

CIFR Second Submission to Productivity Commission¹

Comments on the Draft Report:

How to Assess the Competitiveness and Efficiency of the Superannuation System

31 August 2016

AUSTRALIAN UNIVERSITY PARTNERS



GOVERNMENT PARTNERS



RESEARCH CENTRE PARTNERS



INDUSTRY PARTNERS



¹ This submission represents the views of CIFR's management team, and not the official views or policies of any CIFR stakeholders such as its funders and consortium members.

Comments on Draft Report: “How to Assess the Competitiveness and Efficiency of the Superannuation System”

The Centre for International Finance and Regulation (CIFR) welcomes the opportunity to provide further input to the Productivity Commission (PC) during Stage 1 of its investigation into the efficiency and competitiveness of the Australian superannuation industry. The comments in our second submission respond to the draft report dated August 2016.

This submission contains four parts:

1. *Comments on the draft report* – these are largely general in nature
2. *Response to information requests* – we comment extensively on benchmarking of performance
3. *Suggestion regarding collected data* – we ask the PC to consider making their data broadly available
4. *Additional CIFR research* – we note some relevant recent research that CIFR has provided to the PC

1. Comments on the Draft Report

CIFR would like to congratulate the PC on what is an excellent and comprehensive piece of work. We largely agree with the broad thrust of what the PC is proposing. More importantly, the report conveys a strong impression that the PC is across the key issues, and reveals a deep appreciation for the difficulties that bedevil what is an extremely challenging task. CIFR is also appreciative of the citations within the report of many CIFR-funded research projects, as well as its own first submission.

CIFR has a few comments to offer on the PC’s proposed approach and methods:

- **Working from stated objectives is a sound idea** – The approach of specifying objectives to provide context and focus for the analysis is commendable. The stated objectives seem sensible, and fit for purpose.
- **Too ambitious, so filter?** – When it comes to the assessment criteria and the related indicators, CIFR holds concerns that the PC may be attempting to cover too much ground. The approach seems to aim at identifying all relevant aspects, with the intent of addressing each as comprehensively as possible. While this is laudable, we question whether a more targeted approach might be more ‘efficient’, especially given the sheer size of the task being undertaken. Our experience as researchers warns that analysis always turns out to be harder, is more complex, and takes longer, than initially anticipated. CIFR recommends that thought be given to filtering the list of criteria and indicators in a disciplined manner, in order to make the job more manageable. Criteria such as ‘relevance’, ‘feasibility’ and ‘cost’ might be used to rank the proposed evaluation criteria and related indicators.
- **Combining measures** – Noticeably absent from the draft report is any detailed discussion of how analysis ranging over multiple dimensions is to be combined to form an overall ‘system’ view. (The discussion around the Sparrow Tiered Performance Framework on page 53 comes closest.) Perhaps the question of combining the findings was purposefully left out of the draft report, with the intent of addressing aggregation at a later stage. In any event, this important and difficult issue ultimately needs to be tackled; and a range of methods are possible.
- **Fund dispersion / frontier analysis (p117-118)** – CIFR queries the value of this style of analysis when the objective is to measure overall system efficiency. We express two misgivings. First, frontier analysis (including DEA) is ultimately a measure of relative efficiency between funds, and hence does not *directly* speak to the level of efficiency across the system. A high degree of clustering by funds is at best an indirect indicator, and could also be observed in inefficient systems. Second, the ultimate problem is the lack of an effective measure of output, given that this largely comprises investment returns which have high variability both across time and individual funds due to the nature of investment markets. It is hard to see how frontier analysis helps to overcome this core problem. Indeed, such methods may be comparatively exposed to measurement issues, given that the variability of investment returns will

manifest in high variability of fund returns, and will add much noise to any estimates. That is, dispersion across funds could reflect interaction between the vagaries of markets combined and differences in investment exposures, as much as indicating inefficiency among some funds.

2. Response to Information Requests

CIFR offers comments on two areas where the PC has requested further information, and where it has something to add: (1) benchmarking of investment returns, and (2) other sources of data and evidence. Many of the other information requests contained in the draft report are better left for those with greater expertise in the specific topics. To assist the PC, we also note some people known to CIFR who might be worth consulting on specific matters, if the PC has not already received their input. This information is provided in a confidential attachment.

Information requests on investment benchmarks (p115, p117, p139)

In its first submission, CIFR discussed various problems with measuring the value-add from investments, and recommended evaluating MySuper (i.e. default) funds using a reference portfolio approach. In this submission, we respond to the PC's request for more information on benchmarking methods. We also expand our discussion on the assumption that the PC would prefer to evaluate a wide range of products.

Benchmarking for Asset Class and Multi-Asset Products

We initially comment on the information requests appearing directly below. We do so by proposing a unified approach that spans all four areas listed.

- *Reference portfolio for long-term net returns, both system-wide and market segments (p115)*
- *Asset class benchmarks (p117)*
- *Data on net returns, including fees and tax (p117)*
- *Benchmarks for optimal allocation by age cohort (p139)*

The philosophy underpinning our recommended approach is that members choose to invest in (or are assigned to) a particular class of product on the expectation that it delivers exposure to investments that accord with their objectives. As a consequence, it is appropriate to calculate performance at a product class level, and then aggregate the results to a system-wide level. We further recommend that the performance of various product classes be benchmarked against readily investable, low cost, passive investments. The unit of measurement for this analysis should be net returns, allowing for taxes and fees, applied to *both* the fund products and their benchmarks. The remainder of this section sets out our suggestions for how such analysis might be undertaken.

Product Sample

To implement the recommended approach, the first task is to identify a sample of products to be included in the analysis. The PC's draft report observes that registerable superannuation entities (RSEs) offer around 40,000 investment options (p30). Many retail providers have menus of in excess of 100 products; and some RSEs offer scope for members to invest directly in individual managed funds or listed securities. The most effective approach is to extract a sample from this quagmire that is representative of the major product classes offered by the superannuation industry. CIFR suggests that analysis be conducted on the key multi-asset or 'pre-mixed' products (e.g. conservative, balanced, growth, high growth; and MySuper balanced strategies plus their predecessor defaults); as well as selected asset class products. Some analysis would be required to establish the appropriate scope for the latter; but it is envisaged that coverage would at least include Australian equities, international equities, listed property, fixed income and cash.²

² In fixed income and listed property, an issue is whether analysis is based around Australian and international, versus globally diversified products, given that these asset classes are offered in both formats within the industry.

Whether life-cycle products should be analysed is more problematic. The approach we outline can accommodate performance evaluation for life-cycle strategies (as outlined further below). Nevertheless, a number of issues arise. First is a lack of history for this class of product, which makes analysis of realised performance fraught for the foreseeable future. Second is that some life-cycle strategies are based on switching members between existing multi-asset products as they age. This strategy employs products that should be captured by the PC's evaluations in any event. Cohort-based strategies represent the main distinct class of life-cycle products, and comprised 16 out of 23 life-cycle products in early-2014 according to Chant et al. (2014).³ Third, Chant et al. also point out that a substantial portion of life-cycle products (approximately 40%-50%) are built from passive investments. This diminishes the value in comparing their performance with a benchmark that is built from passive funds. Arguably it is more relevant to form a judgement on whether life-cycle strategies represent an efficient innovation in product design relative to traditional balanced strategies,⁴ rather than evaluate their (limited) historical performance.

Product Net Returns

Estimating net returns on the product sample would ideally entail collecting data on declared returns and fees on a product-by-product basis.⁵ Tax realizations should be embedded in the declared return series. The most effective approach to calculating net investment returns might be to account for investment fees only. This avoids issues around administration fees, which are typically expressed in fixed dollar amounts, and are conceptually intended to cover non-investment costs in any event. Given the focus is system efficiency, ideally the product returns would be aggregated to the product class level through weighting by funds invested.⁶

Benchmarks

CIFR recommends identifying a set of passive index funds or ETFs that can be used as not only benchmarks for the asset class products, but also form the 'building blocks' for construction of reference portfolios to be used as benchmarks for multi-asset products. It is envisaged that the set of passive funds would include Australian equities, international equities (both hedged and unhedged),⁷ Australian and international listed property, Australian fixed income, international fixed income (hedged), and Australian cash.⁸

In the case of products based around single asset classes, identifying the benchmark will often be straightforward, particularly for 'core' asset classes like equities.⁹ Problems can arise where an asset class product relates to unlisted alternative assets, as investible passive indices are typically unavailable. Further, series such as the Mercer/IPD Property Fund Index and IPD Unlisted Infrastructure Index reflect assets held by the major superannuation funds, and hence do not amount to independent benchmarks. Fortunately, unlisted alternatives are rarely offered as stand-alone products.¹⁰ Hence they can probably be excluded without greatly biasing any system-wide evaluation. Another option might be to benchmark these products

³ Chant, W., Mohankumar, M. and Warren, G. "MySuper: A New Landscape for Default Superannuation Funds", *CIFR Research Working Paper*, No. 020/2014, July 2014

⁴ Unfortunately, there is no clear answer to the question of whether life-cycle strategies as they are currently designed are superior to balanced products. The answer may also vary across members.

⁵ Industry research houses should have much of the required data.

⁶ Industry research houses might be able to provide their data in aggregated form for product classes.

⁷ The benchmark currency hedge ratio is an important variable, and can be influential in performance measurement. A benchmark hedge ratio can be incorporated into the analysis via adjusting the benchmark weights on unhedged versus hedged international equities. A 50% hedge ratio for international equities (and listed property) should be a defensible assumption. A fully hedged position in international fixed income is the industry standard.

⁸ Vanguard, for instance, currently offers products in all these categories. Where available history is limited, a proxy return might be constructed using the indices themselves combined with fee and tax assumptions.

⁹ While most asset class benchmarks will be straightforward, difficulties may arise in fixed income as product configurations vary across the industry. One approach might be to use sovereign bonds as the benchmark, in which case activities such as credit exposure and 'core-plus' investing is effectively treated as a potential value-add activity.

¹⁰ This reflects their illiquidity, which does not sit well with member investment choice. Some funds nevertheless offer access to options that house illiquid alternatives, e.g. the AustralianSuper and HostPlus property options.

against their passive listed counterparts; although this is not particularly satisfactory given that the basis of pricing and valuation is quite different (i.e. traded prices versus appraisal valuations.)

For multi-asset products, a reference portfolio approach is recommended. This involves identifying the portfolio comprised of investable, low cost, passive investments that reflect a static strategic asset allocation (SAA) according with the long-term objectives for members who invest in the product class in question. The reference portfolio approach effectively measures the additional value created by fund providers arising from a combination of: (a) dynamic asset allocation; (b) active investment management; and, (c) investment in assets outside the reference benchmark, including alternative assets.

Reference portfolios could be based around representative SAAs for each multi-asset product class. For example, Chant et al. (*op. cit.*) report that the average SAA for MySuper balanced funds in early-2014 was 72% growth and 28% defensive (including 14% and 6% in growth and defensive alternatives respectively). The reference portfolio for these products might comprise (say) 32% in Australian equities, 30% in international equities, 10% in listed property, 23% in fixed income and 5% in cash.

If the PC is intent on evaluating the performance of life-cycle products, a reference portfolio for various age cohorts might be constructed from passive building blocks in the following manner. First, a representative 'glide path' for growth versus defensive assets is specified. This could mimic the industry average,¹¹ on the basis that this is what members have committed to.¹² Second, the passive index building blocks are used to construct 'growth' and 'defensive' component benchmark portfolios. Third, the weight between these two component portfolios is then varied in accordance with the glide path to form a series of reference portfolios against which life-cycle product returns can be compared for specific cohorts.

Benchmark Net Returns

In addition to declared gross returns on the passive investments, it would be necessary to collect both the management fees paid by large institutional investors and tax distributions. Hopefully passive funds providers would assist. Estimation of net benchmark returns would be formed by deducting fees and applying taxes at the rate for superannuation funds.

Performance Evaluation

Performance evaluation would occur by comparing the net return history for each product class with the matched benchmark return, and then aggregating. A long data history would be desirable, ideally extending back at least 10-15 years in order to span multiple cycles. Analysis of performance might be supplemented by an examination of return trends over time, in order to assess dynamic efficiency.

Attempts should also be made to account for the riskiness of product class portfolios versus their benchmarks. The aim might be to gain insight into whether the industry has protected members from downside risk at the margin, relative to investing in basic, passive investments. Measures of absolute risk are most relevant in this context, rather than risk measured relative to benchmark such as tracking error. Preferred risk measures might be based around distributions of cumulative payoffs over longer holding periods, e.g. 3-, 5- or 10- years. Measures of shortfall might be calculated for the products (e.g. losses; outcomes relative to target), and compared with similar measures for the benchmark. Volatility and Sharpe ratios might also be calculated for the product classes, and compared to the benchmarks. However, volatility-based risk measures are not preferred for two reasons. First, volatility is a questionable measure of risk when objectives relate to long-term payoffs. Second, volatility can be attenuated where products include unlisted asset classes with smoothed returns, resulting in misleading impressions of lower risk.

¹¹ Chant et al. (*op. cit.*) report the average glide path for MySuper products as at early-2014.

¹² Another approach could be to specify an 'optimal' glide path, if the PC has a sound basis for doing so.

CPI + X: What is the appropriate level of 'X' (p115)

CIFR holds misgivings about evaluating fund efficiency through a direct comparison of realized returns with real return targets. The main reason is that outcomes are dictated by the performance of markets over which funds have no control (and the industry at large cannot avoid).¹³ Also, an infeasibly long time period is required to form conclusions based around returns relative to target with any confidence. Nevertheless, the PC may well intend to examine real return levels, given that they strike at the prime system objective to support income in retirement. CIFR suggests that the most straightforward approach is to compare real returns versus product class level targets, then aggregate. Real return targets might be based on the average of those published by superannuation funds for specific product classes. Using an average helps mitigate the problem that return targets can differ across the industry, depending on the assumptions. This approach accords with the philosophy that members have chosen to invest in (or are assigned to) a particular product under these indicated return targets, and hence they constitute a valid yardstick. It also dovetails with the benchmarking approach we recommend above, given that the analysis involves substituting target returns for the passive fund benchmarks.

What other benchmarks should be used? (p115)

CIFR views the methods referred to within the draft report as the best available. No other compelling benchmarks come to mind. We note that the ideal would be to benchmark against an optimal portfolio, rather than low-cost passive investments. Unfortunately, identifying the optimal portfolio is impractical. This is for a range of reasons, the most important of which is that no ubiquitous system-wide optimal solution exists, as the optimal portfolio is likely to vary significantly across members.

Other sources of data and evidence (p165)

We recommend that the PC investigates tapping into administrative data. As the draft report observes, some of the administrative functions in the industry – such as custodian and member administration – are out-sourced to a handful of suppliers. These suppliers will warehouse large amounts of data, and together should provide a broad footprint due to their concentration. They hence might provide an efficient way of gathering information on investment activities and members, if the data can be accessed.

3. Suggestion Regarding Collected Data

The draft report makes mention of collecting a broad range of information, including potentially conducting surveys where existing data sources are inadequate. The data collected could also serve a broader purpose if it were made available to academics and others if feasible to do so, perhaps in de-identified form. As well as extracting additional community benefit from the effort in compiling the data, having the capacity to share data with other parties may present the PC with an avenue to commission research around topics of interest. In the experience of CIFR, one of the most effective ways to incentivise academics is to provide them with access to unique data, and allow them to use the findings in pursuit of an eventual publication. SIRCA may be worthwhile talking to, given it is a body with connections to the university sector that warehouses data, and makes it available to the academic community.

4. Additional CIFR Research

Following CIFR's first submission in April 2016, a number of research outputs have been subsequently provided to the PC given their relevance to competitiveness and efficiency in the superannuation industry. The research, which is described below, was sent to the PC during August 2016.

- **CIFR Project T004 on Superannuation Fees** – This project examines the relation between fees and superannuation fund performance, allowing for the influence of asset allocation on the expected level of investment fees. Two working papers from this project were finalised in July 2016. The research

¹³ The recommended benchmarking approach, as discussed above, controls for market effects.

uncovered no clear relation between the level of fees and performance, such that it was not apparent that higher fees were necessarily detrimental to members. The research employed benchmarking methods that the PC might want to review.

- **Capacity in Fund Management** – CIFR has generated a draft paper on capacity in funds management, a topic that is directly relevant to the question of efficiency. The aim of the research project is to produce a ‘how it is done’ guide to analysing capacity for the benefit of the fund management industry at large. In doing so, the report addresses issues that relate to scale more broadly, which is an area on which the PC appears to be seeking further input. At the time of writing this submission, the draft paper had been sent around the investment industry, seeking comments before finalising.
- **In-House Investment Management** – CIFR has also provided its paper on in-house management, on the basis that this issue is mentioned in the draft report without indication that the CIFR research had been noted. A trend towards in-house management is evident among larger superannuation funds, and has implications for both efficiency and competition. Much of what is motivating larger funds to in-source relates to efficiency; whereas the ‘threat’ that a fund may take the assets in-house adds another element to the competitive dynamics with external investment managers. The report itself describes how the industry (specifically not-for-profit funds) view and approach in-house management. CIFR believes this research might provide useful background on this piece of the puzzle.

List of Additional CIFR Research Papers Supplied to the PC (also available on CIFR website and SSRN)

“Are Superannuation Industry Fees Too High?”, A. Ainsworth (USYD), S. Akthar (USYD), A. Corbett (USYD), A. Lee (UTS), T. Walter (USYD), July 2016 (*CIFR-funded targeted research; project T004*)

“Superannuation Fees and Asset Allocation?”, A. Ainsworth (USYD), S. Akthar (USYD), A. Corbett (USYD), A. Lee (UTS), T. Walter (USYD), July 2016 (*CIFR-funded targeted research; project T004*)

“Evaluating Fund Capacity: Issues and Methods” (Draft), August 2016, by M. O’Neill (Investors Mutual) and G. Warren (CIFR) (*Internal CIFR research – not yet publically available*)

“In-House Investment Management: Making and Implementing the Decision”, March 2016, by D. Gallagher, T. Gapes and G. Warren (*Internal CIFR research; project F016*)

Authorship: Geoff Warren (Research Director), with input from David Gallaher (CEO) and Tim Gapes (Centre Director)