustees in setting fund investment policy and for designing defaults, in particular.

Time diversification — the notion that extending the investment horizon reduces risk — has been one of the most hotly contested ideas in finance since its first formal treatment by Samuelson (1969). Since then, nearly 100 scholarly papers have been produced on the subject, all to no avail. That the relationship between risk and investment horizon remains unresolved surely confirms the status of time diversification as one of finance's most enduring puzzles.

Conventional wisdom suggests that investment risk decreases as the time horizon increases. A large number of empirical studies support this idea, finding that the standard deviation of annualised returns falls over time. Entire books have been dedicated to communicating to a popular audience this idea that risk is tamed by time: for example, Siegel's (1994) *Stocks for the long run: A guide to selecting markets for long-term growth.* The widespread acceptance of the supposed inverse relationship between risk and time has led many in academe and industry to suggest that time diversification is more than just conventional wisdom, and has instead graduated to become a 'stylised fact' of modern finance.

So if this inverse relationship between risk and time horizon is so far beyond doubt then where is the puzzle? While many studies — including this one — confirm that the standard deviation of annualised returns decreases over time, studies also find that the standard deviation of cumulative returns does not diminish over time. In fact, if we frame risk in these terms, we find that dispersion actually increases over time. Figure 1, for example, shows these first two conceptions of risk plotted against

investment horizon. Samuelson (1969) began the debate by asserting that there is no relationship between risk and investment horizon, arguing that risk is constant with investment horizon being a function of risk preferences.

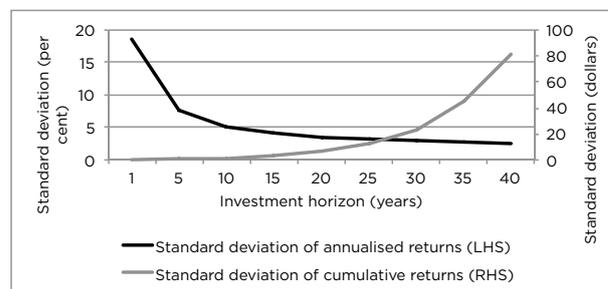**FIGURE 1: Contradictory evidence**

Figure 1 plots the standard deviation of annualised returns, in percentage terms, on the left vertical axis and the standard deviation of cumulative returns, in dollar terms, on the right vertical axis. Both series are plotted against investment horizon.[2]

This contradiction strikes at the heart of the time diversification puzzle. How is it that two versions of the one statistical measure of dispersion — standard deviation — give us opposite answers regarding the relationship between risk and time? Are risk and time negatively related — the received position — or positively related? Or is Samuelson (1969) right? What causes these seemingly contradictory outcomes? This paper is motivated by this contradiction at the heart of the time diversification puzzle, a contradiction for which the academy has no satisfactory answer as yet.

The time diversification literature

After Markowitz's (1952) groundbreaking work on portfolio choice in a one-period setting, scholars began to consider the portfolio selection problem in a multi-period setting like that encountered with practical investment problems. Chief among these scholars was another Nobel Prize winner Paul A Samuelson who considered the problem in a multi-period setting using expected utility theory. Samuelson's (1969) work is of particular interest to us for two reasons. First, he was among the first to bring the genius of Markowitz's (1952) work into a multi-period setting which, by itself, is remarkable.[3] Second, and particularly germane to this paper, Samuelson (1969) initiates the time diversification debate by considering whether the concept of diversification works with time, in the same way as it does among assets or securities (cf. Markowitz, 1952). In order to study the existence of time diversification, Samuelson (1969) selects the classical expected utility theory as his framework of choice. Expected utility theory is thus the point of departure for this debate, and all other competing streams or schools of thought tend to emerge at least in part as a reaction to Samuelson's (1969) work.

Samuelson (1969) isolates the relationship between risk and time by observing the optimal allocation to risk assets with horizon, based on three assumptions. While a number of proponents confirm the mathematical certainty of his findings, even more scholars — including some who are otherwise advocates of expected utility theory — call into question Samuelson's (1969) assumptions. In fact, it is Samuelson's (1969) three assumptions that provide later scholars with oxygen to keep the time diversification debate burning. While a comprehensive review of the literature identifies four schools of thought, this paper will only consider Samuelson's (1969) original framework — because of its three assumptions — and what was described by Booth (2004) as the 'applied' stream.[4]

The 'applied' stream in the time diversification debate is defined more by what it's not, than what it is. While the applied stream is a somewhat nebulous confection of studies, there is the faint semblance of a unifying theme. Scholars who pursue this path tend to approach the problem of time diversification empirically, and without resting on a theoretical edifice in the way that Samuelson (1969) does. Simulation techniques are also a common methodological choice as Booth (2004) suggests. Parallel to the time diversification debate, a rich literature on risk measures has emerged. Leaning on this literature, applied scholars tend to define risk in a certain way — for example, value at risk — and then proceed to estimate their selected risk measure over a number of horizons of different lengths. Scholars then draw conclusions about the presence or otherwise of time diversification by applying reasoning to these estimates. Naturally, it is possible to define risk in many ways and so the applied stream has tended to grow as new conceptions of risk emerge. Some scholars have even developed measures purely for the purposes of analysing the time diversification question.

Through time, in the time diversification literature we have seen a quest for the measure of risk that properly isolates the relationship between risk and time horizon. At the turn of the century, Kritzman (2000, p. 50) remarked wistfully that 'for many the time diversification debate has degenerated into a referendum on the meaning of risk'. We agree that the debate has become, and remains, a referendum on risk and that, in Kritzman's (2000, p. 50) words, such a referendum is to some extent '... futile'.

On the other hand, is a focus on risk necessarily a bad thing? We argue that such a focus on risk is desirable provided trustees resolve their attitude to risk, set investment policy with this frame in mind, and measure and monitor performance in a way consistent with the risk frame. Before we consider the most appropriate risk frame we briefly outline why this debate should be of interest to trustees.

The relevance of the time diversification debate for trustees

While a review of the literature shows that the puzzle remains largely unresolved, there are several conclusions that have emerged from the debate that should interest trustees. First, as we have shown, much of the time diversification debate is about risk, and how it is framed. In fact, a review of the time diversification literature shows that the relationship between risk and investment horizon depends on one's view of risk.

It is thus important that trustees resolve how they conceive risk before turning their minds to other important aspects of their role like setting investment policy and fulfilling the investment governance function. For example, if one views risk as the standard deviation of returns — as much of the industry appears to — this might lead to a different approach to investment policy than for a trustee that sees terminal wealth adequacy as the objective.

How one views risk also has implications for investment governance. If trustees see adequacy as the key objective of their fund, then performance expressed in terminal wealth terms is more informative than returns-only measures. According to this model, trustees would also need to consider other determinants of terminal wealth — like contribution rates — in addition to returns. In this

scenario, while wealth becomes the central focus, returns remain an important measure of the success for the underlying investment program. Governance may therefore need to become multi-dimensional.

The real debate

With very few exceptions, the entire time diversification debate is conducted in a returns-only framework. Risk is thus seen through the narrow lens of returns. Contributions and other factors (e.g. salary growth) are almost completely overlooked. Trustees of Australian superannuation funds — whose members contribute at a minimum rate of 9 per cent per annum — would identify this as a significant deficiency. As trustees well know, pension finance is about wealth (the outcome), not only about returns as some literature would have us think.

Without considering realistic accumulation models, the time diversification debate, despite its understandable focus on the relationship between risk and investment horizon, largely ignores recent pension finance research. Basu and Drew (2009), for example, highlight the so-called 'portfolio size effect' which sees a rapid rise in portfolio size as retirement approaches, due to the combined effects of returns, contributions and salary growth. This portfolio size effect magnifies the potential effects of sequencing risk (Macqueen and Milevsky 2009; Basu, Doran and Drew 2012, Doran, Drew and Walk 2012): the risk of experiencing an inopportune sequence of returns.

Therefore, the dynamics of superannuation investing means that a minus 25 per cent return, for example, has different wealth impacts depending on the timing of the return. For example, the impact of the global financial crisis (GFC) — an example of sequencing risk realised — on those in their late 50s/early 60s has been devastating in wealth terms. In the next section we will provide a simple example of the dynamics which we argue the time diversification debate ignores; dynamics which are of great interest to trustees.

A practical example

In order to bring to life these dynamics we consider the experiences of two hypothetical investors:

> **Late 30s** — considers the experience of an individual who begins their retirement saving in their mid-20s with an account balance of zero dollars, and contributes at a rate of 9 per cent per annum for 13.5 years until their late 30s. Over the period their salary grows at a rate of 4 per cent per annum from $40,000 to $68,579.

> **Near retirement** — considers the experience of an individual who continues saving from their early 50s with an account balance of $100,000, and contributes at a rate of 9 per cent per annum

for 13.5 years until near retirement age. Over the period their salary grows at a rate of 4 per cent per annum from $60,000 to $102,868.

The only differences between these two examples are their starting salary levels and their initial wealth. Comparing these hypothetical investors allows us to consider the differential impacts of identical returns at different stages of the investing lifecycle. To draw out the importance of the sequence of returns we look at three accumulation paths derived from the one set of synthetic balanced fund returns:

> **Actual path** — the actual path uses balanced fund returns as they occurred over the period January 1999 to June 2012.

> **Reverse path** — the reverse path uses the actual returns but in reverse order. The GFC would therefore have occurred early in the accumulation phase of the hypothetical investor in question.

> **Average path** — the average return path uses the periodic arithmetic mean for the return series for each period. In this sense, it is as if the investor earned the average return for each period.

We plot these three paths for each of the hypothetical members in Figures 2 (Late 30s) and 3 (Near retirement).
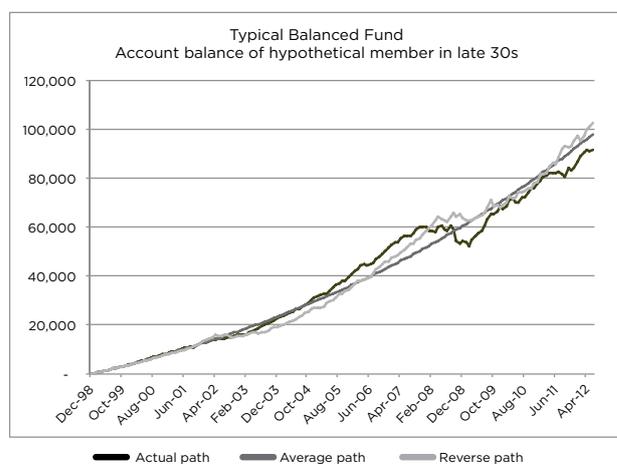
FIGURE 2: Late 30s

Typical Balanced Fund
Account balance of hypothetical member in late 30s

Actual path   Average path   Reverse path

Figure 2 plots three separate accumulation paths for a person who begins their retirement saving at age 25 with a salary of $40,000 and an account balance of zero dollars. Over the 13.5 year accumulation period the salary grows at a rate of 4 per cent per annum such that their income in their late 30s is $68,579. Our hypothetical member contributes at the rate of 9 per cent per annum. The three accumulation paths are as follows: 'Actual path' uses synthetic balanced fund returns as they occurred over the period January 1999 to Jun 2012; 'Reverse path' uses the actual returns in reverse order; and 'Average path' uses the arithmetic mean for the return series for each period.

FIGURE 3: Near retirement

Typical Balanced Fund
Account balance of hypothetical member near retirement age

Legend: Actual path — Average path — Reverse path
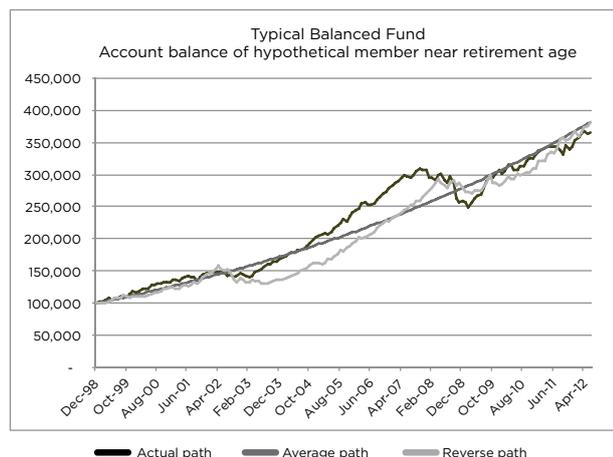
This figure plots three separate accumulation paths for a person who in their early 50s earns a salary of $60,000 with an account balance of $100,000. Over the 13.5 year accumulation period the salary grows at a rate of 4 per cent per annum such that their income near retirement (i.e. mid 60s) is $102,868. Our hypothetical member contributes at the rate of 9 per cent per annum. The three accumulation paths are as follows: 'Actual path' uses synthetic balanced fund returns as they occurred over the period January 1999 to June 2012; 'Reverse path' uses the actual returns in reverse order; and, 'Average path' uses the arithmetic mean for the return for each period.

In Figures 2 and 3 we observe at least two common features. First, terminal wealth at the end of each 13.5 year path differs significantly, in most cases. These differences are shown in wealth and percentage terms in Tables 1 (Late 30s) and 2 (Near retirement). Despite each return path having identical arithmetic means, we can see terminal wealth can be significantly different. This clearly shows the limitations of arithmetic mean returns in measuring the performance of superannuation funds.

Second, we can observe that throughout the 13.5 year accumulation paths, wealth can differ significantly between paths. For example, in Figure 3, between December 2004 and December 2007 the actual and reverse paths are approximately $50,000 apart. Furthermore, in Figure 3, where there is more wealth at stake, the amplitude of the paths is more significant. A risk is that the 'roughness of the investment ride' might induce investors to make decisions that are suboptimal (e.g. to move out of risk assets too early, or too quickly).

TABLE 1: Late 30s

Measure	Actual	Average	Reverse
Arithmetic mean	Identical	–	Identical
Terminal wealth (%)	-6.4%	–	4.6%
Terminal wealth ($)	-$6,311	–	$4,506

Table 1 summarises in percentage and dollar terms the relative outcomes of the three paths shown in Figure 2.

TABLE 2: Near retirement

Measure	Actual	Average	Reverse
Arithmetic mean	Identical	–	Identical
Terminal wealth (%)	-4.2%	–	0.1%
Terminal wealth ($)	-$15,941	–	$284

Table 2 summarises in percentage and dollar terms the relative outcomes of the three paths shown in Figure 3.

But perhaps the most vivid example of the path-dependency of an investor's experience is a comparison of the actual and the reverse paths for each investor:[5]

> **Late 30s** — Actual path is 10.5 per cent lower ($10,817) than the reverse path.

> **Near retirement** — Actual path is 4.3 per cent lower ($16,225) than the reverse path.

We can therefore see that to understand risk in superannuation investing we need to consider realistic accumulation models that incorporate factors like contributions in addition to returns.

Investment policy needs to take account of the accumulation model and the dynamics it introduces. Research suggests that constant asset allocations don't look promising, especially when considered in light of the portfolio size effect and sequencing risk.

Implications for trustees

A number of implications for trustees present themselves from this analysis. First, returns-only measures don't capture wealth dynamics and therefore cannot shed light on the relationship between risk and investment horizon for a superannuation investor. A resolution to the time diversification puzzle thus hinges on an analysis of realistic accumulation models incorporating all relevant variables.

Second, retirement outcomes are highly path dependent. The member gets a single sequence of returns, not smooth 'average' returns. This reality highlights the importance of risk management especially in the latter half of the accumulation phase as the portfolio size effect manifests, and sequencing risk emerges as a serious risk. Investment policy needs to take account of the accumulation model and the dynamics it introduces. Research suggests that constant asset allocations don't look promising, especially when considered in light of the portfolio size effect and sequencing risk. Some researchers have presented evidence to suggest that dynamic strategies may offer assistance (Basu, Byrne and Drew 2011).

Investment governance may need to become dual-focused in order to be comprehensive. In its simplest form, this dual focus would see the following two questions being addressed:

> Are we achieving the member's wealth goals (a wealth-based or money-weighted consideration)?

> Are our managers delivering performance in accordance with expectations and their mandates (generally a time-weighted consideration)? ∎

Notes

1. All errors remain the sole responsibility of the authors.

2. Each of these series is the product of 10,000 trials for each of nine investment horizons using a stationary bootstrap simulation method. The standard deviation of annualised returns is calculated for each horizon by computing the mean monthly return for each of the 10,000 return paths, annualising each monthly mean, then taking a standard deviation of the 10,000 annualised means. The standard deviation of cumulative returns is calculated by applying the simulated return path for a given horizon to a starting value of $1. This is repeated for each of the 10,000 paths. The calculation is completed by taking the standard deviation of the 10,000 cumulative returns. Kritzman (1994), in his Figure A (p. 14), shows a 95 per cent confidence interval for annualised returns. This is the confidence interval equivalent of our standard deviation of annualised returns series.

3. Others include Tobin (1965) and Merton (1969).

4. The remaining two streams are the Black-Scholes-Merton Option Pricing Theory stream, beginning with Bodie (1991, 1995), and the behavioural stream. The option pricing theory approach of Bodie (1991, 1995) apparently emerged because of an unrelated breakthrough in economics, not as a result of a specific critique of Samuelson's (1969) work. Only later did others highlight that Bodie's (1995) approach appeared to offer an objective measure of risk in contrast to Samuelson's (1969) normative treatment of risk. Perhaps the most substantial critique of Bodie's (1995) work was that it was conducted in a risk-only framework. Behavioural economists are among the most vocal opponents of any framework that tends to see economics as (hard) science, as opposed to social science. These two visions of economics mix like oil and water. Behavioural economists introduce the richness of humanity to economic problems, often in qualitative terms, whereas 'scientists', of whom Samuelson (1969) was most definitely one, prefer to take approaches characterised by theoretical formality and the rigour of mathematical reasoning, even if it means making simplistic assumptions about human behaviour. In these few sentences, we have briefly outlined both the behaviouralists' principal critique of Samuelson (1969) — the inappropriateness of his underlying assumptions — and our critique of the behavioural stream of literature — the lack of framework, and negative approach to the problem. While the influence of the behaviouralists is limited to providing critiques of the other streams of the literature, they do provide some compelling insights relating to the selection of risk measures.

5. Recall that the period being considered is 13.5 years versus a typical accumulation phase of 40 years. It is possible that over a longer period the divergence could be even more significant.

References

Basu, AK, Byrne, A and Drew, ME 2011, 'Dynamic lifecycle strategies for target date retirement funds', *Journal of Portfolio Management*, vol. 37, no. 2, pp. 83–96.

Basu, AK, Doran, B and Drew, ME 2012, *Sequencing risk: A key challenge to sustainable retirement incomes*, Finsia (Financial Services Institute of Australasia), Sydney.

Basu, AK and Drew, ME 2009, 'Portfolio size and lifecycle asset allocation in pension funds', *Journal of Portfolio Management*, vol. 35, no. 3, pp. 61–72.

Bodie, Z 1995, 'On the risks of stocks in the long run', *Financial Analysts Journal*, vol. 51, no. 3, pp. 18–22.

Bodie, Z 1991, 'Shortfall risk and pension fund asset management', *Financial Analysts Journal*, vol. 47, no. 3, pp. 57–61.

Booth, L 2004, 'Formulating retirement targets and the impact of time horizon on asset allocation', *Financial Services Review*, vol. 13, no. 1, pp. 1–17.

Doran, B, Drew, ME and Walk, AN 2012, 'The retirement risk alone: A baseline study', *JASSA, The Finsia Journal of Applied Finance*, iss. 1, pp. 6–11.

Kritzman, MP 2000, *Puzzles of finance*, John Wiley & Sons Inc, New York.

Kritzman, MP 1994, 'What practitioners need to know … about time diversification', *Financial Analysts Journal*, vol. 50, no. 1, pp. 14–8.

Macqueen, AC and Milevsky, AM 2009 *Pensionize your nest egg: How to use product allocation to create a guaranteed income for life*, John Wiley & Sons Canada Ltd, Mississauga.

Markowitz, HM 1952, 'Portfolio selection', *Journal of Finance*, vol. 7, pp. 77–91.

Merton, RC 1969, 'Lifetime portfolio selection under uncertainty: The continuous-time case', *Review of Economics and Statistics*, vol. 51, no. 3, pp. 247–57.

Samuelson, PA 1969, 'Lifetime portfolio selection by dynamic stochastic programming', *Review of Economics and Statistics*, vol. 51, no. 3, pp. 239–46.

Siegel, JJ 1994, *Stocks for the long run: A guide to selecting markets for long-term growth*, Irwin Professional Publishing, Burr Ridge, Illinois.

Tobin, J 1965, 'The theory of portfolio selection', in *The theory of interest rates*, eds FH Hahn and FPR Brechling, Macmillan, London.