

Estimating the cost of capital for implementation of price controls by UK Regulators

An update on Mason, Miles and Wright (2003)

Stephen Wright, Birkbeck, University of London

Phil Burns, Frontier Economics

Robin Mason, University of Birmingham

Derry Pickford, Aon Hewitt

Acknowledgements

We gratefully acknowledge the assistance of Mike Huggins, Rob Francis and Michael Yang (Frontier), Donald Robertson (Cambridge), Ken Hori (Birkbeck) and James Fernandes (Aon Hewitt) in writing this report.

Meline Danielewicz (Birkbeck) and Rubens Morita (Birkbeck) provided invaluable research assistance.

We are grateful to Ian Rowson (formerly Ofgem, now CAA) and other members of the UKRN Cost of Capital Working Group for helpful discussions while preparing this report. We also received very helpful advice from James Wardlaw (Campbell-Lutyens) and Catherine Waddams (UEA), both members of the UKRN Expert Panel.

Disclaimer

All of the authors are writing in a personal capacity. The views expressed in this report should not be interpreted as necessarily reflecting the views of our respective institutions.

1	Introduction	4
1.1	Our overall approach.....	4
1.2	Some key return concepts.....	6
1.3	Estimating the CAPM-WACC: Our 10 Key Recommendations	7
1.4	CAPM-WACC, RAR, RER, and bid premia.....	12
1.5	Using the CAPM-WACC: where we agree	14
1.6	Using the CAPM-WACC: where we disagree.....	15
2	Is CAPM the best available model for estimating the cost of equity for regulated companies? ..	16
2.1	Overview	16
2.2	A short history of the CAPM and subsequent developments	17
2.3	Does CAPM capture the key determinants of the cost of equity for regulated companies?.....	19
2.4	The Low Beta puzzle.	21
3	A simple benchmark CAPM-WACC framework	22
3.1	The vanilla WACC and the asset beta.....	22
3.2	The CAPM-WACC vs the CAPM(E)-WACC	23
3.3	The impact of leverage on equity beta, and the role of the risk-free rate	24
3.4	Insights from multifactor models.....	25
3.5	CAPM and the Dividend Discount Model	25
4	How should regulators estimate the key components of a market-consistent CAPM-WACC? ...	28
4.1	Investment Horizon.....	28
4.2	The appropriate price index for constructing real returns.	30
4.3	The risk-free rate (RFR).....	31
4.4	The expected market return (EMR).....	36
4.5	Beta Estimation Issues	49
4.6	Leverage (gearing) and its implications for asset beta, the assumed cost of equity, and the CAPM-WACC	55
4.7	The market cost of debt.....	58
4.8	Alternative historical estimates of the market-consistent CAPM-WACC for regulated utilities 62	
5	Implementing the CAPM-WACC in regulation: the Regulatory Expected Return (RER) vs the Regulatory Allowed Return (RAR).....	64
5.1	The Regulatory Expected Return (RER) and the Regulatory Allowed Return (RAR): some definitions.....	64
6	Attempting to understand recent market valuations of regulated utilities (Burns).....	66
7	Areas of Agreement and Disagreement	68
8	Using CAPM-WACC estimates for regulation (Mason, Pickford and Wright).....	68
8.1	A simple conceptual framework	68

8.2	Aiming Up: Why the Regulatory Expected Return (RER) should be higher than the CAPM-WACC	70
8.3	A simple implementable and defensible framework for setting the RAR and the RER.	73
8.4	Some possible objections to our ideal framework	75
8.5	Two possible compromises	77
8.6	Conclusions: where next?.....	79
9	Implementing the CAPM-WACC in regulation: the Regulatory Allowed Return (RAR) (Burns)..	
9.1	Introduction	80
9.2	Should RAR differ from the CAPM-WACC?	81
9.3	Should a mechanism be incorporated into the regulatory price-setting process to force RER to converge on RAR?.....	86
	Appendices.....	A-89
A	Are Investor Investment Processes Consistent with the CAPM? (Derry Pickford).....	A-89
B	Factor models of Equity Returns (Derry Pickford).....	B-98
C	The asset beta for an unlevered firm	C-107
D	Choice of price index and adjustment of real returns (Derry Pickford)	D-109
E	Updating evidence of the historic market return.....	E-125
F	Beta estimates (Phil Burns)	F-126
G	Beta Estimation for CAPM-WACC at Long Horizons (Stephen Wright and Donald Robertson) G-139	
H	Adjusting bond yields for default risk (Ken Hori, Birkbeck)	H-154
I	Aiming up: analysis of asymmetry in setting the RAR (Robin Mason).....	I-158
J	What drives bid premia for regulated utilities (Phil Burns).....	J-165
K	Evidence of Outperformance (Stephen Wright)	K-178
L	Terms of Reference	L-182
	Works Cited	185

1 Introduction

1.1 Our overall approach

This report takes a fresh look at the role of the cost of capital in the regulation of UK utilities. An earlier report (Mason, Miles and Wright, 2003, hereafter MMW) addressed a number of issues relating to the measurement of the components of the weighted average cost of capital (WACC) – primarily focussing on the measurement of the cost of equity. This report updates this analysis but also takes a broader perspective.

The authors were commissioned jointly by the CAA, Ofcom, Ofgem and the Utility Regulator¹, (collectively referred to as “the regulators”) to answer a number of questions framed in the terms of reference (ToR) for this study, included as Appendix L. The ToR make reference in particular to the current environment of low interest rates, as well as recent elevated premia over regulated asset values. The ToR identify two key questions: first, how should regulators determine the WACC given this environment? Secondly, what factors might explain the recent premia?

To address this task, we have assembled a group of authors and other contributors who are academics, practitioners and consultants, from a range of relevant disciplines and perspectives. We believe these diverse perspectives have created some positive creative tension in the study and should help to make it more relevant. It has also highlighted some areas where we have not been able to agree. We hope that areas of agreement and disagreement will help to stimulate a further dialogue between regulators, interested parties and experts. All authors, and hence the views and recommendations expressed, act in an independent capacity and the conclusions of this report has not been endorsed by the regulators.

We have adopted four principles in formulating our view and recommendations:

- Both current academic evidence and investment practice must be taken into account.
- All aspects of the expected return for regulated companies must be considered.
- The overall approach must be both implementable and defensible.
- The degree of discretion that can be applied by regulators should be limited, to no more than necessary to ensure consistency of treatment over time, which is necessary to promote the credibility of the regime and to manage investors’ perception of regulatory risk.

We have taken great care throughout to ensure that our recommendations are consistent with the statutory duties of the regulators. These statutory duties are rather complex and vary across the regulators; and indeed, for a particular regulator, can vary by activity.² We focus in this report on two

¹ Although this subgroup of UKRN members commissioned this report, representatives from other UKRN members, notably Ofwat, were involved in its governance.

² For example, Ofgem’s statutory duties vary according to whether it is gas (determined by the Gas Act 1986) or electricity (determined by the Electricity Act 1989) under consideration (both Acts updated by the Utilities Act 2000). Broadly, the principal objective is to protect the interests of existing and future consumers. Subject to this principal objective, Ofgem must have regard to the need to secure that companies are able to finance activities. Ofcom’s duties, set in the Communications Act 2003, are (a) to further the interests of citizens in relation to communications matters; and (b) to further the interests of consumers in relevant markets, where appropriate by promoting competition. It must also have regard to the desirability of encouraging investment and innovation in relevant markets. The specific objectives and functions of the CAA are set out in the Civil Aviation Act 1982, the Airports Act 1986, the Transport Act 2000 and the directions given under section 66 of the Transport Act 2000. They include regulating civil aviation safety; determining policy for the use of UK airspace so as to meet the

duties that are broadly common to the regulators: to (i) further the interests of both existing and future consumers; (ii) secure that regulated firms can operate efficiently. Hence when assessing how and at what level the cost of capital should be determined, we pay close attention to the balance between the benefits to current consumers flowing from low prices, and the incentives for investors to provide financing for efficient levels of investment, to the benefit of future consumers. Ofwat, in its statement of the final methodology for the 2019 price review (Ofwat (2017)), has summarised the issue well:

“If the cost of capital is set too high, bills may be higher than customers may reasonably expect, company profits may be seen as excessive and the legitimacy of the regulatory regime may be called into question. If the cost of capital is set too low, companies’ ability to raise the finance necessary to deliver services that customers expect might be put at risk.”

We also pay close attention to incentives for efficiency, and their relation to expected returns.

Given the unavoidable tension that can arise between these two common duties, we emphasise especially the need for any recommendations to be both implementable and defensible. By implementable, we mean an approach that does not entail excessive cost or complexity to carry out. Implementability is one of the reasons why we prefer the capital asset pricing model, CAPM, to e.g., multi-factor asset pricing models, despite the caveats that apply to the CAPM. By defensible, we mean that an approach has to be robust to reasonable criticism. For example, defensibility lies behind our recommendation to base the estimate of the expected market return on long-run historic averages. These two criteria are clearly related to the five regulatory principles emphasised in UK regulation: that regulatory activities should be transparent, accountable, proportionate, consistent and targeted only at cases in which action is needed. We find it helpful, however, to recast certain of these principles for the case of cost of capital, given some of the particular issues that arise in this area of regulation. We highlight these issues throughout the report.

We therefore recognise that the regulators are not estimating the cost of capital as an academic exercise. They are generally mandated by law to further specific objectives or exercise specific duties. Generally, these are framed around the interests of existing and future consumers, although other objectives and duties are also relevant, and regulators estimate the cost of capital and other relevant inputs to price controls with these objectives and duties in mind. How regulators estimate the cost of capital can be particularly relevant for capital intensive activities because it can, via the price setting process, affect how much consumers pay and the decisions taken by investors and competitors. The authors of this study therefore recognise that the cost of capital estimation process requires regulators to make judgements that may go beyond a simple application of finance theory and data.

needs of all users, having regard for national security, economic and environmental factors, while maintaining a high standard of safety; economic regulation of airports and air traffic services; licensing and financial fitness of airlines and licensing of air travel organisers. Other principles apply, subject to these primary principles. The Utility Regulator’s duties are specified in the Energy (Northern Ireland) Order 2003; the Water and Sewerage Services (Northern Ireland) Order 2006; the Gas (Northern Ireland) Order 1996; the Electricity (Northern Ireland) Order 1992; and the Electricity (Class Exemptions from the Requirement for a Licence) Order (Northern Ireland) 2013. Its objectives include protecting the short and long-term interests of electricity, gas, water and sewerage consumers, and promoting a robust and efficient industries. Ofwat’s statutory duty, as set out in section 2 of the Water Industry Act 1991, include the primary principles to: (a) further the interests of both existing and future water consumers (the ‘consumer objective’); (b) secure that water companies properly carry out their functions; and (c) secure that they are able to finance those functions, in particular, by securing reasonable returns on their capital. All regulators are also expected to have regard to the principles of best regulatory practice, including ensuring transparency, accountability, proportionality, and consistency.

We make a number of recommendations, on which all co-authors are (largely) agreed, about how the WACC should be estimated. On issues relating to implementation of these estimates, we have also been able to achieve broad agreement on a conceptual framework. We do, however, differ on the nature of the recommendations that arise from this framework.

1.2 Some key return concepts

The WACC is an expected return. In finance theory, this expected return provides the benchmark against which the (risk-adjusted) expected return on any investment is judged. If the risk-adjusted expected return is not at least equal to the WACC, then the investment should not take place; if it exceeds the WACC, then the investment is expected to yield above-market returns.

Regulators need to estimate this expected return, and we devote the bulk of our report to the issue of how best to estimate it, and its key components. But we also address its use in regulation. To assist our analysis, we distinguish between different return concepts:

- **The WACC**, a weighted average of the prevailing expected market cost of equity, adjusted for risk, and the cost of debt. The true WACC can never be known, but must be estimated.
- **The CAPM-WACC**: the *estimate* of the WACC that comes from using the capital asset pricing model (CAPM).
- **The Regulatory Allowed Return (RAR)** corresponds to what has in the past typically been referred to as the “allowed WACC”, which we argue is a misnomer. The RAR is the return on the regulatory asset base before allowing for the impact of outperformance or underperformance on cost or service level.
- **The Regulatory Expected Return (RER)**, which is the RAR plus any expected increase in returns, mostly arising from outperforming the cost and service targets set by the regulator.³ To the extent that any outperformance is predictable, the RER will be higher than the RAR. We introduce the RER as a distinct concept because, as an expected return, it is directly comparable with estimates of the CAPM-WACC.

We argue that it is important to be clear on the distinction between these concepts. A subset of the authors would argue that there have been a number of problems in implementation of regulation that have stemmed from a conflation of these quite distinct concepts.

Both in theory and in practice the RER and the RAR may differ both from each other and from the regulator’s best estimate of the WACC. We also note that the methods and principles that regulators use to set the RAR and RER can feed back into the WACC through their impact on the risks faced by investors in regulated companies.

It is helpful at this point to spell out further the distinctions we are making between, in particular, the CAPM-WACC, the RAR and the RER. The RAR may in practice differ from the CAPM-WACC. A primary reason for this “regulatory wedge” has been a concern about disincentivising investment, along with an asymmetric loss function which makes under-investment costlier (in a sense that we make precise in Appendix I) than over-pricing, at least at the margin. While we recognise some validity in this point, we argue that it can justify only a limited gap between the RAR and the CAPM-WACC.

³ Strictly speaking it might also adjust for any expected reduction in returns arising from (predictable) underperformance relative to regulators’ targets, or other risk factors.

We also note that we cannot consider the determination of the “regulatory wedge” in isolation. The RER may be, and in UK regulatory practice often has been, greater than the RAR. The primary reason for this is the form of price-cap regulation used in the UK, which uses the prospects of additional returns during a price review period to give regulated firms an incentive to reduce costs or improve performance. That is, a positive gap between the RAR and the RER (which we term the “informational wedge”) is a deliberate feature of the regulatory system in the UK. As with the regulatory wedge, the key question is how large it reasonably should be.

There is some evidence that in recent years, realised returns on regulated capital have been higher than the RAR. In addition, market valuations of certain regulated firms show significant premia over the value of the regulatory asset bases: by over 40% in some cases, and on average 24% over the period 1998-2016. In terms of the concepts that we have introduced, we see two main reasons for these premia:

- Investors expect regulated firms’ allowed returns to be set higher than actual current financing costs i.e., the RAR is above the CAPM-WACC.
- Investors expect regulated firms to receive payments through cost or service outperformance that boost their ex-post returns and their expected returns i.e., the RER is above the RAR.

The authors differ on the weights they give to these two factors.

1.3 Estimating the CAPM-WACC: Our 10 Key Recommendations

The weighted average cost of capital (WACC) is a theoretical construct: a weighted average of the prevailing expected market cost of equity and cost of debt. A central feature of the WACC is that the expected cost to the company is an expected return to the investor. This is the fundamental basis for the way that we attempt to measure it.

We make 10 recommendations.

The true WACC can never be known. We can only attempt to measure it by reference to a particular model of capital markets. In Section 2, we argue

Recommendation 1 (CAPM): The Capital Asset Pricing Model remains (despite numerous caveats) the best available model. We do acknowledge the insights that other models can provide, which may at least make them useful as a cross-check on CAPM-WACC estimates. But we argue that competing models all fall foul of one or both of our twin criteria of implementability and defensibility.

In Section 1.2, we set out the key features of the WACC as applied in that model. This leads us to derive estimates of the CAPM-WACC in Section 4. Crucially we estimate the WACC as a *market* concept, which we view as a quite distinct exercise from the issues of implementation in regulation via the RAR and the RER. But we stress that, to be of any use to regulators, any such estimates must satisfy the twin requirements of implementability and defensibility, and that these requirements limit the degree of discretion that can be applied by regulators.

We make 9 further specific recommendations on the following issues

Recommendation 2 (Horizon): On balance, we are in favour of choosing a fairly long horizon, for example, 10 years, in estimating the CAPM-WACC. But we would argue that, more important than the choice of horizon per se is that all components of the CAPM-WACC are estimated using a methodology that is *consistent* with the chosen horizon.

Recommendation 3 (Price Index): There is a strong case for regulators choosing a measure of inflation for estimating the CAPM-WACC that is consistent with that chosen by HM Treasury and implemented by the Bank of England for inflation targeting (currently, CPI). Ideally, we would like to use the same index used in setting the Regulatory Allowed Return (RAR). We can, however, see merit in regulators postponing any decision to change index until there is clarification on whether the measure of prices used for inflation targeting by the Bank of England will be changed.

Recommendation 4 (The Risk-Free Rate): Regulators should use the (zero coupon) yield on inflation-indexed gilts at their chosen horizon to derive an estimate of the risk-free rate at that horizon. For consistency with the treatment of inflation outlined above, the published yield needs to be adjusted for the predicted wedge between the inflation rate used for estimating the CAPM-WACC (and ideally also setting the RAR) and that used for indexing gilt (currently RPI inflation).

We are conscious that this recommendation (which is essentially a restatement of the original recommendation made by MMW) implies a nontrivial change compared to regulatory practice in recent years. Most regulators have made reference to market gilts data in their assessments but then have set a different value when setting the RAR.

While we accept the argument that the market for risk-free debt may possibly be distorted, we do not regard this as a justification for ignoring the prices that arise from this market.

Recommendation 5 (The Expected Market Return): We recommend that regulators should continue to base their estimate of the EMR on long-run historic averages, taking into account both UK and international evidence, as originally proposed in MMW. We suggest a modest downward adjustment of the original range proposed by MMW, to a range of 6-7%, primarily reflecting a smaller adjustment from geometric to arithmetic returns. We acknowledge fairly strong evidence, from academic research, from investor surveys, and from some careful implementations of the Dividend Discount Model, that the EMR has some degree of predictability. But our criteria on implementability and defensibility again come to the fore here. We are unable to point to a methodology that would capture predictability of the EMR that would be as straightforward to implement as the existing approach, nor—given the still-contested and partial nature of the evidence for predictability—that would be robust to criticism. We do however recognise the value of other approaches in at least providing a cross-check on our preferred approach.

Recommendation 6 (Beta Estimation): Regulators should make more use of robust econometric estimates of equity beta. They should derive these estimates from sound econometric evidence and practice, utilising all available data for relevant listed companies. Betas for unlisted companies should be derived from estimated equity betas from the closest available comparator listed companies.

Appendix F, contributed by Phil Burns, summarises the standard approach to estimating equity betas, using the rolling OLS approach on daily, weekly and monthly data for 10 comparator companies applied to 2, 5 and 10-year estimation windows. This currently points to equity beta estimates of between 0.77 and 1.03. These are broadly consistent with the estimates derived by UK regulators at previous price control reviews.

While the authors can all agree on the crucial role of econometric estimates, they differ on the conclusions to be drawn from that evidence

Mason, Pickford and Wright's views on beta estimation

MPW believe that the relatively high values of equity beta assumed in certain recent price controls (for example, the values of 0.8 and 0.9 chosen by Ofwat and Ofgem respectively) are inconsistent with econometric evidence. In Appendix G, Stephen Wright and Donald Robertson point to strong evidence that the recent rise in daily beta estimates is temporary, and likely to be reversed. They argue for estimation on longer-term data and at lower frequencies (which they argue provide evidence more relevant to the long horizons used by regulators), which results in distinctly lower equity beta estimates. These are in the range 0.3-0.5, for the only two pure-play UK regulated companies that are currently listed: distinctly lower than values assumed in recent price controls. While these results are preliminary, MPW argue that regulators should take very seriously the implications of lower values of equity betas, and hence asset betas.

Burns' view

Burns argues that whilst the approach taken by MPW is of interest, it is but one econometric exercise of just two stocks in an already voluminous literature. MPW's results based on higher frequency data are recognisably similar to the existing regulatory estimates over the relevant time-frames, but MPW also adopt the highly unusual practice of estimating the CAPM on quarterly data, which is the key factor that drives the lower estimates of beta. It should be noted that using quarterly frequency data is unusual in academic studies and is not used by any commercial provider. As it stands it does not satisfy the criteria used in this report.

In the light of recommendations 4 and 5, the estimation of beta is the one component of the cost of equity where the regulator must use its judgement and discretion. This places an obligation on regulators to examine the evidence as a whole, not simply relying a single approach that results in outlying estimates, in order to retain the benefits of a stable and transparent approach to setting the RAR. This approach has successfully driven down the WACC over the past 25 years as the perception of regulatory risk has diminished, and this stability has also contributed to a stable commercial environment within which operators have made significant dynamic efficiency improvements.

The time may now be ripe to more comprehensively explore the GARCH approach and its many variants. That research properly undertaken would then form part of the evidence base that regulators could draw upon that would sit alongside pre-existing regulatory approaches. In other words, it would be a complement to pre-existing techniques, rather than a substitute for them. Until that research is more comprehensive, regulators should continue to use the CAPM on a wide range of comparator stocks, using higher frequency data (subject to testing for thin-trading and serial correlation), over different sample sizes, and interpret that body of evidence judiciously, in line with practice to date.

Recommendation 7 (Adjusting beta estimates for leverage): Regulators should exercise care in allowing for the impact of leverage, in deriving asset beta estimates and in “re-gearing” to derive equity betas based on assumed levels of regulatory gearing. If beta estimates are derived from lower frequency estimation (as proposed by MPW), then the treatment of debt betas also needs to be consistent.

Burns' Appendix F carries out the approach currently applied by regulators of “de-gearing” estimated equity betas to derive asset betas, and then, where regulator's notional leverage differs from observed

market leverage, they then “re-gear” to derive the equity beta to be inserted into the CAPM model actually used by regulators.

On this issue, again the authors cannot fully agree.

Mason, Pickford and Wright’s views on de-gearing and re-gearing

MPW broadly agree the methodology used to estimate asset beta; but are very sceptical of the assumptions underlying “re-gearing”. They note that, as currently applied to listed pure-play regulated companies, this results in an assumed value of equity beta distinctly higher than the preferred econometric estimate. They argue that estimated equity betas determine the marginal cost of equity, and should be used directly where companies are listed. For unlisted companies, or listed companies with unregulated activities, the methodology should be as close as possible to that applied to the pure-play listed companies.

Burns’ views on de-gearing and re-gearing

Burns disagrees with MPW’s interpretation of this recommendation. What de-levering and re-levering achieve is to take the actual gearing level of the beta comparators out of the beta estimates and to assess what the equity beta of a company with similar risk would be, at a consciously determined ‘notional gearing’ level. The marginal cost of equity finance will vary with the level of gearing. For example, suppose the listed companies happened to have been geared up to well over the notional gearing level (e.g., 90%). By inference from finance theory, the equity beta of these companies would be much higher. It would be inappropriate for the regulator to use the observed high marginal cost of equity in the cost of equity calculation at the notional gearing level. Equally, suppose the comparators have very little debt. The marginal cost of equity would be very low, and it would also be inappropriate to use that equity beta in combination of a notional gearing level of 65% to calculate the cost of equity.

This issue is currently resolved by the regulator in each sector determining the notional gearing level. The rationale of setting a notional gearing level has become well understood and carries good properties such as encouraging efficient financing by leaving the financing decisions and responsibilities to the companies. Further, we note that the notional gearing level of recent decision from Ofwat and Ofgem have been informed by a combination of the average actual gearing level in the sector and the reasonable gearing range proposed by credit rating agencies for an investment-grade credit rating. To convert equity beta values from actual to these notional gearing levels it is appropriate for regulators to use standard models of re-gearing, embedded in well-established finance theory, as the basis for this conversion. This approach avoids the arbitrary outcome that would result from setting the cost of equity for regulated businesses by reference to the capital structure of a single, unusual comparator.

That being said, the letter of the recommendation is valid. In particular, in situations where there is a material difference between actual and notional gearing, regulators should carefully consider the specific method for re-levering, for example through consideration of the appropriate debt beta to apply to de-levering and re-levering.

A related issue that has been raised is the situation where the actual level of gearing is low relative to the notional level as a result of the valuation of expected future regulatory outperformance (in other words the market value of the regulated business is above the Regulatory Asset Base, RAB). It could

be argued that the outperformance element of the value may have a higher beta (for example because regulators are more likely to target the outperformance during times of economic recession). In that situation the regulator could seek to focus on the underlying beta that relates to the ‘core’ RAB return and not the outperformance component. This situation is analogous to a decomposition of an equity beta to adjust for non-regulated activities, rather than being a challenge to the standard re-levering approaches. It is noteworthy that the implementation of this decomposition raises some practical challenges. First, estimating the extent to which outperformance value carries greater risk than ‘core’ RAB value is not straightforward. Second, as we discuss in the transaction analysis (Appendix J) and the discussion of RER there is an element of market value premium that relates to the ongoing cost discovery and innovation on cost and services that is central to an incentive based regime and this element should be treated in the same way as the ‘core’ RAB.

Recommendation 8 (Estimating Default Risk on Corporate Debt): For consistency with the definition of the CAPM-WACC as an expected return, cost of debt estimates that feed into estimates of the CAPM-WACC should include an adjustment to corporate bond yields to convert these to expected returns. This should ideally take account of evidence of time variation in default risk.

Recommendation 9 (The Cost of Debt): The term “cost of debt” should relate to the expected return (correctly adjusted for default risk) on a traded corporate bond, at the regulator’s chosen horizon, and with risk comparable to regulated utilities at regulators’ chosen leverage. It should be clearly distinguished from the *allowed return* on embedded debt.

Recommendation 10 (The CAPM-WACC Methodology): The term “WACC” should be restricted to the concept of an expected market return on capital of a given degree of systematic risk. It should *not* be used to refer to an allowed return. Our use of the term “CAPM-WACC” is deliberately designed to emphasise the necessary link between the CAPM model and a *methodology* for estimating the WACC. This methodology satisfies the twin requirements of implementability and defensibility.

One conceptually simple and easily implementable option for measuring the WACC, is the “pure” (or equity-only) CAPM-WACC:

$$\text{CAPM-WACC} = RFR + \beta_A (EMR - RFR) \quad (1)$$

This estimate can be rationalised in two ways:

- It can be derived straightforwardly by assuming that the CAPM applies to both equity *and* bond markets (but is subject to the caveat that risk premia on corporate bonds appear to be distinctly higher than the CAPM would imply);
- But it can also be rationalised as an estimate of the WACC for an equity-only firm. As such, if some positive leverage is optimal (which must imply a *lower* WACC), this estimate should represent an upper bound for the WACC of a firm with an optimal level of debt.⁴

⁴ This is in turn subject to the caveat that asset beta β_A is itself not invariant to the level of leverage. Robin Mason addresses this issue in Appendix C to the report and concludes that the implied adjustment to β_A are likely to be small.

All the authors of this report are agreed that this estimate of the WACC should at the very least be used as a cross-check on more conventional estimates of the WACC that are a weighted average of a CAPM-based cost of equity and the observed cost of debt (as defined in Recommendation 9). (Since this conventional approach only applies the CAPM to equity, we refer to this estimate as the CAPM(E)-WACC.)

1.4 CAPM-WACC, RAR, RER, and bid premia

The observed bid premia on regulated companies can be related to our different return concepts.

First, and straightforwardly, if we assume markets use the (unobservable) WACC to discount expected cashflows from regulated companies, then there is a necessary equivalence between bid premia and a gap between the Regulatory Expected Return (RER) and the WACC. Since, as noted above, any difference between the two implies benefits that accrue only to equity holders, the implied gap between the expected return on regulatory equity and the true market cost of equity must be distinctly larger (i.e., must scale up by a factor of $1/(1-g)$ where g is leverage (gearing)).

Observed bid premia for certain recent transactions imply that these differences must be very large indeed. Calculations by Burns (see Appendix J) imply that if we express the differences in perpetuity-equivalent terms, then the return on regulatory equity expected by the buyers of the company at the point of sale must be anything up to around twice the true market cost of equity.⁵ (This calculation assumes that there is no winner's curse or control premium—which Burns argues is likely to be a strong assumption.)

Since neither the true expected return on regulatory equity nor the true market cost of equity can be directly observed, we can never be certain of what the explanation is for this gap; but we *can* be sure that such a gap must exist. However, once we commit to an estimation strategy based on the CAPM-WACC, as outlined above, we can at least attempt an explanation.

We argue below that observed bid premia can be re-expressed in terms of our three return concepts; we show that

- the Regulatory Allowed Return (RAR) has systematically exceeded the CAPM-WACC (as estimated using our preferred methodology); and
- the Regulatory *Expected* Return (RER) has in turn systematically exceeded the RAR, because regulated companies have systematically exceeded their cost targets.

The authors of this report differ on the extent to which they attribute bid premia to these two causes:

Mason, Pickford and Wright's views on bid premia

MPW argue that the gap between the RAR and the CAPM-WACC and the gap between the RAR and the RER have both made significant contributions to investor expectations. Indeed, recent bid premia imply such large differences between the RER and the WACC that it is very hard to provide an explanation that does *not* rely on both making a material contribution. Regardless of the source of bid premia—be it an allowed cost of capital that is too high, or outperformance payments—we argue that

⁵ Note that this need not imply that markets actually expect this gap to persist indefinitely; but it does imply that if they expect it to narrow over time, then the current gap must be even larger than this.

regulators need to be clear what overall expected return (RER) they are aiming for. Since we have clearly not considered the regulatory system in the round, we cannot (and do not) say at what level the RER should be set. This is a matter for regulators, balancing the need to incentivise firms towards efficiency, versus ensuring greatest benefits (returned) to consumers. In doing this, we think the contrast between the expected returns on regulated firms, and the realised returns elsewhere in economies, documented by e.g., Jorda et al. (2017), should be borne in mind. The endpoint should be a much tighter, and more transparent link between the returns regulated companies can expect to earn and our best estimates of their cost of capital.

Burns' views on bid premia

Analysis by Burns in Appendix J focusses on the recent transaction premia for energy networks. He finds that outperformance of cost and output targets and the cost of debt allowance could account for close to 40% RAB premium for an optimistic investor that expected current levels of outperformance to persist (and it is the most optimistic bidder who ultimately wins the transaction). Expected outperformance on the cost of equity could also contribute to the observed premium, and to explain a material proportion of the premia the optimistic investor would need to expect significant cost of equity outperformance to persist. Other factors such as further financial restructuring, tax arbitrage, control premium and winners curse are likely to also be of significance. Identifying a satisfactory explanation for high observed transaction premia is a challenge. It can be argued that expecting current levels of cost and incentive outperformance in the energy networks to persist is unrealistic (even though it has been a persistent feature since privatisation) and equally it can be argued that regulators would act to remove any cost of equity outperformance over time.

Fundamentally, the analysis highlights the challenges that arise in seeking to use transaction premia evidence to make inferences about the cost of equity. Different drivers of outperformance are at play and multiple combinations of various drivers can explain observed premia. The role of expected outperformance means that the premia may result from unobserved investor assumptions that may be considered unrealistic or optimistic but are nevertheless the reality behind the premia. Evidence from Market-to-Asset ratios (MARs) for quoted pure-play utilities are generally not subject to the issues of control premium and winners curse, though there remains the challenge of understanding the unobserved investor assumptions.

Overall, evidence from transaction premia is less reliable and much harder to interpret than other sources of evidence on the cost of equity.

The gap between the RAR and our estimated CAPM-WACC has come about in a number of ways.

- Almost all regulators have been slow to bring their estimates of the risk-free rate down as market rates have decreased.
- The regulatory allowed return on debt, which attempts to capture the cost of embedded debt, has also been persistently higher than the market cost of debt. It is worth emphasising that for most of the period since privatisation, both nominal and real interest rates have been falling and these reductions have been largely unanticipated by markets. Therefore, to the extent that the allowed return on debt as a component of the RAR has been set in a way that reflects yields on

long-term fixed rate debt (i.e., embedded debt) this has driven a wedge between the RAR and the CAPM-WACC.

- Mason, Pickford and Wright would also argue that regulators’ assumed values of equity beta have been too high, compared to the plausible range of econometric estimates. Allowing for this would imply either a distinctly lower cost of equity (in the CAPM(E)-WACC) or, via a lower asset beta, a lower overall WACC (as implied by our “pure” CAPM-WACC estimate).

But we would all agree that, while the gap between RAR and CAPM-WACC is large, it is not *sufficiently* large to explain the entirety of bid premia. Thus Appendix J points to the need to assume a continuing gap between the RER and the RAR to explain the magnitude of recent observed premia. We argue that this is also consistent with the evidence available to us (which does not cover all regulated sectors) on historic return outperformance.

1.5 Using the CAPM-WACC: where we agree

While the primary remit of our report was to focus on estimation of the cost of equity, it is clear that a number of issues raised in the terms of reference can only be addressed by considering the way CAPM-WACC estimates are *used* by regulators: i.e., in our terminology, how they translate to the RAR and hence to the RER.

Why do we focus on the link between the CAPM-WACC and the RER, rather than the RAR? We have debated this issue at some length (and have not yet fully achieved consensus). As discussed above, the RER is, by definition, an expected return, and thus directly comparable with the CAPM-WACC. And, as noted above it is the gap between the RER and the WACC that must determine bid premia.

Clearly the RAR (previously known as the allowed WACC) plays a crucial role, since it is what regulators can directly control. The way the RAR is set is also fundamental for managing perceptions of regulatory risk and creating a stable commercial environment within which to operate efficiently.

But the RAR is *not* an expected return (if it were, we would not need the distinction). Thus, it is not directly comparable with the CAPM-WACC. The RAR and the expected rewards from exceeding targets set by the regulator⁶ *jointly* determine the RER. Observing the RAR in isolation does not tell us about expected returns unless we also know the nature of the targets. Thus, a regulator could in principle lower or raise the RAR, but simultaneously change targets by some amount such that the RER was left unchanged. However, for *given* targets a rise or fall in the RAR clearly equates to a rise or fall in the RER.

We argue that a number of the decisions made by regulators in recent years that really should have been about *estimating* the WACC have been motivated by issues relating to *implementation* of those estimates in regulation. Primary examples have been the treatment of the risk-free rate, the cost of debt and (MPW would argue) use of beta estimates. We show (see Figure 4.13 below) that our CAPM-WACC estimates have fallen markedly in recent years, but UK regulators have not updated their estimates of the WACC to reflect all of this reduction (in large part due to their treatment of embedded debt).

We argue that regulators should treat the activities of estimation and implementation as quite distinct.

⁶ Which may include explicit incentives for information quality (e.g., Ofgem’s IQI mechanism).

We should stress that we all recognise that both the RAR and RER can *legitimately* differ from the CAPM-WACC:

- There is a legitimate—albeit limited—case for “aiming up”, by setting the RAR in such a way that the (regulator’s best estimate of the) RER lies above the midpoint estimate of the CAPM-WACC. All the authors of this report agree with this objective, although different authors emphasise different aspects of the rationale. Thus MPW focus on arguments based on the asymmetric risks of getting the WACC estimate wrong (see Appendix I). Burns also points to the need to have regard to financeability and to signal stability of the regulatory approach (with the aim of reducing the CAPM-WACC over time).
- The RER is also likely to exceed the RAR by at least some amount. In a world of perfect information, regulators would be able to set incentivised targets equal to expected values. The RER and RAR would thus be equal. Firms would then not be able to earn any economic rent arising from any informational advantage they have over the regulator, or benefit from cost discovery through profit incentives leading to improved productive and dynamic efficiency. But in practice, we acknowledge that these rents are—to some extent at least—unavoidable. The key issue, then, is the size of the informational rents: or, by another, related measure, the size of the expected payments made to regulated firms for outperforming targets.

All the authors also agree that in setting the RAR, regulators should also provide an estimate of the RER, and provide clarity over the drivers of the differences from the estimated CAPM-WACC. In this respect, the approach Ofgem has taken since DPCR5, for example, which is to analyse how RoRE can be expected to change in response to different levels of company performance, is a step in the right direction.

We note however that, to achieve these objectives, regulators would need to assemble a much more consistent dataset on historic *realised* returns than is currently easily accessible. At present, we can observe the RAR and our estimates of the CAPM-WACC, but any attempt to estimate the historic levels of RER, both prospective and historic, would currently fail for lack of consistently collected data over time.

Thus far the authors are (largely) in agreement.

1.6 Using the CAPM-WACC: where we disagree

The authors have failed to agree on the precise mechanism by which CAPM-WACC estimates should feed into the regulatory process. On this topic, we present our views separately at the end of the report. We present a brief summary here.

In Section 8, MPW argue that a crucial feature of any regulatory regime should be explicit numerical targets for the two “wedges”: the regulatory wedge (the gap between the RAR and the estimated CAPM-WACC) and the “informational wedge” (the gap between the RAR and the RER). They suggest a top-down approach, in which the crucial target should be the *sum* of the two wedges, which should be set to achieve an explicit “aiming up” objective. They note that a logical implication of aiming up is that on average at least, bid premia should be positive; but argue that it would be extremely difficult to defend the magnitude of recent bid premia on grounds of aiming up alone. MPW then argue that, since a positive “informational wedge” is also almost certainly an unavoidable feature of incentive-based

regulation, the “regulatory wedge” should then be set to achieve a desired degree of aiming up, *conditional* upon an explicit value of the informational wedge.

MPW also argue that, in an ideal world at least, regulators should use the “pure” CAPM-WACC as the basis for setting the RAR. However, they acknowledge the practical and political difficulties of implementing regulation in this way. They argue that a compromise approach, in which the allowed return on equity was updated in line with its implied CAPM value, but regulators continued to set the allowed return on debt on the basis of a mechanistic estimate of the cost of embedded debt, would at least have the merit of linking the RAR to the estimated value of the CAPM(E)-WACC, on average, over longer periods of time.

In Section 9, Burns’ position is similar to the compromise position that MPW settle on, but he starts from a different perspective. He argues that the pure CAPM-WACC approach would not benefit customers and may even be detrimental, relative to the current regime. The wedge that exists between the RAR and the CAPM-WACC is almost entirely due to the regulatory treatment of embedded debt costs, and as MPW acknowledge, the regulatory approach to embedded debt is a quite tightly constrained dragging anchor. So, whilst customers may currently be “overpaying” relative to the CAPM-WACC, in future they could be underpaying. It is unclear whether there would be any benefit for customers over the long term, since the approach to remunerating debt is a smoothed long-run approach. In contrast, there would be likely to be costs associated with a CAPM-WACC approach, in terms of raising beta, and weakening the stability of the regulatory regime which could both increase risk and undermine incentives for future performance. In summary, the case for a switch to the CAPM-WACC has not been made. Consequently, Burns recommends allowing embedded debt and allowing the default premium on debt.

MPW also argue in Section 8 that there should be an explicit mechanism that links the RER and the RAR. Burns argues in Section 9 that such a proposal is likely to be detrimental to customer interests in undermining efficiency incentives and/or increasing risk. Concerns about excess returns appear to stem from the unusually high outperformance of cost and output targets, particularly amongst some energy networks following the first round of RIIO price control reviews. There are many mechanisms to ensure that profits lie within a socially acceptable range, but some of these mechanisms would be extremely detrimental to customers, whilst others would be much less so. In Burns’ view, regulatory action on outperformance should apply to the cost and output targets not to the RAR – the RAR should be focussed on the cost of equity and debt and minimising regulatory risk implies that this should be clear and transparent. An arbitrary adjustment factor applied to the RAR would only add to regulatory discretion and risk and undermine efficiency incentives.

2 Is CAPM the best available model for estimating the cost of equity for regulated companies?

2.1 Overview

In this section, we revisit the issue of whether the Capital Asset Pricing Model (CAPM) remains the most appropriate model for use in assessing the cost of capital for regulated companies. We focus in

particular on the issues of a) whether CAPM is sufficiently rich to capture key features of observed investor behaviour; and b) the impact on the finance industry of asset pricing models developed since CAPM (and typically arising from academic evidence of CAPM “anomalies”).

Our overall conclusions are:

- CAPM is still widely used particularly as the still-dominant model in calculating an assumed cost of capital in capital budgeting by chief financial officers;
- Investors’ decision-making processes are often informed by models that are more sophisticated than CAPM, both in terms of the assumed statistical properties of returns (including allowing for the impact of “factor models” and in their flexibility in dealing with investor preferences (for example in relation to loss aversion).
- But we also point to strong evidence that investors still *behave* as if CAPM is their benchmark model.
- We conclude that while these developments *have* changed financial markets nontrivially, there is no obvious respect in which these additional subtleties impinge on the estimated cost of capital for regulated companies in particular.
- Thus we advocate continuing to use the CAPM as the basis for estimating the WACC; although we do recommend that regulators should be aware of the possible implications of multifactor models for beta estimation.

These issues are dealt with in more depth in Appendices A and B.

2.2 A short history of the CAPM and subsequent developments

The Capital Asset Pricing Model (CAPM) developed in the 1960s remains the most important analytical building block in both portfolio theory and corporate finance.⁷ The key features of the CAPM are that a) only systematic or undiversifiable risk is remunerated by market risk premia, and b) that systematic risk can be captured simply by the excess return on the “market portfolio”. As a result, the expected return on any security or portfolio of securities over and above the risk-free rate is simply the product of how much systematic risk, or beta, it has and the expected market risk premium (the difference between the expected return on the market portfolio and the risk-free rate).

This relationship is given by the CAPM Security Market Line (SML):

$$E(R_i) = R_F + \beta_i (E(R_M) - R_F) \quad (2)$$

where $\beta_i = \frac{\text{cov}(R_{it}, R_{mt})}{\text{var}(R_{mt})}$ captures the sensitivity of asset i to the market portfolio.

The SML can be expressed more compactly, in terms of risk premia, as

$$RP_i = \beta_i RP_M \quad (3)$$

⁷ One indicator of its central role is that, along with Ross’s Arbitrage Pricing Theory (1975), the CAPM is one of two main equilibrium models for estimating expected returns on securities included in the Chartered Financial Analyst’s (CFA) syllabus. Further references are given in Appendices A and B.

where $RP_i = E(R_i) - R_F$ is the risk premium on security i , and $RP_M = E(R_M) - R_F$ is the market (or equity) risk premium (often denoted ERP). Expressed in this form, the SML makes it clear that CAPM implies that β_i is the *only* source of a risk premium on asset i .

This is a very strong result, and its strength has helped to underpin the influence of the CAPM. One way of seeing this is in relation to the following statistical representation of returns in terms of a regression equation:

$$R_{it} - R_{Ft} = \alpha_i + \beta_i (R_{Mt} - R_{Ft}) + \varepsilon_{it} \quad (4)$$

This is in principle a very general representation. The excess return on asset i in period t , relative to the risk-free return can *always* be decomposed into three components:

- The “alpha” of the security, α_i represents any non-zero mean element in returns that is unrelated to the market;
- The second term $\beta_i (R_{Mt} - R_{Ft})$ captures the systematic component in asset i 's return;
- The residual component ε_{it} captures any zero-mean random component in returns that is *not* correlated with the market (frequently interpreted as the “idiosyncratic” component)

A comparison of (4) with (3) reveals one crucial characteristic of the CAPM: if we take expectations of (4) it can only be consistent with (3) holding if $\alpha_i=0$: the “no alpha” result. The CAPM implies that if a security *does* have positive alpha ($\alpha_i > 0$) this should be a source of arbitrage profits, and would therefore be eliminated.

A second key feature of CAPM can be derived by a simple variance decomposition of equation (4):

$$\text{var}(R_{it}) = \beta_i^2 \text{var}(R_{Mt}) + \text{var}(\varepsilon_{it}) \quad (5)$$

So return variance for any security is made up of two components: the systematic component $\beta_i^2 \text{var}(R_{Mt})$ and the idiosyncratic component $\text{var}(\varepsilon_{it})$. But since idiosyncratic risk $\text{var}(\varepsilon_{it})$ does not enter the security market line, (3), the CAPM implies that it is not priced: only systematic risk, from beta, affects the risk premium.

Both of these features have very strong empirical implications, and thus have been the source of an extensive literature (see Appendix A) that has examined CAPM “anomalies”, most of which reduce to the claim that some types of assets (or more frequently, portfolios of assets) *do* have alpha. But a more indirect testimony to the power of CAPM is that the term alpha is extensively used (and pursued) in the finance industry as a metric of success by investors (for evidence on this point see Appendix B).

In an early follow-up to the original CAPM, Roll (1977) pointed out that the market portfolio is difficult to identify but the equation will still hold if it the market portfolio is substituted with a mean-variance efficient portfolio. It is common in most applications of the CAPM for the market return to be substituted with an equity index, with typically a local index being substituted in for both the expected market return and the beta. Provided this proxy portfolio is mean-variance efficient (it lies on the Markowitzian efficient frontier) then the expected return on the security will still be correctly estimated.

Early tests of the CAPM involved two-stage regressions. Firstly, the beta on a stock or portfolio of stocks was identified. Then in stage two, subsequent returns were regressed on these estimated betas,

with the assumption that differences in realised returns should reflect differences in expected returns and hence risk premia. Such approaches suffer from an “error in variables” problem: the explanatory variables are not true betas but estimated ones. Idiosyncratic risk in equities means that beta will be measured with error, and therefore the coefficients in the second stage will be biased. The standard approach in the literature to try to minimise the error in variables problem has been to estimate the beta on portfolios of stocks rather than individual equities, following on from the approaches of Blume (1970) and Fama and MacBeth (1973). We discuss the challenges in beta estimation in more detail in Section 4.5 in the main body of the report.

The possibility that not all investment opportunities might be able to be identified, and therefore that it may be impossible to know if a portfolio is mean-variance efficient, led to alternative formulations of, and extensions to, the CAPM, which rely only on the absence of arbitrage opportunities and the development of the Arbitrage Portfolio Theory of Ross (1977). This opened up the possibility of there being additional systematic risk factors. Crucially such factors would appear to be a potential source of “alpha” in the context of CAPM, but in the APT these simply represent additional sources of non-diversifiable risk, hence for which even well-diversified investors demand a risk premium.

Since Fama and French’s (1992) three factor model was published, this has become the most referenced extension of the CAPM in the academic literature, and it has also had a significant impact on the behaviour of professional investors. The first factor in the Fama-French model is the market factor, as in the basic CAPM. Two additional factors are added that are based on the fundamental characteristics of the stock: value measured by price to book, and size are added. We discuss this, along with other factor models developed more recently, in detail in Appendix B. In the next section we explore whether these factors impact companies’ cost of equity.

2.3 Does CAPM capture the key determinants of the cost of equity for regulated companies?

The capital budgeting process requires firms to estimate their cost of capital. Since Graham and Harvey’s (2002) survey of US Chief Financial Officers (CFOs), it has been well known that the majority of large businesses use the CAPM model when assessing the cost of capital. Although less extensively used, the CAPM is also the dominant model for European CFOs (see Bancel and Mittoo (2004)). A weakness of survey data is that results may be biased by non-responses and what management claim they do might not match reality. In the original paper 73.5% of CFOs claimed to use the CAPM always or almost always, and 34.3% to use CAPM with extra risk factors, 39.4% used historic returns, 15.7% backing out from a dividend discount model, and 13.9% what their investors tell them and 7% what their regulators tell-them (NB: Total responses sum to more than 100% as CFOs use multiple methods). How CFOs move from multiple approaches to a final number likely involves a high degree of discretion. Intriguingly, where CFOs do adjust the cost of equity on the basis of multi-factor models it is usually on the basis of “macro factors” (see Appendix B). Macro factors make up the first four common additional factor after the market factor, with size, the least robust of the Fama-French factors, being the only fundamental factor entering the top eight.

The “dual” of the question of which is the right model for estimating the cost of equity, is to ask: what is the model that investors actually appear to use to assess the required return on equities. Few investors will actually say that they use the CAPM as their model for estimating asset returns. Furthermore there is empirical evidence both in terms of the behaviour of individual stock behaviour and the overall asset allocations of individual investors which is inconsistent with the CAPM. However for the purposes of regulators, the key issue is whether these apparent differences actually invalidate the use of CAPM as

a model of investor's desired returns, in ways that are relevant to the cost of capital of regulated companies.

In Appendix A we assess whether investor investment processes are consistent with the CAPM. Investor processes are typically multi-stage. At the asset allocation stage (how much investors will hold of broad asset classes, such as domestic equities, international equities, government bonds, corporate credit, and so on) the methods used are in principle more nuanced than the CAPM. In particular models allow for greater loss aversion and more complex distributions of returns.

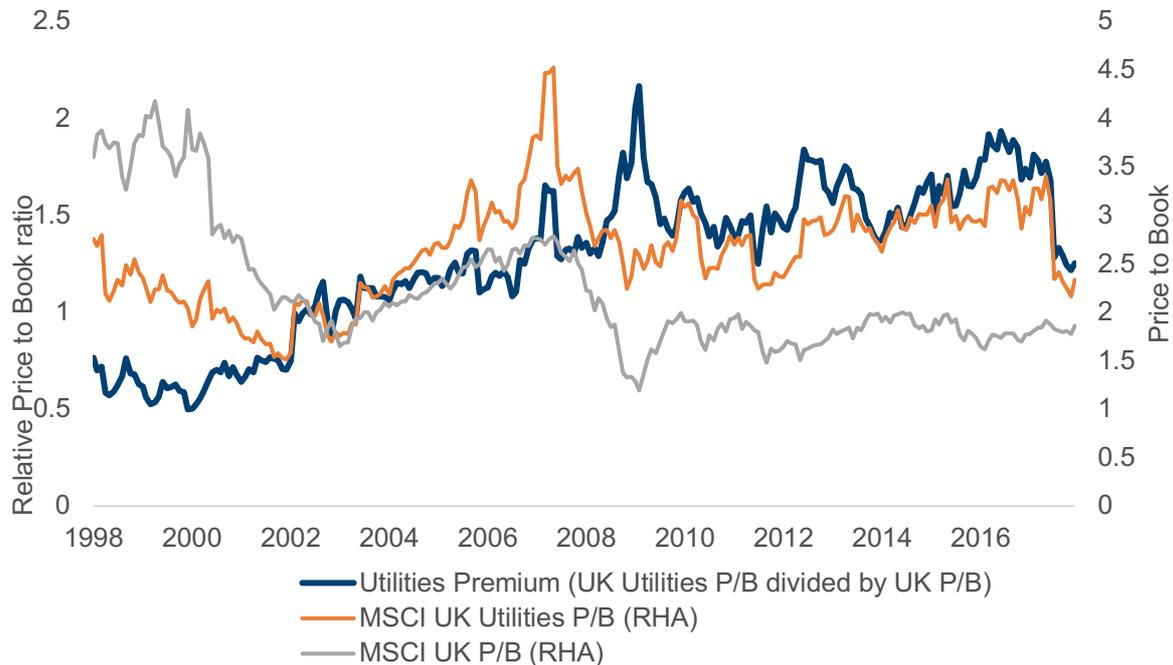
However, once investors have decided how to allocate a given amount to equities the detailed decisions on allocations to individual equities or types of equities is usually subcontracted to specialist equity managers. When we look at aggregate investor behaviour in terms of how funds are allocated to equity managers, the CAPM appears to reflect aggregate investor behaviour well. More specifically, ex-post returns on an equity portfolio can be decomposed into how much was driven by the broad market factor and by other factors. Survey results suggests that investors are indifferent between whether additional return relative to market risk is added from "factor tilts" (increasing exposure to factors such as stocks with positive price momentum, low price to book ratios, and "small cap" stocks) or alpha after all risk factors are allowed for. Thus the evidence suggests that investors pursue "CAPM alpha", consistent with the underpinnings of CAPM. We therefore conclude that whilst factors might influence investor expectations they do not change the cost of equity.

In Appendix B we look at factor models of equity returns in more detail. We argue that whilst factors are relevant for forecasting returns on portfolios of equities they are less relevant for forecasting returns on individual equities, as for specific equities factor exposures are often transient. Our key conclusions are:

- For individual stocks the impact of a given factor is often transient. While factors may impact on the expected returns on portfolios of stocks, the same is not true for individual stocks, with factor loadings actually changing sign quite frequently.
- Fama and French identify two other factors other than the market factor: value (based on priced to book ratios) and size. Size effects are now widely disputed. Value effects are possibly robust although the long-period of underperformance from late 2009 to mid-2016 has cast doubt on this too.
- The value effect is also a clear example of how factor loadings can change sign. Although pre-2002 UK utilities have in the past traded at lower price to book multiples than the broader market this is no longer the case, as Figure 2.1 below shows. Current price to book valuations would not therefore justify any upward modification in the allowed cost of equity.

Figure 2.1

UK Utilities' Price to Book Ratio Premium over UK equities



Source: MSCI, Thomson Reuters

Additionally, companies often have significant discretion in determining their book values. If regulators were to start to adjust for the value factor this could encourage listed companies to inflate book values in order to lower price to book multiples and raise allowed returns on equity.

Although use of CAPM is therefore far less ubiquitous amongst investors than it is amongst CFOs we believe it still approximates investor behaviour well, at least when it comes to equity selection, which is what impinges directly on the cost of equity for individual regulated companies.

2.4 The Low Beta puzzle.

There is however one aspect of factor models that does appear to have more relevance to regulated companies: namely, the evidence on the performance of low beta stocks, which do appear to represent something of an ongoing puzzle.

Multi-factor extensions to the CAPM are usually consistent with the Arbitrage Portfolio Theory of Ross (1977). No arbitrage requires that each unit of exposure to a particular factor would therefore earn a constant amount of premium. This approach excludes the possibility of low beta stocks earning an excess return, since the reward for market risk should be the same for each unit of risk, and not decrease for higher beta stocks. Thus most academic models exclude “low beta” as a risk factor, despite some evidence that low beta stocks have tended to earn “alpha”, and thus represent a CAPM anomaly. In contrast, as outlined in Appendix B, commercial factor funds quite frequently focus on it as a source of outperformance.

We discuss reasons for the apparent outperformance of low beta in Appendix B and suggest it could be a function of contracting difficulties in the market for fund management services; although it is not clear that these difficulties have continued into the recent past.

The basic rationale underlying the CAPM result that low beta stocks should *not* earn alpha is that, for example, a well-diversified portfolio of stocks with a beta of one half can in principle be geared up by borrowing, so that a portfolio with 50% gearing (and by assumption zero debt beta) would have a beta of one. Thus any alpha on low beta stocks should in principle be eliminated by arbitrage.

However, this feature relies crucially on the assumption that those carrying out the arbitrage can fund their activities by risk-free debt. In practice leverage may be restricted or cost sufficiently more than the risk-free rate to make the exercise not worthwhile. Traditionally leverage has been limited to hedge funds (offshore investment vehicles which are subject to less stringent regulation). The historic limitations on leverage may help to explain the observed excess returns on low beta stocks.

Since 2011, the development of sophisticated UCITS⁸ under the UCITS IV framework has opened up leverage to onshore funds. This has also occurred at a time where there has been substantial fund flow into “smart beta” funds, quasi-passive funds which overweight exposure towards stocks with positive risk factors. According to Morningstar (the fund data provider), assets in “smart beta” exceeded \$1 trillion by the end of 2017 (having been below \$300 billion at the end of 2012).

Additionally, we would argue that a quite strong case can be made that the *potential* for arbitrage profits arising from low beta assets has been one of the key explanations of why low beta assets have become attractive targets for private equity investors, who then lever up balance sheets, and thus achieve the arbitrage CAPM would predict.

Again we explore this in more detail in Appendix B, but overall we conclude that, while the historic excess returns on low beta stocks do represent a challenge to CAPM, and one that is particularly relevant to regulated companies, we believe that the impediments to arbitrage via leverage have now been sufficiently eliminated that there would be no basis for regulators to assume that low beta stocks will continue to earn excess returns (and thus face a higher cost of capital) in the future.⁹

Recommendation 1 (CAPM): Regulators should continue to use CAPM, but should be aware of the implications of more recent multifactor models for beta estimation in particular.

3 A simple benchmark CAPM-WACC framework

3.1 The vanilla WACC and the asset beta

To simplify, we focus for now on the standard definition of the vanilla WACC:¹⁰

$$WACC = gE(R_D) + (1-g)E(R_E) \quad (6)$$

⁸ The Undertakings for Collective Investment in Transferable Securities Directive 2009/65/EC is a consolidated EU Directive, that allows collective investment schemes to operate freely throughout the EU on the basis of a single authorisation from one member state.

⁹ We note also that in the beta estimation exercises carried out on regulated stocks in Appendices F and G, estimated alpha values have been insignificantly different from zero in samples covering at least the last decade.

¹⁰ Or equivalently, we can assume that the differential impact of tax shelters has been allowed for in the definitions of the returns on debt and equity.

where g is leverage (gearing), R_D is the return on debt and R_E is the return on equity, and we explicitly note here that the WACC is defined as an *expected* return.¹¹ The definition of the asset beta is as follows:

$$\beta_A = (1 - g)\beta_E + g\beta_D \quad (7)$$

Note that in general the asset beta will be a function of the leverage of the company g . We typically assume¹² that β_A is a U-shaped function of g , reaching a minimum at the optimal level of gearing, g^* . The CAPM in turn implies that the same property must hold for the WACC.

3.2 The CAPM-WACC vs the CAPM(E)-WACC

If we apply the CAPM security market line *just* to the expected return on equity, then we can write the vanilla WACC as:

$$\text{CAPM(E)-WACC} = gE(R_D) + (1 - g)[R_F + \beta_E RP_M] \quad (8)$$

where our rather cumbersome notation is chosen deliberately to reflect the selective way in which the CAPM is being applied here. This is the formula as commonly applied by regulators in the UK.

Manipulation of this formula yields an important insight. The formula for the asset beta implies

$$\beta_E = \frac{1}{1 - g}(\beta_A - g\beta_D) \quad (9)$$

and, substituting into the definition of the CAPM(E)-WACC, we then have

$$\begin{aligned} \text{CAPM(E)-WACC} &= gE(R_D) + (1 - g) \left[R_F + \left(\frac{1}{1 - g}(\beta_A - g\beta_D) \right) RP_M \right] \quad (10) \\ &= R_F + g(E(R_D) - R_F) + (\beta_A - g\beta_D) RP_M \\ &= \{R_F + \beta_A RP_M\} + g[RP_D - \beta_D RP_M] \end{aligned}$$

Note that all that we have done here is to manipulate an identity, assuming the CAPM applies only to the cost of equity. But in doing so we gain important insights:

- The first term, in {}, equals the hypothetical WACC of an equity-only (unlevered) firm with systematic risk given by β_A ,¹³
- And, crucially, the second term, in [], simply measures the deviation of the risk premium on debt from the value implied by the CAPM.

If therefore the CAPM is assumed to apply to both the cost of equity *and* the cost of debt, then the term in square brackets is zero, and we have a consistent definition of the “pure” CAPM-WACC given by

¹¹ The crucial distinction between actual and expected returns has been quite frequently ignored in contributions to the discussions of the WACC in the context of UK regulation.

¹² See for example, Korteweg (2010).

¹³ Note that this hypothetical construct does *not* represent the WACC of an actual equity-only company engaged in the same operations, since if a company chooses its leverage optimally its actual asset beta will be lower than if it were an equity-only firm. For a more detailed discussion, see Appendix C.

$$\text{CAPM-WACC} = R_F + \beta_A RP_M \quad (11)$$

We have argued above that, for all its faults, the CAPM is the best available asset pricing model. Regulators have in the past at least paid lip service to the concept, but have typically paid much more attention to its implication for the cost of equity than the cost of debt. They have thus typically calculated the CAPM(E)-WACC rather than the pure CAPM-WACC.

As noted above, this approach can at least partly be defended by the fact that is not just regulators who do not apply the CAPM to the cost of debt; academic research has also followed this tendency. Nonetheless we would argue that the pure CAPM-WACC, with this very simple form, is something that regulators should calculate, at least for reference purposes, not least because it is a very simple and transparent calculation.

Note also an important contrast between the pure CAPM-WACC and the CAPM(E)-WACC, that can be demonstrated by the comparison between Equations (10) and (11). To the extent that the risk premium on debt RP_D is greater (or less) than the value implied by the CAPM (i.e., if the square bracketed term in Equation (10) is positive (or negative)), an increase in leverage, g will increase (or decrease) the estimated CAPM(E)-WACC, whereas the pure CAPM-WACC in (9) will be entirely unaffected, consistent with standard finance theory.¹⁴

3.3 The impact of leverage on equity beta, and the role of the risk-free rate.

An advantage of the pure CAPM-WACC definition given in Equation (11) above is that it brings out the crucial role of the risk-free rate, R_F . In contrast, this role can quite easily be obscured in the more commonly used CAPM(E)-WACC.

This is particularly so when equity beta, β_E is assumed to be close to 1, as has been the case in some recent regulatory decisions. Indeed, in the special case that β_E is *precisely* equal to unity, as noted by MMW, the CAPM security line simply reduces to $E(R_i) = E(R_M)$: thus, the risk-free rate disappears from the cost of equity, and hence from the conventionally calculated CAPM(E)-WACC.¹⁵

Furthermore, for any given level of asset beta, β_A , there is always some value of leverage such that β_E is *precisely* equal to 1. By manipulation of the formula for equity beta we have

$$g = \frac{1 - \beta_A}{1 - \beta_D} \Rightarrow \beta_E = 1 \quad (12)$$

But note that this critical value of leverage¹⁶ is inversely related to asset beta. Thus, the lower the systematic risk of the company's operations, as captured by β_A , the higher is the required leverage at which equity beta is equal to unity, and hence the higher is the implied weight on the cost of debt in the

¹⁴ See for example the textbook treatment in Brearley and Myers (2016).

¹⁵ As was indeed the assumption made by Ofgem until 2007.

¹⁶ Note that strictly speaking we should write this expression as

$$g = \frac{1 - \beta_A(g)}{1 - \beta_D(g)} \Rightarrow \beta_E = 1$$

where here we explicitly acknowledge the dependence of both betas on leverage so the particular value of g that satisfies $\beta_E = 1$ is only defined implicitly.

CAPM(E)-WACC calculation. In the pure CAPM-WACC the risk-free rate would still enter the formula directly; but this is obscured in the CAPM(E)-WACC calculation, where it only enters indirectly via its influence on the cost of debt (see discussion in Section 4.7.2).¹⁷

We revert to this special case further in later sections.

3.4 Insights from multifactor models

Multifactor models imply that the random variation over time in *actual* returns can be written as

$$R_{it} - R_{ft} = \alpha_i + \beta_1 F_{1t} + \beta_2 F_{2t} + \dots + \beta_k F_{kt} + \varepsilon_{it} \quad (13)$$

where the k factors F_{it} , $i=1..k$ capture mutual correlation between returns on different assets, such that any residual component in ε_{it} is purely idiosyncratic. Conventionally the first factor, F_{1t} is usually assumed to be the market return, i.e., $F_{1t} = R_{Mt} - R_{ft}$.

The standard argument then follows that the generalised asset pricing relation becomes

$$E(R_i) = R_f + \beta_M RP_M + \beta_2 RP_2 + \dots + \beta_k RP_k \quad (14)$$

Where the “zero alpha” result assumes that if α_i is *not* equal to zero, there will be potential arbitrage profits.

By comparison with the standard CAPM as set out above, it is evident that in principle the CAPM subsumes any multifactor model in which non-market factor risk premia, i.e., $RP_2 \dots RP_k$ are zero.

Thus while the process for returns on individual securities or portfolios may in principle be much better *described* by a multifactor model than by the CAPM (indeed there is plenty of evidence that this is the case), this only matters for the cost of capital if there is very strong evidence that the additional factors are *priced*.

Clear examples of such additional zero-risk-premium factors in Fama-French are the impact of bond market factors in determining stock returns. In Appendix B, Derry Pickford discusses the evidence that the evidence of sustained positive excess returns on other factors, such as the “small firm” effect appears also to be quite weak, and argues that, even where excess returns do appear to have been sustained, this is more likely to be ascribed to behavioural factors than to genuine risk premia.

But in principle multifactor models may still impact on the way CAPM betas are *estimated*. If additional factors have explanatory power for the cross-section of returns then omitting these factors from the standard CAPM regressions (as carried out routinely by Bloomberg, for example, and as reported by Burns in Appendix F) may result in missing variables bias. Donald Robertson and Stephen Wright carry out some preliminary investigations of this issue in Appendix G.

3.5 CAPM and the Dividend Discount Model

It is worth stressing that the CAPM and the Dividend Discount Model (DDM) are in principle perfectly mutually consistent. Indeed the CAPM does not actually live up to its description as a *pricing* model

¹⁷ We note this feature purely as a mathematical implication, rather than as endorsement of the assumption that β_E is precisely equal to 1.

unless turned into an expression with price on the left-hand side. Solving this model leads to some form of Discounted Cash-Flow (DCF) model, with the DDM being one particular type of DCF which can be derived.

However the linkage typically does involve restrictive assumptions that are by no means always made explicit. Thus, from the definition of returns, the CAPM implies (inserting time subscripts for precision)

$$P_{it} = \frac{E_t(P_{it+1} + D_{it+1})}{1 + RFR_t + \beta_i RP_{M,t}} \quad (15)$$

where D_{it} is dividends per share for company i . The numerator represents the expected total cash payoff in period $t+1$, while the denominator is simply $1+$ the CAPM cost of equity for firm i .

In many applications, for example in a recent Bank of England application of the DDM¹⁸ this expression is iterated forward, substituting for future prices, to derive price as the discounted sum of future dividends. But this process is not actually straightforward.

In order to make the calculation viable as a standard DDM, the discount factor must be known at time t , and hence non-stochastic.¹⁹ The textbook approach typically skates around this problem by assuming that everything in the denominator is constant, in which case the expression becomes²⁰

$$P_{it} = \sum_{j=1}^{\infty} \frac{E_t D_{it+j}}{(1 + RFR + \beta_i RP_M)^j} \quad (16)$$

While this is a useful expository device to illustrate the importance of systematic risk in increasing discount rates, it can never precisely describe the true fundamentals-based price, since in practice the term in brackets in the denominator is *not* constant: so this approach can (at best) be used as an approximation.

Much of the academic literature that applies the DDM uses a closely related approach, by calculating the “internal cost of capital” for firm i at time t , ICC_{it} , such that for a given path of expected dividends per share, the following expression holds *exactly*:²¹

$$P_{it} = \sum_{j=1}^{\infty} \frac{E_t D_{it+j}}{(1 + ICC_{it})^j} \quad (17)$$

But note that, in contrast to the CAPM cost of equity, the implied value of ICC_{it} does *not* represent the expected return on the firm’s stock over any particular horizon: it is simply a summary statistic of expected returns over all possible horizons, *conditional* upon an assumed path for expected dividends.²²

¹⁸ See for example, Dison and Rattan (2017).

¹⁹ More modern approaches that assume a stochastic discount factor (e.g., Cochrane (2005); Campbell (1991)) can also produce a similar approach by some combination of approximations and distributional assumptions.

²⁰ Assuming that the impact of future discounted expected prices decays to zero.

²¹ See for examples see: Gebhardt, Lee and Swaminathan (2001); Claus and Thomas (2001); Li, Ng and Swaminathan (2013).

²² The Bank of England’s approach (see Dison and Rattan, *op cit.*) is closely related to the ICC approach above, but instead of solving for the internal cost of capital, they derive a DDM in terms of a *set* of known risk-free rates, along the yield curve. They then assume that the market risk premium at time t is assumed constant along the yield curve (so the implicit expected stock return over different horizons follows the same profile). This makes it hard

There is a clear analogy with the yield to maturity of a coupon bond; but with the caveat that, in contrast to the coupons on a bond, the path of dividends per share the stock will generate is not known, but needs to be estimated.

In principle the estimated internal cost of capital for any given firm provides an alternative estimate of the cost of capital, *independent* of the CAPM. This approach has indeed been widely used as an alternative to the CAPM by regulators in the United States: see Sudarsanam, Kalttenbronn and Park (2011). But it is heavily reliant on, and sensitive to, the assumed path of future dividends for the individual stock.

To illustrate, take the simplest possible application of the DDM, the Gordon Growth Model (Gordon, 1962), in which, at time t , firm i 's dividend per share is expected to grow at a constant rate, G_{it} , in perpetuity.²³ Combining this assumption with the internal cost of capital expression above gives

$$P_{it} = \frac{(1 + G_{it})D_{it}}{ICC_{it} - G_{it}} \quad (18)$$

where note that everything on the right-hand side is t -dated, and both P_{it} and D_{it} are assumed known (thus the approach clearly only works for quoted stocks). For any given assumed value of G_{it} , therefore, there must be an implied value of ICC_{it} , the internal cost of capital, given by

$$ICC_{it} = (1 + G_{it}) \frac{D_{it}}{P_{it}} + G_{it} \quad (19)$$

But note that, transparently, this is actually a necessary relationship between two magnitudes, *both* of which are unknown. Thus, there is a clear element of conditionality in any application of the DDM: the implied cost of capital is *always* conditional on an assumption about growth of dividends per share. Any error in estimating G_{it} translates more or less one-for-one into an equivalent error in estimating ICC_{it} (itself, as noted above, only a summary statistic for expected returns over all horizons²⁴). Furthermore, the implicit assumption, in making such conditional forecasts, is that dividend growth forecasts are more reliable than return forecasts.²⁵

This is not to argue that the DDM approach is valueless; indeed, the necessary relationship between expected growth and the internal cost of capital provides a helpful cross-check on *any* approach to estimating the cost of capital. We revert to this issue further below, in Section 4.4.

to derive an exact estimate either of the internal cost of capital, or of the EMR from their results, so in our discussion below we need to rely on approximations.

²³ By assumption this assumed growth rate may change over time, hence the time subscript. More recent applications as in, e.g., Damodaran (2016), modify this assumption by splitting the expected growth path into two or more segments, using, e.g., broker forecasts of near-term growth, supplemented by longer-term constant growth assumptions, but the essential insights remain. Indeed for a given internal cost of capital, any non-constant growth path always has an equivalent constant growth path.

²⁴ Note that a version of Equation (19) is quite frequently written (as in the original exposition in Gordon (1962) with the expected return on the left-hand side. However, this expression is only valid under the strong restriction that the dividend yield is constant. Thus, any such representation is at best valid as a long-run description, under the assumption, for example, that both the expected return and the dividend yield are mean-reverting (the rationale applied, for example, Fama and French (2002).

²⁵ There is indeed a quite extensive academic literature suggesting that the reverse is the case: see for example, Cochrane (*op cit*) and Cochrane (2008). Note that, in contrast to the application of the DDM cited, above the Bank of England has also exploited the predictability of returns to infer growth forecasts, see Chin and Polk (2015).

4 How should regulators estimate the key components of a market-consistent CAPM-WACC?

At this stage, we focus purely on how regulators should estimate these key components, *not* on the way they are implemented in regulation: i.e., we distinguish between the problem of estimating the CAPM-WACC, as a market concept, and the impact these estimates should have on the Regulatory Allowed Return (RAR) and the Regulatory Expected Return (RER), to which we revert in Section 5.

The elements we need to consider, to ensure a fully consistent treatment, are:

- Investment Horizon
- The appropriate price index for constructing real returns
- The risk-free rate (RFR)
- The expected market return (EMR) and the equity risk premium (ERP)
- Asset beta and equity beta
- Gearing (leverage)
- The cost of debt

4.1 Investment Horizon

The CAPM is usually presented, as in Section 2.1 above, as a two-period model, but the length of this period is not clearly specified. Some academic representations assume an instantaneous return (if in continuous time); while much empirical estimation is at a monthly frequency. As a result, CAPM is often viewed as a model of expected short-horizon returns and risk premia (especially given the standard definition of the risk-free rate, which historically, at least, did not really exist at longer horizons until the advent of indexed bonds).

However, at any given maturity over a wide range, given the existence of indexed bonds in modern financial markets, there does always exist a rate that is risk-free *at that maturity*²⁶ (but *not* at shorter maturities) so there is in principle no obstacle to applying CAPM at longer maturities.

What is the appropriate investment horizon to choose in estimating a market-consistent CAPM-WACC?

- **Instantaneous (or very short-term, e.g., monthly) horizon?** This approach is consistent with the great majority of empirical implementations of asset pricing models in the academic literature. It is also consistent in principle with the behaviour of many financial market practitioners who are constrained professionally to have relatively short investment horizons (although there are exceptions—most notably private equity, which typically has at least somewhat longer horizons, and has come to play a major role in infrastructure finance). The counter-argument to this approach that is usually advanced is that it would introduce a clear disconnect between the horizons of the (notional) investor and the expected life of the assets employed.

²⁶Strictly speaking this statement is only true of zero coupon indexed gilts. No such bonds exist in practice (though they can in principle be constructed using swaps), but their implied prices and hence yields can be calculated from yields on coupon bonds: the Bank of England publishes these on a daily basis. We do not at present regard default risk on UK sovereign debt as worth worrying about, despite the occasional flurries of interest in this topic in the financial press.

- **Investment horizon= end of price control period?** This has clear attractions, at least in terms of debt finance, since in principle if this were standard practice companies could (if they wished) align their debt issuance to the same maturity date. However, a counter-argument is that the assets in which most regulated firms invest are distinctly longer-lived, as are, typically their liabilities (though this latter choice may well be endogenously determined).
- **Investment horizon \approx debt liabilities?** This has typically been the choice of regulators and is usually defended in terms of ensuring a reasonable stable path for the cost of capital, since long maturity real rates are typically more stable than short rates.²⁷ It also gets closer to matching the expected life of companies' physical assets.²⁸ An additional justification is that the underlying investors whose assets are being managed by practitioners arguably have similar investment horizons, since they are typically investing, through pension funds, over their life cycle, and hence might be expected to have an investment horizon of roughly half a working life, hence around 20 years. A possible counter-argument to the option is a concern about the degree of uncertainty in longer-term forecasts, if the average maturity of debt liabilities is very long-dated (e.g., more like 30-40 years, rather than 20 years).

While we have discussed the issue of the optimal horizon at some length, we have not yet come to a consensus. In part this may reflect the fact that there can be no clear-cut answer. In the benchmark case of informationally efficient financial markets the choice of horizon would be immaterial, since long-term interest rates would simply match expectations of future short rates.

Hence while a case can be made for both choices A and B, we can see no obvious *objection* at least to choice C, which typically appears to be the regulators' preferred choice. Thus it can be at least defended, if not on any particular grounds of principle, at least in terms of continuity of practice.

However, we *are* in agreement on a key caveat: that, whichever horizon is chosen, the components of the cost of capital should, as far as practically possible, be estimated in a way that is *consistent* with the chosen horizon, since without this consistency we cannot view our CAPM-WACC estimate as a true expected return. We shall argue that this has not always been the case for the choices made by UK regulators. We revert to this topic at various points below, in considering the individual components of the WACC.

Recommendation 2 (Horizon): On balance, we are in favour of choosing a fairly long horizon, for example, 10 years, in estimating the CAPM-WACC. But we would argue that, more important than the choice of horizon per se is that all components of the CAPM-WACC are estimated using a methodology that is *consistent* with the chosen horizon.

To illustrate the importance of consistency, consider the case that regulators choose a 10-year horizon. Then the appropriate estimate of the risk-free rate and cost of debt should be derived from yields on bonds with maturity of ten years (see Recommendations 4, 8 and 9), the estimate of the Expected Market Return should be based on the properties of an appropriate return at a 10-year horizon (see Recommendation 5), and the estimates of equity and/or asset beta should represent the best econometric

²⁷ Although we would give this argument more weight if regulators actually used even long-term market interest rates in their WACC estimates – an issue we revert to below.

²⁸ Note however that while the replaceable components in a network may have finite lives, in some important respects the network operates as a single asset which is continuously maintained and enhanced, and thus is effectively infinitely lived. The life of replaceable components in the network may therefore not have the same role in the economics of network regulation as asset lives do in other sectors. (We are grateful to Ian Rowson, CAA for pointing this out).

estimates that best capture the relationship between individual stock returns and the market return over a ten year horizon (see Recommendation 6).

4.2 The appropriate price index for constructing real returns.

Regulators have historically chosen (we believe correctly) to focus on real returns to capital, and to index firms' allowed expenditure to inflation (which may include allowance for 'Real Price Effects' for inputs that are subject to different inflation influences). A direct implication is that both historical and forward-looking analysis of the CAPM components should be carried out in real terms.

However, we would argue that the price index originally used as the basis for regulation, the RPI, has a number of defects that detract from its value as the basis for estimating the CAPM-WACC in real terms. Our key criticisms of the RPI (which Derry Pickford outlines at length in Appendix D) are:

- The RPI has become increasingly divorced from the choice of price index in monetary stabilisation, with the Bank of England focussing on the CPI. To the extent that inflationary expectations are anchored by the monetary regime, they are anchored to the CPI, not the RPI.
- The methodology for constructing the RPI is in a number of respects inconsistent with standard practice in constructing measures of consumer price inflation in most countries, whereas the CPI follows standard practice. The resulting differences between inflation as measured by the RPI and the CPI have been quite large, and quite volatile, and are expected to continue so in the future.
- The historical time series for the RPI has been subject to significant changes in its data construction over time (most notably in its treatment of housing costs) which make it hard to use the RPI to derive consistent historic estimates of real returns. In contrast, the ONS and the Bank of England have published consistent historical estimates of the CPI. We understand from discussions with the ONS that similar long-term series for CPIH is also under development.

As against these (to us) clear disadvantages, a clear advantage of the RPI is that (for largely historical reasons) inflation-indexed bonds issued by the UK government are still only indexed to RPI inflation. However, we argue that this situation is not likely to persist, and indeed Appendix D notes that there is a strong latent demand for CPI-indexed bonds that would be facilitated by a regulatory shift towards the CPI.

A number of regulators (e.g., Ofcom; and from PR19, Ofwat) have already signalled their intention to move away from RPI-X to alternative measures of consumer price inflation. We believe this choice to be correct; however,

Recommendation 3 (Price Index): There is a strong case for regulators choosing a measure of inflation for estimating the CAPM-WACC that is consistent with that chosen by HM Treasury and implemented by the Bank of England for inflation targeting. Ideally, we would like to use the same index used in setting the Regulatory Allowed Return (RAR). We can however see merit in regulators postponing any decision to change index until there is clarification on whether the measure of prices used for inflation targeting by the Bank of England will be changed.

Thus, while there are undoubtedly arguments in principle for a consumer price index that includes housing costs (the rationale for the CPIH), we would argue that, for as long as the Bank of England continues to use the CPI (which excludes housing costs) in its inflation targeting, regulators should act

consistently with the BoE: essentially sub-contracting the choice of appropriate price index to the institution charged with price stability.

Quite apart from arguments about the pros and cons of different price indices (some of which Derry Pickford addresses in Appendix D), one key advantage of this approach is that it helps both in looking forward (given the BoE's stabilising remit) *and* in looking backward over the history of returns. The Bank of England has invested considerable effort in constructing the "Millennium" dataset, which projects the real economy backwards using the best available historical sources. We believe that their back-projected measure of the CPI should be used in constructing and analysing historic real returns (raw data for which are all measured in nominal terms), and that this is distinctly superior to the RPI.

Appendix D analyses the properties of this historic series, and compares it with the historic properties of the RPI, as well as other commonly used price indices; it also includes some analysis of likely wedges between RPI and CPI inflation in the future.

We would argue that this "subcontracting" approach to the measurement of inflation and to the indexation rules in regulation meets the twin requirements of implementability and defensibility set out above.

A corollary of our recommendation is that if, at some stage, the Bank of England judges that monetary stabilisation should be conducted with respect to a *different* price index (for example, the CPIH), then regulators should (as soon as is feasible) switch to using the BoE's preferred new measure of inflation.²⁹

Burns agrees with the logic of the position outlined here but notes that regulators may balance other objectives in choosing the appropriate index to apply. For example, other indices may show a better correlation to input prices in that sector. In his view, it would be reasonable for a regulator to take these other factors into consideration.

4.3 The risk-free rate (RFR)

Here we have little to add to the conclusion of MMW (2003): at any given horizon, the risk-free rate is directly observable in nominal terms, and observable in real terms, from yields on indexed bonds, subject to the caveat that the indexation of such bonds is in terms of the RPI. Thus any reported real yield based on RPI indexation should be adjusted to be in terms of the Bank of England's chosen CPI series (and hence be consistent both with the historic evidence and, we would hope, with future indexation practices of the regulators).³⁰

Recommendation 4 (The Risk-Free Rate): Regulators should use the (zero coupon) yield on inflation-indexed gilts at their chosen horizon to derive an estimate of the risk-free rate at that horizon. For consistency with the treatment of inflation outlined above, the published yield needs to be adjusted for the predicted wedge between the inflation rate used for estimating the CAPM-

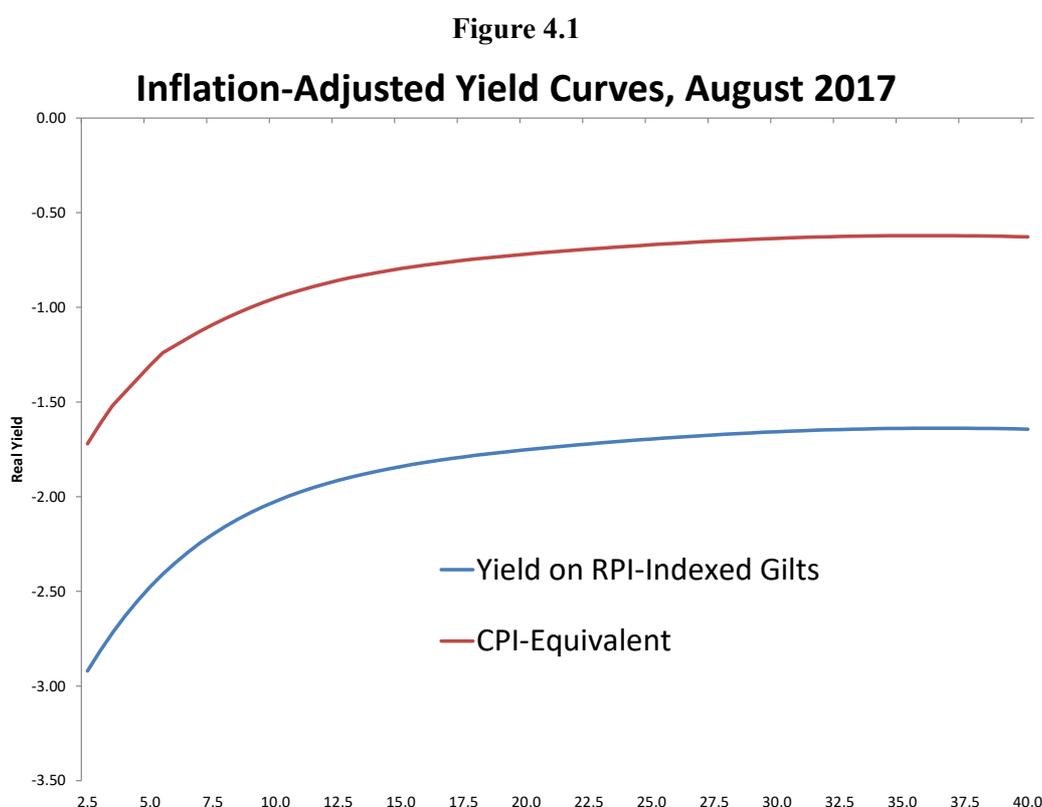
²⁹ In principle this could require some adjustment to estimates of historic returns; however we would not expect the BoE to switch without having first documented the history of the chosen new index; we would also expect their choice of target to adjust, if necessary, to any sustained differences between the new inflation measure and the measure they currently use, thus any switch by regulators would not (in expectation) imply any discontinuity of practice.

³⁰ A caveat here is that the so-called "breakeven inflation rate" calculated from the gap between the yield on nominal vs indexed gilts, is *not* a reliable indicator of market expectations of RPI inflation, or, by extension of CPI inflation, once corrected for the prospective gap between the two inflation rates. Since nominal and indexed gilts have different risk characteristics, this calculated rate is likely to be a biased measure of expected inflation. See Appendix D.

WACC (and ideally also setting the RAR) and that used for indexing gilts (currently RPI inflation).

Appendix D discusses alternative methodologies for the required adjustment from RPI inflation to CPI inflation. While most give very similar results, at present, on grounds of implementability and defensibility, we would advocate a further application of the “sub-contracting” approach outlined in relation to the choice of price index: namely using published OBR forecasts. In due course, however, we would hope that this adjustment will be rendered redundant by the emergence of liquid markets in CPI-indexed bonds.

Figure 4.1 shows the impact of correcting published indexed gilt yields for the predicted difference between RPI and CPI inflation, consistent with our Recommendation 3, using the methodology explained in detail in Appendix D.

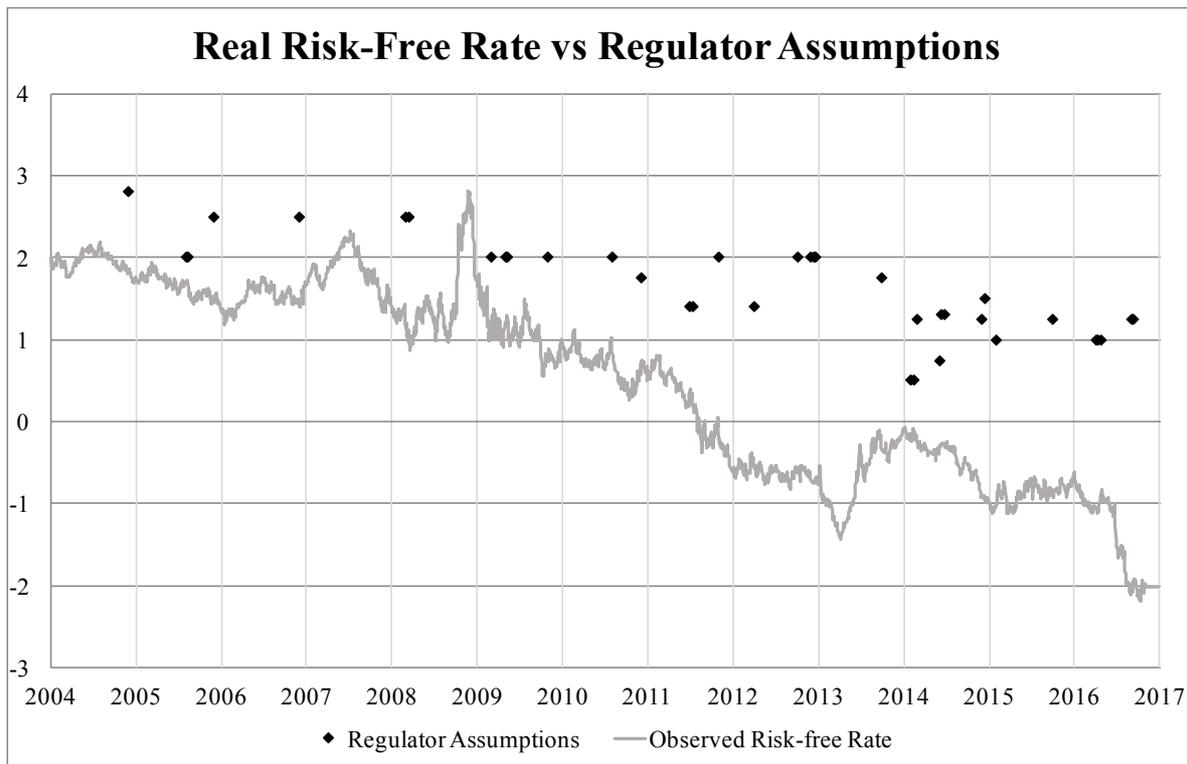


In making this recommendation we are very conscious that this approach to the risk-free rate is *not* the approach taken by UK regulators over the past decade or so. As Figure 4.2 shows, while regulators have gradually lowered their risk-free rate assumptions³¹ over this period of unprecedentedly low risk-free rates, these reductions have been at a distinctly slower rate than the fall in market rates, leaving an increasingly large gap between what markets say is the return on risk-free assets, and what regulators assume.³²

³¹ Note that the measure of the RFR chosen by regulators has varied considerably, so that the points shown in Figure 4.2 are not always strictly comparable.

³² Note that the real interest rate shown here is the implied yield on 10-year zero coupon indexed bonds, from the Bank of England), which should be clearly distinguished from the 10-year forward rate used by some regulators. Note also that this yield is *not* adjusted for the CPI-RPI wedge, since we do not currently have a consistent time

Figure 4.2



Different justifications for this approach have been advanced:

- **Gradualism: the “dragging anchor” approach**

The argument here is that regulators will adjust their assumptions in line with markets in due course, but, in the interest of regulatory stability, do not respond instantaneously to every movement in the market rate: the “dragging anchor” metaphor.

We discuss this argument further in later sections of the report but here we would simply note two key caveats.

First, if there *is* a case for gradualism, then MPW would argue that any dragging anchor should be applied to the Regulatory Expected Return (RER) as a whole rather than component by component, since if this approach is not applied consistently across components, not only might this reduce the impact of any smoothing, but may also induce distorted choices. (Burns dissents from this view, however.)

series for the predicted difference between the two; but Figure 4.1 shows that in the most recent data the yield needs to be adjusted upwards by almost exactly one percentage point. For historical comparisons it would be desirable, and relatively straightforward to produce an adjustment along the lines of that proposed in Appendix D by some combination of time series methods and published forecasts. The regulator assumptions are derived from a database of 42 regulatory decisions our research assistant Meline Danielewicz has constructed for the purposes of this report, and are all expressed in RPI terms. An alternative version of Figure 4.2, which shows regulatory decisions on the x-axis, instead of time, is also available. All regulators involved in the commissioning of this report are included in the database.

Secondly, taking the metaphor of the dragging anchor seriously, the rope on the anchor cannot have infinite elasticity. Thus a dragging anchor in the form of, say, a fixed length moving average, at least has the merit that any sustained shift in the market measure will ultimately be passed through into the measure used by regulators. Inspection of Figure 4.2 gives serious room for doubt that the assumed value of the RFR used by UK regulators will *ever* average out to the same value as the market rate.³³

Thirdly, to the extent that a parallel is drawn between the approaches to assumptions on the RFR and the expected market return (EMR), discussed below, this seems to us to be a prime example of the conflation between problems of measurement and of implementation. Thus, PWC's recent report to Ofwat states:

*"In setting the cost of equity, UK regulators have tended to take a more through-the-cycle view of required returns. This involves the calibration of the cost of equity towards long-run historical averages for parameters such as the RFR, EMRP and TMR (EMR in our terminology).... Such an approach smooths out any short-term, cyclical variations and supports investor confidence by providing long-term stable returns on their long-term investments."*³⁴

This quotation offers a very different rationale to the original motivation proposed in MMW which related to a strategy for *estimating* the EMR with reference to historic averages (which we discuss further below). There is a clear need for such estimates because the EMR cannot be directly observed. But in contrast, the RFR can clearly be observed, with minimal error, so we do not need such an estimation strategy. Whether there are additional benefits to setting a stable *allowed* return on equity is a completely separate issue.

- **"The market for risk-free debt is distorted"**

It is often argued that a justification for regulators setting a value of the risk-free rate different from market rates is that various institutional features of the market for UK government debt (e.g., pension regulations, etc.) have led to yields on gilts (and especially indexed gilts) being artificially depressed (indeed the terms of reference—see Appendix L—include some of these arguments).

This argument may or may not be correct; we would argue that it is simply irrelevant: the market price of indexed debt (and hence its implied yield) is simply what it is. The following extract from Cochrane (2011) makes the point very succinctly and vividly:

When you shop for a salad, all you care about is the price of tomatoes. Whether tomatoes are expensive because the trucks got stuck in bad weather or because of an irrational bubble in the tomato futures market makes no difference to your decision.

We see no reason to treat the market for indexed debt differently to the market for tomatoes. Pursuing the analogy, if the price of tomatoes falls (for whatever reason), this is likely also to impact on the price of substitutes for tomatoes, such as cucumbers or other salad components. In the case of the risk-free rate, a fall (for whatever reason) must also have an impact on the price of other low-risk investments.³⁵

³³ Note that here (as discussed below) there is a distinct contrast with the approach currently applied to both the expected market return and, for some regulators at least, the cost of debt.

³⁴ PWC (2017), Appendix B, p 75.

³⁵ The only possible qualification to this conclusion would be a second-best argument: that if a case could be made that the distortions in the market for risk-free debt result in an outcome that is sufficiently sub-optimal, then regulatory policy might mitigate the effect of this distortion. But it is not easy to see how allowing regulated companies to earn higher returns would rectify this distortion.

- **“The risk-free rate is negative, hence unsustainable”**

Here, again, the price of tomatoes argument applies. There is no economic principle that rules out a negative risk-free rate; furthermore, there have been extended periods in the past when risk-free rates have been negative (see Figure 4.4 below). We would also note that, on the basis of both historical evidence and a priori arguments, even if the risk-free rate is less than zero this will not usually result in the cost of risky capital being negative.

- **“A negative risk-free rate is irrational”**

The argument here is that it cannot be rational to accept a loss for holding an asset with zero risk: an investor would be better off just not investing/borrowing than purchasing assets with a negative real yield. We accept that a negative (real) RFR is relatively rare; but it is not irrational and it is consistent with a standard decision-making model. In the benchmark optimising intertemporal model, the quite general equilibrium condition is

$$r \approx \theta + \gamma g - \frac{\gamma^2}{2} \sigma_c^2 + f$$

where r is the real risk-free rate, θ is the subjective discount factor (a measure of impatience), γ is a measure of risk-aversion (equivalently of aversion to shifting consumption over time), g is the expected growth rate of consumption, σ_c^2 is the variance of consumption growth, and f is a measure of financial market frictions.

If we assume away risk or financial frictions (and hence set the last two terms to zero) *and* we assume that consumers are at least somewhat impatient ($\theta > 0^{36}$) *and* that expected consumption growth, g , is non-negative, then yes, real risk-free rates should be positive. However, if consumption growth g is expected to be negative, the standard model would imply that consumers would attempt to save more to boost consumption in the future, thus boosting the price of the risk-free asset, and hence depressing its return. Furthermore, even in the absence of negative growth expectations, uncertainty about future consumption ($\sigma_c^2 > 0$) implies an additional “precautionary motive” for saving, that leads to depressed returns, and can in principle be large enough to offset the first two terms. Finally, and probably most importantly in quantitative terms, the last term, f , captures the impact of financial frictions, of which the most obvious is the “no short sales constraint”, i.e., only a very restricted class of borrowers can *issue* risk-free debt (essentially, fiscally sound governments and private borrowers with abundant collateral). Young investors, in particular, may wish to, but be unable to borrow. If they *could* issue risk-free debt, this would depress bond prices and boost risk-free returns. But if this constraint binds for the marginal investor, then this cannot happen, depressing the risk-free return ($f < 0$). There is a very substantial literature that suggests that the so-called “equity premium puzzle” is in fact explicable mainly by the risk-free rate being too low, relative to a frictionless world (i.e., $f < 0$), rather than the equity return being too high, precisely because “junior can’t borrow”. See Constantinides et al. (2002) There have indeed been claims that recent low real risk-free rates have actually been primarily driven by a global shortage of borrowers who actually have the capacity to issue truly risk-free debt; see Caballero et al. (2017).

³⁶ There is in fact a quite extensive literature in finance that claims to find evidence against this assumption, although most micro-based evidence would strongly support it.

The importance of accepting that the RFR can be negative is even more important if recommendation 3 (on the CPI) is rejected. Calculation of the real rate, both future and historic, can only be done in the context of a price index. The real risk-free rate derived from RPI linked bonds is unlikely to represent a real risk-free rate for the vast majority of individuals, because RPI inflation significantly overstates the true inflation rate for the majority of households. Thus, a negative real rate as measured using the RPI could reflect bias in the measure of inflation rather than a negative risk-free rate. Retention of RPI for indexation purposes and placing a zero lower bound on the real risk-free rate would represent a compounding of errors.

To conclude, we do not believe that any of the above defences of recent regulatory assumptions on the risk-free rate hold water. And, we would note in contrast that a simple strategy of simply picking a given horizon and using the market risk-free rate at that horizon clearly satisfies our two key criteria of implementability and defensibility.

Our conclusion does however have important implications. As outlined in Section 3.3, the RFR does not matter very much when Beta is close to 1, and this may have led regulators to largely ignore it. However, given our conclusions on the RFR and EMR, the setting of beta is the sole component of the cost of equity that the regulator has significant discretion over, and this will need to be exercised with care, as we discuss later in the report.

4.4 The expected market return (EMR)

4.4.1 The MMW approach revisited

In contrast to their treatment of the risk-free rate, UK regulatory practice has typically (with some recent exceptions) been much closer to the recommendations of MMW.

MMW proposed a methodology in which Expected Market Return (EMR) (that is, the expected real return on investments in the equities of a firm with a CAPM β of precisely one), should be assumed constant, and set in the light of realised historic real returns in a range of stock markets, over long samples.

As stressed above, this methodology is about deriving an *estimate* of the EMR. It does *not* claim to be a precise description of the *actual* EMR (which is of course not directly observable).³⁷

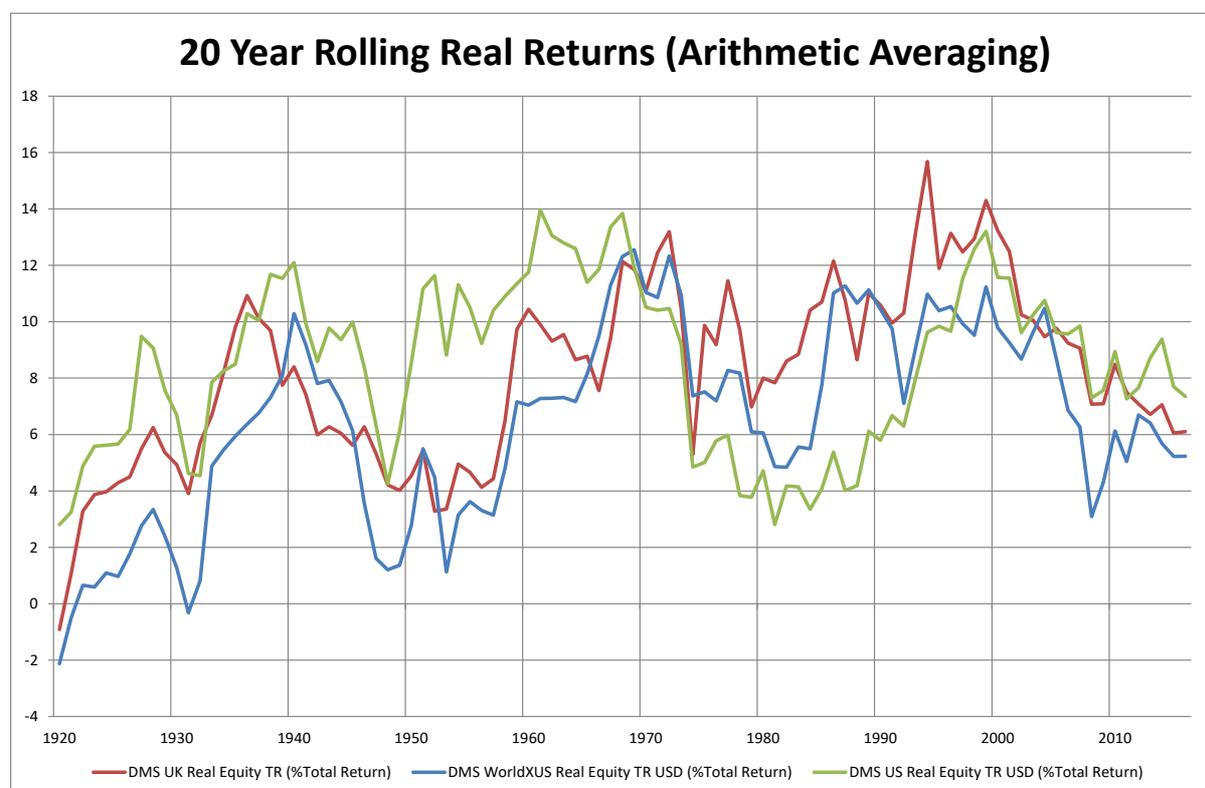
In Appendix E (which largely follows the original MMW methodology as updated in Smithers and Wright (2013)), we consider the impact of recent data on these historic averages (including the impact of our proposed adoption of CPI inflation as the basis for adjustment into to real terms). Unsurprisingly, given the relatively small amount of new data to be added to the original samples, the numerical answers do not change very significantly.

Figure 4.3 and Figure 4.4 provide a simple illustration of the basis for the MMW methodology: the apparent stability of long-run real returns on mature stock markets, especially in comparison with the considerably *lesser* degree of stability of returns on other major asset classes.

³⁷ Indeed, in the original CAPM formulation the market return should be the return on a portfolio consisting of all the assets in the economy, of which the stock market actually represents only a small fraction. The longstanding tradition of proxying the market return by the return on a broad stock market index is itself a significant compromise based on the need to derive estimates of an unobservable true magnitude.

Figure 4.3 illustrates the similarity of long-run average returns on the UK and US stock markets, and (especially in the post WW2 era) also of returns in world markets excluding the US.³⁸

Figure 4.3



While 20-year returns to 2016 (the last complete year shown) were relatively low, this is arguably to be expected given that we are comparing returns with the latter period of the bull market of the 1990s; 20-year returns during the late 1990s were, as is also to be expected, unusually high. But, despite this point of comparison, the most recent 20-year returns are well within the historical range.

Figure 4.4 takes an even longer historical perspective, and updates a chart originally included in MMW: it brings out the remarkable stability of the US stock return over more than two centuries. The issues behind Figure 4.4 are discussed in considerable depth in MMW (Section 2.4), who also noted the consistency with long-run data on real stock returns from a wide range of other markets.³⁹ But we highlight one key feature: the distinct *lack* of stability of long-term returns on competing asset classes, the real return on “cash” (i.e., short-term deposits, hence Figure 4.4 shows the ex-post real interest rate), and on bonds (where the return shown includes capital gains and losses).

While MMW obviously did not predict the prolonged fall in real rates that has been observed since the financial crisis, they *did* point out that such prolonged shifts—in either direction—were if anything the

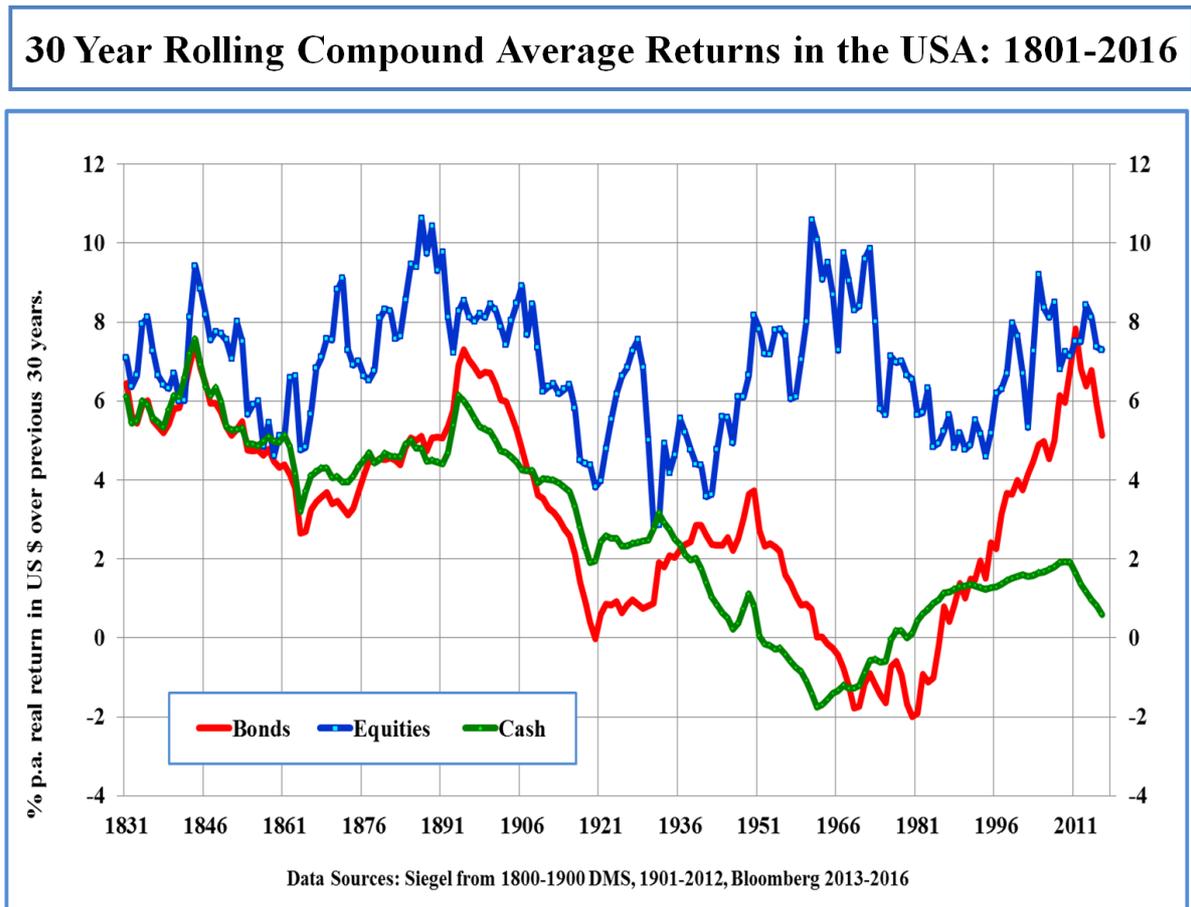
³⁸ Returns on European markets during the early part of the twentieth century were arguably depressed by exceptional, and one-sided shocks: most notably the hyperinflations of the 1920s and the Second World War. For a detailed discussion of this issue, and of a number of other aspects of long-run return series, readers are referred to the discussion in MMW, Section 2.4.

³⁹ The academic evidence on this score appears to be progressively building, including indications that the stability of real returns appears to be a general feature of *real* assets, not just stock markets: thus there is growing international evidence of stability of returns on housing, at similar levels to the long-run stock return. See e.g., Jorda et al. (2017).

norm, rather than the exception. This was already quite evident at the time of the original report in 2003. Thus Figure 4.4 is also a reminder that, while the recent period may be historically unprecedented in nominal terms, it is far from unprecedented in *real* terms: i.e., the ex-post real interest rate was distinctly more negative in the 60s and 70s than it has been in recent years.

By implication at least, ex-post the equity premium has been very far from stable; and the long-term historical swings have been so large that it would very hard to reconcile such a pattern with a stable forward-looking equity premium.

Figure 4.4



Thus, 15 or so years on from the original MMW report, we do not see any obvious evidence in the history of returns themselves to cast doubt on the key evidential basis for the treatment of the EMR: that long-run stock returns are stable in real terms.

An important caveat, discussed extensively in the original MMW report,⁴⁰ is that using the longer horizon preferred by regulators does point to a smaller adjustment from geometric to arithmetic returns, in light of evidence of univariate return predictability. So this is a case where the choice of horizon does have non-trivial quantitative implications.

⁴⁰ See MMW Section 2.4.2.

A more substantive question, and one which has been raised in a number of recent contributions to the debate (particularly the most recent Ofwat deliberation), is whether developments in both academic work and in financial markets have undermined the MMW methodology itself.

We consider three alternative approaches that have been proposed elsewhere:

- Alternative A: Assume the equity risk premium (ERP) is stable.
- Alternative B: Update the estimate of the EMR (and hence ERP) in line with econometric evidence of return predictability.
- Alternative C: Estimate the EMR using alternative models such as the Dividend Discount Model.

4.4.2 Alternative A: Assume the equity risk premium (ERP) is stable

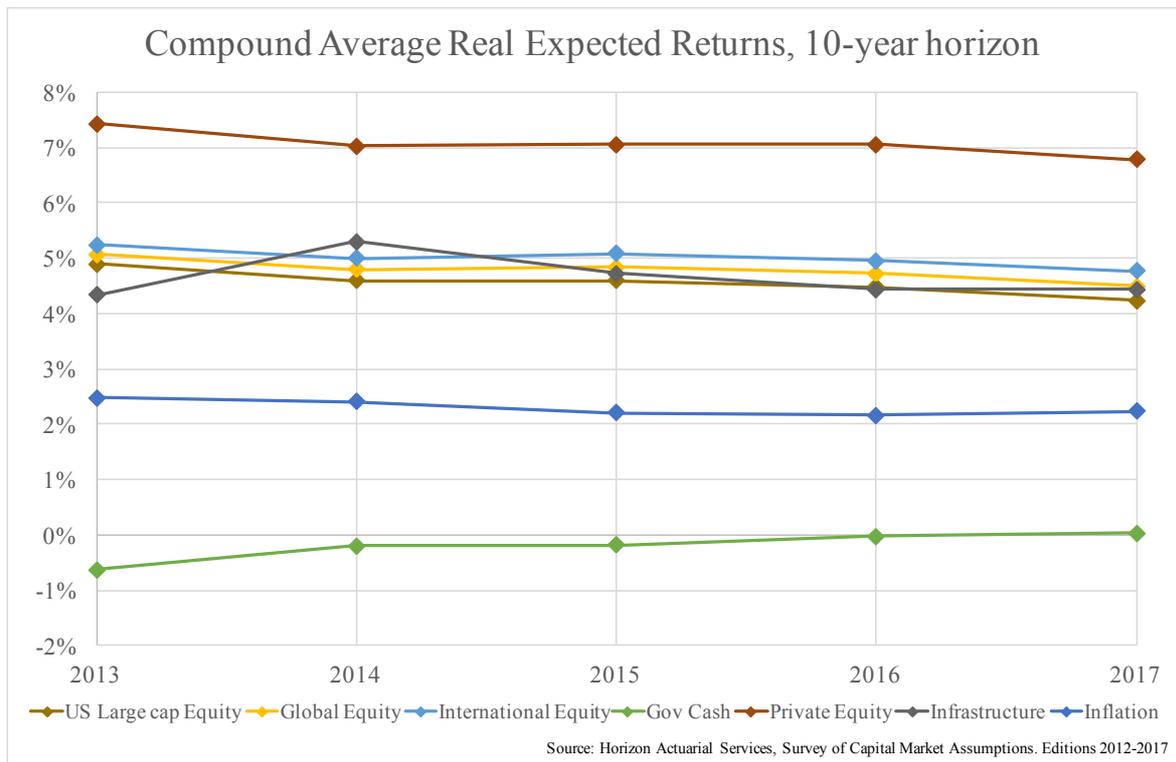
The empirical basis for this assumption was always weak, compared to the MMW methodology; but it has been further undermined by more recent evidence that risk premia are countercyclical. Smithers and Wright (2013) discuss this issue at some length, and examine the academic evidence. Here we simply stress again that, while evidence for counter-cyclical risk premia is strong (and is supported, indirectly, by some recent applications of Alternative C⁴¹), this should not be taken as a claim that the ERP instead moves *precisely* one-for-one in the opposite direction to the RFR. We discuss this issue further below.

4.4.3 Alternative B: Update the estimate of the EMR (and hence ERP) in line with econometric evidence of return predictability

In contrast, we recognise that the evidence base for Alternative B (which was explicitly acknowledged in the original MMW analysis) has strengthened. Most notably, the academic evidence of return predictability has, to a considerable extent, been taken on board by financial practitioners. Figure 4.5 provides some evidence of a consensus amongst financial practitioners that expected returns on global stock markets have been weakening in recent years (although the extent of this weakening is not especially dramatic):

⁴¹ See, for example, the discussion of the Bank of England's estimate of the equity risk premium in Section 4.4.4 below.

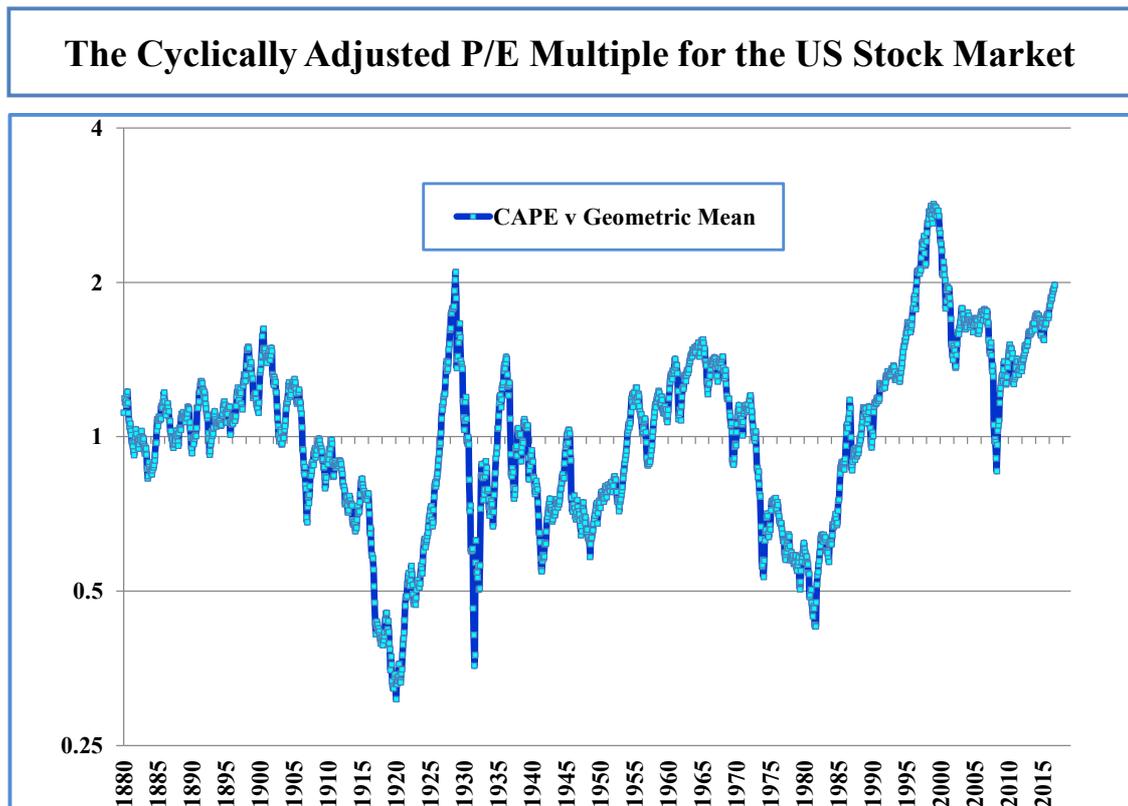
Figure 4.5



The basis for these falls in expected returns is increasingly driven by the recognition (informed by academic research) that valuation indicators such as the cyclically-adjusted P/E multiple shown in Figure 4.6 (popularised by Robert Shiller) appear to have predictive power for (at least) the US stock market.⁴²

⁴² This has certainly been helped by the presence of evidence presented by Shiller and John Campbell to a congressional sub-committee in the late 1990s, which used the cyclically adjusted P/E to (correctly) predict the end of the 1990s bull market, see Campbell and Shiller (1998).

Figure 4.6



Source: Robert Shiller to September 2017

Does this evidence undermine the MMW methodology? The simple answer is: yes, it does; and arguably all the more so if regulators wish to evaluate the EMR at a long horizon, since most of the evidence for predictability is at long horizons, with short-term return behaviour being much harder to distinguish from the random walk model.

Yet the practical problem for regulators is what to do about this.

While the *qualitative* evidence for return predictability (and for this predictability being embodied in market expectations—which is not necessarily the same thing) is quite strong, it is much harder to point to an agreed *quantitative* methodology that could be employed to capture this feature in a methodology that is both implementable and defensible.

On implementability: if regulators are looking for a well-defined mechanism for updating the EMR they would need to look to econometric evidence. Here there is a distinct contrast with, e.g., beta estimation. While some of the authors argue below that current regulatory practice on beta estimation is far from ideal, we would all agree that it is possible to rectify any problems by a careful application of well-established and standardised econometric techniques. In contrast, there is as yet no clear consensus on the quantitative link between e.g., valuation ratios and future returns (indeed the bulk of the available evidence relates only to the US stock market). Furthermore, while the majority of academic studies do point to return predictability, it remains a controversial topic. As a result, the defensibility of any assumed mechanism would be open to serious doubt (we suspect that it would not be hard for regulated

companies to find academic economists who would be prepared to debate the econometric methods of any mechanism that was chosen).

It is also worth noting that, to the extent that there is evidence of econometric predictability, this evidence *is* consistent with the MMW approach being correct at least in the long run, since pretty much all such evidence relies on the long-run stability of valuation ratios (like the P/E multiple) which would not be stable unless returns themselves were stable. Thus the MMW approach does at least satisfy the “dragging anchor” criterion that, as long as returns are indeed stable, it will get the EMR right on average.⁴³

4.4.4 Alternative C: Estimate the EMR using alternative models such as the Dividend Discount Model

In recent price control periods Ofwat, in particular, has drawn from the Dividend Discount Model, amongst other approaches, to inform its estimate of the EMR.

There is indeed no doubt that applications of the DDM are widely used by reputable organisations, such as the Bank of England, and in a quite large body of academic research.⁴⁴ And, as outlined in Section 3.5 above, there is in principle no conflict between the CAPM and the DDM: if carefully applied, they are mutually supportive and complementary.

One immediate caveat is that, typically, the *purposes* for which the DDM is used are often quite distinct from regulatory ones. The Bank of England’s most recent application for example, as discussed in Section 3.5, uses the model as an accounting procedure to explain shifts in the stock market after the event, *not* to predict returns. Indeed, the authors of the article state quite clearly a number of caveats:

“As the ERP cannot be observed, any estimate of it is necessarily subject to uncertainty....The inherent uncertainty about the true value of the ERP is reflected in the wide dispersion of ERP estimates in the literature. ...Given the uncertainty associated with measuring the ERP, the Bank’s analysis tends to focus less on the precise level of the ERP and more on changes in the ERP over time or on the level of the ERP relative to historic averages.”

As noted by Sudarsanam et al. (2011), except for the USA, international regulators rely primarily on the CAPM to calculate the cost of equity of regulated entities.⁴⁵ Subject to this caveat, what can the DDM contribute to the process of assessing the EMR, in particular?

We start by noting a key feature of the DDM approach. As analysed above, in Section 3.5, if aggregate stock prices change, then the internal cost of capital approach to the DDM implies that this *must* be explicable in terms of some combination of changes in the internal cost of market equity⁴⁶ or in the expected growth rate of dividends per share. If, as is typically the case, estimates of growth update only

⁴³ Note also that the evidence of return predictability does imply that the use of historical averages may in principle be consistent with taking some account of the possibility that these averages themselves may be at least somewhat distorted by recent market movements. Thus the original MMW report (see Section 2.4.5) discussed the impact of the extreme valuations of the late 1990s on long-term return estimates. Given long samples these effects were modest, even at the peak of the market in 2000, but not negligible (of the order of half a percentage point).

⁴⁴ See references cited in Section 3.5 above.

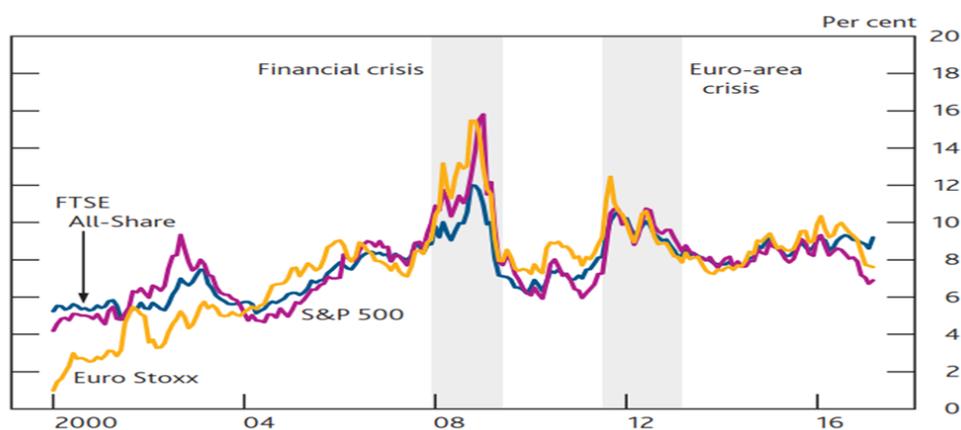
⁴⁵ BERC’s 2017 report indicates that no European telecoms regulator relies on DDM alone to estimate the ERP.

⁴⁶ As discussed in Section 3.5 this magnitude is not a direct estimate of the EMR since it is a summary statistic of expected returns at all horizons.

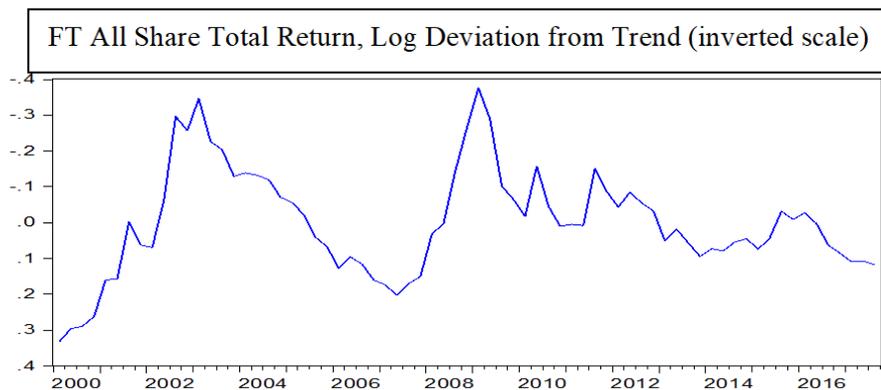
relatively gradually, then almost by necessity a rise in the stock market implies a *fall* in the implied internal cost of market equity.

This is illustrated by Figure 4.7. It takes a chart of estimated equity premia for the UK, the US and the Eurozone from Dison and Rattan (2017). Below this we insert a chart of the path of the FT All Share total return index over the same period, linearly detrended, and on an inverted scale: thus when the line in the lower chart is high, the FT was *below* its trend over this sample.⁴⁷ The correlation with changes in the Bank of England’s estimated premium for the UK is quite marked.⁴⁸ (We would *not* expect it to be perfect since both real interest rates and growth rate forecasts were also changing over this period.)

Figure 4.7



Sources: Bloomberg, IMF *World Economic Outlook*, Thomson Reuters Datastream and Bank calculations.



To a nontrivial extent, given the algebra set out in Section 3.5 above, this must be a feature of *any* application of the DDM: any recent sharp movement in stock prices will typically show up as a

⁴⁷ The implied trend return over this period was only just over 6% nominal, but note that the starting point was the market peak in 2000, so we would expect this to be a distinctly downward biased estimate of the true trend return.

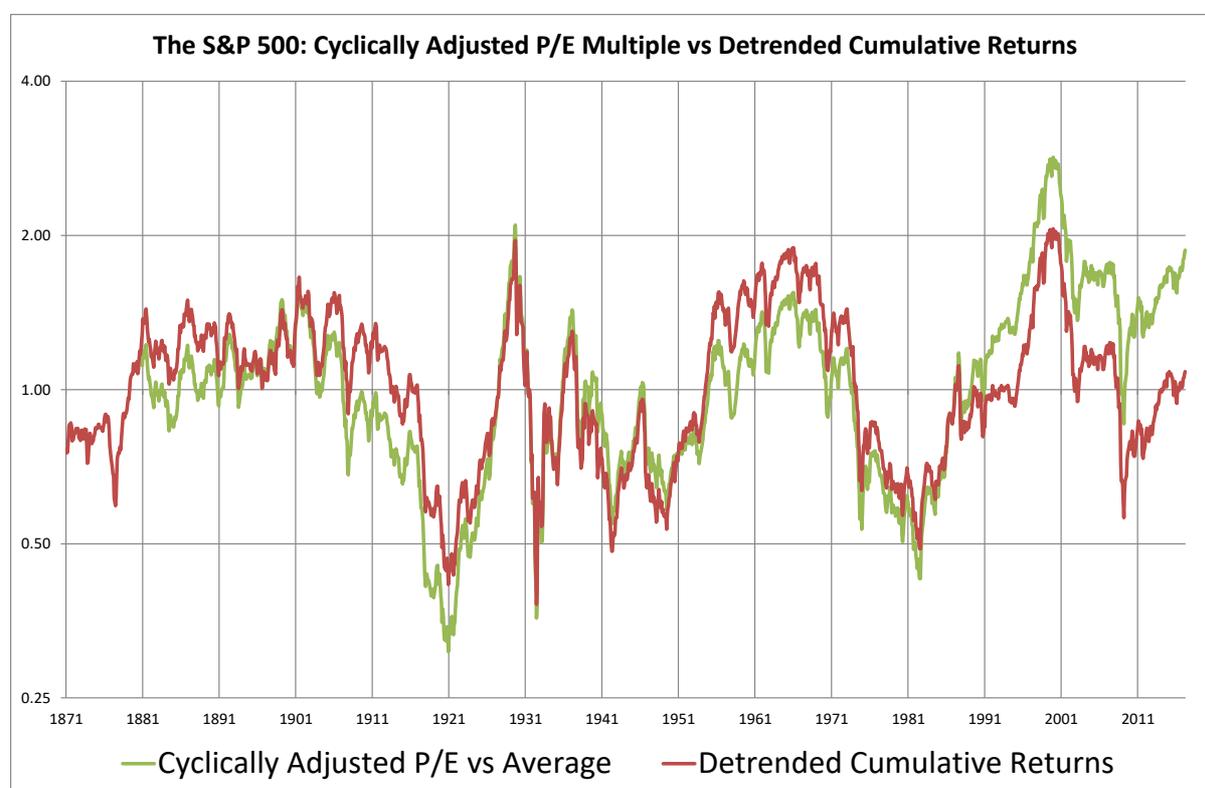
⁴⁸ Note that, while the Bank of England’s model is closely related to the internal cost of capital approach it differs somewhat, since rather than assuming a single internal cost of market equity, it assumes a fixed risk premium over observed market yields along the yield curve, and thus estimates this risk premium, rather than the internal cost of market equity itself. As a result, while the article reports a DDM-style representation (their Equation (6)) the authors have confirmed in email correspondence that this is in fact only an approximation for the true expression.

movement in the implied internal cost of equity in the opposite direction.⁴⁹ So in some sense, for all its forward-looking element, the DDM is also driven by history – it is just driven by more recent history.

This also helps to explain why at least some applications of the DDM have predictive power for subsequent returns.⁵⁰ As shown above, the internal cost of equity is a summary statistic of future expected returns at all horizons implied by the model, with a bias towards longer horizons. So a fall in the internal cost of equity implied by the DDM after a rise in the stock market implies a prediction of lower returns over long horizons. Campbell and Shiller, and many others, have also shown that valuation indicators like the cyclically adjusted P/E multiple mainly have predictive power for longer horizon returns. But in practice movements in the cyclically adjusted P/E (CAPE) (and most other valuation indicators) are themselves dominated by recent movements in returns.

Figure 4.8 illustrates, by comparing movements in CAPE with the deviation of cumulative total real returns from its long-run trend: the resemblance for most of the sample is very striking.⁵¹

Figure 4.8



⁴⁹ Mechanically, in Equation (19) a rise in the stock market, for example, will imply a fall in the prospective dividend yield. For any given value of G_{it} , the estimated internal cost of equity, ICC_{it} , must fall in line with the prospective yield.

⁵⁰ Albeit that most such exercises are built up from bottom-up applications of the DDM to individual stocks: see for example Li et al. (op cit).

⁵¹ It is also notable that the pessimism about future stock returns signalled by CAPE in recent years is not evident in the pattern of detrended returns.

Thus in practice much (and possibly all) of the predictability found in the academic literature *also* arises from the simple property that recent rises in the stock market tend to be bad news for future returns.⁵²

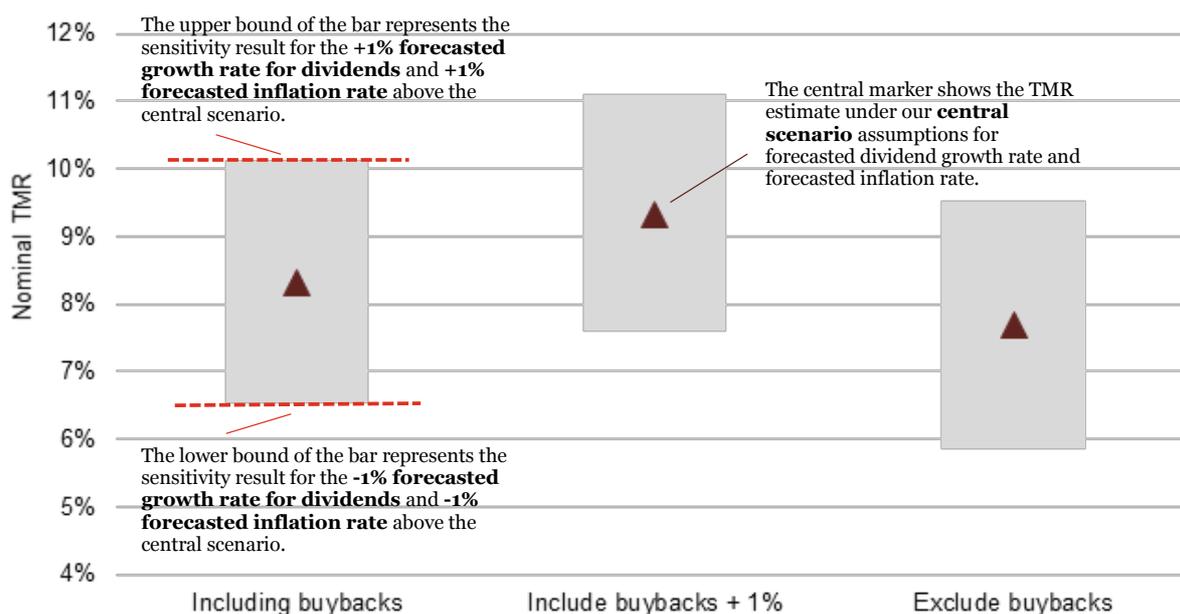
Thus, to the extent that the DDM does have some predictive power, at least some of that predictive power would appear to be consistent with the other evidence on return predictability.

But of course, this is not *just* what the DDM does: clearly there is also a contribution from the forecast growth rates of dividends per share that also feed into the model. But here, we note the clear point of logic discussed in Section 3.5 above: the calculated internal cost of market equity, *conditional* upon these forecasts, is only as reliable as the forecasts themselves. Errors in growth forecasts feed more or less one-for-one into errors in the estimated cost of market equity.

We can illustrate the difficulties that may arise here with reference to one recent application of the DDM: PwC’s 2017 report to Ofwat, although we note Ofwat referred to a wide evidence base and placed limited weight on DDM. PwC’s Figure 26 is reproduced below (Figure 4.9). This shows sensitivities of their EMR (here denoted TMR) estimates to changes in assumptions feeding into their model. These are very wide ranges indeed: considerably wider than the range of long-run historic average returns.

Figure 4.9

Figure 26: DDM spot output sensitivities



Source: PwC analysis

The ranges of roughly 2 percentage points either side of the midpoint shown in Figure 4.9 are indeed exactly in line with what Equation (19) in Section 3.5 would predict, since they allow for errors of 1 percentage point in both (nominal) growth of dividends per share and in inflation (and hence real

⁵² This feature is often referred to, rather confusingly, as “mean reversion of returns”, but this strictly refers to the feature that *cumulative* returns have a mean-reverting component (as opposed to being a pure random walk with drift).

growth). Admittedly an error as large as 1% in predicting inflation *in perpetuity* would arguably be on the extreme side, in an era of inflation targeting; however, it is perfectly plausible that the range of errors in forecasting real growth of dividend per share growth might well be larger, so that the ranges shown are probably not unrepresentative.⁵³

We do not wish to detract from the value that DDM-based analysis may provide, particularly as a cross-check for our preferred methodology. However, if we focus purely on the objective of estimating the EMR, we revert, yet again, to our twin aims of implementability and defensibility.

On the issue of implementability, on the face of it our Recommendation 2, that the choice of inflation index should be subcontracted to the Bank of England, might appear to have a natural extension: could not regulators simply make use of the Bank's DDM to generate estimates of the EMR?

There would indeed be no logistical obstacle to doing so. While the Bank only publishes implied ERP numbers these could relatively easily be converted into implied EMR estimates.⁵⁴ However, Figure 4.7 above suggests that had regulators actually implemented these estimates over the period since 2000, there might well have been serious difficulties persuading consumers that doing so was consistent with their statutory duties.

Thus, the Bank's estimated ERP for the UK market peaked during the crisis at 12 percentage points. At this point UK indexed yields were still mildly positive so the total implied EMR would have been even higher. It is hard to imagine that giving regulated companies a bonus of such a magnitude, even if only shortlived, at the time of the global financial crisis, would have been politically acceptable.⁵⁵

Of course, we are well aware that no regulator is actually proposing to subcontract the process of estimating the EMR in this way. But even this hypothetical exercise does illustrate the practical difficulties in implementation.

We have also already highlighted the difficulties involved in *defending* such an approach against the wide range of criticisms of possible methodologies for predicting growth rate of dividends per share, and the resulting extremely wide range of uncertainty about the resulting estimates.⁵⁶

⁵³ Thus, as just one example, both PWC and the Bank of England make the common assumption that growth of dividends per share over the long-term will be driven by growth of GDP. There is no basis for this assumption either in theory or in the data: see, for example, Alessandri, Robertson and Wright (2008) who draw on the original analysis of Miller and Modigliani (1961) to show that while steady state growth may anchor the growth rates of *aggregate* dividends, the same does not apply to the growth of dividends *per share*. This is also evident in the data: the long-term growth rate of real dividends per share in the United States (over the period 1900-2013) was only 1.2% (calculated using data from Robert Shiller's website) compared to the growth rate of real GDP of 2.6% over the same period. Alessandri et al., along with other authors, show that correcting dividends for other forms of cash payments to shareholders actually increases the volatility of total payments so much that long-term growth rates become even harder to estimate.

⁵⁴ As noted above, the methodology would imply that the EMR would have a maturity profile identical to the term structure of interest rates, with a constant premium at all maturities at any point in time; but this premium would update with movements in the stock market, as outlined above for the internal cost of capital.

⁵⁵ Stocks of quoted regulated utilities in any case weathered the storm with considerably less pain than the market index, The FT all share fell by nearly 40% from peak in 07 q2 to trough in 09 Q1, whereas both UU and SVT fell by less than 25%.

⁵⁶ The Bank of England for example cite a survey of ERP estimates for the S&P 500 that vary between -1% and 14.5%."

Thus we cannot point to a methodology for using DDM to calculate the EMR that would be simultaneously implementable *and* defensible.⁵⁷

Despite all of the above reservations, it is perhaps worth noting, as a final point, that both the longer-term properties of DDM-based estimates of the EMR and the implied ERP, *and* values in the recent past, typically actually look very similar to those that arise from the MMW methodology. Thus the Bank of England's most recent ERP estimate of just over 8 percentage points is broadly consistent with a long-run historic average of around 6-7% real, and a real risk-free interest rate of around -2%. Furthermore, the rise in their estimated ERP since 2000 (the earliest point in the chart) has been by around 4 percentage points, almost exactly matching the fall in the RFR over the same period. So, both on the most recent data, and in terms of longer-term trends, at least since 2000, the Bank of England's methodology is largely consistent with MMW's methodology.⁵⁸

4.4.5 Secular Stagnation?

We should note briefly here an important contrast between any criticisms of the MMW approach that might be made based on return predictability (which as noted above are predicated on the long-run stability of returns), versus arguments based on the "secular stagnation" hypothesis: that long-run returns on all asset classes, including stocks, may have fallen permanently.

On this point, it is important to stress (following on from the discussion around Figure 4.4 in Section 4.4.1) that the sustained falls in nominal and real interest rates since the financial crisis are *not*, in themselves, evidence of a secular shift in the long-run value of the EMR. Figure 4.4 showed that similar, or larger sustained shifts in real returns on risk-free assets and bonds have occurred historically, *without* any evidence of similar shifts in the mean stock return. Indeed, the *lack* of a stable mean value of the real risk-free rate was a key part of the motivation for the original MMW approach based on the stability of real stock returns. Subsequent events have only reinforced this original conclusion.

This is not to deny the logical possibility that the hypothesis of a sustained, or even permanent fall in the EMR may ultimately be proved correct. But the problem with such arguments is that, whether or not we find them persuasive, they can neither be supported nor refuted by data until after the elapse of long periods of time. We can directly *observe* the fall in the RFR; but, we reiterate, the EMR is not observable. As MMW noted in their original report, even over extremely long samples we can only *infer* it from the long-run average of *actual* stock returns.⁵⁹ Given the volatility of market returns it would take decades to provide clear statistical evidence of a structural shift in the mean real stock return: i.e., to reject the null hypothesis of *no* shift. As such, attempting to take account of such secular shifts (in the face of more than two centuries worth of evidence, from multiple markets, in favour of stable stock returns) would clearly fall foul of our twin criteria of implementability and defensibility.⁶⁰

⁵⁷ At least given the current state of knowledge. We cannot of course rule out the logical possibility that such a mechanism might in due course emerge.

⁵⁸ Nor should we be surprised by this, since the Bank make it clear that their ERP estimates are systematically cross-checked against long-run data on realised returns.

⁵⁹ See discussion in MMW Section 2.4.1.

⁶⁰ Note also that, as discussed in Section 4.4.3 above, the academic evidence of return predictability is almost invariably crucially dependent on the assumption that the long-run average return is constant; yet this directly contradicts arguments that there may have been a secular downward shift in all expected returns.

4.4.6 Implications for the assumed equity premium.

A direct implication of our recommendations on the EMR and the RFR is that the implied market risk premium (MRP)⁶¹ automatically follows as the difference between the two numbers. Thus, since the assumed EMR is extremely stable, the implied MRP moves in the opposite direction (and one-for-one) with the RFR.

This issue was discussed at some length in Smithers and Wright (2013), and we do not have any new substantive arguments to add to their analysis, which we simply summarise here.

The most crucial point we wish to stress again is that this methodology is *not* designed to be precisely descriptive of the path of the true MRP. The *calculated* MRP implied by our methodology has expanded one-for-one in recent years, in line with the gap between the historic average EMR and the market risk-free rate. But the evidence of lower market returns that we alluded to in our discussion of the EMR above would imply that the *actual* MRP may not have expanded so much.⁶²

However, Smithers and Wright (2013) point to fairly strong evidence that, while this feature may not be precisely descriptively accurate, it does at least capture the qualitative feature that risk premia tend to move counter-cyclically

In practical terms, the key feature of our proposed approach is the contrast with Alternative A (as set out in the previous section) in which the MRP is assumed constant, so that the market cost of equity, and hence the implied CAPM-WACC would simply track the risk-free rate down. In recent data, given the fall in the RFR, this would have resulted in dramatically lower costs of capital even for companies with β close to 1, whereas our proposed implementation only impacts on CAPM-WACC estimates for companies with beta not equal to (and in practice, typically less than) 1.

4.4.7 Conclusions on the EMR.

Having considered a number of different approaches, we conclude that the treatment of the EMR proposed by the original MMW 2003 report remains valid.

Recommendation 5 (The Expected Market Return): We recommend that regulators should continue to base their estimate of the EMR on long-run historic averages, taking into account both UK and international evidence, as originally proposed in MMW. We suggest a modest downward adjustment of the original range proposed by MMW, to a range of 6-7%, primarily reflecting a smaller adjustment from geometric to arithmetic returns.

That is, the best means—one that satisfies the twin criteria of implementability and defensibility—to estimate the EMR is to assume that it is constant, and to look at realised historic real returns in a range of stock markets and over long samples. As DMS (2017) conclude (p. 41):

⁶¹ Commonly (and more precisely) referred to as the ERP (the equity risk premium) but for consistency with our use of the term EMR for the expected market return we use the term MRP to represent the risk premium given $\beta = 1$.

⁶² Although we noted above that in practice, e.g., the Bank of England's estimate of the MRP has seen a longer-term shift almost exactly in line with that implied by the MMW approach.

“For practical purposes, it is hard to beat extrapolation from the longest history available when the forecast is being made.”⁶³

4.5 Beta Estimation Issues

In their original report in 2003, MMW took it as given that the CAPM should be applied to the cost of equity with estimated equity betas, and set out an econometric methodology for deriving such estimates. Subsequent reports to regulators have suggested amendments to the methodology, including estimating betas using time-varying parameter models, and investigating multifactor models.⁶⁴

In recent deliberations both Ofwat and Ofgem have assumed values of equity beta quite close to 1 (with values of 0.8 and 0.9 respectively). For values of equity beta *precisely* equal to 1, as noted in Section 3.3 above, the risk-free rate drops out of the cost of equity formula. As a result, the implied CAPM(E)-WACC (the calculated WACC when CAPM is assumed to hold only for the cost of equity) at least, becomes a weighted average of the expected market return and the assumed cost of debt. So when regulators use values of equity beta close to 1, this effectively minimises the role of the risk-free rate as a determinant of the cost of equity⁶⁵. In a period during which the RFR has shifted so dramatically, this has potentially major consequences, which suggests that the estimation of beta should be critically reviewed.

While the authors are in full agreement that beta estimation is a crucial issue, they have not been able to reach a consensus on the quantitative conclusions to be drawn from econometric analysis. The quantitative analysis in this section focuses on beta estimation for energy and water networks. The discussion on beta estimation methods and implementation should be of relevance to all regulators.

In Appendix F, Burns summarises the results of estimating equity betas for a range of UK energy and water utility companies, as well as some international comparator companies. Whilst there is no single “standard approach” universally adopted by the UK regulators, the methodology used by Burns to estimate the betas broadly reflects the approaches taken by regulators and by the CMA in recent determinations.⁶⁶ using different (2 to 5 years) samples of data at varying frequencies (usually daily, weekly or monthly). For comparative purposes, Burns also provides equity beta estimates over the most recent ten-year sample. An example of the results for five-year equity betas “geared up” at 65% notional gearing is set out in the table below, the first three columns show betas estimated against the UK all share index, the final column against a world index.

⁶³ We were asked specifically (see the Terms of Reference in Appendix L) to consider two issues noted by DMS: (i) historical relative outperformance in the UK equity market; and (ii) expansion in the price/dividend ratio. On the former issue, our analysis in Appendix D argues that the historical relative outperformance documented by DMS can largely be attributed to the price index that they use. Correcting their price index brings the real return of UK equities closer to the average historical experience across the world: see Table D.4. On the latter: as we note above (see section 4.4.3), using various financial ratios (e.g., price-earnings, price-dividend or others) to estimate the EMR presents considerable difficulties for regulators who want to use an approach that is both implementable and defensible.

⁶⁴ See, for example, Wright (2005), Wright et al. (2006), Imrecon/ECA (2012).

⁶⁵ Though it would still flow through into the cost of debt, tempered initially any ‘dragging anchor’.

⁶⁶ For example, CMA Bristol Water 2015.

Five-year equity betas at 65% notional gearing

	5-year daily	5-year weekly	5-year monthly	5-year daily with world index
Severn Trent PLC	0.87	1.05	1.20	0.76
United Utilities Group PLC	0.83	0.90	1.04	0.67
National Grid PLC	0.97	1.05	1.08	0.76
Pennon Group PLC	0.97	1.11	0.96	0.83
SSE PLC	1.37	1.25	1.06	1.28
Terna Rete Elettrica Nazionale	0.92	0.79	0.64	1.41
Snam SpA	0.96	0.86	0.59	1.42
Enagas SA	0.85	0.84	0.80	1.25
Red Electrica Corp SA	0.95	1.06	1.05	1.34
REN - Redes Energeticas Nacion	0.53	0.53	0.74	0.58
Average	0.92	0.94	0.91	1.03

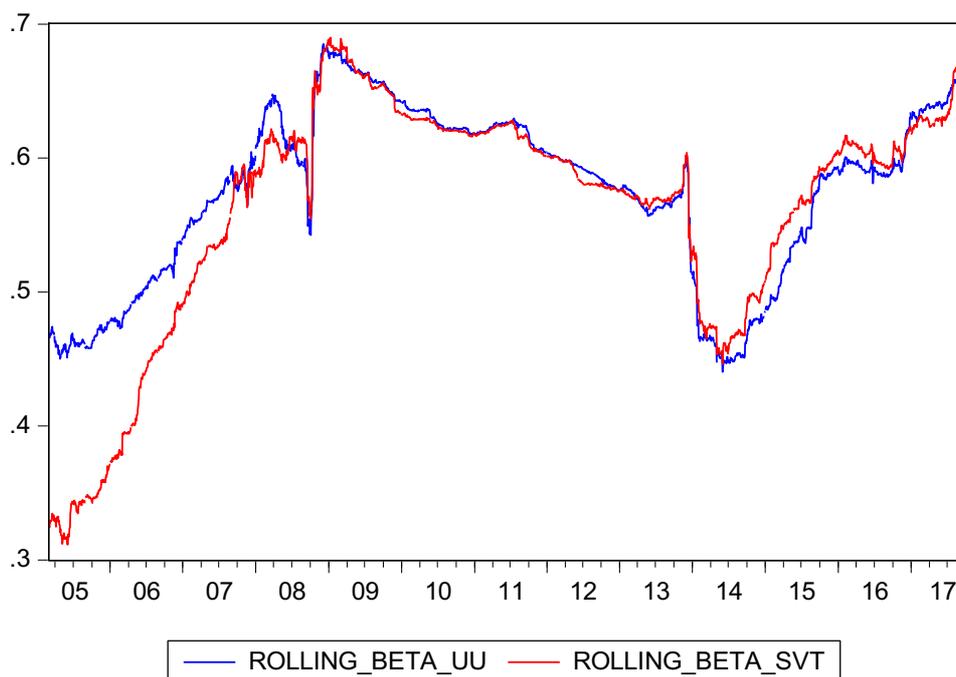
Source: Bloomberg data, Frontier Analysis

This approach is often supplemented by examining any evidence from “rolling beta estimation”, in which beta is estimated over a window of data of a fixed length (usually 2 to 5 years), but examining the pattern of betas that appears to emerge as the window moves through a longer sample of data. The assumption here is usually that this may reveal if the estimated beta has a tendency to drift over time.⁶⁷

Figure 4.10 illustrates, using data for two quoted companies, United Utilities and Severn Trent, that, at least since around 2009, have been close to pure-play regulated water companies. The chart shows raw beta estimates from daily data. At any point in the chart, the estimate shown is based on the five years up to that point (thus the chart actually exploits a full sample of data going back to the start of 2000).

⁶⁷ The econometric basis for using this approach is actually fairly shaky: in particular all parameter standard errors are invalidated by this methodology.

Figure 4.10
Rolling 5-Year Equity Beta Estimates: Daily Data



A first key feature to note in the chart is that at no point during the entire sample did the estimated value of equity beta exceed 0.7. In Appendix F, Burns shows that, on some samples (most notably the most recent five-year monthly sample), a few estimates lie outside the range in the chart; they also show that there is a wider range of estimates for comparator companies (especially European comparators, using a world market index). But virtually all the raw beta estimates⁶⁸ lie below 1, and average values are typically in the range 0.6 to 0.7 on the most recent five-year sample of daily data, and somewhat lower over the most recent ten-year sample.

This difference between 5-year vs 10-year estimates can also be related straightforwardly to a second key feature of the chart above, which shows that the rolling 5-year beta estimates have been tending to drift upwards in recent samples. So a clear issue worthy of investigation is whether these recent shifts towards a value of around 0.7 are likely to be sustained, or whether there is evidence that beta may revert to the lower average values seen in earlier samples.

But there is a further, more basic issue relating to beta estimation. If regulators wish to estimate the CAPM-WACC appropriate to a relatively long horizon (say, 10 years), is it appropriate to estimate beta over such a short sample (often distinctly shorter than the horizon itself) and using high frequency (daily or weekly) data?

We should start by noting that there is one important benchmark case in which the length or frequency of the sample used in estimation should *not* affect beta estimates, and indeed in which high frequency estimation may be preferred. If returns on both the market and the individual stock are serially uncorrelated (i.e., cannot be predicted from the history of either return) and have volatilities and a correlation that are constant over time,⁶⁹ then an econometrician should expect to estimate the same

⁶⁸ That is, beta estimated directly and before being re-levered to a common gearing level (as reported in the table above).

⁶⁹ Formally if both returns are jointly IID at all frequencies, and thus are a joint Markov Process.

value of beta in any sample, at any frequency – and since high frequency data offers more observations for a given amount of calendar time, the resulting beta estimates should in principle be more precisely estimated.

By implication, however, if there is evidence that these conditions are *not* satisfied, then the length of sample and the frequency with which returns are measured *may* matter. In statistical terms the assumptions may be violated in two distinct ways: volatilities and correlations may not be constant (heteroscedasticity), and/or future returns may be predictable from past returns (serial correlation).

To give an illustration by analogy, consider the case of a 10-year risk-free indexed bond. Since market yields fluctuate on a high frequency basis, the daily holding return on such a bond has nontrivial volatility. Furthermore, some of this volatility may be correlated with changes in stock markets, so typically on daily data such a bond would have a non-zero estimated market beta. However, if held to maturity, such a bond clearly has a beta of precisely zero. Furthermore, returns on such a bond are “heteroscedastic” (since its volatility declines as it approaches maturity, as must its correlation with the market return) *and* serially correlated (since if e.g., the price of a coupon bond rises above its redemption value, it must in due course revert to its redemption value).

This example is designed to illustrate the principle that the length of the investment horizon may in principle affect the nature of systematic risk over that horizon, and hence, within a CAPM framework, the cost of capital. Clearly, shares in regulated companies are not bonds. But the same point in logic applies: if the returns are heteroscedastic or serially correlated (or possibly both), then systematic risk at longer horizons for these companies may differ from systematic risk at very short (daily or weekly) horizons. If regulators are concerned with long-horizon risks, then they should ideally be attempting to estimate long-run, rather than short-run values of beta, that take account of these properties.

In Appendix G, Stephen Wright and Donald Robertson examine all of these issues, focussing on data just for the two pure-play UK regulated companies illustrated above.

They point first to strong evidence that the volatilities and correlation structure of returns *do* change over time: that is, returns *are* heteroscedastic. This is indeed a well-established property of most financial returns, particularly in high frequency data. They show that these features are also evident in the dataset used to construct the beta estimates shown above. This implies that there may indeed be a distinction between “short-run beta” (capturing these short-run correlations) and “long-run beta” (which takes into account longer-term correlations).

Furthermore Appendix G presents evidence that the pattern of estimated short-run betas that explicitly take account of heteroscedasticity⁷⁰ quite closely tracks the pattern of rolling betas for the two companies shown in the chart above. Crucially, there is strong historical evidence that short-term shifts in volatility and correlations do *not* persist indefinitely. As a result, Robertson and Wright conclude that the most recent rolling beta estimates are very likely to prove temporary. Furthermore, they conclude that there is evidence that both short- and long-run beta estimates appear to have been quite stable, and that there is therefore a quite strong *prima facie* case to use all available data to estimate, beta, *not* just a relatively short recent sample. This suggests that the most recent beta estimates shown in the chart above are very likely to be over-estimates.

⁷⁰ Using an econometric technique, Generalised Autoregressive Conditional Heteroscedasticity (GARCH) that is very commonly used in academic research, and readily available in econometric software.

A further advantage of using a longer sample of data is that it makes it possible to estimate beta on lower frequency (up to quarterly) data (at which frequency the shifts in volatility and correlations are much less evident) and investigate whether this results in different estimates. Appendix G shows that on quarterly data, using all available data, results using all techniques are (as expected) much more consistent, and point to even lower estimates.⁷¹

The overall conclusion of Appendix G is that estimation on longer-term data suggests that equity beta estimates are likely to lie in the range 0.3-0.5.⁷²

The authors of this report are divided on the conclusions to draw from these results:

Mason, Pickford and Wright's views on beta estimation

MPW take these results as *prime facie* evidence that regulators should consider seriously the implications of lower values of equity beta than they have assumed up until now.

Burns' views on beta estimation

Burns remains unconvinced and provides a reaction to the Robertson and Wright analysis in Appendix G. There are several questions that the regulator needs to address in estimating beta:

- What sample of comparator information should be used? Regulators have to balance the advantages of a larger sample against the relevance of the sample (e.g., unregulated assets, or assets that are regulated by another regulator).
- Over what time period should the betas in the sample be estimated, and should outliers that relate to highly unusual periods be removed from the sample?
- Over what period should the betas be averaged?
- What frequency of data should be used?
- What estimation method should be used?

Answers to these questions have generated a voluminous amount of research on beta that has led to the estimates currently being applied by the regulators. Quite simply, the Robertson and Wright analysis is but one econometric exercise of just two stocks in that vast body of work. MPW's results based on higher frequency data are recognisably similar to the existing regulatory estimates over the relevant time-frames, but MPW also estimate the CAPM on quarterly data, which is the key factor that drives the lower estimates of beta. It should be noted that using quarterly frequency data is unusual in academic studies and is not used by any commercial provider.

The impact of our recommendations 4 and 5 means that the estimation of beta is the only component of the cost of equity where the regulator must use its judgement and discretion. This reinforces the obligation on regulators to examine the evidence as a whole, not to rely solely on outlying estimates, in

⁷¹ Note that this feature may also be indirect evidence of the second feature noted above, namely serial correlation. But the direct evidence on this score is—in contrast to the evidence of heteroscedasticity—not statistically significant.

⁷² Both Appendices F and G also investigate the impact of using a broader measure of global returns as a proxy for the market return. Appendix G also considers the impact of including a role for interest rates, as a simple application of the multifactor approach.

order to retain the benefits of a stable and transparent approach to setting the RAR. As we argued earlier, the weakness of the DDM for regulatory purposes is that any error in estimating the growth rate of dividends translates more or less one-for-one into an equivalent error in estimating the cost of equity. By the same token, a single piece of research, using an unusual data frequency that yields unusual results would not be the basis on which to overturn the practice to date.

That said, the analysis does suggest a worthwhile avenue for future research. As Wright, Mason and Miles (2003) have noted:

“Use of an explicit technique to handle time-varying variance and co-variances is in some respects the ideal solution.....But a major drawback of the technique is that it is susceptible to over-fitting and can reveal apparent signs of time variation where none exist, especially if complicated models of time variation are used. At a more practical level, it involves use of techniques that are highly non-linear and not widely used amongst practitioners who estimate betas. So there would be problem of getting a beta estimated with a time-varying technique to be widely accepted as a standard estimate – this is partly because there are many different ways to model time variation (GARCH, EGARCH, GARCH in mean and many newer variants).”

The time may now be ripe to more comprehensively explore the GARCH approach and its many variants. For regulatory price-setting purposes, more work would be required to establish the robustness of the technique, which would include testing a wider set of variants of the GARCH model than just the one estimated, and using Monte Carlo simulation analysis to test how well each specific GARCH model actually estimates the unconditional betas for each stock under consideration (which would be more than the two estimated in the Robertson and Wright Appendix G). As part of this ongoing research the appropriate frequency of estimation could be further considered. Indeed, it would also be appropriate to estimate alternative functional forms of the model itself – especially since the Robertson and Wright beta estimates increase significantly if the risk-free rate is included in their model.

That research properly undertaken would then form part of the evidence base that regulators could draw upon that would sit alongside pre-existing regulatory approaches. In other words, it would be a complement to pre-existing techniques, rather than a substitute for them.

Until that research is more comprehensive, regulators should continue to use the CAPM on a wide range of comparator stocks, using higher frequency data (subject to testing for thin-trading and serial correlation), over different sample sizes, and interpret that body of evidence judiciously, in line with practice to date. The continuation of this holistic approach is important because it has been an important contributor to regulatory stability which itself has contributed to reducing investor risk premiums over the past 25 years.

On one key point, however, the authors are agreed. The results from both Appendices F and G definitely do *not* point to values close to zero. Past research commissioned by Ofgem has put forward the case that on the basis of *a priori* reasoning the risk profile of cashflows for regulated businesses is almost entirely idiosyncratic.⁷³ But we would argue that, while there may be *a priori* arguments as to whether regulated companies *should* have positive betas, Cochrane’s (2011) “price of tomatoes” argument again

⁷³ See Imrecon/ECA (2012).

applies: the key issue is whether there is evidence that their market prices reflect a market *perception* of systematic risk. The evidence for this is very strong.

Recommendation 6 (Beta Estimation): Regulators should make more use of econometric estimates of equity beta. They should derive these estimates from sound econometric evidence and practice, utilising all available data for relevant listed companies. Betas for unlisted companies should be derived from estimated equity betas from the closest available comparator listed companies.

Given the importance of this issue, we would strongly recommend regulators to devote more resources to this area.

4.6 Leverage (gearing) and its implications for asset beta, the assumed cost of equity, and the CAPM-WACC

4.6.1 Deriving Asset Beta Estimates

In Appendix F, Burns employs standard techniques to “de-gear” equity betas to derive implied values of asset beta from their estimated equity betas (using a tax-adjusted version of the formula in Equation (7)).⁷⁴ This results in values of asset beta that are mostly in the range 0.3-0.4.

Clearly the evidence from Appendix G would appear to suggest lower values. However, some caution needs to be applied here. Both companies examined in Appendix G have gearing of around 50%, but simply dividing the range of equity betas of 0.3-0.5 by two, giving a range of 0.15-0.25 would arguably be too crude.

Tax treatment of debt makes only a modest difference, as in the standard approach applied in Appendix F. But, in common with most applications of this approach, Appendix F simply assumes a debt beta of zero, whereas strictly speaking the asset beta is a weighted average of the equity beta and the debt beta (see Equation (7) and related discussion in Section 3.1).

Here the need for consistency in estimation again applies. On daily data the data do indeed suggest that debt betas are barely different from zero. But the lower end of the range of estimated equity betas in Appendix G is derived from full sample quarterly estimates. At this frequency, in contrast to daily data, the debt beta (derived from iBoxx data as used in Section 4.7) below) appears to be nontrivially different from zero.⁷⁵ Thus the impact of de-gearing “long-run” equity beta estimates could in principle result in smaller proportional downward adjustments to asset beta.

Thus a corollary of any shift towards lower equity beta estimates—particularly if these are based on long-run beta estimates, as recommended by MPW—would be that care would need to be applied in deriving the implied asset beta estimates.

Finally it is worth noting the caveat that, however carried out, “de-gearing” (and indeed “re-gearing” as discussed in the next section) of necessity involves some nontrivial simplifying assumptions. The

⁷⁴ The use of the tax adjusted formula (the Hamada formula) does not imply a preference or recommendation for one form of standard technique over another. We note that some regulators have used the Hamada formula while others have used the version set out in Equation (7).

⁷⁵ We hesitate to quote precise parameter estimates, since we do not have pure return series for the iBoxx data. But changes in yield (which will determine returns) are strongly correlated both with changes in gilt yields (positively, as is to be expected) and negatively with market returns (implying a positive return beta).

standard result in finance that underlies Equation (7), and hence provides the rationale for de-gearing, is that the beta of a portfolio is equal to the market-weighted average of the betas of the elements of that portfolio. However, this result 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. These conditions are very far from holding for regulated companies, but it is not clear how important the simplifying assumptions that are, of necessity, made in deriving asset beta estimates may add additional elements of uncertainty.⁷⁶

Again, we would argue that this area merits further investigation.

4.6.2 “Re-Gearing” Asset Betas, and unlisted companies

Burns’ Appendix F also illustrates the standard technique of “re-gearing” asset betas to derive equity beta estimates consistent with levels of leverage assumed by regulators. In recent years, some regulated companies have traded at a premium to their regulated capital values. Thus leverage is higher, relative to regulated capital values, than it is in comparison with total market value. Appendix F shows that this can result in implied values of “re-gearred” equity betas that are higher than directly estimated values.

The authors do not agree on the validity of “re-gearing”:

Meredith, Pickford and Wright’s views on “re-gearing”.

On the face of it the practice of re-gearing might appear to provide a rationale for the fact that, as noted above, in recent decisions UK regulators have set equity beta assumptions at values well above the range of estimates derived directly from estimation. However, while this approach is standard, we are very sceptical of its validity.

A simple test of the rationale for re-gearing is to apply it to the two quoted stocks, Severn Trent and United Utilities, whose estimated rolling equity betas were illustrated in Figure 4.10 above. In Appendix F, Burns presents estimated values for equity beta for both companies for a range of estimation windows and frequencies. To pick just one figure for illustrative purposes, their Figure G.3 shows that over the most recent 10-year sample of daily data the estimated value for equity beta for both companies is 0.6.⁷⁷ This estimate is then used (in their Figure G.6) to derive an estimate of asset beta, using observed gearing for these companies. Finally, following standard regulatory practice, the asset beta is “re-gearred” using Ofgem’s assumed leverage of 65%. Since, as noted above, both companies had average measured gearing lower than this, in relation to their actual market value, this boosts their notional re-gearred equity betas (shown in Figure G.9) to 0.84 and 0.83 respectively.

MPW would argue that in the case of these two companies this calculation is indeed purely notional. The equity beta for these companies can be, and has been, directly estimated in the data, and it is this value, not the notional value quoted above, that determines their marginal cost of equity capital in the CAPM framework. If, for example, United Utilities needs to raise equity capital to finance an increase in its Regulated Capital Value, it is the equity beta of its quoted shares that will determine the expected

⁷⁶ For example, leverage figures are usually based on the book value of net debt, and the de-gearing is usually based on an assumed value of leverage that is constant over the sample. Pension fund assets also complicate the calculation for at least some regulated companies. And, of course, limited companies are not simply portfolios.

⁷⁷ The analysis of Appendix G would suggest that this is probably at the upper end of the range of plausible estimates, but the precise value chosen is not material to the argument.

return on this new equity, and hence the cost of new equity capital. Thus MPW argue that “re-gearing” does not constitute a valid argument for assuming values of equity beta outside the range of econometric estimates.

Furthermore, MPW argue that since it seems clear that re-gearing assumed asset betas using notional leverage values is inappropriate for listed companies, it is hard to argue that it is an appropriate technique for unlisted companies. Since unlisted companies make up the great majority of UK regulated companies this is a potentially important issue.

We should however lay particular stress on the word “potentially”. Even when equity beta is assumed to be close to, or even equal to unity (and thus out of line with estimated values), this does *not* imply that the degree of systematic risk of the underlying assets has no impact. As discussed above, in Section 3.3, the lower is asset beta, the higher must be the assumed rate of leverage consistent with a value of equity beta close to 1, and hence the higher is the weight of the cost of debt in the WACC. We present strong evidence below that the market cost of debt for regulated companies is very strongly linked to the RFR, so, at least indirectly, as long as the cost of debt is estimated correctly, the lower is asset beta, the stronger is the (indirect) impact of the RFR on the CAPM-WACC, even when it barely enters directly through cost of equity estimates.

So the most crucial issue is to get the asset beta right, which requires careful estimation of equity beta, as set out in Section 4.5, as well as care in adjusting these estimates to derive asset beta, as set out in Section 4.6.1. Indeed, MPW note that an advantage of the estimate of the “pure” CAPM-WACC is that *only* the asset beta matters, so it is entirely unaffected by considerations of “re-gearing”.

However, to the extent that regulators continue to use the CAPM(E)-WACC that only applies CAPM to the cost of equity, the assumed value of equity beta *does* matter; so we would advocate caution in applying re-gearing techniques.

Burns’ views on “re-gearing”

Burns disagrees with MPW’s interpretation of this recommendation. What de-levering and re-levering achieve is to take the actual gearing level of the beta comparators out of the beta estimates and to assess what the equity beta of a company with similar risk would be, at an independently estimated ‘notional gearing’ level. It is therefore not equal to the actual observed equity beta of comparators unless those comparators happen to have an actual gearing level equal to the ‘notional gearing level’.

MPW’s suggestion of directly using the estimated equity betas of comparators cannot be combined with a notional gearing level to calculate the cost of equity. To see this more clearly, suppose Severn Trent and UU happened to have been geared up to well over the notional gearing level (e.g., 90%). By inference from finance theory, the equity beta of these companies would be much higher (possibly higher than 2 or 3)—or to put it another way, the marginal cost of equity within the CAPM framework would be very high. It would be inappropriate for the regulator to use the observed high marginal cost of equity in the cost of equity calculation at the notional gearing level. Equally, suppose the comparators in our example have very little debt. The marginal cost of equity would be very low, and it would be inappropriate to use that equity beta in combination of a notional gearing level of 65% to calculate the cost of equity.

Practically, the effect of the MPW proposal is that the regulator's estimate of the cost of equity becomes conditional on company-specific levels of gearing for those companies which are listed, which creates potentially significant endogeneity problems and increases scope for regulatory gaming. For companies that aren't listed, but where their beta is set by reference to comparator stocks, their allowed cost of equity is directly influenced by another company's capital structure – which may be a peer in the sector, or maybe from another sector or jurisdiction altogether. For those companies, their allowed cost of equity becomes essentially arbitrary.

This issue is currently resolved by the regulator in each sector determining the notional gearing level. The rationale of setting a notional gearing level has become well understood and carries good properties such as encouraging efficient financing by leaving the financing decisions and responsibilities to the companies. Further, we note that the notional gearing level of recent decisions from Ofwat and Ofgem have been informed by a combination of the average actual gearing level in the sector and the reasonable gearing range proposed by credit rating agencies for an investment-grade credit rating.

That being said, the letter of the recommendation is valid. In particular, in situations where there is a material difference between actual and notional gearing, regulators should carefully consider the specific method for re-levering, for example through consideration of the appropriate debt beta to apply to de-levering and re-levering.

A related issue that has been raised is the situation where the actual level of gearing is low relative to the notional level as a result of the valuation of expected future regulatory outperformance (in other words the market value of the regulated business is above the Regulatory Asset Base, RAB). It could be argued that the outperformance element of the value may have a higher beta (for example because regulators are more likely to target the outperformance during times of economic recession). In that situation, the regulator could seek to focus on the underlying beta that relates to the 'core' RAB return and not the outperformance component. This situation is analogous to a decomposition of an equity beta to adjust for non-regulated activities, rather than being a challenge to the standard re-levering approaches. It is noteworthy that the implementation of this decomposition raises some practical challenges. First, estimating the extent to which outperformance value carries greater risk than 'core' RAB value is not straightforward. Second, as we discuss in the transaction analysis (Appendix J) and the discussion of RER there is an element of market value premium that relates to the ongoing cost discovery and innovation on cost and services that is central to an incentive based regime and this element should be treated in the same way as the 'core' RAB.

Recommendation 7 (Adjusting beta estimates for leverage): Regulators should exercise care in allowing for the impact of leverage, in deriving asset beta estimates and in “re-gearing” to derive equity betas based on assumed levels of regulatory gearing. If beta estimates are derived from lower frequency estimation, the treatment of debt betas also needs to be consistent.

4.7 The market cost of debt

The cost of debt (CoD) was not explicitly included in the terms of reference for this study, but we would argue that it is impossible to consider the impact of the assumed cost of equity on CAPM-WACC estimates without addressing the cost of debt, given the interactions described above, and given the high assumed rates of leverage both assumed by regulators and as actually practiced by regulated firms.

The original MMW study gave very little coverage to the cost of debt (just a single page); but here, as with the treatment of the risk-free rate we would argue that this was because not very much actually *needed* to be said. With only modest caveats, this is because the market cost of debt, like the RFR, can be directly observed with a fairly high degree of precision.

We stress that here we focus on the market cost of debt as a component of a market-consistent CAPM-WACC. As such this is an estimate of the *marginal* cost of new borrowing for a regulated company, at a given chosen horizon. (We postpone for now any consideration of the case that, in *implementing* CAPM-WACC in regulation, regulators may wish to allow for the cost of embedded debt – an issue we revert to in later sections.)

4.7.1 A caveat: yields vs expected returns.

The original analysis by MMW, while concise, did make space to note a key caveat on the observability of the CoD. What regulators can observe directly is the market *yield* to maturity on debt of a similar maturity and risk characteristics to that of regulated companies with leverage at the level targeted by regulators. Regulators have indeed made considerable progress over the past decade in using appropriate market benchmark yields⁷⁸ – with the iBoxx yield used by Ofgem in recent price controls representing a very valuable benchmark, which indeed we shall exploit below.

But, as noted by MMW in 2003, the cost of debt is *not* equal to the yield to maturity on that debt. The yield to maturity is the internal rate of return of the bond (and, for a zero-coupon bond, the expected return to maturity), but calculated *on the assumption of no default*. Yet default on corporate debt both happens (indeed the limited liability of corporations is precisely defined for that eventuality) and, crucially (and logically distinct from the true risk of default) is quite clearly *priced* by markets.⁷⁹ As a result the true expected return on corporate debt is strictly less than the yield. (Hence, as a direct implication, the yield *spread* between corporate bonds and risk-free bonds does not represent a true risk premium.)

This point of logic, while inescapable, has typically been observed only in the breach by UK regulators. How much does this neglect of default risk matter?

In Appendix H, Ken Hori discusses an application of techniques derived from recent work by Afik and Benninga (2010, 2014) who derived implied expected returns on US corporate debt from market yields, using the S&P's matrix of default and downgrade probabilities. They show that these calculations resolve the puzzle of an apparently weak relationship between credit ratings and yield spreads; in contrast there is a strong, and plausible link with expected returns.

Appendix H applies the same techniques to UK bond yield data. These results are preliminary, but suggest some reassurance to regulators: while they have been qualitatively wrong in equating yields to expected returns (before any issuance fees), the *quantitative* significance of this error appears, at least on these preliminary results to be quite modest for bonds of comparable rating to those of UK utilities

⁷⁸ Following from earlier work by Baskaya et al (2005) for Smithers and Co Ltd.

⁷⁹ The pricing of the risk of default indeed represents a long-standing empirical puzzle in finance, since observed credit spreads appear hard to explain in terms of any combination of observed frequency of default and plausible risk premia (spreads are typically too high). We note however that the puzzle applies only to “buy and hold” investors. Ilmanen (2011) summarises the differences between the low ex-post credit premia at the index level versus the high ex ante “yield advantage” as being due to “much less visible negatives”. Ng and Phelps (2011) attribute this to forced selling of downgraded bonds, or “Fallen Angels”.

with adjustments of less than 10bp for an A-rated 10-year bond, and only around 20bp for a BBB rated bond.

However, we do not regard this result as definitive, not least because the calculation is somewhat sensitive to the assumptions made,⁸⁰ and can only be carried out at particular points in time (see, for example, the discussion in the next section of the increase in spreads during the financial crisis).

Hence,

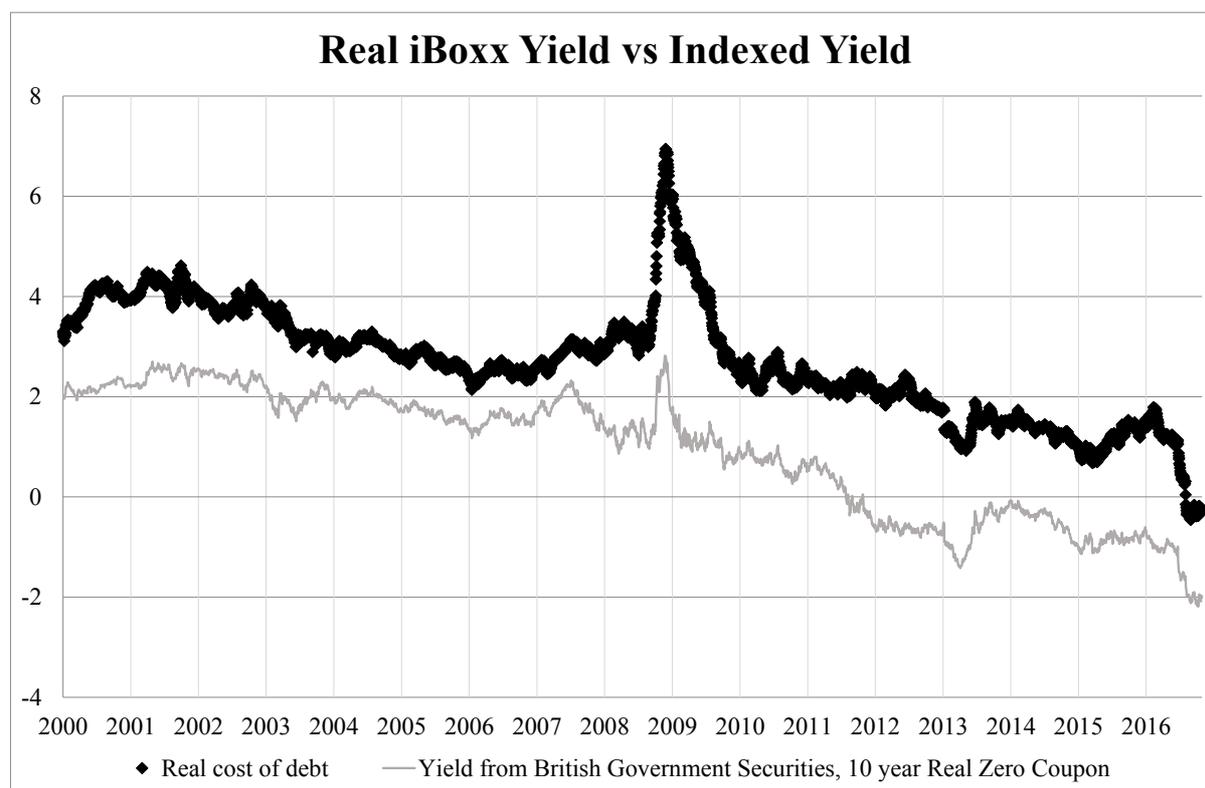
Recommendation 8 (Estimating Default Risk on Corporate Debt): For consistency with the definition of the CAPM-WACC as an expected return, cost of debt estimates should include an adjustment to corporate bond yields to convert these to expected returns. This should ideally take account of documented time variation in default risk.

Burns agrees with the analysis set out above in terms of the theoretical WACC but notes that the adjustment for default risk should not be applied to the RAR. This is discussed in Section 9.

4.7.2 The market cost of debt, the risk-free rate, and the allowed return on debt.

Figure 4.11 shows the recent history of a simple average of two benchmark iBoxx real yields, alongside the real yield on indexed gilts.⁸¹

Figure 4.11



⁸⁰ We note in particular that the calculations assume a buy-and-hold investor; as noted in the previous footnote risk of de-rating may also add to spreads. But here we again face yet another application of John Cochrane’s “price of tomatoes” critique: if markets *perceive* credit risk (in either form) as high, and price it as such, that is what actually matters for the market cost of debt.

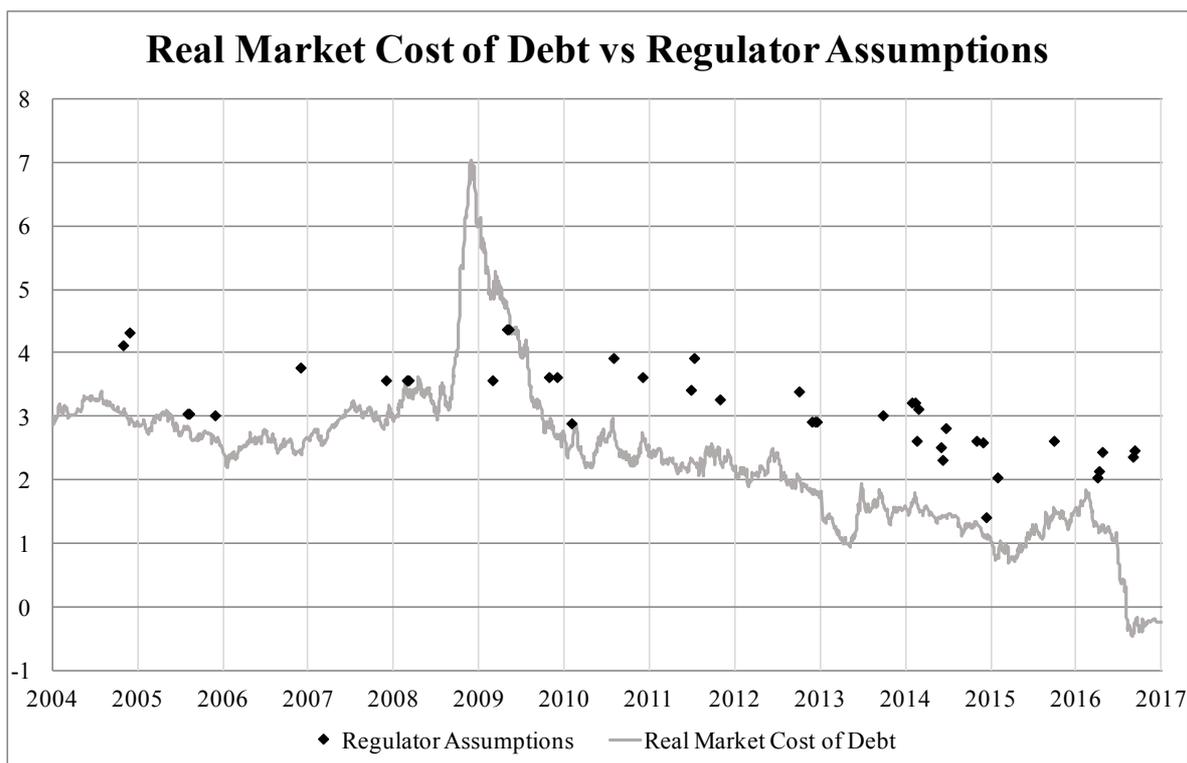
⁸¹ Note that both yields are expressed on an RPI basis. For the adjustment in terms of CPI inflation as applied to the risk-free rate, see Section 4.3

Clearly the dominant feature in recent data has been the fall in the RFR, with only relatively modest shifts in the credit spread. While there was a clear pattern of expansion of the spread in the financial crisis, it then fell back, but to levels distinctly higher than in the years preceding the crisis. It seems fairly clear that the expansion of the spread during the crisis would have reflected some combination of increased risk of default together with, almost certainly a counter-cyclical expansion of the risk premium.⁸²

It is notable that the fact that the spread continues to be higher than in the pre-crisis period (but not in the earlier part of the sample shown) does not appear to reflect any relaxation of the credit risk assessment of the benchmark companies by ratings agencies—indeed we have been advised that these have tightened up compared to the pre-crisis period.⁸³ If anything it would appear to be more consistent with the pre-crisis period having been something of an aberration, in which both perceived credit risk and risk premia reached exceptionally low levels.⁸⁴

While the evidence is very strong that the market cost of debt has been pulled down very strongly by the risk-free rate, Figure 4.12 shows that the same cannot be said for estimates of the notional cost of debt used by regulators in setting the allowed WACC, which (as with the RFR) have adjusted distinctly more slowly.

Figure 4.12



However, while the difference is striking, there is crucial difference in terminology. Our definition of the cost of debt is a market concept, defined in terms of an expected return, and thus is consistent with the market concept of the WACC. In contrast the value often *referred* to as the cost of debt by regulators

⁸² Smithers and Wright (2013) discuss the academic evidence pointing to counter-cyclical risk premia on bonds, which is broadly consistent with the evidence on the equity risk premium.

⁸³ Advice provided by in conversation with James Wardlaw.

⁸⁴ A feature that was indeed commented on at the time in the analysis by Baskaya et al (op cit), who also compared the relatively short history of UK credit spreads with the much longer available dataset from the USA.

is a quite distinct concept: it is the *allowed return* on debt (RAR), consistent with concept of the RAR introduced earlier.

Figure 4.12 shows that the gap between the RAR and the market cost of debt has become increasingly striking in recent data. We argue below that this gap has been a crucial determinant of bid premia.

Recommendation 9 (The Cost of Debt): The term “cost of debt” should relate to the expected return (correctly adjusted for default risk) on a traded corporate bond, at the regulator’s chosen horizon, and with risk comparable to regulated utilities at regulators’ chosen leverage. It should be clearly distinguished from the *allowed return* on embedded debt.

4.8 Alternative historical estimates of the market-consistent CAPM-WACC for regulated utilities

Figure 4.13 provides an illustration of alternative approaches to estimating the CAPM-WACC, given actual data on the real risk-free rate and the real cost of debt (both at a horizon of approximately 10 years⁸⁵). It shows that the “allowed WACC” set by regulators (in our terminology, the RAR) has routinely exceeded our two alternative estimates of the CAPM-WACC:

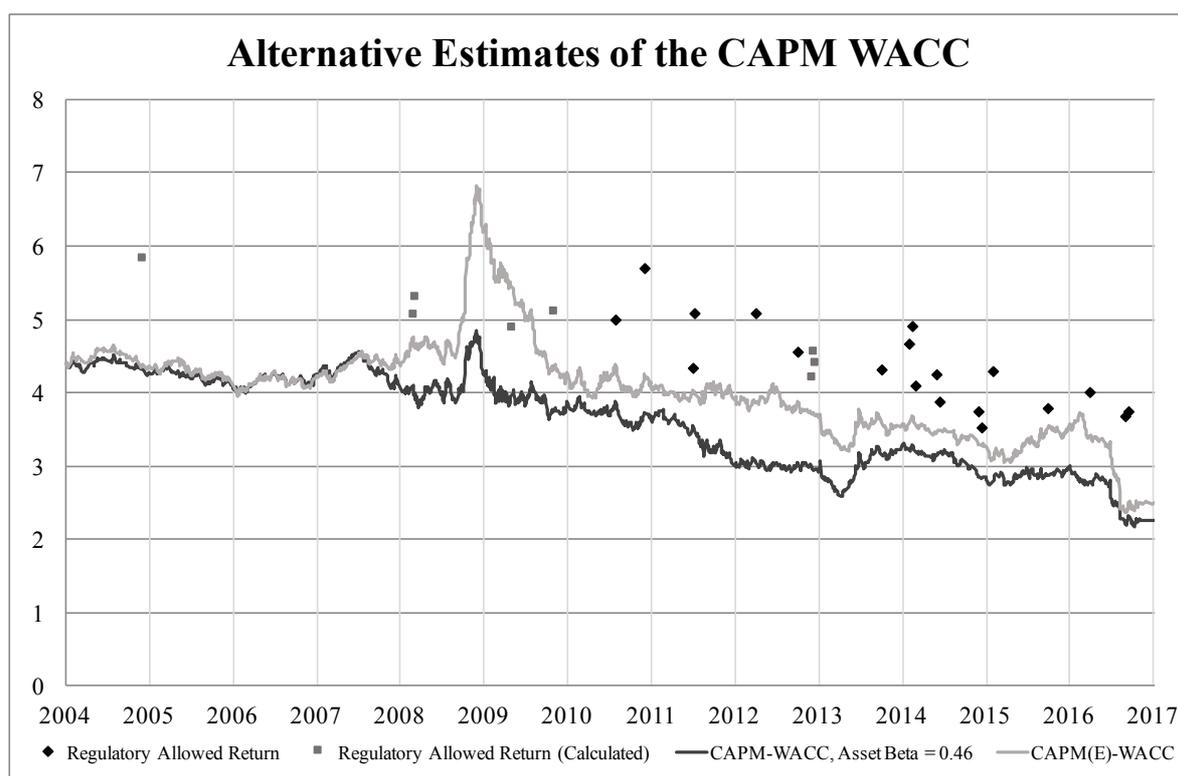
- The “pure” CAPM-WACC (i.e., assuming CAPM applies to both equity and debt markets, as defined in Section 3.2)
- The CAPM(E)-WACC (i.e., assuming CAPM only applies to equity markets, hence - in this respect at least - as typically calculated by regulators).

To accentuate the comparison, we assume that the assumed value of leverage is set such that the resulting value of equity beta is *precisely* equal to unity, so that the resulting figure is simply a weighted average of the EMR and the CoD.

Note that, crucially, both estimates of the CAPM-WACC shown in the chart are calculated using the average value (0.46) of the asset beta assumptions that were fed into the regulatory decisions shown. So the differences shown do *not*, for example, reflect the impact of assuming the lower values of beta that would be implied by Appendix G.

⁸⁵ The approximation relates to the iBoxx estimate of the cost of debt which is for maturities of 10 years or greater. We are not clear as yet what the average maturity is; however, at such long maturities the impact on the resulting figure would be quite small since the term structure has only limited slope at this point.

Figure 4.13



We stress that the estimates we show above are both preliminary and illustrative. But they are deliberately designed to illustrate a methodology that satisfies our twin aims of defensibility and implementability. While the resulting estimates vary over time, the methodology does not.

Several features of the chart are clearly worthy of note:

- As we would expect, both estimates vary markedly more than the Regulatory Allowed Return (previously “allowed WACC”) set by regulators. We should not be surprised by this. The CAPM-WACC is a market-consistent concept, which we would, therefore, *expect* to be relatively volatile). The Regulatory Allowed Return is, both in theory and in practice, quite distinct, even though in principle they must be very closely related concepts. UK regulators have historically preferred the RAR to be stable. An issue on which the authors of this report cannot agree is whether this stability is desirable.
- It is also striking how similar the two different estimates of the CAPM-WACC are, despite their quite different methods of construction. We take this as indirect evidence that most of the time, at least, the CAPM is telling a similar story about the cost of debt as it is about the cost of equity.
- To the extent that at certain points the two estimates diverge (most strikingly during the financial crisis) we suspect that the chart overstates the true divergence, since the expansion of corporate yields almost certainly reflected sharply increased market perceptions of default; thus the true cost of debt during this period would be lower, and with it the CAPM(E)-WACC.

Recommendation 10 (The CAPM-WACC Methodology): The term “WACC” should be restricted to the concept of an expected market return on capital of a given degree of systematic risk. It should *not* be used to refer to an allowed return. Our use of the term “CAPM-WACC” is

deliberately designed to emphasise the necessary link between the CAPM model and a methodology for estimating the WACC.

While setting out a methodology for estimating a market-consistent estimate of the CAPM-WACC is we argue, a key part of the regulatory process, it is clearly only a part. We have not yet addressed the issue of how such estimates should be transmitted, via the regulatory process, to the Regulatory Expected Return. To this issue we now turn.

5 Implementing the CAPM-WACC in regulation: the Regulatory Expected Return (RER) vs the Regulatory Allowed Return (RAR)

Thus far in this report we have restricted ourselves to the relatively narrow task of how to estimate a market-consistent estimate of the CAPM-WACC. The regulatory aspect of how these estimates has only entered into our deliberations to the extent that we have argued for departures from a purist approach to this task with reference to our desired twin properties of implementability and defensibility.

But we are very conscious that the mapping from this narrow (if crucial) task of estimating the CAPM-WACC to the actual act of regulating companies is not straightforward. We need to look systematically at the *way* such estimates of CAPM-WACC are used.

We start by defining some key concepts.

5.1 The Regulatory Expected Return (RER) and the Regulatory Allowed Return (RAR): some definitions

In the introduction to this report we have already raised the possibility that, even given our preferred methodology for estimating of the CAPM-WACC as set out above, the expected return that regulators allow on the assets of regulated utilities may in principle justifiably deviate from the CAPM-WACC.

At the cost of introducing new acronyms into an area already awash with such terms, we believe it is worth explicitly introducing the two distinct concepts of the Regulatory Expected Return (RER) and the Regulatory Allowed Return (RAR) to differentiate them from the CAPM-WACC on which we have been focussing thus far.

5.1.1 The Regulatory Expected Return (RER)

We define the RER to be the *expected* return on the assets of regulated utilities. For comparability with the CAPM-WACC we assume this to be the expected return on the total regulated asset base.

While intimately linked, the RER and the CAPM-WACC are in principle quite distinct, although in regulatory practice they have often been treated as if they were identical (MPW would argue, at the cost of considerable confusion).

Thus, to clarify,

- The RER is *not* necessarily equal to the CAPM-WACC. Clearly this expected return must be informed by the CAPM-WACC but it need not be (and in practice may often not be) equal to it at any particular point in time.

- The RER, as an expected value, is clearly an *ex-ante* concept. Conceptually it is the return expected by investors in regulated firms, which cannot be directly observed. However, as a practical concept the RER can be defined conditional upon the information set of the regulator, i.e., taking account of any statistical evidence of predictability in past outperformance.⁸⁶

5.1.2 The Regulatory Allowed Return (RAR)

Here we introduce a new term for an established concept: the Regulatory Allowed Return corresponds precisely to what has in the past typically been referred to as the “allowed WACC”.

To change terminology in this way may on the face of it seem unnecessary, but we believe that terminology matters a great deal, particularly since the concept of the WACC has such a different meaning in the corporate finance literature.

Precision also matters. So here we work in terms of formulae:

$$RAR = g RARD + (1-g) RARE \quad (20)$$

where

- g is the (notional) leverage (gearing) in the regulatory asset base;
- $RARD$ is the allowed return on the notional debt component;
- $RARE$ is the allowed return on notional equity *conditional on hitting targets* set by the regulator.

Thus clearly the RAR is an *ex-ante* concept. We therefore require a corresponding ex-post concept. We refer to this simply as the Regulatory Return (RR). We define this by

$$RR = g RARD + (1-g) (RARE + X) \quad (21)$$

where $RARD$ is set in advance, and X is the outperformance incentive component in the *ex-post* return on regulatory equity⁸⁷

5.1.3 The RER defined in terms of the RAR

It follows directly from the above definitions that we have

$$\begin{aligned} RER &= E(RR) && (22) \\ &= g RARD + (1-g) (RARE + E(X)) \\ &= RAR + (1-g) E(X) \end{aligned}$$

Thus, straightforwardly, the wedge between the two is driven by expected outperformance.

A key feature of the above definition is that any element of outperformance accrues only to equity holders. Thus a rise of one percentage point in the RER increases the expected return to equity holders by $1/(1-g)$ percentage points.

⁸⁶If we assume that the regulator’s information set is strictly dominated by the firm’s (which seems plausible) then the Law of Iterated Expectations applies: the regulator’s expected return is also equal to the regulator’s best estimate of the firm’s *expected* return.

⁸⁷ In terms of current terminology, therefore, $RARE + X = RoRE$. We do not object to the continuation of this terminology, although for consistency and compactness we would lose the lower case “o” and define it as RRE .

5.1.4 Why bother?

We believe that the distinctions between the CAPM-WACC, the RER and the RAR both help to explain some aspects of regulatory behaviour as implemented in the UK since privatisation, but also allows us to write down what the different authors argue are desirable features of the relationship between the three concepts.

Perhaps understandably, these distinctions also allow us to pin down differences between the perspectives of different authors of the report. We can all agree on the definitions; where we disagree is on what to do with them. We set out our different views on this in Sections 8 and 9 below.

6 Attempting to understand recent market valuations of regulated utilities (Burns)

Recently, positive private transaction premia and stock market premia observed on regulated businesses have been used to suggest that allowed equity returns may have been too high. PwC for example, in its 2017 report on Ofwat's incentives at PR19, lists "inferred cost of equity from RCV (RAB) premia" as one of five potential approaches to estimating equity market returns.

Market premia (or MAR: market to asset ratio) can be calculated for publically listed businesses, by comparing their enterprise value to their RAB. Of GB regulated businesses which are listed, only United Utilities and Severn Trent do not have significant unregulated activities, and are can therefore be used to estimate MARs. Calculating the MARs for United Utilities and Severn Trent as of 31 March 2016 gives 28% and 27% respectively.

Larger premia have been observed in some private transactions, such as the sale of a majority stake the National Grid Gas Distribution businesses in March 2017. Estimates of the transaction premium include 49% based on National Grid's estimation of the implied enterprise value of the business and its estimated RAB for 31 March 2017.

However, we would caution the direct jump into conclusion that the existence of large premia must have been caused by the allowed equity return being set too high. There are a large number of different potential drivers of these premia, of which a divergence between the allowed and actual cost of equity is only one. We summarise the main potential sources of transaction premia below.

- Cost outperformance. The Totex Incentive Mechanism in energy network regulation allows operators to retain a given proportion of any underspend they achieve against their totex allowances. The remainder is passed on to customers. Equivalently, any overspend must be partially borne by the operator, and partially by customers. The sharing proportion typically ranges between 40% and 60%. Companies which underspend therefore experience an increase in their return on regulatory equity (RoRE). Similar cost sharing mechanisms are applied by other regulators.
- Output (Outcome) incentive outperformance. Ofgem and Ofwat also have in place incentive mechanisms, which reward or penalise licensees based on various measures of performance. Outperformance on these measures also creates uplift in RoRE.
- Cost of debt outperformance. Ofgem allows energy networks to earn a return on their debt based on the yield from iBoxx indices. If networks can raise debt at a lower rate than this, they retain the difference. This outperformance can also be expressed in RoRE terms. Ofwat allows

water companies a fixed cost of debt upfront, and as the debt market continued to move downward during the last price control, water companies were able to incur debt cost lower than that set by Ofwat.

- Difference between allowed and actual cost of equity. If a regulator determines the allowed equity return based on a long-term measure of market returns, it is inevitable that in the short run there will be deviations from the actual cost of equity.
- Additional value in financial restructuring. There is the potential to restructure the financing of these businesses and increase the gearing level significantly beyond the notional gearing of 65%, up to for example 85%. This can be done by retaining a 65% gearing for the entities within the regulatory ring fence, but by adding a further 20% debt into a holding company outside the ring fence. Potential buyers may believe that they can do this without jeopardizing its credit rating, and combining with a higher allowed debt cost than actual, this can create extra value.
- Tax arbitrage. Investors may be able to out-perform the allowance for corporation tax included by the regulator. This could be achieved through increasing the proportion of debt financing and / or taking advantage of opportunities for arbitrage across different tax jurisdictions (in accordance with transfer pricing rules). Note that regulators have policies for sharing some types of tax out-performance with customers through clawback mechanisms.
- Control premium. Investors are willing to pay a premium to acquire a majority stake in a business. This may be because there are perceived value in holding some real options on taking future strategic decisions, which is absent when holding a minority share in the business.
- Winner's curse. As suggested in classic auction theory, in a private transaction involving many bidders where the winning bid has the highest valuation, it often holds that the winners tend to be the ones who have overvalued the asset and end up with a loss when the true value of the asset is revealed.

Appendix J sets out details our quantitative analysis on how the various components can contribute to transaction premium, based on an example energy network sector—gas distribution.

To summarise Burns' conclusions:

- One of the major contributing factors for the transaction premium is the potential cost outperformance and incentives rewards that investors could assume to achieve going forward. Based on investors expecting current rates of outperformance to continue in the gas distribution sector, this can explain a RAB premium of more than 20%. We note that the current rates of outperformance in this sector are relatively high, for example higher than being achieved by water companies. Expecting these rates to persist could be considered an optimistic scenario although we note that some degree of expected cost and incentive outperformance is inherent to incentive regulatory regimes that encourage cost discovery and innovation.
- Another major contributor could have been expected outperformance on the cost of debt allowance, based on debt refinancing prior to the transaction. Our analysis suggests that for the gas distribution sector, this could have accounted a RAB premium of more than 10%. The additional premium created by the potential of financial restructuring can add another 3% - 4% on top of that.
- The above factors together could account for close to 40% RAB premium for an optimistic investor (and it is the most optimistic bidder who ultimately wins the transaction).
- There is potential for outperformance on the actual cost of equity relative the allowed return. The scale of the outperformance depends on the current differential between actual and allowed and how long it is expected to persist before the gap is closed by the regulator. Given the recent

trends in allowed equity returns by regulators it would be an optimistic investor that expected any gap to persist for a long period (although as note above it is the most optimistic bidder who ultimately wins the transaction).

Different drivers of outperformance are at play and multiple combinations of various drivers can explain observed premia. In addition, the role of expected outperformance means that the premia may result from unobserved investor assumptions that may be considered unrealistic or optimistic but are nevertheless the reality behind the premia. For these reasons we consider that evidence from transaction premia is less reliable and much harder to interpret than other sources of evidence on the cost of equity.⁸⁸

MPW draw somewhat different conclusions from this analysis, as discussed further below.

7 Areas of Agreement and Disagreement

Our key areas of agreement are summarised in Recommendations 1 to 10; where we disagree on particular issues we have highlighted these disagreements. The authors are also broadly in agreement on the conceptual framework set out in Section 5, and the key conclusions to be drawn from Section 6.

Our key areas of disagreement are:

- Have regulated utilities been over-rewarded for achieving efficiency gains?
- Can outperformance be predicted, and hence factored into the RAR and the RER?
- Is it desirable to reward outperformance at a return systematically above the CAPM-WACC?
- Is there a case for gradualism in the link between the CAPM-WACC, the RAR and the RER?

The final two sections of the report set out these views in separately authored sections.

8 Using CAPM-WACC estimates for regulation (Mason, Pickford and Wright)

8.1 A simple conceptual framework

Regulated companies have undoubtedly achieved efficiency gains, to the benefit of the consumer. But in our judgement, the returns that they have earned for achieving these gains have not been commensurate with the risks they face, or with the incentives required to persuade them to invest. The magnitude of bid premia suggests strongly that markets agree, and furthermore that they expect this situation to persist.

We would argue that this situation could have been at least partially avoided if the link between the RER and the estimated CAPM-WACC had been made both more explicit, and tighter.

We start by setting out a conceptual framework which allows us address this link explicitly. We can use this framework in several ways.

First, we argue that this framework provides important insights into the implementation of *any* regulatory system, including the current system.

⁸⁸ As a final caveat, note again that this analysis has been carried out for an example sector: the gas distribution sector. Findings and conclusions may be different for different sectors.

Second, in Section 8.2, we show how the arguments for “aiming up” fit within this framework.

Third, in Section 8.3, we then make the case for using the conceptual framework as the basis for what we would regard as an ideal, and extremely simple framework for setting the RAR, which we argue is both implementable and defensible.

Fourth (because we are realists), in Section 8.5, we consider possible compromises to this ideal framework; but in which the insights of our conceptual framework can still be applied.

We first write down the conceptual framework as follows:

$$\begin{aligned} RAR &= CAPM-WACC + W_R \\ RER &= RAR + W_I \end{aligned} \tag{23}$$

Note that this framework is at this stage simply descriptive.⁸⁹

The first equation decomposes the Regulatory Allowed Return (previously “allowed WACC”) into two components, the estimated CAPM-WACC and the “regulatory wedge”, W_R . This decomposition can represent any regulatory framework, including the current system. If, for example, (as in the current system) the RAR is set at a relatively stable value, then shorter-term changes in the estimated CAPM-WACC are absorbed by the regulatory wedge. Figure 4.13 implies that W_R has almost invariably been substantially positive in recent years.

The second equation carries out a decomposition of the Regulatory Expected Return into two components: the RAR and the “informational wedge”, W_I , which captures expected outperformance. Burns’ analysis of bid premia suggests that in the recent past market perceptions are that W_I has also been strongly positive. In Appendix K we argue that, on the limited available data on realised returns (based on electricity distribution companies over the price control period 2010-15), this expectation appears broadly realistic, given strong evidence of systematic outperformance in the past.

The framework above straightforwardly implies

$$RER = CAPM-WACC + W_R + W_I \tag{24}$$

As discussed in the introduction, the magnitude of recent bid premia implied that the gap between the RER and the true WACC perceived by investors in regulated companies must be very substantial. Taking the CAPM-WACC as our best estimate of the true WACC, this implies that the *sum* of the two wedges, $W_R + W_I$, is currently perceived to be very substantial.

This is where our normative arguments come in. In the next section we shall argue that, viewed as a target of regulatory policy, the sum of the two wedges, $W_R + W_I$, which we refer to as the “aiming-up wedge” *should* be positive; but that it is implausible that it should be anywhere near as high as is implied by bid premia.

We then go on to argue that the informational advantage firms possess over regulators will almost certainly always result in a positive “informational wedge”, W_I between the RAR and the RER, if regulators wish to incentivise cost efficiency.

⁸⁹ Mathematically we can at this stage think of these equations as implicit *definitions* of the two wedges.

Two points stand out:

1. There has been considerable variation across regulators as to where in the range the RAR is set. The range goes from the 2014 Ofgem decision, which chose a RAR of 3.80 from a range 3.79-4.21; to the 2014 CC decision, which chose a RAR at the top of the range (a value of 4.10 from a range 3.30-4.10).
2. For any particular regulator, there is considerable variation across decisions as to where in the range the RAR is set. For example, across 10 Ofgem decisions in our database, the degree of aiming up has varied from very close to the lowest point in the range to 83% up the range.

While clearly contexts and situations change and are different across time and regulators, variations of this magnitude are difficult to rationalise or justify. It is also worth noting that during the period illustrated in Figure 8.1, most regulated companies outperformed the RAR, so that implied values of the RER were almost certainly higher than shown in the chart.

In Appendix I, we consider how the RAR, and hence the RER, should be chosen within the estimated range for the (CAPM-)WACC. This matters, since the range is often large; the difference between e.g., the 67th percentile and the 95th percentile has been as much as 0.84%.

Considerable regulatory, and some academic, attention has been paid over the years to this question. The most comprehensive assessments can be found in:

- The New Zealand Commerce Commission Amendment (2014).
- The associated Oxera study, and its responses to various submissions on that report. The latter can be found at <http://www.comcom.govt.nz/regulated-industries/input-methodologies-2/further-work-on-wacc/>.
- The academic article by Ian Dobbs, a simplified version of which formed the basis of Oxera's analysis for the New Zealand Commerce Commission. See Dobbs (2011).

The basic trade-off considered by all these regulators and studies is reasonably well understood (although not always well explained) by now. The heart of the incentive problem facing regulators is to balance the need to further the consumer objective and avoid excessive prices for consumers, with the need to ensure that regulated companies can finance the proper carrying out of their functions. (Associated with these are the risks of inefficiently high investment, as regulated firms look to increase the size of the regulated asset base; and inefficiently low investment, respectively.) The trade-off was summed up nicely by the Competition Commission in its 2007 review of airports:

"If the WACC⁹¹ is set too high then the airports' shareholders will be over-rewarded and customers will pay more than they should. However, we consider it a necessary cost to airport users of ensuring that there are sufficient incentives for BAA to invest, because if the WACC is set too low, there may be underinvestment from BAA or potentially costly financial distress... Most importantly, we note that it is difficult for a regulator to reduce the risks of underinvestment within a regulatory period."

Similarly, the New Zealand Commerce Commission (2014) states (paragraph X17, *ibid*):

"In our view, it is appropriate to use a WACC significantly above the mid-point estimate for price-quality path regulation. This is because the potential costs of under-investment from a WACC that is

⁹¹ We quote verbatim, but note that we would interpret the term WACC used here, and in the extract that follows, as referring to the RAR. Indeed, as we have argued above, it is a logical impossibility for regulators to *set* the WACC, which is determined by markets. We also note the implicit assumption in both quotations that there is no outperformance i.e., that the RER and the RAR are equal.

too low are likely to outweigh the harm to consumers (including any over-investment) arising from a WACC that is too high.”

Despite the long-standing recognition of the issue, there has been a lack of clear exposition of the trade-off involved. Dobbs does develop an analytical framework for the trade-off; but with a number of moving parts to his analysis, it is not altogether clear which aspects are key. (Moreover, Dobbs' approach requires demand to be elastic, which is at odds with empirical evidence for a number of regulated sectors.)

In Appendix I, we lay out a pared-down framework to assess the trade-off. We take a top-down approach and focus on the optimal value of the “aiming-up wedge” $W_I + W_R$, which determines the gap between the RER and the WACC, without consideration of the two individual components (an issue we revert to in Section 8.3).

For new investment (i.e., investment that has not yet taken place), we make the trade-off as stark as possible, assuming that setting the RAR in such a way that the RER is below the true WACC entails an entire loss of surplus. We show that, under a range of parameterisations, the optimal extent of aiming up for new investment is large. For values of the price elasticity of demand in line with those observed empirically for regulated firms (i.e., less than 1), the optimal RAR routinely lies above the 90th percentile; this holds for different function form assumptions for demand, and for different parameters. This is hardly surprising: with relatively low elasticities, the reduction in consumer surplus from setting the RAR, and hence the regulated price, too high is relatively small. In contrast, the welfare loss from setting the RAR (and hence the price) too low is relatively large. This leads to considerable aiming up, as the optimal choice by the regulator.

In contrast, once an investment is sunk, there is (obviously) no risk that investment will not occur if the RAR is set too low. At the same time, regulators are required to ensure financeability of regulated firms. Balancing these two considerations, it is clear that the optimal RER for old (sunk) investment is therefore the expected WACC i.e., the mid-point of the estimated range.

If it is desirable for the RAR to be set such that there is a wedge between the RER and our best estimate of the WACC, how large should this “aiming up wedge” be? We do not have a precise answer; but Appendix I does give some pointers. It argues that if the regulator was only interested in incentivising new investment, for a range of different assumptions on the nature of demand for the regulated good, the RAR would be set at a value such that the RER was above the 90th percentile of the regulator's range of estimates of the true WACC. However, it also argues that for sunk investment future financeability simply requires that existing capital earns the expected WACC— of which our best estimate is the midpoint of the range. So the target value of the RER should reflect the balance between new and sunk investment.

Over any price review period, new investment is relatively small compared to the regulatory asset base. For example, for gas transmission, National Grid Gas (NGGT) during RIIO-T1 (2013-2021) received £1,241 million for baseline load-related and non-load-related capex, with the RAV equal to £4,838m at the start of the period, rising to £6,346m by the end of the period. So, new capex is somewhere between 20-26% of the RAV. If the appropriate percentile on new investment is 90%, say, then the overall percentile choice for the targetted RER should be 52-58%. This is much less than has been the case in past regulatory decision: Ofgem, for example, set the RAR at 83% up the range of its WACC estimates, during a period in which there was very strong evidence of outperformance, so that the implied RER was higher even than this value.

In other words, although we make a case for aiming up, that case is, in our view, a limited one: more limited than appears to have been adopted in a number of past regulatory decisions in the UK. The case is limited further by the extent to which regulators are able to incentivise investment through means other than setting the RAR.

While we do not have a precise estimate of how large the aiming up wedge should be, we can offer one perspective that puts this quantification problem into some perspective. A clear corollary of aiming up is that we would *expect* it to result in some degree of positive bid premia at least on average, since, as discussed in the Introduction, these simply reflect the gap between the RER and the true WACC. But it would be very hard indeed to argue that the magnitude of recent bid premia could be defended on grounds of aiming up alone.

A simple quantification illustrates. In the Introduction, we noted that recent bid premia imply that the expected return on regulatory equity must in perpetuity-equivalent terms be roughly twice the true market cost of equity. To simplify, assume an equity-only firm. Figure 4.13 shows that our best estimate of the “pure” CAPM-WACC in the recent past, which would be the cost of equity for an equity-only firm, would be around 2 ½%, based on regulators’ assumptions on asset beta. The arguments set out above suggest that it would be hard to justify an aiming up wedge that implied an RER above the 60th percentile of our range of WACC estimates. Thus for bid premia to be justified by aiming up, the 60% percentile for our range of relevant WACC estimates (which in this example would simply be the cost of equity) would need to be roughly 5%. By implication there would need to be a 40% probability that the true cost of equity in the recent past was *above* 5%. This seems quite a stretch, to say the least.

8.3 A simple implementable and defensible framework for setting the RAR and the RER.

We have argued, in the context of the conceptual framework set out in Section 8.1, that there is a clear case for “aiming up”; but we have also argued that these arguments cannot be used to justify recent bid premia. We also acknowledge that the informational advantage regulated firms possess will almost certainly result in a positive value of the “informational wedge”, W_I , the gap between the RAR and the RER.

We now argue that there is a strong case for using the framework of Section 8.1 as *more* than just a descriptive framework. We use it to provide us with a set of explicit recommendations.

Recommendation MPW 1: Regulators should set explicit numerical target values for both W_R and W_I , such that the *sum* of the two wedges should be equal to the desired value of the “aiming up” wedge. These values would be periodically revisited at low frequency (probably in light of information emerging over the course of a full price control period), but they would be constant at higher frequency.

Regulators would clearly need to take a view on the values of the two individual wedges. This is clearly novel, but we would argue that it is not an insuperable problem. Clearly also, if, as we would prefer regulators take an explicitly top-down approach by first setting their target value of the sum of the two, the “aiming-up” wedge, then they only need to set an explicit target for either W_R or W_I , but not both.

We argue that on grounds of accountability and statutory obligations to the consumer there is a strong case for setting a target value for the informational wedge. Recall that, crucially, this should represent the regulator’s best estimate of the impact of future outperformance on regulatory returns. Clearly for

this, regulators would need data on past outperformance. While we have noted that current availability of data on outperformance is very patchy, that does not mean that the data do not exist. Hence,

Recommendation MPW 2: Regulators should assemble a systematic and comprehensive database of historic outperformance, to enable them to make their best-informed forecast of the “informational wedge”, W_I , the gap between the RER and the RAR.

There are clear benefits from such an explicit target in terms of accountability. Of course, like all forecasts, regulators will get it wrong. Indeed, a crucial aspect of incentive-based regulation is that the most efficient firms *should* on average earn a higher return than the RER, since this would be set for all firms within a regulator’s remit. But less efficient firms should on average earn *less* than the RER.

But by setting an explicit value for the informational wedge W_I , and hence for the RER, if over the course of time it emerges that regulators have been making *systematic* forecasting errors (i.e., errors that could have been predicted) for the *average* company, then this would represent a failure of their statutory obligations, for which they should be held accountable.

It would be tempting to propose a particularly simple version of our framework, in which regulators should set an explicit value of zero for W_I , implying RER=RAR. (In our first draft of this report, we proposed just that.) However, we are now persuaded that such a zero target would be neither realistic, nor indeed necessarily desirable. By setting an explicit positive value of W_I , regulators would in effect say to regulated companies: *we expect* the average company to outperform, and hence to earn more than the allowed return. It might indeed be desirable to present the RAR as effectively a lower bound for returns; but regulators *would* need to explicitly state the extent of outperformance they were expecting.

What does this imply for the value of the “regulatory wedge”, W_R ? We envisage a top-down approach. There should be an explicit target for the “aiming-up wedge”, which is the sum of the two wedges. So, once regulators set an explicit target for the informational wedge W_I , W_R cannot be set independently. It is in principle logically possible that this might result in a value of W_R equal to, or even less than zero. We do not see this as an obvious flaw in this approach as long as there is an explicit rationale.

Given our definitions of the two wedges as set out in Equation (22), our next recommendation follows directly, as a corollary of MPW 1.

Recommendation MPW 3: Given stable target values of the regulatory and informational wedges, W_R and W_I , at higher frequency the RAR should simply move in line with the best estimate of the CAPM-WACC.

This proposal may well seem naive in its simplicity. Since we are realists, we do not expect it to be adopted, so in Section 8.5 below, we discuss possible compromises below that might make it more palatable in the current regulatory framework. But we shall argue that, even in a system with such compromises, a version of Recommendation MPW 1, and hence MPW 3, should still apply: we would argue that regulators should *still* set target values for the two individual wedges, since even after making these compromises equations these target values should at least be hit *on average*; and Recommendation MPW 2 would also need to be followed in order to make this feasible.

What estimate of the CAPM-WACC should be used in our idealised framework as set out above? On this issue, we are in principle somewhat more flexible: the only key feature of Recommendation MPW 3 is that it should be our *best* estimate.

We would however argue that our estimate of the “pure” (or equity-only) CAPM-WACC has the merits both of a sound theoretical basis (as an upper bound for the true WACC) and of greater simplicity, thus being particularly suited to the twin aims of implementability and defensibility set out earlier in the report. Hence:

Recommendation MPW 4: Regulators should use the “pure” CAPM-WACC to set the RAR.

In terms of direct implementation as an RAR this would therefore imply

$$RAR = RFR + \beta_A (EMR - RFR) + W_R \quad (25)$$

As discussed above, within any given price control period we would not typically expect the regulatory wedge W_R to change. Given Recommendation 5 in the main report, the estimate of the EMR would also be held constant through the price control period. Thus for most of the time the RAR would move up and down with the risk-free rate, with weight $1-\beta_A$.⁹²

We argue that a clear advantage of this approach is that it would reduce considerably the fluctuations in bid premia observed in recent years. We believe that at least a significant proportion of these fluctuations can be attributed to the slow adjustment of allowed returns, at a time when both the cost of equity capital and the cost of debt were falling sharply.

8.4 Some possible objections to our ideal framework

CAPM-WACC is too volatile.

Our first response to this is that it is not *that* volatile, and is distinctly less volatile than the CAPM(E)-WACC for some regulated companies, since the latter uses (by necessity) the market yield on corporate debt as a key component with a high weight (given the high level of gearing assumed for some regulated companies). This is hit by additional shocks which almost certainly reflect cyclical movements in credit risk, and which therefore serve to mask the true value of the cost of debt, and hence of the true CAPM(E)-WACC.

Our second response is: do regulators care more about volatility of the RAR or about the volatility of market valuations? We would argue that the very fact that this report has been commissioned suggests strongly that they do worry about the latter; and a key feature of our proposed regime is that the increase in volatility of the RAR, compared to the current regime, would have a counterpart in *reduced* volatility of market valuations.

Here it is helpful to draw an analogy with bond pricing.

Consider first the case of a 20-year bond with fixed coupon, with coupon equal to initial market yield (so initially at par). If the market yield changes shortly after issue the price of the bond will rise by roughly 5% for every point fall in the yield.

Now consider, second, a notional traded 20-year bond that paid a coupon that was an average of market yields over the past five years. If market yields were constant, then the price of the bond would still fluctuate (but deterministically), according to the ratio of the moving average to the current market yield. But since in practice the price of this bond would also move in response to market yield changes,

⁹² Note that these variations need not necessarily feed into consumer prices at high frequency, if short term movements were credited to/debited from the reserve asset base. See further discussion below.

the inescapable conclusion is that the second bond would be more volatile than the first: yielding capital gains and losses to the issuer of the bond. The source of this additional volatility would be the “dragging anchor” applied to the coupon.

Now consider a third bond as in the second case, but in which the process determining the coupon from past yields was subject to change by some outside authority. This bond would be subject to even more volatility as markets attempted to second guess what this outside authority would do next. (We hope that the analogy here with regulatory discretion is evident.)

Fourth, now consider the (strictly hypothetical) case where the market yield at a 20-year maturity cannot be observed, but can only be inferred indirectly from other evidence. This would be a yet further source of volatility. Here the clear analogy is with the unobservability of the market return.

Finally, imagine a 20-year bond with a floating rate equal to, e.g., the one-month TB rate, which is assumed known. It would trade very close to par all the time.

We would argue that the current regulatory regime has a lot in common with the third and fourth case above, whereas our proposed regime would at least attempt to mimic the behaviour of the last case, the floating rate bond. Clearly it would *not* be possible to mimic this case precisely because, as in the fourth case, the appropriate discount factor for the equity component cannot be observed. But our regime *would* remove the sources of fluctuations equivalent to those in the second and third cases. Furthermore, the additional volatility in these cases is *induced* by the coupon process (which is a metaphor for what regulators do). It does not reflect any true underlying risk. So, we would argue that consumers would almost certainly be better off if it were eliminated.

But even if this point might be debated, a more cynical argument would be that reducing the magnitude of fluctuations in prices of regulated equities would also be politically desirable.

Would short-term fluctuations in the RAR cause problems for regulated companies? We would suspect that any such problems would be quite manageable. Corporations routinely deal with volatility in all sorts of inputs to their decision process. Regulated companies already deal with uncertainty in their ex-post return on regulated equity. Quoted companies—regulated and unregulated—deal with far greater volatility in their stock prices. This does not appear to worry investors in quoted companies. Furthermore, if companies *did* care about volatility in the CAPM-WACC, it would be extremely easy to hedge this risk, since the only source of variation is in the indexed bond yield, which is traded in highly liquid markets. (This is, incidentally, another reason to prefer the “pure” CAPM-WACC to the CAPM(E)-WACC, since corporate bonds may have illiquidity problems, making hedging more difficult.)

The RER is an expectation and therefore cannot be measured.

We reiterate the definition of an expectation: it is the best forecast given available information.

In this context, restating the definition above, $RER = RAR + (1-g) E(X)$ so any value of the RER requires regulators to make a forecast of outperformance. This does *not* mean that regulators are expected to have perfect foresight. It simply means they should to their best in *predicting* outperformance. In practice, prediction can only be carried out with reference to history, hence our Recommendation MPW 2 above.

Using CAPM-WACC would make prices to consumer more volatile.

This seems to us to be a completely separate, and avoidable issue, which could be dealt with straightforwardly. If regulators wish prices to consumers to be stabilised over (for example) the full price control period, unpredicted fluctuations during the price control period in the CAPM-WACC, and hence in the RAR, could simply be credited to/debited from the Reserve Asset Base.

We are grateful to Phil Burns for pointing us to this neat solution (though he is clearly not personally responsible for this recommendation). In his terminology, the RAB is an IOU from consumers to regulated companies. If, for example, at some point in the next few years, the CAPM-WACC turns out higher than expected because real interest rates recover, the impact of this rise on consumers could be delayed by crediting the additional element in the allowed return RAR to the RAB. But if the rise were sustained, in due course firms would be guaranteed that prices would ultimately rise.

The RAR would need to adjust on a daily basis.

Our response to this is: so what? In principle, this would be possible (banks seem to manage to credit microscopic amounts of interest on a daily basis). Equally if transactions costs were deemed too high then a lower frequency could in principle be chosen; but the basis for this choice would need to be clear.

Using the “pure” CAPM-WACC provides no protection to companies from embedded costs of debt.

We are unconvinced by the arguments that regulated companies should be given this protection (we give some reasons below).

However, we do recognise the strong views held by regulators that the RAR should reflect embedded debt costs. So we now consider two possible compromises. These would break the direct link between the RAR and the CAPM-WACC implied by Recommendations MPW 3 and MPW 4; but crucially they would *still* be consistent with regulators implementing (versions of) Recommendation MPW 1 and MPW 2.

8.5 Two possible compromises

We acknowledge that a framework of the ideal (or even idealistic) form set out above is unlikely to appeal to regulators (if at all) except as a longer-term objective. Two possible compromises are at least worthy of consideration:

8.5.1 Use CAPM(E)-WACC?

The “pure” CAPM-WACC does not include the observed premium element in the cost of debt that is unexplained by the CAPM. As a result, as Figure 4.13 above showed, it is typically lower than the CAPM(E)-WACC, that uses CAPM to estimate the cost of equity, but uses bond yields to estimate the cost of debt (although in practice for most of the time the two estimates have moved broadly in line⁹³).

We acknowledge that the unexplained premium component of the cost of debt is a cost companies do face when issuing debt. As such, *if* we take it as a given that regulated utilities should be highly levered, there is a case for incorporating it into the RAR. But this still begs the question: if debt is unwarrantedly

⁹³ A clear exception was during the financial crisis, as discussed above, in Section 4.8.

expensive, why do regulated companies borrow? And why do regulators encourage high levels of gearing, when, by doing so, they need to pass this additional cost on to consumers?

We would also point to a clear disadvantage of CAPM(E)-WACC in terms of implementability, at least as currently applied: the very significant spike in the CAPM(E)-WACC during the financial crisis did not, we would argue, represent a rise in the true cost of debt of anywhere near this magnitude (particularly for regulated companies), since it was almost certainly driven by a sharp rise in the market-perceived probability of default. A similar spike in the future would require consumers to compensate companies, without any justification in terms of the true cost of debt.

8.5.2 Allow for the costs of embedded debt?

We acknowledge that, at least as currently practised by Ofgem, for example, the treatment of embedded debt does have the merit that the allowed return on debt will *on average* be very close to the market cost of debt, so that movements in the RAR due to this practice should (we argue below) average out close to zero.

But we remain sceptical of the arguments for the current practice, for a number of reasons:

- Unregulated companies do not receive this kind of insurance from their customers.
- By reducing risks for regulated companies, allowing for the embedded cost of debt may indeed lower the cost of debt (and possibly also the cost of equity, and hence the WACC itself) but if it simply passes the risk on to consumers, it is far from obvious to us why this should be desirable.
- We would also point to the fact that, in its purest form, allowing for the embedded cost of debt, such that companies interest payments are simply reimbursed by the consumer, is an arbitrage opportunity unless the cost of debt is set precisely equal to the risk-free rate. Even in its more sophisticated forms, such as Ofgem’s “Trombone” there is an arbitrage strategy that would eliminate risk for a company that borrowed a share of its regulator debt allowance precisely in line with the construction of the weighted average.
- Additionally, when the allowed return on debt differs from the market cost of debt - all the benefits of which go to equity-holders – this gives another source of fluctuations that would be avoided by our simple approach.

8.5.3 A compromise approach

While we are sceptical of both these compromises, we acknowledge that both modifications, if tightly defined, *could* still result in a regulatory framework in which the link between the RAR, the RER and the CAPM-WACC was much more tightly defined, with any resulting differences having an explicit justification. Thus consider the formulation below for the RAR that embodies these two compromises:

$$RAR = g.RARD + (1 - g)RARE + W_R \quad (26)$$

where

$$RARE = RFR + \beta_E (EMR - RFR)$$

and *RARD* is set in a formulaic way to reflect the regulator’s best estimate of embedded debt costs for the average firm (for example, using the moving average approach currently applied by Ofgem).

It follows straightforwardly that we can also express the compromise approach in Equation (26), using Equation (8), as

$$RAR = \text{CAPM(E)-WACC} + W_R + g(RARD - E(R_D)) \quad (27)$$

If we compare this with our ideal framework it is evident that our two compromises imply two differences. Most evidently, we replace the “pure” CAPM-WACC with the CAPM(E)-WACC; and the treatment of embedded debt implies the additional term, in brackets.

We would prefer this term not to be there. However, a formulaic approach to embedded debt would set RARD in terms of a weighted average of a market measure of corporate yields (such as the iBoxx yield shown in Figure 4.11). Furthermore, Appendix H concluded that most of the time, the adjustment of corporate yields to expected returns is only quite modest. Thus, this additional term should average out to a value close to zero, and – crucially – in a mechanistic way. As a result these two compromises would *not* rule out regulators implementing versions of Recommendations MPW 1 and MPW 2: i.e., they could (and we would argue) should still set explicit values for the regulatory wedge W_R and the informational wedge W_I , consistent with an explicit objective of “aiming up”.⁹⁴

8.6 Conclusions: where next?

This report has presented both points of principle, together with some evidence, in using the CAPM-WACC in regulation. We do not regard the evidence we have presented on some aspects of the CAPM-WACC as definitive. We would recommend that regulators put resources into the following areas in particular:

- Beta estimation
- Measuring outperformance

The second of these gaps in the data may actually be of the greatest quantitative significance. The limited evidence on outperformance that we present in Appendix K suggests strongly that the true Regulatory Expected Return set by UK regulators since privatisation has been distinctly higher than the Regulatory Allowed Return (or notional “allowed WACC”). But, on the available evidence, we do not know by how much.

But, subject to improvements in the evidence base, we would argue that we have set out a methodology for estimating a market-consistent CAPM-WACC, on which all the authors of this report can more or less agree.

We (MPW) have also made proposals for *implementing* those estimates in the form of the Regulatory Allowed Return (RAR) and the Regulatory Expected Return (RER), that we believe are both implementable and defensible. These proposals would require further work before they could be implemented, particularly our Recommendation MPW 1, which would require the calculation of explicit numerical values for the two “wedges”. But we would argue that *without* such values the link between CAPM-WACC and the returns regulated companies actually earn will remain opaque.

⁹⁴ This would only be a version of Recommendation MPW 1 since a) the sum of the two wedges $W_R + W_I$ would not be equal to the gap between RER and CAPM(E)-WACC at any point in time, but only on average; and b) W_R would need a (small) adjustment to offset the calculated difference between the yield and the expected return.

The evidence that we have collected points fairly clearly, as of late 2017, to two features that have been implicit in recent UK regulatory behaviour:

- The RARs set by (at least some) UK regulators have overstated our estimate of a market-consistent CAPM-WACC, by an amount which has been increasing in recent years.
- At the same time, the RARs set by (at least some) UK regulators have typically *understated* the true Regulatory Expected Return, since they have not taken into account observed predictability of outperformance.

To be clear: we were not asked to review the broader UK regulatory system of price-cap regulation etc.—and nor have we. This report has focussed primarily, as specified in the terms of reference, on estimating the cost of capital. However, we were also asked to consider enterprise valuations and bid premia, which, as we have shown, are driven by the gap between the WACC and the expected returns for regulated companies. The evidence points to significant bid premia: see Burns’ Appendix J; and to substantial returns: see Wright’s Appendix K. These returns stand in contrast to realised returns elsewhere in the economy, as documented by e.g., Jorda et al. (2017), who find that residential real estate and equities have shown very similar real total gains of about 7% per year on average, in a database covering 16 advanced countries, over nearly 150 years.

In our view, this evidence suggests that it is likely that investors in regulated companies have been betting on both continued outperformance, and continued over-estimation of the CAPM-WACC. Since we have clearly not considered the regulatory system in the round, we cannot (and do not) say at what level the RER should be set. This is a matter for regulators, balancing the need to incentivise firms towards efficiency, versus ensuring greatest benefits (returned) to consumers. We do argue, however, that regulators need to be clear what overall expected return they are aiming for. In doing this, we think the contrast between the expected returns on regulated firms, and the realised returns elsewhere in economies, documented by e.g., Jorda et al. (2017), should be borne in mind.

We are very conscious that if regulators were to attempt to address both of these features (the sizes of the regulatory and informational wedges) simultaneously, and with immediate effect, then the result could be a severe shock to some regulated companies. So, any transition to a regime along the lines we have presented above would need to be handled carefully. In this limited respect, there is a case to be made for gradualism. But we would argue that any transitional arrangements should not lose sight of the endpoint, which should be to maintain a much tighter, and more transparent link between the returns regulated companies can expect to earn and our best estimates of their cost of capital.

9 Implementing the CAPM-WACC in regulation: the Regulatory Allowed Return (RAR) (Burns)

9.1 Introduction

The two fundamental questions raised by Mason, Pickford and Wright (MPW) are these:

- Should RAR differ from the CAPM-WACC?
- Should a mechanism be incorporated into the regulatory price-setting process to force RER to converge on RAR?

MPW argue that RAR should be set at the level of the CAPM-WACC; and that a mechanism should force convergence of the RER onto the RAR. We evaluate both these proposals from the perspective of whether they would be likely to work in the best interests of customers. On the first point we conclude that MPW have not demonstrated that their proposal can benefit customers' interests, and we would argue that allowing the RAR to vary from CAPM-WACC—particularly in making proper allowance for the cost of debt—can be in customers' interests. We also conclude that any historical issues arising from differences between RER and RAR should be addressed through a robust and evidence-based approach to setting cost and output target and incentive rates, not through applying an arbitrary adjustment mechanism to the RAR.

9.2 Should RAR differ from the CAPM-WACC?

To evaluate whether it is beneficial to set the RAR equal to the CAPM WACC it is useful to consider the issue initially assuming that no embedded debt exists in the regulated sectors. In this context, MPW's objections 1, 3, 4 and 5 are relevant.

Objection 1: CAPM-WACC is too volatile.

Objection 3: Using CAPM-WACC would make prices to consumer more volatile.

Objection 4: The RAR would need to adjust on a daily basis.

Objection 5: Using CAPM-WACC provides no protection to companies from embedded costs of debt (in this context we refer to embedded debt costs that companies might incur in future).

The over-arching response of MPW to these objections can be summarised as “so what?” Indeed, we have some sympathy with that view in the sense that:

- Companies and investors should be able to manage the additional uncertainty created by a CAPM-WACC regime (objection 1 and 4);
- Regulators can find mechanisms to deal with any possible price instability (objection 3);
- Companies would re-optimize their debt-equity decisions in the light of the new regime that removes any embedded debt protection (objection 5).

However, this response is not the relevant test for whether it is sensible or desirable to move the regime to a CAPM-WACC regime. The relevant test is whether a CAPM-WACC regime would be a better regime than the present one, which applies the CAPM to the cost of equity and has a different treatment of the cost of debt. MPW do not undertake this test, which we attempt in overview here.

MPW appear to identify two benefits to their proposal: first, that setting the RAR equal to the CAPM-WACC has a certain purist neatness from a finance perspective; second, that the RAR has exceeded the CAPM-WACC over the past 10 years, and so therefore customers are overpaying. The first of these is unlikely to represent a significant benefit for customers.

As far as the second potential benefit is concerned, the wedge that exists between the RAR and the CAPM-WACC is almost entirely due to the regulatory treatment of embedded debt costs. And as MPW acknowledge, the regulatory approach to embedded debt is a quite tightly constrained dragging anchor.

So, whilst customers may currently be “overpaying” relative to the CAPM-WACC⁹⁵, in future they could be underpaying.

As is well understood, regulators have applied industry benchmarks for the embedded debt allowance as this preserves the incentives to raise debt at efficient interest rates⁹⁶. In Ofgem’s case this has evolved into a rolling debt indexation method using the iBoxx indices. Within each sector there will be winners who have beaten the regulatory yardstick, but also losers. On average, and over time, customers pay, broadly speaking, the average level of embedded debt costs in the industry, which encourages companies to finance themselves efficiently so they don’t find themselves on the wrong side of the benchmark.

This system, even in its more rudimentary states of evolution, worked uncontroversially until the global financial crisis (GFC), encouraging both efficient financing decisions based on the information available at the time, and ensuring that customers did not overpay for the debt over the medium term. Since the GFC, the cost of debt has fallen significantly but higher-cost embedded debt has remained in the regulated sectors, and it is simply this that drives the wedge between the RAR and the CAPM-WACC. In the event that there is a shock in the other direction in future then customers who will have borne the high embedded debt costs of the sector to date will receive the benefit of low embedded debt costs currently being incurred as debt costs rise in the future.

Consequently, it is not obvious that there are any long-term benefits for customers associated with the MPW proposal.

Their proposal does, however, have potential costs that could make customers materially worse off, which have not been addressed by MPW. These include: the likely increase in beta associated with such a regime; the replacement of factual evidence with inferences to determine the RAR, with an associated increase in discretion in making those inferences which would undermine the stability and transparency of the regulatory system; and the impact on the optimal finance structure of the businesses.

A CAPM-WACC approach would, as MPW acknowledge, alter the risk characteristics of the sector. However, they are silent on whether the beta would change in response to the shift, and indeed their calculations of the CAPM-WACC use estimates of beta derived from analysis of companies regulated according to the present regime, not their proposed regime. It has been established both in theory and in practice that the beta moves with the regulatory regime, and it seems likely to change if the CAPM-WACC regime were applied in future. A company’s risk profile can be influenced by the nature of the regulatory system it faces,⁹⁷ and this conclusion is well supported by empirical evidence.

The channels through which RAR approach can affect WACC include:

- The extent to which the method introduces volatility or cyclicalities into the RAR, which can impact on the beta and debt premium;
- The extent to which the method allows discretion in the implementation can impact on perceived regulatory risk; and

⁹⁵ The “overpayment” may have been exaggerated by MPW, as we discuss below.

⁹⁶ MPW argue that embedded debt allowances create a risk-free arbitrage opportunity for the utilities. This is incorrect and is based on a misunderstanding of the debt regimes and use of yardsticks.

⁹⁷ See, for example, Parker, D. (2003); Guthrie, G. (2006)

- The process of changing from one RAR approach to another can impact on perceived risk of further changes⁹⁸.

On the first point above, we borrow the ‘four bond’ analogy from MPW regarding volatility of returns. MPW explains that a bond with a floating coupon equal to some kind of benchmark market return (e.g., the one-month Treasury Bill) would be trading close to par valuation, while a fixed coupon bond or a bond whose coupon is related to some historic average market return would have a more volatile valuation. MPW then draw parallel of their proposed CAPM-WACC approach to the floating-coupon bond as a way to minimise volatility in asset valuation.

While we acknowledge the mechanism in this analogy, one crucial point that is overlooked is the fact that the floating coupon bond would have a return that is much more correlated to the market return. This is indeed not an issue in MPW’s example as it refers to fixed income assets. But when applied to equity assets, where systemic risk (i.e., correlation to the market return) is a fundamental determinant of asset valuation, this analogy breaks down. In fact, we would argue that the more closely the return on regulated assets follows the market return, the higher the systematic risk on these assets, and therefore the higher the cost of capital. In practice, we would expect to see higher betas on these assets.

More generally, the approach to regulation can influence levels of systematic risk. The price-cap regulation of the form applied in the UK provides firms with incentives to improve efficiency. At each periodic price review, the regulator resets its forecast of efficient costs, informed partly by outturns from the previous period. This has the effect of passing on some of the benefits from efficiency gains through to customers. The efficiency incentives provided by a price-cap system exposes regulated firms to cost outturns deviating from forecasts. In order to be compensated for this risk, investors will demand a higher rate of return, compared to lower risk rate of return style regulation, though potentially lower than in a comparator unregulated sector. This is supported by empirical evidence, as shown in the box below.

The relationship between the regulatory regime and Beta

The length of time between regulatory reviews can affect systematic risk, as shown by Gandolfi *et al* (1996). Using a conceptual model, and simulated data, they show that under a price-cap system, shortening the price control period reduces companies’ betas because prices are re-set more frequently. This effectively ‘buffers’ firms against cost shocks. As the control period is lengthened, the exposure to cost risk increases, thus driving betas up.

Different regulatory systems create different risk-sharing arrangements between the firm and consumers, and different incentives for the firm. Under price-cap regulation, firms tend to be exposed to higher systematic risk than those under rate-of-return regulation. This intuition has been confirmed by empirical research by Alexander *et al* (1996), and Alexander and Irwin (1996), who compare asset betas under different regulatory systems, varying from close to pure price-cap regulation to those close to pure rate-of-return regulation. They find that companies facing incentive-based regulatory systems tend to have higher asset betas than those that do not. This holds true across all investigated industries.

⁹⁸ This third channel is sometimes used as an argument against any change in RAR method. This is clearly not the case but it indicates that changes should be proportionate, transparent and evidenced.

The conclusion of this literature is that the system of regulation and the way it is implemented has a feedback loop back into the CAPM-WACC. UK regulators also acknowledge that their actions can influence risk,⁹⁹ and many have formalised policies to ensure that adaptation of the regulatory system is done in such a way as to minimise uncertainty and regulatory risk.

In summary, this important feedback loop has not been adequately assessed by MPW and it would be reasonable to suppose that a switch from existing RAR approaches to a pure CAPM-WACC regime would likely increase the beta. As a result, it would also be reasonable to assume that the MPW estimates of the CAPM-WACC are understated, so the level of “overpayment” that is claimed to have arisen under the present regime is likely to have been overstated.

The second route by which customers could be made worse off is that the CAPM-WACC approach involves discarding important factual evidence and relying wholly on estimates of unknown quantities. The cost of debt is visible and measurable and covers the majority of the financing costs of the companies. The cost of equity is not directly observable and can only be inferred, and covers only a minority of the financing costs of the companies. In this context, a switch to the pure CAPM-WACC significantly increases the impact of measurement error and would fail the test of implementability and defensibility. It would not be reasonable and proportionate to change radically existing approaches and discard robust and accepted evidence that covers a significant proportion of financing costs.

Connected to this, the third route by which customers could be worse off is that the CAPM-WACC approach would be a less stable and predictable regime. Since the essence of the approach is to discard the only component of the WACC that can be measured, and rely totally on estimates of components that cannot be perfectly measured, this would bring in risk that regulators could be tempted to use their discretion over the key components - beta¹⁰⁰, TMR – in ways that could undermine the stability of the regime. This stability is the cornerstone of the UK regulatory model, where the focus has been squarely on achieving two highly desirable outcomes: maintaining investor confidence in order to keep the WACC low; and stimulating significant dynamic efficiency improvements (in large part through a predictable approach to remuneration of assets and performance).

In 2011 the Government reaffirmed the principles and commitments for effective economic regulation.¹⁰¹ These principles apply to the UK regulators, and were (pp.4-5):

- Economic regulators should have adequate discretion to choose the tools that best achieve these outcomes.
- The framework for economic regulation should provide a stable and objective environment enabling all those affected to anticipate the context for future decisions and to make long term investment decisions with confidence.
- The framework of economic regulation should not unreasonably unravel past decisions, and should allow efficient and necessary investments to receive a reasonable return, subject to the normal risks inherent in markets.

⁹⁹ For example, Ofwat states: “Risk exposure of regulated companies is intrinsically linked to the regulatory framework. This includes the way the companies are regulated and the manner in which they are remunerated for their exposure to risk.” See: Ofwat, ‘Cost of capital and risk mitigants – a discussion paper’, June 2011.

¹⁰⁰ The debate between the authors in this paper around beta estimation itself provides an indication of the scope for discretion in estimating this parameter.

¹⁰¹ https://www.gov.uk/government/uploads/system/uploads/attachment_data/file/31623/11-795-principles-for-economic-regulation.pdf

- The framework of economic regulation needs capacity to evolve to respond to changing circumstances and continue to be relevant and effective over time.

On the principle of *predictability* the government said:

“Economic regulation plays a significant role in establishing the terms under which investment is made. Efficient investment is an important part of promoting the long-term interests of consumers. It is important that the regulatory frameworks avoid adding undue uncertainty to the business environment.

To a large extent this is achieved by building a stable and transparent regulatory environment with a long track record of consistent regulatory decision making. A history of rational regulatory decisions, which can be objectively justified, creates an expectation that a narrow set of outcomes will follow a given set of circumstances. This in turn will help both investors and consumers to predict regulatory decisions. On the other hand, piecemeal, ad hoc or unanticipated changes in policy or regulatory responsibilities are likely to erode investor confidence and increase the cost of capital.”

The CAPM-WACC approach would undermine these principles.

The fourth area of concern with the pure CAPM-WACC proposal is that it does not take account of the wider context for setting regulatory allowances, and specifically the corporate tax allowance. Under the present model, debt financing generates corporate tax savings which can be passed onto customers. The capital structure that regulators apply to utilities balances the benefits of debt financing (e.g., tax advantages) against the costs (transaction and issuance costs and default costs), and consequently aims to ensure a consistent treatment across these components. MPW have not set out how this would be done under their proposals and it is not obvious that it is straightforward to do so.

In summary, the case for a switch to the CAPM-WACC has not been made. It is unclear whether there would be any benefit for customers over the long term, since the approach to remunerating debt is a smoothed long-run approach. In contrast, there would be likely to be costs associated with a CAPM-WACC approach, in terms of raising Beta, and weakening the stability of the regulatory regime which could both increase risk and undermine incentives for future performance.

If, however, the CAPM-WACC were adopted, then the present context would be relevant. If the existing embedded debt is not to be stranded (and MPW has given no reason why it should be), then there would need to a transitional process to ensure the continued remuneration of this embedded debt as part of the evolution to the new regime.

In practical terms, therefore, I recommend:

Recommendation Burns 1: Allow embedded debt. An allowance for embedded debt using benchmark rates (as in Ofgem’s or Ofwat’s approach) does not affect the risk balance between shareholders and customers. Rather it establishes a profile for smoothing changes in financial market conditions into the RAR, and the use of a benchmark (rather than pass-through) encourages efficient raising of debt. To the extent that these promote price stability and predictability this can generate benefits for household and business customers. In any event the methods for allowing embedded debt have evolved into the regulatory methodology and would require substantial transitional arrangements to change.

Recommendation Burns 2: Allow the default premium on debt. Furthermore, for the purely regulated activities that the debt is funding, the difference between the observed and expected return on debt (see recommendation 8) should not be incorporated into the RAR, as the RAR should reflect the cost to issuer not the expected return. Gearing levels are determined by shareholders and when making this decision shareholders trade-off the tax benefits of debt against the costs of issuing debt (default costs and liquidity costs). In this process the shareholder will take account of the cost of debt to the issuer not the expected return to the debt investor. This conclusion would be unreasonable if pricing or pre-tax returns were above those implied by an equity only model (i.e., an equity only entrant could come in and completely wipe away the surplus). So a condition on the above is that the pre-tax WACC based on gearing and the cost of debt to the issuer should be lower than the pre-tax WACC implied by the equity only model. This would be true in unregulated sectors, where it would be a driver of the ‘optimal’ gearing range and we would argue is also true in the regulated sectors. In other words to the extent that the benefit of gearing (in terms of lower corporate tax) are passed onto customers through the regulatory approach (which is generally true) then it is fair and necessary that the costs of gearing (including issuance costs and the default premium) are also passed onto customers (through basing the post-tax WACC on the cost of debt to issuer). To do otherwise would be inconsistent.

9.3 Should a mechanism be incorporated into the regulatory price-setting process to force RER to converge on RAR?

From a pure finance perspective, it is clear that MPW are frustrated that the normal equilibrium conditions of finance theory – that the cost of capital equals the expected returns on that capital – do not appear to be satisfied in the regulated sector. This has motivated them to argue that the RAR should be equal to the CAPM-WACC (discussed above). It has also led them to recommend that:

“any transitional arrangements should not lose sight of the endpoint, which should be to maintain a much tighter, and more transparent link between the returns regulated companies can expect to earn and our best estimates of their cost of capital”.

Whilst MPW are silent on the precise mechanism, the intent is clear which is for an adjustment mechanism to force a fixed relationship – or even a convergence - between the RER and the RAR, so that ideally, RAR would drop out altogether and RER would move in line with the CAPM-WACC.

In a general equilibrium framework with perfect competition and efficient capital markets, then the cost of capital would equal the expected return on capital – along with other theoretically important conclusions such as efficient levels of operations, investment and pricing. However, even in the unregulated sectors of the economy, these conditions do not hold in practice – often for extended periods of time. And in sectors that are regulated, they do not hold for very well-understood reasons. There is a tension at the heart of regulation, between encouraging allocative efficiency (where prices equal costs and there are zero excess returns), and encouraging productive and dynamic efficiency (where costs are as low as they could be given the assets in place, and where those assets are added to and replaced as efficiently as possible). Thus, the US-style of cost of service regulation historically attempted to achieve allocative efficiency, but at the cost of productive and dynamic inefficiency that would not be observed in the idealised competitive equilibrium. In the UK, a built-in feature of incentive-based regulation is to promote information revelation and cost discovery, and this requires that companies are given profit incentives to beat the regulator’s targets. No regulatory system of any significant degree of complexity can be expected to simultaneously satisfy all the outcomes of a theoretical model of perfect

competition¹⁰². It is for this reason that it has never been an objective of UK network regulation that it “mimics competition” – the statutory duties of the regulators do not require this; no regulator has established that benchmark; and recent NAO guidance does not refer to it either. This does not, of course, preclude regulators from seeking to apply the insights gleaned from observations of competitive markets to their task of regulating networks, but that should not be confused with regulators having an objective of mimicking competitive markets.

The first observation on the suggestion by MPW that there should be some mechanism to drive a convergence between RER and RAR therefore follows from this preamble – whatever the mechanism, it will impact on the behaviour of the management of the companies (and of their shareholders) in respect of one or more of their operational, investment, pricing, or forecasting behaviour. The RER is driven by investors’ expectations of how the companies will perform in future—their appetite to manage cost risk, innovate, improve operational performance, and optimise their investment spending (which has both spatial and intertemporal elements), and this performance will be conditional on the regulatory regime itself. If the regulatory regime is now also to include a mechanism to force convergence between the RAR and the RER then this will change behaviour, and lead to a range of unintended consequences, including poorer performance.

An underlying assumption of the MPW mechanism is that the RER can be “set” by the regulator. In a trivial sense it can be, if it is simply defined as the regulator’s estimation of the investors’ expectations of the returns that the company will make under the regulatory regime. But at a deeper level, the investors’ expectations that drive the RER in the MPW framework are outcomes of the process itself, rather than an input to it (depending as they do on the expectation of managerial performance conditional upon the regime itself). Outperformance of the companies of their cost and output allowances is a built-in feature of incentive-based regulation (the primary purpose of which is to promote cost discovery), so it would be reasonable to expect outperformance. However, we disagree with MPW that this is “predictable” - if it were, then regulators would have set their cost and output targets to capture this predictable outperformance. It is noticeable that the pattern of outperformance varies over time, across sectors, and between cost and output performance.

Moreover, except for listed utility businesses, investors’ expectations can only be infrequently observed when businesses are bought and sold, and even then these are only the expectations of the winning buyer. Setting up a mechanism that links an infrequently observable set of investor expectations, to an RER as “set” by the regulator, with an unobservable CAPM-WACC will be fraught with a host of practical issues that are probably insurmountable, and which could quite likely increase regulatory risk.

Consequently, at both a conceptual and practical level, the MPW proposal is likely to be detrimental to customer interests in undermining efficiency incentives and/or increasing risk.

However, that is not to say that regulators are powerless. They already have at their disposal a range of instruments that effect company profitability, most notably the targets that they set the companies and the incentive rates they apply around those targets. Regulators have evolved their usage of these instruments over the years. For example, in the early years of RPI-X regulation in the 1990s, the intentional design features of incentive based regulation – combined with a highly generous set of original price controls implemented at privatisation – were such that by 1994 and 1995, the excess profits made by utility companies, accompanied by large increases in remuneration for senior management,

¹⁰² Regulatory regimes based on yardstick competition have attractive features but even these do not perfectly meet the theoretical competitive benchmark. Moreover, these regimes can only be implemented if there are a sufficient number of operators in the sector.

created significant public disquiet. The lesson of that period was that customers and a wide range of stakeholders did not only care about efficiency. They also cared about the fairness of the distributional settlement between shareholders, the executives of the businesses, and customers. Noticeably, policymakers and regulators still recognised that it was important that the cost discovery incentives were not discarded in favour of a superficially appealing cost-plus regime, because of the significant efficiency impact on customers. Research at the time utilised an option pricing framework to define the optimal compromise between cost-plus regulation and high-powered regulation¹⁰³, and over the same period regulators moved very quickly to jettison pure-form RPI-X regulation and develop more effective methods of customer protection, while attempting to preserve the efficiency properties of the original Littlechild RPI-X model. This process still continues today, as regulators attempt to provide a stable commercial and incentivising framework within which businesses can make efficient operating and investment decisions, whilst protecting customers.

Our appendix on the drivers of transaction premia is not driving us towards a conclusion that the protection of customers' interests would be enhanced by fundamental changes to the way that regulators have set the cost of capital. If there are concerns about high premia and associated excess returns, these appear to stem from the unusually high outperformance of cost and output targets, particularly amongst some energy networks following the first round of RIIO price control reviews.

There are many mechanisms to ensure that profits lie within a socially acceptable range, but some of these mechanisms would be extremely detrimental to customers, whilst others would be much less so. The art of regulation is to promote incentives for efficiency for the long-term benefit of customers, whilst achieving a fair settlement in the shorter term with a minimum of disruption to the longer term goal.

In our view, regulatory action on outperformance should apply to the cost and output targets not to the RAR – the RAR should be focussed on the WACC and minimising regulatory risk implies that this should be clear and transparent. An arbitrary adjustment factor applied to the RAR would only add to regulatory discretion and risk.

Recommendation Burns 3 (The CAPM-WACC and the RAR): The RAR should include the cost of debt as well as an estimate of the cost of equity, and should not include an adjustment factor that links the RER with the RAR. Regulators should however be able to justify their treatment of the cost of debt within the RAR, and should also ensure that they are able to broadly understand the sources of outperformance amongst the companies they regulate in order to take those insights into future price control reviews.

¹⁰³ See, for example, Burns and Weyman-Jones (1996) and Burns, Turvey and Weyman-Jones (1998).

Appendices

A Are Investor Investment Processes Consistent with the CAPM? (Derry Pickford)

Summary

In part two of the main report we argue for the retention of the CAPM for assessing the cost of equity. In this appendix we try to explain how observed behaviour of investors is consistent with the CAPM. We argue that when it comes to asset allocation it isn't, but it is consistent when it comes to selecting equity managers. We argue that it is this latter part of investor behaviour that is crucial for determining the relative cost of equity of different companies. We argue equity fund flow data suggests that investors pursue alpha as measured by the CAPM rather than by multi-factor models (see Appendix B). This incentivises fund managers to seek out equities that they believe will produce the highest CAPM alpha rather than multi-factor alpha. We argue that this implies cost of equity is best assessed in terms of the CAPM. Inevitably there is a high degree of overlap with Appendix B, which discusses these multi-factor models in more detail.

Our main conclusions from this appendix are that:

- Asset allocation models are on the whole far richer than the CAPM would suggest.
- Investors tend to dislike really bad outcomes more than the CAPM might imply. This is known as “loss aversion” and may explain some financial market anomalies, in particular why corporate bonds appear to earn higher returns than some formulations of the CAPM may suggest.
- Some CAPM anomalies likely arise from difficulties in identifying the right “market portfolio”. Identification of the “market portfolio” is also necessary for estimating the CAPM beta. For UK based investors a reasonable proxy for the “market portfolio” is the £45 trillion of “return seeking” assets summarised in table A1.
- Equity security selection is usually delegated to fund managers who act as agents.
- Equity fund flow data suggests that investors appear to assess their equity fund managers by how much Jensen's alpha (the realised return minus the return the CAPM model would suggest should have been earned given the market performance) is added.
- This feature of the flow data therefore suggests investors are indifferent as to whether net alpha is added from factor tilts or multi-factor alpha (the left-over return after all factors are accounted for). Investors consider deviations from CAPM as exploitable opportunities.
- Investors and their managers often use models different from the CAPM for forecasting returns on specific equities. These tend to incorporate subjective views and often suggest a difference between the cost of equity and the expected return on equity (i.e. they imply market inefficiency). Even when multi-factor approaches are used these are heterogeneous. When it comes to assessing the “average” expected return on *individual* equity securities investor behaviour can still be reasonably proxied by the CAPM.

(i) Asset Allocation

The original versions of the CAPM imply investors should simply hold combination of the market portfolio and the risk-free asset. In reality investor portfolios are extremely heterogeneous. A further complication is that investment processes are often characterised by a long chain of principal-agency relationships. The term “investor” can often end up referring to either the principal (the ultimate owner and beneficiary of the investment) or their agent which may be a fund manager, or a trustee acting on behalf of the ultimate beneficiary.

Calverley, Meder, Singer and Staub (2007), and Sharpe, Chen, Pinto, and McLeavey (2013) provide good overviews of respectively how many institutional investors form their expectations about asset returns and then asset allocate. These analyses form key parts of the Chartered Financial Analyst (CFA) syllabus and therefore are likely to be influential in terms of actual practice amongst practitioners. Calverley et al argue that capital market expectations (CMEs) are the starting point for portfolio construction processes. They dichotomise CMEs into either macro CMEs regarding broad asset classes, or micro CMEs regarding individual securities. For larger investors it is the former which gets the most attention. These are typically developed in-house or in conjunction with their investment advisors. CMEs amongst investment advisors are more commonly known as capital market assumptions or CMAs. Asset class CMAs not only provide expected returns (usually the compound average or geometric) over different horizons (usually between 10 and 30 years) but will also estimate a covariance matrix of all major asset classes. CMAs will also occasionally make assumptions about higher moments, i.e. the level of skew and kurtosis of returns too, although this is sometimes accounted for directly in the risk modelling process.

How granularly the asset classes are defined will also vary but typically these will cover returns for equities, sovereign bonds, and credit (made up of corporates bonds, securitised consumer debt, and loans) for most regions. The US, as the largest equity market by market capitalisation (representing 59% of MSCI World index and 52% of the MSCI All Countries Index at the end of 2017), is often further subdivided into small and large cap. Alternatives (real estate, commodities, private equity and debt, hedge funds and infrastructure) are also modelled.¹⁰⁴

The CAPM would suggest that once the expected return and risk of the global market portfolio is identified, which is summarised in Table A1, investors should simply decide where on the capital market line (i.e. how much of the risk-free asset and the market portfolio) they want to be. Consultants sometimes subdivide the assets summarised in table A1 into return seeking (55%) and “safe” assets (45%). These are not “safe” in the CAPM sense. Firstly, they can exhibit significant nominal (and real) short-term price volatility. However, as we discuss below whilst the CAPM is a one period model whereas investors have a much longer time horizon. For most investors the value of their liabilities (whether explicit, in the case of a defined benefit scheme, or implicit in the case of an individual saving towards retirement) will be positively correlated with these “safe” assets and will be demanded as a hedge.

Secondly, although investment grade bonds are considered as a safe asset they still have a degree of default risk.

¹⁰⁴ The Horizon survey of investment advisors shows the minimum level of granularity that forecasters provide.

Table A1**Size of global invested capital market measured in sterling**

Asset Class/Category	Market Size in GBP (Trillions)	Proportion of Global Invested Capital Market (%)
RETURN SEEKING ASSETS		
U.K. Equity ex REITs	1.77	2.1%
Non-U.K. Equity (Developed) ex REITs	24.55	29.8
Emerging Markets Equity ex REITs	3.61	4.4
Frontier Markets Equity ex REITs	0.11	0.1
Private Equity	2.07	2.5
Private (Unlisted) Infrastructure	0.31	0.4
Timberland	0.04	0.1
Private Real Estate Debt ¹⁰⁵	0.47	0.6
Private Real Estate Equity	4.13	5.0
Public Real Estate Equity	1.36	1.7
Commodities ¹⁰⁶	0.21	0.3
High Yield Bonds	1.91	2.3
Bank Loans	0.93	1.1
Emerging Market Bonds (Sovereign; USD)	0.62	0.8
Emerging Market Bonds (Sovereign; Local Currency)	1.50	1.8
Emerging Market Bonds (Corporate; USD)	0.68	0.8
Insurance-Linked Securities	0.06	0.1

¹⁰⁵ We include only private real estate debt market accessible to institutional investors, estimated to be 10% of the entire private real estate debt market

¹⁰⁶ We include commodities as part of the global invested capital market while acknowledging that a large part of commodity investing takes place through derivatives (e.g., futures contracts) and the total net position in derivatives is zero.

“SAFE” ASSETS		
U.K. Bonds (Investment Grade)	1.43	1.7
Non-U.K. Bonds (Developed)	30.40	36.9
Inflation-Linked Bonds ¹⁰⁷	2.04	2.5
Money Market/Cash Equivalents	4.13	5.0
Total Global Invested Capital Market Excluding Hedge Funds	£82.35	100.0%

Source: Aon Hewitt

Which assets should be used in assessing the “market portfolio”? As we discuss in the main body of the report any mean-variance efficient portfolio solves the CAPM equation and therefore will give unbiased estimates of the expected return on a particular asset. However, identifying a mean-variance efficient portfolio requires knowing the expected returns of the different asset classes with small permutations in assumptions making large differences to the optimal portfolio. The Black-Litterman model (discussed below) is one way of solving this. An alternative approach is to take only return seeking assets in proportion to their market capitalisations as the market portfolio. This creates a diversified universe of assets of £45 trillion.

Another “real-life” discrepancy from CAPM is the home bias puzzle: investors tend to own far more of their own equity market than its share in global markets. The standard explanation of this is that the CAPM assumes no taxes, transaction costs or other investment barriers, the reality is of course very different. Withholding taxes may mean that the overall tax on a foreign holding maybe higher than domestic holdings, while overseas investments may also have higher transaction costs.¹⁰⁸

French and Poterba (1991) tried to quantify these additional costs of overseas investment and argued that it was, nonetheless, difficult to explain this “home bias” puzzle. For large investors investing in the main developed market equity markets these should not be major barriers. French and Poterba’s methodology is unable to adjust for other potential barriers to overseas investment. When investing away from home, investors may worry about whether there are information asymmetries, and, where property rights are weak. There may be a tail risk that foreign investors might be unfairly treated. However, once accepting that there is a home bias what does this mean for the assessment of the cost of equity?

The CAPM would suggest that the beta of a security should be assessed with respect to the global portfolio. However, if home bias was extreme a case could be made for using the local index instead. Whilst acknowledging the home bias puzzle we would argue that assessing at the global level is still best practice. Even where managers are given global mandates, active stock pickers are more likely to invest in local stock markets, as documented by Couerdacier and Rey (2013). Whilst the stock of equities is more typically invested into active vehicles, the flow is becoming increasingly passive. Investing at a global level is considered to be best practice by most investment consultants since it maximises diversification.

¹⁰⁷ These include catastrophe bonds, ILWs, sidecars, and collateralized re and others in the total asset under management of Insurance –Linked Securities market.

¹⁰⁸ Exposure to real exchange rate risk is also seen as an additional problem, although it is far from clear that this represents additional risk. See Black (1989).

This is particularly true for investors from countries with smaller local capital markets. Secondly, since prices and therefore expected returns are set at the margin we need to look at the behaviour of the marginal investor. Whilst the marginal investor is not identifiable, home biased investors are likely to invest in their home market regardless and are consequently less likely to be the marginal investor.

Currently the Horizon survey¹⁰⁹ of investment advisors' CMA's does not differentiate between the returns on active and passive equity investing or the returns from different "factor styles" (see Appendix B). Although most investment advisors will attempt to model how much additional return can be potentially added from either active strategies or "smart beta".¹¹⁰ This will often be done at a client level since the amount of active risk (how much a portfolio will differ from a benchmark) will be dependent on the client beliefs and preferences about the benefits of active versus passive.

Once CMEs are established, these are then used in the formation of a strategic asset allocation. Approaches to strategic asset allocation are often subdivided into asset-liability modelling (ALM) and asset-only approaches. ALM is most associated with defined benefit pensions schemes and insurance companies who have little flexibility about their future payouts. However, other institutional investors such as charities or educational establishments' spending needs can be seen as having quasi-liabilities, which can be modelled in a liability framework.

As most investors will typically have a reasonably long investment horizon, the modelling is typically long-term and is used to derive a long-term strategic asset allocation. The CAPM by contrast is a single period model. As the model is for only a single period, the arithmetic mean and geometric mean return over the investment horizon are identical. Where solely historical data are used, the arithmetic average of returns will be an unbiased estimator of the one period return. By contrast most investment consultants will focus on their expectation of the compound average or geometric return over the investment horizon. This can sometimes create some confusion as whilst the historical arithmetic return is an unbiased estimator of the arithmetic return, the historical geometric return is not an unbiased estimator of the geometric return. This is because the average geometric return is a concave function of the cumulative return and therefore, by Jensen's inequality, the geometric average return of the expected cumulative return is greater than the expected geometric return.¹¹¹

The shift to focusing on the expected geometric rather than the arithmetic return isn't the only consequence of shifting to a multiple-period model. Investors don't just care about the distribution of real returns but how the investment opportunity set is likely to evolve too. This is the main insight of the Intertemporal CAPM (ICAPM) developed by Merton (1973). This alongside the difficulty of identifying the market portfolio, and the wide heterogeneity of portfolios, has led to alternative approaches. Ross's (1976) Arbitrage Pricing Theory argued instead for a variety of risk premia. Ross focused on the sensitivities of portfolios to macroeconomic surprises. Whilst these factors are used in some risk models, in terms of expected returns the focus is predominantly on factors which are derived from fundamental characteristics of underlying securities. This is explored in Appendix B.

¹⁰⁹ The Horizon survey is a survey of 35 large investment advisors about their capital market assumptions http://www.horizonactuarial.com/uploads/3/0/4/9/30499196/horizon_cma_survey_2017_v0822.pdf.

¹¹⁰ Smart Beta is a quasi-passive strategy where more equities with certain characteristics or factors are held, again this is discussed in more detail in Appendix B

¹¹¹ If the entire information set on which expected returns are conditioned are historical returns over T years, there is no serial correlation of returns, and the time horizon of the investor is N years then a nearly unbiased estimate of the expected compound (geometric return) is given by:

$$E(R_N) \approx (aA_T^N + (1-a)G_T^N)^{1/N}$$

where A_T is the historical arithmetic return, G_T is the historical geometric return and $a = (T-N)/(T-1)$. See Blume (1974) and Cooper (1996) <http://faculty.london.edu/icooper/assets/documents/ArithmeticVersusGeometric.pdf>

The dominant tool used for ALM is Monte Carlo simulation. For ALM this enables the joint modelling of how the net present value of liabilities and assets may evolve. Monte Carlo simulation is also used for asset-only approaches. Monte Carlo simulation allows the use of far more realistic asset return paths than can be done with analytical approaches. The left-hand side of Figure A1 shows 400 different return paths, however, simulations will not stop there and thousands will be run to create the fan chart or “tunnels of doubt”, shown on the right-hand side.

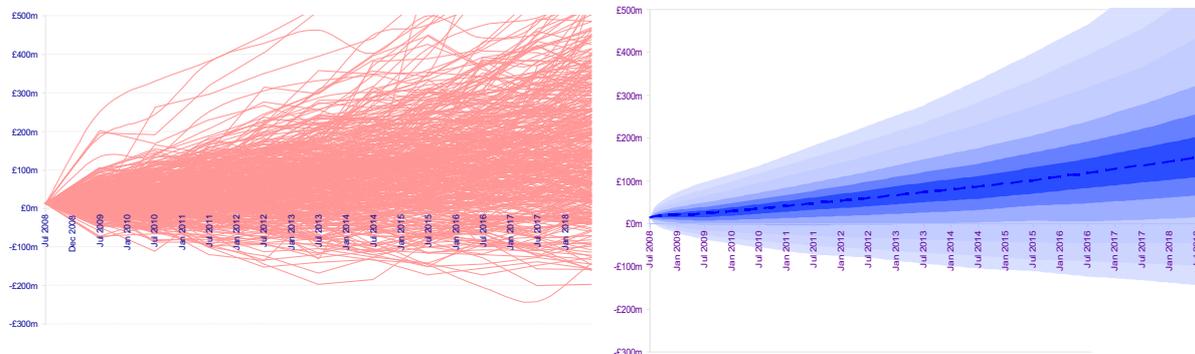
An investor can then choose the portfolio which they believe reflects their preferences over risk and return best. Economists often assume that preferences can be modelled by quadratic utility functions which imply a simple trade-off between the variance and return on a portfolio. However, investors may care more about extreme poorly returns or, in the case of investors building assets against a liability, large deficit scenarios. Monte Carlo simulation therefore enables investors to express more complex preferences, than a simple trade-off between the mean return and the variance of returns.

Monte Carlo simulation of portfolio returns is also able to capture features of how asset class correlations can change in risk events. These risks can be missed when using a stable covariance matrix. In particular, as we discuss in Appendix H returns from corporate bonds, when held to maturity, appear higher than the CAPM would suggest. Whilst most of the time returns on corporate bonds are not closely correlated, performance in very bad times is. Combined with greater extreme loss aversion than standard mean-variance optimisation suggests, investors rationally require higher returns from corporate bonds than the CAPM suggests.

Figure A1.
Monte Carlo simulation can be used to build up return distributions
for the value of a portfolio or a scheme surplus.

400 Return Simulations

Fan Chart Showing Probability Density of Different Outcomes



Source: Aon Hewitt

Before Monte Carlo simulation became more common place, quantitative investment strategy analysis tended to use mean-variance optimisation (MVO). MVO was developed by Markowitz (1959) and is one of the key building blocks of CAPM. Whilst CAPM typically assumes leverage is possible and therefore that more than 100% of the portfolio can be invested in the market portfolio, mean-variance optimisation enables the modelling of additional constraints. MVO and CAPM both assume returns are normally distributed, and this simplifying assumption enables optimal portfolios to be derived analytically (the minimum variance portfolio for a given level of return, or the highest return for a given level of volatility). Monte Carlo analysis enables more complexity about the return process to be used including the possibility of “fat-tailed” events.

A drawback of MVO is that small changes to assumptions about returns or the covariance matrix can make very large differences to the optimal asset allocation. Whilst there are various permutations of it, the Black-Litterman model typically provides more stable asset allocations than a naïve MVO, and its use is more widespread amongst investors who make tactical tilts as well as multi-asset funds (funds which invest in wide spectrum of assets including equities and fixed income). A typical Black-Litterman process will first attempt to identify the global investable universe of assets. An estimate of the covariance matrix of returns for the investor's time horizon is then derived. This is often the historical covariance matrix; used to ensure that the covariance matrix remains positive semi-definite, although the more granular the asset classes become the less robust this will be. Then the market capitalisation of the various asset classes is assessed. This can be considered as analogous to the CAPM market portfolio. A set of equilibrium CAPM expected asset returns is then reverse engineered. An investor can then adjust these "market based" returns using Bayes rule to incorporate their private beliefs to create a "posterior" view about returns. These new return estimates are used in conjunction with the covariance estimate to derive a new efficient frontier. The investor, as with the CAPM is then able to combine a risk-free asset and the optimal risky portfolio to create their desired portfolio. Black-Litterman therefore differs from CAPM in that it allows for heterogeneity of views but in its basic version is a very similar model.

What are our key conclusions from our survey of asset allocation processes?

- CAPM clearly does not capture most of the key features of asset allocation as investors do not simply hold a combination of the global market portfolio and a risk-free asset.
- Time horizons are far longer than the one period that CAPM assumes. Investors care about the compound average return over their investment horizon, rather than the arithmetic average of returns.¹¹²
- Investors care about the entire distribution of returns rather than just the variance.
- Some market anomalies, such as the high returns on corporate bonds, do not remain anomalies once more complex return characteristics and investor preferences around extreme loss aversion are accounted for.
- Home bias is a more difficult puzzle to solve. As UK investors are increasingly allocating equities at a global level, regulators should use global indices rather than domestic for assessing betas.
- Whilst ideally all risky assets should be used in formulating the market portfolio, in practice alternative assets suffer from low levels of price transparency. These assets are often illiquid with return series created through extrapolation or sticky prices. This leads to high levels of serial correlation in return series which will mean that their riskiness will be significantly underestimated.
- Once the asset allocation process has decided how much to allocate to equities, decisions about equity selection are often then delegated to sub-managers. This means that the asset allocation process may limit the direct impact on the expected returns on low beta versus high beta equities. We explore issues and around manager and security selection in more detail in the next section of this appendix.

¹¹² Care also needs to be applied when translating the expected cost of equity into a discount rate, since Jensen's inequality means that the discount factor will be slightly low than the reciprocal of the cost of equity.

(ii) Security Selection

Although some investors may build portfolios from direct securities, the economies of scale involved in building diversified portfolios means that investors will often use funds offered by investment management firms instead. Even larger investors may still delegate to external managers through the use of segregated mandates. This enables them to still get the benefits of the economies of scales from research and manager skill, but remove some of the issues with pooled assets.

Expectations about the returns on individual securities are referred to as micro expectations. The CAPM would suggest that investors should invest passively. However, Morningstar (a fund data provider) estimate that passive funds at the end of 2016 amounted to \$6.7 trillion versus actively-managed assets under management of \$23.9 trillion, albeit with the former growing far more quickly. If indices capture all investable securities and are within an asset class, investors should not in aggregate be able to earn alpha. If an index doesn't allow for costs then aggregate alpha will be negative. This creates the famous Grossman-Stiglitz (1980) paradox on the impossibility of efficient markets. Without the ability to generate alpha there is little incentive to engage in the research necessary for assets to be correctly priced.

Active managers are typically focused on identifying securities which will give superior risk-adjusted returns. Techniques are extremely heterogeneous but are often of the following three forms:

- Discounted cash flow models. These will either be at the per share level (the dividend discount model); the equity level, where the present value of free cash-flows to equity is taken; or the firm level where the present value of cash to debt and equity holders is calculated. The discount rate used will often be CAPM based.
- Multiples based models. Forecasts of the growth of some fundamental such as earnings, earnings before interest tax depreciation and amortisation (EBITDA), and then putting that fundamental on multiple which might be based on current levels for the company or its peers, or historic averages.
- Models which assess large amounts of quantitative information about the stock. These often have a degree of overlap with factor based investing.

Even if CAPM based discount rates are used, these micro-expectations are rarely CAPM consistent since they implicitly assume that there is a deviation between the intrinsic value and the current price of a stock. Consequently few active investors will think that the CAPM provides a good model for the expected return on an individual stock.

Similarly, smart beta investors (funds which increase exposure to the types of factors explored in Appendix B according to fixed rules) usually perceive that the expected return on an equity that they buy is greater than the cost of equity even if some interpretations of factor based investing are that the factor premia are rewards for risk.

If a large proportion of investors think that the expected return and the cost of equity differ; the models for expected return are heterogeneous and often require subjective inputs how do we measure the cost of equity for individual stocks? We believe that equity fund flow data, how much investors are putting in to and taking out of different investment funds, provides useful evidence on this and shows how investors perceive the opportunity cost of their equity performance.

Whilst Carhart (1997) argued that past fund performance, once adjusted for factor exposures, is a poor guide to future performance of fund managers, investors have relatively little else in terms of quantitative data on which to assess manager skill. Barber, Huang and Odean (2016) and Berk and van

Binsbergen (2016) argue that investor flow data can show the “revealed preference” of investors. Berk and van Binsbergen use different models to assess performance for managers investing in US stocks. They first test three models where performance isn’t adjusted for risk:

- Return
- Return above Treasury bills
- Return in excess of the market

They then test two models where performance was calculated using CAPM based models

- Jensen’s alpha with the S&P 500 as a benchmark
- Jensen’s alpha with the broad index as a benchmark (the CRSP value weighted index)

Finally two more sophisticated risk models were also used to assess performance:

- Alpha after adjusting for Fama-French risk factors
- Alpha after adjusting for the Carhart model which includes momentum

Berk and van Binsbergen find, regardless of the time horizon used, that whether a fund receives above-normal net flows or not was best explained by performance judged on a CAPM model with the broad index of stocks. How do we interpret this result? This suggests investors assess whether a manager is doing a good job or not, by its performance versus the CAPM. They are consequently indifferent whether alpha is added from factors or “pure” alpha after adjusting for factor exposures. This indifference suggests that they do not perceive factor premia as a reward for risk. However, even if factors are a reward for risk this doesn’t mean that cost of equity should be adjusted for the presence of factor exposures, as we explore in Appendix B.

We conclude that whilst equity managers, whether qualitative or quantitative, will rarely form their expectations of the return on an equity using the CAPM, they are rewarded (in terms of new flows) on the basis of how well their stock selections perform on the basis of the CAPM. This suggests the opportunity cost of a particular equity is its CAPM expected return.

Whilst investor preferences over the distribution of returns suggest that at the asset allocation level are more complex than suggested at the CAPM level, when it comes to security selection the CAPM is a reasonable model with deviations seen as exploitable opportunities rather than equilibrium risk premia which need to be priced.

B Factor models of Equity Returns (Derry Pickford)

Summary

In section 2 of the report we argued for the retention of the CAPM (Capital Asset Pricing Model) in assessing cost of capital. This appendix explores factor models of equity returns. The author of this appendix believes factor models are a better model for predicting the expected return on *portfolios* which target specific levels of factor exposure, although this isn't a conclusion that all the authors agree with. However, where the authors do agree is that CAPM should be retained for cost of equity calculations. We argue in this appendix that this is because:

- For an individual security, factor exposures are often transient in nature, it might be long a factor at one point in time, and then subsequently neutral, or even short. Factors therefore apply to the expected returns on *portfolios* but they are unlikely to make a significant impact on expected returns on *individual equities*, particularly at longer horizons.
- A wide universe of potential factors, termed “the factor zoo”, exists. Proponents of factor-investing suggest multi-factor approaches, in order to reduce the reliance on any one specific factor, which may prove ethereal. Any specific equity would likely be long some remunerated factors and short others and even if the factor exposures didn't change judging what the overall impact of factors at an individual stock level would be close to impossible.
- The robustness of individual factor premia is contentious, with some academics arguing that they are an artefact of data-snooping. Adjustments to the expected return on an individual equity on the basis of its factor loading would therefore be very difficult to defend.
- We believe implementing a factor framework where factor premia could be accurately assessed would likely be resource intensive. The size of adjustments would be small and whether the accuracy of the WACC estimate would be increased is highly debatable.
- Non-linearity of factors, and in particular the outperformance of low beta strategies, in theory open up the possibility of arbitrage. Any framework which allowed a low beta premium would likely be arbitrated by changing the capital structure.

We organise this appendix around six key questions:

- (i) What are equity factors?
- (ii) What causes these deviations?
- (iii) Are factor premia robust to data-snooping concerns?
- (iv) Do factor premia vary over time?
- (v) Do factors apply to individual stocks as well as portfolios?
- (vi) Do factor exposures impact cost of equity?

(i) What are equity factors?

According to the CAPM framework only an asset's sensitivity to the global market portfolio is relevant to determining how much more the asset should be expected to earn relative to the risk-free rate. Moreover, that expected return is an equilibrium one. This means that no investor can raise their expected welfare by changing their asset allocation given the set of expected returns. Proponents of factor investing by contrast suggest that investors can improve investment performance by systematically making sure that portfolios have certain quantitative characteristics or factors.

As discussed in section 3.4 of the main report, the *actual* cross-section of returns on individual securities over and above the risk-free rate, can be written as:

$$R_{it} - R_{Ft} = \alpha_i + \beta_1 F_{1t} + \beta_2 F_{2t} + \dots + \beta_k F_{kt} + \varepsilon_{it}$$

Where R_i is the return on security i , the k factors $F_{jt}, j=1..k$ capture mutual correlations between returns on different assets, such that any residual component in ε_{it} is purely idiosyncratic, and β_j refer to the exposures that security i has to factor j . The first factor, F_{1t} is usually assumed to be the market return, i.e. $F_{1t} = R_{Mt} - R_{Ft}$.

The generalised asset pricing formula becomes:

$$E(R_i) = R_F + \beta_M RP_M + \beta_2 RP_2 + \dots + \beta_k RP_k$$

where $RP_2 \dots RP_k$ are non-market factor risk premia. The formula implies that stocks have zero expected alpha once these other risk factors are accounted for. It also assumes that all risk premia are priced linearly, a stock with twice as much exposure to risk factor two will earn twice as much risk premium two. These two features mean that they are consistent with the assumption of no arbitrage. However, they're unable to account for some empirical features of returns such as the outperformance of low beta stocks versus high beta stocks versus relative to the performance predicted by the CAPM.

These factors may be macro in nature such as, growth surprises, levels of bond yields or inflation as discussed Ross's (1975) Arbitrage Portfolio Theory (see Appendix A); or they may reflect features of the stock such as those identified by Fama and French's models (1992) which we discuss in this appendix.

Similarities between stocks with certain types of characteristics mean that multifactor model inevitably are able to explain a greater proportion of the cross –section of realised returns better than the CAPM can. However, this only matters for expected returns if there is very strong evidence that the additional factors are *priced*. If $RP_2 = \dots = RP_k = 0$ then the model is completely consistent with CAPM, only the market factor is useful in predicting returns.

Factors might include how much of a company's activities are in a particular industry. Including an industry or sector factor may help us to "explain" the cross-section of realised returns but these only help us to improve the predictions of returns on a stock if industry exposures generate risk premia.

The finance literature from the 1970s onwards uncovered numerous anomalies which have suggested that some factors might indeed be priced. Haugen and Heins (1975) showed that low beta stocks tended to perform better than the model predicted and high beta less well. Basu (1977) showed that stocks with low price-to-earnings ratios tended to outperform stocks with high price-to-book ratios. Outperformance of stocks with high book-to-market equity ratios (the ratio of the book value of a common stock to its market price) was identified by Stattman (1980) and Rosenberg, Reid and Lanstein (1985) who argued that this outperformance could not be explained simply by being higher beta. Banz (1981) argued that a portfolio of the smallest companies in the stock market has delivered superior risk adjusted returns relative to the market cap index over the long term. This view was also tested and supported by Reinganum (1981). Size and value, together with the CAPM market return became the 3 building blocks of the highly influential three-factor model of Fama and French (1992). The first factor

is the standard CAPM market factor, the second represents the performance of being long a portfolio of high “book-to-price” stocks and short a portfolio of low “book-to-price” stocks, and the third factor being long a portfolio of small stocks versus being short a portfolio of short stocks. Fama and French argued that the factors in their model represented equilibrium risk premia and were entirely consistent with efficient markets.

Whilst value on the whole became accepted, the size effect has been unclear. Hsu and Kalesnik (2014) claimed that size has not been observed in the United States since the early 1980s and does not exist outside the U.S. This highlights a danger with anomalies; once discovered they frequently stop working (we discuss this in more detail in part ii of this appendix). Shumway and Warther’s (1999) examination of the size effect, led them to conclude that it was likely driven by failure to account for the performance of small stocks that delisted (i.e. there was survivorship bias in the results). Despite this mixed academic evidence, practitioners often argue that size is a good “enabler” of other factors, e.g. for value stocks, the outperformance is more pronounced in smaller companies. Another criticism of size is that any excess returns disappear once higher relative transaction costs (from lower liquidity) are accounted for. Whilst this might mean that net expected returns from investing in smaller illiquid companies may be no higher, the pre-transaction cost return, which will be equal to the cost of equity for the issuer, will be. Small companies may therefore have a higher cost of equity to reflect the higher transaction costs. This might impact the cost of equity of the very smallest independent utility companies. However, in a competitive market companies should find their optimum size. It would be perverse for regulators to reward inefficient scale. Furthermore, whilst higher relative transaction costs might apply to public listed companies it might not apply for other ownership structures.

Following the Fama and French paper, the search for additional risk factors took-off with the resulting universe of identified factors often being referred to as a “factor zoo”. The two additional factors which are most widely used are: momentum (the tendency of stocks that have done well over 6-12 months to continue to outperform) and various measures of quality (the tendency for companies with low capex and high return-on-equity to outperform). Of course, the take-off in the search for factors means that even if individual researchers are careful not to “data-snoop”, the collective effort of researchers does (see section ii).

Jegadeesh and Titman (1993) put forward momentum as a factor and provided evidence that stocks that do well relative to the market over a three to twelve month period tend to continue to do well for the next few months, and that stocks that do poorly continue to do poorly. Momentum was further evidenced by Carhart (1997) who demonstrated that by including momentum as an additional fourth factor to the Fama and French model, it had greater explanatory power.

Sloan (1996) argued that there are excess returns to high earning, quality stocks; Dechow, Ge and Schrand (2010) supported this view with measures including earning persistence, smoothness, and loss avoidance. In recent years, Bender and Nielsen (2013) further evidenced the quality factor with their paper on the accrual effect (stocks which have high accruals as a proportion of total assets tend to underperform). Asness, Frazzini, and Pedersen (2013) added to this with their publication on stable growth, profitability and high pay-out ratios. We summarise these six factors in table B1.

Fama and French (2015) then expanded their three factor model to five factors, adding two measures of “quality”: how profitable the firm is and whether it has to invest a lot (companies with high growth in assets being considered as low quality as this typically means less free cash flow for shareholders). These additions were far from novel, building on work by Aharoni et al (2013), and Novy-Marx (2013). Many factor index publishers combine several quality measures into a single factor. Whilst in the

original three factor model Fama and French provided skeleton arguments as to why value and size factors might be justified as risk premia, why a high quality company should be higher risk is less obvious.

Although Fama and French did not add low beta (their approach excluded the possibility of low beta stocks earning an excess return, since the reward for market risk should be the same for each unit of risk, and not decrease for higher beta stocks), and momentum (again momentum is hard to reconcile within an efficient market framework) these are easy to calculate with public information, and can be added to their factors which they publish on-line. We summarise the six main factors summarised in Table B1.

Table B1 A Summary of Equity Factors

Factor	Paper	Factor Definition/Data Set
Low Beta	The Capital Asset Pricing Model: some empirical tests. Black, Jensen & Scholes (1972)	Trailing five year beta CRSP US equity data 1926-1966
Size	The relationship between return and market value of common stocks. Banz (1981)	Market capitalisation CRSP US equity data 1926-1975
Value	The cross-section of expected stock returns. Fama & French (1992)	Book equity to market value CRSP US equity data (ex financials) 1963 - 1990
Momentum	Returns to buying winners and selling losers: implications for stock market efficiency Jegadeesh & Titman (1993)	Historical price performance (various past periods up to 12 months) CRSP US equity data 1962-1989
Quality Investment	- Stock returns and the Miller Modigliani valuation formula: revisiting the Fama-French analysis Aharoni, Grundy & Zeng (2013)	- Annual change in total assets/total assets CRSP US equity data 1963-2009
Quality Profitability	- The other side of value: the gross profitability premium Novy-Marx (2013)	Gross profits (i.e. revenue minus cost of goods sold)/total assets CRSP US equity data (ex financials) 1962 - 2009

Source: Aon Hewitt

(ii) What might cause factor premia?

These anomalies have led both practitioners and financial economists to wonder whether the CAPM framework is a realistic one. However, even the existence of robust premia doesn't necessarily negate the CAPM. Firstly, all tests of the CAPM are necessarily joint tests of CAPM as a description of equilibrium returns and market efficiency. Factor premia may be a function of market inefficiency rather than that CAPM is bad model of what the expected returns on equities *should* be. Secondly, tests of the CAPM have to identify the "market portfolio". As Roll (1977) points out, as the market portfolio is

unidentifiable it is possible that discrepancies from CAPM may in fact just be picking up misspecification of the market portfolio.¹¹³

There are three main competing explanations for these anomalies. Fama and French (1992) argue that these are all equilibrium risk factors. Whilst arguments can be created as to why Value and Small stocks might be higher risk, this is harder to do for the additional quality factors in their five factor model.

Lakonishok, Schliefer, and Vishny (LSV) (1994) put a framework based on behavioural anomalies as the driver of these anomalies and in particular, investors' tendencies to overestimate free cash flows of growth companies. Institutional factors are also likely to have a role, fund managers who are unable to use leverage have an incentive to gravitate towards high beta stocks in order to try to get to the top of fund league tables during bull markets. This will mean that they will typically be overpriced relative to "dull" low beta stocks. Equity managers are sometimes rated on the basis of quartile performance. Top quartile performance when markets are performing well tends to sell product. However, whilst running a low beta portfolio will lead to strong relative performance when markets are weak, the weak absolute performance will mean that product sales are likely to be low anyhow. Vayanos and Woolley (2016) provide a principal-agency model of the relationship between investors and fund-manager which causes them to concentrate in higher beta stocks. A higher than CAPM predicted return on equity for low beta stocks, implies that factor premia are non-linear and that consequently arbitrage is possible. We therefore suggest even if empirical evidence suggests that low beta is robust and unlike other factors tends to be longer lasting (a low beta stock over the last 5 years will likely be a low beta stock in future in the absence of any change in the capital structure) return on equity should not deviate from the CAPM. This is because even if the regulator does not allow additional leverage at the regulated company level the leverage can be increased at the holding company level in order to achieve the lowest cost capital structure.

(iii) Are factor premia robust to data-snooping concerns?

Many academics since Black (1993) have cautioned about the role that data-mining might have in uncovering supposed market anomalies. Strong academic and financial incentives mean that literally 1000s of permutations of different factors are likely to be tested. Setting a confidence level of say 5% will mean even if none of these have any predictive power, we would still expect dozens of apparently significant factors to be discovered.

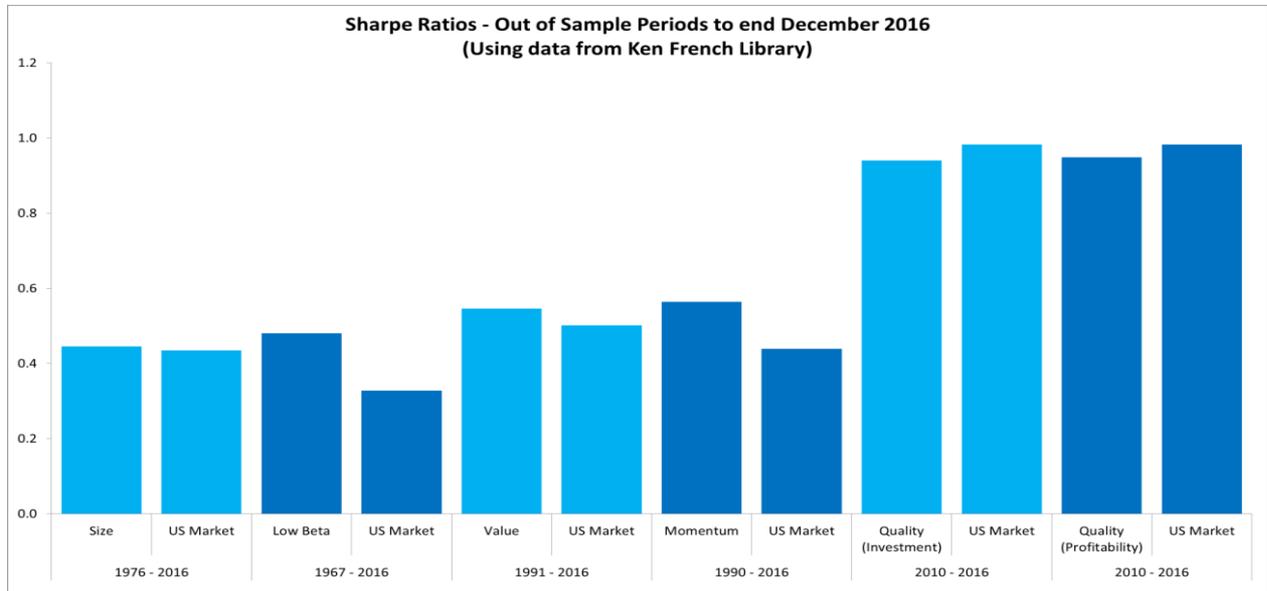
Ensuring that factors are reasonably simple and that there is some form of robust theory behind why the factor would be expected to work should help guard against data-mining or snooping. Out of sample testing may help avoid the dangers of data mining but given the large number of potential factors tested spurious factors may survive even this.

Most assessments of factors follow Fama-French and test factor performance on the basis of long-short portfolios: a long portfolio of equities where the equities which have more of the factor than the market are bought, this is financed by the sale of a short portfolio of equities that have less of the factor than the market. In practice, these portfolios may be difficult to exploit and realised returns may be lower once the costs of borrowing the short portfolio of equities are accounted for. Furthermore the portfolio in reality cannot be constructed at zero cost, as significant collateral would have to be held.

¹¹³ Roll also points out that any mean variance efficient portfolio should satisfy the CAPM equation where that MV portfolio replaces the market portfolio.

Aon (2017) therefore conducts tests of out-of-sample factor outperformance by looking at long-only portfolios versus the broader market. For many investors this is a more implementable strategy, than long-short portfolios. Figure B1 shows the Sharpe ratios of different factors versus the Sharpe ratio of the market since each respective factor was first published by an academic journal.¹¹⁴

Figure: B.1 Out-of-Sample Performance of Factors



Source AonHewitt, Ken French

This suggests that stocks with higher levels of quality have not outperformed since the factor was discovered in 2010. Evidence of a size premium is also weak. Proponents of factor models argue that the discovery dates precede when they might have been published by an academic journal. The idea that value stocks will outperform has been around since at least the 1930s (Graham and Dodd’s famous “Security Analysis” was published in 1934 but the ideas long preceded it). The importance of quality was emphasised in Graham’s 1949 book, “Investment Analysis”.

Perhaps a more valid concern is about investors learning about the value of factors and subsequently increasing their exposure to factor strategies. With a greater weight of money chasing a strategy its outperformance could end up being eliminated.

¹¹⁴ With the exception of size, portfolios are created from the largest 30% of stocks in the Kenneth French database. This is guard against accounting and liquidity issues with the large number of small stocks in the database. Equities are ordered by factor ranking and a portfolio is created of the highest ranking 30% on that criteria and weighting them equally weighted. The portfolio’s Sharpe ratio is than compared to the Sharpe ratio of the largest 30% of stocks. The Size portfolio takes the performance of the middle 40% of the stocks in the database.

(iv) **Do factor premia apply to individual stocks?**

The factor premia literature is generally in the context of portfolio performance rather than individual stocks. If we take value as an example, with 6 month rebalancing, Aon (2017) estimates that turnover for the value factor varies from 44% in the US to 57% in developed ex-US. This compares to turnover on the MSCI AC World index of 2.7%.¹¹⁵ A high value stock will therefore have a slightly higher expected return at a short horizon. However, the higher expected returns of high value stocks are driven by these stocks rebounding and no-longer remaining value stocks. It is the transient nature of factors which both drives the additional return and makes it inappropriate to modify the longer-term cost of capital of a high value stock. At a long-term horizon, the additional expected return from an individual stock might be negligible. Campbell, Lo and Mackinley (1997) survey the evidence on individual stock returns, and whilst they find positive returns at an individual stock level the returns are nowhere nearly as strong as at the portfolio level. This would suggest caution would be needed when applying risk premia identified at the portfolio level to individual stocks.

(v) **Do factor premia vary over time?**

Another consideration is that even if factor premia are valid, how confident can we be about the size of future premia? Ex-post there is significant variation in how well these strategies do (see figure B2). This has inevitably led to a literature on whether factor returns can be predicted. See Arnott (2016, 2017) arguing that valuation ratios (such as price to book) can be applied to factor strategies and Asness (2016, 2017). Arnott argues that the large increase in funds seeking out factors, particularly low beta and quality have made them expensive and that returns will be disappointing. Taking the opposite side of the debate Asness argues that valuation ratios applied to factor premia have not been able to forecast realised factor premia, arguing that the transient nature of factors means that applying valuation ratios to constantly evolving portfolios is futile. Where both authors agree, however, is that historic realised factor premia are not necessarily a good guide to future factor premia, and that diversification of factor strategies is needed given a lack of robustness in any one individual strategy.

¹¹⁵ Turnover is defined as total sales and purchases, so if all stocks were sold and a new set were purchased then turnover would be 200%.

Figure B2: Factor Portfolios can have long periods of underperformance

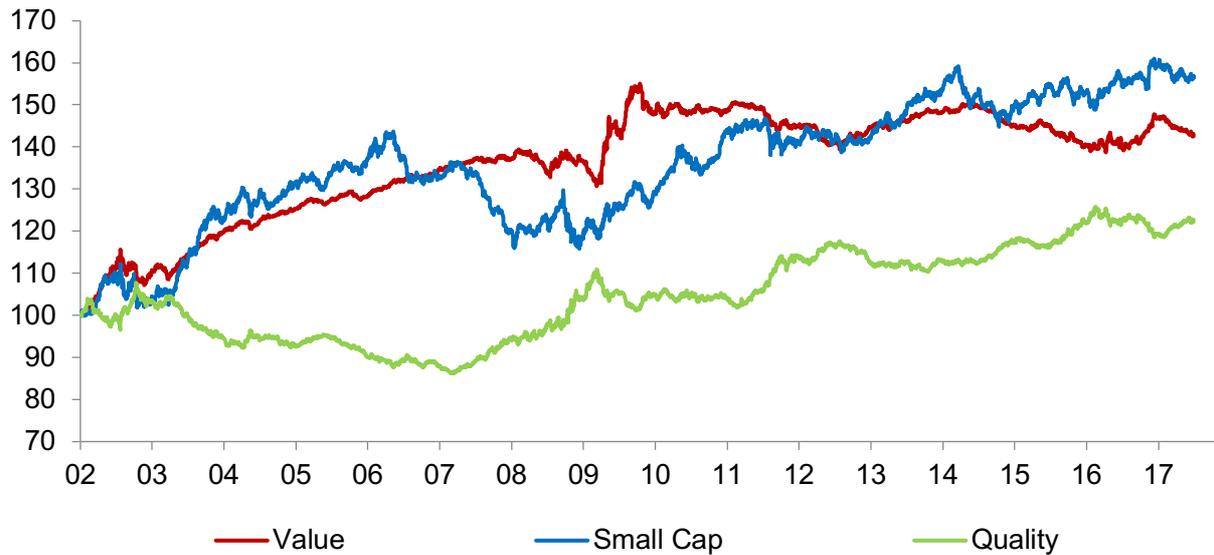
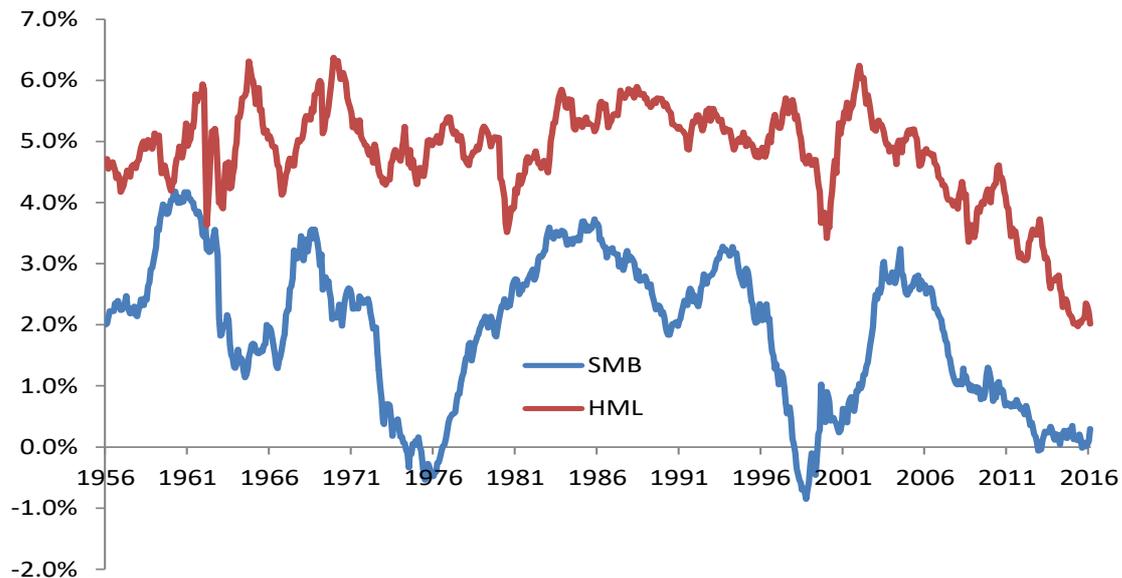


Figure B3: 30 Year Annualised Compound Return from Long-Short “Small minus Big” (SMB) and “High Value minus Low Value” (HML) Excess Returns



Source: Ken French Database, Aon

Fig. B3 shows the 30 year rolling annualised compound premium from being long high value stocks versus low value stocks (HML-“high minus low”) and long small cap versus short large cap (SMB – “small minus big”). In the early to mid-2000s, an investor looking at the history of premia might reasonably believe that on a 30 year horizon returns for the “HML” strategy would be reasonably stable, with the 30 year return mean reverting around the 5% level. However, the 30 year rolling return has subsequently been on a downward trend, raising doubts about whether the value risk premium is stable.

(vi) **Conclusions: Do factors impact the cost of equity?**

Do factors influence the cost of equity? We argue they do not. Firstly, parts (iii), (iv) and (v) suggest that forward looking factor premia estimates might be hard to defend, particularly for individual securities.

Secondly, part (ii) suggests that the two best possible explanations of these factor premia could be behavioural errors by investors rather than equilibrium risk compensation, and contracting problems. The second issue may have been addressed now that funds are increasingly able to take leverage. If we believe that the behavioural biases are likely to persist even though they have become widely highlighted, this might seem to conflict with the argument made in section 4.3 in the main report where we used Cochrane's "cost of tomatoes" to argue the cost of capital used should always be its actual cost rather than what the cost would be if it wasn't for some form of perceived economic distortion (such as secular stagnation, or quantitative easing).

However, we believe that there is an important difference between a financial market price which is directly observable, and clearly applies to risk-free interest rates, and a risk premium whose existence is far from certain.

Thirdly, as Appendix A argues, fund managers get rewarded on the basis of the CAPM alpha of their portfolio not the alpha after adjusting for factor exposures. This means that they're unlikely to differentiate between alpha from factor exposures and "pure" alpha and consequently the cost of equity should be no more adjusted for the additional returns from factor exposures as the additional returns the manager perceives will come from the alpha on the stock.

C The asset beta for an unlevered firm

In a Miller-Modigliani world, the CAPM cost of equity for an equity-only firm is

$$R_F + \beta_A ERP.$$

and adding debt should make no difference to the firm's cost of capital. So the equity-only cost of capital would then be the CAPM-WACC.

But we do not live in an MM world: there are (net) benefits to taking on debt. Assuming that firms do not take on sub-optimal levels of debt, this should mean that, if an equity-only firm gears up, its asset beta, and hence its WACC, should fall, at least initially. Hence, if we can use the right components in the expression above, the (equity-only) CAPM cost of equity should represent an upper bound for the firm's true CAPM-WACC.

But we do need to use the right components; and critically, we need to use the right asset beta. In an MM world, asset beta would be unchanged by gearing, so we have denoted it simply as β_A . But as Korteweg (2010) makes clear, we need to distinguish between the asset beta of a levered firm, denoted β_A^L , and that of an unlevered firm β_A^U . It is clearly the latter that we should use for the (CAPM) cost of equity of an equity-only firm. One challenge, however, is that we typically observe only β_A^L (In fact, we estimate the equity beta of the levered firm, β_E^L , and then infer β_A^L .) We cannot estimate or infer directly β_A^U .

In the technical appendix, we show how to move from β_A^L to β_A^U . For example, with a levered asset beta $\beta_A^L = 0.5$ and a debt beta $\beta_D = 0.1$, we estimate the unlevered asset beta β_A^U to be 0.55. This estimate is, we argue, on the high side of the possible range (given the assumptions that we make to derive the value). But since we are seeking an upper bound for the true CAPM-WACC, this is consistent. In short, we need to uplift the levered asset beta by around 10% (given parameter values typically found in regulatory cases) to arrive at the unlevered asset beta.

Technical appendix: deriving the unlevered asset beta

Following Korteweg (2010), the value of a firm with debt can be written in two, equivalent ways:

$$\begin{aligned} V^L &= E + D \\ &= V^U + B \end{aligned}$$

where V^L is the value of the levered firm; E is the value of equity, D the value of debt; V^U is the value of the (notional) unlevered firm and B is the net benefit of debt: a fictitious security defined as the expected present value of the benefits minus the costs of debt financing. Similarly, the asset beta of the levered firm can be written in two different, but equivalent ways:

$$\begin{aligned} \beta_A^L &= \frac{D}{V^L} \beta_D + \frac{E}{V^L} \beta_E \\ &= \frac{V^U}{V^L} \beta_A^U + \frac{B}{V^L} \beta_B \end{aligned}$$

where β_D and β_E are the debt and equity betas, respectively, of the levered firm; β_A^U is the asset beta of the unlevered firm i.e., the systematic risk of the unlevered assets; and β_B is the systematic risk (the beta) of the levered firm's net benefit of debt financing. Rearranging these equations gives

$$\beta_A^U = \frac{V^L}{V^U} \beta_A^L - \frac{B}{V^U} \beta_B.$$

We are primarily interested in β_A^U . Unfortunately, we do not directly observe this; nor do we observe B , and hence V^U ; and nor do we observe β_B . We only observe V^L and estimate β_A^L . In order to proceed, therefore, we will need to make some assumptions.

First, assume that the systemic risk of the levered firm's net benefit of debt financing, β_B , is equal to the systemic risk of its debt, β_D . Typically, regulators choose this to be in the range 0-0.2. Conservatively, suppose $\beta_D = \beta_B = 0.1$.

Secondly, the empirical evidence points to the net benefit of debt being at most 10% of the value of a firm: see Korteweg (2010) and Van Binsbergen et al. (2010).

Hence with $\frac{B}{V^L} = \frac{1}{10}$, $\frac{V^U}{V^L} = \frac{9}{10}$. Putting all this together gives

$$\beta_A^U = \frac{10}{9} \beta_A^L - \frac{1}{10} \times 0.1.$$

With a levered asset beta of, e.g., 0.5 this gives an unlevered asset beta of 0.55: a 10% increase.

D Choice of price index and adjustment of real returns (Derry Pickford)

Summary

In section 4.2 we argued that the Consumer Price Index (CPI) should be used for assessing the CAPM-WACC. The choice of inflation measure for this is distinct from the question of which measure to use for the purposes of revenue controls and Regulatory Asset Value (RAV) linkage. Although clearly, using a separate price index means that there needs to be careful adjustment of cash flows linked to that alternative index back in to CPI terms. In this Appendix we explore the issues around price index choice in more detail. In particular we ask:

- (i) What are the advantages and disadvantages of CPI, RPI, and other indices?

We argue:

- a) At the time of writing, long-term price indices measured on a consistent basis are only available for the CPI. Long-term CPIH indices are under-construction and we would anticipate will be available before any switch of the Bank of England target to CPIH
 - b) Whilst CPIH is conceptually a better measure of prices than CPI (or RPI and RPIJ), as the Bank of England's target, CPI is the measure which is most focused on by markets and the media.
 - c) The elementary price aggregation methods in RPI create significant and unstable biases between recorded inflation and what it is attempting to measure.
 - d) CPIH and CPI are more comparable to price indices used abroad. Historic international data on real equity returns therefore provide us with information on ex-ante CPI or CPIH based real equity returns.
 - e) Changes to the underlying methodology mean that the RPI is not comparable over time, whereas historical CPI estimates try to match current methodology. Historic equity returns deflated by RPI will therefore have limited informational content about future equity returns deflated by RPI.
 - f) Even in the absence of further changes to RPI, the “wedge” or “basis” between RPI and CPI is likely to be greater and more uncertain than has historically been the case.
 - g) CPI inflation is therefore more likely to be closer to “true” inflation and the difference less variable than with RPI inflation.
- (ii) Why might regulators decide to retain RPI for revenue and RAV indexation purposes, and would CPI still be the best measure for estimating the real cost of equity in such a scenario?

Regulators might retain RPI for revenue controls and RAV indexation. RPI RAV and revenue linkage would mean that the RAR (the regulatory allowed return) – and the RER (regulatory expected return) would then both be RPI linked. This may be more attractive to investors and possibly lower the overall cost of financing, i.e. regulators could set a RPI linked regulatory allowed return so that:

$$RAR^{RPI} \leq RAR^{CPI*} - E(RPI-CPI)$$

where RAR^{RPI} is the regulatory allowed return when measured in RPI terms, RAR^{CPI*} is the optimal regulatory allowed return when measured in CPI terms, and $E(RPI-CPI)$ is the expected RPI-CPI wedge.

Whilst there is no evidence that this is the case from RPI-CPI swaps markets, the low liquidity in this market means we cannot be certain that this would remain the case if there was a big increase in CPI sellers.

- (iii) How should regulators forecast the RPI-CPI wedge?

Whatever approach regulators decide upon, robust estimates of the future RPI-CPI wedge are needed. Key themes in this report have been implementability and defensibility. This suggests delegating the responsibility of forecasting RPI-CPI basis to the OBR (Office of Budget Responsibility). Forecasts of this “wedge” can then be used to derive estimates of the risk-free rate and consequently the CAPM-WACC.

- (iv) How do UK historic real equity return estimates change when measured in consistent CPI terms?

We suggest shifting to using the Bank of England’s long-term CPI series when calculating real returns, rather than Dimson, Marsh and Staunton’s hybrid series. This reduces the long-term real return on UK equities from 5.5% to 5.2%. This brings the UK returns closer to the global equity long-term USD real return of 5.1% (and the UK real USD return of 5.1%).

- (v) How consistent are future measures of inflation likely to be with historic measures?

We outline factors which might increase or decrease the bias between true inflation and measured CPI inflation. This could potentially change the measured real return on equity. Taking the US as an example, we argue that if there were large shifts in measurement error then these would show up in measurements of real equity returns. As MMW (2003) and MPWB (2018) show US real returns have been remarkably stable. This suggests that the shifts in inflation measurement bias must be relatively small. By contrast the difference between UK CPI and UK RPI whilst being 0.14% over the C20th, had risen to 0.7% over 2000-16 (according to Bank of England estimates) and is now projected to be over 1% by the OBR. This suggests that the difference between RPI and true inflation has not been stable over time.

- (vi) Would a potential shift by the Bank of England to CPIH impact our advice?

We argue that CPIH-CPI difference is likely to be second order versus RPI-CPI basis. CPIH is overall a better index, and provided that sufficient back histories of CPIH are provided to create CPIH consistent asset return series that regulators should also shift to CPIH.

- (i) What are the advantages and disadvantages of CPI, RPI and other indices**

This section focuses on the advantages and disadvantages of CPI and RPI. Whilst we conclude CPI is a better measure of prices, it is still not perfect, and RPI has the advantage of being used for index-linked gilts (often referred to as “linkers”). We also discuss CPIH, which addresses one of the short comings of CPI, and RPII which addresses one of the short comings of RPI. Following the

recommendations of the Johnson review the latter no longer remains an official statistic as the focus is instead on introducing CPIH.

No price index is perfect for all purposes. Even for indices aimed at households, different households will lead to a range of inflation experiences. This has led to a plethora of different measures of prices. These tend, although with varying rigour, to focus on the price of current consumption. Economists have long recognised that the utility of households is a function of both present and future consumption. To measure the true purchasing power of money they could therefore include asset prices but this has generally been considered impractical (see Alchain and Klein (1972) and Goodhart (2001)). Whilst this may seem an esoteric academic point, we raise it to emphasise that all calculations of price indices tend to involve a trade-off between theoretical robustness and practical implementability. Furthermore, even if we can agree conceptually on what we wish to measure, the recorded rate of inflation will often be a biased measure of the true inflation rate (see Boskin 1996, Gordon 2006).

Consumer Price Indices attempt to measure how the cost of a broad basket of consumer goods and services evolves over time. The basket is meant to be representative of domestic private consumption (including expenditure by foreigners in the UK but excluding expenditure by UK households overseas). In the UK, the CPI has (when it started to be used for the Bank of England's inflation target) been the same as the Eurostat measure of harmonised consumer prices (HICP) since 2003.

European HICPs differ from the US CPI in that they exclude Owner Occupied Housing (OOH). The best methodology to use in order to estimate OOH is a controversial one, and the lack of consensus amongst statisticians meant that the European statistical agency, Eurostat, has excluded the cost of owner occupied housing altogether from HICPs.

Whilst we recognise inflation experiences will differ, there is an advantage in economic agents using a common measure of inflation and consequently common estimates of real asset returns. If this is the same index as used for the Bank of England's inflation target, currently the Consumer Price Index (CPI), this will have the added advantage of creating medium term stability in expectations of that inflation rate.

As well as being the most focused measure of inflation, as the Bank of England's target measure, it is the measure which is likely to exhibit the greatest long-run stability. We would also argue that it is the measure that currently has the greatest backward and forward comparability. Although the Bank of England's Millennium data sets include estimates for both RPI and CPI going back in history, the CPI estimates are on more a consistent basis whereas the RPI indices match the various switches of methodology since 1947. We also have a double check in the form of ONS estimates, whereas the ONS do not publish estimates pre-1947 for the RPI. This enables us to construct very long-term real return series for UK assets on a consistent basis.

The complete exclusion of owner occupied housing is clearly unsatisfactory and in 2013 the ONS launched CPIH (consumer price index including housing). The ONS's choice of methodology for homeowner costs, "Rental Equivalence" (RE), is the same method used to account for the output of OOH in the national accounts, as well as the approach used in the US CPI and PCE (the Personal Consumption Expenditure Deflator). Rental equivalence uses estimates of what the household would have had to pay if they had rented an equivalent house in the private rental market to provide a weight, and how much these equivalent rents change over time. They are then adjusted for the change in the quality of the houses to derive a price index. As US rents have broadly been stable in real terms over the long-term, US real returns would have not been very different if deflated with a European style CPI.

CPI and CPIH can be contrasted with the RPI and RPIJ (RPI but with greater use of the Jevons aggregation). The RPIs are a hybrid of a consumer price and household cost index (household cost

indices are currently under development by the ONS. Astin and Leyland 2015 show how they can be superior to the CPI for some purposes). RPIs therefore measure the price of (the majority of) regular cash disbursements in a household budget. When it comes to owner occupied housing the RPI measures mortgage interest payments (MIPS), although not repayments of principal or the cost of the deposit. However, one-off costs such as estate agent, legal and conveyancing fees are also included. The inclusion of MIPS created substantial volatility in the difference between the RPI and the CPI. MIPS were only introduced into the RPI in 1975. This makes the index (pre-1975) less volatile versus the CPI. Similarly housing depreciation was only introduced in 1995. Since the RPI is not backwardly revised this makes comparing the RPI over time difficult, if not impossible.

Table D.1 Key differences Between CPI, CPIH, RPI, RPIJ

	CPI	CPIH	RPI	RPIJ
Population base	Includes expenditure by private and institutional households in the UK and also includes expenditure by foreign visitors to the UK.		Includes expenditure by private households but excludes households whose income lies within the top 4% of the income distribution and pensioner households with at least three-quarters of their income coming from state pensions and benefits.	
Expenditure data (or weights)	Final household consumption expenditure (includes foreign visitors to the UK, nursing and retirement homes and university halls of residence).		Weights from Living Cost and Food survey (LCF). Private households only, excluding the highest- income households and pensioner households mainly dependent on state benefits. Includes expenditure abroad.	
Owner occupiers' housing costs (OOH)	OOH is excluded from the CPI	Owner Equivalent Rent	Measured using the payments approach which includes mortgage interest payments (MIPs), house depreciation, buildings insurance and other house purchasing costs, such as estate agents' and conveyancing fees.	
Non-housing coverage	Includes: - university accommodation fees - foreign students' tuition fees - stockbroker fees - unit trust fees - new cars		Includes: - expenditure abroad - council tax and licence fees - used cars	
Official Series Start	1996	2013	1956	2013
ONS Historical Estimates				
Bank of England Historical Estimates	1209	-	1661	-
National Statistic status	Yes	Yes	No	Yes
Elementary aggregate formulae	Jevons – roughly two-thirds of the index Dutot – around 5% of the index Other or weighted formula – roughly one-third of the index.		Carli – roughly one-third of the index Dutot – roughly one-third of the index Other or weighted formula - roughly one-third of the index	Jevons – roughly one-third of the index Dutot – roughly one-third of the index Other or weighted formula - roughly one third of the index
Revisions	Yes		No	Yes
Rounding of Indices	No		Yes	
Internationally Recognised Classifications	Yes		No	
Insurance	Service charge only		Entire premium	

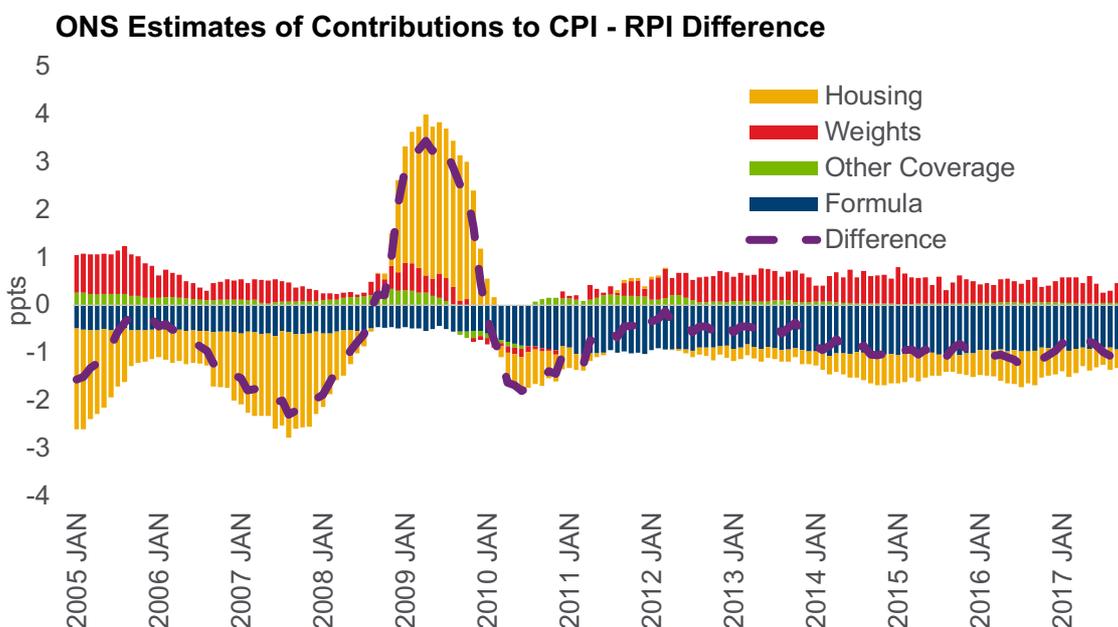
Source Aon and ONS

Since 2005 the most important difference in terms of driving a wedge between the level of the CPI and the RPI (and consequently the average differential) has been the aggregation formula.

Figure D1 breaks down the contributions to the difference between the CPI and RPI as defined by the ONS. The ONS defines these contributions as the rate of inflation for the sub-component multiplied by the difference in weight. This means that for MIPS, the contribution to the RPI is the percentage change in MIPS multiplied by its RPI weight as there is a zero weight for MIPS in the CPI. The corollary of a

zero weighting for MIPS is that collectively other components of the CPI have a higher weight than in the RPI, which typically widens the CPI-RPI gap. An alternative way of accounting for contributions would be to look at the differences in the inflation rate for the subcomponent versus the overall CPI. This approach would diminish both the housing and weights impact on the differential would be diminished.

Figure D1: The CPI-WPI Wedge



Source ONS and Aon

The biggest driver of the difference between RPI and CPI is known as the “formula effect”. This is down to the methods used at the basic level of aggregation of prices (before expenditure weightings are known, this is often referred to as elementary aggregation). The RPI uses the Carli and Dutot formulas whereas the CPI follows international standards in primarily using Jevons (geometric) formula. An advantage of the Jevons formula is that it helps reflect consumer behaviour where there is likely to be significant substitution between very similar products. The Jevons advantage over the Carli index is not limited to this. In particular, the Carli formula can create some peculiar biases particularly when prices are reduced around promotions and then subsequently raised. When the exact product changes after the promotion period, but the prices of the products are linked for price index calculation purposes this can create some well-known incongruities. In 2010, greater differentiation was introduced in clothing price collection where it can be difficult to find matching items. However, this quest for greater matching created more measured dispersion in prices with some shops having low January prices which then rise rapidly after the sales and other shops having more steady prices. This is then compounded by difficulties in finding direct comparisons. This led to very rapid price inflation in certain components of the RPI relative to the CPI. For example, the CPI estimates that women’s strappy tops from independent shops have risen in price by 20% between 2010 and 2016 whereas the RPI estimates this price increase has been 400%.¹¹⁶ Whilst this example is the most extreme it is not unique. As the

¹¹⁶ <https://www.statisticsauthority.gov.uk/wp-content/uploads/2017/02/APCP-T1702-RPI-and-CPI-a-tale-of-two-formulae.pdf>

underlying source data for the CPI and RPI is often the same, future improvements to the process for CPI could create additional (likely upward) volatility in the RPI.

Given these problems it was unsurprising that the Johnson Review (2015, p. 13) was critical of the RPI's use of the Carli aggregation:

This review concurs with the Authority's conclusion that the use of the Carli is statistically flawed and can result in an upward bias in recorded inflation. This is not essentially, as some have argued, because the Jevons method takes account of substitution between goods. Rather, it is because there are basic statistical flaws and biases in using the Carli formula. As we stressed above it is generally hard in this area to come to absolute conclusions. But it is our strong view that the use of the Carli is inappropriate and that the RPI is upwardly biased because of its use.

The problems with Carli aggregation are of course well-known and RPIJ (a national statistic from November 2013 until it was discontinued in January 2017) was launched to address these problems.

The Johnson review (p. 13) argues:

RPIJ corrects the inadequacies of the formula used in the RPI and keeps the rest of the methodology untouched. But it is not just the use of the Carli which is problematic in the construction of the RPI as a measure of consumer price inflation. Issues with the data source of the weights, population coverage and treatment of some goods (like insurance and owner occupiers housing costs) make the RPI less suitable as a measure of overall inflation. RPIJ is problematic for all the same reasons.

We think that Johnson may overstate the coverage issue. The RPI excludes certain households who are unlikely to be representative of the broader population. CPI by accounting for all households including the top 4% by income ends up being more "plutocratic" (see Astin and Leyland, 2015) i.e., it reflects the experience of the median household less well. The treatment of owner occupied housing is more troubling. The majority of UK mortgages are either floating or have relatively frequent resets. Any normalisation of UK interest rates could see mortgage payments rise significantly. Over late 2008 to early 2010, MIPS fell as interest rates were cut sharply. However, any reversal of those changes would mean very rapid MIPS inflation and an increase in the RPI-CPI wedge.

Sensitivity to MIPS, increased granularity of price collection impacting the formula effect therefore make the wedge both higher and more uncertain than historically. We return to this in part (iii) of this appendix.

The Johnson Review's 4th and 5th recommendations (p15) are that:

4. ONS and the UK Statistics Authority should re-state its position that the RPI is a flawed statistical measure of inflation which should not be used for new purposes and whose use should be discontinued for all purposes unless there are contractual commitments at stake.

5. Government and regulators should work towards ending the use of the RPI as soon as practicable. Where they decide to keep using it the UK Statistics Authority should ask them to set out clearly and publicly their reasons for doing so. Where the Authority judges the continued use of the RPI to be inappropriate, it should say so.

Whilst we agree that RPI has some serious shortcomings, there is demand for its continued use. Just as recipients of RPI linked pension benefits will not want to shift to CPI linked payments without some form compensation the same will apply to regulated entities. Where revenue and RAVs are linked to RPI-X, the X may have to be reduced by more than the expected wedge if revenue and RAV is linked

instead to a CPI-X formula. We explore this more in part (ii) and ask whether that would alter our view that the CAPM-WACC should be estimated in CPI terms if regulators decide that consumer interests are best served retaining RPI for revenue and RAV indexation.

(ii) Why might regulators decide to retain RPI for revenue and RAV indexation purposes and would CPI still be the best measure for estimating the real cost of equity in such a scenario?

Our slight preference is to follow the lead of OFCOM and OFWAT and use either CPI or CPIH for RAV and revenue indexation. However, this is not something we have been asked to give advice on and we can also see strong arguments for retaining the use of RPI for this. This part explores how to estimate cost of equity if RPI is retained for revenue and RAV indexation.

The RPI is likely to remain with us for many decades. HM Treasury issue large amounts of long-dated index-linked gilts and defined benefit schemes have RPI linked liabilities. Some regulated entities have also chosen to fund themselves through RPI linked securities. Many of these liabilities have maturities in excess of the review period meaning that they will have RPI liabilities for many years. By contrast there is a lack of a CPI linked debt market and only very limited liquidity in CPI swaps. These (very illiquid) markets would suggest that it would be cheaper, at the margin, for companies to have CPI liabilities (they could do this by still issuing RPI debt but swapping it in to CPI payments) if there was substantial CPI issuance however this cost advantage would be unlikely to remain.

The lack of a liquid market in CPI-linked securities may seem surprising, particularly as almost all defined benefit pension schemes will have some level of CPI exposure¹¹⁷. Even investors without inflation-linked liabilities will care about real purchasing power and securities with CPI linkage will for many investors provide a closer match to their need for real purchasing power than RPI-linked securities.

In the long-run, we would anticipate the development of a CPI-linked debt market and attribute the lack of a market currently as essentially a chicken and egg problem. Whilst market participants have encouraged the DMO to issue CPI-linked securities, the DMO has been hesitant partly due to fears that it will split liquidity in the linker market. However, if CPI linkage for regulatory contracts becomes commonplace we believe that there will be a greater supply of CPI-linked corporate bonds. Price discovery would likely be aided if there was a CPI-linked gilt market from which these could be priced. However, in the absence of liquid CPI linked debt markets, regulators may assess that offering a RPI linked return through revenue and RAV linkages to RPI would possibly lower the overall cost of financing. In such a scenario:

$$RAR^{RPI*} \leq RAR^{CPI*} - E(RPI-CPI)$$

where RAR^{RPI*} is the lowest regulatory allowed return that still ensures financeability when measured in RPI terms, RAR^{CPI*} is the lowest viable regulatory allowed return possible when measured in CPI terms, and $E(RPI-CPI)$ is the expected RPI-CPI wedge.

Regulators may also wish to postpone switching to CPI as there is a possibility that HM Treasury may well shift to CPIH for the Bank of England's target. One of the advantages of CPI is that it should in

¹¹⁷ Almost all UK DB pension funds where benefits accrued between 1988 and 1997, will have liabilities linked to CPI. The Guaranteed Minimum Pension component (GMP) had to rise by CPI capped at 3% and floored at 0%. Components over and above the GMP are often linked to CPI capped at either 5% or 2.5% depending on when the benefits accrued. With implied inflation swaps at 2.8%, there is often little incentive for pension funds to hedge this risk.

the long-run be more stable. However, a shift to a CPIH target could reduce this stability and make CPIH the series with the greatest long run stability. Rather than risk switching the price index twice, waiting and seeing which index HM Treasury decide to focus on may be optimal (see part vi).

Even if RPI is retained for indexation we argue that that WACC-CAPM can still be assessed in CPI terms. RPI cash flows can be converted into CPI ones or the CPI WACC-CAPM curve can be converted into a RPI WACC-CAPM. A set of RPI based discount rates applied to RPI linked cash flows will give approximately the same net present value as CPI linked cash flows with a CPI based discount rate.

Why do we think that, the cost of equity at the first stage has to be calculated in CPI terms? We explain this in more detail in part (iv) but essentially this can be summarised in two key conclusions:

- Firstly, we can calculate a CPI based historic real return on a consistent basis, and then augment this with evidence on international returns which are essentially CPI based real returns.
- Secondly the CPI based cost of equity is likely to be more stable than the RPI based real return.

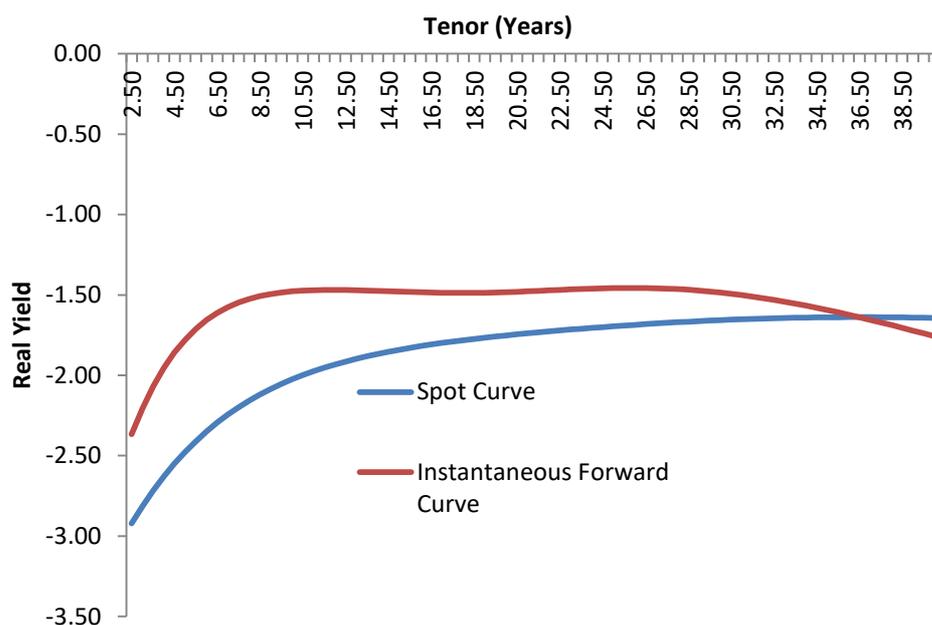
Whatever choice a regulator makes with regard to indexation forecasting the RPI-CPI wedge is essential and we turn to this in the next section.

(iii) How best to deriving Risk-free Real Rates: Forecasting the RPI-CPI wedge.

Participants in UK fixed income markets often refer to the yield on RPI index linked bonds as being *the* real yield. Of course, there is no unique real yield curve but a whole set of real yields curves corresponding to each measure of inflation. We refer to the CPI adjusted real interest rate over a period of t years as r_t^{CPI} , and this is our preferred measure of the real risk-free rate. Similarly we refer to RPI real rates, r_t^{RPI} , which can be observed directly from the index-linked gilt market.

Together with nominal bond rates this gives us RPI break-evens (the average realised inflation rate where the returns from holding a conventional gilt to maturity will be the same as the return from holding the linker). These can be considered as a risk-neutral measure of expected future RPI inflation.

Figure D2: RPI Based Real Yields (as at end of August 2017)



Source: Bank of England

Index linked gilts are not the only source of information about r_t^{RPI} . There is also a reasonably liquid market in RPI zero-coupon swaps, which can be compared to either nominal swaps or government bonds. This market enables participants to hedge inflation risks without locking out significant amounts of capital. Typically this will price inflation slightly higher than break-even inflation. Currently swap break-evens are around 30 bps higher at a 10-year tenor than gilt linkers. This is because investors who wish to hedge RPI risk with linkers will find that this requires large amounts of collateral to do so (and requires shorting nominal bonds). Swaps therefore have a convenience and capital advantage. Because the arbitrage between the swaps market and the linker market is capital intensive the prices are not brought in to line. Given the greater depth of the linker market we propose that real yields and (risk neutral) expected inflation should be derived from this market. However, we still require a robust mechanism to adjust observable RPI based yields into CPI based real yields.

Four alternative mechanisms suggest themselves:

- (i) RPI-CPI basis swaps

Whilst there are no CPI linked gilts, there is a limited market in CPI-RPI basis swaps. On the basis of a limited survey of dealers, a 10Y RPI-CPI basis swap is currently quoted at 60 bps. This implicitly prices CPI at around 80 bps over the middle of the Bank of England's target band and only 20bps below the top end of its target range. We believe that the volumes going through at these prices are extremely limited, and consequently both transaction and quote-based measures would be biased.

Table D2: Market Pricing of Inflation

A	10Y Zero Coupon RPI Inflation Swap:	3.40%	(Source: Bloomberg)
B	10Y Gilt RPI Linker Break-even:	3.10%	(Source: Bloomberg)
C	10Y RPI-CPI Basis Swap:	0.60%	(Source: Dealers)

D	Implied 10Y Zero Coupon CPI Inflation:	2.80%	(A)-(C)
E	Implied 10Y CPI Inflation via Gilts	2.50%	(B)-(C)

Source Bloomberg, Aon (as at 18th September 2017)

(ii) Consensus expectations

There are numerous market surveys of economist expectations for CPI-RPI basis. The Bloomberg survey currently expects a 90 bps differential over 2017 and 2018 falling to 60 bps by 2019. However, the number of participants providing both RPI and CPI forecasts is limited and therefore could be vulnerable to gaming if it was used formulaically. Moreover, participants in surveys are not directly rewarded for the accuracy of their forecasts and resources devoted to the forecasting process are typically lower than public bodies. If individual respondents are asked for forecasts they may face incentives to either herd and be as close to consensus as possible or to deliberately be an outlier in order to generate publicity. The use of consensus forecasts also requires resources to clean the survey for stale forecasts which have not been updated. For these reasons we don't think that consensus expectations is a viable route.

(iii) Office of Budget Responsibility forecasts

A counterpart to delegating the choice of inflation measure to the Bank of England would be to use the forecasts of the Office of Budget Responsibility (OBR), the official body responsible for independent and authoritative analysis of the UK's public finances. As both the RPI and CPI are directly linked to different forms of expenditures and taxes, the OBR has a strong incentive to invest resources in providing accurate forecasts of both. Furthermore, its *raison d'être* is to be impartial and objective. Moreover, its forecasts, unlike private sector ones, are unlikely to suffer from incentives to be biased. Although the OBR has so far overestimated the RPI-CPI differential, it operates a forecast evaluation mechanism which should uncover and correct any systematic biases in the forecasting process.

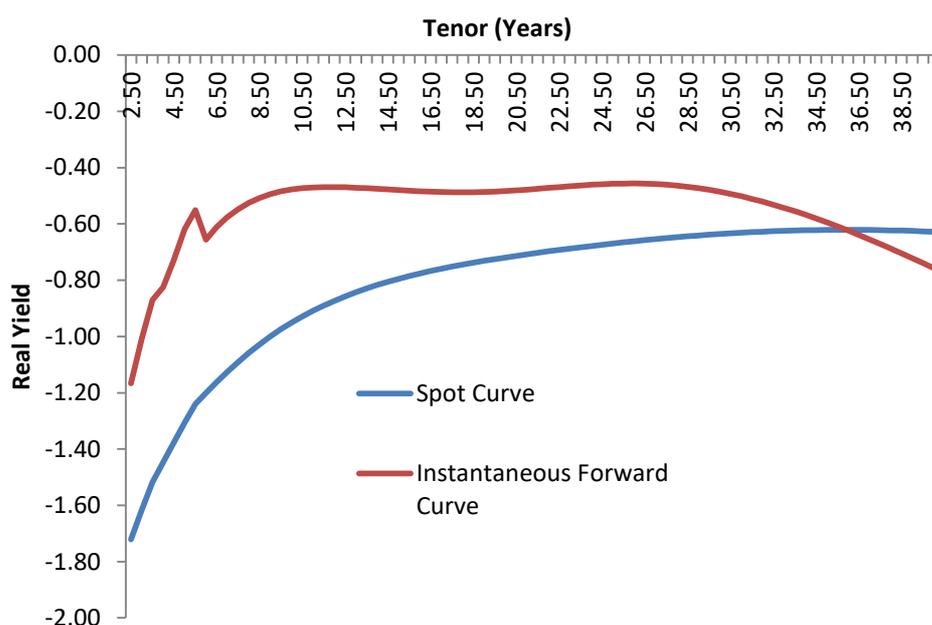
Table D3: Office of Budget Responsibility Implicit Forecasts of RPI-CPI Wedge

RPI-CPI Inflation Wedge for Year															Arithmetic mean of forecast RPI-CPI differential
Forecast	2009	2010	2011	2012	2013	2014	2015	2016	2017	2018	2019	2020	2021	2022	
June-2010	-2.70	1.20	0.90	1.00	1.20	1.40	1.50								0.64
November-2010	-2.70	1.20	0.60	1.10	1.10	1.30	1.60								0.60
March-2011	-2.70	1.30	0.90	1.10	1.50	1.60	1.80								0.79
November-2011		1.30	0.70	0.60	0.80	1.10	1.60	1.80							1.13
March-2012		1.30	0.70	0.40	0.40	0.60	1.60	2.00							1.00
December-2012			0.70	0.40	0.50	0.40	1.10	1.40	1.70						0.89
March-2013			0.70	0.40	0.40	0.40	1.10	1.60	1.90						0.93
December-2013				0.40	0.50	0.60	1.20	1.60	1.70	2.00					1.14
March-2014				0.40	0.40	0.70	1.20	1.60	1.80	1.90					1.14
December-2014					0.40	0.90	1.00	1.20	1.40	1.60	1.60				1.16
March-2015					0.40	0.90	0.80	0.90	1.10	1.20	1.10				0.91
July-2015						0.90	0.80	1.00	1.20	1.30	1.20	1.20			1.09
November-2015						0.90	0.90	1.00	1.10	1.30	1.20	1.20			1.09
March-2016						0.90	1.00	1.00	0.80	1.20	1.10	1.20			1.03
November-2016						0.94	1.08	0.86	1.04	1.08	1.14	1.20			1.05
March-2017						0.94	1.09	1.28	1.25	1.14	1.13	1.16			1.14
November-2017							1.08	0.93	0.95	0.92	0.89	0.94	0.97		0.95
Outturn of the RPI-CPI wedge*	-2.70	1.32	0.74	0.38	0.48	0.91	0.94	1.08							0.39
Arithmetic mean of forecasts before year			0.75	0.95	0.92	0.93	1.37	1.41							1.05
Arithmetic mean of forecast error			0.01	0.57	0.44	0.02	0.43	0.33							0.30

Source: Office of Budget Responsibility, Authors Calculations

Using the OBR's RPI-CPI wedge as of March 2017, we can derive a CPI real risk-free rate curve in Fig D3.

Figure D3: CPI Based Real Yields (as at end of August 2017)



Source: Bank of England, Office of Budget Responsibility, Authors Calculations

Based on data as at the end of August 2017, the OBR has overestimated the wedge by 30bps. The main driver of this error has been assumptions about normalisation of interest rates causing mortgage interest payments to rise more quickly than other prices. Perhaps a more pressing problem with using the OBR’s forecasts is that they are only available twice annually. This can mean that forecasts of the wedge can jump, which would create sudden movements in estimates of r_t^{CPI} . These jumps have at times been as much as 70 bps for current year basis, although jumps are typically smaller for longer horizon estimates of the differential.

(iv) Historical Differentials

Taking the historical RPI-CPI differential will have the advantage that past errors will eventually get compensated for. The approach could be analogous to the “dragging anchor” approach used for setting the risk-free rate. However, we would caution that some of the disadvantages of the dragging anchor approach for setting r_t^{RPI} would also apply as it could set a risk-free rate which was clearly wrong.

(iv) How should nominal asset price returns be deflated in order to derive real returns and what does this mean for estimates of historic real returns and measures of real returns going forward?

Errors in measuring past inflation rates will result in biased estimates of historic real returns. The problem is particularly pronounced for measuring local currency real returns of countries that have experienced episodes of hyperinflation. As discussed in MMW (2013), Dimson Marsh and Staunton (2002) use RPI inflation from 1947 onwards and a much more narrowly defined cost of living index in earlier data, which rose less rapidly than the Bank’s estimates of RPI or CPI. DMS in 2016 updated their methodology to account for the perceived problems with the RPI by using CPI from 1988 onwards.

Rather than make ad hoc adjustments we believe that it makes sense to use an index which is consistent. The RPI is a “protected” index. Once a value is published it is not subsequently changed. Whilst this means that we don’t need to worry about subsequent revisions changing values it does mean that the index is not backwardly compatible. Any changes to the underlying methodology do not get reflected

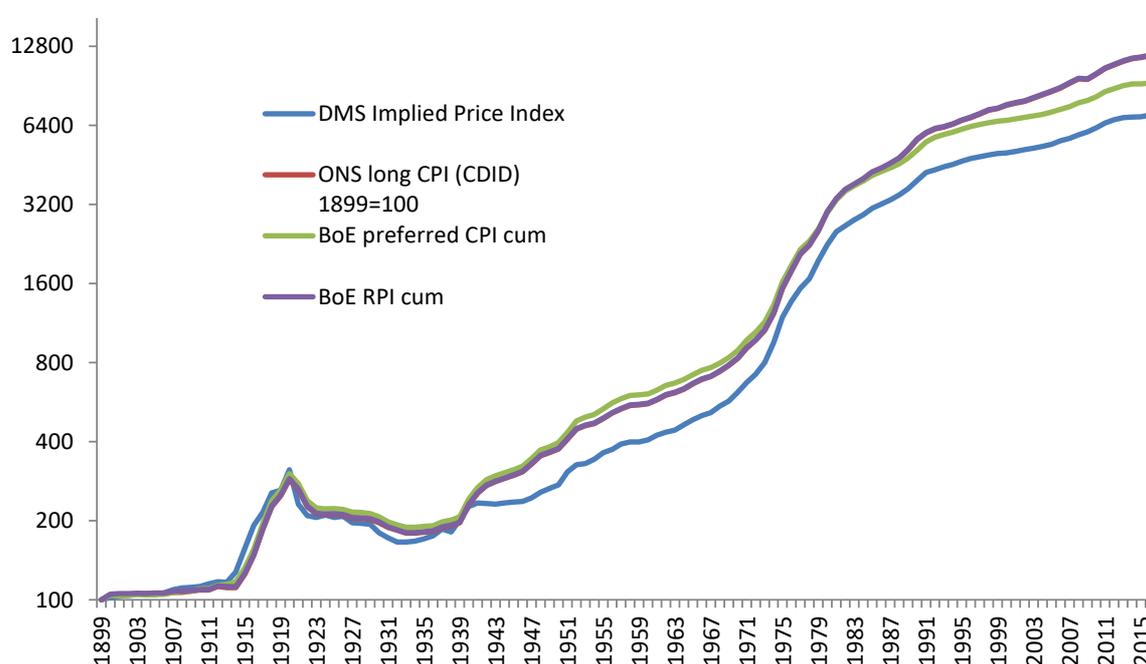
in the historic data. We have argued that for consistency we should follow the Bank of England’s long-term price measures, since the Bank is more focused on providing long-term series which provide relevance to analysis of current economic conditions. Retaining DMS’s nominal total return estimates but switching to these alternative price measures would lower the estimated CPI-based geometric real return to 5.23%.

Table D4: Compound Annualised Returns and Inflation Rates for the UK

	Inflation				Real Equity Returns				
	DMS	ONS Long-term Price Measure	BoE preferred CPI	BoE RPI	DMS UK Real Return, Sterling	DMS UK Real Return, US Dollars	DMS UK Nominal Deflated by BoE CPI	DMS UK Nominal Deflated by BoE RPI	DMS World (inc US) Real Return, US Dollars
1899 to 2016	3.70	4.16	3.95	4.17	5.48	5.07	5.23	5.07	5.05
1899 to 2000	3.96	4.39	4.25	4.39	5.88	5.61	5.58	5.60	5.36
2000 to 2016	2.09	2.75	2.06	2.76	2.97	1.78	3.01	1.78	3.16

As Table D4 shows, correcting DMS’s price index brings the real return of UK equities closer to the average historical experience across the world of 5.1%. Given that large number of stock markets experienced being nearly wiped out during World War 2 or widespread nationalisation without compensation in various revolutions, the global realised historical return may understate ex-ante expected returns.

Figure D4: Four different measures of long-term Consumer Prices in the UK



Intriguingly the real USD return on UK equities was also 5.1%. (Sterling depreciated in real terms versus the USD.) To the extent that equity market returns are determined globally (which is what we would anticipate with free movement of capital) and the average real appreciation/depreciation of Sterling might have been anticipated, then this would suggest that UK returns were in line with the global historical average.

(v) How consistent are future measures of inflation with historic measures?

Provided the bias between measured inflation and the true rate of inflation is stable, the implications for cost of capital estimation are fairly inconsequential. If inflation is overstated by x%, then historic real returns and the risk-free rate will be underestimated by x%, but so will measured future returns.

More complex issues arise if the bias varies significantly over time. For example, improved methodologies may make measured inflation closer to “true” inflation, reducing the upward bias identified in previous studies (e.g., Boksin). Even without changes in methodology, changes in the economic environment can impact the level of bias. For example, changes in the pace of product innovation, may not be perfectly captured in price indices. Biases can also be introduced if the price elasticities of demand for the goods experiencing relative price increases differ from those experiencing relative price declines.

Provided a CPI/CPIH is used we’re not too concerned. As demonstrated in MMW the measured real return on equities in the US is relatively stable. If there were large shifts in inflation biases over time we would anticipate that this would generate greater variation in measured real returns on equities. Since we do not observe this we are reasonably relaxed that the biases remain relatively stable over time. By contrast we are more concerned if RPI is used to estimate real returns. The difference between UK CPI and UK RPI was 0.14% over the C20th. This had risen to 0.7% over 2000-16 (according to Bank of England estimates). It is now projected to be over 1% by the OBR. This suggests that the difference between RPI and true inflation has not been stable over time and we would therefore be concerned that RPI measured real return on equity would not be stable.

(vi) If the Bank of England moves to CPIH what dilemmas could this pose for regulators?

In March, the Office of National Statistics (ONS) decided to make CPIH (Consumer Prices including Housing), their “preferred” measure of consumer prices. It is now the most prominent measure of inflation on its website and the ONS’s inflation releases. If the ONS move is a precursor to a change in the target rate of the Bank of England¹¹⁸ then, according to our outsourcing recommendation this would suggest that regulators should also shift to CPIH for the calculation of the WACC-CAPM as well. Since CPIH is conceptually better than CPI (it covers an important cost for households) this would likely move measured inflation closer to true inflation. However, at the time of writing no long term CPIH series are available and consequently we’re unable to calculate long-term CPIH realised equity returns for the UK. However, we do not think this would represent an insurmountable barrier to shifting to CPIH. US equity returns are more comparable to CPIH than the UK’s CPI and therefore this will improve rather than reduce the relevance of US asset return data. Secondly we understand that the ONS are compiling a long-term CPIH series and would anticipate the Bank of England would do something similar in advance of any target change. Thirdly, as we explain in the rest of the section we believe that the difference between CPIH and CPI will be much smaller than the RPI-CPI differential.

The CPIH can be thought of as being approximately a composite index made up of 83% of the Consumer Price Index (CPI) and 17% of an index measuring the ONS’s estimate of the cost of Owner Occupied Housing (OOH). Official estimates of CPIH go back to 2005 and over this period OOH inflation has been 0.55% pa lower than CPI and CPIH 0.12% pa lower than CPI. However, this period may not be representative. “Backcast” OOH data goes back to 1988 and over the period 1988-2005 OOH(RE) (OOH with rental equivalence) was 0.85% p.a. higher than CPI and this would probably have

¹¹⁸ CPIH might also be used for the indexation of State benefits and pensions, and utilities’ price reviews. OFWAT has indicated that if it moves away from RPI linkage that it might well use CPIH rather than CPI.

translated into CPIH being 0.17% p.a. higher than CPI. Forthcoming research from the ONS will provide more insight on this.

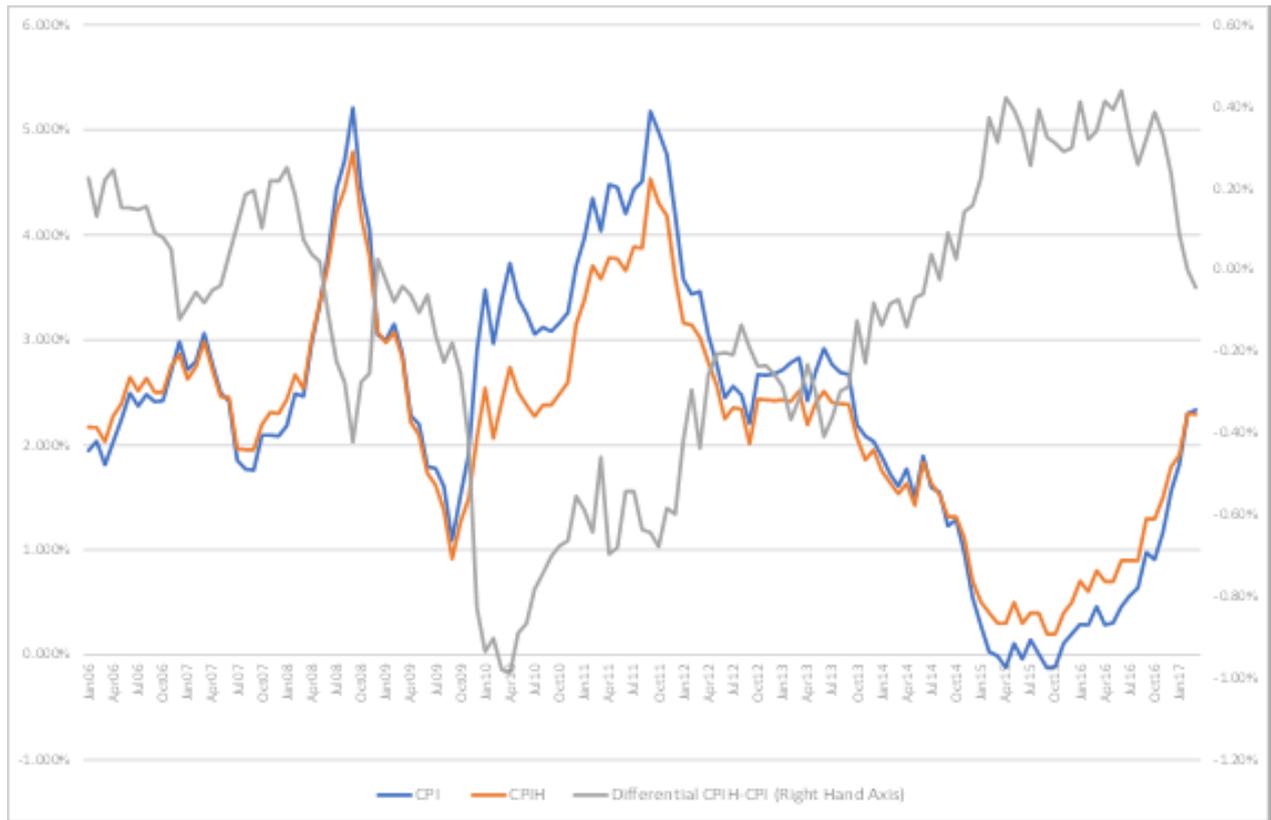
The longer-term differential between CPI and CPIH will therefore predominantly be a function of real (inflation-adjusted) rents. Whilst any major liberalisation of the UK's planning system could see real rents fall, if the current under-supply continues, and the labour market normalises, we could see real private rents rising at a similar rate as in the 1980s and 90s.

Even if CPI and CPIH were expected to average out at the same rate in the long-term there could be implications for regulated entities with CPI linkage. At present, uncertainty about the level of the CPI in the future is reduced because it is the Bank's target. However, if the Bank were to shift to a CPIH target and there were to be persistent differentials between CPIH and CPI then uncertainty about future levels of the CPI would be increased. This could potentially expose regulated entities to undiversifiable risks.

There are some additional concerns about CPIH. The series for OOH is based on administrative data from the UK's Valuation Office Agency (VOA). The ONS acknowledges that their previous estimates of rental-equivalence OOH (OOH-RE) have been inaccurate. The UK Statistics Authority decided back in August 2014 that the CPIH index (which had only been introduced in 2013) should therefore be "de-designated" a national statistic, with the designation only being returned in July 2017. There are also some practical difficulties in that although ONS use stratification (dividing the housing stock into different groups with similar characteristics), finding like-for-like homes in the rental sector can be difficult. Furthermore, the nature of home ownership is very different from a tenancy, with the former providing greater certainty in continued occupation and being able to customise a house for their own preferences, creating an arguably more valuable flow of services.

Whilst alternatives exist to rental equivalence (such as the net and gross acquisition approaches) none is without potential drawbacks. Overall we believe that some form of measurement of OOH is better than completely ignoring it and consequently the eventual shift to CPIH will represent an improvement.

Figure D4: CPIH-CPI Differential



We therefore conclude that ex-ante CPIH and CPI differentials will be sufficiently small that our main conclusions about the nature of historic returns will not be fundamentally changed. We understand that the ONS are in the process of creating long-term historical estimates of CPIH and we would anticipate the Bank of England would do the same in advance of any shift of their target measure. This would enable us to recreate long-term equity return series for estimating ex-ante equity returns which would be consistent with the new inflation measure.

E Updating evidence of the historic market return

Table 1 below shows estimates of long-run returns on a range of different markets from the Dimson, Marsh and Staunton database. It updates a similar table in Smithers and Wright (2013).

The addition of 4 years' worth of data unsurprisingly makes relatively little difference to long-run averages compared to the numbers reported there. The intervening years have boosted returns over the period since 2000 somewhat, but actually there is a somewhat larger change in the long-run real return reported by DMS due to a downward revision in their long-run measure of inflation (we are unclear why).

Notably, however, once we adjust for our preferred measure of CPI inflation, as described in Appendix D, the resulting long-run average is identical, to two decimal places, to the number we reported (on the basis of returns to 2012) using the DMS inflation rate as reported in that vintage of data. A salutary reminder of the imprecision of long-run estimates.

Table 1. Summary of DMS Long-Term Real Returns on UK and Global Stock Markets

	Compound Average Real Returns, % p.a.					
	UK, £ (DMS)	UK, £ (CPI)	UK, \$	World, \$	World, ex US, \$	US, \$
1899 to 2016	5.48	5.23	5.07	5.05	4.33	6.39
1899 to 2000	5.88	5.58	5.61	5.36	4.63	6.79
2000 to 2016	2.97	3.01	1.78	3.16	2.43	3.87

By long tradition, since the CAPM relates to expected returns, market return assumptions are normally expressed in terms of the arithmetic average return. This issue was also discussed at some length in both MMW and in Smithers and Wright (2013). In that discussion we concluded, again, that rather than calculate arithmetic averages directly (which can generate spurious differences, especially when returns are affected by exchange rate fluctuations), it is more appropriate to work from geometric (compound) average returns and add an adjustment of 1 to 2 percentage points, depending on the extent to which regulators wish to take account of serial correlation of returns.

MMW's original proposed estimate for the assumed long-run return on market equity was 5½% (compound average), with a range of 6½-7½% for the arithmetic average. The table above does not give us much cause to bring the estimate of the compound average rate down, although a figure of not much more than 5% on a compound basis appears increasingly persuasive, given the striking similarity of long-run UK returns to long-run global returns.

We would, however, argue that the case for an adjustment to arithmetic averages as large as 2 percentage points (which was implied by the upper end of MMW's range) is distinctly weakened if regulators wish to set returns on a consistent basis at a relatively long (e.g., 10-year) horizon, given that (as noted in MMW) long-horizon returns have distinctly lower volatility than would be the case in a random walk stock market.¹¹⁹

Thus, in light of both these issues, we suggest a modest downward adjustment of the original range proposed by MMW, to a range of 6-7%, primarily reflecting a smaller adjustment from geometric to arithmetic returns.

¹¹⁹ Indeed, Derry Pickford notes that finance practitioners make little or no use of arithmetic averaging.

F Beta estimates (Phil Burns)

This appendix sets out recent analysis on raw equity betas, as well as unlevered and re-levered betas, for the UK regulated utilities. Whilst there is no “standard approach” adopted by the regulators, the methodology used to estimate the beta set out in this paper broadly reflects the approaches taken by regulators and by the CMA in recent determinations, and as recommended by Mason, Miles and Wright (2003). The practical application of the paper is on the betas for the network utilities in energy and gas. We also briefly discuss in this paper the approach set out in another appendix to this report, on the GARCH estimation of betas reported by Wright and Robertson.

We discuss plausible interpretations of the estimated betas, in relation to the risk profile of regulated utilities compared with other assets such as riskier regulated businesses, unregulated businesses, and the risk-free asset. We conclude on whether or not the beta assumptions made in recent regulatory decisions can be supported by robust evidence.

Introduction

As noted in the main report, there are many ways to estimate the cost of equity. Within the family of asset pricing models that originated with the CAPM, there are range of multi-factor models (e.g. the Fama-French three factor model). A key conclusion of the report is that for the purposes of regulatory price setting, the simple CAPM represents the most appropriate model to use, largely because it is so well understood by all stakeholders, which assists in promoting transparency and stability of regulatory practice over time.

However, the estimation of the Beta factor in the CAPM framework presents a number of challenges for the regulator. These include the following:

- The multitude of decisions that need to be made around the method of estimation, this includes: data frequency, sample duration, regression method, use of Bayesian adjustments, method for de-levering, to name a few.
- The difficulty in find pure-play comparators and need to balance sample size of companies against the suitability of each comparator.
- The existence of multi-factor models indicates that there may be other drivers of expected equity premiums and, in this case, the CAPM Beta factor will need to capture any impact of these factors.

All of this suggests that there is no single right answer for the value of the Beta that is estimated, so regulators have always exercised their judgement in weighing up the evidence before them before determining the beta that enters into the allowed cost of equity. In short, there is no single truth that pertains for the beta, but rather a body of evidence that regulators should draw upon to inform their judgement.

What does the Beta measure?

One of the key insights provided by the CAPM is that only non-diversifiable risks should matter to investors – that is, investors should only expect to be compensated for risks that they cannot eliminate by holding a well-diversified portfolio of assets. Beta captures how market movements affect the company in question, which could derive from one of several channels:

- *Investor asset allocation.* A big driver of movements in the market is investors adjusting their asset allocations between equities, bonds, commodities etc. This reflects long-term decisions and short-term perceptions¹²⁰. This factor will create a systematic link between the market and the utility since when investors move out of equities into another asset class they will reduce holdings in all equities including utilities. The correlation may not be one for one since at the same time other investors may switch from higher risk equities into lower risk equities. Nevertheless there will be Beta component relating to asset allocation.
- *Economic cycle.* Information on the state of the economic cycle will drive market movements as it affects demand levels, exchange rates, input prices and interest rates.
- *Long-term growth prospects.* The market will respond to changing confidence about growth, which may result from expectations of the impact of new technology. To the extent that utilities' scope for cost outperformance derives from new technologies there will be a Beta component through this route.
- *Regulatory uncertainty.* There will also be an impact on utilities that comes through regulatory risk and the perception that the degree of regulatory 'toughness' will correlate with economic conditions.
- *Political uncertainty.* Finally the market will move with changes in political uncertainty. So, for example, the prospect of a government that is less "pro-business" would move the market and the shares in privatised utilities.

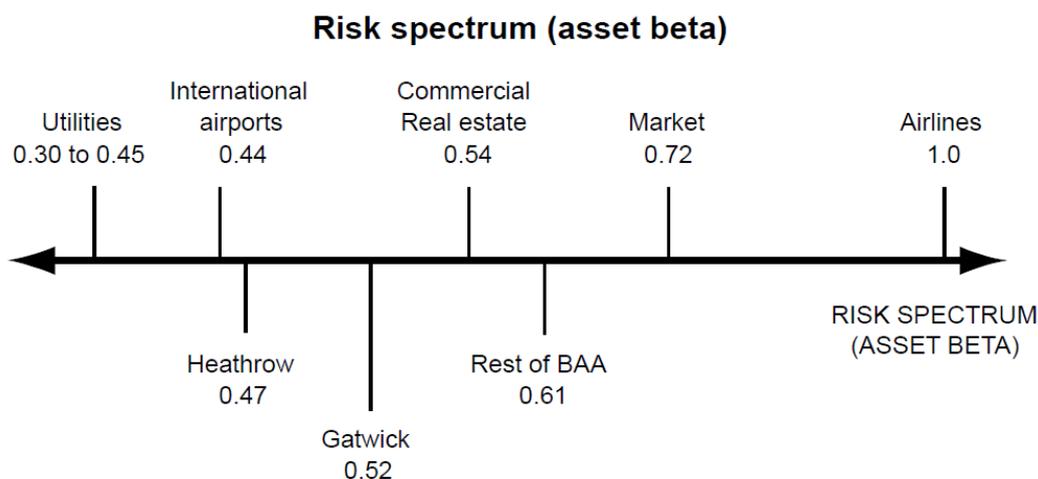
It is worth noting that whilst it is common to hear practitioners discuss diversifiable and non-diversifiable risks as binary, non-overlapping concepts, in practice, the risks faced by a business lie on spectrum in terms of whether (and by how much) the risk can be diversified by holding a portfolio of investments. In practice very few risks lie at the end of the spectrum where they can be completely diversified.

An important influence on the impact of these effects on the business – and therefore on Beta - is the regulatory regime itself. A very low-powered cost-plus regime will tend to diminish the impact of the economic cycle, and possibly also the regulatory and political uncertainty. A higher-powered incentive based regime might be expected to have the opposite effect. Indeed, this is borne out by evidence presented in a paper by Alexander, Mayer and Weeds (1996) which looked at the asset betas of utilities categorised by regulatory regime (low risk cost pass through, hybrid and higher risk price cap), and found that the businesses operating under high-powered regulatory regimes had higher Betas than those operating under lower-powered regimes.

This assessment also conforms to the Competition Commission's view on the asset Beta spectrum of different industries, first outlined in its Heathrow and Gatwick regulatory report (2007) and quoted in subsequent regulatory price controls such as the Competition Commission's Bristol Water (2010) and NIE (2014) determinations. We reproduce this spectrum in Figure F.1 below.

¹²⁰ This could include the 'animal spirits' referred by Keynes in the General Theory of Employment, Interest and Money.

Figure F.1 CMA’s asset beta spectrum



Source: CC analysis in Heathrow and Gatwick regulatory report (2007).

The CMA’s analysis suggests that regulated utilities are less risky than sectors that operate in more competitive (high-powered) environments. Within the Beta range for regulated utilities of 0.30-0.45, there will be companies exposed to different strengths of regulatory regimes. Indeed, there are assets that are even less risky than the standard regulated utilities quoted by the CMA. The energy transmission network operators in Belgium (e.g., Fluxys), are regulated under a cost pass-through for their operating costs and financing cost (cost of debt), and the asset Betas of these companies are usually estimated at around 0.2.

Whilst it would be informative to separately estimate these different risk factors into additive sub-components of beta, this is not possible due to the lack of pure comparators for such a decomposition to be statistically robust. Therefore, we consider it reasonable to rely on what we observe in the market (using robust beta estimation methods) to be our best evidence of the systematic risk these businesses carry.

Estimation methodology

This section summarises the methodology and results of the beta estimations for UK regulated utilities. In essence, this is to:

- Estimate the raw equity betas;
- De-lever those betas using the Modigliani-Miller (or Hamada) formula; and
- Re-levering the betas using the same formula and using regulatory assumptions on the notional gearing level to obtain a consistent measure of the equity beta for regulatory purposes.

We discuss each step in turn.

Raw beta estimation approach

Whilst there is no “standard approach” adopted by the regulators, the methodology used to estimate the Betas set out in this paper broadly reflects the approaches taken by regulators and by the CMA in recent determinations.

We directly estimate the raw Betas using Bloomberg’s unlevered beta estimator. This template allows the user to toggle from daily, weekly and monthly estimation. Additionally, it also allows the choice of index. We do not carry out any Bayesian adjustment on the beta, such as Blume or Vasicek adjustment. We do not carry out serial correlation adjusted method such as Dimson betas¹²¹.

- Comparators: We have chosen the usual utilities from the UK as comparators, as used in the CMA’s determination for NIE in 2014. Additionally, we have included a few regulated network utilities in Europe whose majority of business activities are regulated networks with broadly similar regulatory framework as the ones in the UK.
- Index: We have chosen the local all-share indices for the estimations. In addition, we have also performed a cross check using a global index (MSCI).
- Frequency: We have carried our estimation with various frequencies, including daily, weekly and monthly.
- Estimation window: We have used 2-year, 5-year and 10-year estimation windows. We have not reported the results from a 2-year monthly estimation due to the small sample size. These windows are commonly used by the regulators and by the CMA.
- Debt Beta. We have assumed a debt Beta equal to zero in these calculations.

As noted, this approach is consistent with the broad methodology used by regulators. Figure F.2 summarises the methodology for Beta estimation in some recent decisions. It is not intended to be exhaustive but it provides a fair reflection of recent regulatory practice.

Figure F.2 Recent UK regulators approach to estimating Betas

Regulator decision	Frequency / sample period	Time frame	Sample	Debt Beta	Bayesian adjustment
Ofwat PR14	Daily (2 year), Monthly (5 year)	Daily 2 year average, Monthly 5 year average, results presented back to 2001	Penon, United Utilities, Severn Trent	Zero	Yes
CMA Bristol Water 2015	Daily (2 year), Weekly (2 year and 5 year), Monthly (5 years) ¹²²	Current values and average over past 1, 2 and 5 years	Penon, United Utilities, Severn Trent	Zero	No
CAA Airports 2014	Monthly (5 years)	12 year period	Sample of international quoted airport operators	Zero	Yes

¹²¹ The Bloomberg dataset and associated equity betas have been cross-checked for accuracy.

¹²² CMA also considered lower frequency data but decided against quarterly data because the results were not stable. Referred to academic evidence to suggest that lower frequency data would tend to underestimate Betas.

Regulator decision	Frequency / sample period	Time frame	Sample	Debt Beta	Bayesian adjustment
Ofcom LLCC 2016	Daily (2 year)	Daily 2 year average, current value	BT	0.1	No

Source: Various regulators determinations and supporting documentation

In most cases the Beta values are sourced from established data providers such as DataStream (Ofwat) or Bloomberg (CMA and Ofgem). In none of these cases was the regression method explicitly described but we can infer that an OLS method was used.

The approach we adopt is also broadly consistent with Mason, Miles and Wright (2003), albeit we also report longer term estimates of the beta than they would have recommended:

“Our recommendation is that using between one and two year periods with daily data will generally give low standard errors and that if the one year betas and two year betas are little different, the time variation problem is unlikely to be significant.”

Raw beta estimation results

The estimation results as of 25 September 2017 are as shown below.

Figure F.3 Two-year raw equity betas

	2-year daily	2-year weekly
Severn Trent PLC	0.61	0.86
United Utilities Group PLC	0.58	0.76
National Grid PLC	0.51	0.61
Pennon Group PLC	0.62	0.77
SSE PLC	0.85	0.81
Terna Rete Elettrica Nazionale	0.58	0.45
Snam SpA	0.57	0.43
Enagas SA	0.44	0.48
Red Electrica Corp SA	0.46	0.59
REN - Redes Energeticas Nacion	0.65	0.65
Average	0.59	0.64

Source: Bloomberg data

Figure F.4 Five-year raw equity betas

	5-year daily	5-year weekly	5-year monthly	5-year daily with world index
Severn Trent PLC	0.63	0.76	0.86	0.55
United Utilities Group PLC	0.62	0.68	0.78	0.50
National Grid PLC	0.61	0.66	0.68	0.48
Pennon Group PLC	0.62	0.71	0.61	0.53
SSE PLC	0.72	0.66	0.56	0.68
Terna Rete Elettrica Nazionale	0.59	0.50	0.40	0.90
Snam SpA	0.61	0.54	0.37	0.90
Enagas SA	0.53	0.52	0.49	0.77
Red Electrica Corp SA	0.58	0.64	0.63	0.81
REN - Redes Energeticas Nacion	0.50	0.49	0.70	0.55
Average	0.60	0.62	0.61	0.67

Source: Bloomberg data

Figure F.5 Ten-year raw equity betas

	10-year daily	10-year weekly	10-year monthly
Severn Trent PLC	0.63	0.76	0.86
United Utilities Group PLC	0.62	0.68	0.78
National Grid PLC	0.61	0.66	0.68
Pennon Group PLC	0.62	0.71	0.61
SSE PLC	0.72	0.66	0.56
Terna Rete Elettrica Nazionale	0.59	0.50	0.40
Snam SpA	0.61	0.54	0.37
Enagas SA	0.53	0.52	0.49
Red Electrica Corp SA	0.58	0.64	0.63
REN - Redes Energeticas Nacion	0.50	0.49	0.70
Average	0.60	0.62	0.61

Source: Bloomberg data

Unlevered asset beta estimation

In order to compare the estimated equity betas in a meaningful way, we de-lever them using the Modigliani-Miller formula (Hamada approach), which incorporates the tax treatment in de-levering the beta. The equation is as follows:

$$Asset\ beta = \frac{Equity\ beta}{1 + (1 - t) * \frac{D}{E}}$$

where D denotes net debt, E denotes the market capitalisation of equity, and t denotes the effective tax rate.

We calculate D , E and t using the average of the corresponding beta estimation window. For example, to de-lever a 2-year beta estimated as of 25 September 2017, we take the average D , E and t for the period between 26 September 2015 and 25 September 2017. We do the same for the five year and ten year betas.

We assume a debt beta of 0 for the purpose of this analysis. The resulting asset betas are shown below.

Figure F.6 Two-year asset betas

	2-year daily	2-year weekly
Severn Trent PLC	0.35	0.49
United Utilities Group PLC	0.31	0.41
National Grid PLC	0.33	0.40
Pennon Group PLC	0.39	0.48
SSE PLC	0.62	0.59
Terna Rete Elettrica Nazionale	0.36	0.28
Snam SpA	0.36	0.27
Enagas SA	0.29	0.31
Red Electrica Corp SA	0.32	0.42
REN - Redes Energeticas Nacion	0.29	0.29
Average	0.36	0.39

Source: Bloomberg data, Frontier Analysis

Figure F.7 Five-year asset betas

	5-year daily	5-year weekly	5-year monthly	5-year daily with world index
Severn Trent PLC	0.35	0.42	0.48	0.31
United Utilities Group PLC	0.33	0.36	0.42	0.27
National Grid PLC	0.39	0.42	0.44	0.31
Pennon Group PLC	0.39	0.45	0.38	0.33
SSE PLC	0.55	0.50	0.43	0.52
Terna Rete Elettrica Nazionale	0.37	0.32	0.26	0.57
Snam SpA	0.39	0.35	0.24	0.57
Enagas SA	0.34	0.34	0.32	0.50
Red Electrica Corp SA	0.38	0.43	0.42	0.54
REN - Redes Energeticas Nacion	0.21	0.21	0.30	0.23
Average	0.37	0.38	0.37	0.41

Source: Bloomberg data, Frontier Analysis

Figure F.8 Ten-year asset betas

	10-year daily	10-year weekly	10-year monthly
Severn Trent PLC	0.34	0.39	0.23
United Utilities Group PLC	0.34	0.36	0.29
National Grid PLC	0.38	0.38	0.28
Pennon Group PLC	0.35	0.42	0.30
SSE PLC	0.49	0.46	0.35
Terna Rete Elettrica Nazionale	0.31	0.29	0.26
Snam SpA	0.28	0.27	0.21
Enagas SA	0.38	0.38	0.42
Red Electrica Corp SA	0.38	0.42	0.43
Average	0.36	0.37	0.31

Source: Bloomberg data, Frontier Analysis

Note: REN is missing in this analysis as it started trading less than ten years ago.

Re-gear equity beta

Finally, we re-gear the asset betas using the Modigliani-Miller formula to derive asset betas consistent with the regulator's assumptions for gearing.

$$\text{Equity beta} = \text{Asset beta} * (1 + (1 - t) * \frac{D}{E})$$

We use a D/E ratio consistent with Ofgem's assumed 65% gearing level of the network price controls, and a 19% corporate tax rate for 2017. The resulting re-gear equity betas are shown below.

Figure F.9 Two-year equity betas

	2-year daily	2-year weekly
Severn Trent PLC	0.87	1.22
United Utilities Group PLC	0.77	1.02
National Grid PLC	0.83	0.98
Pennon Group PLC	0.96	1.20
SSE PLC	1.53	1.47
Terna Rete Elettrica Nazionale	0.91	0.70
Snam SpA	0.90	0.68
Enagas SA	0.71	0.77
Red Electrica Corp SA	0.80	1.03
REN - Redes Energeticas Nacion	0.72	0.72
Average	0.90	0.98

Source: *Bloomberg data, Frontier Analysis*

Figure F.10 Five-year equity betas

	5-year daily	5-year weekly	5-year monthly	5-year daily with world index
Severn Trent PLC	0.87	1.05	1.20	0.76
United Utilities Group PLC	0.83	0.90	1.04	0.67
National Grid PLC	0.97	1.05	1.08	0.76
Pennon Group PLC	0.97	1.11	0.96	0.83
SSE PLC	1.37	1.25	1.06	1.28
Terna Rete Elettrica Nazionale	0.92	0.79	0.64	1.41
Snam SpA	0.96	0.86	0.59	1.42

	5-year daily	5-year weekly	5-year monthly	5-year daily with world index
Enagas SA	0.85	0.84	0.80	1.25
Red Electrica Corp SA	0.95	1.06	1.05	1.34
REN - Redes Energeticas Nacion	0.53	0.53	0.74	0.58
Average	0.92	0.94	0.91	1.03

Source: Bloomberg data, Frontier Analysis

Figure F.11 Ten-year equity betas

	10-year daily	10-year weekly	10-year monthly
Severn Trent PLC	0.85	0.97	0.58
United Utilities Group PLC	0.84	0.90	0.72
National Grid PLC	0.93	0.94	0.70
Pennon Group PLC	0.86	1.03	0.74
SSE PLC	1.22	1.14	0.86
Terna Rete Elettrica Nazionale	0.77	0.71	0.66
Snam SpA	0.70	0.66	0.53
Enagas SA	0.93	0.96	1.04
Red Electrica Corp SA	0.95	1.04	1.07
Average	0.90	0.93	0.77

Source: Bloomberg data, Frontier Analysis

Note: REN is missing in this analysis as it started trading less than ten years ago.

Conclusions

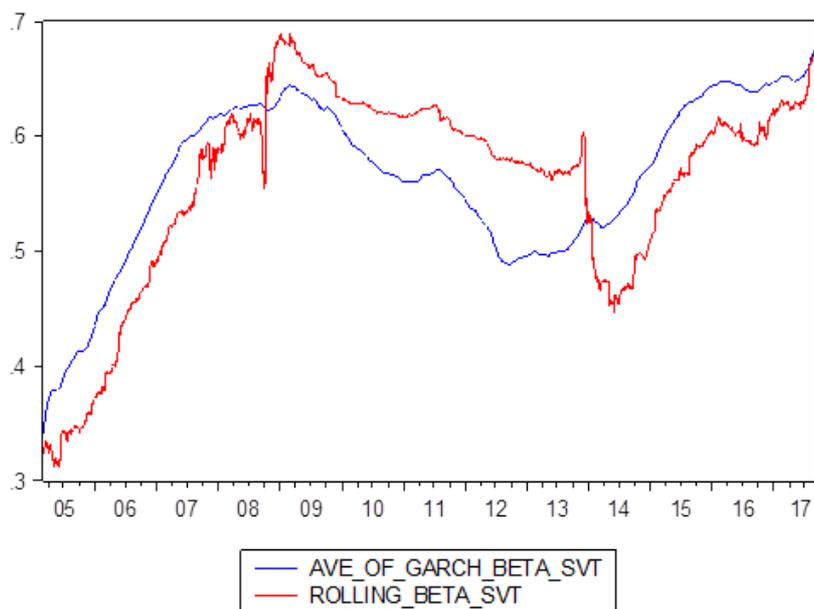
Across all the estimates from different methods, the re-gear regulated equity betas range from 0.77 to 1.03. The recent regulatory decisions in the UK have been largely within this range.

Alternative estimation approaches

In a separate appendix, Wright and Robertson report the results of some GARCH estimation of the betas. In this section we offer some brief comments on their analysis. At a general level it is always useful to explore different methods for estimating Betas and extending the understanding of best practice in this area. GARCH based approaches have been well-understood in the academic literature but have not been adopted in regulatory assessments to date. The questions to be considered are how this evidence fits alongside the existing body of evidence and established approaches.

First, the analysis does not reveal anything about the time-varying pattern of the Betas that is not already well-known. As the figure below shows, the 5 year moving average beta derived from a rolling OLS estimation is very similar to the time-varying GARCH profile of the beta, and the later Beta estimates are in line with the raw equity Betas reported in the tables above.

Figure F.12 GARCH and OLS Beta estimates



Source: Wright and Robertson

Second, as outlined in this note there are a myriad of options available for Beta estimation, of which the estimation method is just one. Therefore, the key questions for regulators (irrespective of whether the Beta has been estimated using OLS or GARCH) remain:

- What sample of comparator information should be used? Regulators have to balance the advantages of a larger a sample against the relevance of the sample (e.g. unregulated assets, or assets that are regulated by another regulator).
- Over what time period should the betas in the sample be estimated, and should outliers that relate to highly unusual periods be removed from the sample?
- Over what period should the betas be averaged?
- What frequency of data should be used?

Whatever the estimation technique, these questions would need to be addressed by the regulators. Robertson and Wright argue that the second question can be answered by estimating Beta over as long a time period as possible; and that the unconditional long run Beta would represent the appropriate average (provided that the GARCH process is stationary, which it seems to be for these two stocks). However, they do not provide a sufficient set of diagnostic tests to judge the appropriate frequency. There is an important point, though, which is that this estimation method itself is just one of many that could in principle be used, and which could yield a range of different results, as noted by Wright, Mason and Miles (2003):

“Use of an explicit technique to handle time-varying variance and co-variances is in some respects the ideal solution.....But a major drawback of the technique is that it is susceptible to over-fitting and can

reveal apparent signs of time variation where none exist, especially if complicated models of time variation are used. At a more practical level, it involves use of techniques that are highly non-linear and not widely used amongst practitioners who estimate betas. So there would be problem of getting a beta estimated with a time-varying technique to be widely accepted as a standard estimate – this is partly because there are many different ways to model time variation (GARCH, EGARCH, GARCH in mean and many newer variants).”

The complexity of the GARCH approach and its multiple variants may explain why it has not been adopted in regulatory or commercial Beta estimation to date. Robertson and Wright correctly describe their work as preliminary. For regulatory price-setting purposes, more work would be required to establish the robustness of the technique, which would include testing a wider set of variants of the GARCH model than just the one estimated, and using Monte Carlo simulation analysis to test how well each specific GARCH model actually estimates the unconditional betas for each stock under consideration (which would be more than the two estimated in the Robertson and Wright appendix).

As part of this ongoing research the appropriate frequency of estimation could be further considered. Robertson and Wright indicate that lower frequency results are more consistent with the overall framework proposed in this report. We consider that such an argument would relate to the timeframe of the estimation period rather than the frequency of returns. Our assessment has covered a range of frequencies (daily, weekly and monthly) and we note that the established regulatory practice (following Wright, Mason and Miles) has largely focussed on shorter frequencies because these have much smaller standard errors.

Interestingly, the recent research of Gregory et al. (2016) suggests that despite the advantages of higher frequency data, these estimates understate some risk factors and so are biased downwards. This suggests that the question of appropriate frequency would also be a fruitful area of future research.

Finally, if future research is to include the method of estimation, and the appropriate frequency of the dataset, in order to pursue, with renewed vigour, the search for a better econometrically determined beta, it is logical to also reconsider the functional form of the model itself. It is well-known that the CAPM model may perform less well than other multi-factor models from an econometric perspective. In their work, for example, Robertson and Wright extend the CAPM to include the risk-free rate and find that the beta estimates increase by 33%. A renewed research agenda into the statistical analysis of utility risk and returns would need to consider all these factors. That research properly undertaken would then form part of the evidence base that regulators could draw upon that would sit alongside pre-existing regulatory approaches. In other words, it would be a complement to pre-existing techniques, rather than a substitute for them.

Until that research is more comprehensive, we would argue that regulators should continue to use the CAPM on a wide range of comparator stocks, using higher frequency data (subject to testing for thin-trading and serial correlation), over different sample sizes, and interpret that body of evidence judiciously, in line with practice to date.

Conclusion on beta estimates

The findings of our assessment of Betas are as follows.

- Estimating Beta values in the CAPM framework requires a multitude of choices around sample size, frequency, timeframe, regression method and so on. Therefore, there is no single right

answer for the value of the Beta that is estimated. Regulators have exercised their judgement in weighing up the evidence before them before determining the Beta that enters into the allowed cost of equity.

- A qualitative assessment of relationship between the factors that drive market returns and how these factors affect utility return would suggest that utilities are low risk. This has been borne out by previous regulatory decisions. In addition, the extent to which utilities are low risk with low asset Beta will depend on the form of regulatory control, with low-powered cost-pass through regimes observing asset Betas in the region of 0.2.
- We have estimated Betas using a methodology that broadly corresponds to that used by UK regulators. In brief this involves daily, weekly and monthly estimates over a period of years, estimated using OLS methods often from well-established data providers. We have extended the sample size and timeframe to provide a broader picture.

Our results are, perhaps not surprisingly, consistent with recent regulatory decisions. Asset betas in the range 0.31 to 0.39 and equity Betas (at 65% gearing) in range 0.77 to 1.03.

G Beta Estimation for CAPM-WACC at Long Horizons (Stephen Wright and Donald Robertson)

Motivation

Our Recommendation 2 (Horizon) emphasised that, while the authors of the report do not have strong views on the choice of horizon regulators make in estimating the CAPM-WACC, it *is* crucial for all components to be estimated in a way that is as far as possible consistent with that horizon. Recommendation 6 stressed that this is particularly crucial in the context of beta estimation, because if we are concerned to assess the nature of systematic risk at long horizons, we should ensure that our estimation techniques are consistent with that horizon.

Burns has helpfully summarised in a separate appendix the most recent results of what is now standard practice in beta estimation: the use of relatively short (2- 5 year) samples of, usually daily data. To some extent this standard practice (available at the click of a button on Bloomberg) reflects the relatively short-term objectives of most users of estimated betas in the finance industry.

But for regulators, who deliberately pick long horizons, it appears at first sight to be distinctly counter-intuitive to use such a short samples of high frequency data to assess the systematic component of equity returns over long horizons. The conclusion of this appendix is that on closer examination there appears to be a lot of support for this initial intuition.

We should stress that we do not present the analysis below as (anything like) the final word on beta estimation. But we do argue that our results cast significant doubt on the standard approach.

First, crucially, we argue that regulators should **use** estimated betas. We argue below that the figures assumed by regulators in recent decisions lie at the upper extreme of plausible ranges derived from estimation.

But second, we argue that beta estimates should, wherever possible, make use of all available evidence, and take into account some key properties of returns: most notably the clear evidence of variations over time in volatilities and correlations. We show below that this helps to explain most of the variation in rolling betas over shorter samples.

Third, and more cautiously, we point to evidence that betas appear to fall at lower frequencies.

We focus our analysis solely on the two quoted stocks that are closest to purely regulated utilities, United Utilities (UU) and Severn Trent (SVT). We focus primarily on estimating their CAPM betas using the FT All Share as the market index. We do however comment briefly on the impact of choosing a global index, and on allowing for the impact of interest rates.

Summary of key data features

Figure G.1 shows cumulative log returns for UU and SVT compared to that of the FT all share (code ASX), on both daily and quarterly data.¹²³ Both stocks had a period of strong outperformance in the

¹²³ There is an odd temporary upward blip in SVT driven by a one-off return of 13.8% on 14th May 2013, which was almost exactly unwound by falls of 6% and 8.8% on 10th and 12th June. This disappears on quarterly data.

early part of the sample. They have continued to show some average tendency to outperformance (albeit distinctly less evident) since the pre-crisis peak.

Figure G.1
Cumulative Returns from 1/1/2000, Daily

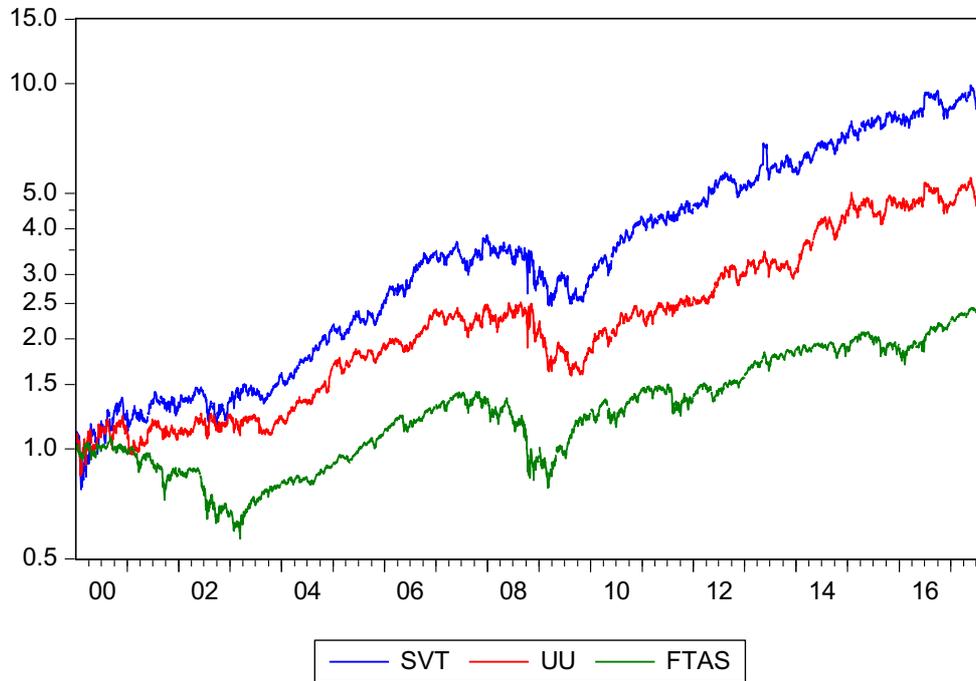
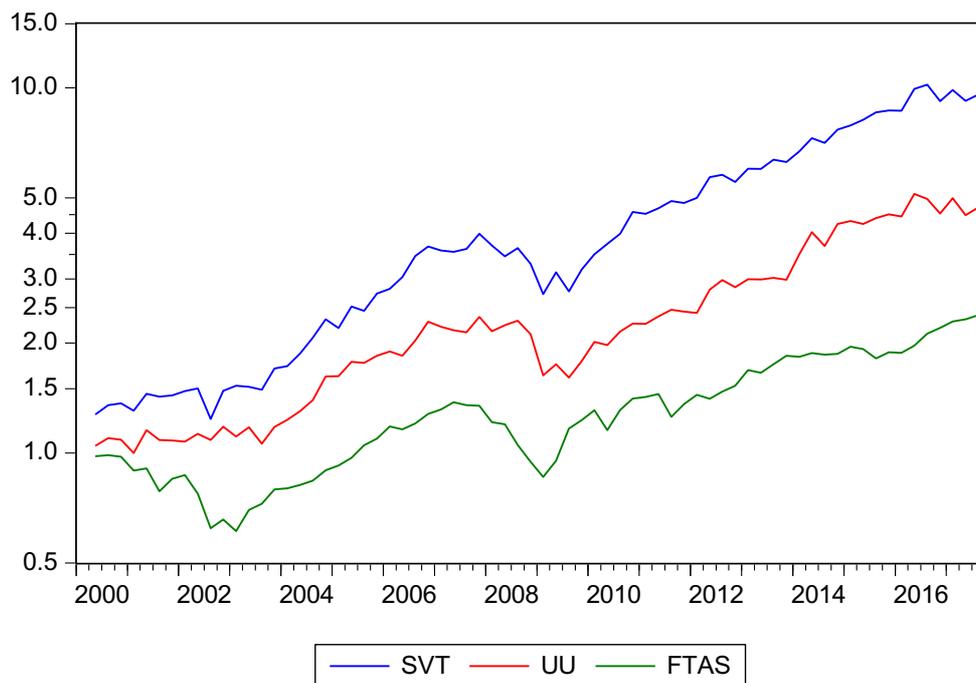


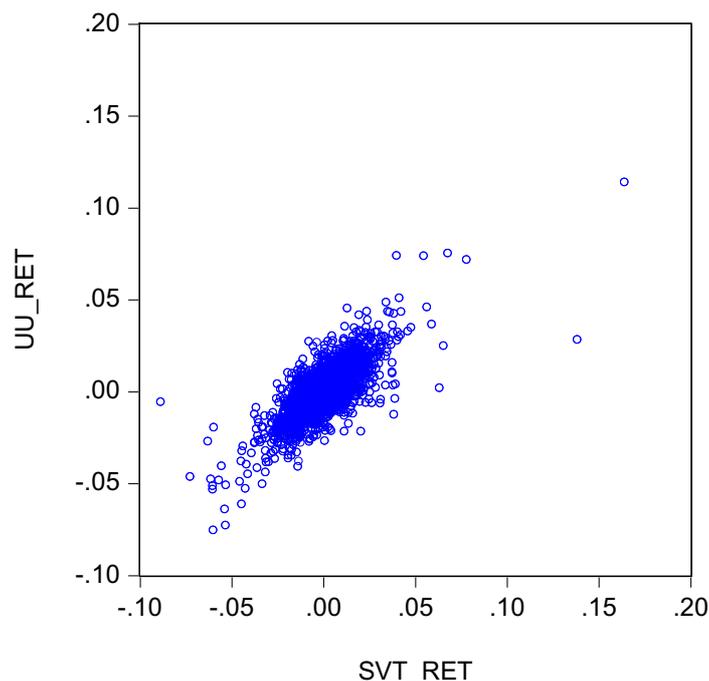
Figure G.2
Cumulative Returns from 2000 Q1, Quarterly



At all frequencies, the two stocks are very strongly correlated, especially since 2006, with a correlation between returns of 0.77 on daily data and 0.83 on quarterly data.

To a fairly good approximation, most of the time they are therefore close to behaving as a single stock. The scatter diagram of daily returns below (Figure G.3) shows that, on data since 2006, but for a few outliers (presumably representing occasional company-specific news) the line of best fit between returns would have a slope very close to unity, with most points close to the line.

Figure G.3
Scatter Diagram of Daily Returns: UU vs SVT



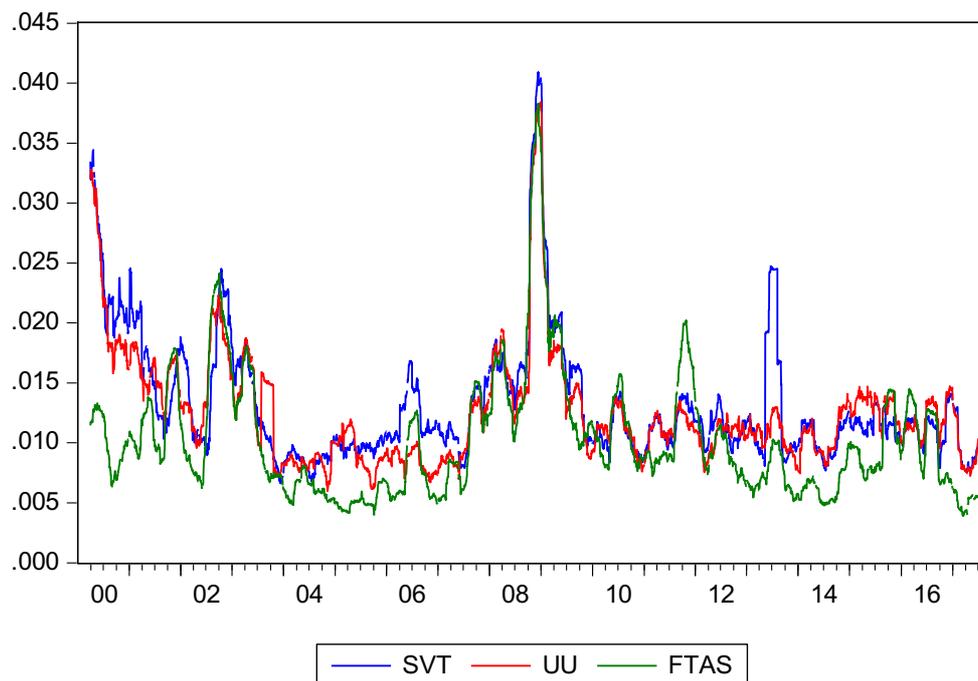
Time series properties

Neither the two stock returns nor the FT index return show any statistically significant evidence of serial correlation, at least that can be captured on standard metrics. We do however comment below on evidence suggestive of some modest (as in, not statistically significant) negative correlation at the highest frequencies, resulting in lower-than-expected volatility at lower frequencies - with possible further implications for betas at these lower frequencies.

However, there is extremely strong evidence of time variation in volatility, particularly at high frequencies. Figure G.4 shows 60 day moving averages of volatilities for the two stocks and the FT All Share. The chart makes clear that as well as being strongly correlated the two stocks mostly have very similar patterns of variation in their volatility.¹²⁴ This pattern is distinctly different from that of the market index: most of the time the two stocks have higher volatility (reflecting their idiosyncratic risk); but at points when market volatility has risen most sharply (most notably in the crisis) it has temporarily been higher than that of the individual stocks.

¹²⁴ The upward blip in 2013 for SVT is driven by the sequence of exceptional returns noted in the previous footnote.

Figure G.4
60 Day Volatilities (Moving Standard Deviations)



This time variation in volatility can be captured by a GARCH (Generalised Autoregressive Conditional Heteroscedasticity) model (see technical appendix).

Figures G.5, G.6 and G.7 show the estimated conditional volatilities of return that come out of the model show considerable persistence in daily data, which quite closely matches the pattern of the simple 60-day volatility.¹²⁵

¹²⁵ This is to be expected since the GARCH model constructs the conditional variance as a geometric weighted average of actual squared returns, with weight given by the estimated GARCH coefficients (which are all highly significant). Mean lags for the GARCH processes for the three series are all around one month.

Figure G.5
SVT: 60 Day Volatility vs GARCH Estimate

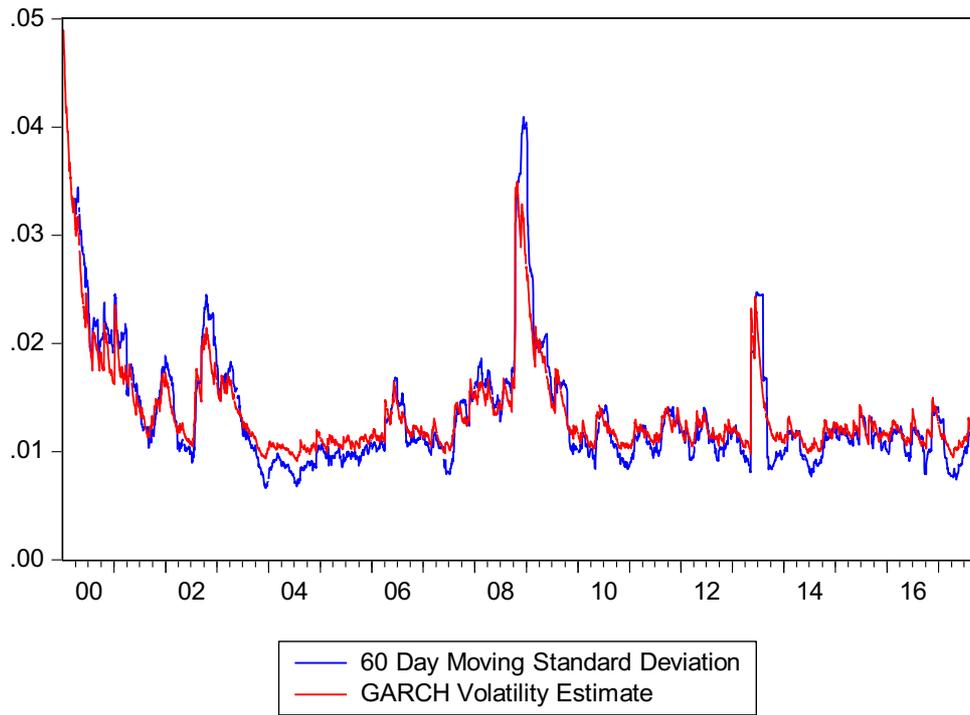


Figure G.6
UU: 60 Day Volatility vs GARCH Estimate

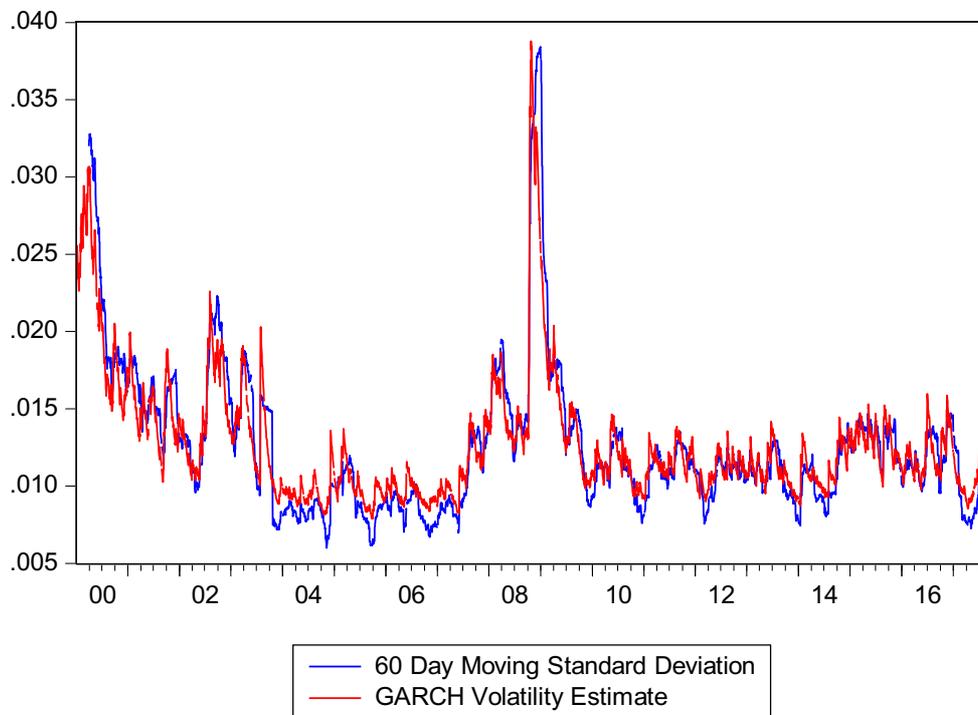
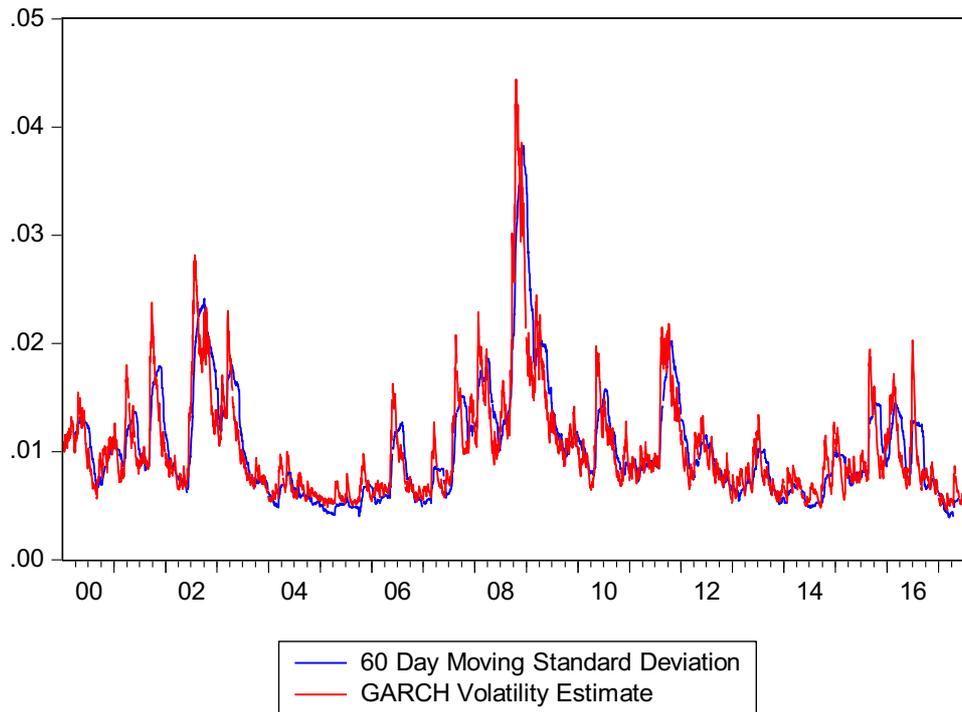


Figure G.7
FTAS: 60 Day Volatility vs GARCH Estimate



Time-varying betas from multivariate GARCH vs rolling betas

In a multivariate GARCH framework both the variances and covariances vary over time, implying

that the conditional beta, $\beta_{SR,t} = \frac{\text{cov}_t(R_{it}, R_{mt})}{\text{var}_t(R_{mt})}$ also varies. There is a well-established academic

literature (see technical appendix) that uses GARCH and related techniques to estimate these conditional (or short-run) betas.

Figure G.8 and Figure G.9 show that if we first calculate the implied conditional betas from the GARCH estimates, and then take five-year moving averages, the resulting patterns quite closely resemble those of the rolling 5-year beta estimates for both stocks.¹²⁶

¹²⁶ Note that the sharp spikes in the rolling beta estimates in 2008 and in 2013/14 also match the pattern of short-run variation in conditional betas (e.g., the spike down in 2013 for UU corresponds to a sharp fall in the conditional beta in the new observations progressively added at the end of the sample, *relative* to its value in the earlier observations progressively omitted).

Figure G.8
Rolling Beta Estimates vs 5-year average of GARCH Short-Run Beta: UU

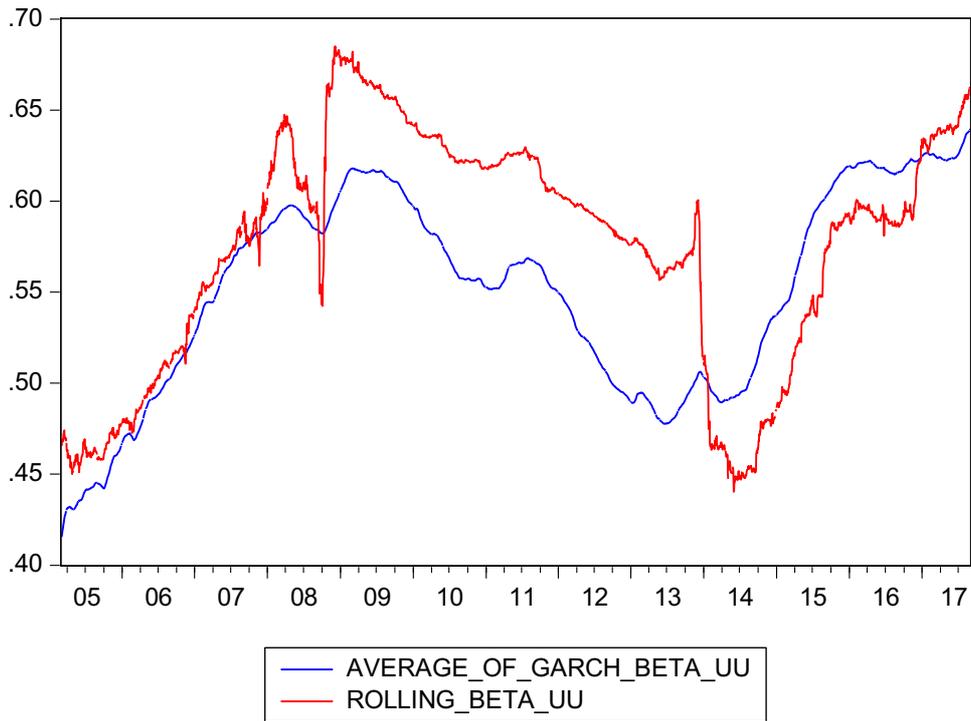
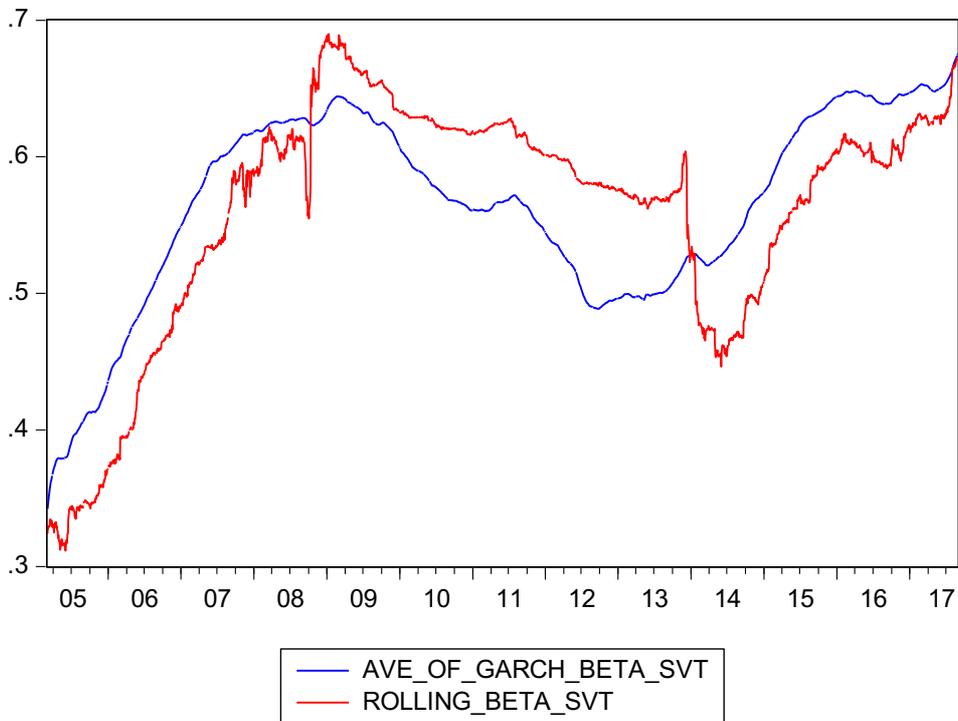


Figure G.9
Rolling Beta Estimates vs 5-year average of GARCH Short-Run Beta: SVT



Thus the pattern we observe in rolling betas in high frequency data appears to be largely, or quite possibly entirely, driven by the GARCH properties at these higher frequencies. However a crucial difference between the two estimation approaches is that the GARCH estimation explains the time variation in rolling beta by estimation over the *full* sample. Furthermore, the estimated GARCH processes are stationary for all three series, implying that the variances and covariances, and hence conditional betas, converge to stable long-run values.

Thus the GARCH approach suggests that rather than taking a snapshot of the rolling beta in the most recent rolling sample (whether 2 or 5 year), we should examine the properties implied by the full sample. This contrast is particularly important, since, as the chart reveals, in recent samples rolling beta estimates have been on a rising path, and latest observations are at a near-high compared to earlier samples. Current practice effectively treats this shift as permanent, whereas the explanatory power of GARCH suggests strongly that it is likely to be only transitory.

Implications for lower frequencies and “long-run” beta.

An implication of the above analysis is that, if daily data over the whole sample suggest a stable GARCH process, but with short-term variations that explain the pattern of rolling betas, this pattern may well be less evident in lower frequency estimation. Furthermore, since it suggests no obvious evidence of instability over the full sample, the usual obstacle to estimation at lower frequencies – the small number of observations in relatively short rolling samples – largely disappears. Even on quarterly data, we have a fairly respectable 70 observations.

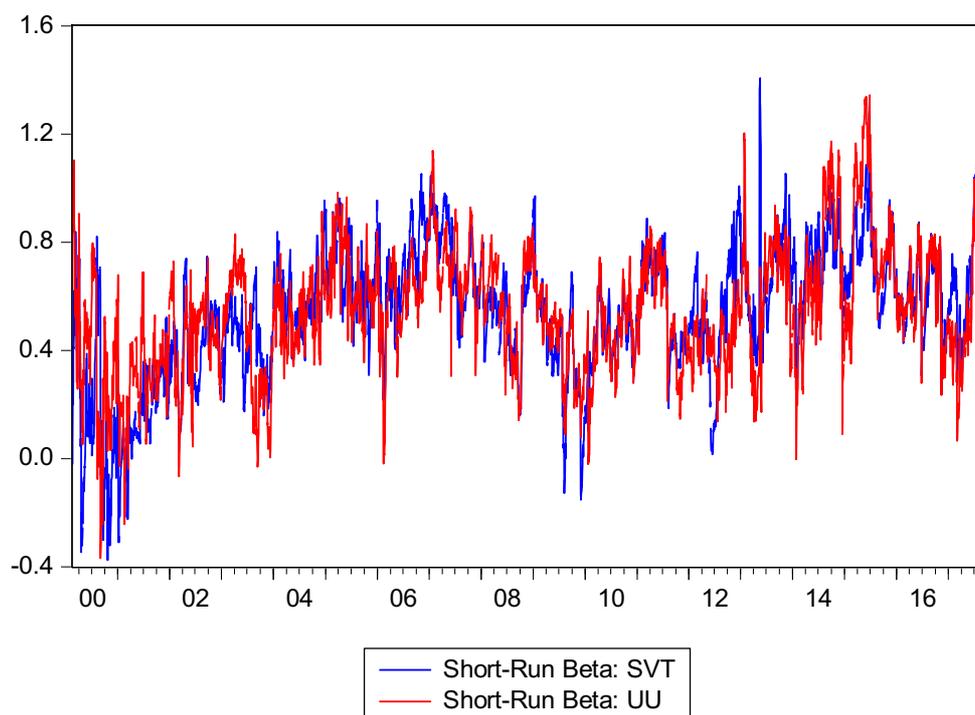
The clear pattern that emerges from estimation at lower frequencies is that evidence of GARCH is distinctly weaker for the two individual stocks, with little sign of persistent shifts in volatility,¹²⁷ and much more limited for the market index (which simply shows a spike around the global financial crisis). As a result, as we show below, broadly speaking results at lower frequencies from the full sample are fairly similar, whether using OLS or GARCH.

Which beta?

A further issue that needs to be considered when looking at beta estimates derived from GARCH is: which beta should we pick? This is particularly important given the long horizons preferred by regulators. The underlying estimated conditional betas that, when averaged as above, match up to rolling betas, are very volatile, as Figure G.10 shows. (Note the considerable similarity of the two estimates). Nor would we in any case wish to pick a value at any particularly point in time, since the conditional beta captures only quite transient shifts in the relationships between daily returns.

¹²⁷ This is consistent with the good match between 60 day volatility and daily GARCH estimates, implying that most fluctuations in volatility disappear within two months.

Figure G.10
Short-Run Beta Estimates from GARCH



But GARCH estimates also have implications for longer horizons, since a stationary GARCH process implies that in expectation all elements of the covariance matrix converge to their unconditional values. Thus, as well as a conditional beta, we would ideally like to estimate the unconditional (or “long-run”) beta, which is the ratio of the *unconditional* covariance to the *unconditional* variance of the market return. It is long-run beta that will determine the impact of systematic risk over the horizons relevant to regulators.

Note that, crucially, (for reasons set out in the technical appendix) the long-run beta is *not* equal to the average (or expected value) of the conditional beta.¹²⁸

The technical appendix discusses two alternative methods of estimating long-run beta. One, which we denote β_{LR} below, is calculated using the underlying estimated parameters of the GARCH process (and hence may not be unbiased in small samples); but a second, simpler approach, which is unbiased, and which we denote β_{av} below, is simply calculated as the ratio of the average conditional covariance to the average variance.¹²⁹ We show below that (as we would expect) at high frequencies these alternative estimates of long-run beta are fairly consistently lower than the average of the short-run beta, since the impact of GARCH on the short-run beta is strong; whereas at lower frequencies the results are more consistent.

Summary properties of beta estimates.

Table 1 compares beta estimates from GARCH and OLS estimation at 4 different frequencies, daily, weekly, monthly and quarterly.

¹²⁸ See also the graphical evidence presented below.

¹²⁹ NB: the ratio of averages is not equal to the average of the ratios.

Table 1
UU and SVT betas

Three methods considered at various frequencies

SVT	β_{LR}	$\beta_{av} = \frac{\overline{Cov}}{\overline{Var}}$	$\beta_{SR} = \left(\frac{\overline{Cov}}{\overline{Var}}\right)$	β_{OLS}	$\overline{\beta_{roll}}$	$\beta_{roll,final}$
Daily	.44	.48	.53	0.53	0.56	0.67
Weekly	.36	.43	.51	0.46	0.50	0.74
Monthly	.40	.40	.54	0.36	0.41	1.0
Quarterly	.30	.30	.43	0.29	0.23	0.20

UU	β_{LR}	$\beta_{av} = \frac{\overline{Cov}}{\overline{Var}}$	$\beta_{SR} = \left(\frac{\overline{Cov}}{\overline{Var}}\right)$	β_{OLS}	$\overline{\beta_{roll}}$	$\beta_{roll,final}$
Daily	.43	.53	.54	.57	.57	0.66
Weekly	.37	.46	.53	0.50	0.52	0.75
Monthly	.37	.34	.49	0.36	0.41	0.90
Quarterly	.32	.30	.35	0.30	0.25	-0.00

There is a distinct contrast between whole sample results and results in the last rolling five-year period, shown in the last column. In the rest of the table, full-sample results display a fairly consistent pattern: beta tends to fall both moving down any given column (as frequency falls) and moving leftwards in any given row (but with the latter fall being much more marked for daily and weekly data).

At any given frequency, short-run beta estimates are higher than long-run estimates, which can be as low as 0.3 (at lower frequencies). OLS estimates tend to lie midway between short-run and long-run estimates, but, especially at quarterly frequency, all full-sample estimates are fairly similar (which we would expect, since GARCH effects largely disappear within any given quarter). Furthermore, we can have a reasonable degree of confidence in the precision of these lower frequency estimates, since if we use all (or at least a large part of the available dataset) we have a reasonably large sample even at lower frequencies: e.g., even at a quarterly frequency we have 70 observations.¹³⁰

Thus we get a reasonably consistent picture across the frequencies: pretty much all full-sample beta estimates lie in the relatively narrow range 0.3 to 0.5. Note that we would *not* expect complete consistency across frequencies except under quite restrictive assumptions. We discuss this issue further below.

In contrast the last five-year sample shown in the final column (which captures current standard practice in summary form) gives markedly contrasting results across the frequencies – most strikingly the contrast between the exceptionally high monthly estimates of the five year rolling beta and the quarterly estimates, which are both close to (or exactly) zero. Here the results are however driven by the small numbers of observations, so that individual observations can have a highly significant impact on the estimates. Thus, for example the rolling monthly beta estimate jumps from 0.66 to 0.92 as a result of adding a single observation (June 2017), and subtracting the June 2012 observation.¹³¹ While the latest higher frequency rolling estimates were also unusually high (compared to full sample values), the

¹³⁰ It would almost certainly be possible in principle to increase precision of lower frequency estimates by using overlapping higher frequency returns. Eg, we can use daily data to estimate beta on 90 day, or even year-on-year returns. Point estimates are unbiased; and while standard errors are contaminated by overlapping observations, the gain in precision almost certainly dominates. Preliminary investigations suggest very similar results to those on pure non-overlapping quarterly data.

¹³¹ Both UU and SVT fell sharply (by 13% and 10% respectively), at the same time as the FT All Share fell, but only by a very modest amount, thus boosting the rolling beta estimate.

analysis above suggests that this is consistent with the evidence from GARCH estimation. Furthermore, as the chart above shows, the “spot” estimates of the conditional betas on daily data had already fallen back to close to their long-run average values by the end of the sample.

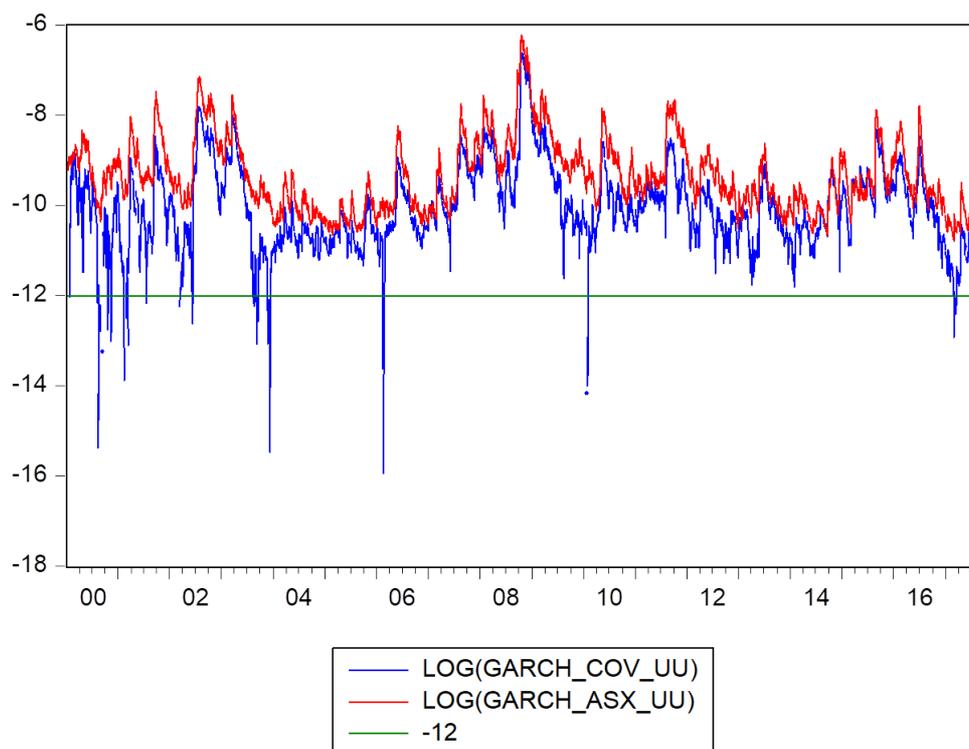
It should also be noted that these last (clearly, exceptionally high) estimates of beta, derived from the last rolling monthly sample, are the *only* values in the table lending any support to the values of 0.8 and 0.9 assumed by Ofwat and Ofgem, respectively, in recent deliberations.

A graphical interpretation of the results.

Figure G.11 plots, for United Utilities, the time path of the estimated covariance, and the variance of the market return, on daily data.

A first key feature that the chart reveals is that both appear to be stationary, so that the concept of the long-run-beta as the ratio of their long-run means does indeed appear sensible.¹³²

Figure G.11
UU: GARCH Estimates of Instantaneous Variance and Covariance with FTAS



A second, quite striking, feature is that the covariance is almost invariably positive. There are only 52 negative values in a total of 4462 observations.¹³³ So the result that beta is positive is very robust indeed: there is strong evidence for at least some systematic risk.

¹³² Note that this does not appear to be an artefact of GARCH estimation: for example the simple 60 day rolling covariance and variance show a very similar pattern.

¹³³ These are censored out automatically in the chart above, since this is in logs. However their absence is barely visible. In contrast the occasional sharp dips in the log(covariance) shown in the chart occur when it gets very close to zero. However, while these appear quite dramatic in the chart, the combined impact of zero, or very small covariances on beta estimates is very small: censoring the regression by excluding all observations in which the

A third clear feature is that the covariance and the variance have tended to rise together, but peaks in the covariance have been sharper, thus raising estimates of short-run beta at times of relatively high volatility. Even when averaged out these higher estimates tend to dominate the average of the short-run beta.

Considerations for sample selection

The same chart also sheds some light on sample selection. As noted above, the clearest feature is long-run stability, suggesting a strong *prima facie* case for using all available data. However, the chart does suggest at least visual evidence that, while the volatility of the market was fairly stable, the covariance fluctuated around a somewhat lower value during the early part of the sample, suggesting a lower beta. Some prior knowledge of the history of UK regulation might also suggest that the early part of the sample was somewhat atypical. The period up until around 2005 was, as the first two charts showed, a period in which both UU and SVT appeared to systematically outperform the index. It might possibly be argued that this was a period in which companies and markets were still learning about the nature of UK regulation, and that these two stocks were more strongly driven by this than by their relationship with the market.

We should however be careful about arguments of this nature, which can be perilously close to data mining. While visual inspection might seem to suggest that both stocks had “alpha” during this period the estimated constant term is not actually statistically significantly different from zero in this sub-sample at any frequency. And it should be borne in mind that both companies had been engaged in regulated water activities for a decade before the beginning of our sample. Thus, before censoring the sample, it would arguably make more sense initially to attempt to extend the dataset further back, to see if evidence of an apparent shift becomes stronger, or possibly weaker, on a longer evidence base. It should also be subjected to standard tests of the null hypothesis of structural instability.

Why does beta appear lower at lower frequencies?

While all the full-sample beta estimates lie within a relatively narrow range of 0.3-0.5, this still leaves a puzzle: why the difference between high and low frequency betas? But over the full sample of data the pattern is pretty clear: the lower the frequency the lower the beta. Given the long horizons chosen by regulators, this difference is potentially quite important.

In the simple benchmark case in which all returns were serially independent we would not expect to see any pattern in beta across different frequencies. However, while, as noted above, standard tests on high frequency data do not reject the null hypothesis of serial independence, there *is* some evidence of negative serial correlation: i.e., that periods of exceptionally high or low returns tend to be followed by some reversal in succeeding days.¹³⁴

A simple way to illustrate this feature of the data is to compare the variance of returns at different frequencies. In the serially independent case, if there are n_F days in a period of frequency F (with $F \in \{D, W, M, Q\}$) then we should have $\text{var}_F(R_t) = n_F \text{var}_D(R_t)$. E.g., for monthly data $n_M \approx 21$, so the variance of monthly returns should be 21 times higher than the variance of daily returns, implying

covariance is either negative or less than -0.12 in logs (as illustrated in the chart) the OLS beta is unchanged to 2 dp.

¹³⁴ The importance of this point had been noted in an internal memorandum by Ian Rowson, then of Ofgem.

that the standard deviation should be $\sqrt{21} \approx 4.6$ times larger. In fact the ratio of variances of monthly vs daily returns is only around 12 for both individual stocks and the market.

The “variance ratio” (Cochrane 1988) is a summary property that compares the ratio of the variance of returns at any given frequency to the implied variance if returns were serially independent at the highest frequency.¹³⁵ If the variance ratio is less than 1 at lower frequencies this indicates negative serial correlation. For all 3 series, the variance ratio (VR) declines at lower frequency, nontrivially. For example, at a monthly horizon the VR (based on daily returns) is around 0.6 for both UU and ASX, at a quarterly horizon it is around 0.5 for UU but remains at 0.6 for ASX. So, along with a clear indication of a declining VR there is also a slight (though probably not statistically significant) differential pattern in this decline.

A declining variance ratio need not of itself imply that beta must change at lower frequencies, but it *does* imply that we cannot rule it out. Further investigation of this issue would be required to establish whether there is any clear evidence of a systematic pattern that is also statistically significant.¹³⁶

Using Global Market Indices

We have replicated the techniques summarised above, replacing the FT All Share with returns on global indices. Results are broadly similar, but, as also shown in Burns’ analysis, betas on global market returns tend to be somewhat lower. The fact that there is only a minor impact may reflect the extent to which the FT All Share is itself a globally diversified portfolio. The Bank of England estimates that the exposure of the FTAS to the UK in the most recent data has been less than 30%: see Dion and Rattan (2017).

Multifactor estimation? The role of yield changes.

As discussed in the main report, multifactor models can provide a helpful cross-check on standard techniques for estimating CAPM betas. In particular, various papers by Fama and French have noted that in principle bond market factors may influence equity market returns, and vice versa. If there are correlations between these additional factors and the market return, then in principle there may be missing variables bias in a standard CAPM regression. However, if such factors have no statistically significant risk premia associated with them, they disappear in expectation, leaving only the impact of the expected market return.

We also have a particular reason to be interested in the role of bond yields. Since estimated CAPM betas for our two stocks are low, this implies that the expected return on both stocks should be strongly affected by the risk-free rate. When this changes, the discount rate applied to future cashflows should change too. Furthermore, as also noted in the main body of the report, changes in risk-free rates are readily observable.

¹³⁵ Since the variance ratio compares unconditional variances it should in principle not be affected by GARCH effects, although this does make small sample inference harder.

¹³⁶ Thus the memorandum by Ian Rowson, noted above, finds some evidence that individual stock returns may respond with a lag to the market return, with negative coefficients on lagged returns, thus bringing down the long-run impact. We have investigated this pattern using simulation techniques, which suggest that it could, with non-trivial probability have arisen by chance, if returns were in fact serially independent. However the pattern he finds is suggestive, and is broadly consistent with the pattern we find in both the variance ratio and in beta at different frequencies.

So we would have a strong *a priori* reason to expect a role for some measure of risk-free rates. Preliminary investigations suggest that we do indeed find a statistically significant impact on returns for both stocks, at all, but particularly at lower, frequencies, along with some shift in estimated CAPM betas.

As illustrative result, if we add the change in the 10-year gilt yield as an additional regressor in our quarterly regression for UU on the FT All Share return, the CAPM beta actually rises somewhat, from 0.3 to 0.4. But the regression coefficient on the change in the yield itself is also interesting. The point estimate for UU is around -5 (and strongly significant): implying that a 1 percentage point fall in the 10-year gilt yield in any quarter would boost the price (and hence the return in that quarter) by 5 percentage points. This is, to a good approximation, equivalent to the impact of a one point change in yield on a 5 year zero coupon bond. Thus these preliminary results are suggestive that these stocks *do* have some “bond-like” features; although the modest associated increase in the CAPM beta is a crucial reminder that systematic market risk also plays a very important role.

Conclusions

We certainly do not wish to claim that the above results provide the final word on beta estimation. But they do both encapsulate econometric best practice, by allowing for the clear evidence of time variation in volatility, and help to explain the observed pattern of rolling betas. As such, since, on the most recent samples the latter end up at exceptionally high values, given the evidence of stable long run variances and covariances, this appears to give quite strong grounds to view recent rolling beta estimates as implausibly high. Indeed our results cast strong doubt on the whole approach of using rolling OLS beta estimation (which has indeed always been known to be difficult to defend econometrically).

Our results would suggest that, for the two stocks we examine, a value for equity beta at long horizons in the range 0.3 to 0.5 would be econometrically defensible. Since over the sample as a whole both stocks had reasonably stable leverage of around one half (expressed in terms of the market value of equities and the reported book value of debt), this suggests a value of asset beta roughly in the range 0.15 to 0.25. These are clearly significantly lower numbers than those currently assumed by regulators.

We certainly do *not* advocate that regulators should immediately switch to particular numerical values for equity and asset beta. Our results point to a number of avenues of further investigation:

- It would also clearly be of interest to apply the same techniques over a wider sample of stocks of regulated and unregulated companies, both in the UK and overseas (as carried out, using the current standard approach, by Burns);
- Equally crucially, the dataset of returns should be extended as far back as possible, to provide as long a sample of data as possible.
- The contrast between results for beta at different frequencies could be investigated further by examining overlapping low-frequency returns on high frequency data. This technique should result in unbiased parameter estimates, but the standard errors from such techniques should be corrected for the induced moving average error process.
- The impact of global indices (both unhedged and hedged) needs to be investigated more thoroughly, taking into account any possible interaction the role of exchange rate and interest rate changes;
- While, as noted above, the *ex ante* case for using all available data is quite strong, there may also be a case for examining the impact of truncating some of the early years of the sample; but

the case for this would need to be robust, and based on clear statistical evidence of structural shifts.

- The potential role of negative serial correlation and lagged impact of market returns would also repay further investigation.
- The regressions discussed above have followed standard practice in simply using total returns, despite the fact that the CAPM relation is clearly specified in terms of *excess* returns over a risk-free asset over the relevant time interval. At high frequencies our preliminary investigations suggest that this omission appears to have a trivial impact on results; but this is less obviously so at longer horizons. An additional advantage of specifying the regression in terms of excess returns would be that estimation could then be carried out under the “no alpha” hypothesis implied by the CAPM (i.e., excluding a constant term in the CAPM regression) and tested against an unrestricted alternative.
- Finally, and crucially, the implications of these lower values for beta, and hence the distinctly lower implied cost of equity, need to be sense-checked in practical terms: e.g., against the implied premia for these companies over their regulated capital values. In the current juncture, values of beta less than one half would imply very low implied costs of equity.

While we stress that our results reported here are preliminary, the distinctly lower range for beta that we find *does* appear more consistent with regulators’ prior views, as expressed in the terms of reference, that the underlying risk of these companies is quite limited. Here, however, we note the distinct contrast between the qualitative views expressed by regulators, which would suggest that beta is low, and the actual numerical values for equity beta they have been assuming, which have been close to unity. We would also note that these high assumed values of equity beta have been chosen *despite* an abundance of evidence, long preceding this report, that the range of estimated values was typically much lower.¹³⁷

Our strongest recommendation, therefore, is that regulators should base their evidence on beta on sound econometric principles. All of the estimation results described in this appendix (including the extensions suggested above) use well-established techniques that are readily available in standard econometric software packages. These techniques would be covered in any MSc or even advanced undergraduate econometrics course. Thus everything we suggest above clearly satisfies our requirement of implementability. If carried out reasonably carefully, without cherry-picking results from outside plausible ranges of estimates, it would also be clearly defensible. Regulators’ conventional reliance on unadjusted simple beta estimates derived from the most recent short daily sample is, we argue, prone to overstate the underlying beta risk and, in light of our analysis, very much harder to defend.

¹³⁷ See, for example, the reports by Baskaya et al (2006); RIIO Financeability Study, Imrecon/ECA, November 2012; Smithers and Wright (2013).

H Adjusting bond yields for default risk (Ken Hori, Birkbeck)

Introduction

Afik and Benninga (2012) make the following observations:

“The yield to maturity (YTM) of a defaultable bond is a promised return based on promised future cash flows, if the bond is held to maturity and its issuer doesn’t default.”

“Asset pricing theory typically focuses on expected returns. Good examples are the single-factor capital asset pricing model (CAPM) and the multi-factor arbitrage pricing model (APT); in these models the expected return is derived from the appropriate risk factor loadings.”

“An example of the inappropriateness of the YTM is the computation of the weighted average cost of capital $\left(WACC = \frac{E}{V}r_E + \frac{D}{V}r_D(1 - T)\right)$. Whereas the cost of equity r_E is typically computed from the security market line and hence represents the expected returns to equity holders, the cost of debt r_D is usually computed as the YTM of the firm’s debt. These two measures are incompatible. A more consistent measure of the cost of debt is the debt’s expected return.”

All of these statements are correct. Here we investigate the effect of applying the EBR instead of the YTM in the cost of debt estimation.

Method

The estimation of EBR follows Afik and Benninga (2014). YTM is the IRR of the promised bond cash flows CF_t when you pay price p for the bond,

$$p = \sum_{t=1}^T \frac{CF_t}{(1 + YTM)^t}$$

The expected bond return EBR is the IRR of the expected bond cash flows that reflects the default probability,

$$p = \sum_{t=1}^T \frac{E[CF_t]}{(1 + EBR)^t}$$

Afik and Benninga call the difference, $YTM - EBR$, the credit risk premium, CRP. This is because, *ceteris paribus*, by law of no arbitrage bonds should have the same EBR. Then higher the default probability, the higher the YTM; in particular if the bond has zero default probability, then $YTM = EBR$.

The expected cash flows $E[CF_t]$ are computed using the transition matrix provided by rating agencies. This gives the probabilities of moving from one credit rating to another over a year. Multiplying the

matrix N times gives the N -year transition matrix. Specifically it gives the probability $\pi_{R,t}$ of defaulting at time $1 \leq t \leq T$, where T is the bond's maturity, if the current rating is R . Then the expected cash flow at time t is,

$$E[CF_t] = (1 - \pi_{R,t})CF_t + \pi_{R,t}\delta,$$

where δ is the principal's recovery rate upon default and

$$CF_t = \begin{cases} C, & 1 \leq t < T \\ 1 + C, & t = T \end{cases},$$

where C is the coupon rate.

For example, S&P Global Ratings (2017) publish the following one-year corporate transition rates (1981-2016) for Europe (Table 22):

	AAA	AA	A	BBB	BB	B	CCC/C	D	NR
AAA	82.64%	11.16%	0.62%	0.21%	0.00%	0.00%	0.21%	0.00%	5.17%
AA	0.29%	84.99%	10.43%	0.58%	0.00%	0.00%	0.00%	0.00%	3.72%
A	0.01%	1.92%	86.69%	6.08%	0.19%	0.01%	0.00%	0.04%	5.06%
BBB	0.00%	0.10%	4.26%	83.71%	3.89%	0.37%	0.10%	0.08%	7.50%
BB	0.00%	0.00%	0.09%	5.22%	72.78%	7.64%	0.41%	0.41%	13.45%
B	0.00%	0.00%	0.04%	0.37%	6.65%	71.22%	4.37%	2.53%	14.82%
CCC/C	0.00%	0.00%	0.00%	0.00%	0.00%	13.62%	41.28%	26.38%	18.72%

Here D is the state of default and NR is the case where no rating has been requested or there is insufficient information on which to base a rating. This 7×9 matrix is converted into a 9×9 modified transition matrix Π by (i) renormalising such that the probabilities from AAA to D add up to one, (ii) adding a new row for default state D, and (iii) adding a new row and a new column for post-default state. Both new rows have probability zero for all columns except the last where the probability is one (e.g., the probability of going from default to post-default state is one). Note the adjustment (i) assumes that for companies concerned, the bonds will always remain rated.

Given this modified transition matrix, the expected cash flows are calculated by the following matrix multiplication,

$$E[CF_t] = S \Pi^t CF_t$$

where S is the initial state 1×9 row vector where all entries are zero except for the column for the initial rating, for which the entry is 1, and CF_t is the 9×1 column vector of cash flows in each of the state {AAA, AA, ..., CCC/C, D, post-D} for $1 \leq t \leq T$.

Once the expected cash flows are computed for all t , the IRR is then computed as the EBR.

Results

First the above method was applied to Southern and Scottish Energy (SSE) 5.875% 2022 bond. The bond specifics are:

Issue date	22 March 1999
Maturity date	22 September 2022
Coupon rate	5.875%
Coupon frequency	Annual
Rating	A

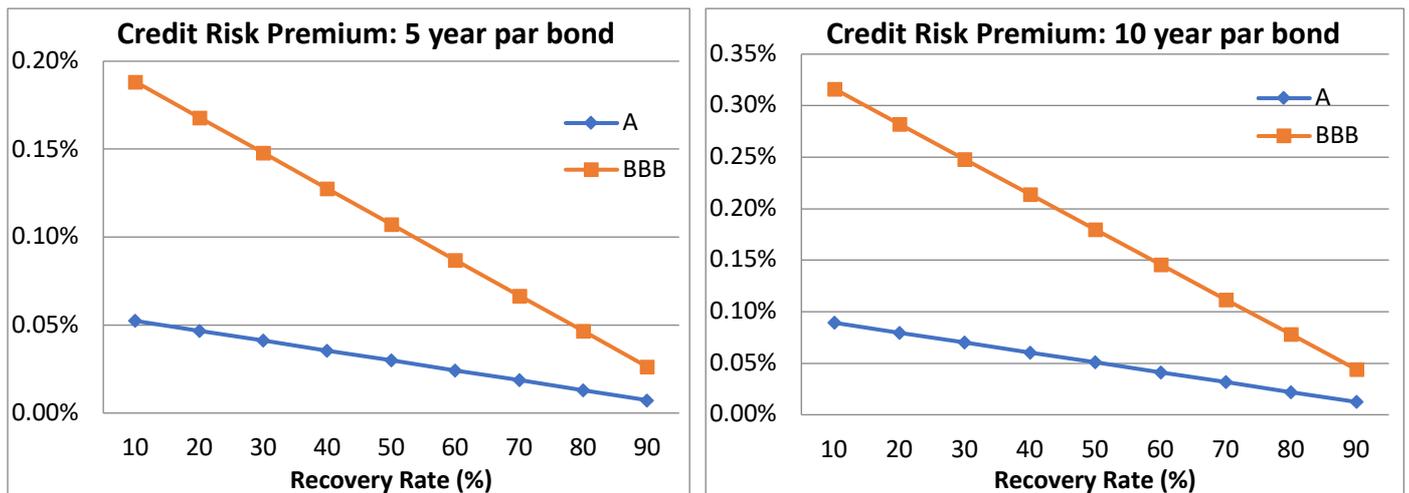
The price of this bond quoted on 26 September 2017 was 120.826, equivalent to the YTM of 1.506%. With the recovery rate of 48.8 (taken from Exhibit 8 of Moody's (2016)), the EBR is calculated to be 1.446%, i.e., the CRP of 6bp. Using Moody's (2016) Average One-Year Letter Rating Migration Rates, 1970 - 2015 (Exhibit 26) instead produces a very similar result.

Next, to investigate the EBR across a range of maturities, hypothetical par bonds with coupon rate (and hence the YTM) of 3% were considered for all credit ratings. The simulated CRPs are:

T	1	2	3	4	5	6	7	8	9	10
AAA	0.000%	0.019%	0.031%	0.038%	0.042%	0.045%	0.047%	0.048%	0.048%	0.049%
AA	0.000%	0.001%	0.003%	0.004%	0.006%	0.008%	0.010%	0.012%	0.014%	0.016%
A	0.023%	0.024%	0.025%	0.028%	0.031%	0.034%	0.038%	0.042%	0.047%	0.052%
BBB	0.047%	0.063%	0.079%	0.094%	0.110%	0.125%	0.141%	0.156%	0.170%	0.184%
BB	0.257%	0.350%	0.441%	0.522%	0.594%	0.655%	0.707%	0.751%	0.787%	0.817%
B	1.610%	1.967%	2.165%	2.271%	2.322%	2.341%	2.340%	2.327%	2.307%	2.283%
CCC/C	17.591%	15.860%	14.305%	12.970%	11.852%	10.926%	10.159%	9.519%	8.980%	8.520%

The results are quite insensitive to the coupon rate chosen. For example, for $c = 5\%$, the CRP for 10-year BBB bond is 18.6bp, while for B-rated bond it is 2.366%.

Finally to investigate the sensitivity of the CRP to the recovery rate, δ was varied from 10% to 90% for 5-year and 10-year hypothetical par bonds with coupon rate 3%. The results are shown in the graphs below:



We note that these numbers are distinctly smaller than those estimated by Afik and Benninga (2014), though of similar magnitudes. The difference in the estimated CRP may be due to the difference in the treatment of NR-ratings in the transition matrix.

Conclusions

In the estimation of the WACC of companies, the bonds' expected bond returns (EBRs) should be used rather than their YTM's. For an A-rated 10-year bond this would lower the cost of debt estimation by less than 10bp (and less if the recovery rate is high), while for a BBB 5 year bond the cost of debt would be less than 20bp lower.

I Aiming up: analysis of asymmetry in setting the RAR

(Robin Mason)

A model of minimising losses in setting the RAR for new investment

This section develops a simple model of losses from choosing a RAR, and hence RER, which is different from the true, but unobserved WACC, for new investment i.e., that has not yet taken place. (We consider sunk investment later.)

To simplify the analysis in what follows we assume away the existence of the “informational wedge” W_I defined in Section 8.1, and thus set $RAR=RER$. In practical implementation, however, clearly the problem solved below is equivalent to choosing the target value of the *sum* of the regulatory wedge W_R and the informational wedge W_I .

Suppose that the true (unobserved) WACC, y , is a random variable with probability density function (pdf) $f(y)$, cumulative distribution function (cdf) $F(y)$ and mean μ . Denote the RAR chosen by the regulator by x . The loss incurred when choosing x is:

$$L(x, y) = W(y) - W(x), \quad x \geq y, \\ = W(y), \quad x < y,$$

where $W(z)$ is the welfare that results when the RAR is set at the level z . In words: there is no loss resulting from choosing the RAR equal to the true (but unobserved) WACC: $x = y$. The loss from over-estimating the true WACC (i.e., choosing $x > y$) is the difference between the resulting welfare, $W(x)$, and the welfare $W(y)$ that would prevail had the WACC been set correctly. The loss from under-estimating the true WACC (i.e., choosing $x < y$) is the welfare foregone as a result of the lack of investment when the RAR is set too low i.e., the entire welfare $W(y)$. (Later, we consider possibilities for the welfare function $W(\cdot)$.)

The expected loss from choosing the RAR to be x is then

$$L(x) = \int_0^x (W(y) - W(x))f(y)dy + \int_x^\infty W(y) f(y)dy.$$

Minimising this with respect to x gives the first-order condition for the loss-minimising solution x^* .¹³⁸

$$-W'(x^*)F(x^*) - W(x^*)f(x^*) = 0.$$

This can be re-written as

$$-\frac{W'(x^*)}{W(x^*)} = \frac{f(x^*)}{F(x^*)}.$$

The left-hand side is the percentage change in welfare for a small (marginal) change in the RAR. The right-hand side is the reverse hazard rate: the probability that the true WACC is close to the RAR, conditional on the true WACC being no more than the RAR. It is more intuitive still to write this in terms of elasticities:

¹³⁸ We need to check that the second-order condition is satisfied. This imposes conditions on the welfare function and the distribution of the true WACC. These conditions are satisfied with the demand and distribution functions used below.

$$-x^* \frac{W'(x^*)}{W(x^*)} = x^* \frac{f(x^*)}{F(x^*)},$$

which balances, for a small change in the RAR, the proportional change in welfare (the left-hand side) against the proportional change in the probability that the RAR is above the true WACC (the right-hand side).

The welfare measure

The choice of welfare measure is not without controversy: in the New Zealand case, for example, there was considerable discussion of whether a consumer or total surplus measure ought to be used. On balance, it seems reasonable to suppose that regulators have a primary duty to consider consumer surplus, subject to ensuring that regulated firms can finance properly their operations (including ensuring that the WACC is appropriate). Let the demand curve be denoted $D(p)$. Then let

$$CS(p) \equiv \int_p D(z) dz,$$

given a price p . Note that $CS'(p) = -D(p)$. The welfare function $W(x)$ is then $W(x) = CS(p(x))$, where $p(x)$ is the regulated price that results from setting the allowed WACC at level x . To specify this further, let the regulated price be $\tilde{p}(x) = c + xk$, where c is the marginal operating cost, k is the per unit cost of capacity, and x is the allowed WACC. It will in fact be easier to deal with the regulated price relative to the capacity cost $p(x) \equiv \frac{\tilde{p}(x)}{k} = \frac{c}{k} + x$. Note that this relative price goes up by one unit when the allowed WACC x goes up by one unit. In this case then, $W'(x) = -CS'(p(x))p'(x) = -D(p(x))$, and so

$$-\frac{W'(x^*)}{W(x^*)} = \frac{D(p(x^*))}{W(x^*)},$$

and so the condition for the optimal RAR is

$$\frac{D(p(x^*))}{W(x^*)} = \frac{f(x^*)}{F(x^*)}.$$

Assumptions to quantify the RAR

To pin down the degree of aiming up or down, we need to make some specific assumptions about demand and the distribution of the true WACC. For the latter, assume

Assumption A1: The true WACC y is normally distributed with mean μ and standard deviation σ .

Assumption A1 is a pretty standard distributional assumption, but it does impose certain conditions on the true WACC. One of these is that we can normalise the term on the right-hand side of the first-order condition (technically, the normal distribution belongs to the family of location-scale distribution functions). To do so, define

$$z^* \equiv \frac{x^* - \mu}{\sigma};$$

that is, z^* is the optimal degree of bias when choosing the loss-minimising RAR x^* , normalised by the degree of uncertainty σ . Then

$$\frac{f(x^*)}{F(x^*)} = \frac{\phi(z^*)}{\Phi(z^*)}$$

where $\phi(\cdot)$ and $\Phi(\cdot)$ are the pdf and cdf of the standard normal distribution. $\Phi(z^*)$ is, therefore, the percentile choice of the optimal RAR x^* : a value of $\Phi(z^*) = 0.5$ means that z^* is 0, and so $x^* = \mu$ i.e., the optimal choice of the RAR is the mean of the distribution of the true WACC. Aiming up means that $\Phi(z^*) > 0.5$; aiming down, the converse.

Next, we have to make assumptions about the demand function. This is a trickier exercise. Dobbs (2011) presents perhaps the most rigorous exploration of expected welfare losses from setting an RAR that differs from the true WACC. One distinct limitation of his analysis is the reliance on an iso-elastic demand curve. In order for welfare measures to be well-defined in this case, it is necessary to assume *elastic* demand: absolute values of the price elasticity of demand greater than 1. But this is at odds with the empirical evidence for regulated commodities such as water, electricity, gas etc.. The essential nature of these goods is such their elasticities tend to be below 1: sometimes, significantly below. For many regulated goods and services, we expect demand to be inelastic ($\epsilon < 1$), and indeed for the price elasticity of demand to be well away from 1: commodities such as water, electricity, gas etc. are essentials. Typical values of the elasticity for electricity, for example, lie in the range 0.4-0.9, depending on whether the demand in question is residential, commercial or industrial.¹³⁹ Typical values for residential water use lie in the range 0.4-0.6 (depending on whether short-run or long-run responses are analysed), although very low estimates (less than 0.25) are derived when looking at e.g., short-run responses during particularly dry spells.¹⁴⁰

So, instead of looking only at an iso-elastic demand function, consider three, fairly standard options:¹⁴¹

D1: Demand is iso-elastic, taking the form $D(p) = a p^{-b}$;

D2: Demand is linear, taking the form $D(p) = a - b p$;

D3: Demand is semi-log, taking the form $D(p) = a e^{-b p}$,
for non-negative constants a and b .

With these different demand functions, the ratio $\frac{D(p(x))}{W(x)}$ (the left-hand side of the optimality condition) is, respectively:

D1: $\frac{b-1}{p(x)}$, noting that this requires $b > 1$, i.e., demand is elastic;

D2: $\frac{2b}{a-b p(x)}$;

D3: b .

The following figures illustrate the outcomes from these different forms of demand, showing the percentile choice of the loss-minimising RAR x^* (i.e., they plot $\Phi(z^*)$), on the vertical axis, against the price elasticity of demand that results, on the horizontal axis. Note that for the iso-elastic demand function D1, the elasticity is greater than 1 (as is necessary in this case); for the other demand functions, we focus on values of the elasticity less than 1, in line with the empirical evidence.¹⁴²

Figure I.1: Percentile choice of the optimal RAR against elasticity for iso-elastic demand

¹³⁹ See Paul, Myers and Palmer (2009).

¹⁴⁰ See Espey, Espey and Shaw (1997).

¹⁴¹ Again, the price implicitly used in these demand functions is relative to the per unit capacity cost.

¹⁴² The parameter values used in these figures are: $\mu = 5\%$, $\sigma = 1\%$, $c = 0.05$.

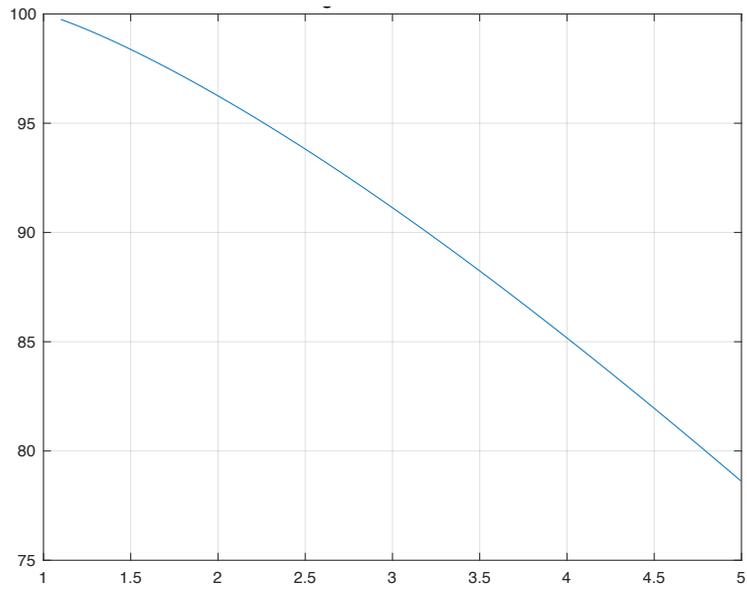


Figure I.2: Percentile choice of the optimal RAR against elasticity for linear demand

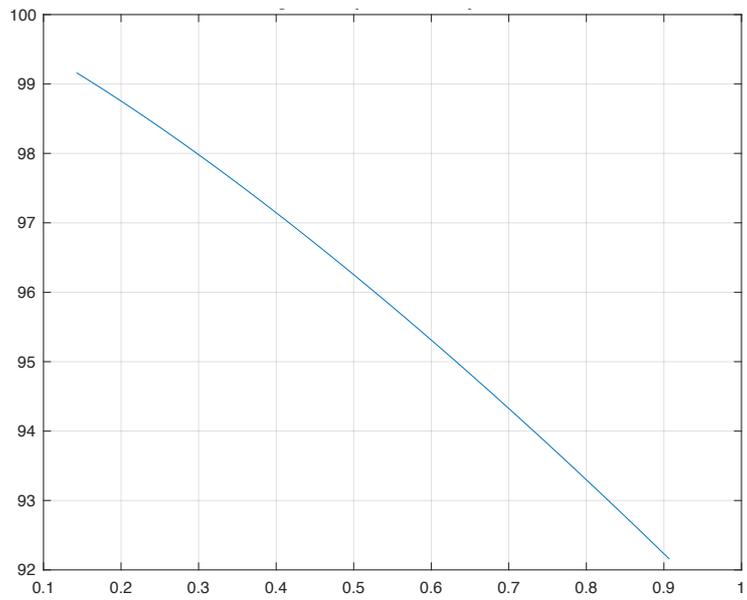
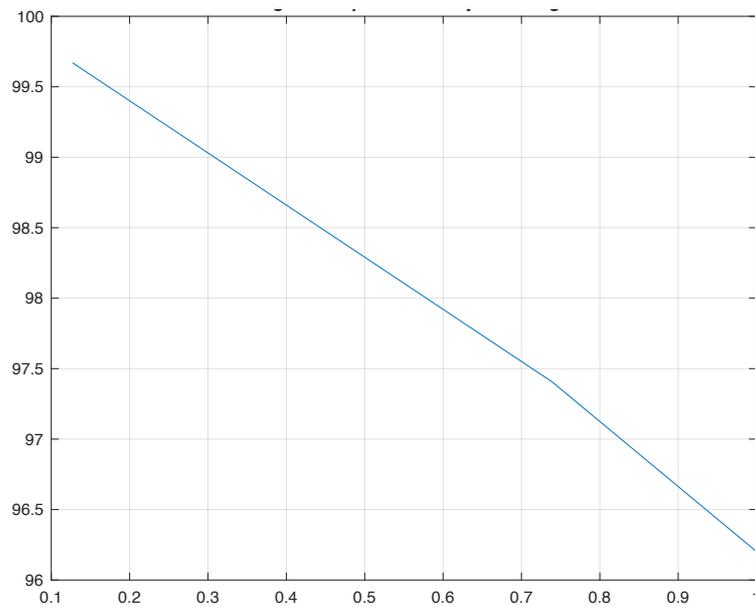


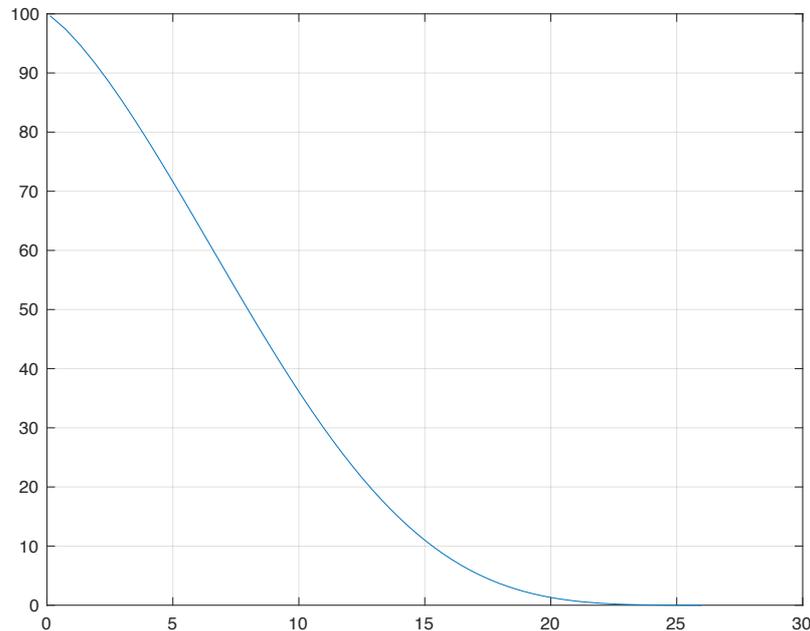
Figure I.3: Percentile choice of the optimal RAR against elasticity for semi-log demand



What the three figures have in common is that the degree of aiming up is high: for elasticities less than 1 with the parameterisation used in these figures, above the 90th percentile. Nor is this dependent on the particular parameterisation: other values yield the same result. (Indeed, with semi-log demand, the optimal percentile is necessarily independent of μ , σ and c , since the first-order condition is $b = \frac{\phi(z^*)}{\phi'(z^*)}$.)

In order to get lower percentile choices, quite extreme values of elasticity are required. The figure below shows this for the semi-log demand function D3 (using previous parameter values). For the optimal RAR to be chosen at the mean level μ , the elasticity has to be around 8: a very high value relative to empirical estimates.

Figure I.4: Percentile choice of the optimal RAR against elasticity for semi-log demand (extended)



This is hardly surprising: with relatively low elasticities, the reduction in consumer surplus from setting a higher price is relatively small. In contrast, the welfare loss from setting the price too low (i.e., setting too low a RAR) is relatively large. This leads to considerable aiming up, as the optimal choice by the regulator.

In summary: when the consequence of setting too low a RAR is a complete loss of investment, the optimal choice of the RAR (and hence, in this simplified framework the RER) is high, in terms of the percentile within the range of distribution of the true WACC. This conclusion does not depend materially on the particular parameterisation used in the analysis.

The loss function revisited for sunk investment

What drives this analysis is the feature that no investment occurs if the RAR is set below the true WACC. This feature matches the situation for new investment, yet to happen; but not old investment that is already sunk. For the latter, there is (obviously) no risk that investment will not occur. As a result, the loss incurred when choosing x is just

$$L(x, y) = W(y) - W(x),$$

and so the expected loss from choosing the RAR to be x is then

$$L(x) = \int (W(y) - W(x))f(y)dy.$$

If the welfare function W is consumer surplus, then strict minimisation of this would lead to a RAR of zero, to ensure the lowest possible regulated price and therefore highest possible consumer surplus. This would, however, fall foul of regulators' requirement to ensure financeability of firms. Balancing

these two considerations, it is clear that the optimal RAR (and hence RER) for old (sunk) investment is therefore the expected WACC, μ .

J What drives bid premia for regulated utilities (Phil Burns)

Introduction

Since the utilities were privatised in the 1980s and 1990s, there has been an active market in corporate control. A useful reference point for the valuation of the business is the ratio of the market valuation (excluding any unregulated activities) to the Regulatory Asset Base (RAB) that buyers have paid for the business, commonly known as the market to asset ratio (MAR).

The MAR can be calculated for publicly listed businesses, by comparing their enterprise value to their RAB or RCV.¹⁴³ Of GB regulated businesses which are listed, only United Utilities and Severn Trent do not have significant unregulated activities, and can therefore be used to estimate MARs that lend themselves to a reasonably straight-forward interpretation that is not overly complicated by the listed entity's wider activities. An analysis of the implications of the MAR for these companies has undertaken by PwC for Ofwat as part of the PR19 methodology statement.¹⁴⁴

Where the businesses are not listed, the MAR can only be estimated at the point at which the businesses are sold – which allows us to estimate the so-called transaction premium. For example, it is estimated that the transaction premium associated with the sale of a majority stake the National Grid Gas Distribution businesses in March 2017 was 49% (i.e. a MAR of 149%).¹⁴⁵

As the figure below (taken from the PwC analysis) shows, the MAR of the energy and water businesses has varied considerably over time¹⁴⁶. Over the past 20 years it has averaged 124%, with episodes where businesses were sold at a discount to the RAB, and periods where there have been significant premia. Over the course of 2016, transaction premia (and market premia for listed companies) were significantly above average.

In this appendix, we identify the potential drivers of premia, and undertake some analysis of their relative importance. The primary aim of this appendix is to assess the transaction premia in energy networks and we have not reviewed the MARs for the listed utilities, though some of the observations and findings may be equally relevant.

¹⁴³ Enterprise value = market capitalisation of equity shares + book value of debt

¹⁴⁴ PwC, Updated analysis on cost of equity for PR19, December 2017.

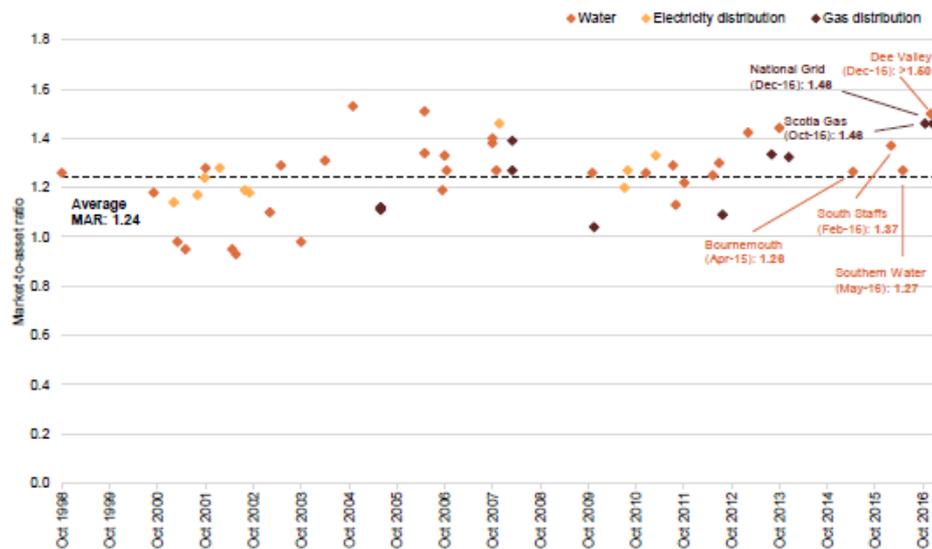
¹⁴⁵ National Grid (8 December 2016) Sale of majority interest in National Grid Gas Distribution. Proposed one-off £4 billion return of capital to shareholders.

Enterprise value adjusted for “positive fair value adjustments of approximately £0.4 billion on debt novated from National Grid Gas plc in connection with the Transaction.”

Depending on the price base used to measure RAV, this could range from 46.6% (using a price base of 2017/18) to 50.7% (using a price base of 2016/17)

¹⁴⁶ The data in the chart relates to transaction premia. We note that transaction premia can be harder to interpret than the market premia of quoted companies, due to a range of other factors that can drive transaction values. This is explored further later in the appendix.

Figure J.1 Market to asset ratio for UK regulated utilities, 1998-2017



Source: PwC Valuations, Macquarie

Note: Average MAR shown excludes multiples for South Staffordshire and Dee Valley for which there is greater uncertainty over the ratio value.

Source: PwC (2017), p.84

The drivers of the premia

The premium over the RAB that investors are prepared to pay is driven by their expectations of the future performance of the business against the assumptions used by the regulator in setting the price controls. If there is no expected outperformance (or expected under-performance) of the regulator's assumptions, then the market value will be equal to the RAB (MAR=100%, premium=0%). If, on the other hand, there is expected to be outperformance of the regulator's assumptions in future, then the MAR will be greater than 100%. It is important to note that these expectations relate not just to outperformance of the price control that currently exists, but also to future price controls as well. Whilst investors may be able to form reasonably informed expectations of outperformance in the current price control period, their expectations of outperformance against price controls that are not yet in place will be more speculative.

There are four broad sources of outperformance.

- Cost outperformance. This is the extent to which the companies are expected to spend less in future than the regulator expects. The incentive-based arrangements that exist in the UK are designed to provide businesses with the incentive to meet and beat the cost targets set by the regulator through allowing them to retain a proportion of their efficiency savings. The efficiency savings so stimulated through the incentive regime will be shared with customers through the cost sharing methodologies employed by the regulators. These incentives for cost discovery – whilst having a strong efficiency motivation that benefits customers in the longer term – also give rise to the ongoing likelihood that companies will usually earn some excess returns by outperforming their cost targets, which will be factored into company valuations.
- Output incentive outperformance. In many sectors, regulators put in place incentive arrangements that reward or penalise companies based on various measures of performance. Outperformance on these measures also creates excess returns which will be factored into company valuations.

- Cost of debt outperformance. All regulators make some allowance for the cost of debt incurred by the businesses they regulate. For example, Ofgem compensates energy networks for their debt costs based on the yield from iBoxx indices. If networks can manage their debt book such that the effective cost of debt is lower rate than this, they retain the entirety of the difference as excess returns during the price control period, and possibly beyond.
- Difference between allowed and actual cost of equity. If a regulator determines the allowed equity return based on a long-term measure of market returns, it is inevitable that in the short run there will be deviations from the actual cost of equity, which will be factored in by investors when making valuations. This could be due to cyclical variations, which regulators may not reflect in their allowances, or structural changes, which regulators respond to but with a time lag.

Given this range of sources of potential outperformance, we have considered whether, using data available in the public domain on existing outperformance, it is possible to identify which of the above may be the most important drivers of premia, including those at or around recently observed levels.

Quantifying the drivers of the premia

To support our assessment of the key drivers of transaction premia we use a simple stylised model provided by Ofgem to evaluate the importance of the four drivers. Cashflows are determined by an assumption on growth and a starting level of RAB, in addition to the chosen parameters for the four outperformance drivers. The premium is then estimated by discounting cashflows and adding the debt value to give the enterprise value, and comparing this to the RAB.

We developed a range of values for outperformance on cost, output incentives and cost of debt based on recent energy outturn outperformance, which can provide a guide for reasonable expectations of future outperformance. We note that in developing these scenarios we have presumed that investors will take the view that outperformance can be broadly sustained into the future. We consider that this is consistent with assumptions commonly made by investors (i.e. that well run companies will be able to beat regulatory targets on a sustained basis), historical evidence which shows that there has indeed been consistent outperformance; and the intended design feature of incentive-based regulation which is to encourage companies to seek out efficiency improvements to beat their targets.

The range of values for cost, output and debt outperformance are summarised below, expressed as a percentage of return on regulatory equity (RoRE). The rationale for the chosen parameters is discussed below.

Figure J.2 Outperformance parameter scenarios

Parameter	Low case	High case
Cost outperformance (RoRE %)	RIIO-1: 2.12%	RIIO-1: 2.89%
	RIIO-2: 2.12%	RIIO-2: 2.89%
	Thereafter: 2.12% ¹	Thereafter: 2.89% ²
Output incentive outperformance (RoRE %)	RIIO-1: 0.96%	RIIO-1: 1.05%
	RIIO-2: 0.96%	RIIO-2: 1.05%
	Thereafter: 0.96% ¹	Thereafter: 1.05% ¹

Parameter	Low case	High case
Cost of debt outperformance (in cost of debt % terms, i.e. average coupon vs. allowance)	RIO-1: 1.96%	RIO-1: 2.16%
	RIO-2: 0.90%	RIO-2: 1.10%
	RIO-3 & 4: 0.61%	RIO-3 & 4: 0.81%
	Thereafter: 0.00% ³	Thereafter: 0.20% ⁴
Cost of debt outperformance (in RoRE % terms ⁵)	RIO-1: 3.64%	RIO-1: 4.01%
	RIO-2: 1.67%	RIO-2: 2.04%
	RIO-3 & 4: 1.13%	RIO-3 & 4: 1.50%
	Thereafter: 0.00%	Thereafter: 0.37%

Source: ¹ RoRE outperformance of Cadent GDNs during RIO-1, as reported in Ofgem's RIO-GD1 annual report data, 2015-16

² Regulatory equity-weighted RoRE outperformance across the GD sector during RIO-1, as reported in Ofgem's RIO-GD1 annual report data, 2015-16, p.18

³ Frontier projections based on iBoxx 10 year trailing average minus NCGD real cost of debt

⁴ As the Low case, plus 20bps

Note: ⁵ This is done by apportioning up the outperformance based on the gearing level, as the RoRE applies to the equity portion of the RAV, whereas the CoD outperformance applies to the debt portion.

We also explore the extent to which premia may be driven by divergence between the allowed and actual cost of equity. To do this we develop illustrative cost of equity scenarios, which we discuss below.

Additionally, there are a number of other benefits at the group level, which investors could factor into their valuation of a business. These include the potential to further restructure the financing of the business, tax arbitrage opportunities, and a control premium if the purchaser is buying a majority stake in the business. We do not explicitly model these effects but will discuss these other factors later.

Cost and incentive outperformance

To develop scenarios for cost and incentive outperformance, we draw data from the gas distribution sector. We develop a lower case cost outperformance based on Cadent's forecast retained total expenditure (totex) underspend in RIO-GD1, of 8.7%.¹⁴⁷ The high case is based on the weighted industry forecast average outperformance, which equates to a 12.3% underspend on allowed totex.¹⁴⁸ For output incentive outperformance, the low and high cases are also based on Cadent and sector outperformance in RIO-GD1.

¹⁴⁷ Underlying figures can be found in Ofgem's 2015-16 RIO gas distribution annual report data, "RoRE" and "Forecast totex" tabs. Calculated by comparing forecast totex over RIO-GD1 to allowed totex over the period. Data accessible from: <https://www.ofgem.gov.uk/publications-and-updates/riio-gas-distribution-annual-report-2015-16>

¹⁴⁸ Ibid.

Below we set out these assumptions in tabular form and provide the premia that result from these outperformance assumptions. We assume at this point that the actual and the allowed return on equity is 6.7% in RIIO-1, and 6% (same as RIIO ED1) thereafter¹⁴⁹. We also assume that there is no cost of debt outperformance. The premia estimated under the Low and High scenarios are reported below and since we have essentially “switched off” outperformance from other sources, the resulting premia are the effect arising from cost and incentive outperformance alone. As the figure illustrates, the estimated premium due to cost and output outperformance varies between 19% and 25%.

Figure J.3 Impact of cost and incentive outperformance on the estimated RAB premium

Scenario	RIIO-1 RoRE %	RIIO-2 RoRE %	Thereafter RoRE %	Premium
Low	3.08%	3.08%	3.08%	19.37%
High	3.94%	3.94%	3.94%	24.77%

Source: Frontier analysis, Ofgem stylised model of risk premium

It can be argued that the scale of cost performance in these assumptions is relatively high and is unlikely to be sustained in future periods. For example, Ofwat data shows that the cost outperformance in the water sectors in the first two years of the current period amounts to 0.7% of RoRE and the incentive outperformance amounts to 0.14% of RoRE¹⁵⁰, suggesting that the RIIO1 experience is on the high side. At the same time, we would note the following.

- A degree of cost outperformance has been a persistent feature of UK regulated networks since privatisation and, as outlined above, is an inherent component of the incentive regulatory model.
- Investors may form their expectations of future outperformance based more heavily on the recent outperformance evidence.
- For an assessment of transaction premia the successful bidder is likely to have the most optimistic assumptions in relation to future outperformance.

We consider the interpretation of these scenarios later in the paper.

Cost of debt outperformance

To develop our cost of debt outperformance scenarios we draw on publicly available evidence from Cadent. As we set out below, there is evidence to suggest that Cadent will outperform significantly its RIIO-1 cost of debt allowance. The primary reason for this is that at the time of the transaction Cadent did not have a book of higher cost embedded debt. Cadent’s debt book was largely refinanced as part of the recent sale process and hence Cadent presently has a debt portfolio with a low average coupon that gives rise to a substantial debt outperformance. Depending on Ofgem’s policy at RIIO-2, this outperformance may persist. Outperformance on this scale might be regarded as somewhat unusual, as most regulators will aim to develop a cost of debt methodology that broadly allows for the recovery of

¹⁴⁹ When assessing cost, incentive or debt performance the assumption on the cost of equity is relevant to the extent that it affects the rate of discounting of expected future outperformance.

¹⁵⁰ Ofwat, Monitoring financial resilience, November 2017, page 13.

reasonably incurred debt costs, with more limited scope for outperformance. The cost of debt analysis is more relevant to an appraisal of high recent transaction premia, rather than MARs for listed utilities.

The low end of the range for the cost of debt outperformance is based on forecasts of Cadent’s cost of debt outperformance reflecting public information on its existing debt book and the likely trajectory of the rolling cost of debt allowance. We assume that when the company replaces its existing low coupon debt in future, it would do so at a rate similar to the iBoxx indices. Hence the present debt outperformance is assumed to reduce over time, but will sustain into the RIIO-GD2 period. We note that the extent of outperformance into GD2 will depend on the cost of debt policy adopted by Ofgem, which may change from RIIO-GD1. Nevertheless, this assumption serves as a helpful illustration of the potential upside.

The high scenario we have developed adds 20 basis points to the low scenario assumptions, reflecting an optimistic investor’s belief on a possible ‘halo effect’ (i.e. that the company can issue debt at a lower rate than the spot iBoxx yield).

We again, at this stage, make the same assumptions about the cost of equity, and assume no totex or incentive outperformance to enable us to isolate the effect of cost of debt outperformance alone. We note that in making this assessment we assume that gearing matches Ofgem’s regulatory assumption (i.e. 65%).

Figure J.4 sets out analysis of these cost of debt outperformance assumptions on RoRE and the premium. The estimated premium that results from debt outperformance alone varies between 10% and 12%. Again it is worth emphasising that the majority of this premium relates to the specific debt structure of the Cadent transaction, with no high cost embedded debt. One would not expect this scale of premium to be replicated across energy and water networks as a whole.

Figure J.4 Impact of cost of debt outperformance on the estimated RAB premium

Scenario	RIIO-1 RoRE %	RIIO-2 RoRE %	RIIO-3 & 4 RoRE %	Thereafter RoRE %	Premium
Low	3.64%	1.67%	1.13%	0.00%	9.70%
High	4.01%	2.04%	1.50%	0.37%	12.04%

Source: *Frontier analysis, Ofgem stylised model of risk premium*

Allowed cost of equity deviation

Next we test the impact of the allowed return on equity deviating from the actual cost of equity. Estimating the impact of this is somewhat more complex than estimating the impact of other outperformance. Whilst the actual cost of equity assumed by the winning bidder for a business is unknowable, we can at least develop a range of possible values (as with the other drivers of outperformance). However, that range of assumed values impacts not just on the range of difference between the actual and allowed cost of equity (outperformance), it is also the rate which is used to discount the outperformance cashflows.

Consequently, in order to make our assessment, we need to make two sets of assumptions, in respect of:

- the actual level of the current cost of equity; and

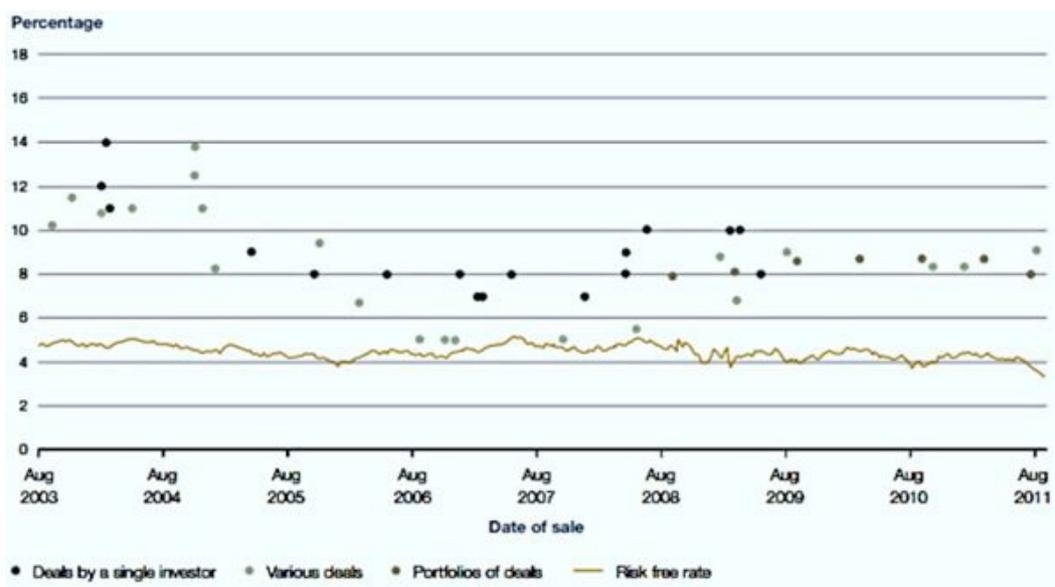
- the extent to which investors may believe that any gap between allowed and actual cost of equity will persist over time.

To inform a range of possible current values of the cost of equity we have considered alternative sources of market-based evidence of required investor returns. We have focussed on this evidence, rather than estimates based on CAPM or DDM methods, for the principal reason that the other methods are fully covered in the main part of the paper. We have aimed to ensure that the estimates are consistent across the different methods.

To start, we draw on evidence from the OFTO tendering experience. The NAO¹⁵¹ quoted OFTOs as bidding equity IRR's (nominal, post-tax) during TR1 in the range of 9-11%. In a later report for Ofgem, CEPA¹⁵² reports that “in subsequent tender rounds (i.e. TR2 and TR3) OFTOs required equity returns have not increased, and in many cases have fallen closer to reported secondary market rates of return in PFI projects”.

CEPA do not present evidence on recent secondary market returns in UK PFI, but historically, as the figure below shows, there has tended to be a premium of around 4% over the risk-free rate, at least over the historical period analysed by the NAO.

Figure J.5. Reported secondary market returns in UK PFI



Source: NAO

A more recent report by the NAO¹⁵³ has investigated investor required returns on deals since 2013. It concluded that most involved a premium above the risk-free rate of between 2% and 4%.

“The 2010 National Infrastructure Plan estimated an indicative cost of capital for PFI as 2% to 3.75% above the cost of government gilts. Data collected by IPA [Infrastructure and Projects Authority] on PFI and PF2 deals entered into since 2013 show that debt and equity investors are forecast to receive a return of between 2% and 4% above government borrowing. However, some

¹⁵¹ Offshore electricity transmission: a new model for delivering infrastructure, NAO, June 2012

¹⁵² Evaluation of OFTO tender round 2 and 3 benefits, CEPA, March 2016

¹⁵³ National Audit Office, PFI and PFI2, January 2018.

2013 deals, agreed when credit market conditions were poor, projected an annual return for debt and equity investors of over 8%; this was more than 5% higher than the cost of government borrowing at the time.” [para 1.17]

However, it is worth noting that the OFTOs and Private Finance projects have very different risk characteristics to most of the network utilities. Compared to established utilities operating under traditional regulatory frameworks, OFTO are awarded contracts that give rise to a regime with the following characteristics:¹⁵⁴

- Fixed 20-year revenue stream, index linked to RPI.
- No price controls, so no regulatory reset risk (although some residual political/regulatory risk may remain should the OFTO model be revised retrospectively).
- No construction risk (at least all existing OFTOs for which evidence is available have been delivered under the “generator build” model under which the OFTO faces no construction risk).
- The OFTO asset value is fully depreciated by the end of initial 20 year revenue stream, implying no terminal value risk.
- Financing can be largely completed upfront, implying very limited refinancing risk (but with some scope for refinancing upside).
- Limited counterparty/bad debt risk, as the counterparty to OFTO contracts is National Grid.
- No exposure to generator performance.

This leaves only (capped) availability risk and operational cost risk to be borne by the OFTO. Private Finance projects involve greater construction risk than OFTOs but like OFTOs face lower risks than regulated utilities in terms of financing, regulatory and counterparty risks.

Consequently, whilst the required returns from OFTO and PFI projects is informative, given the risk characteristics of these projects, they represent the low end of the range of comparable values for network utilities.

The figure below converts the PFI evidence on the overall premium to the government borrowing rate into an implied cost of equity. Using the top of range for the risk premium (3.5% to 4.0%) to reflect the risk characteristics of regulated utilities and assuming a debt spread of 2.0% and gearing of 65% implies a real cost of equity between 4.4% and 5.8%.

Figure J.6 Cost of equity implied from OFTO / PFI data

Risk premium over government bonds	3.5%	4.0%	Top of range implied by PFI data
Nominal government yield	1.25%	1.25%	Average yield on 10 year gilts in 2017
Nominal WACC (post-tax)	4.75%	5.25%	Sum of government yield and risk premium
Nominal cost of debt	3.25%	3.25%	Sum of government yield plus average iBoxx spread 2012-2017

¹⁵⁴ See for example the description provided by KPMG in an investor perspective report prepared for GEMA. <https://www.ofgem.gov.uk/ofgem-publications/85943/offshoretransmission-aninvestorperspective-updatereport.pdf>

Risk premium over government bonds	3.5%	4.0%	Top of range implied by PFI data
Gearing	65%	65%	
Nominal cost of equity	7.53%	8.96%	
RPI inflation	3.0%	3.0%	
Real cost of equity	4.4%	5.8%	

Source: Frontier analysis

Taken together with the earlier OFTO evidence, we consider that this market evidence might suggest a lower bound in the order of 4.5% (real RPI, post tax). We note that in its December 2017 methodology for PR19 Ofwat was provided cost of capital guidance included a cost of equity of 4.0% (real RPI). This figure was based on 60% gearing and converting it to 65% gearing (using Ofwat's assumptions on debt beta etc) would imply a cost of equity of 4.6%, above the lower bound included here.

We consider that these values provide a relevant point of reference to inform the range that an investor might use for the current cost of equity to inform current valuations.

Equally important, if not more important, than estimating potential values for the actual cost of equity we also need to estimate the potential gap between the actual and allowed cost of equity. As explained above the variation in the actual cost of equity could reflect cyclical or structural factors. If they are cyclical then the gap between actual and allowed would be short-lived rather than persistent. If the driver of the gap is structural then investors would probably expect the gap to close over time as regulators respond prudently to emerging evidence. We note in this context that Ofwat has reduced the cost of equity from 7.4% at PR09 to 5.4% at PR14 to a proposed 4.0% at PR19¹⁵⁵. This indicates that regulators have been prepared to revise their cost of equity allowances materially from one review to the next and that it would be reasonable for investors to expect any gap between the allowed and actual cost of equity to close over time. That said, as we identified above for cost outperformance, it is the assumptions and beliefs of the most optimistic investors that drive transaction premia. A relative advantage of MAR analysis of quoted companies over transaction premia evidence is that it allows the regulator to monitor trends in the premium over time, which could assist in distinguishing cyclical from structural changes.

In our view, this suggests that it might be reasonable to model an equivalent permanent gap of 0%, 0.5%, 1.0% and 2.0% in order to evaluate the incremental impact of these differences. Note that the figures represent a permanent gap between the allowed and actual cost of equity. A 1% permanent gap is equivalent in discounted terms to a current 2% gap that is closed gradually by the regulator over a period of 20 years. We believe therefore that range of 0% to 2% covers a wide range of potential scenarios, and should the reader wish to appraise alternative scenarios around differences and future convergence, the premia we present below can be interpolated/extrapolated in a relatively straightforward manner.

We are now able to proceed to estimate how divergence between allowed and actual cost of equity may drive premia. As noted above, the range of assumed values impacts not just on the range of difference

¹⁵⁵ These figures have been adjusted to 60% gearing. The PR09 figure was 7.1% at 57.5% gearing and PR14 was 5.65% at 62.5% gearing.

between the actual and allowed cost of equity (outperformance), it is also the rate which is used to discount the outperformance cashflows.

Consequently, we estimate the impact of the cost of equity on transaction premia by assuming different levels of outperformance at different levels of the actual cost of equity assumed by the bidders under 3 scenarios:

- Zero outperformance from other drivers (to isolate the pure cost of equity effect);
- In conjunction with low-case outperformance from other drivers; and
- In conjunction with high-case outperformance from other drivers.

The results are presented in the series of tables below.

Figure J.7 Transaction premia arising from different allowed and current levels of CoE (zero outperformance elsewhere)

Actual	Allowed v actual			
	0%	0.5%	1%	2%
4.5%	0%	4.12%	8.22%	16.37%
5%	0%	3.68%	7.35%	14.64%
5.5%	0%	3.33%	6.65%	13.23%
6%	0%	3.04%	6.06%	12.06%
6.5%	0%	2.79%	5.56%	11.07%
7%	0%	2.57%	5.14%	10.23%

Source: *Frontier analysis, Ofgem stylised model of risk premium*

Figure J.8 Transaction premia arising from different allowed and current levels of CoE (low case outperformance elsewhere)

Actual	Allowed v actual			
	0%	0.5%	1%	2%
4.5%	37.64%	41.76%	45.87%	54.01%
5%	34.52%	38.21%	41.87%	49.16%
5.5%	31.92%	35.25%	38.57%	45.16%
6%	29.73%	32.76%	35.78%	41.79%
6.5%	27.84%	30.63%	33.40%	38.92%
7%	26.20%	28.78%	31.34%	36.44%

Source: *Frontier analysis, Ofgem stylised model of risk premium*

Figure J.9 Transaction premia arising from different allowed and current levels of CoE (high case outperformance elsewhere)

Actual	Allowed v actual			
	0%	0.5%	1%	2%
4.5%	48.26%	52.38%	56.48%	64.63%
5%	44.06%	47.75%	51.41%	58.70%
5.5%	40.58%	43.91%	47.23%	53.82%
6%	37.66%	40.69%	43.71%	49.72%
6.5%	35.15%	37.94%	40.71%	46.23%
7%	32.99%	35.56%	38.13%	43.22%

Source: Frontier analysis, Ofgem stylised model of risk premium

The following points emerge.

- Some cells seem to be much less relevant because they imply a regulatory allowed CoE that is higher than the rate that regulators are currently setting (e.g., 5% actual plus 2% difference equals 7%). In our view the assumptions underlying these cells would not be reasonable even for the most optimistic of investors. These cells are shaded out but are still reported as they allow the reader to perform their own interpolations and extrapolations.
- The premia in the first column (0% difference) represents the impact of cost, output and debt outperformance, given zero outperformance on the cost of equity. They are analogous to the figures provided in Figures J.3 and J.4, but vary up and down the column because of the discounting effect of different assumed levels of current cost of equity. More generally, for any given difference between the allowed and current return, the impact on premium is greatest the lower the investor's current CoE – this is the discounting effect.
- The greater the difference between allowed and current levels of CoE, the greater the premium, as one would logically expect.
- From both Figure J.6 and moving rightwards along Figures J.7 and J.8, it is clear that the difference between the regulator's allowed return and investors' perceptions of the current cost of equity has a smaller effect on the transaction premia, compared to the impact of cost, output and debt outperformance.
- Cadent-like levels of premia can only be observed in the low outperformance case when the investors' cost of equity is presumed to be 4.5% and where there is an ongoing expectation of 1% outperformance against the regulator's CoE.
- Cadent-like levels of premia can only be observed in the high outperformance case when the investors' cost of equity is presumed to be 5% and where there is an ongoing expectation of 1% outperformance against the regulator's CoE.

Before offering our conclusions on this analysis, we briefly discuss two further areas where investors may see additional value, which have not been modelled here.

Additional value in financial restructuring

There is the potential to restructure the financing of these businesses and increase the gearing level significantly beyond the notional gearing of 65%. This could be done by raising debt within the

regulatory ring fence, or by retaining a 65% gearing within the regulatory ring fence, but by adding a further debt into a holding company outside (above) the ring fence so as to increase the effective level of gearing from the perspective of the equity investor.

In introducing further leverage into a business, investors will need to consider whether this changes materially the risk profile, such that their required returns change and/or that the risk profile of business changes in the eyes of lenders. However, if it were possible to add leverage without materially changing the perception of investor risk, this could unlock additional value for equity investors. To illustrate the potential scale of further leverage, we have considered the effect of adding a further 20% of debt in our simple model. We scale the estimated premia based on cost of debt outperformance from 65% to an additional 20% of the RAB.

Scaling up the Low estimate of 9.70% gives an additional 2.98% in premium. Scaling the High premium of 12.04% gives an additional premium of 3.70%.

Other potential contributing factors to transaction premia

The above exercise has not quantified a number of other important factors that can generate value for investors and add to the premia they are willing to pay for a business. Key factors include the following.

- Tax arbitrage. Investors may be able to out-perform the allowance for corporation tax included by the regulator. This could be achieved through increasing the proportion of debt financing and / or taking advantage of opportunities for arbitrage across different tax jurisdictions (in accordance with transfer pricing rules). Note that regulators have policies for sharing some types of tax out-performance with customers through clawback mechanisms.
- Control premium. Investors are willing to pay a premium to acquire a majority stake in a business. This may be because there is perceived value in holding some real options on taking future strategic decisions, which is absent when holding a minority share in the business.
- Winner's curse. As suggested in classic auction theory, in a private transaction involving many bidders where the winning bid has the highest valuation, it often holds that the winners tend to be the ones who have overvalued the asset and end up with a loss when the true value of the asset is revealed. (We note that evidence from MARs for quoted pure-play utilities are generally not subject to the issues of control premium and winners curse.)

Conclusion

In this appendix we have sought to investigate the potential sources of outperformance that may lead to premia, in order to understand what may be driving recent observations particularly on the transaction premium of energy networks. We have explored a range of scenarios broadly informed by publicly available information on the extent of existing outperformance, coupled with an assumption that investors may consider it reasonable to presume outperformance will persist into the future.

Our observations on this analysis are as follows:

- If current levels of outperformance on costs and incentives from RII01 were expected to continue then this could justify a RAB premium of between 19% and 25%. These rates of outperformance appear to be relatively high and therefore it would be optimistic of investors to expect them to continue at this rate. At the same time, expectations of cost and service outperformance are integral to an incentive model that is designed to promote innovation and cost discovery; there has been persistent outperformance in these dimensions since

privatisation; and therefore an investor expectation of significant on-going outperformance is not unreasonable.

- The specific features of the Cadent transaction includes significant initial outperformance of the allowed cost of debt, which would be expected to reduce over subsequent periods. This could add 10% to 12% to the RAB premium. Combining our appraisal of cost and incentive outperformance with our appraisal of the (more specific) analysis of debt outperformance results in a substantial premium (mid 30%). This is above the levels observed over the longer term, but still below the levels observed in the recent Cadent transaction.
- Differences between the actual and allowed cost of equity can also be a driver of the transaction premium. The scale of the premium depends on the size of the difference and how long investors expect the gap to persist. In order to explain a material proportion of the observed transaction premium the difference between actual and allowed cost of equity would need to be large (e.g. 2% or over) and sustained. If the difference has arisen from a structural shift in financial market conditions then past experience suggests that regulators would close the gap over time, reducing the contribution to explaining the premium. In the light of regulatory experience it would be an optimistic investor to expect the gap not to be closed over time.
- Our modelling suggests that the observed Cadent premium could be explained by the expected debt outperformance combined with ongoing outperformance on cost and output incentives, outperformance on the cost of equity, coupled with other non-modelled factors, such as further financial restructuring, tax arbitrage, control premium and winners curse.
- If the investors expected the current rate of cost and incentive outperformance to continue and the non-modelled factors also played a role, then the contribution from cost of equity outperformance would be modest.

What is evident from this analysis is transaction premia alone do not provide sufficient evidence to make inferences about the cost of equity. Different drivers of outperformance are at play and multiple combinations of various drivers can explain observed premia. In addition, the role of expected outperformance means that the premia may result from unobserved investor assumptions that may be considered unrealistic or optimistic but are nevertheless the reality behind the premia.

For these reasons, we consider that evidence from transaction premia is less reliable and much harder to interpret than other sources of evidence on the cost of equity.

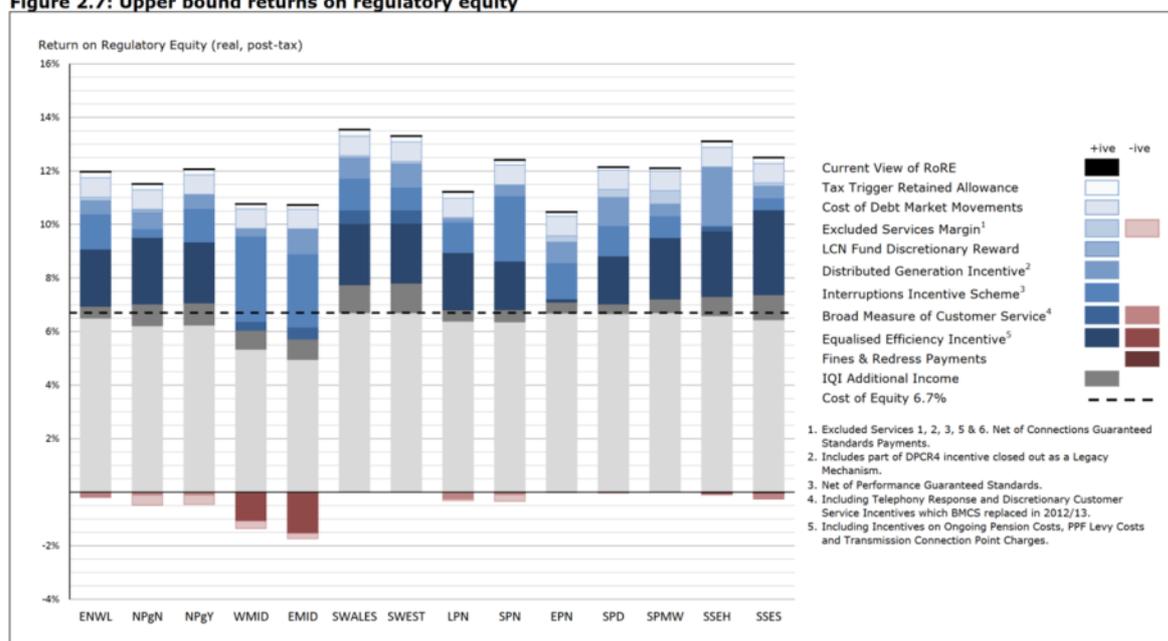
K Evidence of Outperformance (Stephen Wright)

The history of UK regulated companies is relatively short. It is also not well documented. While evidence has been presented in a range of reports on individual price control periods (we present some examples below) there appears to be an absence of any systematic dataset of realised returns on regulated asset values on consistent basis over time. We cannot see any obvious justification for this information not to be readily available in the public domain.

So, on the basis of what information we do have, we must of necessity be cautious in drawing conclusions. But this limited amount of evidence is quite suggestive.

The chart below shows returns on regulatory equity (as opposed to returns on total RAV) for electricity distribution companies over the price control period 2010-15.

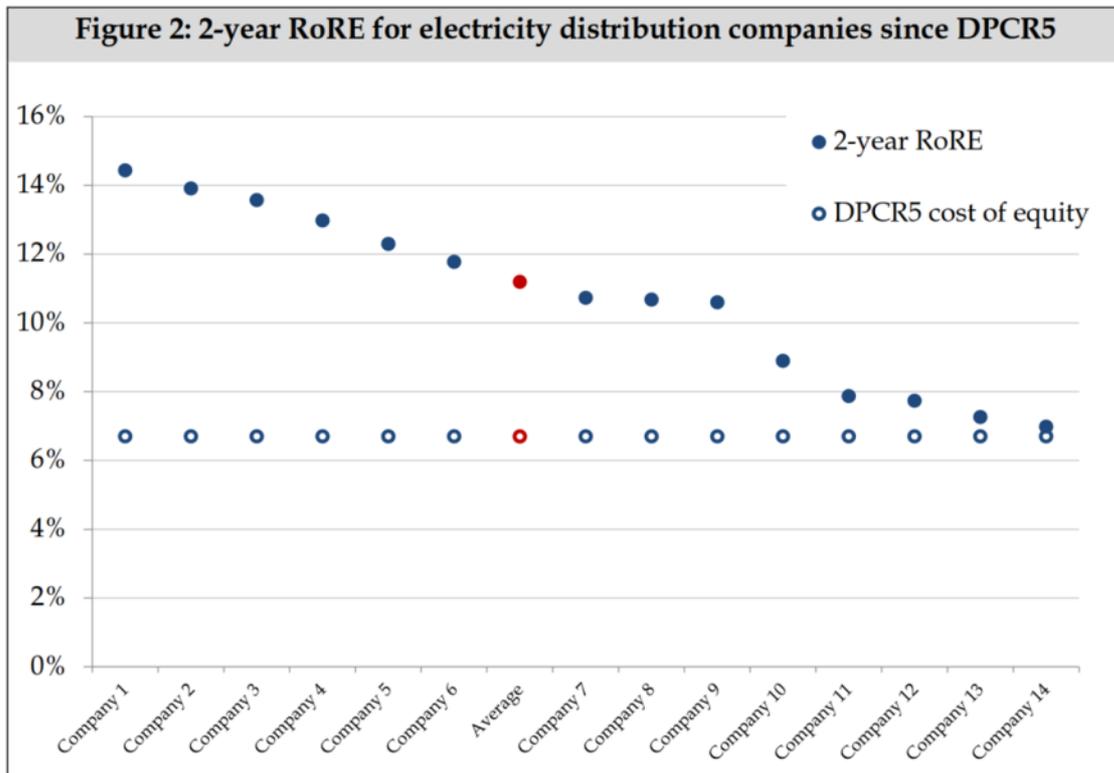
Figure 2.7: Upper bound returns on regulatory equity



Source:

https://www.ofgem.gov.uk/sites/default/files/docs/electricity_distribution_company_performance_2010-2015.pdf

The next chart gives a similar picture for an earlier price control period.

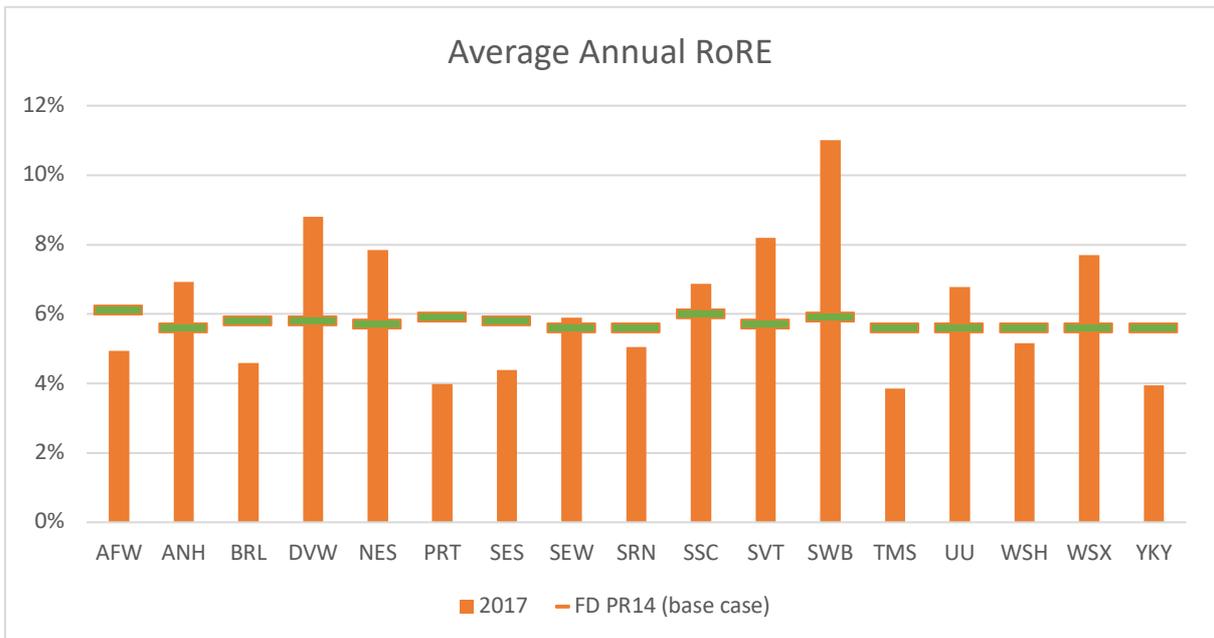
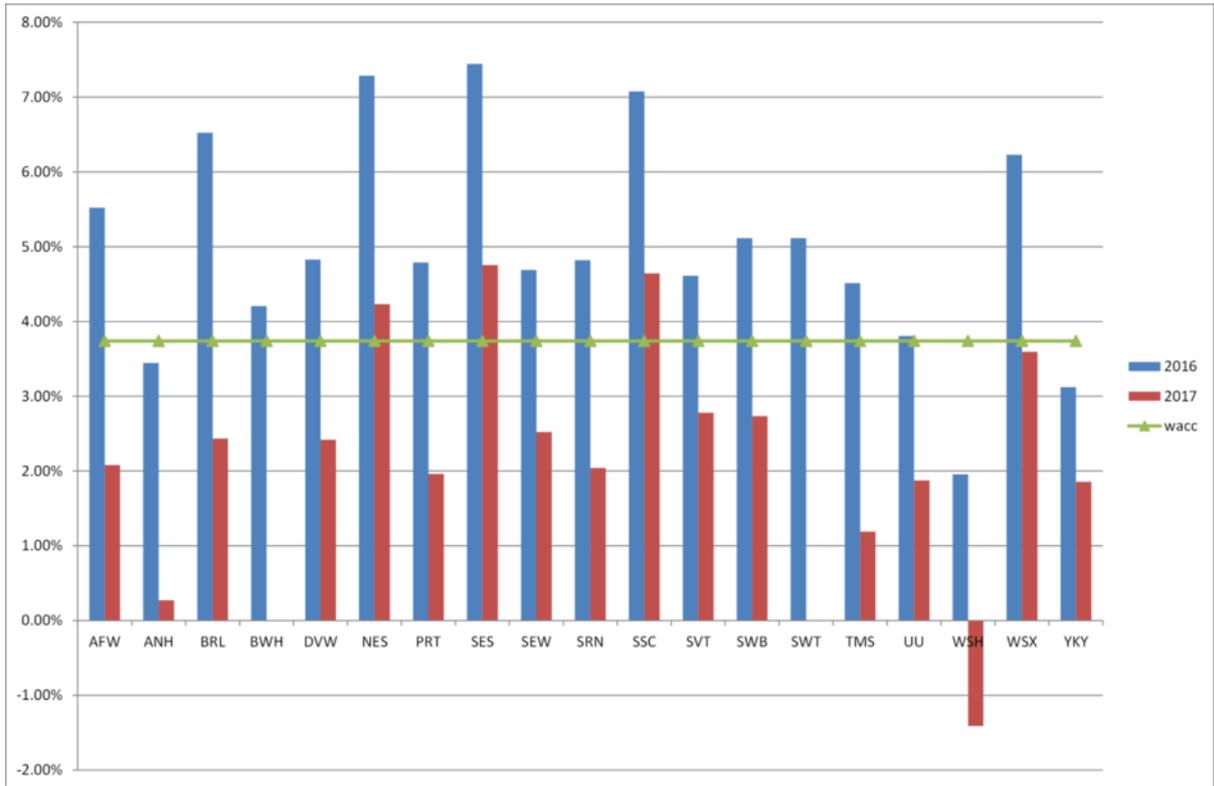


Source: Imrecon, RIIO reviews Financeability study, November 2012

The two charts above suggest that the regulatory allowed return on equity was, at least ex post, very far from an unbiased forecast of the actual realised return: indeed, it has been very much closer to being a lower bound.

We do not (yet) have similar evidence for other regulators. We have been provided thus far with some data from Ofwat: from the most recent price control period (the first chart below), and the cumulative RORE outperformance for the first 2 years of PR14 (the second chart). This gives a more mixed picture, shown below, but, given the very short-term nature of the evidence, and the fact that it is both preliminary and from an incomplete control period, it is hard to draw firm conclusions.¹⁵⁶

¹⁵⁶ We have been advised that the pattern of returns within the control period is at least in part under the control of the regulated companies.



While the evidence above is distinctly patchy, and the sample of data is quite short, nonetheless it is at least suggestive of fairly strong conclusions:

- Regulated companies have been provided with strong incentives to carry out cost savings and service improvements, which they have responded to.

- Regulators have (correctly) pointed to the benefits to consumers of such improvements, although the evidence for this is quite difficult to disentangle given the difficulties of any counterfactual.¹⁵⁷
- But the asymmetry of information between regulators and the companies they regulate has inevitably given an advantage to the companies in judging how difficult it is for them to exceed regulators' targets (a problem that has been explicitly acknowledged by regulators)
- As a result at least on the basis of historical evidence, regulated companies would reasonably have expected to earn returns at least on average, well above their notional allowed returns. Thus as implemented by UK regulators the term WACC has been a misnomer. In terms of the concepts we have introduced in the main report, the implicit Regulatory Expected Return (RER) used by UK regulators has in practice been distinctly higher than the notional allowed WACC set in price controls.
- This expectation of continued outperformance has almost certainly contributed to the recent unprecedentedly high bid premia as documented by Burns' analysis in Section 6.
- From the consumers' perspective, the issue is not really about whether companies have achieved cost savings and service improvements: it is about whether their shareholders have earned economic rents from doing so. The data suggests that, at least ex post, they have.
- In terms of our terminology this has been both because the RAR (aka the "allowed WACC") has been well in excess of a plausible estimate of a true CAPM-WACC, but also because of the element of systematic outperformance, so that the RER has been significantly higher than the RAR.

¹⁵⁷ Burns lays great stress on this in their recommendations above, but the academic evidence on the benefits to consumers are not clear-cut. Examples of recent studies finding ambiguous evidence are: Borghi, E., Del Bo, C., and Florio, M. (2016), Hyland, M. (2016) and Polemis, M. L. (2016).

L Terms of Reference

These are the terms of reference for a proposed study commissioned jointly by the CAA, Ofcom, Ofgem and, the Utility Regulator, collectively referred to as “the regulators”. All are members of the UK Regulators Network (UKRN).

The need for this study emerged from discussions within the UKRN Cost of Capital Working Group acknowledging that key decisions that the regulators will be making in price control decisions for capital-intensive sectors would benefit from authoritative guidance on certain aspects of the cost of capital. Those aspects arise from the evolution of financial market conditions since the financial crisis of 2008. The markets of particular interest are the markets for risk-free assets and the market for equity investment in regulated businesses.

The regulators’ methodologies for estimating the cost of capital have been significantly informed by an earlier jointly commissioned study¹⁵⁸ (the 2003 study). The proposed new study is intended to guide the regulators in assessing whether and, if so, how it would be appropriate to adapt their methodologies in the light of new market conditions. We expect the study to focus on the cost of equity.

To this end, the regulators wish to commission an authoritative study and a report on its findings by a team of academics and consultants who are recognised authorities on the cost of capital for regulated utilities. The report should explain how its conclusions relate to those in the 2003 study, the main drivers of any changes and why such changes are now justified. We do not expect the report to extend beyond about 50 pages.

The study should address the following questions:

General

1. The 2003 study concluded that there was no one clear successor to the Capital Asset Pricing Model (CAPM) for practical cost of capital estimation. In light of the evolution in financial market conditions and any significant empirical or theoretical developments, does this conclusion remain valid?

Long run total equity return and risk-free rate

Regulators conventionally look at a variety of sources when informing their assessment of prospective equity market returns, but have traditionally placed considerable weight on long run historical evidence, for example as shown in the annual DMS analysis¹⁵⁹. Regulators conventionally refer to historical returns on Treasury Bills and gilts and forward yields on index-linked gilts to inform estimates of the long-run risk-free rate.

2. In estimating the prospective long run total equity market return, should regulators adjust the historical evidence of UK or international equity market returns for i) factors suggested in the DMS analysis, including historical relative outperformance in the UK equity market and

¹⁵⁸ ‘A Study into Certain Aspects of the Cost of Capital for Regulated Utilities in the UK’, Stephen Wright, Robin Mason and David Miles, 13 February 2003.

¹⁵⁹ Analysis of investment returns since 1900 in ‘Credit Suisse Global Investment Returns Yearbook 2017’, Elroy Dimson, Paul Marsh and Mike Staunton, published by Credit Suisse Research Institute.

- expansion in the price/dividend ratio, ii) longer-run dividend yields/growth suggested by Fama and French¹⁶⁰ or iii) any other factors?
3. What would be the implications, if any, of a continuing low yield environment for regulators' estimates of the prospective long run equity market return? Similarly, would there be implications for estimates of the long run risk-free rate?
 4. How should regulators take account of structural changes in price indices arising from changes in measurement methodologies in 2010 reported by the ONS in interpreting the history, making real-terms estimates or converting real estimates into forward-looking nominal estimates?
 5. Taken together, do these issues (or any others) indicate an update of the range indicated in the 2003 study of between 6.5% to 7.5% for the equity return?

Medium term estimates of equity returns and the risk-free rate

The regulators' cost of capital decisions can apply to short or medium term periods. Market expectations of the risk-free rate for such a period can be derived from the gilts market at any one time. Regulators have no equivalent direct market evidence for prospective equity returns, though in an environment where medium term market returns are expected to be below the long term historic trend, there may be a need for regulators to increase their focus on prospective equity returns that are expected to prevail through the price control. Regulators are concerned to ensure that the intended effects of their incentive-based price control regimes are not materially distorted by differences between regulatory estimates of the cost of capital and the actual cost of capital that influences investor and corporate decisions.

6. In light of their incentive objectives, please comment on what considerations may be appropriate for a regulator to consider before deciding whether to adopt medium term or long term estimates of the cost of equity.
7. Has the low yield environment since 2008 or any other factor provided evidence to qualify or update the conclusion in the 2003 study that the equity return has been fairly stable over time, implying an inverse relationship between the short and medium term risk-free rate and the equity risk premium? We note that the range suggested by the Competition Commission in its 2014 NIE determination, in part reflecting current market conditions, was 5% to 6.5%¹⁶¹.
8. In light of the possible drivers for lower risk-free rates (eg policy, supply-demand, risk-aversion, increased uncertainty, reduced confidence in market returns), is there merit in placing more weight on alternative approaches to setting the cost of equity, including academic studies, market expectations or DGM-based methods of estimating TMR? Please provide analysis and commentary on the predictive power of alternative the approaches and criteria for their specification. As a subsidiary question, if using DGM methods of interpreting TMR, should they incorporate adjustments to ensure consistent with arithmetic means of expected returns.
9. Would it be appropriate to adopt a negative real risk-free rate for the estimation of the cost of capital if one is supported by the market evidence?
10. While regulators may have clear evidence of current market expectations for the risk-free rate over a future period of time, there remains uncertainty. Please comment on what considerations may be appropriate for a regulator to consider before deciding to build-in mechanisms to adjust for market movements or make an ex ante estimate of the risk-free rate to ensure its cost of capital estimates remain robust to uncertain future market movements.

¹⁶⁰ Fama, E F and French, K R, 'The Equity Premium', Journal of Finance, April 2002

¹⁶¹ Paragraph 13.147, Final Limited price determination, Competition Commission, March 2014

Interpreting sector-specific evidence in water and energy networks

Recent equity transactions in both water and energy networks have revealed relatively high enterprise valuations, significantly in excess of the respective regulatory values. Some commentators have ascribed these high values to their relative low risk status¹⁶². Regulators are concerned to avoid any systemic mis-estimation of the cost of capital. Ofgem has also received advice and carried out analysis to inform its understanding of the risk profile of its regulated networks. Meanwhile, observed beta in returns to equity in the remaining listed water and energy companies do not indicate any structural reduction in beta risk in recent years.

11. Please reconcile the conclusions the study reaches on the questions above (1 to 10) with the evidence from recent equity transactions, beta observations and Ofgem's analysis of the risk profiles of its networks. We envisage a coherent, robust methodology for triangulating these issues, identifying areas of uncertainty and risks.
12. If it is not possible satisfactorily to reconcile these sources of evidence, please comment on what considerations may be appropriate in ascribing weight to irreconcilable factors in determining an estimate of the cost of equity in these sectors?

¹⁶² For example, in its 12 May 2017 equity research report on National Grid, Barclays noted "recent high regulated utilities valuations have primarily been driven by scarcity of reasonable returns available from other similar relatively low risk investments in the current low bond yield environment."

Works Cited

- Afik, Z., & Benninga, S. (2010). *The Term Structure of Expected Bond Returns*. Retrieved from SSRN: <https://ssrn.com/abstract=1628036>
- Afik, Z., & Benninga, S. (2014). *Expected Bond Returns and the Credit Risk Premium*. Retrieved from SSRN: <https://ssrn.com/abstract=2373845>
- Aharoni, G., Grundy, B., & Zeng, Q. (2013). Stock returns and the Miller Modigliani valuation formula: Revisiting the Fama French analysis. *Journal of Financial Economics*, 110(2), 347-357.
- Alessandri, P., Robertson, D., & Wright, S. (2008). Miller and Modigliani, Predictive Return Regressions and Cointegration. *Oxford Bulletin of Economics and Statistics*, 70(2), 181-207.
- Alexander, I., & Irwin, T. (1996). *Price caps, rate-of-return regulation, and the cost of capital*. World Bank.
- Alexander, I., Mayer, C., & Weeds, H. (1996). *Regulatory structure and infrastructure firms: an international comparison*. World Bank Policy Research .
- AonHewitt. (2017). *Factor Investing*. Retrieved from AonHewitt: http://www.aon.com/unitedkingdom/attachments/aon_hewitt/i/Factor-Based-Investing-July-2017.pdf
- Asness, C. S. (2016). The Siren Song of Factor Timing aka “Smart Beta Timing” aka “Style Timing”. *The Journal of Portfolio Management*, 42(5), 1-6.
- Asness, C. S., Frazzini, A., & Pedersen, L. H. (2013). *Quality Minus Junk*. Retrieved from SSRN: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2312432
- Bancel, F., & Mittoo, U. R. (2004, Winter). Cross-Country Determinants of Capital Structure Choice: A Survey of European Firms. *Financial Management*, 33(4), 103-132.
- Banz, R. W. (1981). The relationship between return and market value of common stocks. *Journal of Financial Economics*, 9(1), 3-18.
- Bender, J., & Nielsen, F. (2013). Earnings Quality Revisited. *Journal of Portfolio Management*, 39(4), 69-79.
- BEREC. (2017). *Regulatory Accounting in Practice 2017*. Retrieved from BEREC: http://berec.europa.eu/eng/document_register/subject_matter/berec/reports/7316-berec-report-regulatory-accounting-in-practice-2017
- Berk, J. B., & van Binsbergen, J. H. (2016). Assessing asset pricing models using revealed preference. *Journal of Financial Economics*, 119(1), 1-23.
- Black, F. (1989). *Equilibrium Exchange Rate Hedging*. National Bureau of Economic Research Working paper 2947.
- Blume, M. E. (1970). Portfolio theory: a step toward its practical application. *The Journal of Business*, 43(2), 152-173.
- Borghesi, E., Del Bo, C., & Florio, M. (2016). Institutions and Firms' Productivity: Evidence from Electricity Distribution in the EU. *Oxford Bulletin of Economics and Statistics*, 78(2), 170-196.
- Burns, P., & Weyman-Jones, T. (1996). *Regulatory review and populist pressure*. Economic research paper, Loughborough University of Technology, Department of Economics.
- Burns, P., Turvey, R., & Weyman-Jones, T. G. (n.d.). The Behaviour of the Firm under Alternative Regulatory Constraints. *Scottish Journal of Political Economy*, 45(2), 133-157.
- Caballero, R. J., Farhi, E., & Gourinchas, P.-O. (2017). The safe assets shortage conundrum. *Journal of Economic Perspectives*, 31(3), 29-46.

- Campbell, J. (1991). A Variance Decomposition for Stock Returns. *101*(405), 79-157.
- Campbell, J. Y., & Shiller, R. J. (1998). Valuation Ratios and the Long-Run Stock Market Outlook. *The Journal of Portfolio Management*, *24*(2), 11-26.
- Carhart, M. (1997). On Persistence in Mutual Fund Performance. *The Journal of Finance*, *53*(1), 57-82.
- Chin, M., & Polk, C. (2015). *A forecast evaluation of expected equity return measures*. Bank of England Working paper No. 520.
- Claus, J., & Thomas, J. (2001). Equity Premia as Low as Three Percent? Evidence from Analysts' Earnings Forecasts for Domestic and International Stock Markets. *The Journal of Finance*, *56*(5), 1629-1666.
- Cochrane, J. H. (2005). *Asset Pricing*. Princeton University Press.
- Cochrane, J. H. (2008). The Dog That Did Not Bark: A Defense of Return Predictability. *Review of Financial Studies*, *21*(4), 1533-1575.
- Cochrane, J. H. (2011). Presidential address: Discount rates. *The Journal of Finance*, *66*(4), 1047-1108.
- Commission, C. (2007). *A report on the economic regulation of the London airports companies (Heathrow Airport Ltd and Gatwick Airport Ltd)*. Competition Commission.
- Commission, C. (2010). *Bristol Water plc: A reference under section 12(3)(a) of the Water Industry Act 1991*. Competition Commission.
- Commission, C. (2014). *Northern Ireland Electricity Limited price determination A reference under Article 15 of the Electricity (Northern Ireland) Order 1992*. Competition Commission.
- Commission, N. Z. (2014). *Amendment to the WACC percentile for price-quality regulation for electricity lines services and gas pipeline services*.
- Constantinides, G. M., Donaldson, J. B., & Mehra, R. (2002). Junior can't borrow: A new perspective on the equity premium puzzle. *The Quarterly Journal of Economics*, *117*(1), 269-296.
- Damodaran, A. (2016, March 5). Equity Risk Premiums (ERP): Determinants, Estimation and Implications – The 2016 Edition.
- Dechow, P., Ge, W., & Schrand, C. (2010). Understanding earnings quality: A review of the proxies, their determinants and their consequences. *Journal of Accounting and Economics*, *50*(2-3), 344-401.
- Dimson, E., Marsh, P., & Staunton, M. (2017). Analysis of investment returns since 1900. *Credit Suisse Global Investment Returns Yearbook 2017*.
- Dison, W., & Rattan, A. (2017). An improved model for understanding equity prices. *Bank of England Quarterly Bulletin*, *57*(2), 86-97.
- Dobbs, I. (2011). Modeling welfare loss asymmetries arising from uncertainty in the regulatory cost of finance. *Journal of Regulatory Economics*, *39*(1), 1-28.
- Espey, M., Espey, J., & Shaw, W. D. (1997). Price elasticity of residential demand for water: A meta-analysis. *Water Resources Research*, *33*(6), 1369-1374.
- Fama, E. F. (1973). Risk, return, and equilibrium: Empirical tests. *Journal of Political Economy*, *81*(3), 607-636.
- Fama, E. F., & French, K. R. (1992). The Cross-Section of Expected Stock Returns. *The Journal of Finance*, *47*(2), 427-465.
- Fama, E. F., & French, K. R. (2002, April). The Equity Premium. *Journal of Finance*, *57*(2), 637-699.
- French, K., & Poterba, J. (1991). Investor Diversification and International Equity Markets. *American Economic Review*, *81*(2), 222-226.

- Gandolfi, M., Jenkinson, T., & Mayer, C. (1996). *Regulation and the cost of capital*. Working paper, University of Oxford, School of Management Studies.
- Gebhardt, W. R., Lee, C. M., & Swaminathan, B. (2001). Toward an Implied Cost of Capital. *Journal of Accounting Research*, 39(1), 135-176.
- Gordon, M. J. (1962). *The Investment, Financing, and Valuation of the Corporation*. Homewood, Ill.: Richard D. Irwin.
- Graham, B., & Dodd, D. (2008). *Security Analysis* (Sixth ed.). New York: The McGraw-Hill Companies.
- Graham, J., & Harvey, C. (2002). How Do CFOs Make Capital Budgeting and Capital Structure Decisions? *Journal of Applied Corporate Finance*, 15(1), 8-23.
- Gregory, A., Hua, S., & Tharyan, R. (March 2016). *In Search of Beta*. University of Exeter Business School, Xfi Centre.
- Guthrie, G. (2006). Regulating infrastructure: the impact of risk and investment. *Journal of Economic Literature*, XLIV, 925–972.
- Haugen, R. A., & Heins, A. J. (1972). *On the Evidence Supporting the Existence of Risk Premiums in the Capital Market*. Wisconsin Working Paper, December 1972. Retrieved from SSRN: <https://ssrn.com/abstract=1783797>
- Hsu, J., & Kalesnik, V. (2014). *Finding Smart Beta in the Factor Zoo*. Retrieved from Research Affiliates: https://www.researchaffiliates.com/en_us/publications/articles/223_finding_smart_beta_in_the_factor_zoo.html
- Hyland, M. (2016). Restructuring European electricity markets – A panel data analysis. *Utilities Policy*, 38(C), 33-42.
- Imrecon working with Economic Consulting Associates. (2012). *RIIO reviews*. Ofgem.
- Insights, E. (2014). *Regulatory Precedents for Setting the WACC within a Range*. Report prepared for New Zealand Commerce Commission.
- J.H., C. (1988). How big is the random walk in GNP? *Journal of Political Economy*, 96, 893-920.
- Jegadeesh, N., & Titman, S. (1993). Returns to Buying Winners and Selling Losers: Implications for Stock Market Efficiency. *The Journal of Finance*, 48(1), 65-91.
- Jorda, O., Knoll, K., Kuvshinov, D., Schularick, M., & Taylor, A. (2017). Rate of Return on Everything, 1870-2015. *Bank of England workshop on Growth, Stagnation and Inequality*.
- Korteweg, A. (2010). The Net Benefits to Leverage. *The Journal of Finance*, 65(6), 2137-2170.
- Li, Y., Ng, D. T., & Swaminathan, B. (2013). Predicting market returns using aggregate implied cost of capital. *Journal of Financial Economics*, 110(2), 419-436.
- Merton, R. (1973). An Intertemporal Capital Asset Pricing Model. *Econometrica*, 41(5), 867-887.
- Miller, M. H., & Modigliani, F. (1961). Dividend Policy, Growth, and the Valuation of Shares. *The Journal of Business*, 34(4), 411-433.
- Moody's. (2016, February 29). *Annual Default Study: Corporate Default and Recovery Rates, 1920-2015*. Retrieved from Moody's: www.moody.com
- Novy-Marx, R. (2013). The other side of value: The gross profitability premium. *Journal of Financial Economics*, 108(1), 1-28.
- Ofwat. (2017). *Delivering Water 2020: Our final methodology for the 2019 price review*. Retrieved from Ofwat: <https://www.ofwat.gov.uk/publication/delivering-water-2020-final-methodology-2019-price-review/>

- Parker, D. (2003). Performance, risk and strategy in privatised, regulated industries: The UK's experience. *International Journal of Public Sector Management*, 16(1), 75 – 100.
- Paul, A. C., Myers, E. C., & Palmer, K. L. (2009). *A partial adjustment model of US electricity demand by region, season, and sector. Discussion Paper*. Retrieved from Resources for the Future: <http://www.rff.org/files/sharepoint/WorkImages/Download/RFF-DP-08-50.pdf>
- Polemis, M. L. (2016). New evidence on the impact of structural reforms on electricity sector performance. *Energy Policy*, 92(C), 420-431.
- PWC. (2017). *Refining the balance of incentives for PR19. Report to Ofwat, June*. Retrieved from PWC: <https://064f1d25f5a6fb0868ac-0df48efcb31bcf2ed0366d316cab9ab8.ssl.cf3.rackcdn.com/wp-content/uploads/2017/07/PwC-Balance-of-incentives-June2017.pdf>
- Reinganum, M. R. (1981). Misspecification of capital asset pricing: Empirical anomalies based on earnings' yields and market values. *Journal of Financial Economics*, 9(1), 19-46.
- Roll, R. (1977). A critique of the asset pricing theory's tests Part I: On past and potential testability of the theory. *Journal of Financial Economics*, 4(2), 129-176.
- Rosenberg, B., Reid, K., & Lanstein, R. (1985). Persuasive evidence of market inefficiency. *The Journal of Portfolio Management*, 11(3), 9-16.
- Ross, S. A. (1977). The Capital Asset Pricing Model (CAPM), Short-Sale Restrictions and Related. *The Journal of Finance*, 32(1), 177-183.
- S&P Global. (2017, April 13). *Default, Transition, and Recovery: 2016 Annual Global Corporate Default Study And Rating Transitions*. Retrieved from S&P Global: <https://www.spglobal.com/our-insights/2016-Annual-Global-Corporate-Default-Study-and-Rating-Transitions.html>
- Shumway, T., & Warther, V. A. (1999). The Delisting Bias in CRSP's Nasdaq Data and Its Implications for the Size Effect. *The Journal of Finance*, 54(6), 2361-2379.
- Sloan, R. G. (1996). Do Stock Prices Fully Reflect Information in Accruals and Cash Flows about Future Earnings? *The Accounting Review*, 71(3), 289-315.
- Smithers, A., & Wright, S. (2013). *The Cost of Equity Capital for Regulated Companies: A Review for Ofgem*. Ofgem.
- Stattman, D. (1980). Book Values and Stock Returns. In *The Chicago MBA: A Journal of Selected Papers* (Vol. 4, pp. 25-45).
- Sudarsanam, S., Kaltenbronn, U., & Park, P. (2011). *Cost of Equity for Regulated Companies: An international Comparison of Regulatory Practices*. UK Competition Commission.
- van Binsbergen, J., Graham, J. R., & Yang, J. (2010). The Cost of Debt. *The Journal of Finance*, 65(6), 2089-2136.
- Vayanos, D., & Woolley, P. (2016). *Curse of the Benchmarks*. Retrieved from LSE: <http://www.lse.ac.uk/fmg/dp/discussionPapers/fmgdps/DP747CurseoftheBenchmarks.pdf>
- Wright, S. (2005). *Beta Estimates and Bond Spread Analysis*. Ofgem.
- Wright, S., Mason, R., & Miles, D. (2003). *A Study into Certain Aspects of the Cost of Capital for Regulated Utilities in the UK*. London: Smithers & Co Ltd report to the UK Regulators and the Office for Fair Trading.
- Wright, S., Mason, R., Baskaya, M., Hori, K., & Satchell, S. (2006). *Report on the Cost of Capital*. Ofgem.

Disclaimers

This document is based upon information available to us at the date of this document and takes no account of subsequent developments. In preparing this document we may have relied upon data supplied to us by third parties (including those that are the subject of due diligence) and therefore no warranty or guarantee of accuracy or completeness is provided. We cannot be held accountable for any error, omission or misrepresentation of any data provided to us by third parties (including those that are the subject of due diligence). This document is not intended by us to form a basis of any decision by any third party to do or omit to do anything.

Any opinions or assumptions in this document have been derived by us through a blend of economic theory, historical analysis and/or other sources. Any opinion or assumption may contain elements of subjective judgement and are not intended to imply, nor should be interpreted as conveying, any form of guarantee or assurance by us of any future performance. Views are derived from our research process and it should be noted in particular that we can not research legal, regulatory, administrative or accounting procedures and accordingly make no warranty and accept no responsibility for consequences arising from relying on this document in this regard.

Calculations may be derived from our proprietary models in use at that time. Models may be based on historical analysis of data and other methodologies and we may have incorporated their subjective judgement to complement such data as is available. It should be noted that models may change over time and they should not be relied upon to capture future uncertainty or events.