Comments prepared for Ofwat on the CMA's Provisional Findings Anglian Water Services Limited, Bristol Water plc, Northumbrian Water Limited and Yorkshire Water Services Limited price determinations: Cost of capital considerations

Final report
26 October 2020

Stephen Wright
Robin Mason[1]

[1] The authors are involved in a personal capacity. The views expressed in the work should not be interpreted as necessarily reflecting the views of any of their respective institutions. They would like to thank PJ McCloskey and Donald Robertson for helpful input to this report, although responsibility for any errors remains with the authors.

1. Introduction

1.1. Ofwat have asked Stephen Wright and Robin Mason to comment on the approach that the CMA has taken to the estimation of the cost of capital of regulated water companies who have appealed Ofwat's PR19 price determination. Specifically, Ofwat has asked the following aspects to be considered:

- The estimation horizon(s);
- The risk-free rate;
- The expected/total market return;
- The equity beta;
- Debt: embedded and new;
- Aiming up;
- The gearing outperformance sharing mechanism.

1.2. While Ofwat has commissioned the work, this report is independent and expresses the views of its authors, which may not necessarily reflect the views of Ofwat. The arguments in the report should be seen holistically (especially in the context of our overall view on aiming up).

2. The estimation horizon(s)

2.1. In our view, there is no clear-cut guide to the appropriate choice of horizon. The main issue is that horizon should be chosen *consistently* across all components of the estimate cost of capital. On the issues of the risk-free rate and the total/expected market return, we see little difference between the CMA and Ofwat, both of whom pick long horizons of 10 to 20 years, and little reason to argue with either approach.

2.2. Where issues of horizon do become more contentious is in beta estimation. Here, our views are:

2.2.1. The rationale for estimating beta using short samples of daily data is inconsistent with the long horizon assumed in other calculations.

2.2.2. There is an inconsistency between using such short samples (implicitly assuming some form of parameter instability in beta) and using OLS estimation that is predicated on a constant value of beta.

2.2.3. The apparent short-term variation in rolling estimates of beta is consistent with a multivariate GARCH process in which the conditional beta varies over time, but has a stable unconditional value.

2.3. Each of these considerations leads us to argue for beta estimates using longer samples, and ideally taking GARCH estimates into account. (Fuller details of our arguments can be found in the UKRN 2018 report[2].) This appears to have influenced recent beta estimates produced by Ofgem, but it appears to have little impact on

[2] UKRN (2018), Estimating the cost of capital for implementation of price controls by UK Regulators. A report by Wright, Burns, Mason and Pickford.

CMA's analysis. We discuss the implications of making use of longer samples of data for beta estimation in Section 5.

3. The risk-free rate

3.1. Ofwat followed the recommendation of the UKRN report in estimating the risk-free rate (RFR) using RPI index-linked UK government bonds i.e., Index-Linked Gilts (ILGs). (The detailed calculation used by Ofwat is summarised in the CMA report at paragraph 9.46.) The point estimate of the RFR used by Ofwat is -1.39%, CPIH-real.

3.2. The CMA's approach is somewhat different. First, it uses ILG data in a slightly different way (summarised in paragraph 9.140 of the CMA report) to calculate a lower bound on the RFR of -1.40%. More importantly, it argues that "the yield on ILGs is likely to sit below the 'true' estimate of the theoretical RFR" because "the government can borrow at rates substantially lower than even higher-rated non-government market participants" (paragraph 9.135). This leads the CMA to look at AAA-rated non-gilt yields, as an upper bound (in their view) for the RFR: see paragraph 9.137. It calculates this upper bound to be -0.81%. Finally, through its aiming up argument (about which we comment further in Section 7), it chooses the 75% point within this range i.e., -0.96%. It argues that this point estimate corresponds to the zero-beta asset expected return within the CAPM framework.

3.3. We first note that this is counter to the approach taken by the CMA in the recent NATS decision,[3] where the CMA concluded "We considered that current ILG rates continued to provide the most appropriate basis for the measurement of a notional investors' achievable risk- free returns" (paragraph 13.258). It is the first time, that we are aware, of a UK regulator (implicitly) advocating that the CAPM be implemented with a distinction between the lending and borrowing rates of the marginal investor.

3.4. If this distinction is to be made, then it would need to be made carefully, rather than approximated by an upward adjustment in the estimate of the risk-free rate, since the framework with different lending and borrowing rates has a number of complications. The CMA outlines its approach in appendix C of its report. There, it states: "In paragraph 9.77 we highlight Berk and DeMarzo's explanation of the CAPM with different savings and borrowing rates available to different investors. Berk and DeMarzo's Figure 11A.1 … demonstrates the different securities market lines (SMLs)[4] available to investors depending on the rate at which they can borrow, and that this leads the market rate to sit a level between the rates achievable by these different classes of investor. For our purposes, *we interpret this as the true RFR sitting between the rates available to our best proxies - government rates, represented by ILGs, and private investors, represented by AAA non-government bond indices*" (paragraph 5, page C3, *emphasis added*).

[3] CMA (2020) NATS (En Route) Plc/CAA Regulatory Appeal – Final Report (NATS/CAA).

[4] This Figure does not in fact show securities market lines: it shows capital allocation lines. SMLs have covariance (β) as the horizontal axis.

Figure 1: CAPM outputs with different savings and borrowing rates available to different investors: taken from the CMA report, appendix C.

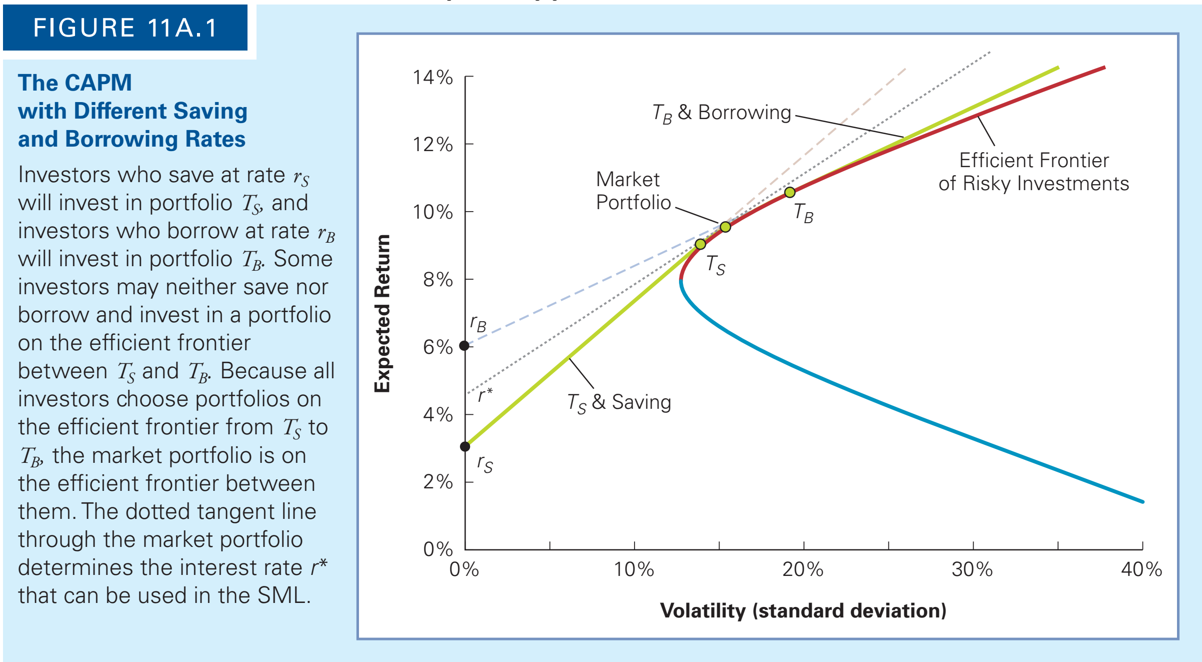

Source: Berk and DeMarzo (2014).

3.5. Berk and DeMarzo's[5] explanation follows from the original papers by Brennan (1971)[6] and Black (1972)[7], extending the CAPM to cases where there are difference between borrowing and lending rates (including when there is no risk-less asset). The capital allocation line (CAL)[8] with different lending and borrowing rates involves 3 segments:

3.5.1. For a lender, the optimal portfolio is along a straight-line segment between the lending rate r_S and a tangent portfolio, denoted T_S in Figure 1, taken from the CMA's report.

3.5.2. For a borrower who can only borrow at a higher rate (r_B in the diagram), the optimal portfolio is at the tangent point T_B.

3.5.3. If the investor neither borrows nor lends, his optimal portfolio lies at any point along the curved section between T_S and T_B.

3.6. It is worth stepping through this framework in more detail. The first step is to consider the CAL for a net lender. We take the lending rate r_S to be the ILG yield, r_f, since lenders are assumed to be able to buy ILGs. The logic of the CAPM is that these investors will allocate their capital between just two assets: the risk-free asset at which they lend (with return r_f) and the tangency portolio T_S, where the straight line from the point r_f touches the efficient frontier. Since this investor is taken to be

[5] Berk, J. and P. DeMarzo (2014). *Corporate Finance*, Third Edition.

[6] Brennan, M. (1971). "Capital Market Equilibrium with Divergent Borrowing and Lending Rates". *Journal of Financial and Quantitative Analysis* 6, 1197–1205.

[7] Black, F. (1972). "Capital market equilibrium with restricted borrowing". *The Journal of Business*, 45(3), 444-455.

[8] The capital allocation line shows graphically the possible combinations of expected return and risk (volatility) an investor can achieve, by investing in available assets. Paragraph 3.6ff explains further.

a net lender, its CAL is just the line between r_f and T_S: no other options are relevant for this investor. See Figure 2.

Figure 2: The capital allocation line for a net lender

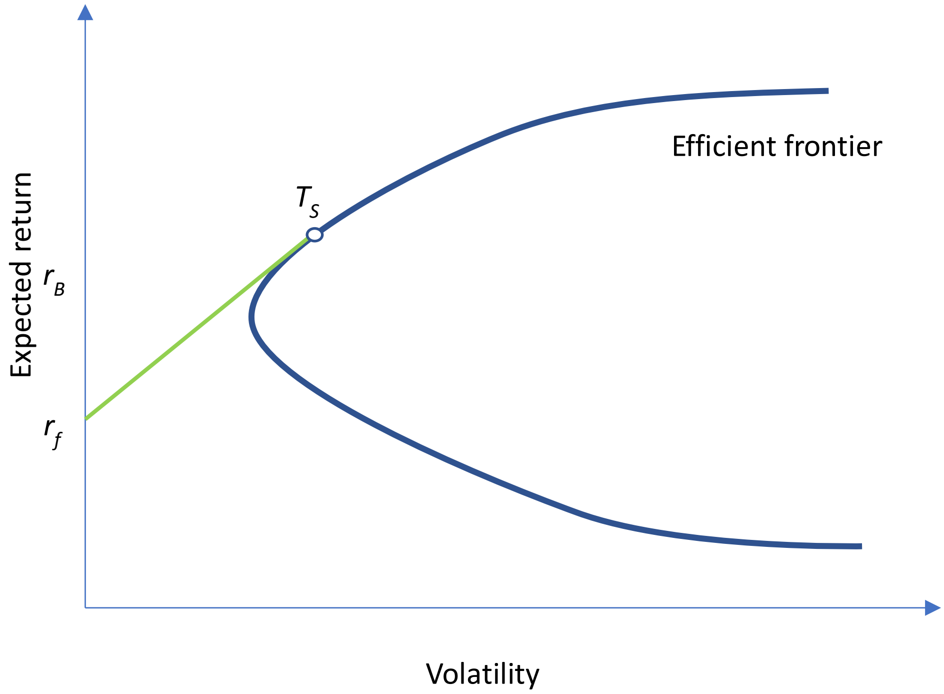

3.7. Now consider the CAL for a net borrower, shown in Figure 3. It is the straight line to the north-east of the tangency portfolio T_B: to the north-east, since this investor is a net borrower. (The portfolio T_B is the point of tangency for a straight line drawn from the interest rate r_B.) Note that portfolios to the south-west of T_B are not available to this investor, since they would involve lending, and the investor is a net borrower.

3.8. Finally, the CAL for an investor who is neither a lender nor a borrower is shown in Figure 4: the section of the efficient frontier shown in red. No other portfolios are available, since they would involve either lending or borrowing, and this investor is assumed to do neither.

Figure 3: The capital allocation line for a net borrower

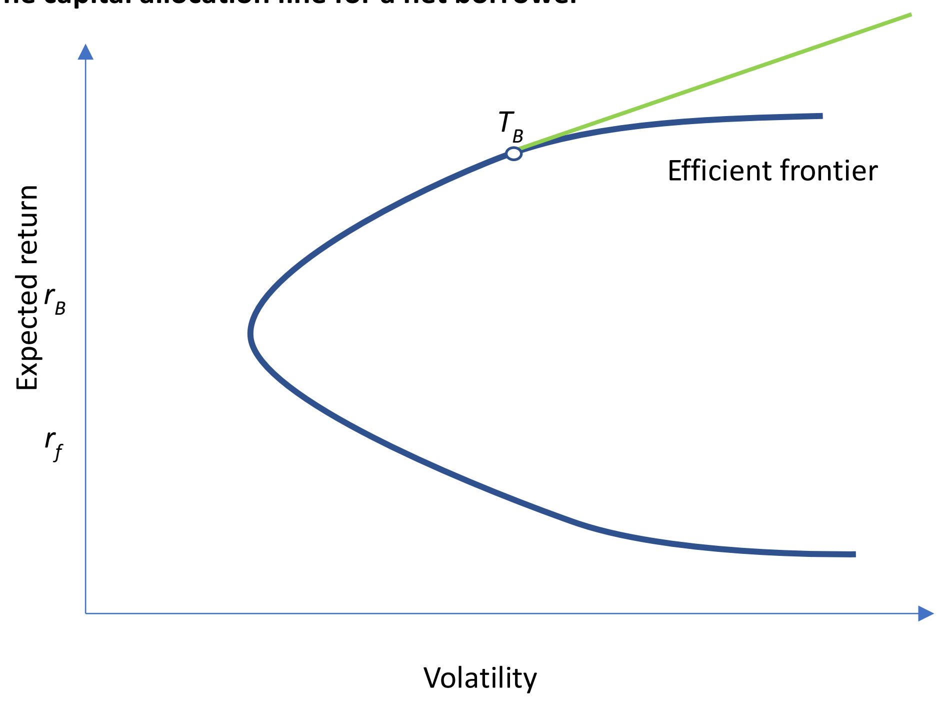

Figure 4: The capital allocation line for an investor who is neither a net lender nor net borrower

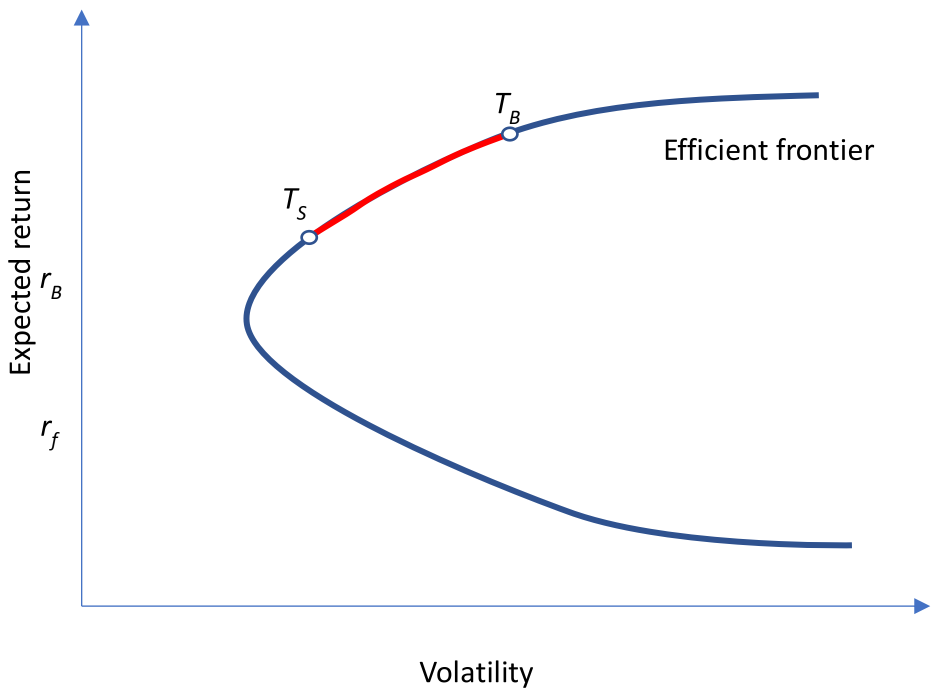

3.9. The key observation at this point that, while the CAPM is an equilibrium model—the total value of demand for an asset must equal the total value of its supply—the important investor in the CAPM framework is the *marginal* investor i.e., the investor who is most likely to be trading actively at the margin and therefore has the most influence on price. Put differently, the expected return on a traded asset will be determined by the investors who are actively trading it. Greater care has to be taken about the marginal investor when borrowing and lending rates are taken to be different, since different types of investors face different capital allocation lines. With equal borrowing and lending rates, all investors face the same CAL and the distinction between the marginal and average investor is not needed.

3.10. The main questions, therefore, concern (a) the type of the marginal investor in UK water companies; and (b) whether this marginal investor is likely to be a borrower, a lender, or neither. The type of the marginal investor is typically assumed to be the same as the type who holds most of the stock. In the case of the four UK water companies in this appeal, the marginal investor appears to be institutional,[9] acting on behalf of (non-institutional) savers.[10] Hence the marginal investor for these companies is effectively a net lender.

3.11. As result, the CAL of the marginal investor in UK water firms is as shown in Figure 2: the line segment between the ILG yield r_f and the tangent portfolio T_S. Note that this line lies strictly below the line formed by the zero-beta asset, with return r^* and the tangency portfolio, as Figure 5 shows. In other words, the CMA has strictly over-estimated the cost of equity capital facing the marginal (i.e., institutional) investor in UK water companies. It is worth being clear that this line used by the CMA is *not* a capital allocation line: it does not represent a feasible set of choices for any investor. It is instead an average of the choices of the three different types of investors: lenders, borrowers and the rest.

[9] For the shareholders of the companies, see https://www.awg.com/investors/, https://www.bristolwater.co.uk/the-company-today/ - :~:text=Group Structure %2F Shareholding,Corporation of Japan (20%25), https://www.nwg.co.uk/about-us/nwl/nwl-structure/#:~:text=Ownership%20of%20Northumbrian%20Water%20Limited,Northumbrian%20Water%20Group%20(NWG), https://www.yorkshirewater.com/corporate-governance-and-structure/.

[10] The marginal, institutional investor can, of course, receive dividends from the regulated company, which may (directly or indirectly) be funded through debt. The limited liability of the firm isolates the institutional investor from this short (borrowing) position, so that the institutional investor is still net long (lending).

Figure 5: The capital allocation line for the marginal investor versus the line used by the CMA

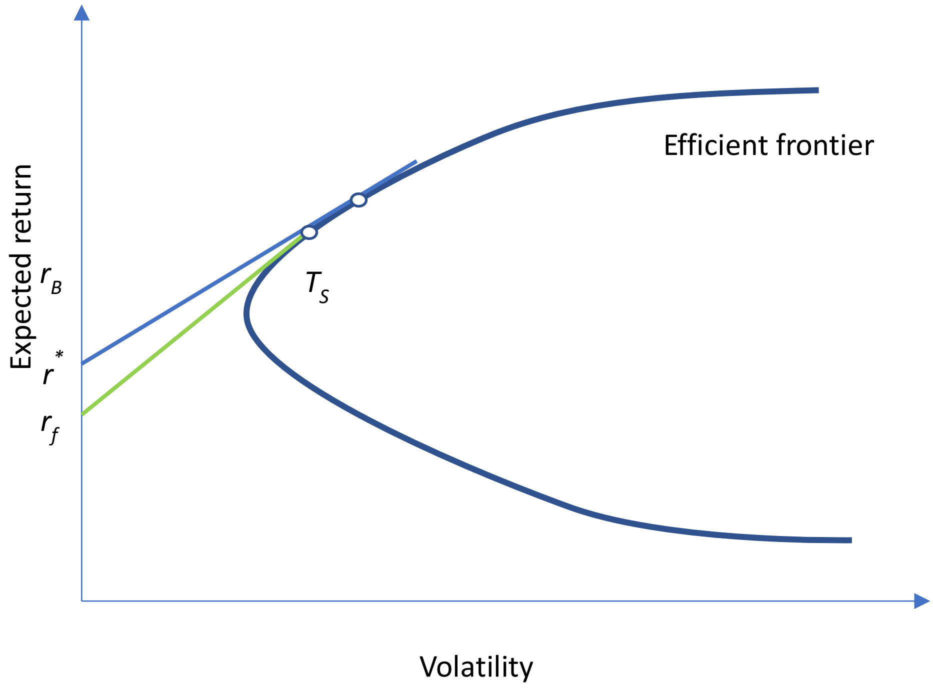

3.12. But even if we were to adopt the CMA's argument—that the CAL should be the line defined by r^* in Figure 5—we think that the CMA has made an error in how it calculates r^*. The CMA calculates r^* as the 75th percentile point between the ILG yield and the AAA-rated non-gilt yield, where the 25/75 weights come from aiming up arguments (see Section 7). But the logic of the CAPM with different borrowing and lending rates means that the zero-beta asset expected return r^* should, in fact, be based on the proportions of investors in UK water companies who are net lenders, net borrowers or neither.[11] In the case of the four disputing companies, it seems that the investors are predominantly net lenders, and so the zero-beta asset expected return r^* should lie very close to the ILG yield.

3.13. In summary: if the CMA wants to use a CAPM with different lending and borrowing rates, then it needs to do so carefully. The identity of the marginal investor becomes critical. In the four disputing water companies, the investor is institutional, and therefore a net lender. Their relevant risk-free rate (for lending) is the ILG yield; and the relevant capital allocation line is defined by the ILG yield. Even if the CMA's argument is adopted, the zero-beta return should be calculated according to the proportions of investors in UK water companies who are net lenders, net borrowers or neither—not by aiming up considerations. This calculation leads to a zero-beta asset expected return r^* that lies very close to the ILG yield. zero-beta asset expected return r^* should lie very close to the ILG yield. Either way,

[11] Similarly, the market portfolio identified in Figure 1 is the tangency portfolio on the efficient frontier defined by a straight line drawn from r^* on the vertical axis.

the CMA over-estimates the RFR. The ILG yield remains our preferred measure, as argued in the UKRN report.[12]

4. The expected/total market return

4.1. The CMA's point estimate for the TMR[13], given in Table 9.26 (p. 674) is 6.95%. The derivation of this figure is convoluted and in some of the details, somewhat obscure. Working backwards, this figure is in turn derived as the notional 75th percentile of the range of 6.20% to 7.21% given in Table 9.4 (p. 558). The CMA derives this estimate in turn from a range of estimates, based on a) average historic returns, and b) ex-ante and forward-looking estimates.

4.2. Before discussing the derivation of this range, we should note that this notional percentile is calculated on the implicit assumption of a uniform distribution within the chosen range. While commonly done, it is worth noting that even this choice has the impact of pushing up the implied estimate further. If, for example, we were to assume a normal distribution centred on the midpoint of the range, an assumed standard deviation of 0.25 would result in 97.5% of the distribution lying within this range. But the 75th percentile of this distribution would only be at a point roughly two thirds across the range, rather than three quarters. A relatively small effect, of perhaps 8 basis points, but, in line with almost all the adjustments, one that pushes the point estimate upwards.

Average historic returns

4.3. In the UKRN report, we argued strongly that all real return estimates should be calculated on a consistent CPI basis, both for historic calculations and for forward-looking estimates feeding into allowed returns. We presented a range of arguments on this issue, which are summarised at the start of Appendix D (p. D-109). Probably the most pertinent of these was the second sentence of our point d), on that page (bold face added for emphasis):

> *"Changes to the underlying methodology mean that the RPI is not comparable over time, whereas historical CPI estimates try to match current methodology.* ***Historic equity returns deflated by RPI will therefore have limited informational content about future equity returns deflated by RPI.****"*

4.4. We should stress that these views (shared by all the co-authors of the report) were not just our views, but based on the considered arguments of a wide range of experts on the subject. Indeed, we argued that the issue of which price deflator should be used ought essentially to be subcontracted to those with specialist expertise—in particular the Bank of England and the ONS.

[12] The latest 20-year RPI-indexed yield is -2.58%, so our latest estimate of the risk-free rate is -1.7%. This widens the gap further to the CMA's preferred risk-free rate of -0.96%.

[13] As noted in the UKRN report, we prefer the acronym EMR, reflecting the fact that what we are attempting to measure is an *expected* return.

4.5. In light of the full range of arguments, we also recommended that regulators should wherever possible shift to replacing RPI with CPI in its indexation process (a process that we understand has been fully carried out by Ofgem, but thus far only partially by Ofwat); but, where some element of RPI indexation remains, historic returns should still be calculated on a consistent CPI basis, with the adjustment to a forward-looking RPI basis carried out only as a final stage.

4.6. This is *not* what CMA has done. Despite an extensive discussion of this issue, in which the CMA acknowledged the problems with RPI in recent years, and the sequence of shifts in the relationship between RPI and CPI, in terms of the range of real return estimates it chooses, the CMA has in effect given more-or-less equal weight to the CPI and the RPI.

4.7. Thus the CMA'S discussion of historic average returns centres on Table 9-3 (p. 549), which presents a range of estimates of what the reasonable reader might infer are CPI-deflated real market returns using different methodologies—the more so since the results in each column are so similar.[14] In fact (at least as far as we can tell) this is not the case. The left column is used to derive an estimate of CPI-deflated returns which are in turn used to derive equivalent forward-looking RPI-deflated returns. The right column is used to derive an estimate of forward-looking RPI-deflated returns *directly*; later on, this estimate is used to derive an estimate of forward-looking CPI returns.

4.8. Thus, from the left column, the CMA takes the geometric mean of CPI-deflated returns, 5.2%, and adds to it an uplift of 1.2% to adjust to an arithmetic mean of 6.2%, which gives the bottom end of their range for the forward-looking CPI-deflated return. From this, the CMA derives the implied forward-looking RPI-deflated real return by adjusting downwards by 90 basis points. Thus far, in line with the approach we advocated.

4.9. But the derivation of the upper end of the range effectively gives equal weight to RPI deflation. From footnote 1406, the CMA uses the right-hand column of Table 9.3 to derive an estimate of forward-looking RPI-deflated returns *directly* from the historic data: i.e., a geometric mean of 5.0%, plus the same uplift of 1.2%, giving a figure of 6.2%. Eight pages later, this RPI-deflated figure is then transformed to a (spuriously precise) equivalent CPI-deflated return of 7.21%.[15]

4.10. It is worth noting that when we compare this supposedly CPI-based real return estimate with the left-hand column of Table 9-3, which shows historical real returns estimates using actual CPI data, only a single estimate (out of 12 shown) is (almost) as high as this figure. This single figure, 7.2%, is[16] the arithmetic average of the 6 observations of the non-overlapping 20-year return, a figure on which the

[14] We note also that this Table does *not* report a data source, nor provide any link to a full description of its methodology. This is one of numerous examples of a lack of replicability of the CMA's analysis.
[15] The apparent extra 11 basis points arise because these adjustments are multiplicative, rather than additive: i.e., 1.05*1.012*1.009=1.0721. We are grateful to PJ McCloskey for elucidating this point.
[16] Or, strictly, appears to be: the precise methodology is again not actually described.

CMA itself comments "Due to the small sample size we place less weight on these results" (note to Table 9-3).

4.11. This figure is in turn higher than the simple arithmetic mean of annual CPI-deflated returns shown in the top row of the left column of Table 9-3, of 7.0%. And *this* figure is in turn higher than all other figures in the same column, other than the non-overlapping return estimate. Nor is this surprising, since all the alternative methodologies are predicated on the property that the sample arithmetic average return is an upper bound for the true expected return. Thus the gap between the arithmetic mean and geometric mean is 1.8 percentage points respectively, close to the upper end of the broad range of adjustments to geometric returns of 1 to 2 percentage points advocated in Mason, Miles and Wright (2003)[17] and again in the UKRN report.

4.12. As far as we can tell, the upper end of the range chosen by the CMA can only be derived by giving equal weight to RPI deflation of historic average returns, despite the well-documented problems with the RPI. In contrast, if the CMA had estimated CPI-deflated historic returns directly, then these estimates would lie within the range 6.1% to 6.9%.[18]

Ex-ante and forward-looking estimates

4.13. Without delving too much into the detail of these alternative methodologies, we simply note the comments the CMA itself makes in paragraph 9.221, that its final chosen range "… is comfortably at the top end of investors' current expectations regarding market returns over the next few years". In broad terms we would concur with this point, but it is worth noting that the range of plausible *lower* values could in principle be very much lower.

4.14. In this context, it is worth stressing that while the UKRN report advocated the use of historic averages as the best method of estimating the expected market return that was both *defensible* and *implementable*, this should *not* be taken as implying the judgment that such estimates will necessary be unbiased at any particular point in time. Indeed, in the UKRN report we presented a range of evidence that, we concluded, suggested quite strongly that market expected returns were below, and quite possibly significantly below, historic averages. These arguments almost certainly have even more force in the current conjuncture. Our concerns were that if regulators attempted to take this into account, i.e., applied discretionary downward adjustments to their EMR estimates, this would potentially

[17] Mason, R.A., D. Miles and S. Wright (2003). A Study into Certain Aspects of the Cost of Capital for Regulated Utilities in the U.K..

[18] We have no particular quarrel on the issue of which historic price index to use to attempt derive consistent CPI-based real return estimates. As stressed in the UKRN report, the methodology for producing the RPI has changed on multiple occasions; in some early sub-samples of data it is at least possible that the RPI may been a better proxy for what we would now call the CPI. But we note only that viewed in this light, the CMA's figures in the right-hand column of Table 9-3 imply *reduced* historic real return estimates, which might well therefore be used as an argument to push the range *down*wards, rather than upwards.

open the floodgates to disputes about the precise magnitudes of these adjustments.

4.15. It seems odd that, despite acknowledging the likely upward bias of historic average returns, in light of forward-looking considerations, the CMA has ended up with an upper limit of its range so obviously at, or beyond, the range even of historic return averages. This choice, in turn, has a disproportionate influence on their final chosen figure, given their choice of the notional 75% percentile.

5. The equity beta

5.1. We should make clear that our views of beta estimation set out below are broadly unchanged from those we expressed in the UKRN report (in particular in Appendix G, co-authored with Donald Robertson), and thus Ofwat and CMA are arguably closer in their approaches to beta estimation than we are to either. But we present our views here for completeness, since they feed into our overall conclusion (on which we *do* agree with Ofwat) that CMA has been overly generous in setting the WACC.

5.2. The core approach taken by Ofwat to the determination of the equity beta (the "Harris-Pringle" approach) is as follows:

5.2.1. Estimate econometrically the "raw" equity beta.

5.2.2. Calculate an unlevered beta: remove the effect of gearing, assuming a debt beta of zero and using observed gearing.

5.2.3. Estimate the debt beta.

5.2.4. Calculate an asset beta, using the unlevered beta, the estimated debt beta and observed gearing.

5.2.5. Calculate a relevered equity beta, using the asset beta, estimated debt beta, and notional gearing.

Ofwat's estimates and calculations, taken from Table 5.10 of their cost of capital technical appendix[19], are shown in Table 1.

Table 1: Ofwat's estimates and calculations for equity betas

	Estimates and calculations
Raw equity beta	0.63
Observed gearing	54.2%
Unlevered beta	*0.29*
Debt beta	0.125
Asset beta	*0.36*
Notional gearing	60%
Re-levered beta	*0.71*

Italicised parameters are calculated from estimated parameters in non-italics.

5.3. The CMA also follow the Harris-Pringle approach. Any differences with Ofwat therefore arise from the values of the components in the approach. The CMA prefer

[19] Ofwat, July 2019. PR19 final determinations: Cost of capital technical appendix.

to present unlevered equity betas as their estimated parameters, although these are clearly based on estimated raw betas. In Table 2, we therefore infer the raw equity betas that the CMA has used. For comparability and transparency, we would prefer the CMA to be more explicit about their preferred raw equity beta estimates.

Table 2: CMA's estimates and calculations for equity betas

	Low	**High**
Raw equity beta	*0.59*	*0.70*
Observed gearing	54.2%	54.2%
Unlevered beta	0.27	0.32
Debt beta	0.15	0.00
Asset beta	*0.35*	*0.32*
Notional gearing	60%	60%
Re-levered beta	*0.65*	*0.80*

Italicised parameters are calculated from estimated parameters in non-italics.

5.4. Given the speed with which this report needed to be produced, we were not able to attempt any direct estimation. We have however been able to draw on estimation results provided for us by PJ McCloskey, on behalf of Ofgem, for which we are extremely grateful.

5.5. As already set out in Section 2, we argued in the UKRN report that ideally there should be a consistency between the horizon over which all parameters of the CAPM are estimated. Clearly there are limitations to the degree of consistency that is possible, since we have considerably longer datasets for market returns data than we do for the returns on regulated companies. But, subject to these limitations, there is a strong case for taking account of as much data as is available. We also have a strong preference for focussing in particular on the directly estimated equity beta, rather than the constructed unlevered and asset beta series. These may provide interesting insights, but we need to bear in mind that they are constructed by assumptions, rather than estimation.

5.6. Figure 6 gives the longest available perspective. It shows rolling equity beta estimates over all available data for United Utilities and Severn Trent. The underlying dataset starts in 1996, but the chart shows rolling beta estimates with a window of two years. We should stress that rolling beta estimates are a legitimate diagnostic tool for addressing the issue of whether the true (and unobservable) beta is stable over time. In contrast, if the true beta is assumed *not* to be stable over time, rolling betas are fraught with problems as estimators of this time-varying value at any point in time – and most notably standard errors (whether OLS or heteroscedastic-consistent) are spurious. Thus we simply present rolling point estimates.

Figure 6: rolling equity beta estimates over all available data for United Utilities and Severn Trent

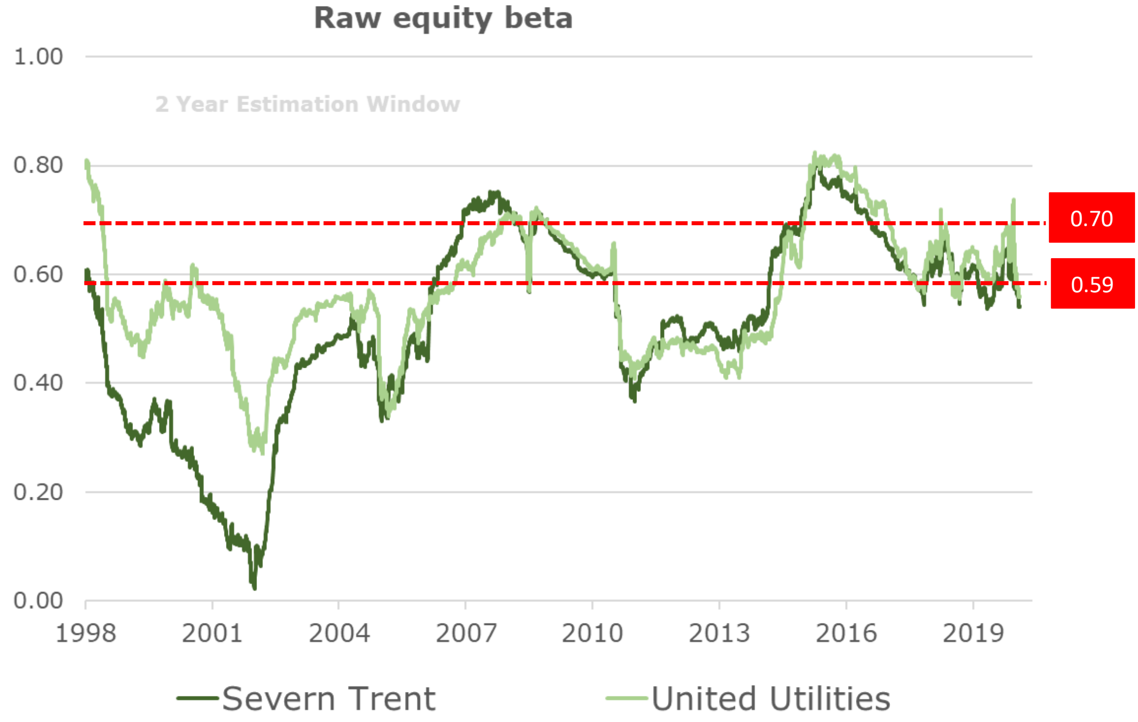

Note: The red dotted lines show the CMA's low (0.59) and high (0.70) estimates of raw equity betas, calculated from their low (0.27) and high (0.32) values of the unlevered beta and the observed gearing level of 54.2%.

5.7. This chart is revealing in several ways. First, on a simple eyeball test, there seems little evidence of parameter instability. Long-run stability appears particularly evident in the case of UU, but even for SVT there has been considerable stability since the mid-2000. [20] Strikingly, also, for the past fifteen years or so the two beta estimates have been virtually identical: essentially the equity market appears to view the two companies as having identical systematic risk.

5.8. The second striking feature of Figure 6 is that the CMA's chosen range of values for raw beta, of roughly 0.6 to 0.7, is very much at the upper end of the range of estimates shown in the chart. Indeed, average values of beta for the two companies over the period since end-2003, the point from which the two betas essentially moved together, are somewhat *below* the value CMA picks as the lower end of its range;[21] and Figure 6 shows that estimated betas have only rarely moved *above* its chosen range. Thus, as in other contexts, even before the CMA applies its aiming up correction by picking the 75th percentile, we argue that it has picked a range that is biased upwards. (We also argue that Ofwat has also been too generous in its own estimates.)

[20] This feature is also consistent with results of "GARCH" estimation, discussed further in footnote 18, below.
[21] To be precise, the averages of two year rolling betas from October 2003 to October 2020 are 0.59 for SVT and 0.59 for UU. If the full sample (starting two years earlier, to allow for the two-year window) is used for OLS estimation, the estimates are 0.58 for SVT and 0.60 for UU.

5.9. Pursuing our preference for looking at directly observable data, an obvious cross-check on the apparent stability of equity beta for the two companies is to look at gearing estimates for the two companies, shown in Figures 7 and 8.

Figure 7: gearing estimates for United Utilities

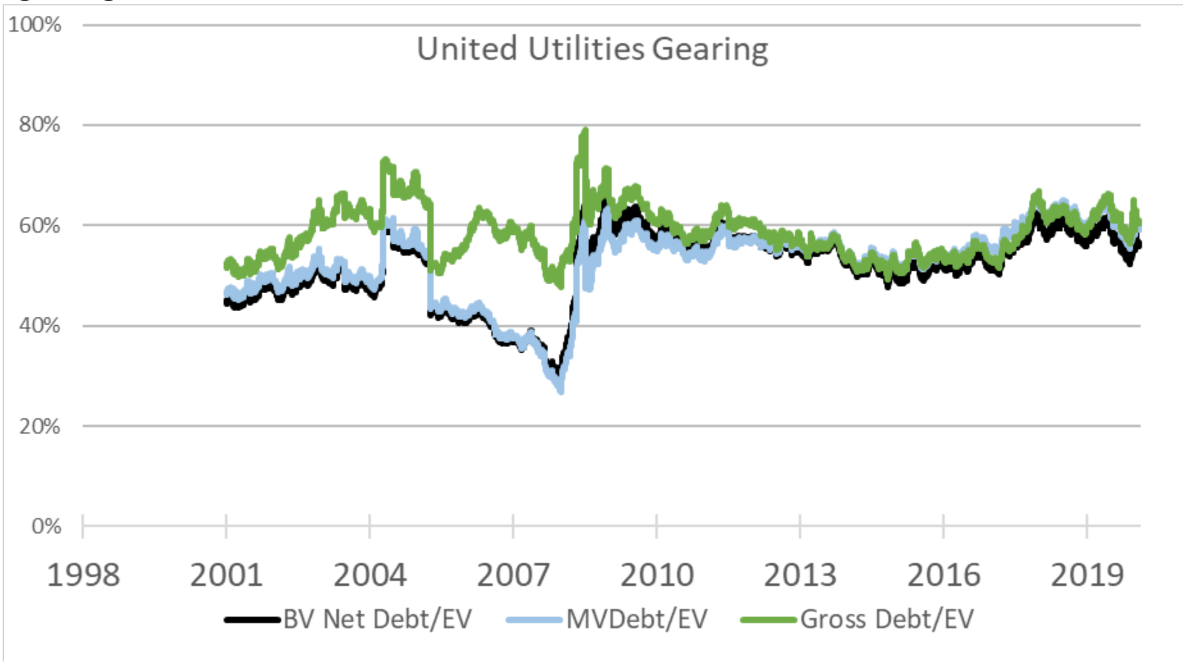

Figure 8: gearing estimates for Severn Trent

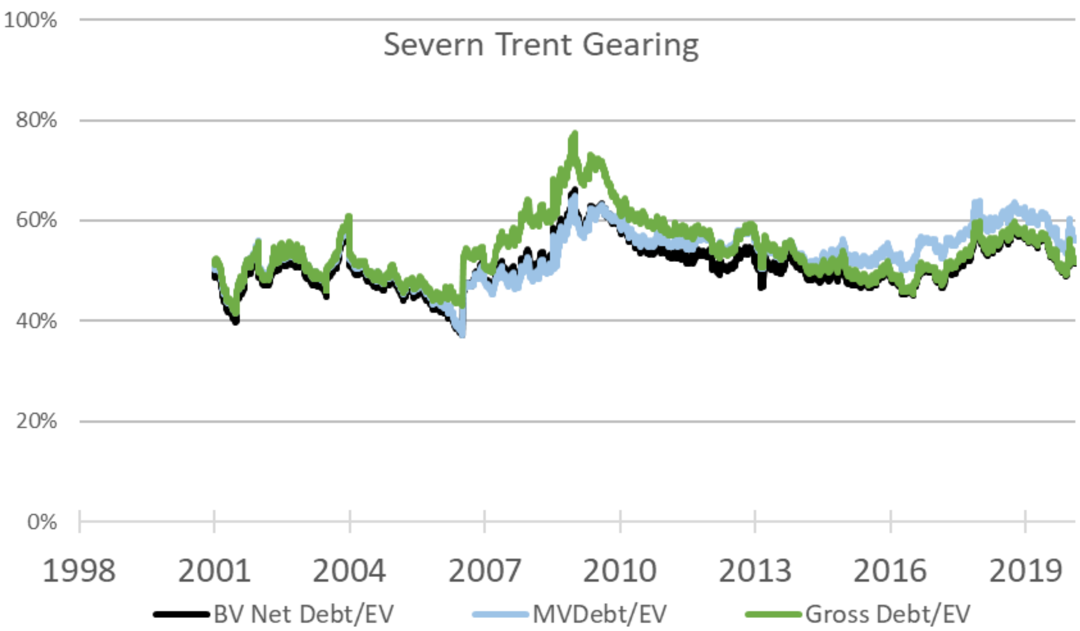

5.10. Both companies show considerable stability in gearing, on three different measures. There is admittedly some element of upward drift, but it is noticeable that the timings of upward movements in gearing are not obviously reflected in the timing of upward movements in equity beta.

5.11. Finally, Figure 9 shows asset beta estimates, constructed using book value of debt and an assumed debt beta of 0.15, the upper value of the CMA's chosen range. While it is important to bear in mind that these numbers are a constructed series, and therefore far from being hard data, they nonetheless provide an interesting cross-check. Again, there seems to be evidence of long-run stability, and, again, the CMA's chosen range of 0.27 to 0.4 looks clearly upward biased.

Figure 9: asset beta estimates for United Utilities and Severn Trent

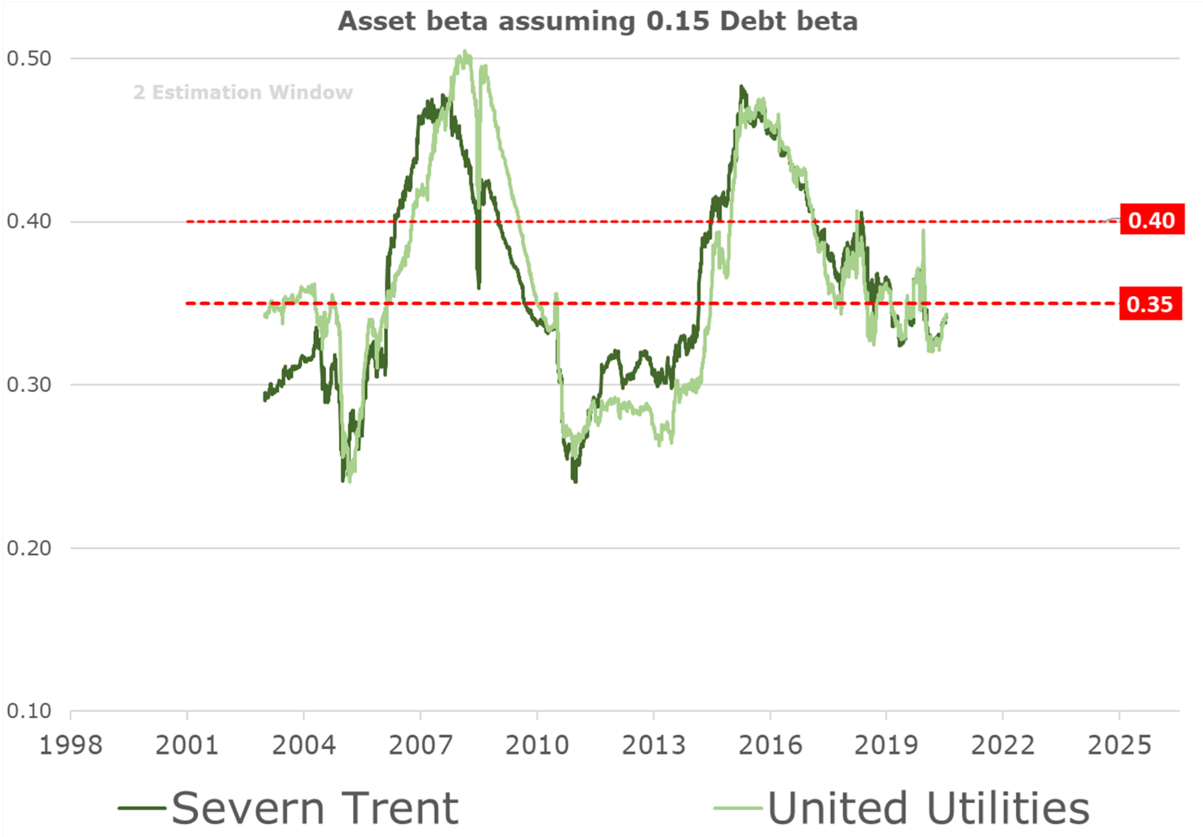

Note: The red dotted lines show the CMA's low (0.35) and high (0.40) estimates of asset betas, recalculated with a debt beta of 0.15.

5.12. To conclude, while we have not attempted any precise quantification, we would argue that, consistent with the CMA's chosen long-term horizon, it makes sense to use as much data as possible to estimate equity beta, and hence derive estimates of asset beta. There would clearly be problems in this approach if there was clear evidence of any long-term drift in beta, but we do not see any such evidence. And, given the use of a longer sample of data, we see clear signs of an upward bias in the way the CMA (and, by implication, Ofwat) have picked their ranges for beta. [22]

[22] We also note that, by contrast, at least one other regulator, Ofgem, in setting its own beta estimates, has taken account of both longer samples, and GARCH estimation (as we also advocated in the UKRN report), which tends to further reduce beta estimates. GARCH estimation takes into account the well-established empirical result that the volatility of both asset and individual stock returns varies over time—in particular, with a tendency for volatility to rise during periods of crisis or information shocks, and then fall back during calmer periods; but also with a tendency to revert to a stable long-term value. This can be captured by

6. Debt: embedded and new

6.1. Ofwat and the CMA disagree on the costs of both new and embedded debt. It is only the latter that is material: given the assumed proportions of new debt issued and the cost difference, the differing views on the cost of new debt make a difference of less than 3 basis points to the WACC.

6.2. On embedded debt, the two material points of difference between Ofwat and the CMA are:

6.2.1. The estimation window: 15 years (Ofwat) versus 20 years (CMA).

6.2.2. The use of an "outperformance wedge" by Ofwat.

Our discussion will focus on these two issues.

6.3. We note, as we have previously in the UKRN report, that we think regulators should start from a position of being sceptical about allowing for the costs of embedded debt, for the simple reason that unregulated companies do not receive this kind of insurance from their customers. We understand the argument that this allowance reduces risks for regulated companies, which may in turn lower the cost of debt—and possibly also the cost of equity—and hence the WACC itself. (We also understand that the risk is not perfectly reduced for the firms, since their embedded debt cost recovery is relative to an index. But as we note in the UKRN report, that just means that regulated firms have to be a little clever, exploiting the arbitrage opportunity by borrowing a share of its regulated debt allowance precisely in line with the construction of the weighted average.) But in the end, it is the customer who pays, and so the risk is simply passed on to them. It is far from obvious to us why this should be desirable.

6.4. Starting from this position provides an alternative to the trailing average, or trombone, currently used by regulators. The trombone anchors embedded debt costs in previous periods where the cost of debt was much higher. In contrast, our prior position anchors the cost of embedded debt to the cost of new debt: which both Ofwat and the CMA acknowledge is much lower. (Indeed, the gap is considerably greater for the CMA: a difference of 2.39%, compared to 1.89% for Ofwat.) This is not to say that firms should be either encouraged or discouraged to take on short-term debt (although elsewhere the CMA seems agnostic about the financial decisions of firms). It simply reflects the fact that non-regulated firms have

General Autoregressive Conditional Heteroscedasticity (GARCH) estimation, where the term "conditional" reflects the feature that variations in volatility are conditional upon the time period in which it is measured. This in turn means that the *conditional* beta (the ratio of a conditional *co*variance to a conditional variance) also varies over time, but, in turn, reverts to a stable long-term value. The details of this approach were set out in Appendix G (co-authored with Donald Robertson) of the UKRN report, where we showed that GARCH estimation produces results that are consistent not just with short-term fluctuations in rolling beta estimates as presented in Figure 6, but also with long-run stable beta values. Subsequent work by Donald Robertson for Ofgem has shown that when these long-run values of beta are estimated directly using GARCH techniques, this also results in somewhat lower beta estimates than OLS estimates, on the same samples. See, for example, Table 15 of Ofgem's RIIO-2 Draft Determinations (Finance Annex), https://www.ofgem.gov.uk/system/files/docs/2020/07/draft_determinations_-_finance.pdf#page=46.

to face the risk without compensation that the cost of debt will change, and we do not see a compelling reason why regulated firms should not also face this risk.

6.5. The problem created by allowances for embedded debt are compounded by the fact that historically, regulators have over-estimated the cost of debt, as Figure 11 (taken from the UKRN report) demonstrates. Put simply, a trailing average/trombone over a period of falling interest rates has led to persistent over-estimates of the cost of embedded debt.

Figure 11: real market cost of debt v. regulator assumptions

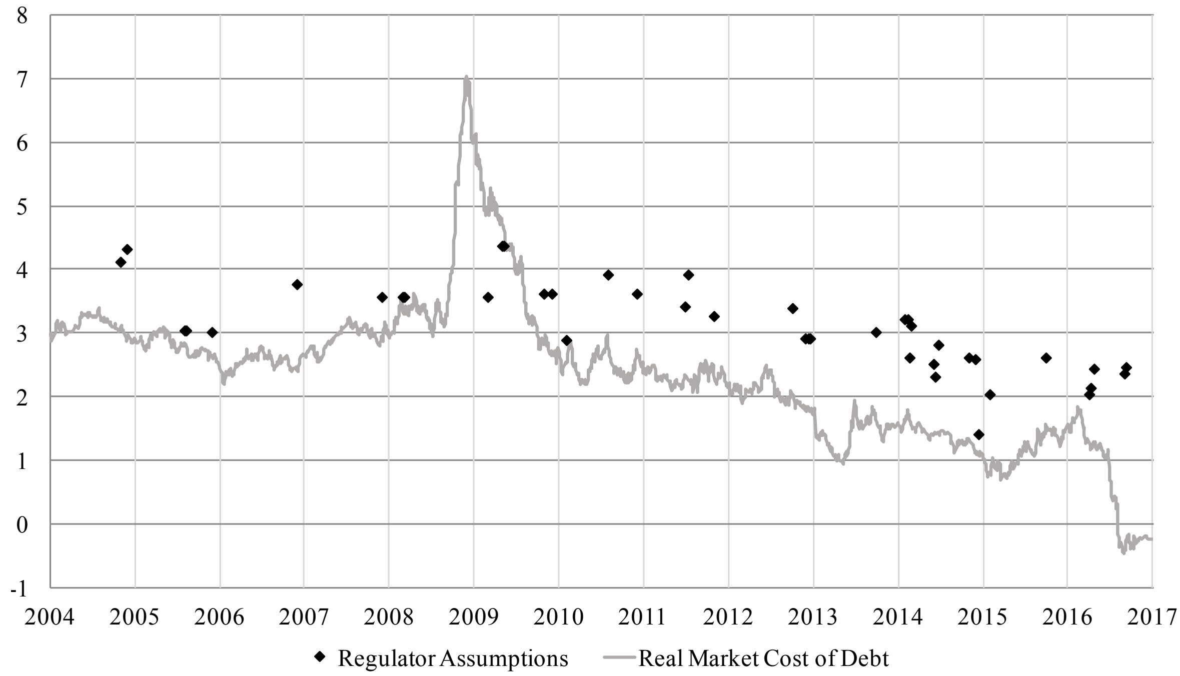

6.6. In light of this, given that we are of the view that embedded debt should not be compensated, our response to the first issue—should the trailing window be 15 or 20 years—is: the trailing window should be zero. If, however, a trailing window has to be used, then ensure that it does not compound the error of allowing for embedded debt in the first place. This would point towards a shorter trailing window; or, if that proves unpalatable, a weighted trailing window with weights derived from the maturity profile of the sectors' embedded debt.

6.7. The second area of major disagreement between Ofwat and the CMA concerns the application of an "outperformance wedge": the observation by Ofwat that water companies' debt tends to be cheaper than the iBoxx A/BBB index by 25 basis points historically. The CMA is persuaded by evidence presented that once tenor and credit ratings are accounted for, there is no evidence for a wedge (paragraph 9.352 of the CMA report).

6.8. Considerable evidence was produced by the disputing companies and their advisors on this question. We do not, therefore, have any empirical evidence to add to the debate. In our view, the difference between Ofwat and the CMA springs from a misunderstanding of the phrase “outperformance wedge”—which we think is misleading. It seems to us that Ofwat and the CMA agree on facts:

6.8.1. The water companies' actual historic cost of debt has been below the index. This is Ofwat's position, but the CMA seems to accept it: or at least, has not produced evidence to counter it;

6.8.2. Once the water companies' debt is compared to debt of similar maturity and credit-rating, it does not appear to be systematically lower. This is the CMA's point (see e.g., paragraph 9.352).

6.9. It seems to us that there are three options to reconcile these two facts:

6.9.1. Continue to use the iBoxx A/BBB index unadjusted and anticipate that the firms’ cost of debt will be lower than it, as it has been in the past. (This is the CMA’s approach.)

6.9.2. Discontinue use of the index and create an index that more closely matches the maturities and credit ratings of the water firms' embedded debt.

6.9.3. Continue with the index but make an adjustment to reflect the historic gap between actual cost of debt and the index. (This is Ofwat’s approach.)

6.10. It seems to us that Ofwat has actually chosen the third option, but have perhaps not made this clear in choosing to call their adjustment an “outperformance wedge”. Using this label has led to an analysis of whether the water companies have actually outperformed other companies with similar characteristics. Rather, the question is: what adjustment needs to be made to the index in order to ensure that it more closely matches the regulated water companies. The facts described in points 6.8.1 and 6.8.2 above show that the index does not match perfectly the key features of the regulated water companies (in particular, tenor and credit rating) that determine the cost of their embedded debt. As such, what Ofwat calls the “outperformance wedge” would be better viewed as a “matching adjustment” to the index.

6.11. The need to adjust the index points to option 2: the construction of an index that is a better match for the water sectors’ tenor and credit rating. This may be something that needs to be considered in the future. For now, the use of the iBoxx A/BBB index is well accepted and it is unlikely to be worth stepping away from it at this point in the price determination process.

6.12. In our view, the first option is the worst of the three. Effectively it requires water customers to reward shareholders because the regulators cannot agree a way to ensure historic debt is costed correctly. A possible counter to this view is that over the long run, it will all “come out in the wash”: periods where the firms’ embedded debt is cheaper than the index will be balanced by periods when it is more expensive. We have concerns about this view. Given the importance of embedded debt for the WACC, and the importance of the WACC in financeability tests, we suspect that in periods where the cost of embedded debt is above the index, financeability will require the higher figure to be used. In periods where the embedded debt cost is below the index, again the higher figure would be used. In short, the asymmetry arising from financeability would prevent it coming out in the wash.

6.13. Our view, therefore, is that there is sufficient evidence for a "matching adjustment" to be applied to the index, of around 25 basis points at this review point. We expect the value of the adjustment to be assessed periodically during the price determination process.

7. Aiming up

7.1. As noted in the UKRN report, aiming up has been regulatory practice in a number of different countries for a number of years. In that report, we give an analytical framework explaining why aiming up may be desirable for regulators. We also give some sense of quantification, concluding that while a case can be made for aiming up, "that case is, in our view, a limited one: more limited than appears to have been adopted in a number of past regulatory decisions in the UK. The case is limited further by the extent to which regulators are able to incentivise investment through means other than setting the RAR" (p. 73).

7.2. The CMA's approach to aiming up is as follows. In its view, the uncertainty surrounding cost of equity estimates is greater than that surrounding the cost of debt. It therefore applies aiming up to the former but not the latter. It chooses to aim for the 75th percentile for each component of the cost of equity calculation: see Table 3. (For the cost of embedded debt, it takes it low estimate; for the cost of new debt and the proportion of new debt, it chooses mid-points.)

Table 3: CMA's low, high and point estimates of key components of the WACC (CPIH-real)

Element	Ofwat PR19	CMA low	CMA high	CMA point estimate
TMR	6.50%	6.20%	7.21%	6.95%
RFR	-1.39%	-1.40%	-0.81%	-0.96%
Unlevered beta	0.29	0.27	0.32	0.31
Debt beta	0.125	0	0.15	0.04
Calculated re-levered equity beta	0.71	0.65	0.80	0.76
Cost of new debt	0.53%	0.21%	0.52%	0.37%
Cost of embedded debt	2.42%	2.76%	3.16%	2.76%
Proportion of new debt	20%	21%	13%	17%
Calculated pre-tax cost of debt	2.14%	2.32%	2.92%	2.45%
Calculated post-tax cost of equity	4.19%	3.56%	5.60%	5.08%

Note: components coloured yellow are at the 75th percentile; blue are at the 50th percentile; green is at the 0th percentile (i.e., lower bound).

7.3. The CMA's rationale for aiming up is given in paragraph 9.665ff. Summarising, the CMA appears satisfied that there are sufficient incentives for investment in new assets in AMP7; but is concerned about investment in the water sector more broadly; asymmetry of returns created by PCs and ODIs; and indications from Ofwat's financeability checks, where a degree of accelerating revenue was employed.

7.4. We have a number of concerns with the CMA's approach:

7.4.1. There is an unacknowledged source of upward bias in the CMA's calculation of the "75th percentile".

7.4.2. We do not agree that the CMA has "not tried to aim up or down when setting the individual metric estimates" (paragraph 9.663). We argue that at key points, the CMA already aims up when estimating the ranges for the CAPM components.

7.4.3. We do not think that there is evidence for broader concerns about investment in the water sector; or that these concerns should materially affect the setting of the WACC in AMP7.

7.4.4. We do not think that adjusting the WACC is the appropriate way to deal with concerns about asymmetry of returns created by performance incentives.

7.4.5. Being of the view that even Ofwat has been too generous in setting the WACC, we are of the view that the CMA has overshot by a wide margin.

7.5. Our first concern is that there is an unacknowledged source of upward bias in the CMA's calculation of the "75th percentile" for the cost of equity. In calculating its 75th percentile, the CMA implicitly makes two strong assumptions:

7.5.1. For each component of the cost of equity, there is uses a uniform distribution within its parameter ranges.

7.5.2. Each component of the cost of equity is uncorrelated with all others.

Both are strong assumptions; neither is stated explicitly by the CMA as forming the basis of its calculations. We will examine each assumption in turn.

7.6. First, we also note that there is in principle a difference—even assuming independent uniform distributions—between aiming up on each of the cost of equity components, instead of aiming up on the final cost of equity. As we argued in the UKRN report, aiming up should ideally be applied at the final stage, once the cost of capital has been calculated. It so happens that in this case, it makes little difference: around 1 basis point, as the CMA notes in footnote 1738. This is demonstrated algebraically in the Appendix to this report, which makes clear that this outcome is highly dependent on the implicit assumption of independent uniform distributions.

7.7. We now turn to the two assumptions of uniform and independent distributions for parameter values. It is not uncommon to assume uniform distributions; nevertheless, it is worth noting that this choice assumes that all values in the range are equally likely. At a number of points, however, the CMA has chosen upper values for ranges that are supported by very few data points. One example is when estimating the TMR, where in paragraph 4.10ff, we argue that the upper bound of 7.2% is supported by very few data points, and moreover is above the theoretical upper bound given by the sample arithmetic average return. Indeed, as we note, even the CMA recognises this when saying "Due to the small sample size we place less weight on these results" (note to Table 9-3). The use of a uniform distribution is acceptable when there is good reason to place equal weight on the top and bottom; it is an error when there is good reason to doubt equal weights. The effect of this is to over-state e.g., the 75th percentile. If, for example, we were to assume a normal distribution centred on the midpoint of the range, an assumed standard deviation of a quarter of the range (i.e., 0.25) would mean that the 75th percentile lies at 6.87%. This point is roughly two-thirds across the range, rather than three-quarters, leading to a gap of around 8 basis points. The value that the CMA takes to be the 75th percentile with a uniform distribution is the 85th percentile with a normal distribution.

7.8. The effect of using a different distribution becomes larger when calculating the overall cost of equity, because uncertainty about each component compounds. To understand conceptually why this is the case, consider the cost of equity:

$$\tilde{R}_E = \widetilde{RFR} + \tilde{\beta}(\widetilde{TMR} - \widetilde{RFR})$$

where the ~ denotes a random variable. Hence the cost of equity, $\tilde{R}_E$, is a combination of (a) the sum of random variables; and (b) the product of random variables.

7.9. Looking first at the sum of random variables, consider the 75th percentile of a random variable z that is the sum of two other random variables x and y: $z = x+y$. If both x and y are uniformly distributed, then the 75th percentile of z is simply the sum of the 75th percentiles of x and y: the uniform distribution is very special in this regard. Suppose instead that x and y are standard normally distributed i.e., with mean 0 and standard deviation 1. The 75th percentile of each of the variables x and y is 0.674; and hence the straight sum of their 75th percentiles is 1.349. In contrast, the 75th percentile of the variable z is 0.954: considerably less. This is because z is itself normally distributed with mean 0, but with a standard deviation of $\sqrt{2}$. Simply adding the individual 75th percentile of x and y assumes that the resulting sum is normally distributed with mean 0 and standard deviation 2. In fact, the sum of the variables has a lower standard deviation, due to the mathematics of adding normal distributions.

7.10. Now consider the product of random variables: $z = x\ y$. Again, if both x and y are uniformly distributed, then the 75th percentile of z is simply the product of the 75th percentiles of x and y: once again, the uniform distribution is a special case. But if x and y are standard normals, the distribution of z is more concentrated than either of the constituent distributions of *x* and *y*.[23] The 75th percentile of *z* is 0.365, whereas the product of the 75th percentiles of *x* and *y* is 0.455.

7.11. These effects are significant. For each component of the cost of equity, assume that, rather than a uniform distribution, a normal distribution holds, with a mean equal to the mid-point of the CMA's low-high range and a standard deviation equal to a quarter of the range (so that the range corresponds to two standard deviations either side of the mean). The resulting cost of equity has a 75th percentile value of 4.78%.[24] The value, 5.08%, that the CMA takes to be the 75th percentile with uniform distributions is the 95th percentile with normal distributions.

7.12. The situation becomes significantly more complicated when the second implicit assumption—that each component of the cost of equity is independent of the others—is varied. Here, we simply note that the practice of estimating the total market return (TMR), which assumes that the TMR is constant, means that the TMR and the RFR are uncorrelated. In contrast, the CMA's calculations assumes that they are positively correlated: the low (high) value of the RFR is used with the low (high) value of the TMR to calculate the low (high) value of the cost of equity. But there has been little work done on how to account for correlations between the estimators of the components of cost of equity in the CAPM, and so we draw no conclusions at this point.

7.13. Our second concern is that the CMA has been generous throughout in its estimates of the ranges of the components of the cost of equity. We have argued elsewhere in this report that the CMA has:

- 7.13.1. Set a range of values of the TMR higher than is higher than can be readily supported by the evidence;
- 7.13.2. Over-estimated the risk-free rate, owing to inconsistent use of the CAPM;
- 7.13.3. Over-estimated the levered equity beta, in the way that it has estimated betas and used debt betas in the re-levering calculation;
- 7.13.4. Been generous in estimating the cost of embedded debt (which is more important than the cost of new debt, where admittedly the CMA's estimate is lower than Ofwat's).

7.14. We argue that there is implicit or explicit recognition by the CMA of at least some of these points. For example, in setting its range for the TMR, the CMA notes that its range "is comfortably at the top end of investors' current expectations regarding market returns over the next few years" (paragraph 9.221), as we observed earlier. Admittedly, the CMA's objective is not to estimate investors' expected returns over the next few years: it is to use of the TMR at its chosen 20-year investment horizon. Furthermore, in the same paragraph, the CMA notes that

[23] The product of two normal distributions is not a normal distribution: instead, the distribution is proportional to a Bessel function of the second kind: see https://mathworld.wolfram.com/NormalProductDistribution.html.

[24] Since the distribution of the random variable determining the cost of equity is complicated, it is easier to perform this calculation by drawing a large number of random variables and averaging. The figure of 4.78% as the 75th percentile is derived by drawing 1 billion random variables, in a process known as Monte Carlo simulation.

it has adjusted upwards the range used by the CMA in its recent CAA/NATS decision, to reflect "the uncertainty over the accuracy of the available inflation data series". That is, the adjusted range already reflects uncertainty.

7.15. As we detail in Section 3, the CMA has, through an inconsistent application of the CAPM, over-estimated the risk-free rate. The key argument is that the marginal investor for UK water companies is a net lender, and hence the relevant zero-beta asset is the ILG. (We also argue that the CMA is incorrect to use aiming up to determine the weights used to calculate the zero-beta asset's expected return.)

7.16. As we noted in Section 5, it appears to us that the CMA's chosen range of values for the raw equity beta, of 0.59 to 0.70, is very much at the upper end of the range of estimates shown in the chart. We argue that the same point holds in the CMA's treatment of debt betas. In paragraph 9.314, the CMA states that Ofwat "provide[s] a compelling case that the regulatory model should include a positive debt beta". But in paragraph 9.315, the CMA appears to set its lower bound for the debt beta in light of "significant calculation uncertainties associated with debt beta".

7.17. These points all add up. By our calculation, taking the mid-points (i.e., 50th percentile) of the CMA's ranges, the calculated cost of equity would be 4.55% i.e., 0.36% above Ofwat's point estimate of 4.19%. The remaining 0.53%—to get to the CMA's final figure of 5.08%—comes from aiming up at the final stage. That is, around 40% of the CMA's higher cost of equity figure comes from lifts in the individual components; the remaining 60% comes from the final uplift.

7.18. Turning now to evidence for broader concerns about investment in the water sector: none is given as far as we can see. Paragraph 9.668 in the CMA's report states "we note that the most common decision has been that some 'aiming up' has been merited in order to promote investment in the sector, and that there may be benefits to consistency – including ensuring investor confidence in the sector". In the next paragraph, the CMA notes countervailing evidence from Ofwat based on market-asset ratios (MARs). In Figure 12, we provide updated data on MARs for United Utilities (UU) and Severn Trent (SVT), using data supplied to us by Ofwat. The Figure shows that there continues to be a significant premium over the regulatory capital value for these two firms: certainly no evidence that we can see of why there should be a broader concern about investment in the sector.

Figure 12: Premium to regulatory capital value (%)

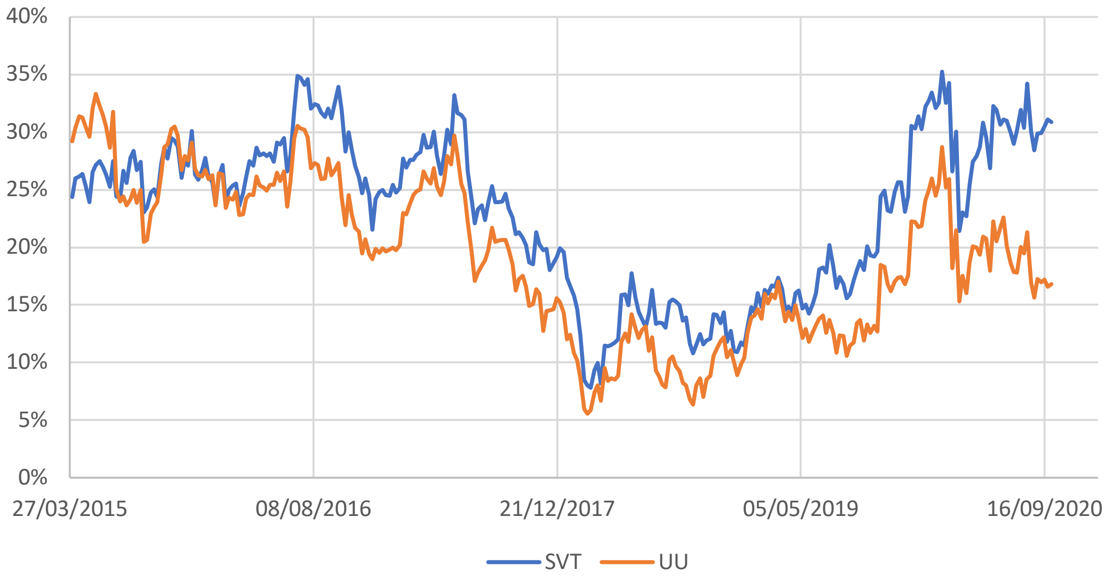

Source: data supplied by Ofwat[25].

7.19. No further arguments are given on this point by the CMA. In our view, that falls well short of sufficient reason for aiming up. In addition, since the CMA is content that there are sufficient incentives to invest in new assets in AMP7, it seems to us that if the CMA has residual concerns about investment, then these could and should be reflected in the next price determination—that, after all, is the point of having regular regulatory reviews.

7.20. The final reason advanced by the CMA for aiming up is a concern about asymmetry of returns created by performance incentives. The CMA largely agrees with Ofwat's package of penalty-only and asymmetric outcome delivery incentives (ODIs). It estimates that the asymmetry could lead to an average performing company facing a potential loss of around 0.1% to 0.2% impact on the return on regulatory equity (RORE). For the disputing companies, the RCV is around 2.5 times greater than the regulatory equity. Hence the 0.1-0.2% effect on regulatory equity is equivalent to a 0.04-0.08% effect on the RCV. By comparison, the CMA's own estimate of aiming up—notwithstanding earlier points we have raised about upward biases—increases the WACC by 0.21% (the difference between the WACC calculated using midpoint estimates, and the WACC using 75th percentiles for the components of the cost of equity). In our view, this is a disproportionate response to reductions in expected returns coming from performance incentives. More importantly, we fundamentally disagree with an approach that, through a higher WACC, compensates firms—indeed, more than compensates them—for missing performance targets that have been agreed to: this seems a perverse approach to regulation.

7.21. In summary: we can see no merit in any of the CMA's arguments relating to aiming up. We think the CMA has anyway aimed up as it has gone along, both

[25] Data kindly provided by Martin Malinowski of Ofwat.

implicitly in its approach, and in setting ranges for the components of the cost of equity which already are high. The CMA present no real evidence for broader concerns about investment in the water sector; what evidence there is shows very healthy premia over RCVs. Its approach to compensating firms for missing performance targets is both perverse and disproportionate. As we will discuss in Section 9, we are of the view that Ofwat has been too generous in setting the WACC at 2.96%. We therefore disagree strongly with the CMA's equivalent figure, based on midpoint estimates, of 3.30%. We see no grounds at all for the CMA's final figure of 3.50%.

8. The gearing outperformance sharing mechanism

8.1. Ofwat's gearing outperformance sharing mechanism (GOSM) requires companies whose gearing exceeds a certain threshold to make a payment based on the gap between the allowed nominal cost of equity and the actual cost of debt. The motivation behind the mechanism is to disincentivise, but not prohibit, higher gearing levels, as part of the regulator's concerns about rising debt levels and reduced financial resilience.

8.2. The CMA objects to the GOSM for a number of different reasons:
- 8.2.1. The effectiveness of a whole-sector approach, particularly given a range of regulatory protections already explicit within the licence conditions (paragraphs 9.588 and 9.590);
- 8.2.2. Insufficient evidence about the benefits that firms receive from increased gearing (paragraph 9.610(a));
- 8.2.3. Setting a regulatory precedent without adequate risk assessment (paragraphs 9.610(b) and 9.628).

8.3. In this Section, we focus on the second of these reasons; and in particular, the CMA's statements that theory (in particular, the Modigliani-Miller, or MM theory) suggests that there are no benefits from gearing, since the WACC is or should be broadly unaffected by gearing (see paragraphs 9.605 and 9.627): "It is consistent with corporate finance theory that the cost of equity will increase with gearing, and the GOSM that Ofwat has implemented, which assumes the cost of equity is broadly stable with gearing above a certain level, is not consistent with this theory" (paragraph 9.627).

8.4. The CMA is, of course, right that the GOSM marks a change in approach by a UK regulator. Until this water price review, the practice has been to draw a sharp divide and to view the financial structure of regulated companies as outside of the regulatory remit. The rationale for that has been partly theoretical i.e., the MM theory itself. It has also been partly for lack of a precedent. In so far as the CMA is arguing that the lack of a precedent militates against the GOSM, then we disagree. Indeed, we would argue that the UK regulators have often led in the adoption of new, eventually to become best, practice. Of course, regulators must have due regard to past decisions, both by themselves and other regulators. But the CMA has demonstrated in a number of places in this appeal that it appears to feel no constraint in setting aside previous decisions that it disagrees with or it feels are not merited in this particular case.

8.5. It is legitimate, however, for the CMA to look for more evidence or justification than usual in making or accepting a regulatory precedent: this is what it asks for in paragraph 9.627. The critical evidence for the CMA on this issue concerns the dependence of the WACC on gearing, and hence departures from the MM world of invariance of the WACC to capital structures.

8.6. In some ways, this is a strange question. The point of the MM theory is to demonstrate what must hold in order for the WACC to be invariant to financial structure. The theory is not and was not intended to be a description of reality, but to elucidate the channels by which financial structure matters for the value of a firm. The key assumptions (taken from Berk and DeMarzo, 2014) required for the MM theorem to hold are:

8.6.1. Investors and firms can trade the same set of securities at competitive market prices equal to the present value of their future cash flows;

8.6.2. There are no taxes, transaction costs, or issuance costs associated with security trading;

8.6.3. A firm's financing decisions do not change the cash flows generated by its investments, nor do they reveal new information about them.

8.7. Just the statement of these assumptions is enough to demonstrate that they will rarely, if ever be satisfied. A long literature has developed to examine the consequences of departures from the MM assumptions. In short, very few who work in the field believe that the MM theorem actually holds. Instead, they use it as a useful thought experiment and to identify the factors that matter for determining the value of a firm. It is true that in the UK's regulated water sector, the treatment of taxes means that companies are not able to increase returns by increasing gearing to outperform their tax allowance. Nevertheless, there are plenty of other departures from the MM theorem, even for these firms.

8.8. Nevertheless, regulators do make much reference to the MM theorem; and in using the CAPM as the basis for setting the cost of equity, rely on it both implicitly and explicitly. At the same time, regulators recognise that there is a problem: the WACC that they calculate is not independent of gearing, contrary to MM. The CMA stated this very explicitly in its NATS decision: "given the parameters that we have estimated for the cost of capital, this approach [the standard one used by UK regulators] has the unexpected effect of resulting in the WACC strictly increasing with gearing" (paragraph 13.112). It is the case here: Figure 13 shows how the CMA's calculated WACC for the water companies increases with the notional gearing.

Figure 13: The CMA's calculated WACC (CPIH-real, %) against notional gearing

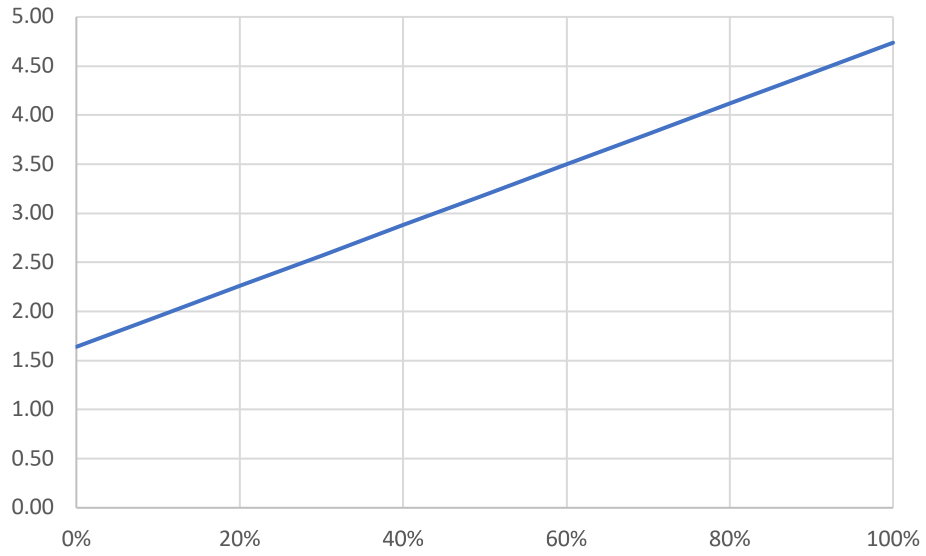

8.9. The presence of regulation can itself provide a reason why the MM theory does not hold. For example, Spiegel and Spulber (1994 and 1997)[26] argue that if regulators cannot commit to a particular regulatory scheme, then regulated firms will have an incentive to finance their investments with debt. A regulator who cannot commit to long-term regulated prices has an incentive to cut prices once the firm's investments are sunk, benefitting consumers at the expense of the firm's owners. But if the firm finances its investments with debt, and if the regulator is concerned about the financial stability of the industry it regulates, then the regulator will set higher regulated prices than it would have done otherwise in order to minimize the risk of financial distress. In short, in the Spiegel and Spulber model, debt financing is a defence against regulatory opportunism. Bortolotti et al. (2011)[27] provide empirical evidence to support this theory from a range of regulated EU utilities over the period 1994-2005. They find that firms in their sample tend to have a higher leverage if they are privately-controlled and regulated by an independent regulatory agency. They also find that leverage Granger-causes regulated prices (but not vice versa): that is, gearing affects prices, but prices do not affect gearing.

8.10. Aspects of the Spiegel and Spulber analysis are, we think, less suited to the current UK context. In particular, given the maturity of the regulatory system in the UK, we think concerns about regulatory opportunism are less pronounced than they would have been when Spiegel and Spulber first wrote their papers. Nevertheless, we argue that the core point—that higher gearing is encouraged by regulation that features financeability concerns—remains valid, for two reasons.

[26] Spiegel Y. and D. Spulber (1994). "The Capital Structure of a Regulated Firm," *RAND Journal of Economics*, 25(3), 424-440. Spiegel Y. and D. Spulber (1997). "Capital Structure With Countervailing Incentives," *RAND Journal of Economics*, 28(1), 1-24.
[27] Bortolotti, B., C. Cambini, L. Rondi and Y. Spiegel (2011), Capital Structure and Regulation: Do Ownership and Regulatory Independence Matter? *Journal of Economics & Management Strategy*, 20: 517-564.

8.11. First, while there is a special administration regime in place should a UK water company be in financial distress, we suspect that there are strong incentives for the regulator to avoid use of this regime. The monopoly position of water companies in England and Wales means that it is not possible for most customers to switch to another provider in the event of failure of a water company. While the regime is intended to minimise disruption to customers, this cannot be guaranteed. The fact is that the special administration arrangements have not been used to date. If correct, then this provides a reason why increased gearing does not have the effect that it does in the MM theory. In that theory, the WACC remains constant as gearing rises because the cost of equity rises by just enough to offset the lower cost of debt. If there is implicit insurance offered by a regulator with financeability concerns, then there is less risk to equity and hence the WACC can be lowered through increased gearing.

8.12. Secondly, we argue that the treatment of embedded debt has encouraged higher gearing. We have already noted that this treatment provides regulated companies with a privilege not enjoyed by unregulated companies. In our view, the allowed cost of embedded debt has generally exceeded the actual cost of that debt: see Figure 11. Assumption 1 of the MM theorem requires that financial assets are traded at competitive market prices; the theorem does not apply when the price of a financial asset is set away from the competitive level. In our view, the treatment of embedded debt has created a clear incentive for regulated firms to increase their gearing.

8.13. These two factors are important, in comparison to e.g., Spiegel and Spulber's argument. In their argument, higher gearing can increase overall welfare, since it is a means to counter regulatory opportunism and hence prevent socially undesirable under-investment. As we have argued, we think regulatory opportunism is not now a primary concern in the UK system. The incentives to increase gearing arise from the benefits to investors, not consumers.

8.14. What estimates are there of the size of benefits that UK water firms gain from higher gearing? On this specific question, there is as yet limited evidence. There are empirical studies about the benefits to leverage: for example, Korteweg (2010)[28] and Van Binsbergen et al. (2010)[29] estimate the net benefit of debt to be at most 10% of the value of a firm. But these studies do not extend readily to the specifics of the UK regulated water sector. The observations that we have made about market-asset ratios in Figure 12 are suggestive, but we cannot say how much of the observed premia relate to the allowed WACC being lower than firms' true cost of capital, as opposed to expected outperformance against performance targets. (See also the discussion in the next Section.)

8.15. We think it highly suggestive, however, that agencies indicate gearing levels that are consistent with particular credit ratings. For example, Moody's has

[28] Korteweg, A. (2010). "The Net Benefits to Leverage". *The Journal of Finance*, 65(6), 2137-2170.
[29] van Binsbergen, J., J. R. Graham and J. Yang. (2010). "The Cost of Debt". *The Journal of Finance*, 65(6), 2089-2136.

indicated that a net debt-to- regulated asset base ratio of at least 55-70% is needed for a credit rating of Baa or higher. [30] This does seem to indicate that if a water firm wishes to go beyond 70% gearing, it is jeopardising its investment credit grade and must therefore see particular benefits to increasing its gearing.

8.16. In this report, we do not provide further assessment of and comment on the fine details of the GOSM. To summarise our view: regulators have for some time now recognised that the MM theory does not actually hold—that the financial structure of a regulated firm matters for value of that firm; and that the gearing of the firm affects its WACC. Moreover, it is certainly that case that credit-rating agencies consider gearing when assessing ratings. In addition, in a number of cases, perhaps starting with the Competition Commission's decision on Bristol Water's appeal in 2010[31], regulators have proposed finance structure responses to financeability issues. The MM theory alone is therefore clearly not sufficient grounds to reject regulatory concern about financial structure.

9. Concluding remarks

9.1. In this report, we have focussed on the main areas of disagreement between Ofwat and the CMA. In this concluding Section, we offer some broader comments beyond just these differences. The thrust of the comments is that Ofwat has likely over-estimated the WACC. In increasing its midpoint estimates, and then aiming up on top of that, the CMA has certainly over-estimated the WACC.[32]

9.2. We start by putting to one side embedded debt and look to find an upper bound for regulated firms' WACC when only new debt is considered. In fact, we even avoid considering new debt by looking at a firm that has no debt: only equity. We will calculate this equity-only WACC. In a Miller-Modigliani world, adding debt should make no difference to the firm's cost of capital; the equity-only cost of capital would then be the WACC. But, as we argued in the previous Section, we do not live in an MM world: there are (net) benefits to taking on debt. Assuming that firms do not take on sub-optimal levels of debt, this should mean that if an equity-only firm gears up, its WACC should fall, at least initially. Hence, the (equity-only) CAPM cost of equity should represent an upper bound for the firm's true WACC. This simple but powerful argument is illustrated in Figure 14.

[30] For example, see Moody's (2018). Regulated Water Utilities: Rating Methodology.

[31] Competition Commission (2010). Bristol Water plc: A reference under section 12(3)(a) of the Water Industry Act 1991.

[32] The comments that follow are explained in greater depth in the UKRN report.

Figure 14: WACC against gearing when there are net benefits to debt

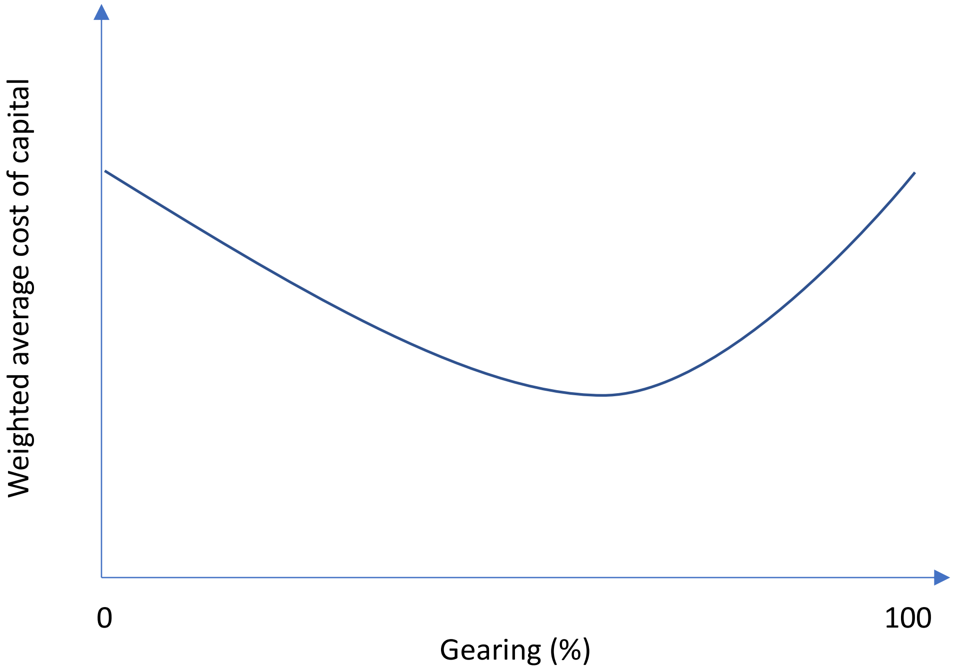

9.3. The equity-only WACC is very easy to calculate using the CAPM: it is

$$RFR + \beta_A(TMR - RFR)$$

where β_A is the asset beta i.e., the beta of a zero-debt firm. It therefore needs only three inputs: the risk-free rate, the expected market return, and the asset beta. In Table 4, we present the values implied by alternative estimates of the key parameters (as shown above in Table 3). Even if we use the CMA's final point estimates for all three parameters (all of which we argue were based on ranges that erred on the upside even *before* applying the "aiming up" adjustment), the resulting estimate—1.63%—is very significantly below any WACC proposed by Ofwat or the CMA, calculated on a consistent basis.

Table 4: alternative calculations of the WACC

Element	Ofwat PR19	CMA low	CMA high	CMA point estimate
TMR =E(R_M)	6.50%	6.20%	7.21%	6.95%
RFR	-1.39%	-1.40%	-0.81%	-0.96%
Asset beta (β_A)	0.36	0.26	0.41	0.33
Equity-only WACC	**1.44%**	**0.58%**	**2.48%**	**1.63%**
Calculated post-tax cost of equity (R_E)	4.19%	3.56%	5.60%	5.08%

Cost of new debt (R_D)	0.53%	0.21%	0.52%	0.37%
WACC with new debt (=0.6*R_D+0.4*R_E)	**1.94%**	**1.53%**	**2.50%**	**2.22%**

9.4. The Table also calculates the WACC with just new debt, using the costs of equity and debt determined by Ofwat and the CMA. While it is higher than the equity-only WACC—reflecting the paradox that debt appears to be "expensive" in CAPM terms, and yet has been used as the dominant means of finance by regulated companies since privatisation—it is still markedly lower than Ofwat's or the CMA's point estimates of the WACC.

9.5. So our starting point is that the WACC is likely to be materially lower than has been allowed by either Ofwat and the CMA. And this is before the regulators' treatment of embedded debt which, we have argued in this report as well as elsewhere, extends a privilege to regulated companies not enjoyed by unregulated companies, and further inflates the allowed WACC. A key point that we make here is that the CMA has used its discretion to push its estimates upwards in a range of other ways, but has entirely ignored the arguments for pushing them downwards.

9.6. We view market prices as evidence for our argument. Figure 15, extending some of the data displayed earlier, shows the estimated market premium over RCV for Severn Trent over an extended period.[33] The premium has been positive for the great majority of the period since privatisation. In recent months it has reached levels close to historic peaks.

Figure 15: estimated market premium over RCV for Severn Trent

[33] Data kindly provided by Martin Malinowski of Ofwat.

9.7. Logically, this premium must reflect a gap between companies' expected return and their perceived cost of capital. A simple but informative calculation is to consider this premium as the proportional difference between each of these returns, if expressed in perpetuity-equivalent form. (If the difference is not expected to persist in perpetuity, the gap must be larger.) This gap could in principle be driven by two factors: the allowed WACC being lower than firms' true cost of capital; or expected outperformance against performance targets.

9.8. To illustrate: at one extreme, taking the most recent premium of 28%, consider the possible case that markets believe that Ofwat will succeed in future in controlling the water companies' historical outperformance, and thus views Ofwat's allowed return of 2.96% as SVT's *expected* return (in perpetuity-equivalent terms). This would imply that markets view SVT as having a perpetuity-equivalent WACC of only 2.31% (i.e., 2.96%/1.28). Note that since this is expressed in perpetuity-equivalent terms, it is consistent with markets' *spot* estimate of the WACC being lower, but with an inbuilt expectation that it will rise over time. This is indeed consistent with the current profile of the index-linked forward curve, which is upward-sloping until a maturity of around 25 years. Thus, we argue that this calculation could in principle quite easily be consistent with our equity-only WACC estimates set out above.

9.9. In the opposite extreme case, markets have the same estimate of SVT's WACC as used by Ofwat. In this case, markets must be building in a perpetuity-equivalent outperformance gap of 83 basis points (i.e., 0.28*2.96%). Again, since this is expressed in perpetuity-equivalent terms, this is consistent with a larger gap in the short term, if markets are building in an expectation that at least at some point regulators will succeed in limiting or eliminating systematic outperformance. If this interpretation is correct, it raises the obvious question of why the CMA recommends aiming up, if markets are already building in a higher expected return.

9.10. We do not attempt to make a judgement about which interpretation of the current market premium is correct. Nor do we need to assume that markets are necessarily getting it right. But the consistency of market premia seems to us to be something that regulators should not simply ignore. It certainly calls into question the CMA's provisional findings to soften Ofwat's price determination.

Appendix

Algebraic analysis of the difference between aiming up component-by-component or at the end

In the cost of equity calculation, there are 4 key parameters: the total market return, denoted R_M; the risk-free rate, denoted R_f; the unlevered beta β_U; and the debt beta β_D. The latter two parameters are combined to calculate the levered beta, β_L. Each of these components has a range of values: the lower value for e.g., R_M is denoted R_{M1}; the upper value is denoted R_{M2}; and likewise for the other components. Treating these parameters as random variables, assume that they are independently and uniformly distributed between their lower and upper values. The lower and bound costs of equity are then

$$R_1 = R_{f1} + \beta_1(R_{M1} - R_{f1})$$
$$R_2 = R_{f2} + \beta_2(R_{M2} - R_{f2})$$

Suppose that in calculating a point estimate, the lower bound is given weight α and the upper bound is given weight 1-α. The point estimate is then $\alpha R_1 + (1-\alpha)R_2$. Simplifying this gives

$$\alpha R_{f1} + (1-\alpha)R_{f2} + \beta_{L1}\left(\alpha(R_{M1} - R_{f1}) + (1-\alpha)(R_{M2} - R_{f2})\right) + (1-\alpha)(\beta_{L2} - \beta_{L1})(R_{M2} - R_{f2}).$$

Alternatively, we can calculate a point estimate component-by-component:

$$\alpha R_{f1} + (1-\alpha)R_{f2} + (\alpha\beta_{L1} + (1-\alpha)\beta_{L2})\left(\alpha(R_{M1} - R_{f1}) + (1-\alpha)(R_{M2} - R_{f2})\right).$$

Simplifying this gives

$$\alpha R_{f1} + (1-\alpha)R_{f2} + \beta_{L1}\left(\alpha(R_{M1} - R_{f1}) + (1-\alpha)(R_{M2} - R_{f2})\right) + (1-\alpha)(\beta_{L2} - \beta_{L1})\left(\alpha(R_{M1} - R_{f1}) + (1-\alpha)(R_{M2} - R_{f2})\right).$$

The magnitude of the difference between these two midpoint estimates is therefore

$$\left|\alpha(1-\alpha)(\beta_{L2} - \beta_{L1})\left((R_{M1} - R_{f1}) - (R_{M2} - R_{f2})\right)\right|.$$

Given the values involved in the calculations, this number is small:

$$\frac{1}{4} * \frac{3}{4} * 0.15 * 0.41 = 0.0115.$$