Financial Conduct Authority

Two plus two makes five?

Survey evidence that investors overvalue structured deposits

Technical Appendix

March 2015

Occasional Paper No.9

Technical Appendix

TA1. Structured deposits

Design

The quantitative model of the expected FTSE100 performance assumed 7.5% (approximately 1.5% pa) expected profit margin, this choice was based on FCA analysis. The assumptions of the quantitative model and the expected issuer margin both constitute a ballpark estimate only and we note that our main results do not depend on these estimates being highly accurate.

The quantitative modelling used data from Totem Valuation, part of Mark-It Partners, on FTSE 100 equity options with strike dates up to, and just beyond, five years in advance and on risk-free rates matched to the maturity dates of these options. The data were quoted at market midpoint prices. Using the put-call parity relationship for puts and calls at triggers of 100%, including the discounted value of dividends, we backed out the expected dividend rate from option prices. And given the dividend rate and the risk-free rate, we worked out the implied volatility.

As the maturity dates were at fixed points in time, using a spline function to interpolate between dates, we backed out the risk-free rate, expected dividends, and implied volatility at exactly 1, 2, 3, 4 and 5 years in the future. We then simulated returns for the different products, choosing the percentage of the FTSE 100 payoff for the 'capped' and 'floor and cap' products and the level of the payoff for the 'Kickout' and 'Cliquet' products so that the implicit charge for the distribution and manufacture of the product is close to 7.5%. We simulated the expected return of the products and the chance that the product pays greater than a 0% nominal return or greater than the risk-free rate.

Complexity assessment

To score the structured deposits employed in this research we used a complexity scoring system designed by FCA's Complex Products Specialists Team specifically for structured deposits in 2013. This scoring system was designed after an extensive review of UK structured product market and aims to span a range between the simplest (score 1) and the most complicated products on offer (score 6).

The complexity score is assigned based on the presence of individual features and combinations of features. The scoring of the five structured deposits was done independently of the results of this research to ensure objectivity. The scores are provided in Table 1.

The key drivers for complexity within structured products identified are:

- **number** of features – the more features a product has, the more difficult it is to understand the impact of the combination of those features
- **complexity of a feature in isolation** – how the feature affects the performance or profile of the product; these can individually vary in complexity, and

- **interaction between multiple features** – some features may function relative to one other; others may be non-correlated.

The complexity scale is based on the baseline measure (the simplest product) as either a FTSE100 product with a binary output or a FTSE100 product with 100% participation. It is not possible for a product to have less features or be based on a more recognised index (the FTSE), our research demonstrated that although there were significant numbers of participants that did not understand the structure of the FTSE, these products were generally the most widely understood.

For each additional feature, there is a degree of corresponding increase to complexity (based on an assessment of the understandable/measurable effect to the potential return). For example, adding a cap to the above product will increase complexity but to a lesser degree than adding a set of path dependent payoffs.

Certain features are determined to be highly complex due to the unpredictable effect (from the point of view of the consumer) they have on the product performance. These include features such as synthetic indexes, which due to their 'unknown' nature, would prevent customers from understanding the impact they may have and their effect on any potential returns.

To establish the effect a feature (or combination of features) may have on complexity, two related exercises were carried out:

- each individual feature was rated on a scale of complexity in isolation (Table 1)

- a mechanism was developed to establish how combinations of features may affect the overall level of complexity (Table 2)

Table 1. Complexity of individual features of structured deposits

Feature	Complexity
CAP	Low
Common Index (Not FTSE100)	Low
Multiple Time Points	Low
Floor/Min	Low
Multiple other than 1	Low
Kick-out/Early exit feature	Medium
Multiple Common Indices	Medium
Clearly weighted basket of companies/commodities	Medium
Returns 'rollover ' if not achieved	Medium
Path Dependant Payoffs	Medium
Look back average	Medium
Synthetic Index	High
Volatility Overlay	High
Active Trend Analysis	High
Complex Clique	High

Table 2. Categories of complexity of structured deposits based on features and their combinations

Category	Description
1	Lowest complexity level – Base structure – FTSE100 with binary output at maturity or FTSE100 with 100% Participation
2	Base Structure + 1/2 low (Green) features
3	Base Structure with several low features or any medium (Amber) feature
4	Multiple simple and medium features; a consumer may be able to calculate approximate return with a high level of financial capability
5	Any complex feature or complex combinations of features; a consumer would be unable to calculate/understand the approximate return
6	Multiple or complex combinations of complex features; a consumer would be unable to calculate/understand the approximate return

TA2. Sample and removals

We removed two main groups of responses from our sample of 514 respondents: respondents who did not pay sufficient attention and respondents whose beliefs are extremely difficult to reconcile internally. Finally, we removed six observations for which our model failed to fit a lognormal distribution due to very high implied standard deviation.

First, we removed observations that indicated insufficient attention to the survey. We removed observations if the reading speed of the first product description as measured from the screen opening time to clicking the 'Next' button was higher than 850 words per minute (wpm). Research indicates values for skimming of plain text in the range of 500 to 900 words per minute. Muter and Maurutto (1991) find that laboratory participants skimmed text at 850 wpm from paper and 500 wpm from screen 'to grasp a general sense of the content or to retain only the main points'. Just and Carpenter (1987) document that speed readers achieved 600 to 700 wpm, but at a cost of comprehension. Our cut-off level is rather conservative to avoid removing potential quick readers.

We used the reading speed for the first product (structured deposit A) because all product descriptions were accessible again through one-click pop-up boxes on subsequent question screens, and it was only for the first product that the respondents would not have known that they will be able to access the description again. Our main findings are robust to removing participants who read faster than 500 wpm and faster than 1000 wpm. When the analysis is done using these cut-offs instead, the reported returns still exceed the implied returns by approximately the same amount (*overestimation is* 1.86% for 500 wpm, 1.81% for 1000 wpm; the difference seems to be caused by slightly lower implied returns).

We demonstrate that respondents removed due to short reading time of the structured deposit A description were significantly less likely to understand the examples of tools (probability questions on dice and cards – see Section 2.3) and the features of the products, and saw more warnings about implausible combinations of reported beliefs than the sample remaining after removals. They also reported more optimistic beliefs about the likely product returns and about the future FTSE performance, but had a slightly lower (no statistical difference) overestimation of expected product returns.

Second, we removed observations that indicated poor understanding of the survey tools. We removed respondents completely if the combinations of reported expectations of the future FTSE100 performance were strongly implausible: expected value higher than 6,500 but the

probability of FTSE growing is less than 40%; and expected value lower than 6,500 but the probability of FTSE growing is more than 60%.

To determine which combinations of reported expected return and probability are internally implausible, we compare the distributions implied with the distribution of historical FTSE100 five-year returns. In the first case, the expectation equal or higher than 6,500 and the probability of FTSE growth being 40% or less imply a probability of at least 25% that the FTSE100 return will be in its lowest 1% of historical distribution. This is highly unlikely.

In the second case, the expectation equal or lower than 6,500 and the probability that FTSE grows being 60% or more imply a probability of at least 99.9% that the FTSE100 return will between -1% pa and +1% pa, which has a likelihood of less than 9% based on historical distribution.

Overall, being removed is strongly associated with lower education, lower self-reported financial expertise, no previous experience with structured products, and poorer understanding of survey tools and products (Table 3). Removed respondents on average reported slightly lower expected returns and had slightly more optimistic expectations about the FTSE100 performance in five years. Being removed due to short reading time was almost uncorrelated with being removed due to implausible statements (correlation coefficient is 0.02).

Table 3. Respodents we removed had lower education, gave poorer quality responses, and somewhat differed in reported expectations

	Removals		Sample after removals		
VARIABLES	**N**	**mean**	**N**	**Mean**	**mean difference***
Demographics					
age of respondent	130	45.6	384	50.4	yes
household income £k	130	95.3	384	97.2	no
education (university)	130	0.59	384	0.68	yes
held structured product	130	0.77	384	0.85	yes
Understanding					
dice chance answer right	130	0.54	384	0.79	yes
card chance answer right	130	0.63	384	0.90	yes
dice exp. answer right	130	0.68	384	0.84	yes
no. of implausibility warnings	130	1.25	384	0.43	yes
Reported expectations					
reported exp return, %pa	130	4.48	384	4.15	yes
implied exp return, %pa	130	3.11	384	2.27	yes
return overestimation, %pa*	130	1.31	384	1.87	yes
exp value FTSE in 5 yrs, %pa	130	3.81	384	2.67	yes

*Indicates whether the mean of the cleaned sample is different from the mean of removals at 5% confidence level.

TA3. Distributions of responses

Figure 1. Investors' expectations of structured deposit returns

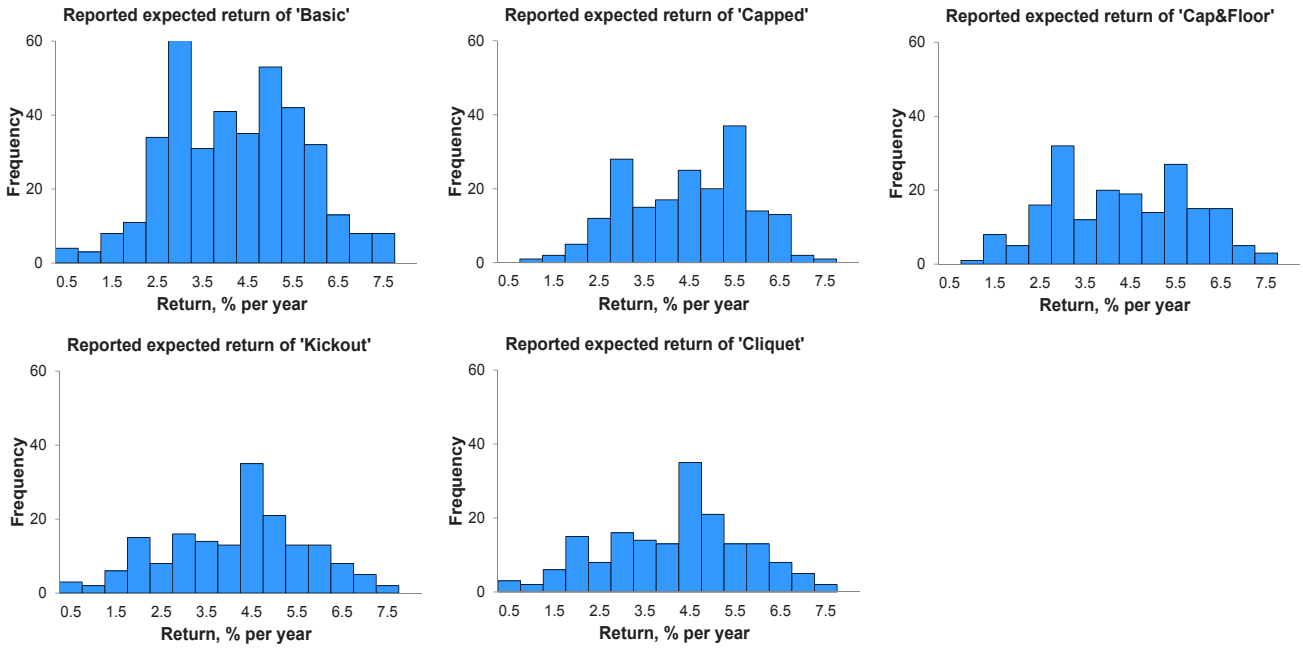

Figure 2. Investors' FTSE100 growth expectations converted to annual growth rates

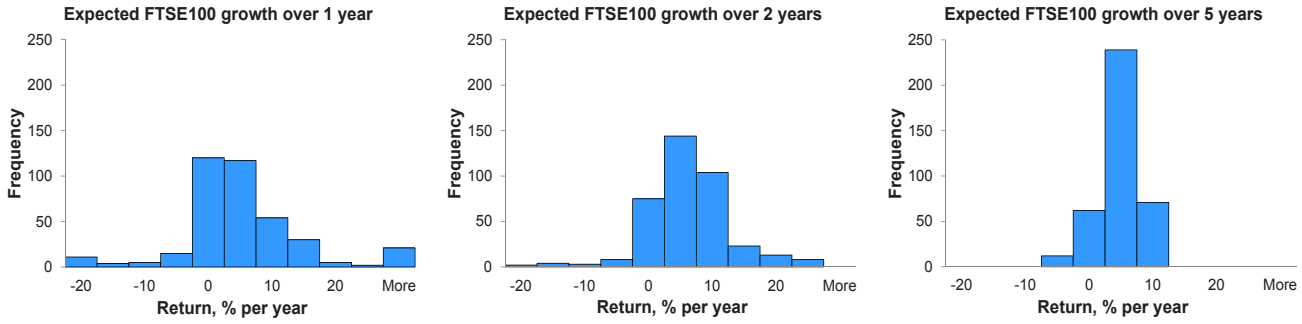

Figure 3. Investors' valuations of products in terms of risk-free rates that deliver equivalent value

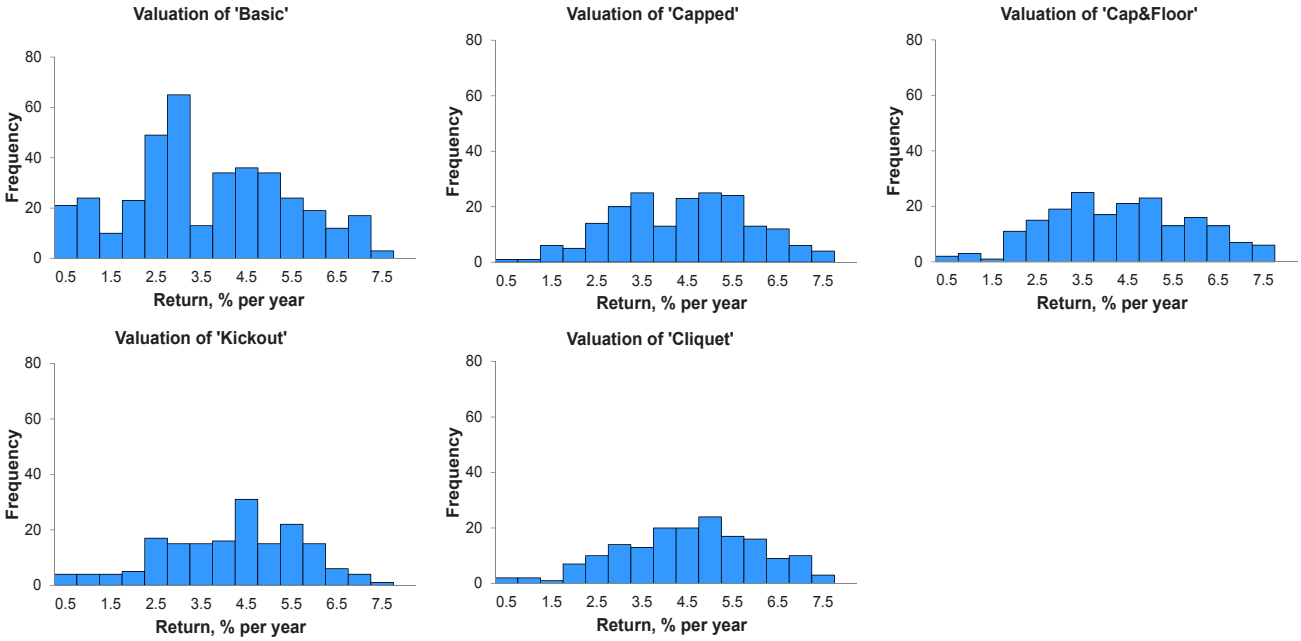

Figure 4. Random distribution of the decision tree responses is skewed towards low values

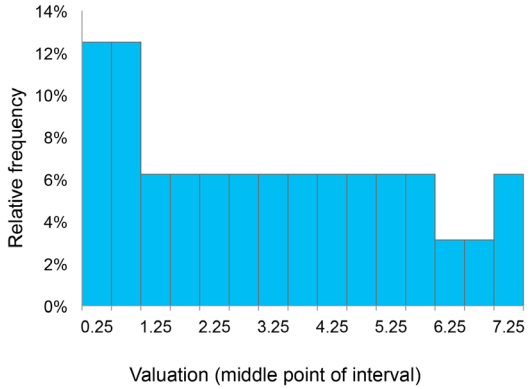

Figure 5. Actual distribution of the decision tree responses is significantly different from random and has a lot of mass in the upper range of values

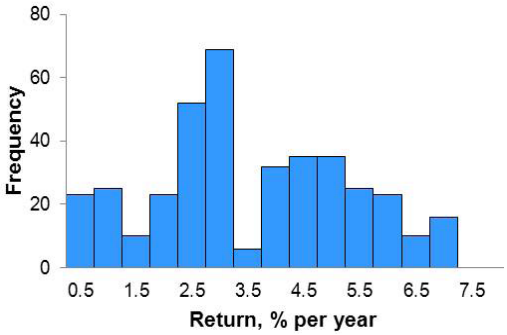

Figure 6. While most respondents did not change valuation after decision tree, those who changed adjustemt mostly upwards.

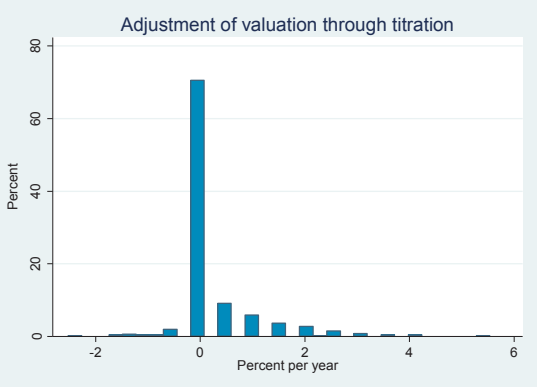

Figure 7: Average valuation in the Ladder tool is slightly above valuation in decision tree and higher than the implied product return

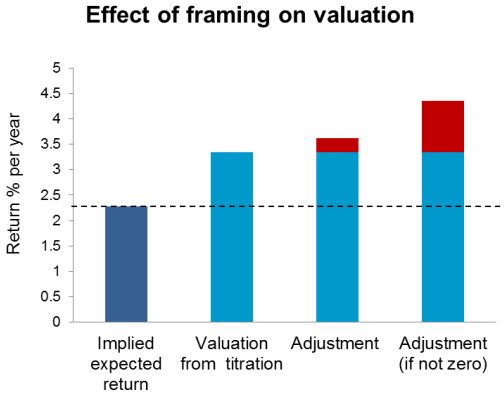

Note: the red parts at the top of the columns in RHS chart signify the adjustments of valuation.

TA4. Modelling implied expected returns

We used the individual beliefs about the future FTSE100 performance to model the individual probability distributions of expected values consistent with the reported estimates. As inputs for fitting, we used the indicated chance of the index value exceeding 6,500 (approximately the value at the time of the survey) and the expected index level at a given time. The chance represents the point in the cumulative density function[1], and the expected index level is the expected value of the probability distribution, where 6,500 is normalised to 1.

Having two data points allows us to uniquely fit a standard two-parameter lognormal distribution with a mean μ and a standard deviation σ. In our main analysis, the lognormal distribution is fully described by two numbers – the mean μ and the standard deviation σ. Our mathematical procedure is designed to calculate the means and the standard deviations of lognormal distributions that are most in line with each of our elicited pairs of data points.[2] We run our procedure for each respondent and for each of the one, two and five year pairs of data points to find the best fitting lognormal distribution. We call them *fitted distributions*.

For each respondent, we can now translate the probability of any FTSE index value in to the probability of the corresponding return of any of the five products using the formulas embedded in the payout profiles of the products. We use the *fitted distributions* of the FTSE values to generate the distributions of structured deposit returns. We call these distributions *FTSE-implied product return distributions*.[3]

1 More precisely, the probability that the FTSE100 value will exceed today's value (normalised to 1), is equal to [1-P(1+R^{FTSE}>1)], where R^{FTSE} denotes the return and P(.) denotes the probability.

2 The indicated probability of the index value exceeding 6,500 (approximately the value at the time of the survey) represents the point in the cumulative density function. The expected index level is the expected value of the probability distribution, where 6,500 is normalised to 1 (see Figure 23 in Annex 3). More precisely, the probability that the FTSE100 value will exceed today's value (normalised to 1), is equal to [1-P(1+R^{FTSE}>1)], where R^{FTSE} denotes the return and P(.) denotes the probability.

3 We generate the FTSE implied return distributions for 'Basic', 'Capped' and 'Cap&Floor' products based on the five year-fitted distribution. For 'Kickout' and 'Cliquet' products we additionally use the one year and two year-fitted distributions, and interpolate the data for every year for 'Kickout' product, and for every half-a-year for 'Cliquet'.

Lognormal specification

The left-hand-side panel in Figure 8 illustrates a lognormal distribution that would have been fit to a probability of 69% that the FTSE will be above its value and an expected value of 7,625 (x1.173 of initial 6,500).

We used the fitted lognormal distributions to generate individual-specific payout distributions of structured deposits (right-hand-side panel in Figure 8). We call these distributions 'FTSE implied product return distributions'. We generated the FTSE implied return distributions for the Basic, Capped and Cap&Floor products based on the five-year fitted distribution. For Kickout and Cliquet products we additionally use the one year and two year-fitted distributions, and interpolate the data for every year for Kickout, and for every half-a-year for Cliquet. The interpolation assumes normal distribution of returns and return independence. We are aware that these assumptions may be somewhat strict. However, our results for Kickout and Cliquet are unlikely to be biased in any particular direction and broadly align with those for the structured deposits Basic, Capped and Cap&Floor.

Figure 8. Left: fitted lognormal distribution. Right: Transformation of fitted FTSE distribution to the distribution of returns of Basic structured deposit

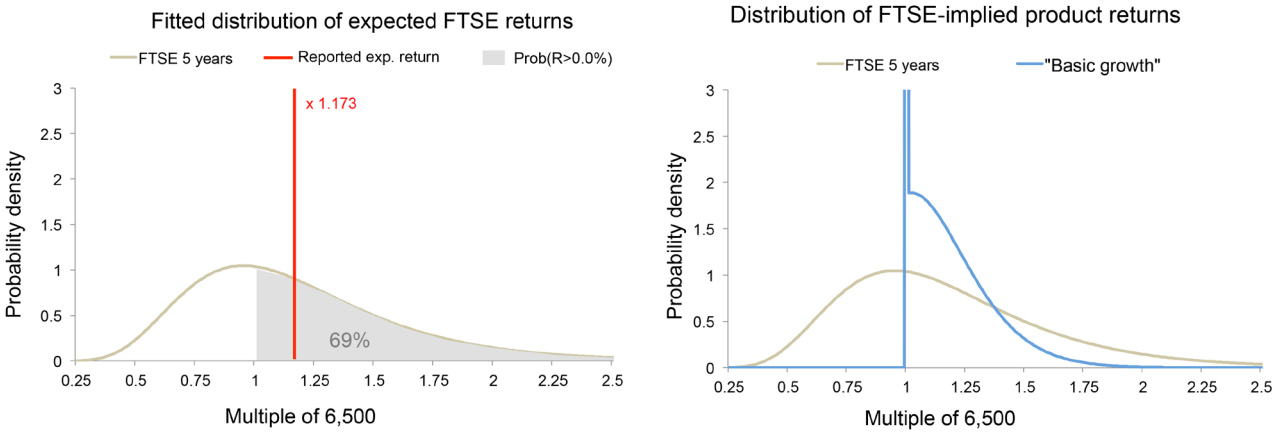

Figure 9. Differences in investors' implied and reported expectations of product returns

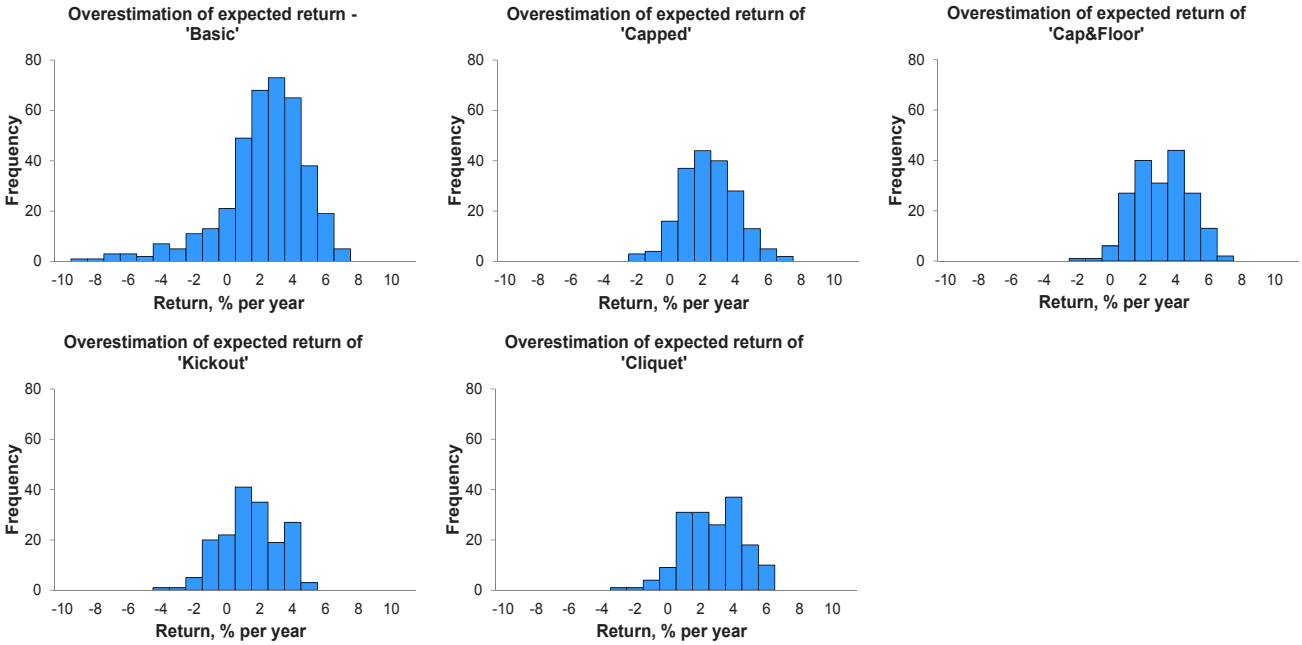

Alternative specifications

We acknowledge that lognormal need not be the best fit for every investor in reality. Respondents' expectations and perception of uncertainty may be more consistent with some different distribution that may or may not be expressed through a mathematical formula.

For example, a sample of structured deposit holders, as in our case, might on average have a more negatively-skewed distribution of expected FTSE values than lognormal and, therefore, may value downside protection more. Fitting the entire distribution to just two values makes it difficult to model investors' expectations of very extreme outcomes. For instance, given that we have only two observed values we cannot rule out that some investors expect the case that would be most challenging for our analytical approach, ie, the case that would maximise the implied returns of structured deposits: all FTSE 100 companies either completely crash or, if not, experience equivalently high growth. This case would be equally consistent with the observed probability and expected value data as a lognormal distribution.

We ran the worst case scenario with the five-year FTSE distribution and three structured deposits.[4] To generate the most extreme distributions we pushed the mass of the probability that FTSE will be below today's value in five years to zero and the remaining mass above today's value, so that the reported expected value (the mean of the distribution) holds. The worst-case scenario approach is extreme: it is equivalent to a situation where everyone in our sample simultaneously thinks that all hundred FTSE companies will fail in five years, with a probability of 41%, or will experience a very high growth with a probability of 59%. This assumption, which is consistent with the average expected FTSE growth, increases the value of downside insurance immensely and maximises the implied returns as well as the value of all structured deposits.

4 Mathematical appendix and spreadsheets are available upon request from the authors.

Under these assumptions, the overestimation was reduced from 10% to 7% for Capped and from 14% to 6% for Cap&Floor over the term of the product. Expected returns of Basic became underestimated by 5% from an overestimation of 11% due to the absence of a return cap. So this test overturns the overestimation for the one uncapped product, but falls short of explaining the overestimation in two other cases. We conclude that it is not our distributional assumptions that cause our findings.

TA5. Robustness: Survey bias

The range and the format of the survey tools, the slider and the ladder, may have had an impact on the distributions of responses about the expectations of the FTSE and of the returns of structured deposits. If the distributions of expectations were biased, the overestimation of returns could be fully or partially induced by this bias. Similarly, the recorded overvaluation of structured deposits could be a result of biased distributions of reported valuations. We examine the distributions of responses and compare the responses in the ladder tool to responses elicited though another less upwards-biasing tool.

In summary, indications of survey framing effects on investors' expectations, while potentially present, seem limited and do not provide a convincing alternative explanation of our findings.

Framing of the survey tools

The prime objection to our ladder tool is that respondents are drawn into the middle of a certain range of possible responses, potentially thinking that the experimenter chose a reasonable range (Andersen 2006). It is conceivable that if some respondents are uncertain about expected values, the range of the ladder tool drew them to the middle point, 7,000 for the FTSE index values (range from 4,000 to 10,000), and to 3.5% for reported returns of structured deposits (range 0% to 7%).

Examination of the FTSE one, two and five-year distributions (Figure 5 in the Paper) reveals that the vast majority of respondents saw today's index value of 6,500 as a reference point and expected some growth over time, but unclear how much exactly because of the potential survey bias. If respondents had no idea about future FTSE performance, they would likely space the one, two and five-year expectations roughly equidistantly throughout the upside part of the range, ie, between 6,500 and 10,000. Under this scenario, we would expect the one-year values to be concentrated around or roughly average at 7,375, the two-year values around 8,250, and the five-year values around 9,125. We observe that the real averages of all three values are very close to each other (6,700 to 7,500) and close to the current value of 6,500, compared to the available range of the ladder. If the ladder bias affected our respondents, the large available range up to 10,000 would have pulled expectations up; but it turns out that not by much, as expectations still remain close to the current value. We conclude that the 'pull-up' effect may have led us to slightly overestimate the FTSE expectations.

Respondents' expected structured deposit returns may be affected by framing too. Average reported returns are close to each other and not far away from the middle point of the ladder, ranging between 4% to 4.3% pa (Figure 1). If the true and unframed expectations of returns were actually lower than the middle of the range, the middle-point bias would have led to an overestimation of returns (or would have reduced it, in the opposite case). We know that the true distribution peaked either above or below the middle point (or exactly at it, in which case the bias would not matter). On whichever side the distribution peaked, the middle-point bias would have pulled responses, and also the peak of the distribution, towards the middle point. Since the middle point acts as a 'centre of gravity', it would not pull the peak over itself in to

the other side. The 'humps' of the distributions of reported returns of all products are all above the middle point of 3.5% (Annex 4, upper panel) and therefore could have only been pulled down towards the middle point from higher values.

Another possible objection is that respondents dragged the values from the top-right corner of the screen down into the ladder tool and therefore dropped them at the first opportunity, ie, in the top part of the ladder. If this effect was strong, the distributions would peak at high levels and definitely above the middle point. There is very limited evidence for this effect, if we look at the FTSE distributions (Figure 5 in the Paper), which had much longer ladders requiring to even scroll down the screen to put the values in, and also were elicited later in the survey when respondents were likely more tired. Most values are concentrated way below the top of the ladder and we only see very small clusters of responses in the top of the range.

Additionally, we obtained each respondent's valuation of Basic structured deposit in two ways: we first asked a set of binary choice questions (decision tree) and then used the ladder throughout the survey. Return overestimation was only slightly lower in the decision tree, which is a more downwards biasing frame than the middle-point and the hump of the distribution of responses was still above the middle point.

The decision tree asked respondents to choose between the Basic structured deposit and a fixed-interest rate, starting at 3%. The fixed-interest rate was always presented above the structured deposit, and went up or down depending on the previous choice to narrow down the valuation of the structured deposit to an interval of 0.5%. The outcome was then shown in the ladder and respondents could adjust the valuation if needed, before placing the other two structured deposits assigned to them in the same ladder.

The valuations reported through the decision tree are only slightly lower compared to the ladder (Figure 5) and still exceed both the implied returns and the 3% starting point. The distribution of adjustments to valuations made in the ladder is presented in Figure 6, indicating that most respondents did not adjust. Those who adjusted did so slightly upwards. We found that respondents, on average, increased their valuations by 0.28% (1% conditional on the change being made). Note that the decision tree tool could have been biasing itself. If respondents did not know the answer and chose randomly between any two presented options, the result would be a binomial distribution with the starting point at 3% pa (Figure 4). The actual distribution (Figure 5) significantly differs from the binomial distribution implied by random responses, suggesting that the expectations we elicited through the decision tree were not random.

The FCA's commissioned consumer research (conducted by 'Ignition House') also suggests that investors' expectations of structured products are generally high. Face-to-face interviews with UK structured deposit holders in 2013 revealed that expectations of returns ranged between 1% and 10% per year, with most values concentrated around 4% to 7% per year. These results support that the middle point of our ladder would likely bias reported expectations down, but should be interpreted with care: the responses were given by product holders who bought their structured deposits during the last three to five years. The growth in the FTSE100 value over the last few years may justify high expected returns, but could have reduced forward-looking expectations. So investor expectations of structured deposit performance could be lower today for the next five years.

Finally, some investors could have had expectations of the FTSE and of structured deposit returns that were outside the available range of the ladder tool, in which case our elicited distributions would have a lower variability than the real distributions. If such observations were

cut-off by the range asymmetrically, the means of the distributions could be biased and could explain some or all of the overestimation we find.

Figure 5 in the Paper shows that no significant parts of the distributions of investors' expectations of the FTSE seem to be cut off, suggesting that the range of the ladder was wide enough not to reduce the true variability. Distributions of investors' expectations of product returns (Figure 1) reveal that the range of the ladder may have led to some high values being cut off on the right hand side, while no values are cut off on the left hand side, because the minimum return of all products was 0%. If we cut off some high values of reported expected returns, we are likely underestimating the mean expected returns of structured deposits reported by investors and therefore underestimating how much the respondents overestimated them.

Effects of low quality responses

It is possible that the survey bias had a stronger effect on respondents who did not have a firm view of their expectations and so were more susceptible to framing. There is a disproportionately large share of respondents who indicated a 50% probability for both the first and the second probability estimate (more than 0.0% return and more than 3.0% return). Research suggests that choosing 'fifty-fifty' as a response to a probability question may be a way of expressing unwillingness or an inability to attach a numeric value to uncertainty (Bruine de Bruin et al, 2000). Such responses may include numerically meaningless entries and using them in the analysis could have biased our findings in a particular direction. We tested the robustness of implied expected returns by removing all observations where respondents reported 49% to 51% chance that FTSE100 would grow in the future.[5] Our main results were not significantly affected, suggesting that this behavioural phenomenon did not bias our responses in any particular direction. Both the implied and reported expected returns are somewhat higher after removing 'fifty-fifty' responses (significant at 5% level), but the magnitude of overestimation is not markedly different from the main sample (1.98% after removing 'fifty-fifty' responses, compared to 1.87% before; not statistically significant at 5% level).

More generally, all poor quality responses could have been subject to survey frame and induced the survey bias. For example, anchoring effects can be lower for more knowledgeable respondents (Wilson et al, 1996). To show that our results did not stem from poor understanding of survey tools, we ran our main analysis with 'top performers' only. We removed from our main analysis sample all respondents who answered at least one of the questions on the examples of the tools incorrectly (slider and ladder), or answered at least one of the five understanding questions on the three products incorrectly, or saw at least one warning about internal implausibility of combinations of their indicated beliefs. The respondents remaining were better educated, had slightly higher incomes and assigned a slightly higher financial expertise to themselves. Our main result of expected return overestimation still holds and is significant across all products. The magnitude of overestimation is not markedly different from the main sample (1.95% with top performers only, compared to 1.87% before; the difference is not statistically significant at 5% level).

5 We allow for an accidental error of +/- 1 percentage point when placing the marker at the '50% ' mark of the slider.