British Actuarial Journal, Vol. 16, part 2, pp. 441–467. © Institute and Faculty of Actuaries 2011
doi:10.1017/S1357321711000171

Asset liability management for individual households

Abstract of the London Discussion

[Institute of Actuaries, 22 February 2010]

Contact
Professor Michael Dempster, MA PhD HonFIA MS FIMA. E-mail: mahd2@cam.ac.uk

This abstract relates to the following paper:
Dempster, M.A.H. & Medova, E.A. Asset liability management for individual households. *British Actuarial Journal*, doi:10.1017/S135732171100016X

The Chairman (Mr R. T. G. Hails, F.I.A.): The main part of this evening is to invite Professor Michael Dempster and Dr. Elena Medova to give the presentation on the subject for this evening, asset liability management for individual households.

Professor M. A. H. Dempster, Hon F.I.A. (introducing the paper): I would like to start by thanking, first of all, the President, Nigel Masters, for setting up this session, and then Mr Hails for chairing it, since the President is not able to be here this evening. I also wish to thank the scrutineers who gave us sharp and interesting questions on the first version of the paper, which has been presented in several different places, so we have benefited from many helpful comments.

We are very pleased to have here with us tonight Richard Foster, from Morgan Stanley, because this project really originated from his ideas. I should also like to thank a large number of staff and students who were involved in the project: staff of Cambridge Systems Associates and students of Cambridge University who helped with various pieces of the complex process of developing this model, solution and system.

Our presentation this evening is going to say something about the background of personal finance and then make some statements about the theory and methodology behind the *i*ALM system. My remarks will be short; it is a very complex product but I will try and give you an overview at least. We use the methodology of *dynamic stochastic programming* (DSP), which comes from operations research. DSP is both a modelling and a solution method. It has been widely used for institutional asset liability management, and the idea of this work is to apply this technology to advising *individual* households.

The system developed represents a new paradigm, we feel, because no existing theory provides comprehensive practical answers for financial planning particularly over long horizons. It also may be used in an interactive mode to do what-if exploration of various decisions such as retirement date, when to buy a house, or purchase an annuity.

First, I will describe the meta-model with its simulation and optimisation parts. However, the system is driven by inputs from the individual household. It does not, however, need input like risk attitude derived through questionnaires or other external evaluations.

https://doi.org/10.1017/S1357321711000171 Published online by Cambridge University Press

I will try to give you an overview, and then I will turn over to Dr. Medova. The purpose of her presentation is to show you the richness of the output that comes from our approach to financial planning and what can be learnt by a household from looking at this output.

Let us begin by outlining the setting for all this. Prescott and his colleagues, as quoted in the Campbell and Viceira (2002) book on long-term asset management for individuals, refer to the *asset allocation puzzle*. This puzzle basically stems from the fact that financial advisers have ignored the theory that has been developed over the past three or four decades in financial economics, starting with Paul Samuelson, Robert Merton, Robert Lucas and others. A common practice used instead by financial advisers is that the equity fraction you should have in your portfolio should be 100 minus your age, or, at the top of the last boom, 110 minus your age. This idea is a simple rule-of-thumb due to John Bogle of the mutual fund division of Vanguard.

In the latest but one issue of the *Journal of Portfolio Management* (2009), Daniel Kahneman talked about the myth of risk attitudes. I will just quote this: "To understand an individual's complex attitudes towards risk we must know both the size of the loss that may destabilise them as well as the amount they are willing to put in play for a chance to achieve larger gains." He went on further to say the following: "Temporary perspectives may be too narrow for the purpose of wealth management… The theories (utility theory and its behavioural alternatives) assume that individuals correctly anticipate their reaction to possible outcomes and incorporate valid emotional prediction into their investment decisions. In fact, people are poor forecasters of their future emotions and future tastes – they need help in this task – and I believe that one of the responsibilities of financial advisers should be to provide that help."

So what we are speaking about tonight is a tool that provides the kind of help sought by Kahneman. The system, which we call *individual asset liability management*, *i*ALM, is a decision support tool based on the theory of stochastic optimisation. It gives a solution to a dynamic multi-stage optimisation problem. This optimal solution gives an annual portfolio balance and full household cash flows from the current planning point to the death of the last principal of the household, which we assume to be one or two persons.

As input to the optimisation we have uncertain stochastic data whose future evolution we must model. We must then simulate this evolution with inflows and outflows of incomes, liabilities, investment returns, and so on.

From the point of view of actuarial practice, this is like dynamic financial analysis but with the optimisation built-in as opposed to changing parameters one at a time and comparing the outputs.

The ideas underlying this project are put together from both behavioural and classical finance and from decision theory. But the key paradigm shift of this approach is the idea that the system runs fast enough that one can interactively re-solve the whole prospective forward plan with changed assumptions. One can then compare the alternatives and choose the one that best suits future tastes and emotions. The system is, in short, an exploration tool.

What makes *dynamic stochastic programming* (DSP), different from *dynamic financial analysis*, with which many of you are familiar? The difficulty with DSP is we may have a wonderful scenario generator, but can we use it to generate scenario *trees*? To do this we run our simulator

https://doi.org/10.1017/S1357321711000171 Published online by Cambridge University Press

conditionally, moving out period by period on a branch of the tree. The tree representation is just a schema; a scenario in the tree represents a path of the multivariate data process. In our particular problem the data process will have market returns and economic variables, like price level, and events, like sickness, long-term care and death.

So, we must simulate all these interdependent processes punctuated by events and optimise over the scenarios in the tree *simultaneously*. To do this we will generate these complex scenarios with many variables: economic, financial and personal. Then we will lay them out as data in a very large-scale problem which we will make linear – actually piece-wise linear – so that we will get a very large-scale optimisation problem that can then be solved within a few minutes with industrial strength optimisation algorithms on current laptops.

What we are doing, practically, is maximising household consumption at each annual decision time subject to constraints, which are very important, such as risk, budget, cash flow balance, and so on. What we are trying to do is to obtain risk managed sustainable wealth maximisation across all years and all generated scenarios simultaneously. That reflects how the optimisation algorithm works.

But this is not easy because there are a lot of different things that have to go into it. What we are showing you, and what we have written in the paper, is not the last word. We acknowledge that, and I am sure that we will have some criticism in the discussion of various elements of the approach. But our main overall point is that we believe this is the first time that all this material has been brought together to solve this kind of problem.

We have to deal with data. We have to deal with econometric and actuarial modelling. We have to simulate the scenario trees which involves quite complex simulators because we have to deal simultaneously with events with economic and financial variables and also with various kinds of liabilities. In addition, we have to have a model which will give us optimally tailored portfolios, goal spending, cash flow balances, etc.

This is a dynamic optimisation model for assets and liabilities. We are trying to maximise risk managed goal spending and there will, of course, be many constraints.

This is only the start, it is only the technical side; there is also the institutional side. It involves formulating the problem properly, getting the right data, simulating it properly, and so on. Then of course the output is extremely complicated. For a young person, the output is a prospective plan over a very large number of years. So we need to be able to visualise the output decisions and we need to allow people to change things and to see the new effects on the plan.

What are the key modelling features of *i*ALM? Portfolio returns are determined by desirable consumption subject to existing and future liabilities. So what is driving all this are the spending preferences of the household which are entered into the system at the beginning.

I might mention that there is also a challenge here from a technical viewpoint. In institutional work, as those of you who work with dynamic financial analysis of corporate problems like capital allocation will know, typically you do a run. You usually do not quite like it, so you play with the inputs and run it again. If something fails, you tweak this or that parameter and try again. We have ourselves done this in much of our work in the past, both academic and commercial.

https://doi.org/10.1017/S1357321711000171 Published online by Cambridge University Press

But the *i*ALM system requires a new kind of paradigm. When the household's data has been entered into the system it has to run *first time*, or it has to tell the user why it did not run. We feel very proud of having been able to achieve that, I would not say with a 100% success rate, but certainly with over a 99% success rate.

So portfolio returns are what are going to generate household wealth, but portfolio risks are managed quite tightly. We constrain portfolio drawdown. This is something that people can understand. How much of a hit can you take to the value of your portfolio? The answer represents a real attitude to risk. Incidentally, if you look at some of the best risk surveys by leading banks this is one of the key questions that has been asked. Can you take a 5%, 10% 15%, 25%, or whatever loss? The answer, of course, is going to depend upon your wealth.

Also, people like to limit their portfolio asset holdings. They do not usually want to put everything in equities or everything in bonds or keep all wealth in cash.

So, what we have in the solution are optimal investment decisions under prospective market conditions and possible outcomes for a number of different choices, so that we can choose a suitable plan for the current year's investment, consumption and other financial decisions, and we also have a full prospective plan so that we can do all kinds of what-if analysis supported by the right kind of graphics.

To move to the behavioural side, decision *framing* has been discussed by Kahneman and Tversky and other behavioural theorists. They have pointed out that any decision problem may be formulated differently: with general or *broad* framing, or with a focus on at most one aspect, so called *narrow* framing. The decisions from different framings do not agree!

We have developed a new approach in which both types of decisions are consistent – simultaneously solving to both formulations, broad and narrow. To achieve this, narrow framing is done with the effects of decisions relative to the solution in the broad framing, which is to provide sustainable spending.

Figure 2 in the paper shows the famous value function of Kahneman and Tversky's Prospect Theory for narrow framing of the problem. You see that the status quo is the point of inflection: penalised on the downside, rewarded on the upside. This is effectively what we have extended for individual goals like annual spending before retirement; annual spending after retirement, sending children to private school or university, buying a boat, a new car, a house, and so on; all of which can be handled by this system.

The basic idea is that we are not going to have one reference point but *three*. One is the *minimum* level that absolutely has to be met. In the case of living expenses that is the poverty line. There is an upper *desirable* level and the middle level is an *acceptable* level. So, for a given goal, like year-to-year living expenses, you will have this three kink utility (or value) function. But in a given year you might have a whole lot of goals, so the actual utility function or value function will be piece-wise linear with a lot of pieces.

Paragraphs 3.2.5 and 3.2.6 in the paper show the situation mathematically. We do not really need to look at this too carefully except to note a couple of things. Basically, total utility, as defined, is essentially expected lifetime spending in present pounds, dollars, euros, or whatever. The objective

https://doi.org/10.1017/S1357321711000171 Published online by Cambridge University Press

is penalised, however, to make sure that taxes are minimised, and also that there is no excess borrowing (at penurious rates with annual rollover) which is essentially a way of modelling *bankruptcy*.

The consumption that figures in this objective will be with regard to two types of goals. Each individual goal can have a specific inflation rate. For example, it is well-known that medical expenses in North America are accelerating much faster than general inflation. Similarly private school fees tend to accelerate faster than general inflation in the UK, and so on. All these involve add-ons to a basic stochastic inflation rate, which we simulate.

We also allow real estate purchase with various types of mortgages, and different maturities. We have not fully handled annuities, except as defined purchase costs and inflows which are then subsequently indexed for inflation.

Wealth generation in this model is really coming from optimal resource allocation over a network of cash flows. In the paper, you will see the network for the main constraint, which is the cash balance constraint. But there are many constraints including: investment limits, portfolio position limits, portfolio drawdown constraints, initial holdings and cash flow constraints. All this is handled in a very large-scale stochastic linear programming problem.

As well as optimising taxes, the model also handles investment asset allocation between alternative tax-shielded accounts: 401(k) in the US and SIPPs and ISAs in the UK.

To summarise, the current model includes about 20 random processes that vary over the household's lifetime and around 200 mathematically formulated constraints per node of the scenario tree. The typical problem size depends on the length of the planning horizon, that is how old the household is, which gives how many more working years remain and how long the principals survive. The typical deterministic equivalent of the dynamic stochastic problem has around 3 million non-zero entries, which is a sizeable large-scale linear programme. But we can still solve such a problem in between two and 15 minutes, depending on the machines used (and 2 to 3 times faster on new multi-core machines). The output provides optimum values for many decision variables: spending, amount of savings, tax-efficient allocation between multiple portfolios, etc. Being able to do this at this speed allows an interactive process for analysing investment and savings alternatives for long-term financial planning.

On that note I am going to handover to Dr. Medova, who will show you an example.

Dr. E. Medova (Cambridge Judge Business School): For testing our system we used publicly available data, particularly from the Financial Times Money Section (from 2004 to 2006). Each week the FT presented a description of a household followed by recommendations from a few professional financial advisers.

Here is a summary of a household described in the FT in July 2006. We consider this household 3 years later in January 2009 and will compare the financial plans obtained for it in 2006 and in 2009.

We have a timeline with the basic information about the Pimlott family. They are in their early 40s. They would like to retire when they are 65. They have two children. They are financially relatively well off and they own their house, which is totally paid for.

https://doi.org/10.1017/S1357321711000171 Published online by Cambridge University Press

Their main goals in life are to have a similar living standard after retirement to the one which they have when they are both working, and to pay private schooling fees for their two children.

Professor Dempster did not have time to explain many details of the model, so I would like to stress in discussing this example that the risk attitude of any individual is a very personal thing. It really depends where you start in life and on what you want from life. Since people often make poor judgements regarding their future prospects we start the analysis with a 'reality check'. This is a rough estimation of the individual requirements for returns from a household's financial assets and indicates their risk attitude. We consider what happens to this family if returns and inflation are constant from year to year (3% inflation with 3.5% saving rate from a retail bank) and we consider what would be the household financial balance at the expected date of death for the oldest member of the household. Then we calculate what kind of expected annual return their existing portfolio has to have at that date if all their liabilities are paid and they spend their total remaining wealth on their desired goals. Note that this 'reality check' is designed for households whose consumption is generally higher than earnings and returns from their savings account are insufficient for attainment of their life goals.

For this example, as we can see in Figure 9 in the paper, the household requires a return of 7.4% per annum for desirable expenditure and 4.8% p.a. for acceptable expenditure. Thus the family or their financial adviser must actively manage their wealth.

In response to our referees, I would like to add some comments regarding the cash flow simulations. We simulate forward all processes for market returns and also simulate the costs of liabilities. The main economic factor which influences the cost of living is inflation; the other specific liabilities are evaluated using a corresponding index, e.g. for valuing the cost of private education.

We also assume that salary has been indexed by a gross rate for different ages and made some assumptions about the cost of borrowing.

There are other assumptions in the model that are different for different countries. These are basically tax allowances and different income tax rates and are examples of what we call rule-based assumptions. Such assumptions are the same for every household in the UK but there would be different rules if they lived in the US.

Again, I would like to stress how portfolio risk management is implemented in *i*ALM, which is on a screen of input assumptions. Risk management is at two levels. One is specified by the risk tolerance regarding portfolio drawdown. The user can input a number which represents the household risk tolerance, e.g. 5% or 15%. This is the limit on drawdown for each simulated scenario in any year and for any asset. We also have another control of risk which represents the financial adviser's opinion on the absolute maximum and minimum holding allowable for a given asset class. In this implementation the user can change these according to the views of the individual household or their financial adviser.

We can now look at this individual family's personal data: their salaries, when they want retire, and all other kinds of information about their existing financial arrangements – pensions, insurance, mortgages, etc.

Another screen shows the life goal specifications. It is important to note that this information is a new feature of our holistic approach to financial and retirement planning. We recognise that it is

https://doi.org/10.1017/S1357321711000171 Published online by Cambridge University Press

difficult to put one number on spending for any future goal but, in general, people have some range of values for a particular consumption item in mind. For example, a *minimum* annual living expenditure may be deduced from the poverty line living standards of the country. An *acceptable* amount for living expenses can come from an analysis of everyday expenses. The *desirable* value is obviously an aspiration. This range of values is given in today's pounds, which will change of course with simulated annual inflation as the prospective problem unfolds. We believe everybody has some view on these monetary values for future goal spending, which depend on one's background, initial wealth and many other personal factors. For example, if you have a high standard of living before retirement, you will want to live at a similar level when you retire.

You can put priority on any goal. For this family a priority is to maintain a standard of living after retirement similar to their pre-retirement level and also they put high priority on giving their children a private education.

Another screen shows their current financial and real estate holdings.

Now we will look at the solution to the problem. As Professor Dempster said, *i*ALM provides values for many different decision variables. It is not just portfolio decisions; there are many other types of decision. In the output they are classified as decisions which relate to portfolios, wealth, life goals and cash flows of incomes and liabilities.

The most important are the current year implementable decisions: the portfolio allocations for taxable and tax-shielded accounts. All the rest of the decisions are indicative prospective decisions which can be used for the purpose of the 'what-if' analysis.

An output screen shows the recommendation for optimal allocation of the Pimlott's portfolio for this year. On the left side of the screen is what the allocation had been in the previous year – the family wealth was mostly held in cash. On the right side of the screen you see the allocation which is recommended by the system. Note how risk management has been implemented. We have some overall limits on asset holdings. For this example, the household cannot invest more than 25% in commodities.

What is most important and new is shown on the following screen. This shows a dynamic view of the prospective optimal expected evolution of the initial optimal portfolio over the household's lifetime.

You can see annual portfolios which are risky when both principals of the household are working, then that the portfolio composition has a large amount of cash at retirement.

One can investigate many other things with the *i*ALM system. You can see where money comes from and is spent for each year over a household's lifetime by looking at the screen called 'sources and uses of funds'. You also can see the previous representation of the prospective dynamic portfolio allocation normalised as a proportion of total value. For example, at the time of retirement, the allocation in the property index is only slightly more than 2%, compared with nearly 18% at the starting point for a 5% portfolio drawdown limit.

The overall message, pertinent to any household, is that it is very important to start saving from the beginning. The resulting optimum allocation over the household's lifetime between the portfolio on which the household is paying capital gains taxes and the tax-shielded portfolios, such as the ISA and SIPP accounts, are shown on the screens of *i*ALM solutions related to the 'wealth' decisions.

https://doi.org/10.1017/S1357321711000171 Published online by Cambridge University Press

This family accumulates money in their SIPP portfolio which is optimal from the point of view of taxes. Another tax-shielded account used is the ISA. These are the results derived when different cash flows are optimised. As Professor Dempster said, this is really a problem of optimal allocation of the household's resources over its lifetime. In a graph we can look at 'sources' and 'users' of funds. You can see that, prior to retirement, salary and income from the financial portfolio are the main contributors to income while, after retirement, the ISA and primarily the SIPP are the main sources of income.

You can see that the area for expenditure demonstrates that the household's requirement that consumption be 'smoothed' over life is satisfied, i.e. there is more or less similar expenditure for living before and after retirement.

An additional useful feature is the cash flow statement. For each year you can check if the balance of cash in-flow and out-flow is correct. You can also look at the statement regarding 'wealth' for best, medium, and worst case scenarios.

To summarise, what we have done in this representation of *i*ALM output is show many individual cash flows separately: financial wealth, living expenses, aggregated liabilities. We may also now look at income tax, capital gains tax and different kinds of borrowing. For the current example there is no borrowing on the house. There is some possibility of bankruptcy, here modelled as 'excess' borrowing but this is not much in this case so it is not serious bankruptcy but rather an emergency loan.

Again, just to give you a bit more insight into what we are modelling in this system, we can consider a few representative scenarios of wealth. The best case scenario is defined from the point of view of highest goal achievement. In the worst case scenario you can see why, when people live too long after retirement, their chances for goal achievement are small and all their accumulated wealth is spent on living expenses. In the median case scenario, the length of life of the oldest principal of the family is approximately predicted. One scenario of the many generated by the simulator shows substantial wealth attained over a long lifetime.

We can also display the main goals of the household. You can see a histogram of goal achievement over different scenarios. This family's acceptable value for annual living expenses was £84,000 a year and desirable was £89,000. They have a higher expectation for consumption than they can really afford. The average across scenarios for living expenses is only £66,500.

They do a little bit better after retirement. The expected value over the scenarios is just slightly over the acceptable value. These values and the presentation of the probabilities of goal achievements is the reason why we say that the whole process of financial advice can move to interactive discussion and consultation regarding a household's future life plans. You can also do another run of *i*ALM, for example with slightly reduced consumption to see how the model generates another option for this modified financial plan. It will not change significantly but it will make the solution less uncertain with the histograms of goal achievement more concentrated around acceptable values.

In anticipation of questions, I want to show one more slide which illustrates that if you look at returns from the assets which we used in this model, you can see from our optimum portfolio allocation that very strict risk constraints on drawdown (5%) are imposed. This is discussed in the paper for a different maximum level of drawdown (15%).

https://doi.org/10.1017/S1357321711000171 Published online by Cambridge University Press

I will finish with a summary and some conclusions. Using this decision support system for financial planning moves to a situation in which the user analyses his preferences and current state of finances over a consultation and devises an optimal risk managed plan according to current market forecasts and life event expectations.

Mr P. D. G. Tompkins, F.I.A. (opening the discussion): I thought this was a very interesting paper which obviously depends heavily on the parameters, construction and running of the model. I was very amused to see Figure 11 in the paper, because I have to confess that I have a graph looking very like that on my PC at home as I have done some modelling of this kind just to get some idea of my own budgeting and investment planning.

One of the things that I noticed was that in Figure 8 you have demonstrated three elements of spending. There is an "acceptable" level and a "desirable" level, which I think we can understand as, respectively, what you are basically going to be comfortable with and what you would like to be able to spend if you had a little bit more spare. But you also have a "minimum" there. I would be interested in some comments from yourself in due course or from others as to the extent to which the constraint, which presumably is a minimum that I would cut back to if times were hard, comes through in the dynamic programming.

We all know the way in which this tends to work in practice. If people see their assets going down – stock markets have fallen or whatever – there tends to be a reaction: "I need to tighten my belt a bit because things are not going so well for me with my savings or my investments." You see people reacting, cancelling a foreign holiday or whatever it might be. Of course, we know logically that we should not look at what happened last year to determine what we should be doing now, we should only look at the state of our assets and our savings and other expenditure at the moment.

So I would be interested in knowing how you deal with the conundrum that the logical thing to do is to look prospectively but the reality is that people's behavioural response will very often be in response to what has happened in the past.

I would also like to ask about the extent to which that minimum is incorporated as a constraint. You commented on trying to maximise total consumption over a lifetime. I believe, for most of us, the gap between the desirable expenditures and the minimum is probably quite large. It is that which gives us the buffer which enables us to invest in more risky investments because we can take that risk knowing that, if things do not turn out, we can tighten our belt. It is because of that that we can take the more optimistic approach of trying to get 7.4% returns, as you have indicated in paragraph 4.4.

I think you are wrong when you say there that: "if they put all their money in a savings account… they will be over £2 million in debt at the end of their life." I do not think they will be. If they put all their money in a savings account and it is not delivering the income that they need for those desirable goals, they will be forced to retrench to the minimum or less than the acceptable goals by cutting other expenditure or aspects of their lifetime spending. I would be interested in comments on that.

Mr M. G. White, F.I.A.: People need sensible financial advice, but it is difficult to formulate, let alone deliver. The demise of final salary pension schemes means that the need for advice is growing

https://doi.org/10.1017/S1357321711000171 Published online by Cambridge University Press

rather than the reverse. This demise potentially makes some actuarial resources available, but the cost of financial advice on a person to person basis is simply too great if it involves more than a few hours of an adviser's time.

To date, the commission system has made it possible for plenty of adviser reward to be extracted at the expense, but without the full understanding, of the retail client. I have found financial advisers to be extremely reluctant to reveal or admit to an hourly rate, preferring to levy a percentage of asset charge and describe it as a fee. To me this is illustrative of the nature of the market for advice and is not encouraging in terms of actuarial jobs in this area. But as a profession we ought to be able to contribute in some way.

Another way of extracting wealth from the client is to sell complicated solutions which lock them in and are difficult to unravel. Smell a rat if anything is described as 'sophisticated', 'specially tailored to the individual', 'modern', or, worst of all, is described as 'using the financial markets to manage risk'. Expect false comfort combined with wealth extraction.

If we go back to the founding of the profession well over a century ago, our roots are in vehicles set up to provide financial security. These were largely on a mutual basis and the motivation was a mixture of self-help and service to the community rather than to make profits for outside owners. And the tradition of public service does still exist within the profession, though today's pressures make this more and more difficult.

It is from our involvement in vehicles to help meet the challenges of life, death, and illness that the motto "certainty from uncertain things" or "certum ex incertis" surely derives. We should always be on our guard to avoid that motto being misinterpreted as a suggestion that we can predict the financial future better than anyone else.

Our training is intended to develop an objective approach to analysing information, and having analysed it, to delivering advice based on it. And for personal financial planning, this must surely include understanding the huge uncertainties involved, and not being misled by the experience of a short period. Choices of parameters, even choices of models, should give significant weight to the experience of the more distant past – and here I think we should be thinking 100 years and more. Unless you force your memory back to include periods, both good and dreadful, which are beyond the experience of anyone working or making financial decisions today, you may miss the essential message that the future may be nothing like the recent past. Tools such as those described in this paper today have a great power to influence people, and represent a danger as well as an opportunity.

I suspect that the greatest value of such tools to the end client is in illustrating the huge funnel of financial and other uncertainty that stretches ahead of all of us, perhaps, combined with forcing the client to face up to the difficult choices to be made, including the present and future emotional constraints.

But from that point on, I become increasingly sceptical. I know that the paper is, in large part, setting out to explain how a model of this type can work, but I can't help focusing on the assumptions going in. I might be being unfair here, but the language of the paper seems to me to be saying that there is a degree of assurance, or comfort, coming out of the model that is worrying. I'll illustrate these concerns in a moment with a few extracts from the paper.

https://doi.org/10.1017/S1357321711000171 Published online by Cambridge University Press

But first, the goal of investment must include obtaining a decent real return, after inflation, tax and expenses. If, by avoiding expensive products, avoiding fund-based fees, and avoiding paying for fund management in a negative sum world, you can enhance your returns by, perhaps, 1% or more per annum, the extra return that will give you over a lifetime will have a dramatic effect. Realistic assessment of the impact of expenses should be an integral part of financial planning.

In paragraph 2.3 and elsewhere, the paper refers to the use of "current market views" for the purpose of projections. So if the market falls and people become more pessimistic, we should use more pessimistic projections? From the market's new lower base? I am not sure about this. Look over long periods of market history and I think you will find a mean-reverting force over time. Having a control mechanism which gets more pessimistic when every one else does, and vice versa, is a prescription for getting poorer in the long term rather than the reverse.

The crisis of 2008–2009 is described as being induced by sub-prime mortgages. It could instead be thought of as the pricking of a bubble, which will happen eventually when asset prices get too high. And one of the reasons asset prices became high was that interest rates were low and borrowing easy. It would be helpful if a financial planning model were to recognise that times of high prices and high recent returns are dangerous, and not necessarily grounds for satisfaction.

In 3.1.10 we see that ten-year sample periods were taken from June 1997 to May 2009. As I have mentioned, I would prefer that the sample periods were taken from a span of 100 years or more. This might lead to a commonsense approach to financial planning, which is to "hope for the best but prepare for the worst." The last speaker referred to that in a sense.

Paragraph 4.3 does lean this way a little, by pointing out that many individuals overstate their earning and spending prospects.

In paragraph 4.4, we read "a target return on investment of 7.4% is achievable under normal market conditions." I cannot see how anyone can say with such apparent confidence that any particular return is achievable, let alone 7.4%. Perhaps there's an implicit assumption that current market conditions will persist. If, on the other hand, as 4.10 suggests, "achievable" is simply meant to mean "possible", with no claims about likelihood, I'm not sure how helpful the language is.

In 4.13, we read that, due to the lower predicted returns from their assets (in 2009, markets and the market mood both having fallen since 2007), our family must invest more aggressively to achieve their goals. Personally, I felt happier the more the market fell, but the prescription here does sound like encouraging a gambler back to the table to try and win back his losses.

However, that's much too negative a point to finish on. I do not want to belittle what the authors have achieved with their model. If a model can illustrate the uncertainties that lie ahead by showing a wide range of potential outcomes, whilst demonstrating the impact of changes in planned savings and consumption levels, it may be an extremely valuable tool. My concerns have simply been about how a model might be used, or misused, and on the assumptions and motivations of the user.

Mr J. D. Harsant, F.I.A.: I am here not because I am retired and I think I know something about retirement, but because I think that this is a very timely paper in that we are considering the state of financial advice given to individuals. I think that there is a lot to work on in this area.

https://doi.org/10.1017/S1357321711000171 Published online by Cambridge University Press

I am not familiar with the mathematical techniques involved. I learnt those 60 years ago and I have spent the rest of the time forgetting them.

In the last 15 years of my employment, I had my own firm, of necessity at the small end of the market, which was providing SIPPs and SASs to local IFAs on a wholesale basis – we did not deal directly with the public in that connection. I found that there are good and bad aspects of advice from IFAs.

I, personally, have had very good advice from an IFA about some of the legal aspects of structuring my affairs, which I paid for on a fee basis. I have found, in other areas, that IFAs are sadly lacking. I find them lacking because, firstly, they are, inevitably, commission-driven. You cannot dissociate advice given from the need to earn commission in order to pay the rent.

The other factor is that I have seen some deplorable lapses of understanding by IFAs of the financial implications of what they were doing. One case I have in mind, where obviously I cannot quote names, concerned a rich man. He was told by an IFA of a well-known company that he was going to have to pay £100,000 tax in respect of something he wanted to do. It was quite clear as an actuary that the insurance broker's IFAs involved, and the insurance company advising them, had advised him completely incorrectly because they had made all the wrong calculations. That is, I suppose, a blatant example; but all along the line the IFA structure in this country is not giving really good advice to individuals.

To conclude, I would say that what we have seen today is a very good tool for advising individuals within a certain section of the community. You are not dealing with the very rich; you are not dealing with the self-employed. The self-employed have many opportunities for accumulating capital which are not generally available to individuals in the public sector or in major companies.

Mr R. G. Thomas, F.I.A.: I have two general points on the paper, three suggestions for a slightly different approach based on heuristics, and finally two questions which I hope might be investigated with the methods in the paper.

My first general point is that I'm very pleased to see a paper which focuses on individuals. Much risk management activity in recent years has been unhelpful to individuals. It has led to risks being transferred away from institutions which can bear them to individuals who cannot. Mr White mentioned an example of this – the demise of final salary pensions. The systemic point I want to make here is that improved risk management for institutions does not have unlimited social benefits, and it is particularly unlikely to be beneficial if risk management for individuals is simply ignored, as it often has been in the past.

My second general point, and here I differ slightly from the authors, is that I question the aspiration in the paper for a scientific approach. That seems to mean a sophisticated optimisation methodology, which the individual then relies on as a black box. I would prefer to aim for simple heuristics, which keep you out of serious trouble, rather than for full optimisation. The way I would use the methods in the paper is to investigate which heuristics are good and which are bad.

Moving on to some possible heuristics, one which I think is useful is the idea of leverage on discretionary wealth. This is promoted by Jarrod Wilcox, who is referenced in the paper. The idea here is that you have a minimum reserve level of, say, 50% of your total wealth which you must not

https://doi.org/10.1017/S1357321711000171 Published online by Cambridge University Press

lose, and so the 50% above that is your discretionary wealth. So, if you invest all your funds in risky assets with no borrowing, you are not unleveraged, you are leveraged 2x on your discretionary wealth. I find that a helpful way of thinking about risk and leverage; I would use the methods in the paper to investigate that heuristic and test it and see if it could be improved further.

A second possible heuristic arises from optimal growth theory, which leads to the idea of maximising expected log return as a reasonable long-term objective. The heuristic here is the formula $E - \frac{1}{2}V^2$ as an approximation for expected logarithmic return. That may seem trivial, but the reason why I find it useful is that, if I am leveraged, the adjustment I need is to just multiply the expected return by leverage and the variance by leverage squared. You can do that in your head and it is very helpful for keeping out of trouble. There are many hedge fund strategies and structured products, split capital investment trusts and other investment propositions, which are just not sensible as soon as you do that check. Again, what I would do with the methods in the paper is to test how robust they are and to see if they can be improved further.

My third heuristic is simply to be suspicious of expert advice. Expert advice has a low value in finance. That's partly because of agency problems, which have already been talked about by earlier speakers, but it is also for structural reasons – an expert consensus is often self-negating in finance. In my view many people would do better if they bought fewer financial products and paid less rather than more attention to financial advice.

The idea of being suspicious of experts and not buying financial products is illustrated in the book I'm writing – title *Free Capital*, subtitle *the world of independent private investors*. The book is about people who have left their jobs in their 30s or 40s to be full-time investors. They generally haven't been highly paid, and they haven't inherited money; they have made money as investors. They have no income other than investments, so they have asset liability issues. On the asset side, they invest in shares and other direct assets. They don't buy structured products, they do not buy collective funds or hedge funds; they do not buy any financial products. On the liability side, they're all very cautious – they live modest lifestyles relative to their wealth, and they avoid inflexible commitments. That approach of not buying financial products and not paying any attention to advice, and being cautious, not trying to be too clever, is not so much a methodology, it's a mindset.

Now on to my two questions that I hope might be investigated with the methods in this paper. First, the paper touches on the allocation of assets to tax-advantaged accounts like ISAs and SIPPS on the one hand, and ordinary taxable accounts on the other hand. My question is what *type* of assets should be allocated to each type of account? There is some US literature on this and it suggests that bonds and high yield shares should go in the tax-advantaged accounts. The rationale for this is similar to the corporate finance rationale for investing the company pension fund in bonds to maximise the value of the tax shield. I'm not sure about that for ISAs and SIPPs. I think it is better to allocate your best share ideas to the tax-advantaged accounts even if they have no dividend yield. I tried to write my own paper about this, which was published in the British Tax Review, but it wasn't entirely satisfactory, because I had no knowledge of dynamic stochastic programming. I would like to see that question investigated rigorously using the methodology in this paper.

The second question I'd like to see investigated is relevant to some of the people in my forthcoming book, and it relates to the timing of charitable giving. Suppose that working with a fund of a few hundred thousand pounds increasing to a few million pounds, you have substantially outperformed market indices for many years. As your funds increase, you expect your advantage

https://doi.org/10.1017/S1357321711000171 Published online by Cambridge University Press

to erode. You're starting to think you have more money than you will ever need. You want to maximise the surplus you give to charity over your remaining lifetime, and you don't want to leave any charitable bequests. Ideally, you want the undertaker's cheque to bounce. When do you start giving, and at what rate? In the framework of this paper, you can think of charitable giving as a form of consumption. But it is not that simple, because giving away money effectively takes time, so you need a time budget; and if you spend more time on philanthropy, that will probably distract you from investing and further erode your investment advantage. The question depends on mortality, the rate at which investment outperformance erodes, and on the time taken to give money away effectively. I have no intuition at all for the answer but I would very much like to know what it is.

References

Thomas, R. G. (2008). Taxable and tax-advantaged portfolio management for UK personal investors. *British Tax Review*, **1**, 34–55.

Thomas, R. G. (2011). *Free capital: the world of independent private investors*. Harriman House, Petersfield.

Mr M. H. D. Kemp, F.I.A.: I should like to welcome this paper. Professor Dempster always comes up with very interesting and thought-provoking material.

The following anecdote may help Mr Thomas. Some years back, I was at a one-week executive training course for the asset management industry. The lecturer asked the question: is it best to invest tax advantaged funds in bonds or in equities? The audience was split down the middle. The lecturer clearly expected everyone to agree that it was best to invest these assets principally in bonds, because of the greater tax advantage accruing to such instruments. However, as you correctly highlight, if you expect to earn a greater overall return from equities (e.g. by clever investment management or by capturing an equity risk premium), then the optimal strategy may be to seek to grow the tax-advantaged pot as quickly as possible.

My main comment touches on some of the comments made by Mr White. Material like this paper is brilliant in terms of the level of detail that it provides, but ultimately the answers that come from the application of such techniques depend on the assumptions that go into them. I notice that the assumptions that the authors are adopting appear to be based merely on historic returns over the last ten years.

In the academic world, there is a concept called "certainty equivalence". This involves taking some historic returns and working out some assumptions to use based on these numbers. Either you think the assumptions are appropriate, in which case you give the assumptions 100% credibility, or you think they are wrong, in which case you bin them and maybe find a different period of time to look at, or come up with some other set of assumptions from somewhere else. But essentially the approach involves a binary decision. Either you like the assumptions or you do not.

I want to highlight that in the real world we do not have that level of certainty 'equivalence'. We just do not know what assumptions we should feed into a model like this with that level of certainty.

That brings me to some other topics. I am also currently writing a book. Mine is looking at extreme events and their impact on portfolio construction. I note that the authors have not taken into

https://doi.org/10.1017/S1357321711000171 Published online by Cambridge University Press

account models that are fat-tailed, even though such models appear to describe the world better than ones that do not exhibit such features. However, they do make the point that it would be possible to refine their framework to do so.

Another topic that I have found interesting to explore when writing my book has been the issue of 'risk' versus 'uncertainty'. In this context, 'risk' means behaviour that we can quantify and understand. The models described in this paper are ones that focus on this type of 'risk'. For example, they quantify risk characteristics using covariance matrices, etc.

However, this is not the only type of behaviour that we might be interested in. There is a quite different type of risk/uncertainty present within finance. This we might describe as 'inherent' uncertainty. In the economic literature this goes under the term Knightian uncertainty (after the economist Frank Knight). It involves situations where we have no idea what the future will hold. It is therefore not capable of being modelled, at least not using the sorts of statistical techniques used in this paper.

I think it would be helpful if the authors could try to explore in the next iteration of their paper how best to take into account the potential existence of inherent (i.e. Knightian) uncertainty.

Another way of making the same point is to note that, if you were clever enough, then you could work out in advance what answer the model would produce for any given set of input assumptions, even ones framed stochastically as here. Intrinsically, the answers are based on those assumptions. So, in cases where we really do not have much idea about what assumptions are correct, then clearly we will also struggle to know how reliable might be the results coming out of the model.

Mr P. E. Read, F.I.A.: I think, to me, what has come out most clearly in the presentation and in the discussion is that this paper seems to be modelled around the needs of a high net worth individual. Certainly the examples which have been used appear to me to represent very high net worth individuals, most of whom have consumption well below the level of their savings. They can save in almost any vehicle they like and will quite readily build up a considerable sum of money.

What concerns me is the proposal that we are going to have to deal with in 2012. We are going to have millions of people compulsorily invested in the new National Employment Savings Trust; they are going to be asked to contribute a relatively small proportion of their salaries, and their salaries are extremely low. I would love to see this model be extended to see how viable this whole concept of the NEST scheme is, in particular to test the accumulation of very small contributions earning relatively conservative rates of return – I do not believe that people with small savings will wish to invest terribly aggressively. And also to test over what period of time must investment be made to avoid falling into the trap of means-tested benefits ultimately subsuming everything that comes out of this so-called savings pot.

My view is that we are heading for a potential mis-selling scandal such as we have never seen before and I am not aware of any work going on to try to stress test that possibility. I think this kind of tool would be wonderful, if I could suggest that as the next stage.

Mr A. D. Smith (student): There is a special reason why I enjoyed this paper, and that is because there have been a fair number of papers discussed at the Institute using dynamic stochastic programming and utility functions for corporate *i*ALM models. And, rather boringly, there is a small cadre of us – myself, John Exley, and others – who will pop up and say, "This does not really

https://doi.org/10.1017/S1357321711000171 Published online by Cambridge University Press

make sense for corporate *i*ALM whose corporations do not have a utility function, and the right way to think about it is…". It would be okay, however, to apply these tools to individual investments; but, of course, that is not what actuaries do. So a lot of people have been wandering around with very sophisticated hammers. I'm delighted that Professor Dempster and Dr. Medova have come up with the right nail for the hammer that many of us have enjoyed developing and building.

Another thing I like about this paper is the degree of granularity of modelling cash flows. I have seen a lot of stochastic programming-type models which assume that your consumption goes up at a certain rate per annum without really fitting it around an individual's priorities, whereas I am nearly in the situation of this family called Pimlott. I have four children instead of two, but many of the considerations are familiar to me. For my children their dates going to and from school and university are largely set by now since they are already born. Therefore I know when they will be 18, and so on. Modelling such cash flows discretely is useful and particularly having the discrete output. So, with this strategy what is the chance of having to pull a child out-of-school or cut back? That seems to me to be a very useful granular level of output, although I entirely understand the concerns a previous speaker has just expressed that it may be inaccessible to the vast majority of people because of the expense of constructing the modelling.

These all seem to me to be very good things about the paper. It would be interesting to add to those other decisions like the increasing flexibility that we have to work harder or less hard, to work longer or less long, having to vary the retirement date, etc.

I have a nagging doubt that some of the things that the model tries to predict are rather less predictable. I give an example: tax rates. Tax rates tend to change. We have public finances that are not in the greatest shape at the moment. It is hard to see that tax rates are not going to rise. The fundamental decisions about pensions saving depend upon the difference between the tax rates after you retire compared with the tax rates when you are saving.

Most of us have probably heard the received wisdom that you are higher rate taxpayers now but once you have retired you will be paying at a lower rate and therefore you should save. Actually, we have just lost the tax relief at the higher rate for saving and I think it is entirely likely that after I am retired we will still be trying to plug a hole in the Government's bucket and I will be one of the people who, whichever Chancellor is in place, will have to help fill the hole.

The trouble is that that decision rests on a knife-edge. If the rate of tax increases, then saving is completely stupid. If the rate of tax decreases, then you should fill your boots. That is certainly a concern. It is not a criticism of the model that the authors have built; but it is certainly something to bear in mind in terms of understanding the sensitivities.

Where I have bigger concerns is about what this model is doing from the point of view of calibrating to individuals going for advice.

So far as I can see you could do one of two things in terms of framing an individual's preferences. You can start off by saying all individuals are fundamentally rational and therefore we will try to calibrate our utility functions to how people actually behave. But then if you do that properly and you do your optimisation, all you get out of the optimisation is the behaviour that you observed in the first place. If you do not get that then people were not acting rationally and therefore it does not make sense to calibrate utility functions as to how people behave.

https://doi.org/10.1017/S1357321711000171 Published online by Cambridge University Press

The alternative is to go down the Kahneman and Tversky route, which is based not so much on observing actual investor behaviour but on people standing on street corners with clipboards and asking you, "Would you rather have £1 million with a probability of 1% or £1000 for certain?" They are all rather hypothetical questions. Then you try to extrapolate from that to build a picture of how individual preferences should behave. The difficulty with that approach is that Kahneman and Tversky have written quite a lot, and a fair amount of what they have written shows that the preferences that people express when faced with market researchers with clipboards are not consistent.

The danger is that you are coding-in irrational behaviour and then just extrapolating that or extending that irrational set of preferences into the future. So it is not obvious to me exactly what service is being provided by calibrating this utility function to certain questions and then coming up with an optimisation.

I have to say that there is one area where I do part company with Mr Thomas. He was concerned about the black-box nature of the system, the mathematical difficulty of understanding dynamic programming and the obstacle that this places to them understanding or being prepared to use this technique. I happen to know that with us today we probably have the top three dynamic stochastic programming experts in the UK. We have two of them on the platform, and I have seen Professor Davis sitting in the audience. For them, and, perhaps, for some of us who have played around with these models a bit, it does not seem quite as black a box. Personally, I find this an attractive way to think about financial planning and would welcome the chance to apply it to my own situation. I come from a position of having spent a little bit of time trying to understand dynamic stochastic programming, and I do wonder how far that could be extended. That seems to me to be a rather significant barrier to consumers or IFAs getting the hang of this.

Mr K. Wesbroom, F.I.A.: I really like this approach because for me it opens up the possibility that there is life after death of the defined benefit scheme for the Actuarial Profession. There could be ways in which we can get involved in using traditional actuarial skills in a broader framework.

I would like to focus on what I think is probably the most important part, which is how this process is communicated to the individuals who are going to be the recipients of this advice.

You will have seen that our President-Elect has said that communication is going to be one of her key themes. I think it should be one of the key themes here as well: the extent to which you can realistically involve the individual in conversations about some parts of this process. I would not even talk to them about the black-box. Do you want to ask the individual what their view is on the covariance of equities and bonds, and all that sort of thing? They are completely incompetent in that area. Let us start with that as the hypothesis.

So what questions, what communication, are we actually going to engage people with to make sure that they appreciate the essential trade-offs that they will be asked to make in this process?

Communication of the output is equally as important. How are we going to get them to appreciate the total ambiguity of the outcomes against some of the background that, for example, Mr Smith so eloquently described?

https://doi.org/10.1017/S1357321711000171 Published online by Cambridge University Press

Robustness of the assumptions is also something upon which previous speakers have commented. One of the things that always used to irritate me about *i*ALM models is if you want to tweak the assumptions a little bit, the thing veers from one end of the scale to the other. That seems to be something that should be addressed in formulating the model: to see how sensitive the outputs are to some of these inputs, and to try to give the advice accordingly.

I agree with the previous speaker about modelling means-testing in the paper. If you can get that in your model, I will buy it from you tomorrow because I do not think anybody understands how means-testing works. So there is encouragement from me because I think this is one of the directions in which the Actuarial Profession needs to go in the post-DB world.

Mr Tompkins: The other point that I wanted to raise was that this is a paper discussing assets and liabilities. In one place in the paper you look at the issue of the capital assets owned by the household. So, for example, their real estate is an asset.

I would just like to question the distinction between spending and investment. Take this example. I might decide to save long-term by buying a chunk of equities or some bonds, some commercial property, or whatever, or some works of art. Works of art are an example of the kind of investments that people might make which gives them pleasure but also gives them an asset which can be redeemed at times when they are short of money. I wonder to what extent you may have thought about that issue which is one which is very real to a lot of people who might, for example, buy some high-quality furniture or some art, or some silver, knowing that in the fullness of time, should they need to do so, they would sell that, as opposed, perhaps, to granny's heirlooms. That is quite an important point, particularly with the types of people that you are talking about in the affluent end of the income and expenditure spectrum. For example: that painting might be what will pay for my long-term care, if I need it, when I will not be enjoying the painting so much.

Mr G. J. Clark, F.I.A.: Like Mr Wesbroom, and a number of other speakers, I very much welcome a possible future for actuaries in something that I think is very, very important.

My own theory, which nobody has quite articulated, is that aspects of this are piece-wise discontinuous, particularly, for example, the retirement age, where I might have a plan today to retire at 60 but actually I'm quite flexible about this. If I reach 54 and the kids have flown the nest, then I may decide I have "enough" assets. My view of "enough" will almost certainly change after the drudgery of commuting to London gets me down and the kids have gone. If I have not reached my then concept of "enough", then I may well carry on beyond 60. The one thing I can be (almost) certain of is that I will not retire at 60.

We have heard that tonight in a number of ways. School fees are an example. I will take a view as to whether I have enough assets to put my children into private education. However, once they are in private education, I would probably have a very high attachment to keeping them in private education. It seems to me that there are a number of discrete life events.

My final comment is that some of the contingent events, for example critical illness, can be insured. There is no point being 5% critically ill in a projection. Either I am critically ill and have to change my circumstances, or I am not. There is both an insurance decision and there is also a more fundamental decision about how to respond to changing circumstances. One of the projections

https://doi.org/10.1017/S1357321711000171 Published online by Cambridge University Press

seemed to suggest a bankruptcy of £500. There seems little point in going bankrupt for £500! I might as well go bankrupt for a few million pounds.

Professor Dempster: This was bad terminology to describe what can actually be described as an emergency loan, as Dr. Medova said.

Mr J. Hunter (a guest): I am a guest here. May I first just make a comment in defence of advisers? A commissioned salesman is going to sell you a product. You cannot possibly complain about that. If you want to complain about anything, you should complain about the regulation that allows a commissioned salesman to call himself an adviser, which has, of course, been going on certainly since 2001, when Adair Turner, who was the first leader of the FSA, put it high on his agenda of things to change. And still we are where we are.

This is a fascinating product and contains lots of really interesting stuff. And I think it is deeply wrong. The objective function is described as maximising future consumption. I do not think that is how individuals think: what individuals want to do is to minimise future pain. The key thing that a complex computer model should try to do is help an individual understand his own attitude to risk and to learn about his own individual response to things that may happen to his investments. To that extent, if I were creating this model, I would include only two investment choices: a cash fund and a risk fund. I would put all the computing effort into interaction with the individual, to help him to learn how he wants to handle the risk management process through his investment lifetime. I would say, in passing, I do not like the word "risk", it tends to mean something that it does not actually mean at all – uncertainty.

I am sure that there is the capacity in there somewhere to do that. I get the impression that this model is heavily stochastic-orientated but in the end it produces one answer. I think that is approaching the problem from the wrong end. There should be a large range of answers which the individual, through interaction with the model, learns to cope with in real life.

Finally, and on a different theme, one of the problems that I have as an individual is trying to prevent my wealth being extracted from me by people who are not providing the service for which I am paying. We know this is very expensive in terms of charges within financial products, so one application for model-builders to consider might be to help the individual evaluate whether the amount he has to pay is actually justified by the benefit he expects to get.

Mr Kemp: I would like to raise one other point and ask the authors to clarify how their model might incorporate effects relating to it.

I was asked over the weekend whether my sister-in-law should invest in a stakeholder pension for her children who are about nine and 11. If she did so then her children would only be able to access this money in maybe 50 or 60 years' time, depending on developments in retirement ages in the meantime.

This type of question raises the issue of how do we best take regard of the different utility we might place on something now versus the utility of something in the future. I think Mr Tompkins has raised this issue in several other guises.

It seems to me that it may be possible to incorporate such trade-offs in this kind of model, but it is not immediately obvious to me how. It involves factors that may be rather difficult to

https://doi.org/10.1017/S1357321711000171 Published online by Cambridge University Press

explain, but may be rather important in terms of impact on solutions to how we should actually invest these types of funds. I would very much value the authors' input as to how they would actually in practice take such effects into account and thus allow for the benefits of cash now versus rolled up cash in the future.

Mr D. I. W. Reynolds, F.I.A. (in a written contribution which was read to the meeting): I much regret that I am unable to attend the discussion of the estimable paper. It is very welcome for a number of reasons.

Personal financial planning has been, and remains, devoid of processes which are broad enough to consider the full range of life events over a lifetime and incorporate both financial and non-financial objectives. This paper starts to get close to doing just that. I say 'gets close' because I do have some concerns about how it can be implemented.

Perhaps, Chairman, you can ask attendees at this sessional meeting to volunteer to attend a day when the model's originators can put them through the advice process. I would be happy to be one such volunteer.

This would clearly not be a full test as the volunteers would be a select group from a non-representational group of the population. But they might be representative of those who would be prepared to pay for such an analysis and advice. It would certainly help see how the interaction between advised, adviser and the model worked.

I am pleased that the model only looks one year ahead for the recommended asset allocation. It ensures regular advice which is important but far too often not occurring in practice.

But this does throw up one of my concerns. How does the model take account of the time taken to readjust the portfolio and the cost of doing so? I presume such costs, and the carrying costs of SIPPs and ISAs, are built-in, though I could not find any reference to costs.

I particularly welcome the 'mean reversion' that the Ornstein-Uhlenbeck model provides. It was a key component of the Wilkie model that initiated much of the stochastic modelling that is now an inherent part of so much actuarial work and also regulatory risk-management requirements.

It is a pity that the FSA does not, in general, accept 'mean reversion'; hence requiring firms to assume a 40% fall in equity markets after such a fall has taken place!

It is only partly encouraging that this might disappear in that its Chairman, Lord Turner, has set out that contra-cyclicality is a vital part of a sound regulatory regime. This will, no doubt, require an assumption of 'mean-reversion' when markets are above the mean, assuming we ever know when that is, and require more capital in those circumstances. But I do not expect the regulator to be brave enough to be consistent and apply mean reversion when markets are low.

Let me conclude with one other minor criticism. The authors should have included in the bibliography the paper by our own Deborah Cooper, which discussed a longitudinal approach to the planning of individual household income, expenditure, borrowing and saving strategies.

https://doi.org/10.1017/S1357321711000171 Published online by Cambridge University Press

Finally I welcome the paper because it represents a flowering of the profession's reaching out to academia and welcoming Honorary Fellows from there into the profession. I look forward to reading the full discussion of the paper and to testing the model in practice.

Reference

Cooper, D. R. (2002). Family fortunes: a guide to saving for retirement. *British Actuarial Journal*, 8, 849–885.

Mr T. V. Boardman, F.I.A. (closing the discussion): Managing investment risk and reward during both the accumulation and the decumulation phases is rapidly becoming one of the critical consumer issues, particularly, as others have commented, since the state and employers are increasingly passing risk back to individuals. These individuals are taking significant investment risk, morbidity risk, and longevity risk. This point was well made by Mr Thomas.

The other critical consumer issue is that we are not getting people to save enough. We have not really touched on this tonight. I think this model should help a lot in terms of encouraging people to save at the right levels as well as in how they manage the savings.

The paper highlights the opportunities that are available arising from the use of decision support systems and the theory of stochastic optimisation. As we have heard, it will allow IFAs and consumers to modify data and to engage in an interactive process. The model runs and within about five minutes you have the answers, so you can engage in "what if" evaluations, and it will allow people to analyse alternatives better.

As we have heard, the model should enable consumers to better understand and make decisions about how much to save and how to invest their assets. In particular, as Mr Wesbroom mentioned, it can aid communication around the impact of different alternatives and what their potential outcomes might be. It will obviously help people to understand the lifestyle that they can afford through retirement; although the point made by Mr Tompkins and others that people can and are generally willing to adapt their lifestyles in a crisis is well made.

Picking up a point from Mr Hunter, I think there will be a need to train both consumers and IFAs to use the model well.

The modelling tool is clearly not going to be appropriate for everyone. I take an interest in behavioural economics. One of my bugbears is that many people assume that everybody wants to be an "econ", and all we have to do is give them more information and more choice and then they will be able to maximise their welfare. The reality, as Richard Thaler and Cass Sunstein in their book, *Nudge: improving decisions about health, wealth and happiness*, and other behavioural economists are telling us, is that most of us are "humans" and actually we do not always want to know all the ins and outs of decisions. So some people would like the black-box that other speakers have spoken about.

However some people, almost certainly a minority, will particularly like the model, have the skills to understand what is being presented to them and, crucially, understand the implications and what actions they can take. These are the people who will benefit the most from the model and its outputs.

https://doi.org/10.1017/S1357321711000171 Published online by Cambridge University Press

I just think we have to recognise that for some people the amount of information that is produced could also confuse and frighten. We need to think about these behaviour outcomes as well.

I have to say, and other speakers commented on this, that the Pimlotts are not a typical UK family. My guess would be that we are looking at no more than 10% or at most 20% of the UK population that would actually be able to afford to use the model or for whom it would be worth using such a tool. I do recognise however that, as others such as Mr Thomas have suggested, the model could be very useful in identifying useful simple heuristic 'rules of thumb' and, as Mr Read suggested, for analysing questions such as the value of saving given the means-tested benefits that are currently available in the UK.

One of the problems we have in the UK is that a lot of wealthy people have their assets in a DB pension scheme and in their housing. Compared to, say, the US, there are not as many people who have a large amount of saved wealth that needs to be personally managed. We have to recognise the differences between different countries.

In the UK, 90% of people have funds of fewer than £50,000 and 90% of them have disposable incomes of below £15,000. Even the mass affluent people are, in a sense, going to find that they may not be able to make as much use of this as, perhaps, the authors have hoped.

Having said that, I think the work on producing this model does point us in a good direction. There has been discussion tonight about the fact that actuaries, particularly DB actuaries, might want to move into advising individuals. This seems a sensible suggestion particularly given the likely negative impact of the Retail Distribution Review (RDR) on existing IFA numbers. If we end up with fewer IFAs than we would like then models like this could be very helpful.

For most pensioners in the UK, buying an annuity is still going to be their best option. They are likely to be able to fall back on the State to pay for their care and get means-tested benefits. For a large number of these pensioners, bequests tend to be random events caused by people leaving housing equity.

So, really, we are talking about the 10% of the population who fall into the affluent pensioner segment. Other speakers have touched on this. Investment strategy is far more complex for the affluent. We also have to recognise, picking up a point made by Mr Reynolds, that investment strategy throughout decumulation is actually much more complicated than accumulation. Pound cost averaging works really well in accumulation, particularly if there is mean reversion. You end up buying more units and they are worth more by the time you come to retire. But of course, when you are in decumulation, you have to sell units or assets to pay for your retirement income, even if we get mean reversion. The reality is that pound cost averaging works against the pensioner.

Unlike the mass market of pensioners, who have a significant proportion of their pension income inflation-protected by the State, the affluent pensioners need to worry about inflation, and that means investing in real assets and this is where the model comes into its own.

Owning a home is a very useful asset in retirement. It gives people somewhere to live. It provides them with reasonable protection against inflation. It also can be the source of funding for bequests and, above all, it allows people to take more investment risk.

https://doi.org/10.1017/S1357321711000171 Published online by Cambridge University Press

A lot of rational people should be adopting the types of models that have been presented today. Clearly, as others have said, the model, as it stands, is very comprehensive and does an awful lot of things. However despite its sophistication, many speakers have identified more options that we might need and want to build onto it. The model that the authors have presented is a good pointer of where we should be heading. There are enhancements, like the ones that Mr Clark raised about being more flexible around retirement age, particularly, more flexible with people phasing into retirement.

There are some issues around long-term care. The model talks about long-term care and points out that in the UK most people have their care funded by the NHS, whereas in reality the segment who are most likely to benefit from the model probably are not going to have their care funded by the NHS. So there are some enhancements that are needed for the UK.

From my perspective annuitisation is very important and I'll say something more on this in a moment.

I should like to pick up and expand on a couple of other themes. One is the role of insurance, which Mr Clark commented on. There are expenditure spikes in retirement. Some of them are uncertain and can have probabilities attached to their likelihood; so there could be suitable insurance solutions. Mr Clark pointed out the fact that, if the model allows for 5% of the cost if there is a 5% chance of increased expenditure, then this is not great for the one-in-20 that actually incur the cost. Therefore I think there is a role for insurance to be brought more into the optimisation process.

At the moment the model allows for insurance but it has to be, as I understand it, manually input rather than the model's rules determining when people should take insurance. As the model develops, I think we should move on from just investment asset allocation to thinking about where it is right for there to be insurance, particularly around the use of annuities to cover longevity risk, life insurance to provide bequests, healthcare costs, and so on.

From a personal point of view, I believe we need to think very hard about what is the role of annuitisation and build this into the model. Again, as I understand it, if people have annuities, the model takes those in as income. It does not look and address when an annuity should be bought. This is a key question I would ask the authors to consider.

In addition, I think at the moment, the model typically produces results based on when people are expected to die rather than building in a randomness about when they might die. This would be another useful enhancement.

My analysis has led me to conclude that as pensioners get older the benefits of purchasing an annuity become greater. Therefore I do not think it should be a question of if people should buy an annuity, rather it is a question of when. I am a supporter of annuitisation later in life. A male at 65 today has about a one-in-1000 chance of living to twice their life expectancy and life expectancy is, perhaps, 21 years. So at 65 are you going to live to 107? Probably not. But when you reach age 85, you have a one-in-ten chance of living to your life expectancy. In the same way that Mr Thomas suggests that we need to take account of the variability of returns on investment assets so, I believe, we may need to look at the variability of longevity to determine the optimal time to buy an annuity.

The other factor with regard to annuity purchase is that, as people get older, they need to move into bonds anyway. The mortality cross-subsidy from the annuity becomes much higher than any equity risk premium. So, this analysis leads to a point where we should really be thinking about

https://doi.org/10.1017/S1357321711000171 Published online by Cambridge University Press

when and how we should buy an annuity. Even people with a lot of money should think about purchasing an annuity because it will allow them to optimise inheritance planning, and reduce the volatility around he amount of money they are able to pass on to their heirs. There should be a role for annuitisation in most retirement income plans.

In terms of insurance, there are obviously consumer propositions that we would ideally like to have around long-term care and health insurance. At the moment these are quite difficult to provide, but I think there must be an opportunity here in the future when the State provision has been clarified.

There are two or three ideas on insurance that come from the United States. Things like people buying an advanced-life deferred annuity, an annuity that will provide your income from maybe age 90 or 95. A number of academics have recommended this in the United States. What makes this deferred annuity conceptually attractive is the fact that you can plan with much more certainty because you know how long the assets need to last, given that you know you have an annuity income stream starting at age 90 or 95 that will pay income for the remainder of your lifetime. Another attractive idea is for insurers and fund managers to provide income guarantees, similar to variable annuity living benefit guarantees, where people only get paid out if they live a long time and their investment performance has been lower than a predetermined return.

Finally, I should just like to return to the issue of attitude to risk. I share the views of a number of speakers who have commented on this tonight. It is very difficult for people to understand what their attitude to risk is, and they only really understand it when they have a negative experience.
The experience of the past two years will have changed many people's attitude to risk quite significantly. I have always wondered and worried about the way in which we ask people about their attitude to risk without them really having a good reference point. This model should help people to understand and simulate what could happen and the implications for them. This should help consumers.

Again, I did get a little bit worried when the authors explained that because it is difficult to understand people's attitude to risk, they use the model to determine what risks people should take. We appear to get to the 7.4% target return because that happens to be the rate of return that is required for the Pimlotts to be able to meet their requirements. I do not necessarily think that telling people they have to earn 7.4% is as informative to consumers as it would be to us in this room. Most people find investment risk and reward difficult to comprehend. However I do accept that anything which will help people to visualise the future will be helpful.

What tonight has brought out is the fact that pensioners do need advice and we need to find tools and ways of providing that advice. Therefore the development of this modelling tool must be a step in the right direction. However, we need to learn from the last two years that models are only models. There were some good comments from Mr Kemp and Mr Thomas around whether the ten-year look-back should be longer and whether we should take longer views. There were discussions, as always, on fat-tails, and so on. The discussion towards the end about risk and uncertainty was also very helpful.

I hope people have found these comments helpful and of some interest.

Reference

Thaler, R. H. & Sunstein, C. R. (2008). *Nudge: Improving decisions about health, wealth and happiness*. Yale University Press, New Haven and London.

https://doi.org/10.1017/S1357321711000171 Published online by Cambridge University Press

Dr. Medova (replying): I was not able to take notes of the discussion, so I will make only a few responses. For example, with regard to Mr Smith's contribution, I have to say that many things in the model are easily changeable. An example is the way tax is treated. If the government introduces different limits or different allowances, these are all parameters which are designed in a way that they can be changed from the user interface. The logic is very general and this also applies to portfolios which have some tax benefits.

Next, I make a very general response to the length of data with which we calibrated our model. I allow myself not to be very theoretical in this regard, but in general I do not believe that we know which periods of data should be used, because for different purposes you need different lengths of data. From this point of view the benefit of stochastic simulation is that it allows the possibility of updating your view on changing markets and easily re-running the model. We chose ten years rather than 100 years because it is a relatively long period, but can still capture the trend and the volatility of the data.

Another comment was whether advising annually is too frequent. In general, it is not necessary to have a new plan every year but it is important to have a new plan when something is happening in your life. This again is a question of risk and uncertainty. We are definitely not in control of our lives. If some event happens unexpectedly, you have to reassess everything. This is the time to introduce new data and update your plan. If nothing new is really happening, then you need not update your plan once a year, you need only do it when you think appropriate.

From our point of view, for example in constructing the model from the input data, choosing the big decision points prospectively is important. For individual household problems this so-called 'branching of the tree' should happen where, for example, people currently think there will be big outflows or inflows of money, like planning to buy a house in ten years' time. This date will then be used by the system to branch the scenario tree, otherwise there will be no branching at this date.

We tried to distinguish between uncertainty and risk in designing *i*ALM. Risk involves small parameter changes; uncertainty is something we cannot control. That is why we are introducing the interactive use of the system.

Professor Dempster (replying): In my view, in response to many of the questions that were dealt with in detail in the discussion, we should emphasise the point that Dr. Medova has just made: this system has been designed to be reused when major events happen, whether they are market events or personal events. Personal events are much easier to deal with. Regarding market events, I think that it is probably worthwhile giving you some background on the testing of this system, which was extensive in the US, and, perhaps, not yet quite so extensive in this country. The UK tests, as those in the US, spanned example households from the really super-rich to those not nearly as rich as the Pimlotts.

There was a related question, for example, in dealing with IFAs in this country, a number of whose clients we used to test this system anonymously. IFAs were very worried about market recalibration issues, and we had some interesting discussions of whether or not prospective returns should exactly reflect current market conditions or should be long run historical averages. This dichotomy has been raised several times in the discussion and is related to the cost of the advice being given. The right answer is a problem, but it is a problem regarding how the system is marketed and used.

There has also arisen the idea of using the system for lower income households. The system has been designed to be stripped down, so it can consider just the retirement decision, or annuitisation,

https://doi.org/10.1017/S1357321711000171 Published online by Cambridge University Press

or whatever. If anyone in the audience has a line to Government we would be very happy to consider this use. There are real questions about what the Government proposes at this point. Of course, these proposals may never come to pass but they are very worrying. We do feel that a tool like this would be very useful to examine the feasibility of these kinds of proposals, which is, I guess, what they remain at the moment.

I should like to say a little bit more about whether we can cope with real assets, insurance solutions, annuities, and so on. We confess that our knowledge of annuities – when to renew, purchase, and so on – is not extensive. However with a built-in annuity calculator, similar to something we have done for mortgages, these are decisions that could be considered by interactive use of the system. Similarly this can be done more easily regarding insurance, which we have done for the US system.

However, there is one thing that is a design feature of this system: its what-if use, which Dr. Medova stressed and I want to stress again. Its what-if use is supported by the fact that the code runs fast enough on laptops. Indeed, on current machinery it runs two to three times faster than the times we quoted earlier because earlier runs were on laptops that were two or three years old.

The main idea is that you can use the system to investigate different alternatives. That was one of the uses of the US model, actually to show the results with and without insurance, which are quite different in many cases, depending on the variations in insurance cost and household wealth. That would be the same with annuity purchase, given suitable annuity calculations. This is really the principal purpose of the *i*ALM system.

With regard to attitude to risk, we may have failed to convey the simplicity of the approach. I am tempted to say that we dressed up what we did with behavioural finance because our basic idea was that any client, super-rich to modest income, actually understands given goals like buying a house or retiring, and how much they are going to spend per annum on average. They understand what they would find an acceptable living standard, what they would like to have and what the minimum they could live on would be.

The system outputs tell the user the likelihood of achieving their goals, including the risks involved. Incidentally, I should make the point that goal priority setting is an extremely important part of how the solution comes out. For things that are absolute – for example, for parents that are absolutely determined to send their children to Eton – then that goal, the cost of sending their children to Eton, is in fact a minimum. The acceptable, desirable and minimum will then all be set to the minimum amount that they will actually need. Then that goal will essentially be the first thing that the system addresses.

In the hundreds of representative households that we have seen on both sides of the Atlantic, it is easier, because of the power of compound portfolio returns, to meet retirement style goals, like buying a yacht when you are 65, than it is to send your children to Eton next year. That is obvious, it seems to me, whatever your level of wealth, and it certainly comes out in the optimal solutions.

We do have a lot of confidence in the stability of these solutions but, of course, not in the parameter settings, and certainly not based on the length of time series that we used in the paper. Solution stability is a natural result of the stochastic approach, which people like Mr Smith, who has played with similar things, will know. When you have many, many alternative scenarios in problems, this

https://doi.org/10.1017/S1357321711000171 Published online by Cambridge University Press

stabilises everything so that optimal decision variables evolve in a smooth way. That is one of the advantages of stochastic solutions.

In summary it is therefore not quite fair to say that we are just making one recommendation with the *i*ALM system. Yes, for today; but all kinds of alternative scenarios can be investigated, from a prospective point of view, and as Dr. Medova said, for a given run of the system with given parameters.

Of course, we can always improve the parameters: they can be reset by user when there is a new estimation, etc.

We could, and, perhaps, should deal with fatter tailed distributions. In this respect I should point out that we are running a simulator that, although it is estimated on monthly data, we are actually running it over 12 steps before we need its output, because we dealing with prospective decisions only on an annual basis. Thus the annual asset return volatilities are somewhat higher than historical. As it is, the optimisation problems solved are very large and if you are dealing with a 25 year-old, as we did, particularly in the US, then they are exceedingly large problems. We can only solve them on an annual basis. From a behavioural viewpoint, that is often the way people think about their plans.

As to how easy the system output is to understand, various speakers have said that this may be easy for some people and not so easy for others. I would respond to this question by saying that a simplified version of it might be quite easy to understand. Basically, the simplified system could be used to investigate questions such as what you can afford to live on up to retirement, and what you can afford after retirement, and when you should retire, and buy an annuity or private health insurance.

I think that summarises our immediate responses.

The Chairman: All that remains for me is to express my own thanks for this paper and, I am sure, the thanks of all of us here, to the authors, the closer, and to all of those who participated in this discussion.

Professor Dempster (subsequently wrote): In our closing rejoinders Dr. Medova and I both ignored two important comments by Mr. Wesbroom. The first was his concern with ALM solution stability to minor parameter or structural tweaks. This is a bad feature of deterministic or unstable DFA simulation models which dynamic stochastic programming models do not share. In DSP models initial and subsequent portfolio decisions are robust to all the remaining scenarios in the model which unfold from the decision point in question. Mr Wesbroom's second, more general, concern was regarding 'means testing'. We share his concern as we find it difficult to understand even what the term means outside of a very specific policy context, and even then different implementations and individuals responsible will make actual practice different. However, we do feel that *i*ALM could potentially be of use both in setting policy and in informing implementation. It would be an exciting prospect to do this in conjunction with Mr. Wesbroom and others.

https://doi.org/10.1017/S1357321711000171 Published online by Cambridge University Press