A report on financial resilience, gearing and price controls

Prepared for Ofwat by Professor Robin Mason and Professor Stephen Wright

3 December 2021

1. Introduction

1.1. Ofwat have asked us to provide economic advice on the potential ways to treat issues of gearing and financial resilience, within price controls and otherwise. The following report gives our views. It has two parts. The first considers the economic rationale for why financial structure can be of concern to regulators; some of the issues that arise when considering financial structure; and some of the options available to address financial structure, if necessary. The second part considers the treatment of gearing in price controls: in particular, problems arising with the estimation of the cost of capital.

1.2. The main purpose of this report is to provide an outside perspective and challenge to current regulatory approaches. It should not be taken as a statement of policy by Ofwat, and nor does it necessarily reflect the views of anyone other than the authors, who have prepared this report in an independent capacity.

1.3. It is standard in UK regulation to leave capital structure decisions to regulated firms. A typical stance in setting allowed returns (at least prior to PR19) is for *actual* financial structure not to be taken into account explicitly; instead, gearing is based on a *notional* gearing assumption. The rationale is three-fold. The first is that, since there is little consensus regarding theories of optimally efficient gearing, regulators typically do not consider calculations of optimal or efficient gearing to be feasible. (Indeed, according to the Modigliani-Miller Capital Structure Irrelevance theorem, there is no optimal capital structure.) The second is to encourage efficient financing by leaving the financing decisions and responsibilities to companies; implicitly, firms are therefore assumed to have superior information on efficient financing; ensuring that they bear the risks associated with financing decisions. Thirdly, and relatedly, the use of a (reasonable) notional gearing avoids the potential to set the cost of equity for regulated businesses by reference to the capital structure of a single, unusual comparator.

1.4. So for example, as stated by the CMA in their PR19 determination final report:

> "9.39 Both Ofwat and the CMA calculate the allowed return on capital with reference to a notional company with a predetermined level of gearing. This notional approach allows companies to make their own choice about their financial structure whilst ensuring that customers only pay for costs associated with the efficient cost of capital for a notionally structured company."

1.5. But there are (at least) two problems with this stance. First, it is not strictly true: capital structure decisions are not left entirely to the management of regulated firms and there are various conditions in place which may constrain managerial choices. Secondly, even in cases where the *level* of gearing is not a specific concern, the *treatment* of gearing is: specifically, the procedure for unlevering and relevering equity betas.

1.6. On the first point, the most obvious constraint on managerial capital choices are licence conditions that require regulated companies to maintain investment-grade

credit ratings. For water companies, Ofwat requires companies to maintain an investment-grade credit rating: a requirement set in light of Ofwat's statutory duty under the Water Industry Act 1991, which includes the duty to secure that water companies are able to finance those functions, in particular, by securing reasonable returns on their capital.[1] Ofwat has also stated an expectation that efficient companies should maintain headroom within the investment grade category; this is also illustrated in its determinations, which most recently focussed on an expectation that the notional structure should achieve Baa1/BBB+ or above (or that justification for Baa2 should be provided). The implication (from Moodys' guidance) of attaining a Baa1 rating is to indicate a limit on RCV gearing (i.e., net debt relative to regulatory capital value, RCV); the limit depends on other financial resilience measures (e.g., covenants that are in place). In the case of NERL, discussed in more detail later in this report, there is a direct cap on its gearing within its licence.

1.7. In short, it is not correct to assert that regulators have left the choice of financial structure entirely to the management of regulated firms. In practice, of course, the regulatory constraints on financial structures may not actually bind; or may be such that firms have sufficient flexibility to manage their finances without risk of breaching the licence conditions. Nevertheless, the constraints are there, and clearly help to shape expectations and behaviour.

1.8. In this report, we look at the question of financial structure and regulation in two distinct but related ways, each related to the problems highlighted above. First, we assess the general case for regulatory concern about financial structures, assessing a market failure that underpins this concern. Identification of the market failure is important for determining the appropriate regulatory responses; we provide a number of different options and their pros and cons. As the previous discussion indicates, there exist already measures to address the market failure; we also consider, therefore, the extent to which existing measures adequately address any market failure.

1.9. Secondly, we look at greater length at the now-standard procedure for moving from the actual gearing of a particular comparator firm or firms, to notional gearing, and the various manipulations of the equity beta. The issues arising with this procedure are now fairly well known; despite this, there is no standard regulatory response to them (other than to note and then ignore them). Again, we provide a number of different options and their pros and cons.

1.10. Before we proceed to analyse these issues in depth, by way of background, we first review some recent regulatory decisions and their implications for these two issues.

---

[1] Strictly speaking, two companies—South West Water and Hafren Dyfrdwy—are currently exempt from the requirement to maintain an investment grade credit rating. Their licences contain provisions for their Board to certify on an annual basis that in the Board's opinion, they "would be able to maintain an issuer credit rating which is an investment grade rating".

We then note an important question about gearing: namely, how it should be defined and measured, and the uncertainty involved in the measurement. We then move on to consider the two issues in detail.

2. Recent regulatory decisions

2.1. In this section, we discuss three recent regulatory decisions and appeals and the particular implications of each for gearing and by extension, financial structures more generally:

1. The NERL appeal to the CMA.
2. The PR19 appeal to the CMA.
3. The RIIO-2 appeal to the CMA.

The NERL appeal

2.2. Since this is discussed in more detail below, our discussion here will be brief. The main gearing issue that arose in the 2019/20 referral from the CAA to the CMA for for a determination on the price controls of NATS is that the weighted average cost of capital (WACC) increases with gearing, when the by-now standard procedure for calculating the WACC is used. (The gearing cap in place on NERL was not at issue in the appeal.) While identifying the problem, the CMA chose not to address it directly. Instead, it opted for a lower notional gearing level than argued for by either the CAA or NERL. The CMA noted that the comparator airports used to estimate NERL's raw equity beta had gearing at around 30%; and that NERL itself had gearing less than 30%. With these observations, the CMA set the notional gearing at 30% (rather than 60% as used by the CAA and NERL). In doing so, the CMA avoided the need to unlever and relever equity betas, and so avoided the feature that the WACC would be increased by using a notional gearing above actual, observed levels. As the CMA stated in §§13.117-118 of its final report:

> "... we decided to set the cost of capital based on a gearing of 30%. This allowed us to use the comparator airports to estimate the equity beta without needing to materially adjust their gearing, since they had similar gearing to that of NERL. We then assumed that it was for NERL to decide whether or not to actually increase gearing to 60%, but we assumed it would only do so if it resulted in a lower cost of capital... We therefore concluded that this approach did not risk us underestimating NERL's cost of capital at its optimal gearing level, whatever NERL ultimately decided was the most efficient financing structure."

2.3. The CMA's approach in the NERL case therefore sets something of a precedent in two ways. First, it indicates that the notional gearing level can be chosen in a way that reduces the problem of a WACC that is not invariant to the level of gearing. (Note however that in the NERL case, the choice of a lower notional gearing was supported by observed gearing levels.) Secondly, it uses an argument that if actual

gearing is different from notional gearing, it must be as a result of firms seeking to decrease their cost of capital. Put differently, this is to argue that a regulated firm is always free to choose to set its gearing at the notional level, in line with the regulator's WACC calculation. If it chooses to depart from this, then it must because doing so is beneficial to the firm.

The PR19 appeal to the CMA

2.4. The main gearing issue that arose in the 2020/21 PR19 appeal to the CMA was Ofwat's proposal to use a gearing outperformance sharing mechanism (GOSM). The choice of the notional gearing level appeared to be uncontentious (Ofwat lowered it slightly from PR14 from 62.5% to 60%). The dependence of the WACC on gearing also did not feature as an issue.

2.5. The CMA was not persuaded by the details of the GOSM. It appreciated the grounds for regulatory concern (e.g., §§9.1194-5). It concluded, however, that

- 2.5.1. There are existing means by which Ofwat can address a concern about excessively-high gearing; and evidence that these means have been, at least partly, effective.
- 2.5.2. The nature of water companies is such that high levels of gearing are sustainable (§9.1202).
- 2.5.3. There is limited or no evidence of consumer harm having occurred in the past in the cases of Enron/Wessex Water and Hyder/Dŵr Cymru (§9.1209); special administration procedures and other licence conditions protect customers so that "These factors do not indicate that customers are likely to bear disproportionately the costs of financial failure" (§9.1210).
- 2.5.4. There is limited or no evidence of (excess) benefits from high gearing structures: for example, in §9.1215, the CMA states "Overall, we have not seen evidence that shows that highly geared companies have achieved overall WACCs substantially lower than the notionally geared company".
- 2.5.5. With an appropriately set cost of capital, firms are financed to invest sufficiently so that customer harm is avoided (§9.1218).

2.6. We return to these issues below when assessing options for addressing concerns about high gearing.

The RIIO-2 appeal to the CMA

2.7. The CMA's summary of its final determination does not indicate any concerns about gearing—either the level, or the effect of unlevering and relevering for cost of equity calculations—in the RIIO-2 Energy Licence Modification Appeals. The fact that it was not a large issue in the appeal may in part be explained by the fact that Ofgem's approach led to a higher estimate of the cost of equity, since their notional gearing of 60% was somewhat higher than observed gearing, which tends to be around 50%

for energy network companies. Ofgem also assumed that the cost of capital was identical at the two different notional gearing levels that it used: 55% for SHE-T, NGET and SPT; 60% for all the others. Since the calculated WACC increases with gearing (as will be discussed later), this benefitted those companies with lower notional gearing.

2.8. Nevertheless, the procedure of unlevering and relevering betas was raised by four of the appellants (Cadent, NGET, NGG and SHE-T; SPT, NGN, SGN and WWU did not object), although not in any great detail, and typically as part of a broader argument as to why the (relevered) equity beta and consequently the estimated cost of equity (and in turn, the regulatory allowed return) should be higher than set by Ofgem.

2.9. We think there is one further aspect of the RIIO-2 appeal that, while not directly related to gearing, is significant. The CMA quashed Ofgem's proposal to apply a -25 (i.e., downward) basis point adjustment to the cost of equity (the "outperformance wedge") in setting the regulatory allowed return: the one cost of capital point on which Ofgem did not succeed. The CMA stated in §28:

> "(c) Given the problems identified in (a) and (b) [errors in the outperformance database; and poor design of the mechanism], there was a realistic possibility that the outperformance wedge, if introduced, might also undermine broader regulatory certainty which could result in increased costs to consumers over time."

3. Measurement challenges

3.1. It is not gearing per se that regulators should be concerned about, nor indeed the overall financial resilience of regulated firms as such; it is the risk that higher gearing and/or weaker financial resilience may increase the risk of social costs from financial failure that are not fully taken into account by management (see below for further discussion of this).

3.2. No single indicator determines financial resilience: hence Ofwat and credit-rating agencies look at a range of measures e.g., interest cover, covenant conditions, etc., themselves correlated with gearing. Nevertheless, gearing is clearly an important metric for determining the equity buffer to withstand cost and income shocks, and so we focus on it specifically in this section. (We consider other aspects of financial resilience in the next section.)

3.3. Despite the frequency of its use, the measurement of gearing seems, to us at least, to be far from straightforward, for at least three reasons. First, should gearing be defined according to book or market values? Secondly, what liabilities should be included when estimating gearing? Thirdly, how should uncertainty in the estimation of gearing be taken into account when estimating equity betas?

3.4. Actual gearing can be defined in a number of ways. For example, the three considered by Ofwat in PR19 for the purposes of delevering and relevering equity betas (see the PR19 Draft Determinations Cost of capital technical appendix) were

1. 'Enterprise value approach': unlevering using the ratio of the book value of net debt to enterprise value,[2] and relevering using the notional gearing assumption.
2. 'Book value approach': unlevering using the ratio of the book value of net debt to RCV, and relevering using the notional gearing assumption.
3. 'Indepen approach': unlevering using the ratio of the book value of net debt to enterprise value, and relevering using the notional gearing assumption divided by an estimate of the market-to-asset ratio (MAR) for the notional company. Indepen use a figure of 1.183 as their estimate.

3.5. These different approaches have significant implications for relevered betas, as Table 3.11 from the cost of capital technical appendix (an extract of which is below) shows:

	Enterprise value approach	Book value approach	Indepen approach
Raw equity beta	0.64	0.64	0.64
Gearing of listed comparator	54.7%	62.8%	54.7%
Asset beta	0.36	0.32	0.36
Notional gearing	60%	60%	55%
Relevered beta	0.71	0.60	0.64

3.6. Ofwat opted for the 'enterprise value approach' in PR19, in which the MAR is greater than 1 when unlevering, but is assumed to equal 1 when relevering. The (implicit) rationale for this seems to be the idea of a "notional" firm is one in which both the firm and the regulator are efficient; in the latter case, this means, inter alia, that the cost of capital is set correctly and all other regulatory allowances are achievable, so that the MAR is equal to 1.

3.7. Our view of this issue is closer to Indepen's. In practice, we know that observed MARs are significantly greater than 1. Some part of this, we would argue, is a result of the regulatory allowed return being set generously (i.e., above firms' true costs of capital). Some part also results from outperformance by regulated firms. This outperformance is indeed an intrinsic and, we would argue, desirable feature of the RPI-X regulatory framework in the UK: the economic rents from outperformance provide incentives for firms to innovate and reduce costs, while meeting output targets. These rents may be lower in water than in other sectors, e.g., because of the greater availability of yardstick competition. The rents are unlikely, however, to be zero—and nor would we want them to be. In the terminology of the UKRN 2018

[2] See below for further comment on the enterprise value used.

report, the informational wedge is and should be always positive. Hence a MAR somewhat above 1 is desirable in principle.[3]

Book and market values

3.8. Even in Ofwat's 'Enterprise value approach', book values are used. (That is, the enterprise value is taken as the market value of equity plus the book value of net debt.) This use of book values needs to be questioned, for two reasons. First, since it is the expected market return that features as the dependent variable on left-hand side of a CAPM-based regression, consistency with the rest of the WACC exercise means that it should be gearing defined in terms of market values that matters. More precisely, the theoretical underpinning for the process of unlevering and relevering relies on the standard result in finance that the beta of a portfolio is equal to the market-weighted average of the betas of the elements of that portfolio. This result in turn relies on the property that all elements of the portfolio are liquid, quoted securities; and only holds exactly at any point in time if market weights are updated continuously.

3.9. So, the standard approach, based on debt measured at book value, is likely to be misleading: it is equivalent to assuming that all debt is on a floating rate. This is clearly inconsistent with the assumed liability structure of the notional firm, which underpins the embedded debt calculation. Ideally, book value debt should always be converted to market value, in a way that is fully documented. The effect of this correction is ambiguous, since it depends on whether the market value of debt is greater (in which case gearing increases) or less (decreasing gearing) than its book value. (If the observed comparator firm *does* have debt for which book value is close to market value, it also suggests a possible weakness as a comparator for the notional firm, which has long-dated debt.)

3.10. A second issue arises: most notably, that other liabilities could arguably be reflected in net debt. The leading contenders are the liabilities of defined-benefit pension schemes (which are still common in the sector), and (mark-to-market valuations of) derivatives.[4] We are, of course, aware of previous debates about whether e.g.,

---

[3] One possible objection to this approach is that incentives might lose power, since expected returns, which are increased as a result of information rents, are reduced by the use of a lower beta, and hence regulatory allowed return. Our view is that while expected returns may be affected, the marginal incentive to outperform targets and generate additional returns will remain (broadly) the same. A second objection is that the notional firm, by definition, will exactly meet performance targets; will therefore not gain additional expected returns; and so will have a MAR of 1. We disagree with this definition of the notional firm. In standard economic models of adverse selection, only the most inefficient firm will obtain zero information rents; all other firms—including, we argue, the notional firm—gain positive information rents in an optimal regulatory scheme with asymmetric information.

[4] For example, the use of index-linked swaps to match borrowing costs to inflation-linked revenue streams may create risks in relation to market or counterparty exposure. In addition, since

pensions should be included in gearing measures: see for example NERA (2017). We are also aware of the complications in estimating pension deficits. Our purpose here is not to repeat these previous (and indeed, ongoing) debates. Instead, our main point is that the appropriate measure of gearing is far from settled—a fact that is typically overlooked in regulators' cost of capital calculations, which tend to treat gearing as completely certain.

Gearing and equity betas

3.11. We will say more below on the procedure whereby "raw" equity betas are unlevered from actual gearing levels and relevered to notional gearing levels. At this stage, we comment on the issues arising due to uncertainty about gearing levels.

3.12. We first note that the use of a notional gearing creates a clear and quite strong incentive for regulated companies to attempt to persuade the regulator and the CMA that the gearing of the comparator companies that are used to estimate raw equity betas (i.e., SVT and UU) is low. As we note below, in PR19 the effect was around 58 basis points on the cost of equity, or 15bp on the overall cost of capital. We have previously advised Ofwat that it should factor in evidence more explicitly that points in the opposite direction.

3.13. It is important to reflect the uncertainty about gearing inherent in the relevering exercise. When actual gearing is measured with certainty, uncertainty in the estimate of the relevered beta comes from only one source: uncertainty about the raw equity beta. Uncertainty arising from the measurement of gearing compounds this, but is typically ignored in the un- and relevering process. We illustrate in Appendix A how these two sources of uncertainty can compound.

3.14. A further complication is that the above calculations ignore the potential for a yet further source of uncertainty, namely the debt beta. An assumed (low) value is almost invariably assumed, but the very fact that it is assumed actually reflects the uncertainty as to the true value. The fact that a lower debt beta leads to a higher value of the re-geared equity beta also has the somewhat bizarre effect of providing companies with an incentive simultaneously to point out that their cost of debt *is* high, and *should* be very low (in a CAPM framework).

Conclusion on measurement challenges

3.15. Our main argument here is that there is considerably more uncertainty when it comes to measuring the observed gearing of the reference companies Severn Trent (SVT) and United Utilities (UU) than is usually recognised. Taking this uncertainty into account leads us to question the relevering exercise. This is especially so for

payments under the index-linked swaps often rank ahead of other senior creditors' claims, the use of swaps may also affect the recovery prospects of other creditors. These liabilities are not accounted for in gearing estimates.

relatively modest differences in actual and notional gearing—differences that are difficult to distinguish empirically, but that have material effects on relevered equity betas and hence on the estimated cost of equity (especially when the equity risk premium is as large as it is currently). An ideal cross-check would be to devise a statistical test of whether notional gearing is statistically significantly different from observed gearing. If the null hypothesis that they are identical cannot be rejected, then the presumption would be that there is no empirical justification for relevering equity betas. (We note that in the NERL case, the CMA chose a notional gearing of 30%, which was deemed "similar to" the gearing levels of comparator companies.) While this is simple to state, we do not underestimate how hard it would be to put into practice. Similarly, when uncertainty about actual gearing is factored in, it is far from clear whether relevered equity betas can be distinguished empirically from the underlying raw equity betas.

4. Market failure and financial resilience

4.1. The potential market failure associated with regulated firms' financial decisions is an externality that creates a gap between the social and private costs of financial distress. Clearly, managers of a regulated firm will consider some costs of failure of the firm they manage; these costs are both financial and reputational, and—more altruistically—may include non-self-interested concerns about customer welfare. In principle, however, they will not consider the full costs of actions—such as increased gearing, covenants, non-debt liabilities etc.—which risk the financial failure of the firm. This is, of course, inherent in limited liability, where investors' losses cannot exceed their investment in the firm. Limited liability makes the firm's payoff function convex, inducing risk-taking behaviour, even by managers who in more normal times would be risk averse (see e.g., Stiglitz and Weiss (1981)[5] and an extensive subsequent economic literature).

4.2. The gap between social and private costs can take a number of forms. Most starkly, it may arise from costs that are incurred when a firm goes bankrupt that are not faced by the management of the firm: for example, customer disruption; the costs of bankruptcy procedure; losses by other stakeholders; etc.. At a lesser extreme, it may be that a firm in financial distress does not go bankrupt, but instead cuts back on areas of expenditure that are important in the longer term, but may be discretionary in the short term. This can occur in a number of ways. Maintenance expenditure may be reduced, which may only be apparent after a number of years. Alternatively, a company might face significant pension fund deficits that could have been addressed with profits that were paid out as dividends. The social-private gap now takes the form of benefits foregone that would have occurred had the firm not been in financial distress.

[5] Stiglitz, Joseph E., and Andrew Weiss (1981). "Credit rationing in markets with imperfect information." *The American Economic Review*, 71(3), 393-410.

4.3. The externality can also take more subtle forms. Spiegel and Spulber (1994 and 1997)[6] argue that if regulators cannot commit to a particular regulatory scheme, then regulated firms will have an incentive to finance their investments with debt—not because debt is cheaper, but because it provides a device to limit regulatory opportunism. In their model, they assume that a regulator cannot commit to long-term regulated prices. In this case, the regulator has an incentive to cut prices once the firm's investments are sunk, benefitting consumers at the expense of the firm's owners. If the firm finances its investments with debt, however, 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. Rao and Moyer (1994)[7] make a similar argument.

4.4. Bortolotti et al. (2011)[8] 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 leads to higher regulated prices (but not vice versa): that is, gearing affects prices, but prices do not affect gearing.[9]

4.5. A final form of externality, related to but distinct from Spiegel and Spulber, comes from any explicit or implicit insurance offered by a regulator when there are concerns about financeability. This leads to lower risk to equity and hence the WACC can be lowered through increased gearing. It is, of course, difficult to collect direct evidence that this behaviour occurs. No regulator would be willing to state the insurance explicitly: the preference will be to assert that no adjustment to the terms of a regulatory contract will occur in the face of impending financial failure, and special administration procedures will sort it out. Investors have little incentive to declare that this is how they perceive the regulator will behave, in case that leads to an adverse regulatory reaction. We argue, however, that the following type of statement, in this instance made by the CMA in its Provisional findings in the PR19 appeal, is strongly suggestive about this form of insurance:

---

[6] 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.

[7] Rao, R. and Moyer, R.C. (1994), "Regulatory Climate and Electric Utility Capital Structure Decisions", *Financial Review*, 29:1, 97–124.

[8] 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.

[9] The notion of causation here is Granger-causality. According to Granger causality, if a variable X_1 "Granger-causes" another variable X_2, then past values of X_1 should contain information that helps predict X_2 above and beyond the information contained in past values of X_2 alone.

"10.95 We consider the choice of WACC, which will be directly related to the level of free cash flow generated by companies that achieve the targets in the determination, is the most important determinant of financeability."

4.6. Figure 1 below illustrates the situation. The Figure looks at one aspect of financial structure: the level of gearing. In this illustration, the social optimal level of gearing, where marginal benefits equal social marginal costs, is given by the intersection of the red and blue lines. The corresponding point with private marginal costs is given by the intersection of the green and blue lines. Since the green line (showing private marginal costs) is everywhere below the red line (social marginal costs), the socially optimal level of gearing is below the level chosen by firms. (For clarity, we consider only the potential gap between social and private marginal costs of gearing. There may also be a gap between social and private marginal benefits.)

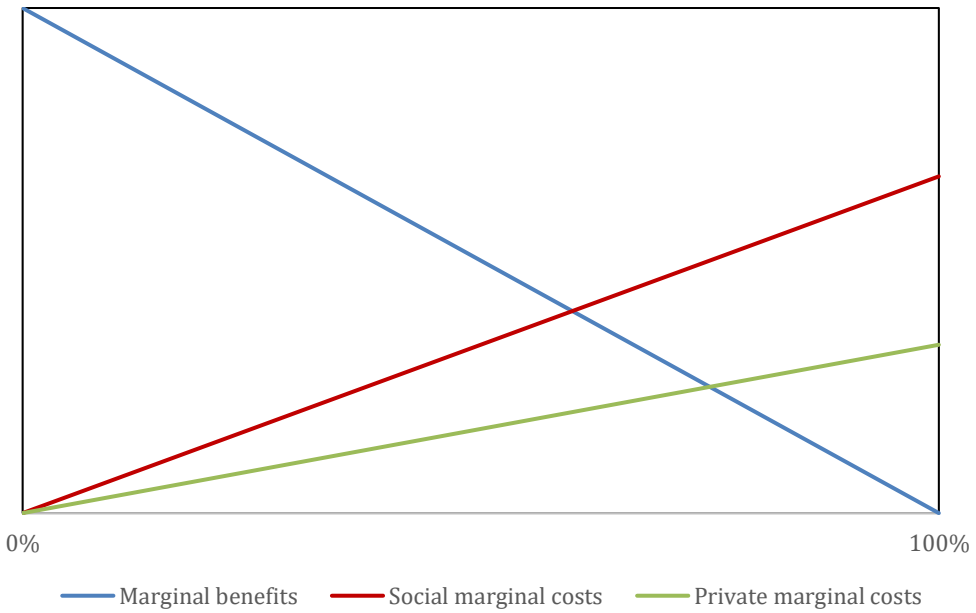

Figure 1: marginal benefits, and social and private marginal costs from gearing (illustrative)

4.7. This Figure is, of course, illustrative, and the question empirically is: how large is the gap between the social and private (marginal) costs of gearing? There are two approaches to answering this question: the empirical evidence about the costs of financial distress; and the particular arrangements that hold for the water sector.

Empirical evidence

4.8. As Branch (2002)[10] notes, the costs of financial distress fall into four categories:

1. Real (i.e., not transfer) costs borne directly by the bankrupt firm;
2. Real costs, borne directly by the claimants (but not by the bankrupt firm itself);
3. Losses to the bankrupt firm that are offset by gains to other entities;
4. Real costs borne by parties other than the bankrupt firm and/or its claimants.

4.9. A negative externality of the type that we refer to arises most clearly in case 4. Unfortunately, this is also the most difficult case to estimate, and the empirical literature concentrates on cases 1-3. Even then, the range of estimates is large, as perhaps might be expected given the differences in bankruptcy regimes.

4.10. To see the difficulties, consider a range of papers (arranged chronologically). Altman (1984)[11] looks at a sample of twelve retailers and seven industrial firms that went bankrupt in the 1970s; he found that total bankruptcy costs amount to, relative to (predistress) firm value, about 15% for industrial firms and around 7% for retailers. Weiss (1990)[12] looked at 37 NYSE and AMEX firms that filed for bankruptcy over the period 1979-86. Average direct costs in this sample were 3.1%. Franks and Torous (1994)[13] found that the average incremental bankruptcy cost is at least 4.5%. Opler and Titman (1994)[14] reported that highly-leveraged firms (in the top leverage decile) in financial distress tend to lose substantial market share: 26% more than those in the lowest leverage decile. In contrast, Kaplan (1994)[15] found in one case that the estimated gains from bankruptcy-induced financial restructuring process exceeded the cost. Andrade and Kaplan (1998)[16], in a sample of 31 large firms that undertook highly-leveraged transactions and subsequently became distressed, estimate costs of financial distress amounting to 10–20% of pre-distress market value. They found that financial distress most frequently led to a reduction of capital expenditure, the sale of assets at depressed prices, and a costly delay before restructuring.

---

[10] Branch, Ben. (2002). "The costs of bankruptcy: A review." *International Review of Financial Analysis*, 11(1), 39-57.

[11] Altman, E. I. (1984). "A Further Investigation of the Bankruptcy Cost Question", *Journal of Finance*, 39, 1067–89.

[12] Weiss, L. A. (1990). "Bankruptcy Resolution: Direct Costs and Violation of Priority of Claims", *Journal of Financial Economics*, 27, 285–314.

[13] Franks, J. and W. Torous, (1994). "A comparison of financial restructuring is distress exchanges and Chapter 11 reorganizations". *Journal of Financial Economics*, 349–370.

[14] Opler, T. and S. Titman (1994). "Financial distress and corporate performance". *Journal of Finance*, 1015–1040.

[15] Kaplan, S. (1994). "Campeau's acquisition of federated post-bankruptcy results". *Journal of Financial Economics*, 35, 123–136.

[16] Andrade, G. and S. N. Kaplan, (1998). "How Costly is Financial (Not Economic) Distress? Evidence from Highly Leveraged Transactions that Became Distressed". *Journal of Finance*, 53:5, 1443–93.

4.11. Summing up, Branch (2002) concludes that, on the average, the par value (i.e., before distress) of the debt of a firm that encounters financial distress is allocated as follows:

1. The loss, which caused the bankruptcy, consumes about 28%;
2. The cost of dealing with distress consumes about 16%;
3. The net value available to distribute to the claimsholder amounts to about 56%.

4.12. To repeat, however: across a number of empirical assessments about the costs of financial distress, two findings emerge. First, there is a wide range of estimates, from around zero up to around 20%. Secondly, there is little (none that we have found) empirical evidence indicating the size of the gap between the social and private costs of financial distress. It is also worth bearing in mind that these studies do not look at regulated sectors or natural monopolies—precisely because financial distress is relatively rare in those sectors, and so there would be limited empirical evidence to assess.

4.13. There is some evidence linking financial structure to operational performance, For example, in a 2018 report, Moody's argues that some key performance measures, including leakage, outcome deliveries and customer service, over the period 2015-18 did not show a negative relationship with financial structures: in their analysis, whether the company was highly covenanted or not. Figure 2 below gives on example from their analysis, looking at leakage rates: firms in the top right are the worst performers, firms in the bottom left are the best performers, against two different leakage measures. For each firm, Moody's notes its financial structure, concluding that there does not appear to be a systematic relationship between highly-covenanted structures and leakage rates.

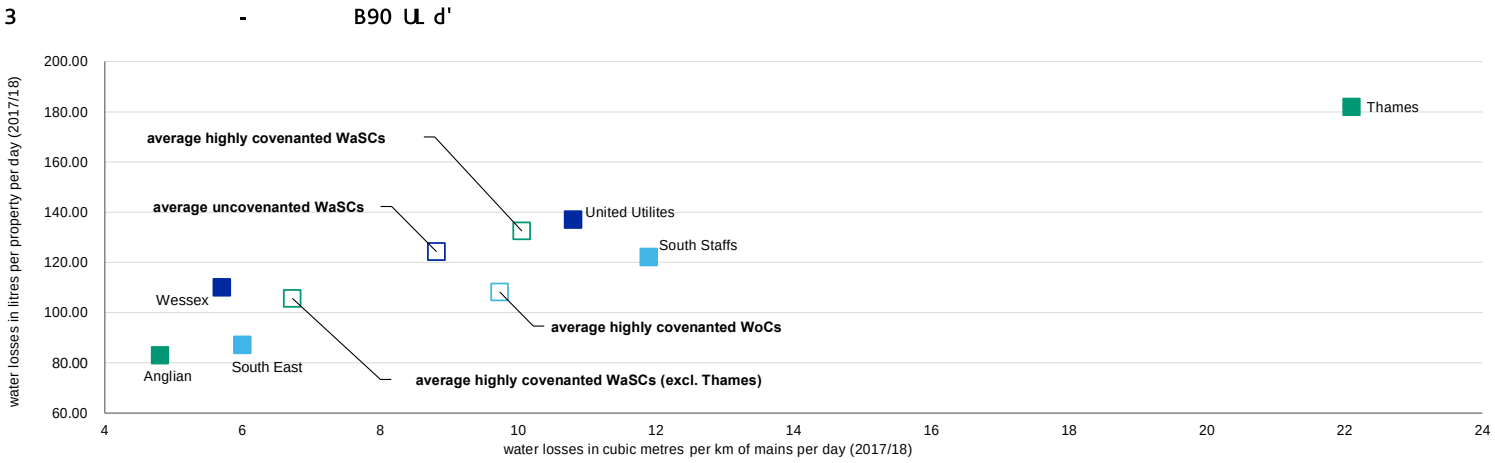

Figure 2: Taken from Moody's (2018).

4.14. More recent evidence gives a counterbalance to this view. The difficulties encountered by Southern Water , the company with the lower credit rating in 2020-21 (see Table 1 below), are in the public domain and well-known. Now-standard reports from Ofwat allow various financial structure metrics (from financial resilience monitoring reports) to be associated with service delivery categories (from service delivery reports). There are, so far, three years of data for categories of service delivery that are determined by Ofwat, based on the information compiled in the service delivery reports. We have conducted a limited analytical exercise to assess the extent of correlation, if any, between credit rating scores and service delivery performance, for the three years of available data. The data we have used is summarized in Table 1 below.

	Service Delivery Report (SDR) categorisation			Lowest credit rating notches above minimum investment grade		
Company	2020-21	2019-20	2018-19	2020-21	2019-20	2018-19
Dŵr Cymru	1	1	1	3	3	4 (Neg)
United Utilities	1	1	1	2	2	2 (Neg)
Severn Trent	2	1	1	2	2	2 (Neg)
Anglian	2	1	2	2	2 (Neg)	2 (Neg)
Portsmouth	2	1	2	2 (Neg)	1 (Neg)	1 (Neg)
Affinity	1	0	1	2 (Neg)	2 (Neg)	2 (Neg)
Northumbrian	1	2	1	2 (RW Neg)	2 (RW Neg)	2 (Neg)
South Staffs	1	2	1	1	1	1
Wessex	1	2	2	1	1	2 (Neg)
Bristol	0	2	1	1	1 (Neg)	2 (Neg)
South East	0	1	1	1	1	1 (Neg)
Thames	0	0	0	1	1	2 (Neg)
SES Water	0	1	1	1 (Neg)	1 (Neg)	2 (Neg)
Yorkshire	1	1	1	1 (Neg)	1 (Neg)	1 (Neg)
Southern	0	0	1	0	0	1 (Neg)
Hafren	1	0	0	-	-	-
South West	1	1	1.5	-	-	-

Table 1: service delivery categories and credit ratings. Source: Ofwat

4.15. To assess the correlation between these (nominal) variables, we have calculated the Cramer's V statistic for each of the three years.[17] Given these data (which have two degrees of freedom), the Cramer's V statistic is said to indicate a small, medium, or large effect size according to whether the statistic is:
4.15.1. Small: around 0.07;
4.15.2. Medium: around 0.21;
4.15.3. Large: around 0.35.

4.16. Table 2 shows the calculated Cramer's V statistics. The results indicate a small correlation between credit ratings and service delivery performance in 2018/19; a large one in the following year, 2019-20; and a medium-tending-to-large effect in 2020-21.

2018-19	2019-20	2020-21
0.1188	0.3591	0.3148

Table 2: Cramer's V statistics

4.17. The fact of the matter, however, is that there is still relatively little data on which to base an assessment of whether there is a robust relationship between measures of financial resilience and operational performance. As Ofwat's database develops, we expect the question to become more amenable to rigorous analysis.

4.18. Finally, the CMA noted in its PR19 Final Report that there have been no observed water company failures over the last 15 years, despite a global financial crisis and a worldwide pandemic. There is, therefore, limited empirical evidence for the UK water sector that can be applied to the issue of the social-private gap in the costs of gearing.

Brief review of Special Administration procedures

4.19. Under the Water Industry Act 1991 (Sections 23–25 and Schedule 3), there are four circumstances under which a Special Administrator can be appointed:

4.19.1. It is deemed inappropriate for the company to continue its Appointment as it has breached or is likely to breach, in a manner sufficiently serious, either any principal duty to maintain a water supply or provide a sewerage system, or a final or confirmed provisional Enforcement Order;

4.19.2. The company is unable, or is unlikely to be able, to pay its debts;

[17] We cannot use e.g., the standard Pearson correlation coefficient, since the data are discrete, corresponding to categories rather than continuous measures (such as e.g., gearing). Cramer's V statistic is designed to deal with this case. As with a more familiar correlation coefficient, it varies from 0 (no association) to 1 (complete association).

4.19.3. If the Secretary of State certifies that it would be appropriate for the company to be wound up, and it would be just and equitable to do so under Section 440 of the Companies Act 1985;

4.19.4. If it is 'certified by the Secretary of State or Ofwat to be necessary in connection with the making of any appointment or variation replacing the company as a statutory undertaker.

4.20. The purpose of the special administration arrangements is to transfer the company's business as a going concern, and to carry out the functions of the company until then. The special arrangements therefore differ from the standard administration regime applying to all companies, where other options are available including the company's assets being sold and the operations of the company terminated. In addition, special administration differs since it can be used where a company has not met its principal duties, as well as for financial reasons.

4.21. Provisions therefore exist for transferring control of a water company in the event of financial failure of the existing operator, and these provisions ought to avoid disruption of services to customers. In the PR19 appeal, the CMA cited the cases of Wessex Water and Dŵr Cymru as examples where transfer of ownership had occurred without customers being adversely affected.

4.22. The more recent example of Southern Water is relevant, however. As detailed in the recent letter from Ofwat's chair, underperformance by Southern Water has continued for a number of years. Southern Water's gearing, including derivative liabilities, has been very high; its credit rating has dropped to the lowest level consistent with (Moody's) investment grade; and it has recently received the largest fine for any water company and is subject to ongoing investigation by the Environment Agency. Nevertheless, the special administration procedures have not been triggered for Southern Water; and indeed, special administration arrangements have yet to be used in the UK water sector. We suspect, therefore, that the conditions under which special administration might be triggered are very demanding, and hence it is so rare for special administration to be invoked. Indeed, one lesson that could be drawn from the recent experience in energy markets, and the case of Bulb, is that it is exceedingly difficult (not to mention costly) for special administration to be triggered.

4.23. If this is correct, then it provides a reason why increased gearing does not have the effect that it does in the Modigliani-Miller 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. This in itself could provide some explanation for

the relatively high gearing levels in the sector, in addition to other, more familiar reasons to favour debt (e.g., tax advantages).[18]

4.24. In short, special administration remains an important mechanism to protect customers from company failure. However, we think it is questionable whether special administration procedures, on their own, or at least in their current form, are sufficient to address the externalities present. It is important for regulators to consider therefore whether other policy responses should be considered to address the externalities.

Options

4.25. There are two general economic policy responses to externalities:

4.25.1. Set a price to close the gap between social and private valuations;

4.25.2. Set quantity controls to ensure optimal levels of activity.

In these terms, Ofwat's GOSM is a price-based policy; NATS gearing cap is a quantity-based policy. On the whole, measures implemented through price controls are (perhaps unsurprisingly) price-based measures; licence conditions are predominantly quantity-based measures, according to this categorisation.

4.26. In a world of complete certainty, these two responses are, in principle, entirely equivalent (although of course in practice, there can be important differences even in the complete certainty case). There is considerable difficulty, however, in measuring the size of externalities associated with gearing choices: indeed, that was one of the CMA's criticisms of the GOSM, and there is a wide range of estimates in the academic empirical literature cited above. Once there is uncertainty, price-based and quantity-based measures potentially yield different outcomes: an observation first made in the seminal paper by Weitzman (1974).[19] Which measure should be used depends on the relative costs of mistakes. Price-based measures make it easy to predict the cost of the policy, but not exactly how much activity will result. Quantity-based measures ensure the amount of activity, but not the exact cost. In terms of gearing: a price-based measure (such as the GOSM) gives greater certainty about revenues raised from the scheme than it does about the levels of gearing that will result. (Again, this formed part of the CMA's criticism of the GOSM: see e.g., §9.1204 of the Final Report.) A quantity-based measure, such as the NATS gearing cap, gives greater certainty about the ensuing level of gearing; but there is uncertainty about how costly it will be if the gearing cap is set sub-optimally.

---

[18] A regulator might argue that there is no such implicit insurance: that a failed firm will be allowed to go bankrupt. The issue, however, is what investors *actually* believe, not just what regulators would like them to believe.

[19] Weitzman, Martin L. (1974). "Prices vs. quantities." *The Review of Economic Studies*, 41(4), 477-491. Although this paper looked at how to regulate pollution, its lessons are of broader applicability.

4.27. We now look at alternative options: the GOSM; a cap on gearing; other licence conditions; and pricing to create a special administration fund.

The GOSM

4.28. We have commented previously on the CMA's rejection of the GOSM. Nevertheless, the GOSM remains in place for the water companies who did not appeal. This might provide, in due course, some information about the effectiveness, or otherwise, of the GOSM.

4.29. There are four key parameters of the GOSM that require empirical justification:

4.29.1. the threshold to trigger the mechanism (starting at 74% in 2020-21 and reducing by 1% each year, ending at 70% in the year 2024-25);

4.29.2. the sharing ratio (set to 50%);

4.29.3. the amount to which this ratio is applied (the difference between notional nominal cost of equity and actual nominal cost of debt);

4.29.4. the reference point for sharing (65%).

Hence the payment required under the GOSM is (Gearing – 65%) x 50% x (Allowed Nominal Cost of Equity – Actual Cost of Debt) x Closing RCV (nominal).

4.30. What should be clear is that this is a relatively complex scheme. In addition, the GOSM is defined in terms of assumed benefits to investors (i.e., the supposed reduction in the WACC that comes from increased gearing) that are to be shared with customers: as Ofwat stated it explicitly in its 2018 Putting the sector in balance position statement:

> "Companies with gearing levels materially above our notional assumption, should propose, in their business plans, outperformance sharing mechanisms that allow customers to share in the returns equity investors achieve from high gearing."

4.31. Both aspects are problematic. The complexity of the scheme presents a challenge to match its components to empirical measures that are difficult to observe and estimate. The definition of benefits immediately raises the question of whether investors do receive higher returns from gearing (inviting recitals of Modigliani-Miller); if so, how much; and whether it is customers who face the cost of higher gearing. Our view is that this framing of the GOSM was unfortunate, even if the underlying motivation may have been sound.

A licence-based gearing cap

4.32. There is at least one example in the UK of a regulator intervening directly to set a cap on gearing. NERL's licence requires it to explain any difference between expected average gearing at 31 March each year and its 'target' gearing level of 60%. In addition, there is a gearing cap of 65%, introduced to NERL's licence in December

2010, which if breached requires plans to be put in place to lower gearing, along with a lock up of dividends and other restrictions. (See e.g., CAA (2020).)

4.33. In the CMA's PR19 final report, it signalled in a number of places that licence conditions offer an alternative option for Ofwat, if high gearing is a concern. For example, in its summing up of its views on the GOSM, the CMA said in §9.1204

> "We do not consider that the GOSM is in any event well-designed to address any risks of financial failure that result from high gearing. The design of the GOSM is based on the sharing of 'excess returns' from high gearing with customers. It is therefore also implicitly sharing the 'risk' of high gearing. While this will create incentives on companies, it does not necessarily prevent companies choosing to run relatively high gearing, as would occur with the imposition of licence conditions limiting gearing. If a company chooses to retain relatively high gearing, the GOSM does nothing to reduce the risks associated with the level of leverage."

And again in §9.1224:

> "There are different options open to Ofwat to address any concerns it may have about the consequences of high gearing and other factors affecting financial resilience. These include licence modifications which could be defined to directly limit gearing."

4.34. A licence-based gearing cap has the advantages, therefore, of a regulatory precedent in the UK; and of ensuring certainty about the level of gearing that will be achieved. Its disadvantage is the flipside of the latter: it does not afford firms the opportunity of choosing a higher gearing level if the gearing cap is set in error, or if there are problems (as we have noted) in accurately measuring gearing.

4.35. If this approach is adopted, it would be advisable to consider whether the licence condition is framed just in terms of gearing, or a broader range of measures related to financial resilience. This leads us to the next option for consideration.

Other licence conditions: cash lock-up

4.36. A gearing cap is not the only (quantity-based) measure that could be taken. In addition to a requirement to maintain an investment grade credit rating, existing licences have other conditions concerning financial resilience. For example, most licences now have cash lock-up provisions. These provisions restrict certain payments, such as dividends, in cash or any other form, when a condition is not met. Currently, cash lock-up is triggered when either

4.36.1. A regulated water company does not hold a credit rating which is investment grade;

4.36.2. One or more issuer credit ratings held by an appointee is not investment grade; or

4.36.3. The rating outlook specified by the credit rating agency which has assigned the lowest investment grade rating is on review for possible downgrade or is on "Credit Watch" or "Rating Watch" with a negative designation, or has been changed from stable or positive to negative.

See Ofwat's *Conclusions on strengthening the regulatory ring-fencing framework*.

4.37. A further option, then, is for cash lock-up requirements to be tightened: for example, by triggering cash lock-up when a company's lowest-rated debt is only one notch above investment grade. (Clearly, there are variations on this theme.)

4.38. This option has the advantage of being a modification of the existing approach. While licence changes require the agreement of both the regulator and the regulated companies, and so are far from automatic, it should be a more straightforward exercise than the creation of new regulatory schemes (such as the next option that we discuss below). The main disadvantage of this approach, as we see it, is the possibility that the agencies apply (even if only implicitly) an uplift to regulated companies' ratings on the basis of a belief that companies will not be allowed to fail. This potential disadvantage can be mitigated by basing cash lock-up on other financial resilience measures, such as gearing levels, adjusted interest cover, etc.. It may also be possible to base cash lock-up on service levels and failures directly—the things that led to the externalities that call for regulation.

Special administration fund

4.39. The final option we review is another price-based mechanism, like the GOSM, but with an important difference: its conceptual basis. In this option, a charge is set that firms with risky financial structures will be required to pay in order to contribute to a fund that is used to pay the costs of the special administration regime. These are the costs that are not considered by the firms' current owners, since the costs are incurred after they have been removed from ownership. Hence, this charge is intended to close the gap between social and private costs, rather than distributing benefits. In this regard, it is sharply different from the GOSM, which was cast as redistributing benefits between equity-holders and customers. It is directly analogous to compensation funds in financial services; and experience from these funds (especially in the US e.g., the Federal Deposit Insurance Corporation, FDIC; but also the UK's Financial Services Compensation Scheme) can inform how this might work, including which organisation holds the fund; what is done with the fund in the absence of any claims against it; etc.. The charge could be levied on gearing levels (as in the GOSM), or be based on other (quantitative) measures of financial resilience e.g., credit ratings.

4.40. The size of this charge would need to be determined: not an entirely straightforward task, given the uncertainty about the size of the externality. Since special administration arrangements have yet to be used in the UK water sector, there is little experience of how costly these arrangements might be, although experience from other sectors can act as a guide. (We note that the CMA did not accept comparisons with Railtrack and Metronet: see §9.1209 of the Final Report.)

On the basis that restructuring of larger firms is likely to be costlier, the charge could be set as a percentage of the RCV. It is likely that in practice, the charge will need to be dynamic, according to the size of the accumulated fund.

4.41. The advantages of this scheme are its relatively simplicity (compared to the GOSM); its conceptual basis of closing the externality rather than redistributing putative benefits; and the flexibility it gives firms to choose financial structures that are different from the notional level. Its disadvantages are the lack of certainty about resulting choices of financial structures; and the lack of precedent within the UK utilities sectors (although there are close analogies with the financial sector). Lessons from the financial sector also indicate the need to pay attention to the potential for moral hazard:[20] for example, customers may be less discriminating about the financial resilience of the water company they use, if they think they are fully insured. The fact that water companies are local monopolies, however, lessens the force of this argument.

4.42. If it proves too difficult to determine the appropriate charge to establish the special administration fund, a potential alternative would be a redesign of the special administration procedure itself, along the following lines:

4.42.1. Tighten further the conditions for triggering special administration such that failure to meet service conditions by some margin, over some time period, is a *sufficient* condition to trigger special administration (analogous to the way that similar failures in rail companies have been used to take away franchises), *irrespective* of the degree of financial distress of the company concerned. (In the case of Southern Water, for example, service failures were extensive and long-standing but did not trigger special administration.)

4.42.2. If special administration is triggered (for any reason), then the cost of maintaining and restoring service levels is a prior claim on the assets, before both equity holders and all creditors. This is the case currently within a price determination period—a firm under special administration is obliged to meet its targets—and it is important that this remains the case. It is untested, however, as to what would apply if special administration were to span two price determination periods. In this case, it will be important for the costs of restoring service levels to be assessed in advance, either by Ofwat or by an independent assessor (in which case the costs of the assessor would also be a prior claim), though it should be expected that cost allowances allowed reflect the efficient provision of services. Under either scenario, as the cost of maintaining and restoring service levels is a prior claim to that of investors, incumbent firms retain a strong incentive to find the more efficient solution themselves, and thereby avoid special administration in the first place.

---

[20] See e.g., Admati, Anat, and Martin Hellwig. 2013. The Bankers' New Clothes: What's Wrong with Banking and What to Do about It. Princeton: Princeton University Press.

4.43. A revised procedure of this form would in principle internalise the externality. There is indeed some risk that it might actually *over*-internalise it, resulting in gearing levels that were below socially optimal levels. But since the regulator's loss function is likely to be asymmetric with respect to service provision, this in itself might not be regarded as too much of a problem.

Conclusion on market failure in the choice of debt

4.44. The socially optimal level of gearing equates the social marginal benefits and costs from debt. While simple to state conceptually, there are formidable empirical challenges to determining this equality. Indeed, there is no consensus about *privately* optimal levels of gearing, never mind the socially optimal ones. This concern formed part, but not the whole, of the reservations about the GOSM expressed by the CMA, and seems unlikely to us to be resolvable to the standards that the CMA seems to have set in the PR19 appeal. This is a particular weakness of price-based schemes, such as the GOSM and the special administration fund.

4.45. While we are attracted to price-based schemes—they allow firms the flexibility to choose gearing when there is considerable uncertainty about optimal gearing levels—pragmatically we recommend that quantity-based measures, such as a NERL-type gearing cap, or tightening of cash lock-up requirements, are considered by Ofwat. To repeat the advantages identified above: for both, there are UK regulatory precedents to which Ofwat can appeal; and they offer certainty about gearing levels. Importantly, both would need to be addressed outside of the price control process, through licence modifications. This is not, of course, without complication.

4.46. In our discussions with ratings agencies, we did not detect concern from them that addressing financial structures would necessarily have an adverse effect on credit ratings for the sector. There was a reaction from the agencies to the GOSM; for example, in Moody's review of the sector in May 2018, which concluded that the GOSM led to a modest deterioration in the stability and predictability of the regulatory regime. In discussions, our understanding is that it was the way in which the GOSM was introduced, rather than the mechanism itself, that caused concern. With appropriate consultation and lead-times, we expect that modifying licences to address financial resilience concerns would not perturb the regulatory regime.

4.47. Finally, we note (returning to a previous theme) the importance of defining clearly and consistently what is meant by gearing, if this is the basis of licence conditions. Our strong preference is for market valuations; and for a comprehensive definition of debt, to include e.g., mark-to-market valuations of derivatives.

5. Relevering beta

5.1. A standard procedure has developed over the years to deal with notional gearing differing from actual gearing. The approach, briefly, is as follows:

5.1.1. Estimate econometrically the so-called "raw" equity beta, denoted $\beta_E(g)$, from data. The notation here makes explicit the fact that the equity beta will depend on the (observed) level of gearing g.

5.1.2. Delever this beta to calculate an asset beta, using the Harris-Pringle formula as noted above and an estimate of the debt beta: $\beta_A = g\beta_D + (1-g)\beta_E(g)$. Both this asset beta and the debt beta are taken to be invariant to the level of gearing.

5.1.3. Relever the asset beta using the notional gearing level g^*. *If* β_D is also invariant to gearing and known, then the relevered equity beta is $\beta_E(g^*) = \frac{\beta_A - g^*\beta_D}{1-g^*}$.

5.2. There are at least a couple of problems with this procedure. First, if β_D is itself a function of gearing i.e., $\beta_D(g)$, then the asset beta formula cannot be used to determine both relevered equity and debt betas: there is one equation but two unknowns. But not very much is known about β_D (despite a growing amount of work on the issue: see for example CEPA (2019); it is typically assumed to be a constant and small (an assumption that the regulated companies have an incentive to support).

5.3. The second problem arises from this latter assumption of a small and constant debt beta: it gives rise to an estimate of the WACC that increases with gearing. This problem has been noted e.g., in the UKRN 2018 report and by the CMA in Appendix D of its final report on the NERL price redetermination. It arises through partial application of the CAPM: typically, the CAPM is used to estimate the cost of equity; but it is not used for the cost of debt, which instead is taken from market data. The formula for the (estimated) WACC[21] is therefore $WACC(g) = gR_D + (1-g)R_E(g)$. Here, as noted above, the cost of debt R_D (presumed to be observed from market data) is assumed not to vary with the gearing level g, while the (estimate of the) cost of equity $R_E(g)$ does. The notation makes explicit the potential for the WACC to depend on the level of gearing. Some straightforward algebra shows that the derivative of the WACC with respect to gearing, $WACC'(g) = R_D - (R_f + \beta_D ERP)$, where R_f is the risk-free rate and ERP is the equity risk premium. Note that the expression in brackets on the right-hand side of this expression is the CAPM-derived cost of debt. Typically, the cost of debt R_D is above this CAPM level; or put differently, the required debt beta to ensure that the cost of debt is equal to the CAPM level is typically assumed to be implausibly high. The immediate implication of this is that the WACC is increasing with gearing. (In terms of the algebra, if $R_D > R_f + \beta_D ERP$, then $WACC'(g) > 0$.)

5.4. This feature is currently compounded by the standard regulatory practice of allowing for embedded debt. We have commented extensively elsewhere about this practice; see e.g., the UKRN 2018 report. For now, we simply observe that with

[21] We shall now just refer to the WACC, rather than repeating "estimated WACC".

embedded debt allowances being well above the allowed cost of new debt,[22] the final WACC is strongly increasing with gearing levels.

5.5. A WACC that increases with gearing is counter to the Modigliani-Miller (MM) theorem, which establishes that as long as certain conditions hold, the capital structure of a firm should be irrelevant for its value. Of course, it may be the case that the MM theorem does not in fact hold: a possibility we discuss further below. The problem with that position, however, is that the relevering (Harris-Pringle) formula itself is based on the MM theorem. Hence there is a profound inconsistency in the methodology: the MM theorem is implicitly used to delever betas; but the subsequent regeared beta leads to a WACC that violates the MM theorem. Perhaps more importantly, a WACC that increases with gearing seems at odds with the observed behaviour of regulated firms, particularly in the water sector, that have chosen gearing levels higher than the notional level used by the regulator.

5.6. There have been, so far, three regulatory responses to this situation. In its 2007 review of Heathrow Airport, the Competition Commission introduced a debt beta (of 0.1) into the WACC calculation. This had the effect of flattening the dependence of the WACC on gearing, although it did not remove it altogether. In the CMA's redetermination for NERL, discussed above, the CMA found that it could not use this approach to flatten the WACC curve without, in their phrase, "using a debt beta which is implausibly high". Instead, the CMA used a lower notional gearing than the regulator—30% rather than 60%—largely based on the fact that the three airports in their data sample all had gearing averaging 30% (relative to enterprise value), and hence the estimated beta did not need to be un- and re-levered. The third regulatory response has been to note the issue, but accept it. This response may be acceptable if the effect on the WACC of the regearing procedure is relatively small. In PR19, observed gearing was 54% and notional gearing 60%. The point estimate of the raw equity beta, at 54% gearing, was 0.63; the relevered beta, at 60% gearing, was then 0.71. With an equity risk premium of 7.89% (i.e., a total market return of 6.50% and a risk-free rate of -1.39%), regearing beta adds 58 basis points to the cost of equity.

What are the alternatives?

5.7. We now detail five alternatives to the standard procedure. Two of the approaches avoid relevering, but in different ways; two of the approaches ensure that the WACC is invariant with respect to gearing; and the final approach moves away entirely from MM assumptions. In all the alternatives explored below, we shall assume that the cost of (new) debt is invariant to gearing. This is probably acceptable for small differences between observed and notional gearing, but may be more problematic for larger differences. In the absence of an explicit debt pricing model, however, we

[22] In PR19, Ofwat set the cost of embedded debt at 2.42% and the cost of new debt at 0.53%, both CPIH-real; the weight on new debt was set at 20%, giving a blended cost of debt of 2.14%, including 0.1% issuance costs.

are unable to determine exactly how gearing affects the cost of debt. In all of the WACC figures given below, we calculate the WACC without and with embedded debt, using the PR19 parameters outlined in footnote 22

I. Set notional gearing equal to observed gearing

5.8. In this first option, used by the CMA in the NERL redetermination, the need to un- and re-lever is removed (reduced) by setting the notional gearing (close) to the observed gearing levels of the reference companies. In PR19, a (appointee) WACC estimated on this basis would have been 2.812%: 15bp lower than Ofwat's estimate of 2.959% (both figures CPI-H real and including embedded debt). Obviously, with a WACC that would otherwise increase with gearing, the effect of setting the notional gearing equal to a lower observed level is to decrease the WACC and hence, all else equal, the regulatory allowed return.

5.9. The move to set notional gearing equal to actual gearing can, in principle, be defended on the grounds given by the CMA in the NERL redetermination. But there is an equivalent way of viewing this option, and that is to assume the WACC to be invariant to the level of gearing: an option that we describe further below.[23] Either approach—lowering the notional gearing, or making the WACC invariant to gearing—leads to the same WACC. But in the former, an explicit argument has to be made about notional gearing, and hence the level of gearing the regulator is assuming or signalling to be appropriate for the sector; and whether the lower gearing of listed companies is due to them having higher enterprise values, and hence should not be applied to the rest of the sector. In the latter approach, the argument relies on consistency with the MM theorem (and hence the whole approach of unlevering and relevering).

II. Use the CAPM for both debt and equity

5.10. In this second option, discussed at some length in the UKRN 2018 report, the CAPM is used to price both equity and debt. The result is what we termed a "pure" CAPM-WACC given by $CAPM - WACC = R_f + \beta_A ERP$. This WACC has the two virtues of: consistency with the general principle accepted by regulators (and us), that the CAPM is the best available asset pricing model; and invariance with respect to gearing levels—gearing does not appear in the formula for the CAPM-WACC. (Of course, it runs counter to the usual approach, which states that the cost i.e., expected return of equity must be estimated while the cost of debt can be observed.) Using Ofwat's PR19 figures, the CAPM-WACC would have been 1.432%. This figure should

[23] Note that this is not the same as the approach that Ofgem took in RIIO-2. There, as we noted earlier, Ofgem assumed that the same WACC applies at a notional gearing of 55% as at 60%. Since the calculated WACC actually increases in gearing, this approach is somewhat generous to the firms to which the 55% notional gearing is applied. The approach we describe later—see option D—works in the opposite direction: the WACC estimated at a lower (actual) gearing level is held constant when moving to a higher (notional) gearing level.

be compared with the PR19 figure excluding embedded debt, which is 2.051%.[24] The CAPM-WACC is therefore some 62bp lower than the WACC for PR19 with embedded debt excluded.

5.11. Note that in this approach, the asset beta β_A must still be estimated. Interestingly, since there is no relevering, we would expect the regulated firms to argue that the debt beta is higher than set by the regulator: unlike now, where their incentives are to argue for a lower debt beta.

5.12. We acknowledge that the CAPM is a poor model of debt returns (while noting it is not the best empirical model of equity returns, but nevertheless is used to estimate them). We continue to advise, however, that the pure CAPM-WACC is something that should be calculated for reference purposes, not least because of its simplicity, transparency and consistency with the CAPM.

III. Use the raw equity beta

5.13. This option combines the estimated cost of equity, computed using the observed "raw" equity beta, with the cost of debt, at whatever notional gearing level the regulator chooses. One interpretation of this option is that the cost of equity is taken to be the same at both the observed and notional gearing levels. We view this option as being anchored in observed data and avoiding a reliance on a relevering formula, with the various problems with this pointed out in this report. In particular, as we have argued, we are not confident that actual gearing is measured well or used correctly. We also view it as problematic to use a relevering formula that relies on MM to be used, but that gives a resulting cost of capital that is inconsistent with MM (i.e., that increases with gearing). In PR19, this approach would have led to a WACC of 2.727%: 23bp lower than Ofwat's figure.

5.14. It is worth noting that this approach leads to an underlying WACC that *decreases* with notional gearing. (In the following explanation, we exclude the cost of embedded debt to make the point clearly.). To see why, note that the WACC can be written as $WACC(g) = gR_D + (1-g)\hat{R}_E$. In this expression, g is the notional gearing level that is chosen; R_D is the cost of new debt: as noted above, assumed constant with respect to gearing; and $\hat{R}_E$ is the estimated cost of equity at the actual (not notional) gearing level. Since $R_D < \hat{R}_E$, this leads to a WACC that declines with notional gearing. This effect is offset, of course, by including embedded debt in the final WACC estimation. This inconsistency with MM matters less in this option, however, since the option does not involve un- and re-levering.

[24] Note that 0.1% issuance cost has been included.

IV. Assume that the WACC is constant

5.15. In this approach, the WACC is forced to be invariant to the level of gearing: that is, for any (notional) level of gearing g, and given an actual/observed level of gearing of $\hat{g}$, we require

$$WACC(g) \equiv gR_D + (1-g)R_E(g) = WACC(\hat{g}) \equiv \widehat{WACC}.$$

In this equation, R_D denoting the gearing-invariant cost of (new) debt, and $R_E(g)$ the (gearing-dependent) cost of equity. The WACC estimation in this case is particularly simple: whatever the notional gearing level used, the WACC is that estimated at the observed gearing—in PR19, 2.812% (including embedded debt).

5.16. It is worth noting the implications of this approach for the cost of equity and the asset beta. Simply rewriting this equation then gives

$$R_E(g) = \frac{\widehat{WACC}}{1-g} - \frac{g}{1-g}R_D.$$

Some further straightforward algebra shows that this implies an asset beta of

$$\beta_A(g) = g\beta_D + \frac{1}{ERP}[\widehat{WACC} - R_f - g(R_D - R_f)].$$

The variation of this asset beta with respect to gearing (i.e., the derivative) is given by

$$\beta_A'(g) = \frac{R_f + \beta_D ERP - R_D}{ERP}.$$

Since $R_D > R_f + \beta_D ERP$ (the cost of debt is greater than implied by the CAPM), the asset beta is decreasing in gearing: in contrast to the standard (MM-based) assumption of a constant asset beta.

5.17. So in this option, the cost of equity at the actual gearing level is 3.60%, while at the notional gearing level, it is 4.18% (using PR19 parameter values).

5.18. Of course, this approach gives the same WACC (2.812%, including embedded debt) as option I in which the notional gearing is set at the observed level. The justifications of the two approaches are, however, quite different, as we have noted.

5.19. It is worth noting that Ofgem used a variant of this approach in RIIO-2: the cost of capital was assumed to be constant *at the level calculated for a notional firm with gearing of 60%*. This value of the cost of capital was then applied to a different notional firm, with gearing of 55%: Table 3 below summarises, using figures taken from Table 13 of Ofgem's RIIO-2 Final Determinations – Finance Annex.

Component	NGET & SPT	GT, SGN south & Cadent
Notional gearing	55.00%	60.00%
Cost of equity	4.25%	4.55%
Expected outperformance	0.22%	0.25%
Allowed return on equity	4.02%	4.30%
Allowed return on debt	1.82%	1.82%
Allowed return on capital	2.81%	2.81%

Table 3: Ofgem cost of capital in RIIO-2

5.20. Because the WACC, as we have noted, increases in gearing, this Ofgem approach yields a higher cost of capital. To see the difference, consider the raw equity beta estimated by Ofgem. This can be inferred from the figures given in Table 9 of the Finance Annex: it is 0.622, at the observed gearing level of 50% (compared to the relevered equity beta of 0.759 at a notional gearing of 60%). Hence Ofgem's cost of equity using this raw equity beta (and other CAPM parameters from Table 11 of the Annex) would have been 3,45%, rather than 4.55%. The resulting allowed return on capital would have been 2.51%, rather than 2.81%.

5.21. We do not recommend Ofgem's approach. It assumes the MM theorem holds twice: when relevering the equity beta; and then when assuming that the WACC is constant at the level calculated for the notional firm with 60% gearing. MM conditions do not hold, however: the WACC is not invariant to gearing, and the process of relevering the equity beta leads to a higher WACC. In our option IV, this contradiction is mitigated by estimating the WACC *at the observed gearing*, and so not carrying out the equity beta relevering process.

V. Departing from MM: explicit recognition of the U-shape in the WACC

5.22. In this final approach, we recognise explicitly the departure from MM assumptions: that gearing increases the value of the firm, and that there is an optimal gearing level. The cost of equity for the unlevered firm must therefore be an upper bound for the levered firm's cost of capital, since the WACC at the optimal gearing level must be (weakly) lower than the WACC of the unlevered firm—otherwise, why would the firm choose non-zero gearing? Hence the aim is to find the cost of equity for the unlevered firm, which gives an upper bound for the true WACC of a firm with (privately) optimal gearing. This is illustrated in Figure 3.

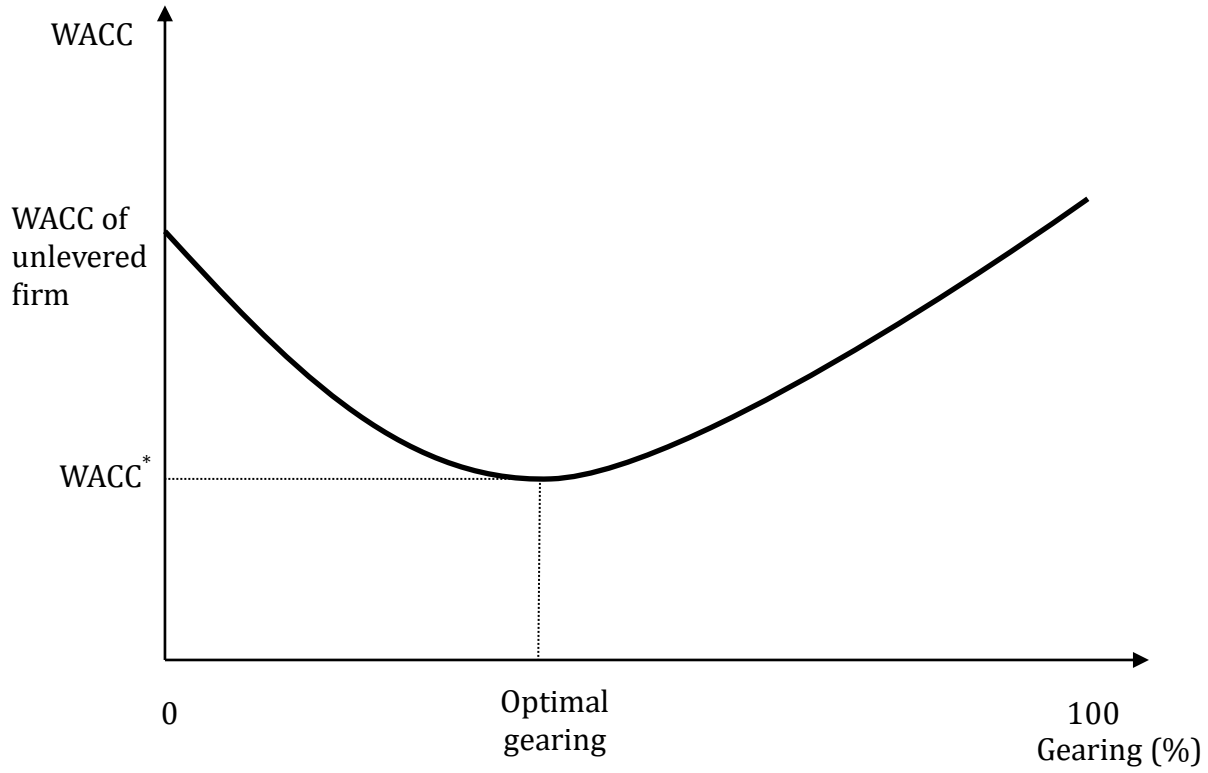

Figure 3: the WACC v. gearing when Modigliani-Miller does not hold (illustrative)

5.23. Since we are no longer in an MM world, however, we cannot use the Harris-Pringle formula to find the beta of the unlevered firm. In Appendix B, we show how to derive the non-MM beta for the unlevered firm. We show there that this beta, denoted β_U, is (approximately) equal to a weighted average of the equity beta and the debt beta—reminiscent of the familiar Harris-Pringle formula for the asset beta. Unlike the Harris-Pringle expression, in which the weights are given by the gearing level (i.e., are 1 – g and g), the weights for the beta β_U are given by $1 - \lambda$ and λ, where λ is a parameter that depends on the gearing level and the ratio of the net benefits of debt to the unlevered firm's value. That is

$$\beta_A = (1 - g)\beta_E + g\,\beta_D;$$
$$\beta_U = (1 - \lambda)\beta_E + \lambda\beta_D,$$

where we can show that $\lambda < g$. The non-MM beta of the unlevered firm will therefore be greater than the Harris-Pringle asset beta.

5.24. In practice, however, the beta of the unlevered firm is close to the Harris-Pringle asset beta. Academic studies indicate the ratio of the net benefits of debt to unlevered value to be no larger than around 11%. As we show in Appendix B, with observed gearing of 54.2%, this implies that $\lambda = 0.485$. Hence β_U has a weight of 48.5% on the debt beta and 51.5% on the equity beta; the standard asset beta has weights of 45.8% and 54.2%. Using Ofwat's PR19 parameters, the asset beta is 0.458

* 0.633 + 0.542 * 0.125= 0.3576; the unlevered beta is 0.485 * 0.633 + 0.515 * 0.125= 0.3867.

5.25. In PR19, this approach would have resulted in a WACC of 1.661%: 39bp less than Ofwat's PR19 figure that excludes embedded debt (this being the appropriate comparison). It should be borne in mind that this value is an upper bound for the true WACC of a firm with optimal gearing: see Figure 3. An alternative way to view this would be that in order for this non-MM WACC to equal Ofwat's PR19 WACC (of 2.051%, excluding embedded debt), the use of debt would need to constitute 25% of the value of the firm: well above empirical estimates (and recalling that in an MM world, that number would be 0).

Conclusion on relevering beta

5.26. Table 4 below summarises the WACCs that would have obtained using PR19 final determination inputs and these different approaches:

Observed gearing A	54.2%	
Notional gearing B	60%	
Risk-free rate C	-1.39%	
Equity risk premium D	7.89%	
Cost of debt (with embedded) E	2.14%	
Cost of debt (new only) F[25]	0.63%	
Raw equity beta G	0.633	
Relevered equity beta H	0.7065	
Asset beta I	0.3576	
Non-MM beta J	0.3867	
	With embedded debt	**New debt only**
PR19	2.959% B*E + (1-B)*(C + H*D)	2.051% B*F + (1-B)*(C + H*D)
I. Notional = actual gearing	2.812% A*E + (1-A)*(C + G*D)	1.992% A*F + (1-A)*(C + G*D)
II. CAPM-WACC	1.432%	C + I*D
III. Raw equity beta	2.727% B*E + (1-B)*(C + G*D)	1.819% B*F + (1-B)*(C + G*D)
IV. Constant WACC	2.812% A*E + (1-A)*(C + G*D)	1.992% A*F + (1-A)*(C + G*D)
V. Non-MM	1.661%	C + J*D

Table 4: WACCs using different approaches

[25] Both debt costs include 0.1% issuance cost.

5.27. All alternative approaches lead to a WACC that is lower than the one used in PR19: sometimes markedly so. The CAPM-WACC and the non-MM figures are the lowest. We continue to argue that the CAPM-WACC provides a useful benchmark.

5.28. On grounds of internal consistency, we would argue for the non-MM WACC, since it is the only measure that explicitly allows for the possibility of the MM theorem not holding. While it introduces one new parameter that needs to be estimated (the net value of debt in firm value) there are strong grounds for assuming that this parameter must live within a relatively narrow range, and thus that the resulting WACC estimate is relatively robust to the value chosen. Additionally, by construction it dispenses with the need to make an assumption about notional gearing. The primary arguments against this estimate are, clearly, its unfamiliarity, and lack of regulatory precedent; and the need to estimate the ratio of net debt to value.

5.29. Of the three remaining options, option III—using the raw equity beta in conjunction with a notional gearing level—has two benefits. First, as we have noted, it is anchored in observed data and avoids a reliance on a relevering formula that has various problems. Secondly, it retains the use of notional gearing which is well-established in UK regulation. Of course, companies may argue that the raw equity beta is not appropriate for them—and that, presumably, it should be set at a higher level. We think the merit of using observed data outweighs this possible counter-argument.

5.30. If this option is not favoured, then we recommend using option IV: mandating a constant WACC. This results in the same WACC as option I, in which notional gearing is reduced to the observed gearing of the reference companies. It avoids, however, the need to argue about the choice of notional gearing, and is readily justifiable in terms of ensuring consistency with MM.

5.31. Whichever option is chosen, we recommend that all these options are looked at, since they all provide valuable cross-checks. As we have made clear throughout the discussion, no one approach is without its limitations. Even our preferred method—option V, recognising the departure from the MM theorem—faces a challenge in estimating key variables. Any one approach to estimating the WACC benefits from comparison with other approaches. We think it instructive that all options that we consider imply a WACC below the level estimated by Ofwat in PR19.

5.32. We also think it important to bear in mind any consequences for the two reference companies used to estimate equity betas, Severn Trent and United Utilities. In both option I and IV, SVT and UU would have an incentive to raise their gearing (or at least appear to raise their gearing), in contrast to the current approach, which gives an incentive for lower observed gearing. Option III, which avoids relevering altogether, is relatively neutral in this regard. In both options II and V, again the incentive is to lower observed gearing, to raise the estimate of the asset or unlevered beta used in the calculation. We note also that all options are exposed to the risk that SVT and UU might delist.

5.33. Finally, to repeat a theme that we have developed during this report: we are concerned that the standard regulatory approach to equity betas that has developed over a number of years is now in danger of becoming too much of a notional exercise. The raw data on which the exercise is based come from a limited number of companies. Unlevering is based on imperfect and uncertain estimates of gearing; relevering (as currently done) is based on a further problematic definition of gearing, and the MM theorem which the resulting WACC does not satisfy. All of this makes a significant difference to the regulatory allowed return, particularly because the equity risk premium is currently so large. In the words of John Kay, "If the answer to a question is very sensitive to a number that is almost impossible to predict, it is time to ask a different question".

6. Answers to Ofwat questions

6.1. Ofwat posed to us the following questions, to which we now give direct responses based on the analysis and recommendations earlier in this report.

1. How far is there a case for setting the notional gearing not only below the actual gearing (relative to RCV), but reducing it below the current 60%?

6.2. Our recommendation, above, is to use observed, "raw" equity betas rather than un- and re-lever based on actual and notional gearing levels. Adoption of this recommendation would mean that the notional gearing level affects the WACC in one way only: the weight on the cost of debt in the WACC. In this case, therefore, lowering the notional gearing has the effect of increasing the WACC. We do not see a case for doing this if our preferred approach is adopted. If instead our alternative approach (which maintains a constant WACC) is adopted, the notional gearing—by construction—does not directly affect the WACC, so again the choice of notional gearing is moot. In summary, therefore, our preferred approaches do not involve choosing notional gearing to affect that regulatory allowed return.

2. Should we expect the regulatory allowed return to be invariant to capital structure—as implied by the weighted average cost of capital in the Modigliani Miller Theorem?

6.3. Consistency with the MM theorem in estimating the WACC is most important when equity betas are un- and re-levered, since this procedure itself relies on the MM theorem (this is implicit in the use of the Harris-Pringle formula). If, as we recommend, relevering is not carried out, then invariance of the WACC with respect to capital structure is less critical. Indeed, our recommendation implies that the WACC decreases with (notional) gearing, as noted above. If our recommended approach is not used, then we suggest that Ofwat adopts the approach that ensures the WACC is invariant to capital structure, for the reasons of consistency just given.

3. If regulatory models of the cost of capital are not consistent with Modigliani Miller Theorem, then how should regulators respond?

6.4. Consistency with the MMT is most important when employing the (Harris-Pringle) procedure to un- and re-lever betas. In making our recommendation about how to use equity betas, we have placed most importance on being anchored in observed data, and as a result, avoiding reliance on a relevering formula that has various data and theoretical problems. This approach is, in one sense, not consistent with the MMT, in that it leads to a WACC that decreases with (notional) gearing. But the approach does not invoke MMT in order to the relever equity beta, and therefore does not suffer from a theoretical inconsistency.

4. What is your view of the standard regulatory approach of de-levering the observed equity beta (from observed gearing – based on enterprise value) and re-levering at the notional gearing level (defined as net debt to RCV)?

6.5. We think there are a number of problems with this standard approach:

6.5.1. We would prefer—if relevering must be done—to relever with a measure of gearing that reflects the fact that MARs are and should be greater than 1, even for the notional firm.

6.5.2. The definitions of gearing should, to be consistent with use of the CAPM, be based on market not book values (for debt).

6.5.3. The approach does not take into account the uncertainty that is inherent in the measurement of gearing. We suspect that for relatively modest differences in actual and notional gearing, it is difficult to distinguish empirically between them, given this uncertainty. Similarly, when uncertainty about actual gearing is factored in, it is far from clear whether relevered equity betas can be distinguished empirically from the underlying raw equity betas.

6.5.4. The standard approach suffers from the inconsistency that it requires MM to un- and re-lever equity betas (via the Harris-Pringle formula), but results in a WACC that is increasing in gearing (and hence counter to the MM prediction).

6.6. These reasons lead us to recommend a different approach.

5. What are the implications of an alternative approach of taking observed equity betas estimated at the observed gearing and setting the regulatory allowed return from this?

6.7. As we have noted above, the effect on the WACC of using this approach is to lower it, from the PR19 implied level of 2.959%, by around 15bp, to 2.812%.

6. How could the GOSM be adapted to address the CMA's concerns? Are there alternative approaches (within a price control framework or using other regulatory levers such as the regulatory ringfence) that could address concerns with risky financial structures?

6.8. We do not think it is helpful to cast the GOSM as a means to share between equity investors and customers benefits accruing to the former from higher gearing. Instead, the correct economic framing for regulatory intervention is to close the gap between the social and private (marginal) costs from higher gearing. While seemingly a small change in language, this is an important change in motivation and rationale.

6.9. Based on this changed framing, we look at two other options: one quantity-based, the other price-based. The quantity-based option would impose licence conditions related to financial resilience: most directly, a gearing cap as applies to NERL; but also potentially other measures. The price-based option is to impose a charge on firms with gearing above notional levels to create a fund that would meet the cost of special administration, if ever that is triggered. This is analogous to compensation funds used in financial services.

6.10. Our sense is that the quantity-based, licence condition route, while not without complication, is most likely to be effective.

Appendix A: Uncertainty about gearing

A.1. The Harris-Pringle formula gives

$$\tilde{\beta}_E^L = \alpha(\tilde{g}, g)\tilde{\beta}_E + (1 - \alpha(\tilde{g}, g))\beta_D,$$

where $\alpha(\tilde{g}, g) \equiv (1 - \tilde{g})/(1 - g)$. In this expression, ~'s denote random variables. So for example, $\tilde{\beta}_E$ is the raw equity beta which has to be estimated; $\tilde{g}$ is actual gearing, again which has to be estimated. $\tilde{\beta}_E^L$ is the relevered equity beta that comes via the Harris-Pringle formula, and so is a function of two random variables, $\tilde{\beta}_E$ and $\tilde{g}$, as well as two variables that are (assumed) known with certainty: notional gearing g and the debt beta β_D. For expository purposes we shall for the sake of simplicity assume that β_D is known, and equal to zero. We revert briefly to this additional complication below.

A.2. When actual gearing is measured with certainty, uncertainty in the estimate of the relevered beta comes from only one source: uncertainty about the raw equity beta. Hence in this case, if the estimate of $\tilde{\beta}_E$ is e.g., uniformly distributed on an interval, the estimate of the relevered beta is also uniformly distributed on an interval scaled by the value of α. To be very clear, if $\tilde{\beta}_E$ is uniformly distributed on [0.60,0.65], with a mean of 0.625; and if actual gearing is 55% and notional gearing 60%, then the relevered beta is uniformly distributed on [0.675,0.731], with a mean of 0.71. Importantly, the distribution of the relevered beta does not overlap the distribution of the raw equity beta: it is unambiguously higher.

A.3. But now suppose that actual gearing is also uncertain: it is uniformly distributed on the interval [0.50,0.60], say, with a mean value of 55%, as in the previous paragraph. (Arguably, this is a very modest degree of uncertainty about gearing. Once pension liabilities and swaps are incorporated in gearing calculations, the range of possible values of gearing extends above 90% for some water companies.) The relevered beta is now distributed as shown in Figure 1 below: the mean is the same as before i.e., 0.71 but it is now (non-uniformly) distributed with a range [0.600,0.813]. This illustrates the simple fact that uncertainty about actual gearing increases uncertainty about the relevered beta—in this case, by so much that its distribution overlaps that of the raw equity beta estimate.

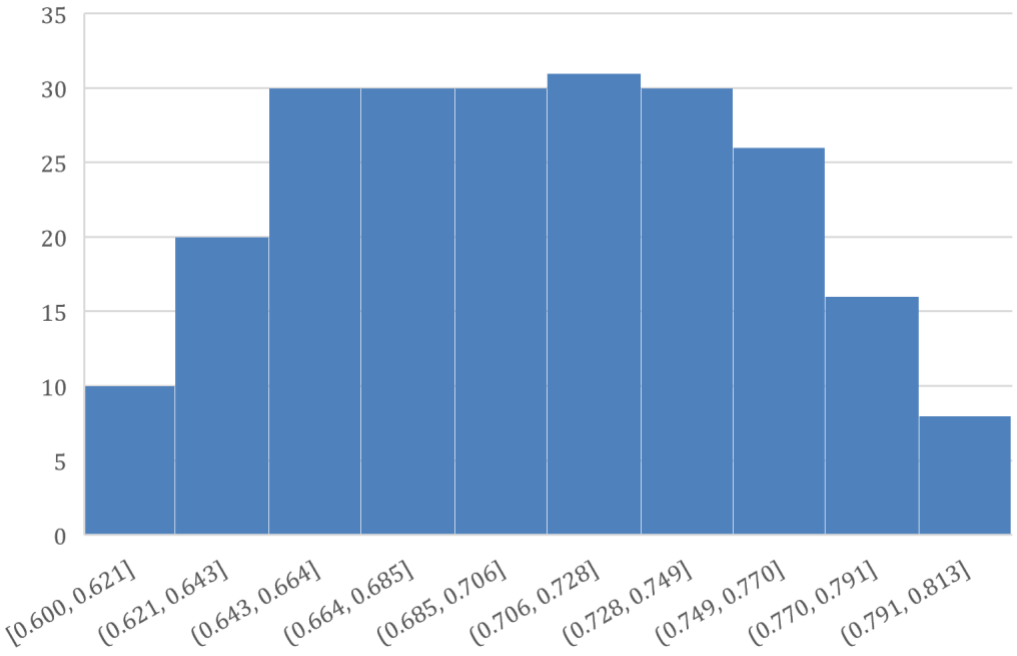

Figure A1: distribution of the relevered equity beta

Appendix B: The beta of an unlevered firm when Modigliani-Miller does not hold

B.1. In this Appendix, we show how to derive the unlevered beta. The value of a firm can be written in two distinct ways. In the first, the value of the firm V is the sum of the value of equity and debt: $V = E + D$. This is an accounting identity and does not rely on the MM theorem. In the second, the value of a firm with gearing is the sum of the value of an unlevered firm, denoted U, and the (net) value of borrowing B: $V = U + B$. B can be viewed as a fictitious security, the value of which is defined by the expected present value of the benefits minus the costs of debt financing. This too is an accounting identity, and so also does not need the MM theorem.

B.2. The MM theorem does have a bearing, however, on the value of B. In an MM world, $B = 0$ and so $V = U$; outside of this world, both B and V are non-zero; when debt adds value, they are both positive.

B.3. The mathematics of betas then mean that we can write $\beta_V = (1 - g)\beta_E + g\beta_D$; and likewise, $\beta_V = (1 - \gamma)\beta_U + \gamma\beta_B$, where $\gamma \equiv \frac{B}{V}$. In the latter expression, β_U measures the non-diversifiable risk of the unlevered firm; β_B measures the non-diversifiable risk of the fictitious security B; and g is the observed gearing of the comparator firms used to estimate β_E. Rearrangement gives

$$\beta_U = (1 - \lambda)\beta_E + \lambda\beta_D + \frac{\gamma}{1 - \gamma}(\beta_D - \beta_B),$$

where $\lambda \equiv \frac{g - \gamma}{1 - \gamma} < g$.

B.4. So, we have an expression for the non-MM beta of the unlevered firm; but it involves several additional variables that have to be estimated: in particular, β_B, the beta of the (fictitious) security representing the firm's debt, and γ, the proportional value of debt.

B.5. In some circumstances, we can avoid estimating β_B. Suppose that β_D and β_B are roughly the same. A case in which this holds exactly is when there is a constant marginal corporate tax rate, no non-tax benefits to debt, and no costs of financial distress: see e.g., Bierman and Oldfield (1979).[26] If departures from this benchmark are not too large, or more generally if $\beta_D \approx \beta_B$, then we can ignore the last term in the expression for β_U.

B.6. With this assumption, $\beta_U = (1 - \lambda)\beta_E + \lambda\beta_D$; that is, the beta of the unlevered firm is a weighted average of its equity beta and its debt beta. This appears to be very similar to the Harris-Pringle expression for the asset beta i.e., $\beta_A = (1 - g)\beta_E + g\beta_D$. But note that the weights for the beta β_U are given by $1 - \lambda$ and λ, rather than the weights $1 - g$ and g that appear in the Harris-Pringle formula; and so the non-MM beta of the unlevered firm β_U is greater than the asset beta β_A.

[26] Bierman, Harold, and George S. Oldfield, (1979). "Corporate debt and corporate taxes", *Journal of Finance*, 34, 951-956.

B.7. We are left with the task, then, of estimating λ, and hence estimating γ: the (unobserved) ratio of the net benefits of debt to the unlevered firm's value. Academic studies looking at the empirical evidence of the net benefit of debt find B to be no larger than around 10% of total value, so that B/V (i.e., γ) is no larger than $0.1/0.9 \approx 11\%$. To provide intuition, recall that B is the *net* addition to firm value from all outstanding debt, after allowing for the cost of debt issuance. Thus this estimate implies that, on average, for every £1 of debt issued, the firm increases its total value by 11p. If optimal gearing is positive, then the shadow value of an extra £1 of debt will be greater than £1, but at optimal gearing it will, by definition, be precisely £1, so this average estimate is consistent with a (possibly distinctly) higher marginal valuation for the unlevered firm.

B.8. With observed gearing of 54.2%, this gives weights of $\lambda = 0.48$ on the debt beta and 0.52 on the equity beta. Consequently, the non-MM beta of the unlevered firm will be somewhat, but not very much, greater than the asset beta (which has weights of 0.458 and 0.542).