



**Chris Drew,
ACSRF**

Chris Drew is an investment manager, public markets at Australian Catholic Superannuation and Retirement Fund (ACSRF).



**Raewyn Williams,
Parametric**

Raewyn Williams is the managing director, research at Parametric.



**Mahesh Pritamani,
Parametric**

Mahesh Pritamani is a senior researcher at Parametric.

OPTIMALITY IS IN THE EYE OF THE BEHOLDER

The Science and Art of Equity Portfolio Construction

Chris Drew, Raewyn Williams & Mahesh Pritamani

Having the right ingredients for equity investing is mission-critical for superannuation funds. On average, APRA-regulated funds invest almost half of members' capital (47%) into listed equities portfolios. Much research exists on the ingredients funds can use to construct a physical equity portfolio—matters like the attractiveness of investing onshore or offshore (in developed versus emerging markets), whether to adopt an active or passive management philosophy, the exploitation of factor risks, different weighting strategies, how to define risk and manage risk budgets, measuring trading costs and the impact of tax on equity investing. Far less research exists on how to combine these singularly important ingredients into the right mix from a taxable superannuation fund investor's perspective.

In this paper we argue that how these ingredients are combined matters. We consider a hypothetical superannuation fund with the following global equity ingredients at its disposal:

- A set of investment beliefs
- A passively managed core portfolio
- Four actively managed, diversified satellite portfolios

- A tax-management option
- A risk budget
- A fee budget

How would a fund combine these ingredients in an optimal way? In the theoretical, frictionless world of 'best ideas', the task is simple: Funds that believe strongly in efficient markets and the 'free lunch' of diversification would maximise allocations to a core passive portfolio, while funds that believe in active management would maximise allocations to the multi-manager satellites. This task is much more complex in the real world, where funds are confronted with nuances like their level of conviction in active management, how much benchmark-relative and peer risk is appropriate to pursue excess returns (alpha) with conviction, and trading off the certain costs of taxes, brokerage and commissions with the uncertainty of alpha. Overlaying this will be the fund's ultimate reality check: a sensible fee budget for pursuing its equity portfolio objectives.

For our hypothetical fund, we adopt a core-satellite structure and show how the fund would identify an optimal combination of an 80% allocation to a passive core and 20% allocation across the satellite active managers, spanning developed and emerging-market equities. This allocation provides the highest potential to maximise returns on a risk-adjusted basis, post-tax, while requiring only that the active



The quote

Higher conviction does two things: It emboldens the fund to add tracking error because the expected alpha payoff is higher, but it also strengthens the incentive to add tax management to the overall solution

managers be moderately skilled and the passive core be tax-managed to provide some damage control over the tax impacts of active management.

This optimal blend keeps overall tracking error (to our hypothetical fund's benchmark) well below 2% and overall fees well below 0.5%—our fund's annual risk and fee constraints—which we believe is very competitive for a multi-manager all-countries equity portfolio with active-investment and tax-management components. A key aspect of our research is showing how our hypothetical fund then steps back from this analysis—what we call the science—to add the art, or the wider lens of judgement, industry awareness and member-centricity, to ultimately settle on an optimal equity portfolio design.

Because optimality is in the eye of the beholder, how well these findings guide superannuation funds will depend on whether they relate to the investment beliefs, objectives and sensitivities of our hypothetical fund. While our fund profile is realistic, even funds that consider themselves very different (for example, not peer-sensitive, not fee-constrained or completely indifferent to tax) will find useful insights.

First, we demonstrate the importance of articulating investment beliefs and reflecting these, as far as practical, in the way equity portfolios are constructed. Second, any fund can use the model we follow in this paper to help blend particular equity portfolio ingredients in an optimal way. A third application is to deepen funds' understanding about how much tax can eat away at alpha programmes—something obscured in the pre-tax performance reporting many funds use—and how a simple device like a tax-managed core portfolio can help. Fourth, this research questions the wisdom of funds setting fee and risk budgets as hard, rather than soft, constraints—many of the optimal portfolio possibilities we identify are simply off the table for funds investing under such strictures.

Defining 'optimal' for our hypothetical fund

Our research cannot feasibly identify a single, universally optimal combination of equity managers and strategies for superannuation funds because investment philosophies differ from fund to fund. What is an optimal mix of equity portfolio ingredients for a fund depends on how that particular fund defines *optimal*. A fund could view its equity portfolio design as optimal by delivering a portfolio with the greatest potential to maximise any of:

- Pre-tax returns or excess returns (alpha)
- Returns or alpha post-tax
- Risk-adjusted returns (Sharpe or information ratios, pre- or post-tax)

A fund could also use a dashboard approach, which looks at how well a set of these investment objectives are achieved in combination.

Our hypothetical superannuation fund believes in active management and the power of diversification (while being wary of over-diversification) and that it can pick moderately skilled active managers. It is concerned with

balancing the risk and return dimensions of its portfolio and thinks of risk primarily as tracking error to a market-cap-weighted benchmark. The fund wishes to align itself with what matters to members—after-tax returns—and therefore considers taxes in the design and management of the equity portfolio. With an eye to future RG 97-type obligations, the fund also views favourably styles that more naturally limit transaction costs. The fund is large enough to access direct accounts and does not invest via pooled funds. Finally and importantly, our hypothetical fund allocates annually a total equity portfolio risk budget (tracking error) of 2% and fee budget of 50 basis points. The fund recognises that limiting tracking error is also helpful in mitigating peer risk.

This set of investment beliefs and sensitivities leads our hypothetical fund to define as optimal the equity portfolio that maximises the probability of generating the highest after-tax excess returns (alpha), adjusted for risk, annualised over a long-term time horizon. It also leads the fund to consider a core-satellite equity structure, in which an active-passive combination can meet its fee and tracking-error requirements, and to select all four active-manager 'ingredients', with allocations reflecting its all-countries large- and mid-cap (ex-Australia) benchmark, for diversification benefits. We believe this setup is one to which many funds will relate, even though neither the risk nor fee budgets are generous.

With this realistically profiled fund and equity structure, we have established the ingredients our hypothetical fund has at its disposal to realise its investment beliefs and objectives: a passively managed core portfolio; four actively managed, diversified satellite portfolios; a concern for tax impact; a risk budget and a fee budget. We now turn to the key question: How can the fund blend this core-satellite portfolio in an optimal way?

Designing our core portfolio

Using the research methodology of Bouchev and Pritamani (2016), we design our core strategy as a passive portfolio to track the S&P BMI Global ex-Australia market-cap-weighted universe of large- and mid-cap stocks. Details regarding the construction of this benchmark are provided in the next section, since it also serves as the benchmark for the multi-manager satellite structure. Tax management is included to boost its after-tax return while tracking its benchmark. This requires a small risk budget (like all sources of active return); hence our core portfolio allows active stock positions of +/-0.25% against the benchmark weight. We also allocate 10 basis points out of our annual fee budget to this core exposure.

The low-risk nature of the core portfolio (based on how our fund defines risk) ensures it will not deviate dramatically from the benchmark on a pre-tax basis. We are, however, keen to see how well our tax-managed core can produce a tax cushion to help with our active-management programme. We simulate the tax-managed core portfolio returns (net of management fees and transaction costs)

Table 1. Simulation results: tax-managed core (annualised)

	Pre-tax (after fees)	After-tax (after fees)
Portfolio return (%)	6.62	6.65
Benchmark return (%)	6.71	6.42
Excess return (%)	-0.09	0.22
Information ratio	-0.14	0.33
Impact of taxes (%)		0.32
Tracking error (%)	0.68	
Portfolio turnover (%)	7.92	
Benchmark turnover (%)	7.00	

Source: Parametric. Simulated returns from 1 January 2008 to 31 December 2017, annualised, net of transaction costs. Tax calculations apply complying superannuation fund tax rates and rules on a pre-liquidation basis. Simulated performance is hypothetical and does not reflect the results of any investor. All investments are subject to risk of loss. See Disclosures for additional information.

over a 10-year period from the beginning of calendar year 2008 to the end of 2017. Table 1 suggests this small allocation of our risk and fee budget will pay off for our hypothetical superannuation fund.

We see that the core portfolio uses tracking error of 68 basis points per annum to produce a very small negative pre-tax excess return for our hypothetical superannuation fund, after management fees. The fund views this, essentially, as the ‘no surprises’ ingredient in its total equity portfolio mix, with the added bonus of a payoff on the tax-management component of the strategy. This is reflected in the positive 22-basis-point after-tax excess return earned by the core portfolio relative to a pre-tax excess return of -9 basis points each year. This difference of close to 32 basis points represents the positive impact of tax management on performance, since realised losses in the core portfolio can offset taxable gains in the satellites. Overall, our tax-managed core promises to be a useful ingredient in our optimal combination search, providing broad market exposure, using only a small amount of our hypothetical fund’s risk and fee budgets and improving the overall tax efficiency of the combined core-satellite portfolio.

Designing our satellite portfolios

Our hypothetical superannuation fund has modest conviction in relation to active management, so we now turn to the potential active ingredients in its equity portfolio mix. Our satellite portfolios come from four diversified manager styles as shown in the table below.

Table 2. The four satellite portfolios

Capitalisation	Style	Region	Number of stocks held	Tracking error (%)*	Initial weight (%)
Large and mid	Neutral	Developed markets ex-Australia	100	4.1	43.80
Large and mid	Growth	Developed markets ex-Australia	75	4.9	22.43
Large and mid	Value	Developed markets ex-Australia	75	5.0	21.37
Large and mid	Neutral	Emerging markets	100	5.1	12.40

*Tracking error is measured against a cap-weighted universe of stocks from the same style and region.

Our hypothetical fund’s belief in the value of diversification (but not over-diversification) leads it to select a four-manager structure, one from each style, resulting in a multi-manager active satellite exposure covering developed and emerging markets. The initial weights are driven by the weights of the styles in the benchmark at the beginning of the analysis. The individual satellite strategies are drifted (that is, not rebalanced) throughout the analysis period. The benchmark for the multi-manager structure is also constructed similarly as a composite of the underlying style universes with initial weights set at inception and drifted thereafter. As mentioned earlier, the tax-managed core portfolio also tracks this benchmark.

While the absolute levels of tracking error for each individual strategy may be too high for some funds to tolerate, the diversifying nature of the multi-manager approach—a key motivation for superannuation funds adopting this structure—suggests an overall tracking error of 2.39% to 2.71% each year for the combined satellite portfolios, depending on the skill level of the individual actively managed satellite strategies. This seems realistic for funds with a risk budget to pursue active management and is a useful ingredient in our hypothetical superannuation fund’s optimal equity portfolio design.

The four manager satellite portfolios’ performance is simulated with varying degrees of skill using the framework in Sorensen et al. (1998) and applied by Bouchev and Pritamani (2016) and Bouchev, Pritamani and Williams (2017). *Skill* is defined as the number of stocks the active manager selects that outperform the benchmark. Outperforming stocks are those in the top half of the universe based on the next one-year returns. At year-end the individual strategies are reconstituted to maintain the same level of skill over the following year. Skill levels vary from 40% (poor skill) to 70% (excellent skill), with 50% indicating no skill and a benchmark-like performance (before trading costs).

The satellite portfolios in each style universe are simulated as follows:

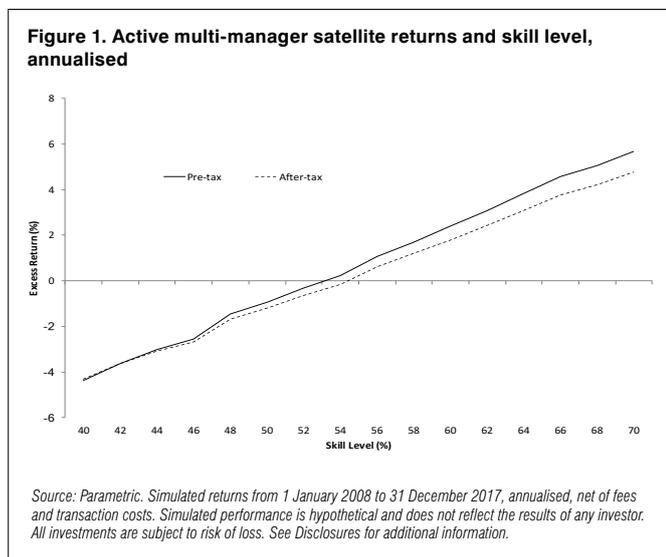
- At the beginning of each year, the stocks in each style universe are classified as winners or losers based on their performance over the coming year.
- We pick the number of stocks presented in the table above and use a stratified sampling approach so the portfolio picks the correct number of winners to reflect its skill level.
- Stocks are held in the portfolio for the entire year unless they drop out of their style universe, in which case they are replaced with random stock picks from their style universe.
- At year-end, stocks are reclassified as winners or losers based on their performance over the following year. At this point the portfolio will no longer match its skill level and is reconstituted to retain its skill characteristics. For instance, if the skill level 60 portfolio

(implying the manager picked winners 60% of the time) held only 52% winners at year-end after reclassification, we would sell 8% of the stocks classified as losers and replace them with winners. We also introduce turnover at year-end by selling 25% of the stocks held by the growth and value managers and 50% of those held by the style-neutral managers. The stocks sold remain part of the eligible universe and can be bought back, so the net turnover introduced could be lower. On the whole the portfolios experience turnover each year-end as part of the reconstitution process and during the year when there are membership changes in the underlying style universes. Annual turnover in the portfolios in the developed markets ex-Australia style universes is around 55% to 60% and around 40% for the emerging-markets universe.

- The stocks are cap-weighted.

We assign a total blended management fee of 60 basis points per annum to these satellite portfolios. As a final dose of realism, the simulated portfolios incur annual transaction costs of 25 basis points per dollar traded.

We simulate the four-manager satellite portfolio returns over a 10-year period from the beginning of calendar year 2008 to the end of 2017. All four managers have the same skill level (ranging from holding 40% to 70% of stocks that outperform). For each skill level, we simulate 100 such four-manager combined portfolios and report the average performance across the 100 simulations. Figure 1 shows the returns (net of management fees and transaction costs) generated by the satellite portfolios collectively at the various skill levels on a pre- and after-tax basis and illustrates a very intuitive relationship between skill and returns.



As expected, the worst performance comes from skill level 40, with an annual pre-tax excess return of -4.36%, while the best performance comes from skill level 70, with an annual pre-tax excess return of 5.67%. Our hypothetical superannuation fund believes it can find moderately skilled active managers for its satellite programme, so it focuses on points on the above lines around skill level 56 on the x-axis. (Later we explore the impact of relaxing this conviction around ‘moderate skill’.) This suggests that the fund’s active-man-

Table 3. Potential annual payoffs from moderately skilled active managers

Skill level	56
Excess return (pre-tax)	1.07%
Excess return (after-tax)	0.63%
Tracking error	2.39%
Information ratio (pre-tax)	0.45
Information ratio (after-tax)	0.26

Source: Parametric. Simulated returns from 1 January 2008 to 31 December 2017, annualised, net of fees and transaction costs. Simulated performance is hypothetical and does not reflect the results of any investor. All investments are subject to risk of loss. See Disclosures for additional information.

agement programme can provide the following payoffs each year: Table 3, along with Figure 1, again underscores the importance of anchoring portfolio design in the fund’s investment beliefs. A pre-tax-focused fund sees a potential annual reward of 1.07% alpha, pre-tax, for a risk budget of 2.39%, giving an information ratio of 0.45. But our hypothetical fund has an after-tax investment focus and is interested in the innate trade-off between the investment alpha the satellite managers generate and the tax paid on these returns, noting that the higher the skill level, the higher the amount given up in tax. As Figure 1 shows, from skill level 42 upward, the after-tax excess returns are consistently lower than the pre-tax excess returns. This difference increases with the level of skill because the highly skilled managers generate higher levels of taxable capital gains.

Focusing on moderately skilled managers (skill level 56), the annual alpha payoff for our fund’s 2.39% allocated risk budget reduces to 63 basis points after tax (almost half of the adjudged pre-tax payoff), reducing the information ratio to 0.26. While investment philosophies differ, as we have noted, it seems clear that our hypothetical fund is right to focus on after-tax investment outcomes given how different these portfolios look when real-life tax impacts are factored in.

As we turn now to the key task of combining these core and satellite elements in an optimal way, we can already see how the mounting impact of a triad of sensitivities—risk, fees and tax (we might also add transaction costs)—will skew a fund’s preferences toward a higher allocation to a passive core relative to funds without such sensitivities. In other words, the natural upper bound to an active satellite allocation will be lower the more sensitive a fund is to these real-world costs.

The science of optimising equity portfolio ingredients

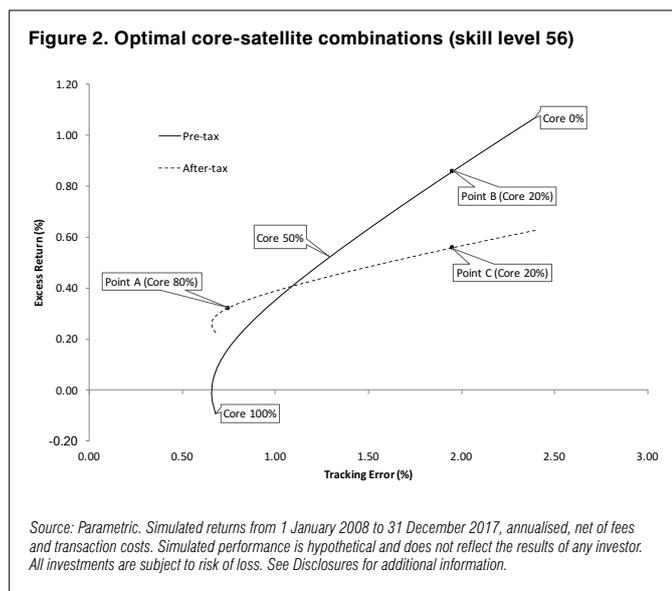
Our hypothetical superannuation fund is looking to maximise risk-adjusted after-tax excess returns by combining in an optimal way its core and satellite ingredients with the following expected properties:

As a reminder, our hypothetical fund hopes to stay broadly in line with peers and operate with 2% tracking error (risk) and 50 basis points in fees per annum.

For any given level of active skill where there is an expectation of positive alpha, we can plot a set of core-satellite combinations and their expected after-tax excess returns for different levels of tracking error (or, in fact, any other valid return objective). Our hypothetical superannuation fund has a range of scenarios available to blend its moderately skilled active satellites with its tax-managed core portfolio, as shown in Figure 2 (pre-tax return combinations also shown for comparative purposes), to fit within its risk and fee constraints.

Table 4. Expected properties

Portfolio	Tracking error
Core	0.68%
Satellites	2.39% to 2.71%



The optimal blend for our hypothetical fund sits on the dashed curve (representing after-tax outcomes) at the point that maximises our fund’s information ratio (excess return divided by tracking error). We find this to be at point A — a blend of 80% core and 20% satellites — which produces an information ratio of 0.43. This core-satellite blend spends only 20 basis points of the fee budget and only 74 basis points of the risk budget each year to produce an excess return, after fees, taxes and transaction costs, of 32 basis points. The relative flatness of this after-tax frontier explains why our hypothetical fund’s optimal blend sits so far to the left (leaving much of its risk and fee budget unspent): The return payoff from adding tracking error is much smaller when the tax erosion to the extra alpha is factored in.

Clearly, our hypothetical fund’s personal expression of its investment mission (an alpha focus, after accounting for risk and tax) has a crucial impact in determining its optimal core-satellite combination. Funds that define optimality differently could instead identify as optimal (within the same risk and fee constraints):

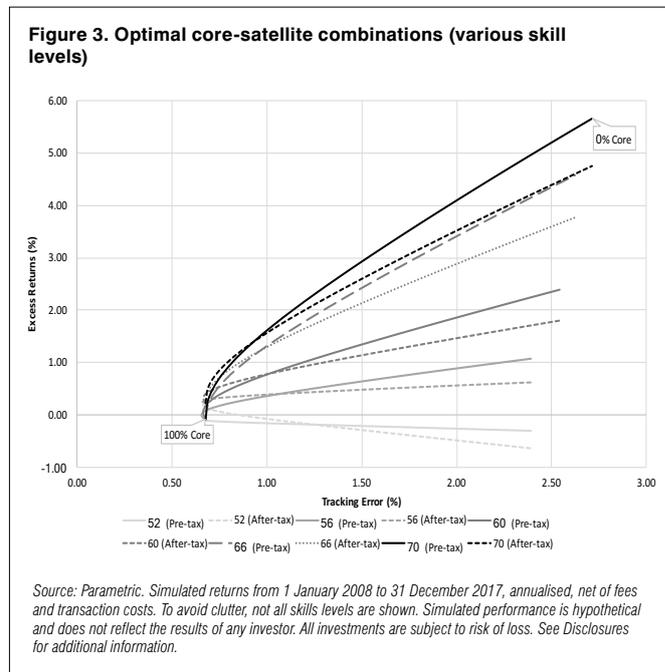
- Point B, which maximises pre-tax returns (total and excess) and the pre-tax information ratio, or
- Point C, which maximises total returns after tax.

A dashboard of objectives may mean some optimal point that is different, again, to points A, B or C. While these are legitimate objectives, we do not advocate focusing on pre-tax returns (point B) since this ignores the impact of tax that members of our fund have to bear, which pushes point B down to Point C. Further, we have a preference for point A over point C due to the lower variability in outcomes being driven by a lower allocation to active management. At point A we are giving up only 24 basis points in expected post-tax alpha each year for a much higher degree of assurance over the outcome.

A wider perspective: where science meets art

While we could at this point end our demonstration of how a superannuation fund can optimally blend equity portfolio ingredients, we think there is an important further area of exploration required—what we might call the art, or judgement, overlaying this scientific (quantitative) approach. Thinking of this as a mission-critical real-world task, we must ask: Is our hypothetical fund’s commitment to its fee budget really a hard constraint? Does its risk budget deny our fund the opportunity to meaningfully outperform its fund peers? Can our hypothetical fund really specify that its confidence in active management is perfectly captured in skill level 56? We believe many funds could start with an optimisation model like ours as a baseline proposition but would benefit from exploring what other possibilities exist if they are more flexible around certain aspects of their portfolio design.

Let’s take our moderate-skill assumption as an example. If our hypothetical superannuation fund can strengthen its active-management conviction through time (say, by building up a good performance track record, changing asset consultants or even insourcing), it may prefer a set of core-satellite possibilities reflecting higher skill. Alternatively, it could lose confidence in its active-management programme, decide alpha is too seasonal or unpredictable or decide to focus more on what it can control (taxes, transaction costs and fees). This would require the fund to contemplate a set of core-satellite possibilities reflecting lower skill. Different expectations around what active skill can deliver will change the optimal mix of ingredients for our hypothetical fund. We capture this in Figure 3 while still limiting ourselves to scenarios in which managers have some skill.



We see here a set of intuitive results: Our hypothetical fund’s frontier of possibilities becomes flatter as conviction in active skill declines and steeper as it increases. Higher conviction does two things: It emboldens the fund to add tracking error because the expected alpha payoff is higher, but it also strengthens the incentive to add tax

management to the overall solution (because the gap between pre-tax and after-tax outcomes widens as skill level increases, seen here and earlier in Figure 1).

In fact, in this real-world scenario, we see tax working much like a fee or risk budget to limit the amount our hypothetical fund is ultimately willing to allocate to an active satellite program. As skill increases, the impact of taxes on excess returns becomes increasingly negative and, yearly, would go from -44 basis points for a 100% skill level 56 satellite portfolio to -91 basis points for a skill level 70 satellite portfolio. This translates to giving up over 40% of excess returns through tax drag for skill level 56 satellite portfolios (16% for skill level 70 satellites) and clashes with our hypothetical fund's belief in the value of implementation efficiencies achieved through portfolio construction techniques, an after-tax focus and fee controls.

By allocating to the core, the fund limits this tax leakage in two

ways—by having less of its overall equity portfolio exposed to the tax damage done in the active programme and by employing tax management in the core portfolio to help shelter the gains generated in the multi-manager satellites from tax. Ironically, our hypothetical fund will get the best out of its active managers by reserving some allocation for a tax-managed passive core.

Let's now open up our findings to a panoply of possibilities if our hypothetical fund is willing to embrace the 'art' of portfolio design and be more flexible with its risk and fee budgets. Rather than stopping at a specific point that scientifically blends our hypothetical fund's equity portfolio ingredients—a set of investment beliefs, a core passive portfolio, four active diversified satellite portfolios, tax management and a risk and fee budget—optimally, let's step back to identify in Figure 4 a sweet spot of possibilities to consider.

The combinations shaded represent the eligible opportunity set

Figure 4. Optimal core-satellite combinations (various skill levels, fee and risk budgets), reflecting various investment objectives

Core Allocation (%)	Mgmt. Fee (%)	Skill Level							
		40	50	52	54	56	58	60	70
Panel A. Tracking Error (%)									
100	0.10	0.68	0.68	0.68	0.68	0.68	0.68	0.68	0.68
90	0.15	0.63	0.66	0.67	0.65	0.66	0.68	0.68	0.74
80	0.20	0.68	0.75	0.75	0.73	0.74	0.78	0.80	0.92
70	0.25	0.81	0.90	0.90	0.89	0.90	0.95	0.98	1.14
60	0.30	1.01	1.10	1.08	1.08	1.08	1.16	1.19	1.37
50	0.35	1.24	1.32	1.28	1.30	1.29	1.38	1.41	1.61
40	0.40	1.49	1.56	1.49	1.53	1.51	1.60	1.64	1.84
30	0.45	1.76	1.80	1.71	1.76	1.73	1.83	1.87	2.07
20	0.50	2.04	2.04	1.94	1.99	1.95	2.06	2.10	2.29
10	0.55	2.33	2.29	2.16	2.23	2.17	2.29	2.32	2.51
0	0.60	2.64	2.55	2.39	2.47	2.39	2.51	2.54	2.71
Panel B. After-tax Excess Returns (%)									
100	0.10	0.22	0.22	0.22	0.22	0.22	0.22	0.22	0.22
90	0.15	-0.17	0.09	0.14	0.19	0.28	0.34	0.41	0.82
80	0.20	-0.58	-0.05	0.06	0.16	0.32	0.45	0.59	1.37
70	0.25	-0.99	-0.19	-0.03	0.12	0.37	0.55	0.76	1.88
60	0.30	-1.42	-0.33	-0.11	0.09	0.41	0.65	0.92	2.36
50	0.35	-1.87	-0.48	-0.20	0.05	0.45	0.75	1.08	2.81
40	0.40	-2.32	-0.62	-0.29	0.01	0.48	0.84	1.23	3.24
30	0.45	-2.79	-0.77	-0.38	-0.03	0.52	0.93	1.38	3.65
20	0.50	-3.28	-0.91	-0.46	-0.07	0.56	1.02	1.52	4.04
10	0.55	-3.78	-1.05	-0.55	-0.11	0.59	1.10	1.66	4.41
0	0.60	-4.31	-1.20	-0.64	-0.15	0.63	1.19	1.80	4.76
Panel C. After-tax Information Ratio									
100	0.10	0.33	0.33	0.33	0.33	0.33	0.33	0.33	0.33
90	0.15	-0.28	0.13	0.21	0.30	0.42	0.50	0.61	1.11
80	0.20	-0.88	-0.08	0.07	0.21	0.43	0.57	0.74	1.50
70	0.25	-1.26	-0.22	-0.04	0.13	0.40	0.58	0.78	1.66
60	0.30	-1.44	-0.31	-0.12	0.07	0.37	0.56	0.78	1.73
50	0.35	-1.54	-0.37	-0.17	0.03	0.34	0.54	0.77	1.76
40	0.40	-1.58	-0.41	-0.21	0.00	0.32	0.52	0.75	1.77
30	0.45	-1.61	-0.43	-0.23	-0.02	0.30	0.51	0.74	1.77
20	0.50	-1.63	-0.45	-0.25	-0.04	0.28	0.49	0.73	1.77
10	0.55	-1.64	-0.47	-0.27	-0.05	0.27	0.48	0.72	1.77
0	0.60	-1.65	-0.48	-0.28	-0.06	0.26	0.47	0.72	1.76
Panel D. Impact of Taxes (After-tax Excess Returns - Pre-tax Excess Returns) (%)									
100	0.10	0.32	0.32	0.32	0.32	0.32	0.32	0.32	0.32
90	0.15	0.30	0.26	0.25	0.25	0.24	0.23	0.22	0.16
80	0.20	0.28	0.20	0.19	0.18	0.16	0.14	0.12	0.01
70	0.25	0.26	0.15	0.12	0.11	0.08	0.05	0.02	-0.13
60	0.30	0.23	0.09	0.06	0.04	0.00	-0.04	-0.08	-0.26
50	0.35	0.21	0.03	-0.01	-0.03	-0.08	-0.12	-0.17	-0.39
40	0.40	0.18	-0.03	-0.07	-0.10	-0.15	-0.20	-0.26	-0.50
30	0.45	0.15	-0.09	-0.14	-0.17	-0.23	-0.28	-0.35	-0.61
20	0.50	0.12	-0.15	-0.20	-0.24	-0.30	-0.36	-0.43	-0.71
10	0.55	0.09	-0.21	-0.27	-0.30	-0.37	-0.44	-0.51	-0.81
0	0.60	0.05	-0.27	-0.33	-0.37	-0.44	-0.51	-0.59	-0.91

Source: Parametric. Simulated returns from 1 January 2008 to 31 December 2017, annualised, net of fees and transaction costs. Simulated performance is hypothetical and does not reflect the results of any investor. All investments are subject to risk of loss. See Disclosures for additional information.

that satisfies the initial annual risk and fee budgets of our hypothetical superannuation fund:

- Material positive excess alpha
- Material information ratio
- 2% tracking error or less
- Fees of 50 basis points or less

The difference is that the specific conviction around active skill level in the satellites is now relaxed. To state this bluntly: if our hypothetical fund can't confidently express its view of active skill, rather than settling on its best guess (a number most models won't probability-weight) it may prefer at the design stage to rove within the sweet spot of shaded combinations, which cover an array of skill levels but still satisfy its risk and fee constraints.

Within this sweet spot, the dark-grey-shaded combinations represent the allocation to core (versus satellites) that maximise the after-tax excess returns for a range of possible skill levels. Deemphasising risk considerations makes any of these potentially optimal allocations to core. The more risk considerations start to drive portfolio design, the more the fund looks to the light-grey-shaded combinations, which offer better after-tax information ratios than the dark-grey-shaded combinations. The light-grey-shaded combinations represent higher allocations to core because the core portfolio helps reduce tracking error and negate the impact of taxes on performance. Finally, the bolded combinations show our hypothetical fund's optimal blend of core and satellite at 80% allocation to core (at skill level 56) as discussed previously.

Figure 4 is like looking at Figures 2 and 3 (where we scientifically plotted excess risk-return possibilities) through a much wider lens to help funds truly assess the trade-offs. For our hypothetical fund, stakeholders could reasonably ask why we would not seek to maximise our after-tax return for a given skill level. Figure 4 helps us answer this question.

First, to maximise after-tax excess returns requires higher tracking error. Within our portfolio constraints, this would require moving to a 20% allocation to core and pushing annual tracking error up to 1.95%. Under a 2-standard-deviation event, expected excess returns blow out from a range of -1.16% to 1.81% to a range of -3.34% to 4.46% each year. Thus, pinpointing an optimal combination of 80% core and 20% satellites has significantly reduced the variability of portfolio outcomes for our hypothetical fund at the expense of just 24 basis points in after-tax excess returns per annum.

Second, our hypothetical fund believes in portfolio efficiencies, and the 80% allocation to the core portfolio maximises the after-tax information ratio at skill level 56. This, by definition, is the most efficient portfolio our hypothetical fund could create if we believe in this level of skill. Tax impacts on our fund's optimal portfolio (Panel D) provide the final balancing act between the trade-offs we are making. Purely containing tax leakage would see the portfolio move to 100% core, but this would result in a 10-basis-point reduction in after-tax excess returns for only a 6-basis-point reduction in tracking error each year. Conversely, maximising after-tax excess returns by allocating 20% to core results in a negative tax impact (-30 basis points) that is, for our hypothetical fund, an unacceptably large proportion (35%) of the excess returns being generated each year. We are comfortable that our hypothetical fund's optimal portfolio is expected to deliver a moderate, positive performance contribution from tax without forgoing potential excess returns and without taking undue active risk.

Figure 4, while complex, is powerful in showing our hypothetical fund (and any fund using this kind of approach) what other portfolio options are back on the table if conviction in active management is recognised as more fluid (perhaps as much art as science) and if fee and risk targets are not set as hard constraints. In our scenario, our hypothetical fund cannot avoid an allocation to a passive core of at least 20% to remain in its fee and tracking-error budget. But, arguably, this allocation becomes harder to justify as the fund's confidence in its active programme increases.

At skill level 70 (where the active satellites get their stock bets right an impressive 70% of the time), our simulations suggest the fund would be forgoing up to 152 basis points in after-tax excess returns every year just to remain within its fee and risk budgets. A more sensible approach in this scenario, we suggest, would be to relax our hypothetical fund's annual fee budget from 50 to 60 basis points and its annual tracking error budget from 2% to 2.5%–2.75%. This would permit the fund to contemplate a core allocation below 20% or no passive allocation at all, which may sit well with a high, unwavering conviction in active management.

Ultimately, we believe it is healthy for a superannuation fund to approach scientifically the task of optimising a blend of equity portfolio ingredients, solving for a specific set of fund objectives and constraints. But it may be unhealthy for a fund to finish there without stepping back to survey the broader sweet spot of equity portfolio combinations available with a little more flexibility.

Conclusion

The importance of multi-manager equity portfolio design for superannuation funds hardly needs stating. But in this paper we have attempted to address the imbalance between research on the particular ingredients available and research on how these ingredients combine in a solution that a fund considers optimal. Far less research exists on the latter, especially when the task is approached from a real-world perspective that must navigate risk and fee budgets, peer sensitivity, fluid conviction around active management and, of course, day-to-day implementation frictions like taxes and transaction costs.

We established a hypothetical superannuation fund with realistic investment beliefs, preferences and sensitivities and showed how this fund could determine an optimal blend of ingredients available by assessing different core-satellite mixes plotted along a risk-return frontier. Despite our observation that optimality is in the eye of the beholder, we identified some general principles: The stronger the fund's sensitivities are to the real-world issues of peer and benchmark risk, investment management fees, taxes and transaction costs, the higher the fund's allocation to a passive core in an optimal blend. A high conviction in the skills of active managers (and the fund's skill in choosing such managers) can counteract this effect; however, if the fund's active-management conviction is flagging as well, the optimal allocation to a passive core will be higher still. High conviction in active management also creates something of a paradox unless the fund wants to ignore the impact of taxes (which we do not advocate): It strengthens the case for a tax-managed core allocation to provide some damage control over the tax impact of active management. Otherwise only a resolutely pre-tax-focused fund could see paying away a sizable portion (for our hypothetical fund, over 40%) of the fund's excess returns in tax as optimal.

We suggest that, before committing to a baseline 'scientific' de-

termination of what is optimal, funds step back to consider what portfolio design possibilities lay just a little beyond the limits of their specifications. This is where the ‘art’ comes in. We would like to see analyses like ours sponsor good discussions within superannuation funds about what possibilities are simply off the table because of hard constraints imposed on equity portfolio design. Mental strictures that could be challenged include:

- An unwillingness to revisit alpha expectations
- A need to consider the trade-offs between uncertain alpha and controllable costs like tax, brokerage fees and commissions
- Targeting headline fee outcomes rather than net outcomes for members
- The appropriateness of tracking error as a definition of risk
- The relevance of a fund’s standing in its peer universe

The power in acknowledging that what is optimal could sit in a broader sweet spot of possibilities rather than one specific combination of equity portfolio ingredients is obvious: It is a way to advance the science of equity portfolio design beyond the realms of specialised (and valuable) quantitative techniques and into the heartland of a superannuation fund’s mission—and what really creates a best-fit, enduring, real-world equity portfolio for its members. **FS**