Comments on ENA/Oxera

Stephen Wright and Robin Mason

5 November 2020

1. Government Bonds as a Risk-Free Asset (ENA 5.7, Oxera 2.3)

a) Negative Beta

1.1. This analysis seems to derive from a confused application of the concept of the risk-free rate.

1.2. A risk-free rate can only ever be derived at a given horizon. It is a standard result that the yield on a default-free bond zero coupon bond is, *by construction*, risk-free if *and only if* it is held to maturity. If held to maturity, both payoff and price are known, and there is no price risk. Furthermore. if the bond is index-linked, there is no inflation risk.

1.3. At any horizon less than maturity, the price of the bond may vary if market yields change—so a long-dated bond is *not* a risk-free asset at these shorter horizons. This can in principle encompass the possibility that some bonds may have negative beta at very short horizons (as Oxera find) if their short-term price fluctuations are negatively correlated with market returns.

1.4. This observation just means that if we wanted to find a risk-free asset at these short maturities, if we look at long-dated indexed-bonds, then we are looking in the wrong place. Short-dated bonds *are* risk-free at short horizons and are traded in liquid markets. If regulators wanted an estimate of the risk-free rate at a very short horizon, the data would be readily available (and would almost certainly point to a *lower* value than used by either CMA or Ofwat).

1.5. But both regulators and the CMA are in agreement that this is *not* what they want: that the appropriate horizon for estimating the WACC is a long horizon. And, as noted above, at any given chosen horizon, the yield on a bond with maturity equal to that horizon *is* risk-free.[1]

b) The "Convenience Premium"

1.6. Oxera also make an argument based on the "convenience premium" of government bonds (we take this to mean a price premium, hence a lower yield). Again, this feature may well apply at short horizons, and is presumably the counterpart of the "liquidity premium" (in terms of yields) that, it is often asserted, applies to corporate bonds. But at any given horizon, if a long investor wishes to invest in a risk-free asset, then the only candidate asset is a default-free government bond with maturity equal to horizon (since

[1] As noted in the UKRN report, strictly speaking this result only applies to zero coupon bonds, but the Bank of England infers implied zero coupon yields from traded coupon bonds.

corporate bonds may have both liquidity and default risk). The yield on such a bond is risk-free at that horizon, hence has zero beta by construction. An investor could of course invest in higher yielding bonds, but in so doing would take on some market risk, and hence would be at a higher point on the Capital Market Line.

2. Averaging Historic Returns

2.1. On this issue there are actually two distinct questions:

1. Whether we should *adjust* the estimated arithmetic mean for parameter uncertainty in setting allowed returns.
2. How best to *estimate* the true arithmetic mean market return = E(R_M).
   We interpret the CMA's para 9.178 (p547) as drawing this distinction.

2.2. The Oxera/ENA approach appears confused about whether it is attempting to answer question 1 or 2. Many of the points it discusses actually relate to 2, but reveal a misunderstanding of the issues. We address each issue in turn. In an annex to this note, we comment by annotation on a number of statements that reveal this confusion.

Should we adjust the arithmetic mean estimate for parameter uncertainty?

2.3. This is the issue addressed by Blume, Cooper, JKM and hence Schaefer. This literature raises a range of issues that are *potentially* of interest to regulators. But the only clear-cut result that we can glean from these various contributions is that in the presence of uncertainty about the true mean, you may wish to adjust the arithmetic mean down or up, depending on a) which problem you want to solve, and b) as JKM point out, what your loss function is. For example, on point a) JKM show that different adjustments should be applied depending on whether you wish to estimate a terminal value, or make an asset allocation decision, and that this latter decision in turn depends on your assumed risk aversion.

2.4. It is possible that there may be yet another optimal adjustment if we were to set up the regulator's decision. But it is clear that this would be pretty complicated, and the problem has not, as far as we are aware, ever been set up fully. Indeed, to do so would require a full specification of the stochastic properties of the *actual* return earned by regulated companies, conditional upon any allowed return. A minimal requirement to even begin on this process, as we argued in the UKRN report, would be that regulators should explicitly account for the difference between the regulatory *allowed* return and the regulatory *expected* return, by taking into account (strong) historic evidence of outperformance. If regulators and potential investors *were* to behave like discounters, they could not even begin to set up the problem without allowing for this. But we are still some way even from this stage.

2.5. Going beyond this stage could indeed raise a range of potentially interesting questions. Schaefer's paper refers to a (completely correct) distinction between compounders and discounters, but does not actually take any position on which regulators should be. Instead, Shaefer's recommendation is that "all the CMA needs to do is to provide an unbiased estimate of the arithmetic return. Discounters and compounders will then make different adjustments to this rate" (paragraph 18). We think this fails to reflect what the regulator needs to do. What would be needed to capture this more fully is a model of how a rational investor compares the prospective returns offered by regulated companies, conditional upon an allowed return set by the regulator, to the returns offered on competing investments. Consistent with the approach of all the contributions to this literature, this model would need to take into account that such a rational investor would not know the true WACC, and presumably would have the same parameter uncertainty as the regulator. The regulator's problem would then condition upon the investor's problem, and would need to take into account both parameter uncertainty, *and* a loss function due to asymmetric social costs of over- vs under-investment (hence the solution could presumably involve an element of *conditional* aiming-up).

2.6. This would be an interesting problem to solve; but it simply has not been solved yet. It is also worth noting that any such investigation would, as in any economic analysis, only result in a *model* of optimal behaviour, conditional upon assumed behaviour by the notional rational investor. Any economic model needs also to be consistent with the data. While we may not know exactly what the predictions of the model might be, we *do* have strong evidence that actual market investors appear to value the two quoted water companies at significantly more than their RCV. As we have argued elsewhere, this suggests strongly that if Ofwat is getting it wrong, it is erring on the upside in setting allowed returns. Whether this would be consistent with a fully specified model is, at this stage, simply unknowable.

2.7. Thus we would argue that, in the absence of a properly specified model, the current practice of implicitly simply assuming away parameter uncertainty at least has the merit of simplicity and familiarity.

How should we estimate the true arithmetic mean?

2.8. Many of Oxera/ENA's arguments appear to confound the above issues with the simple question of how best to estimate $E(R_M)$. The CMA's Table 9-3 possibly accentuates this confusion, since it includes both alternative estimates of the arithmetic mean, *and* adjusted estimates such as JKM or Blume, which, as discussed above, arise from the solution of a problem that may or may not be relevant to the regulator's problem. However, it also contains estimates of the mean annualised arithmetic return at longer horizons, based on overlapping and non-overlapping data. These are simply *alternative* estimates of the arithmetic return; and it is notable that these are always lower (albeit not much lower) than the arithmetic mean annual return calculated directly (and, incidentally, at a 10-year horizon, almost identical to the JKM estimate).

2.9. The rationale for using estimates at longer horizons is (despite what Oxera claim) exactly because, in the presence of some predictability of returns (and *hence* some negative serial correlation), the arithmetic mean of the annual return will be an upper bound for the true arithmetic mean. Given this, there are alternative methods of calculating adjustments.

2.10. Both MMW and the UKRN report argued the case for working in terms of the geometric mean return, and adjusting *up*wards, rather than working from the arithmetic mean return, and adjusting *down*wards. In this context, it is notable that *all* the academic papers cited in this debate (including Schaefer's) take an approach that is consistent with this, since all work on the assumption that returns data are generated by an underlying process for the *log* return (often, but not invariably, assuming log returns are normally distributed). The log return model has the advantage both of plausibility (since returns themselves are bounded below) and simple time aggregation. This means that annualised mean log returns (and hence geometric returns) are invariant to the frequency of data used. The same feature only applies to annualised arithmetic mean returns if returns are serially uncorrelated at any given frequency.

2.11. The Oxera report appears to be predicated on the assumption that the arithmetic average of annual returns has some special significance. In a log-return world, it does not: annual*ised* returns can be calculated from data at any frequency, whether longer or indeed shorter time intervals are used.

2.12. To illustrate, we could in principle use much higher frequency data to derive estimates of geometric and arithmetic mean returns. This would not affect the geometric mean return, but can produce distinctly different answers for the arithmetic return. Thus, for example, if we use the daily FT All share return data (from 2001-2017) used in the beta estimation for the UKRN report, the arithmetic average daily return is almost exactly twice the geometric average. In annualised terms this would imply a gap between geometric and arithmetic averages of around 3.5 percentage points. Nor is this surprising, since it is a well-known result that annualised daily stock volatility significantly overstates the volatility of annual returns.

2.13. Since all those involved in the discussion appear to agree that we should work on the basis of an investor horizon that is of the order of 10 to 20 years, presumably no one would actually recommend using daily data to construct annualised arithmetic averages. But, by implication, nor is there anything special about the arithmetic mean of the annual return—and, as noted above, the evidence produced in the CMA's Table 9-3 shows that taking arithmetic averages over longer horizons does indeed bring down the estimated annualised arithmetic mean.

2.14. Is this a sufficient reduction? In both MMW and the UKRN report we suggested a quite broad range of *upward* adjustments, of 1 to 2 percentage points, to the geometric average return. In the CMA's table the gap between the arithmetic and geometric annual averages is 1.8 percentage points (on our preferred consistent CPI/CED basis). Simply using the twenty-year average of overlapping returns would bring this down to

1.5 percentage points (i.e., 6.7%-5.2% on the same CPI/CED basis), thus, in the centre of our proposed range.

2.15. Is there a case for an estimate being closer still to the geometric mean? In logic yes. If we see evidence of negative serial correlation of returns (which we do), then this is indirect evidence of *potential* predictability of returns using a bigger information set. There is indeed an extensive literature that claims to finds evidence of such predictability.[2] In MMW we showed that allowing for this could in principle result in more precise (and lower) estimates of the arithmetic return (since the adjustment is driven by a lower estimate of the *conditional* variance of returns). We would not wish to draw strong conclusions from the estimates that arise from this approach[3] since the same extensive literature also points to considerable controversy as to the reliability of evidence of return predictability. It was for this reason, and on the grounds of implementability and defensibility, we advocated in the UKRN report that, despite evidence of predictability, regulators should rely on historic averages, rather than predictive models.

2.16. Nonetheless the point does remain, which in turn implies that we cannot rule out the possibility that the best estimate of the arithmetic mean is indeed closer to the geometric mean than would be implied by simple arithmetic averaging of long-horizon returns.

---

[2] Discussed at some length in the UKRN report.

[3] See for example PWC, 2019

Annex: Detailed responses to ENA/Oxera

ENA

4.28 The CMA's TMR is artificially skewed downwards due to its approach to averaging historical returns. It has incorrectly dismissed the use of the arithmetic average.

Here as elsewhere ENA appear to equate the arithmetic average to the arithmetic average of annual returns. Our interpretation is that the CMA has, like MMW and the UKRN report, *not* dismissed the use of the arithmetic average, but has taken (some) account of the arguments that the average *annual* return is upward biased as an estimate of the true arithmetic average return.

4.31 The CMA's contention that the arithmetic average is upwards biased is wrong, and is predicated on a false dichotomy between an investor perspective and a capital budgeting perspective, as well as an incorrect conclusion that serial correlation is a reason for calling the use of the arithmetic mean into question

Here the ENA appears to be conflating two distinct arguments, as discussed in the main document.

4.32 In reaching this conclusion, the CMA appears to have misconstrued the analysis submitted by Professor Stephen Schaefer. Contrary to the CMA's suggestion that Prof Schaefer considered that 'the most weight should be given to the capital budgeting perspective,' Prof Schaefer noted that some weight should be given to this perspective

We see no evidence that Prof Schaefer took a view on the issue at all – he simply drew the distinction between "compounders" and "discounters".

4.34 The role of Ofwat, and by extension the CMA, is to set prices by including an allowance for the rate of return. Investors can then use a discount rate that is either lower or higher than the arithmetic average according to whether they are estimating future or present values.

Thus far we agree, this is a fair summary of Schaefer's note.

…. The CMA's proposal to set cash flows by using a rate of return lower than the arithmetic average has the result of embedding a downward bias to the value of the regulated business and under-compensating investors.

This has no support from Schaefer's analysis *unless* we assume that regulators are discounters.

4.35 Given the need to take account of the arithmetic average, if estimators are used, these should not all have a systematic bias in one direction away from that value. The CMA, however, has used only downwards-biased estimators (namely, the JKM and Blume estimators) and in doing so has wrongly excluded the Cooper estimator which applies an upward adjustment to the arithmetic mean (and thereby offsets the bias of the JKM and Blume estimators to some extent).

It is far from clear how much weight the CMA has actually put on JKM and Blume, given that these give very similar adjustments to the (legitimate) downward adjustments to the annual arithmetic return. We would argue that it is premature to use *any* of these alternative estimators, without a clear rationale based on the regulators's problem.

Oxera

P6

While there is debate about which is the more appropriate averaging method in any given context, in standard corporate finance textbooks the arithmetic average is generally adopted for estimating the ERP to use when computing required equity returns. Indeed, DMS themselves make the following statement

This [the arithmetic mean risk premium] is our estimate of the expected long-run equity risk premium for use in asset allocation, stock valuation, and corporate budgeting applications.

The appeal to the authority of DMS is weakened by the fact that, under parameter uncertainty, JKM show that the arithmetic average risk premium is *not* the correct estimate to use in asset allocation (a downward adjustment is required); and Cooper shows that is *not* the correct estimate to use in corporate budgeting applications (an *upward* adjustment is required).

Two submissions made by Professor Stephen Schaefer to the CMA have highlighted that this statement is incorrect to focus on the role of serial correlation. Professor Schaefer states that:

the difference between the arithmetic and geometric mean return is given by one half of the variance. Bound up in the assumption of normality are further assumptions that both the expected return and the variance of returns are constant over time and that returns are not serially correlated.

And then shows based on analysis of the DMS data that:

despite this, the difference between the arithmetic and geometric means is indeed well approximated in the data by one half the variance.

In other words, the empirical evidence does not justify deviating from the arithmetic mean based on arguments concerning serial correlation

Professor Schaefer's empirics in the paper cited are entirely based on the properties of the annual return. He thus simply does not address the issue of the impact of negative serial correlation at longer horizons. We would note that he *does* say, after deriving this result:

Of course the assumptions of constant expected return and constant volatility do not hold in practice[4]

[4] Schaefer, 2020, Deriving unbiased discount rates from historical returns, Prepared for Energy Networks Association, 14 February 2020