Systematic risk and the role and measurement of equity beta: A report to the AER Consumer Reference Group

Glenn Boyle
Director, Sapere Research Group
Christchurch, New Zealand

Kieran Murray
Managing Director, Sapere Research Group
Wellington, New Zealand

June 29, 2021

Executive Summary

I. The systematic risk of any asset reflects the relationship between returns on that asset and overall investor wellbeing, i.e., the extent to which it offers low returns in times when investor wellbeing is low. In a mean-variance (CAPM) world, this simplifies to the relationship between returns on that asset and returns on the market portfolio, and is known as beta (β).

II. This risk can be decomposed into two sub-betas: cash flow beta and discount rate beta. Traditional discussion of beta typically focuses only on the former but this is incomplete and potentially misleading.

III. Estimating beta for use in practical applications is challenging and subject to many ambiguities. We propose three criteria for evaluating different approaches to estimation: accuracy (in the sense of being able to predict actual beta during the relevant forecast window), simplicity (in the sense of being readily understood and easily replicable), and consistency (in the sense of being consistent with the underlying model and not invoking contradictory assumptions). In practice, most choices involve tradeoffs among these criteria.

IV. Applying these criteria to fundamental principles of financial economics and econometrics, and appealing to recent empirical evidence, our main conclusions are as follows:

- The choice between raw and excess returns is immaterial if daily (and possibly weekly) returns are used, but not otherwise;
- Daily (or possibly weekly) data are preferred to lower-frequency data;
- If data evidence suggests beta is a constant or fluctuates randomly around a constant, then the estimation period should be set as long as possible; if instead data evidence suggests beta is time-varying and mean-reverting, then

the choice is more complicated and depends on the underlying cause of the reversion;

- A domestic index proxy is preferred to a world index;
- The use of foreign comparators may occasionally be unavoidable, but doing so introduces a number of estimation problems;
- The so-called "low-beta bias" should not be accommodated by adjusting beta;
- Significant company-specific risks should not be accommodated by adjusting beta;
- Beta shrinkage may be useful if regulated firm betas are believed to be mean-reverting, but not otherwise.

V. Overall, the critical unknown is the time series behaviour of beta. If beta is a constant, or mean-reverts because of periodic investor mis-pricing, then the best approach would seem to involve unadjusted OLS estimates obtained using daily data from a long time period benchmarked against a domestic index with minimal recourse to foreign comparators. If instead beta varies through time in a manner that reflects true variation in risk, then a short estimation period is likely to be more reflective of relevant market conditions, and various adjustments, including shrinkage, may be necessary. In the absence of clear evidence for rational variation in beta through time, then acting as if beta is constant may well be a reasonable working assumption for regulators.

1 Introduction

1. In this report, we provide an overview of issues that arise in the estimation of beta, with particular reference to the Australian electricity and gas sectors that are supervised by the Australian Energy Regulator (AER). Our list of issues is unlikely to be exhaustive, but it is, hopefully, representative of the major themes that arise in practice. Our aim is to provide an independent perspective on beta estimation issues relevant to the AER and the Consumer Reference Group (CRG).

2. Our approach is to examine the identified issues from the perspective of basic principles of financial economics together with relevant empirical evidence that sheds light on the applicability of these principles. Although we attempt to make recommendations where it is possible to do so, our approach is not overly-prescriptive as, unfortunately, finance theory and evidence has not evolved to the point where this is possible. Instead, we focus on identifying the tradeoffs involved in attempting to resolve the various issues.

3. In the next section, we describe the basic theoretical principles underlying beta: what it is, its interpretation, why it matters, its determinants, its place in a more general framework of asset pricing, and how it relates to regulation. Section 3 turns to the empirical estimation of beta: how this is achieved, what issues arise, and the criteria that can be used to evaluate the possible approaches to resolving these issues. Section 4 then discusses each issue in turn and attempts to evaluate the possible approaches by applying finance theory and evidence to the previously-identified criteria. Section 5 offers some concluding remarks.

2 Background theory and intuition

2.1 Australian regulatory context

4. The starting point for our discussion of systematic risk and the role and measurement of equity beta is the national electricity and gas objectives; these objectives guide the Australian Energy Regulator (AER) in regulating the revenue of electricity and gas network service providers. The objectives are:

   - NEO (National Electricity (South Australia) Act, 1996, s 7): to promote efficient investment in, and efficient operation and use of, electricity services for the long term interests of consumers of electricity with respect to- (a) price, quality, safety, reliability and security of supply of electricity; and (b) the reliability, safety and security of the national electricity system.
   - NGO (National Gas (South Australia) Act, 2008, s. 23): to promote efficient investment in, and efficient operation and use of, natural gas services for the long term interests of consumers of natural gas with respect to price, quality, safety, reliability and security of supply of natural gas.

5. It would seem that although there may be debate about how some aspects of these objectives should be interpreted, there is general agreement that these objectives provide for the promotion of efficient investment in the long-term interests of consumers. The Australian Competition Tribunal has summarised the economic foundation of the national electricity objective and the revenue and pricing principles as follows:[1]

[1] Australian Competition Tribunal, *ElectraNet Pty Limited* (No 3), [2008], paragraph 15, cited by the Federal Court of Australia (above) as principles "all parties appear to embrace", Federal Court of Australia, *Australian Energy Regulator v Australian Competition Tribunal* (No 2) [2017], paragraph 496.

Consumers will benefit in the long run if resources are used efficiently, i.e. resources are allocated to the delivery of goods and services in accordance with consumer preferences at least cost. As reflected in the revenue and pricing principles, this in turn requires prices to reflect the long run cost of supply and to support efficient investment, providing investors with a return which covers the opportunity cost of capital required to deliver the services.

6. A regulatory objective of providing investors with a return that covers the opportunity cost of capital required to deliver an efficient level of service in turn requires consideration of risk and the pricing of that risk.

2.2 General principles of risk and asset pricing

7. Risk is a fairly straightforward concept to grasp, but one that is difficult to precisely define and measure. Intuitively, it would seem to refer to the potential for the loss of something that one values. This could be, for example, one's income or consumption or employment status or wealth. An asset that can potentially reduce the value of any of these by a large amount is then more risky than another asset that leaves them relatively unaffected.

8. Modern asset pricing theory (see Cochrane, 2005, pp15-17) makes this idea concrete by introducing the concept of a pricing kernel (also known as a stochastic discount factor). The kernel y represents the value of an additional dollar of returns. In prosperous states of the world (i.e., "good" times), investors have high wealth and consumption and so an additional dollar of returns has little value, i.e., y is low; in impoverished states of the world (i.e., "bad" times) such as recessions, investors have low wealth and consumption and so an additional dollar of returns has high value, i.e., y is high. The kernel can be thought of as capturing

the value of investors' economic wellbeing or "overall portfolio", i.e., the index of what investors care about.

9. Now consider a single asset i with uncertain future returns R_i. The randomness or uncertainty associated with these future returns can be split into two components: the part that is correlated with the pricing kernel y and the remainder that is independent of y. The former is known as *systematic* risk because of its systematic relationship to the measure that investors care about (y); the latter is known as *idiosyncratic* risk.

10. The distinction between systematic and idiosyncratic risk is important. By definition, only the former is related to, and can therefore affect, the measure that investors care about (y). For example, an asset that yields high returns when y is low and low returns when y is high is one that does well when additional returns are of little value to the investor and does poorly when returns have high value. Such an asset is, in effect, magnifying, or adding to, the risk of the investor's overall portfolio. By contrast, idiosyncratic risk that is uncorrelated with y obviously has no implications for that portfolio.

11. Modern finance theory typically assumes investors are risk averse. That is, they require compensation, in the form of a higher expected return, for bearing risk. But not all risks are created equal. The systematic risk of an individual asset requires an expected return premium because it contributes to the loss of something investors care about (y). The idiosyncratic risk of an individual asset has no consistent effect on y and so cannot command an expected return premium.

12. To summarise, investors care about their overall wellbeing, captured by a pricing kernel y. An individual asset is risky, and merits an expected return premium, to the extent that its returns are correlated with y, which is known as systematic risk. The component of returns that are uncorrelated with y do not affect investors'

overall wellbeing and so do not warrant an expected return premium.

13. The theoretical expected return premium for systematic risk is (see Cochrane, 2005, p14):

$$E[R_i] = R_f\{1 - cov(R_i, y)\} \tag{1}$$

where R_i is the 1-period return on risky asset i and R_f is the 1-period return on the riskless asset. Equation (1) confirms the view, articulated above, that asset risk depends on the correlation between its returns and good/bad times: an asset whose returns are high in good times and low in bad times ($cov(R_i, y) < 0$) adds to the risk of investors' overall wellbeing and so commands a higher expected return premium ($E[R_i]/R_f$) than an asset that hedges overall wellbeing by providing returns that are low in good times and high in bad times ($cov(R_i, y) > 0$). Any volatility in returns that is uncorrelated with overall wellbeing ($cov(R_i, y) = 0$) commands no risk premium.

14. To a first approximation, essential services such as electricity and gas would seem likely to have a low correlation with y and hence a low expected return premium. In bad times (e.g., redundancy, divorce, dog run over, and so on), households have little flexibility to respond by eating or heating less, so the demand for these services is largely independent of y ($cov(R_i, y) \approx 0$) and investors in these services thus bear little systematic risk.[2]

15. However, attempts to apply (1) in practice encounter problems. First, y is not uniquely identified, and there are obviously many candidates. Second, the more plausible candidates (such as aggregate consumption) are not very accurately measured. This has led to reliance on a simplified version of (1) that largely avoids these problems.

[2] Although, as we discuss in section 2.4, this simple story is potentially incomplete. Note also that "$\approx$" denotes approximately equal to.

2.3 The CAPM

16. Suppose that all investors (i) rank portfolios according to single-period means and variances of returns, (ii) perceive the same set of ex ante mean-variance opportunities, and (iii) consume all their wealth at the end of the period. In such a world, all investors hold the market portfolio and all they care about is the return on that portfolio. It is straightforward to show that in these circumstances:

$$y = a - b(R_m - R_f) \tag{2}$$

where R_m is the 1-period return on the market portfolio, and a and b are positive constants. That is, because investors care only about the market portfolio return, a high $R_m - R_f$ is good times (low y) and a low $R_m - R_f$ is bad times (high y).

17. Substituting (2) into (1) yields the the so-called Capital Asset Pricing Model (Sharpe, 1964; Lintner, 1965), or CAPM:

$$E[R_i] = R_f + \beta_i\{E[R_m] - R_f\} \tag{3}$$

where:

$$\beta_i = \frac{cov(R_i - R_f, R_m - R_f)}{var(R_m - R_f)} \tag{4}$$

is the "beta" of asset i (we use the terms β and beta interchangeably throughout this report). In words, the CAPM states that equilibrium asset prices are set so that the expected return on each asset i equals the riskless rate plus a risk premium equal to the product of asset i's beta (β_i) and the market portfolio risk premium ($E[R_m] - R_f$). In the CAPM, systematic risk arises from a correlation between an asset's returns and market portfolio returns.

18. This simplification allows an additional interpretation of systematic risk as non-diversifiable, risk. Oxera (2015, p2) provide a nice explanation of this idea:

> (A)n ice-cream stall typically performs well on sunny days, and less well on rainy days. An umbrella stall does well on rainy days, but not on sunny days. By investing in both, an investor can diversify some of their total risk — but not all of it, because if there is a massive economic downturn people will buy both less ice cream and fewer umbrellas, regardless of the weather. Because this systematic risk cannot be eliminated, investors need to be compensated for it — with the scale of the required compensation being proportionate to the beta.

19. Intuitively, investors will not pay a premium to avoid idiosyncratic (i.e. firm-specific) risk as it can be eliminated via do-it-yourself diversification. Portfolio diversification allows investors to eliminate all risk except the risk of the market as a whole – a systematic risk. Thus, only the latter kind of risk can command compensation via an expected return premium. The contribution of the CAPM is that it quantifies this insight. If, for example, $\beta_i = 1/3$, then asset i has 1/3 as much systematic risk as the market portfolio and hence commands 1/3 of the risk premium.

20. Why does β_i work as a quantitative measure of the systematic risk of asset i? Investors care only about the risk of their *portfolio*. The answer is that the risk of an individual asset in isolation doesn't matter because it needn't be held in isolation; instead, the risk of any asset matters only to the extent it contributes to the risk of the portfolio. Consider an asset whose returns are positively correlated with those of the portfolio – it provides high returns when the rest of the portfolio is also doing well and low returns when the rest of the portfolio is also doing badly. Such an asset "magnifies", or exacerbates, the risk of the portfolio. By contrast, an asset whose returns are negatively correlated with those of the portfolio provides high returns when the rest of the portfolio is doing badly and low returns when the rest of the portfolio is doing well, i.e., it "hedges" the risk of the portfolio.

The first asset is clearly riskier in a portfolio context since it adds to the risk of the portfolio, whereas the second asset decreases portfolio risk. It follows that the risk of any asset i depends on the correlation of its returns with the returns on the portfolio in which it is held. If all investors hold the market portfolio (which they do in a CAPM world), then the risk of any asset i depends on the correlation of its returns with the returns on the market portfolio, i.e., β_i.

2.4 Underlying sources of beta

21. What determines β? It is common to think of β as arising from the correlation between a company's cash flows and market cash flows. That is, companies and projects that generate high cash flows when most other companies are doing the same have high betas while companies and projects with cash flows that fluctuate independently of aggregate economy cash flows have low betas (see, for example, McKenzie and Partington, 2012, pp 5-7). Cornell (1999, p184) nicely articulates this view:

> (W)here does systematic risk come from? The common answer to that question is that it arises from the correlation between a company's cash flows and market cash flows. For instance, it is often assumed that if a company's "investment projects" consist of a series of cash flow gambles that are uncorrelated with any fundamental economic or market factors, then the systematic risk of its equity should be zero and its stock should be priced accordingly. As Brealey and Myers (1996) so artfully describe in their best-selling textbook, "Lone prospectors in search of gold look forward to extremely uncertain future earnings, but whether they strike it rich is not likely to depend on the performance of the market portfolio. Even if they do not find gold, they do not bear much market risk.

Therefore, an investment in gold has a high standard deviation but a relatively low beta."

22. Yet, as Cornell goes on to explain, this logic is potentially misleading. Betas reflect common variation in returns, but an unexpectedly high return can arise either because cash flow expectations have been revised upwards *or* because future return expectations (i.e., discount rates) have been revised downwards. Using the Campbell and Shiller (1988) decomposition to apply this insight to β, Campbell and Mei (1993) show that β actually has two components, or sub-betas:[3]

$$\beta_i = \beta_{ci} - \beta_{ki} \tag{5}$$

where β_{ci} is the component of beta arising from shocks to future expected cash flows (the cash flow beta) and β_{ki} is the component of beta arising from shocks to future discount rates (the discount rate beta). Equation (5) makes it clear that β is driven not only by common variation in cashflows (β_{ci}), but also by common variation in discount rates (β_{ki}).

23. Because discount rates tend to rise and fall together, β_{ki} will typically be *negative.* That is, when market discount rates rise (and hence market returns are low), firm i's discount rate will tend to rise as well. So the firm i discount rate will be high when market returns are low, i.e., $\beta_{ki} < 0$. Thus, a greater sensitivity of firm i discount rates to market returns *increases* β_i. To illustrate, suppose $\beta_{ci} = 0.2$ and $\beta_{ki} = -0.1$. Then $\beta_i = 0.2 - -0.1 = 0.3$. If $\beta_{ki} = -0.3$, then $\beta_i = 0.2 - -0.3 = 0.5$.

24. Equation (5) is potentially important for regulators because it reveals the mechanism by which regulation can affect β and the limitations on this influence. Using price caps or revenue caps or rate of return restrictions, regulators seek to control the cash flows of regulated entities, which potentially affects β_{ci} but is unlikely to have any effect on β_{ki}. Intuitively, regulation can affect market views about

[3] This treatment is a simplification of Campbell and Mei (1993).

future expected cash flows, but is unlikely to influence future market sentiment about the pricing of those cash flows, since the latter depends on factors such as investor risk aversion that are largely immune to regulator rulings. Thus, even if regulation resulted in the cash flows becoming riskless and hence $\beta_{ci} = 0$, the overall β_i could still be substantially positive if β_{ki} is sufficiently negative.

25. This point is often overlooked in discussions of regulatory β. For example, McKenzie and Partington (2012, p15) claim,

> Taken together, the previous conceptual discussion clearly provides evidence to suggest that the theoretical beta of the benchmark firm is very low. While it is difficult to provide a point estimate of beta, based on these considerations, *it is hard to think of an industry that is more insulated from the business cycle due to inelastic demand and a fixed component to their pricing structure. In this case, one would expect the beta to be among the lowest possible* and this conclusion would apply equally irrespective as to whether the benchmark firm is a regulated energy network or a regulated gas transmission pipeline. (emphasis added)

26. The references to "business cycle", "inelastic demand", and "pricing structure" make it clear that McKenzie and Partington are implicitly focusing on β_c, not β. But for regulated electricity and gas firms to have "lowest possible" betas (which we assume means close to zero), it is necessary that *both* β_{ci} and β_{ki} be very small (in absolute value), not just β_{ci}. Discussions that focus solely on β_{ci}, whether arguing for a high or low value, are incomplete.

27. In paragraph 14, we suggested that the demand for essential services such as electricity and gas is likely to have low systematic risk and so such assets would command a small expected return premium. However, we also cautioned in footnote

2 that such an argument is potentially incomplete. The reason for this caution is now obvious: our initial argument about demand for electricity and gas services corresponds, in a CAPM world, to recognising the cash flow beta while ignoring the discount rate beta.

28. In the context of electricity and gas regulation, the discount rate beta cannot be easily dismissed. Regulated electricity and gas networks are typically made up of long-lived assets with stable cash flows, which as Cornell (1999) points out are exactly the kind of assets that have high discount rate sensitivity and hence high β_k (in absolute value). Intuitively, this occurs because long-lived asset prices are particularly sensitive to discount rate changes and so their returns covary strongly with market movements.

29. To summarise, the beta decomposition in equation (5) has three principal lessons. First, relative to the average firm (which has a total-beta equal to 1), regulation insulates company cash flows from market returns and so induces a relatively small cash flow beta. Second, however, because regulation is unlikely to affect overall financial market sentiment, the regulated company's discount rate beta can be bigger or smaller than the discount rate beta of the average firm. Thus, the total-beta difference between the regulated firm and the average firm is, at a theoretical level, indeterminate.[4] Third, high-level theoretical discussions of regulatory firm beta invariably focus solely on the cash flow beta and so should be viewed with some caution.

2.5 β Review

30. β is a highly intuitive, theoretical risk measure that follows from some fairly restrictive assumptions. Of course, as Friedman (1953) long ago pointed out, the

[4] In Appendix A, we lay out this point more rigorously.

validity of any theory is determined not by the realism of its assumptions but by the accuracy of its predictions. There is strong and ongoing debate about the empirical efficacy of β (see, for example, Fama and French, 2004; Brown and Walter, 2013), but it continues to be widely used by practitioners and regulators (Jacobs and Shivdasani, 2012). That being the case, it is important to understand the issues involved in estimating β for use in applied situations and, to the greatest extent possible, identify the "best" β.

31. In the next section, we set out some thoughts on obtaining this "best" β. First, we describe the fundamentals of beta estimation. Second, we identify a number of complications, or issues that need to be resolved, in applying the fundamental approach. Third, we suggest some criteria that can be used to help resolve these issues.

3 Estimating β

3.1 Fundamentals

32. Estimation of the β defined in equation (4) is obtained by applying ordinary least squares (OLS) regression to the linear equation:

$$R_{it} - R_{ft} = \alpha + \beta_i\{R_{mt} - R_{ft}\} + \epsilon_{it}$$

where t denotes time, and ϵ is an error term. This regression yields (ignoring time subscripts):

$$\hat{\beta}_i = \frac{c\hat{o}v(R_i - R_f, R_m - R_f)}{v\hat{a}r(R_m - R_f)} \tag{6}$$

where a ˆ denotes an estimate.

33. Alternatively, noting that (4) can be rewritten as:

$$E[R_i] = R_f\{1 - \beta_i\} + \beta_i\{E[R_m]\},$$

a common approach is to apply OLS to raw returns (rather than excess returns):

$$R_{it} = \alpha + \beta_i R_{mt} + \epsilon_{it}$$

which yields:

$$\hat{\beta}_i = \frac{c\hat{o}v(R_i, R_m)}{v\hat{a}r(R_m)} \tag{7}$$

34. In the CAPM, the riskless rate R_f is non-stochastic and so equations (6) and (7) are equivalent. But in empirical applications, the riskless rate, and its empirical proxies, are stochastic and so the betas can differ (Campbell et al., 1997).

35. Regardless of whether (6) or (7) is used, the β that regulators and practitioners want to identify is that prevailing over some *future* time period. Unfortunately, future realisations of returns are unknown and so β has to be estimated from historical data. Moreover, the CAPM is a single-period model where the single period is of arbitrary duration. These complications mean that the estimation of β is far from being perfectly prescribed and so requires a number of choices to be made, each of which can have a potentially significant effect on the estimate.

3.2 Choices

36. The choices that potentially have to be made, or issues that have to be resolved, in attempting to apply equation (6) (or (7)) include:

- Excess or raw returns
- Length of estimation period
- Frequency of data during estimation period
- Length of forecast period
- Domestic or world proxy for R_m
- Domestic or foreign comparators for firm i (if necessary)

- "Low-beta bias"
- Non-synchronous regulatory and beta periods
- Adjustment for specific risks
- Statistical issues
- Indirect beta estimation
- Beta endogeneity

37. Making such choices first requires identification of a set of criteria for determining which option is superior, a task we turn to next.

3.3 Criteria

3.3.1 Accuracy

38. When faced with a large number of choices in estimating β, arguably the most important criterion is that the choices settled on provide the most "accurate" estimate, in the sense of being closest to the subsequently realised β.[5] A number of recent research papers have examined this issue in detail, e.g., Hollstein (2020), Hollstein et al. (2019) and Levi and Welch (2017).

39. To evaluate the accuracy of different β estimators, Hollstein (2020) and Hollstein et al. (2019) calculate the Root Mean Squared Error (RMSE) of the estimator:

$$RMSE_i = \sqrt{\frac{1}{N}\sum_{i=1}^{N}(\beta^a_{it} - \hat{\beta}_{it})^2}$$

where $\hat{\beta}_{it}$ is an estimate of beta for the period from t to $t+h$, $\beta^a_{it} = \frac{\sum_{\tau=t+1}^{t+h}, R_{i\tau}R_{m\tau}}{\sum_{\tau=t+1}^{t+h} R^2_{m\tau}}$ is the realised, or actual, beta for the period t to $t+h$, and N is the number of out-of-sample observations of realised and estimated betas.

[5] This objective raises some potentially complicated issues in a regulatory setting, which we address in paragraphs 42–44 below.

40. Intuitively, the RMSE can be thought of as measuring the typical magnitude of a beta estimate's prediction error ($\beta_{it}^a - \hat{\beta}_{it}$). To illustrate, suppose there are only two observations, one of which has a prediction error of 0.3 and the other -0.3. Then the average prediction error is zero, but the RMSE is 0.3, reflecting the magnitude of the typical prediction error.

41. Levi and Welch (2017) adopt a slightly different approach. They estimate the regression equation

$$\beta_{it}^a = \gamma_0 + \gamma_1 \hat{\beta}_{it} + \mu_{it} \tag{8}$$

and use the regression coefficient of determination (R^2) to determine the ability of the beta estimate ($\hat{\beta}_{it}$) to predict the realised beta (β_{it}^a). Note that the standard use of $\hat{\beta}_{it}$ as a direct prediction of the true beta is equivalent to imposing the restriction that $\gamma_0 = 0$ and $\gamma_1 = 1$.

42. Both these methods for assessing the accuracy of beta estimates is that they assume actual, or realised, beta to be the appropriate benchmark. This may not always be appropriate in a regulatory setting, for two reasons. First, the objective underlying the standard regulatory approach is to estimate the beta perceived by the market, which need not be identical to the actual beta. However, the usual approach to financial economics, and certainly the branch of financial economics appealed to by regulators, assumes that the market of investors behaves as if it has rational expectations, i.e., it perceives the true opportunity set of future returns. In this case, the "market" beta converges to the actual beta. So any divergence between the two is not strictly admissible in the framework considered by regulators. Moreover, even if the market does have biased expectations, Stein (1996, p437) shows that it is optimal for *long-run* investment decisions to be made on the basis of actual beta., i.e., the appropriate benchmark for regulators is the realised beta.

43. Second, regulators typically, and certainly AER, set as their goal the estimation of the cost of capital (and hence beta) of the hypothetical "benchmark efficient firm", rather than the beta of any individual firm (or groups of firms). Thus, a particular beta estimation approach may yield accurate predictions of a regulated firm's beta, but inaccurate predictions of the benchmark efficient beta if the regulated firm operates in an imperfectly efficient manner. Unfortunately, the benchmark efficient firm, and hence its beta, is unobservable, so the likelihood of such divergence must be considered from first principles.

44. Suppose an inefficient firm j has time t stock returns:

$$R_{jt} = R^*_t - \delta_{jt}$$

where R^*_t is the time t return on the equivalent benchmark efficient firm and δ_{jt} is a positive-valued random variable representing the time t costs (to shareholders) of firm j's inefficiencies. Then firm j's beta is (using raw returns and ignoring time subscripts for brevity):

$$\begin{aligned}
\beta_j &= \frac{cov(R_j, R_m)}{var(R_m)} \\
&= \frac{cov(R^* - \delta_j, R_m)}{var(R_m)} \\
&= \beta^* - \frac{cov(\delta_j, R_m)}{var(R_m)}
\end{aligned}$$

where β^* is the beta of the benchmark efficient firm. Thus, the extent to which the latter differs from the firm j beta depends on the extent to which firm j inefficiency costs co-vary with market portfolio returns. In the absence of any compelling reason to suspect a systematic relationship, the default option is $cov(\delta_j, R_m) = 0$, and hence $\beta_j = \beta^*$, i.e., so long as inefficiency costs are independent of market portfolio returns, then the own-firm beta is the same as the benchmark beta. Inefficiencies may drive other important wedges between inefficient and hypothetical-efficient firms, but the beta is unaffected.

45. All else equal then, in a horse race between competing approaches for estimating beta, the approach that performs best according to the above two criteria (i.e., is the best predictor of realised beta) is preferred.

3.3.2 Simplicity

46. In general, unless there are large differences in accuracy, simple beta estimation approaches should be preferred to more complex approaches. By "simple", we mean an approach that can be easily followed, understood and if necessary replicated by someone with only a basic knowledge of finance and statistics. In our opinion, this currently rules out more specialised statistical techniques based on Generalized AutoRegressive Conditional Heteroskedasticity (GARCH) and Long Memory models for regulatory purposes.

3.3.3 Consistency

47. In general, unless there are large differences in accuracy, consistent approaches to beta estimation should be preferred to inconsistent approaches. That is, the chosen theoretical framework should be applied consistently to data that are relevant for the purpose. In particular, the chosen approach should not contain mutually inconsistent or contradictory assumptions, or assumptions that violate the CAPM (although the need to shoehorn the single-period CAPM into a multi-period world can make it difficult to be too pure about this).

4 Evaluating the choices available to regulators when estimating β

In this section, we work our way through the issues listed in section 3.2, first describing each choice and any arising complications, followed by an evaluation, according to the section 3.3 criteria, of the various options available.

4.1 Raw returns or excess returns

48. The first choice that a beta estimator has to make is whether to proceed using raw returns (R_i) or excess returns ($R_i - R_f$), i.e., should β be estimated using:

$$\hat{\beta}_i = \frac{c\hat{o}v(R_i, R_m)}{v\hat{a}r(R_m)}$$

or

$$\hat{\beta}_i = \frac{c\hat{o}v(R_i - R_f, R_m - R_f)}{v\hat{a}r(R_m - R_f)}$$

49. In the CAPM, the riskless rate of interest is a constant and the two specifications are theoretically identical. In practice though, the riskless rate proxy will vary over time and the two specifications will yield different estimates (Campbell et al., 1997).

50. This becomes less of a problem when using high-frequency data (e.g., daily) as then the riskless return can be reasonably assumed to approximate zero at all dates and the two specifications are essentially identical. If monthly or annual data are used, however, this "solution" is not available and a choice has to be explicitly made.

51. Precisely because the riskless rate is assumed to be non-stochastic, (6) collapses to (7) in a CAPM world, so consistency with the CAPM would suggest a preference for raw returns. In practice, both appear to be widely used, with academic

researchers possibly favouring excess returns (Campbell et al., 1997, p182) and practitioners and regulators more often settling for raw returns.

52. The simplicity criterion would (slightly) favour the use of raw returns as this eliminates the need for riskless return data, which are not always readily available. We know of no work comparing the accuracy of the two specifications, although this could be readily undertaken using the RMSE approach above. However, this would be superfluous if the optimal data frequency is daily (or even weekly).

53. Overall, although raw returns potentially yield different beta estimates to excess returns, any such difference is neutralised by employing daily or weekly returns. In that case, the simpler approach based on raw returns can be confidently used.

4.1.1 International practice

54. AER appears to use raw returns (without explanation), although their original advisor (Henry, 2009) also refers to excess returns. Nevertheless, given that they also favour weekly returns, any divergence is probably small. Anthony et al. (2020) report that regulators in the Netherlands, the UK and the US also employ daily or weekly data in estimating beta, so the issue seems likely to be of minor importance in these countries as well.

4.2 Estimation period: length and data frequency

55. The second choice faced by a beta estimator is, together, the optimal length of the estimation period (e.g., 1 year, 3 years, 5 years) and the frequency of the data to use in that estimation period (e.g., daily, weekly, monthly). In practical terms, the two are interdependent: a decision to use a short estimation period necessitates higher-frequency data, while opting for low-frequency data forces the use of a longer estimation period.

56. High-frequency (daily and possibly weekly) data:

    - offer more observations and hence more precision in estimation (Morse, 1984);
    - eliminate the need to choose between raw and excess returns (as discussed in section 4.1);
    - eliminate the need to choose the return interval (e.g., month-end versus mid-month returns)

    but

    - contain more "noise";
    - may generate stock return observations that are correlated over time (Lehmann, 1990), causing problems for standard estimation methods;
    - tend to under-estimate beta for infrequently-traded securities (Grinblatt and Titman, 2002, p157).

57. Similarly, a long estimation period offers more observations and hence more precise estimates, but also risks introducing irrelevant information into the estimation process. As Levi and Welch (2020, p434) put it:

    > In the absence of time variation in the underlying data-generation process, betas should be estimated with as much historical return data as possible. In the presence of time variation in the underlying betas, the best historical estimation interval...becomes a trade-off between the desire to have more days (to reduce the estimation noise) and the desire to predict the future beta with more relevant recent data.

58. The implication is that if returns have a jointly-stationary distribution, then beta is constant and it should be estimated using as much data as possible, since this draws on the greatest amount of information about the (unknown) distribution. On the other hand, if beta varies through time, then shorter estimation periods

may be preferred because data from the more distant past can provide misleading information about the current beta.

59. Several recent studies have examined the accuracy of beta estimates (see section 3.3.1) using different combinations of data frequency and estimation period. First, using US stock return data between 1927 and 2014, Levi and Welch (2017) conclude that for a 1-year forecast window (i.e., 1-year ahead betas): (i) betas estimated from daily data provide better forecasts than those obtained from monthly data, and (ii) a 1-year estimation period yields more accurate forecasts than those obtained from 3- and 5-year periods, although the improvement over 3-years is small.

60. Second, using 1963-2015 US data, Hollstein et al. (2019) arrive at similar conclusions: for a 6-month forecast window, daily data work better than monthly or quarterly data and a 12-month estimation period outperforms both shorter (down to 1 month) and longer (up to 5 years) alternatives.

61. Third, Hollstein (2020) conducts a similar exercise using post-1980 data from 48 countries, including Australia. For both 6- and 12-month forecast windows, he finds that daily return data and a 24-month estimation period is the best combination in Australia.

62. Taken together, these three studies suggest a strong preference for estimating beta via high frequency data (daily) and a relatively short estimation period (12-24 months). However, some caution is advisable in attempting to directly apply these results to the regulated Australian energy sector. First, their accuracy criterion — predicting realised betas over relatively short forecast windows — may, as we explain in the next two paragraphs, differ from that of regulators. Second, the recommendation of a relatively short estimation period suggests the presence of slow, but real, time variation in beta, which may not be applicable to regulated

firms.

63. In the CAPM, beta is a parameter drawn from the joint distribution of single-period future returns, where the length of the single period is unspecified. In applications, the appropriate single period length most naturally corresponds to the length of time over which the associated cost of capital is to be used. However, the studies of Levi and Welch (2017), Hollstein et al. (2018) and Hollstein (2020) only consider 6- and 12-month forecast windows, which are typically shorter than the period of interest to regulators. Thus, the forecasting objective they consider is different to the horizon of Australian regulators (who use a 5-year regulatory period). Different results may emerge from a focus on longer forecast windows.

64. To illustrate, suppose the most accurate predictor of 6- or 12-month ahead realised betas is obtained using a year's worth of daily data, and is estimated to be 0.5. If the regulatory period is 1-year, then that beta should clearly be used. But suppose the regulatory period is actually 4 years — how much faith can be placed in 0.5 as an accurate predictor of the relevant beta over this much longer time period?[6]

65. Of course, if beta is constant, the issue is moot — the beta of 0.5 estimated for the next 12 months must also apply to the next 4 years. However, time variation in beta may cause the two to diverge. For the typical US firm, Levi and Welch (2017) find that the beta estimate obtained from 1 year of daily data is a much less accurate predictor of 1-year ahead realised beta in 5 years, suggesting mean

[6]One might wish to obtain forecasts of beta over a 4-year period by using 4-year returns data, but as Robertson (2018, p37) points out such an approach is impractical: "Notice that the strategy of tailoring the estimation window to the desired forecast window (so for example if one is interested in the one month ahead beta one could estimate by OLS on monthly data and this would provide a measure of the (average) beta over a one month interval) . . . becomes infeasible if the forecast horizon goes much beyond a quarter, there just isn't a sufficient run of returns data to estimate accurately such models."

reversion in beta. However, the same estimate is, after suitable adjustment for mean-reversion (we address this point in more detail in section 4.7.1), almost as accurate a predictor of the 5-year ahead realised beta. Whether prediction of the latter could be improved further by using more years of data is not addressed.

66. An important factor to be considered in the presence of mean-reverting beta is the cause of the reversion. If this reflects true variation in systematic risk, then the regulator's objective must be to estimate the currently prevailing beta as it is this that is indicative of current risk conditions.[7] That is, if systematic risk is currently below its mean, and lower than it has been in the past, then the appropriate beta is also low. This implies the use of a short estimation period in order to best capture the current low-risk conditions, e.g., 1 year (since even 3 years may be too long if the rate of reversion is not too slow) of daily data.

67. On the other hand, it is possible that mean-reversion in beta merely reflects under- and over-pricing. In this case, a regulator may wish to "look through" short-to-medium run variation in beta, and instead estimate the unconditional mean, or "long-run", beta, as in Stein (1996). This implies the use of a long estimation period in order to have the best chance of evening out the irrationally high and low values and obtaining instead an estimate of the underlying fundamental beta.

68. To summarise, the accuracy criterion for beta estimation seems to be best served by the use of high-frequency daily or weekly data. Also, the use of daily or weekly data avoids the need to decide whether, for example, monthly returns should be measured on a mid-month or end-month basis, and thus better satisfies the simplicity criterion.[8] The situation with respect to estimation period is less clear. The

[7] Blume (1975) suggests rational mean-reversion in beta could arise because firms tend to adopt projects with less extreme risk characteristics as they age. This reasoning seems unlikely to apply to regulated energy networks.

[8] By making riskless returns effectively zero, it also eliminates the need to decide on the relevant

studies of Levi and Welch (2017), Hollstein et al. (2018) and Hollstein (2020) all recommend a relatively short 12-24 month window, but this is based on evidence from short forecast windows; whether future betas corresponding to the regulatory period are better forecast with longer estimation periods is not considered. Whether a long or short estimation period is best ultimately comes down to, first, whether beta is constant or mean-reverting and, second, the underlying cause of any mean reversion.

4.2.1 Is beta constant or time-varying?

69. How likely is time variation in the beta of a regulated firm? From equation (5), a firm's beta can change either if its cash flow beta changes or its discount rate beta changes; the former could occur because a firm alters its business or financial risk; the latter could occur because a firm adjusts the duration of its assets.[9] While either case is possible for the "average" firm in the three studies cited above, both seem *a priori* less likely for regulated firms. Consistent with this view, Wright and Mason (2020) estimate rolling betas for two UK regulated utilities during the 1998-2019 period and argue that these show little evidence of time variation.

70. Of course, resolution of this issue for regulated Australian energy firms is ultimately an empirical matter that is beyond the scope of this report, but we can offer some suggestions for how the question might be addressed.[10] The simplest

riskless interest rate — see para 50. If excess returns *are* used to estimate β, then the riskless interest rate should be the same as that used elsewhere in estimating the cost of equity, i.e., in the R_f and $E[R_m] - R_f$ components of the CAPM.

[9] Of course, over a long enough time period, the composition of the market portfolio may change sufficiently to induce a change in a firm's beta even when the firm's own activities, structure and organisation remain unaltered. However, this seems likely to occur only very slowly for a diverse index such as the ASX200.

[10] Determining the time series properties of a variable is a complex statistical issue, so our comments should be viewed as purely suggestive rather than representing a definitive view on the matter.

approach is that adopted by Wright and Mason (2020) who apply an eyeball test to time series plots of regulated entity beta estimates. Because the plots appear to fluctuate randomly around a constant (at least since 2005), this is supportive of beta stability.

71. A slightly more sophisticated approach involves adding a time variable function to a regression of estimated beta on a constant. For example, one could estimate for each regulated firm i:

$$\hat{\beta}_{it} = a + bt + ct^2 + \mu_{it}$$

where $\hat{\beta}_{it}$ is the estimate of firm i's beta at time t and μ is an error term. The null hypothesis (that beta is a constant) corresponds to $b = c = 0$. If the latter can be rejected, then a constant beta is called into question.

72. If panel data on regulated firms are available (which we understand could be a stretch for Australia), then time fixed effects could be included. That is, one could estimate:

$$\hat{\beta}_{it} = a + \lambda_t + \mu_{it}$$

where λ_t is a time dummy. An F-test that all time fixed effects equal zero could then be applied.

73. Other more complex approaches are also possible. For example, discrete changes in beta could potentially be identified using structural break tests.[11] Similarly, formal tests for mean reversion in beta could be undertaken via, for example, the approach developed by Kwiatowski et al. (1992).

74. What happens if, as is quite possible, such tests provide no unambiguous conclusion? In our view, if regulators are unable to definitively identify and measure

[11] For a description of how to do this, see `https://www.stata.com/stata14/structural-breaks/`

rational variation in beta through time, then vague and ad-hoc attempts to incorporate time variation may end up being counter-productive. As such, acting as if beta is constant may well be a reasonable working assumption for regulators.

4.2.2 International practice

75. According to Anthony et al. (2020, pp42-43), the AER primarily uses weekly data applied to estimation periods of various length, including "some older data and some very long estimation windows." The use of weekly data is broadly consistent with the findings of Levi and Welch (2017), Hollstein et al. (2018) and Hollstein (2020). The use of long estimation windows is also consistent with a constant beta and, depending on the cause of the reversion, possibly with a mean-reverting beta as well.

76. Like AER, US Regulators favour the use of weekly data, but those in the Netherlands and the UK prefer daily data. Unlike the AER, all these countries use shorter 3-5 year estimation periods.

4.3 Domestic vs world considerations

77. The extent to which beta estimation should be viewed as a domestic or global exercise confronts two main questions:

    - Should estimation use a domestic or world market index as a proxy for the market portfolio?
    - When there is a dearth of domestic listed firms subject to regulation, should foreign comparators be used for the purpose of estimating beta?

    We address these two issues in turn.

4.3.1 Domestic or world market index

78. In the CAPM, the market portfolio is the portfolio of all traded and non-traded assets in the relevant market. In practice, this is usually proxied by an index of traded securities, so this raises the question as to whether the appropriate index is domestic (e.g., ASX 200) or global (e.g., FTSE All-World).

79. As is widely recognised, the theoretical answer is that it depends on whether individual-country capital markets are *integrated* (i.e., effectively form one global market) or *segmented* (i.e., are many separate markets, each with their own pricing factors). If the former, then the appropriate market portfolio proxy is a world index, because all investors hold that portfolio in an international version of the CAPM; if the latter, a domestic index is more suitable.

80. The question of whether world capital markets are integrated or segmented has long been a matter of debate amongst researchers. As Ibbotson et al. (1982, p82) put it almost 40 years ago:

    > Since international investment occurs, markets cannot be totally segmented. But, interest rates and equity returns ... appear to differ substantially from country to country. We view the world market as partly segmented and partly integrated.

81. Despite considerable liberalisation of capital markets in the following decades, Boyle (2009) reaches much the same conclusion 27 years later, as, more recently still, does Orlowski (2020) for the European Union. Moreover, attempts to deepen financial market integration continue to interest policymakers (e.g., European Central Bank, 2020), suggesting they continue to see a not-insignificant amount of segmentation.

82. While imperfect capital market integration might on its own suggest a preference

for use of a domestic index, there is also a deeper reason for such a choice. Recall from section 2.3 that the CAPM is a special case of a more general model where pricing is with respect to a "good-bad times" variable rather than the market portfolio return. Even if stock markets are largely integrated, business and interest rate cycles can still diverge, often substantially, and a domestic index is likely to be more correlated with domestic economic conditions than a world index. Thus, use of a domestic index is also more consistent with broader asset pricing theory.

83. A domestic index also better meets the simplicity criterion. World index returns are calculated in a specific currency, so an exchange rate adjustment will typically be required when used in a different country. Moreover, there are a range of extant world indices, all constructed in different ways, so a choice would have to be made between them; by contrast, the choice of domestic index is usually more restricted and its construction easily understood. Finally, a world index can only roughly replicate the tax situation faced by investors in different countries.

84. There is also little evidence that the use of a world market index (and the associated international CAPM) is any more accurate than the simpler domestic approach. For example, Harris et al. (2003) find that the domestic and international CAPMs fit US data equally well and hence conclude that the choice of domestic or world index is largely immaterial. Koedijk et al. (2002) and Bruner et al. (2008) arrive at a similar conclusion from analyses of developed and emerging countries.

85. Overall, given the lack of empirical support for the view that use of a world market index improves the accuracy of beta estimates, the greater simplicity (and consistency with the imperfect integration of global capital markets) afforded by a domestic index should be preferred.

4.3.2 Foreign comparators

86. When there is a shortage of comparable domestic firms to draw on, it is often tempting to look to comparable foreign firms in order to, hopefully, obtain more reliable and more precise beta estimates. While in some circumstances there may be no option but to follow this path, it should be treated with considerable caution as there are significant dangers.

87. First, if the foreign comparators have different asset portfolios to their domestic counterparts, then both their cash flow and discount rate betas may also differ. In this case, the use of foreign comparators risks introducing bias into the estimate of the domestic firm's beta.

88. Second, even if the foreign comparators have identical asset portfolios to the domestic firm(s) of interest, different compositions of the respective market indexes are likely to result in different betas, and hence an incorrect beta estimate for the domestic firm. For example, there is no good reason to believe a given firm would have the same beta with respect to an index heavily weighted in tech stocks as it would to an index heavily weighted in textiles or natural resources, but this is what the use of foreign comparators implicitly assumes.[12]

89. Third, different groups of foreign comparators yield different estimates of beta. For example, Lally (2005) argues that US electric utilities and gas distribution firms have an estimated asset beta of 0.3 whereas the equivalent estimate for UK regional electricity companies is 0.5. The choice of appropriate foreign comparators is likely to be highly subjective, and can potentially result in very different beta estimates.

90. The problems identified in paras 87–89 reflect a more general issue. Regulators faced with the need to determine beta for n firms often have relevant data for only

[12]Of course, if the international CAPM, and hence a world market index, is used, this problem disappears. But, as discussed in section 4.3.1, a world index creates other problems.

a subset of these firms. For example, in the current Australian energy situation, we understand only three of the regulated firms are still listed, and regulated activities form only a small part of the activities of one of these, thereby reducing the number of firms with useful data to just two. In these circumstances, it seems to be common practice to argue that such a sample is too small to yield reliable beta estimates and hence foreign comparators are needed.

91. Such an approach sounds reasonable, but has some drawbacks. If one wished to obtain an estimate of the 15 June temperature in Sydney, one would not claim that data from one city (Sydney) is too small a sample to yield a reliable estimate, or that it requires supplementing with additional June data from Perth, Darwin, Brisbane, Adelaide and Melbourne. Yet this is what the procedure described in para 90 implies.

92. If the regulator believes all n firms have the same beta (which would be implied by a decision to assign the same beta to all of them in determining their cost of capital), then having usable data for just one regulated firm is sufficient. All that's required is to use that data to obtain the best possible beta estimate for that one firm, with consistency then implying that it applies equally well to the remaining $n-1$ firms. No foreign comparators are required.

93. If the regulator does not believe the n firms have homogenous betas, then, obviously, it should not assign them all the same beta. Instead, a separate beta needs to be estimated for each firm, or sub-group of firms. In this case, if some firm groups do not contain any listed firms, then foreign comparators may be required, but, as described above, caution is required.

94. The general point is that averaging beta estimates over a group of chosen comparators introduces an additional source of estimation error. We document this rigorously in Appendix B, but the general idea is straightforward. Consider the

case where a regulator wishes to determine the beta for a particular firm i. To do so, it collects a sample of comparator firms and estimates the beta for each. If the average of these estimates is used to estimate the beta for firm i, then there are two sources of estimation error. The first arises from the sampling error in each individual estimate, as each comparator firm's estimated beta differs from its true value; this also applies in the case where the beta is estimated from firm i data alone. The second arises from the firm-specific characteristics that cause firm i's true beta to differ from the average of the other firms' true betas. That is:

$$\text{Total estimation error} = \text{sampling error} + \text{intrinsic variation}$$

When using data that are not directly related to firm i beta (i.e., foreign comparator data), both components need to be considered.

95. As we show in Appendix B, the estimation error resulting from intrinsic variation does not vanish as the number of comparator firms becomes large (which is typically the reason for resorting to the use of such firms). Indeed, even when there is an infinitely large number of comparator firms that have zero intrinsic variation error on average, there is no guarantee that the beta estimate they yield is more precise than an estimate obtained from firm i data alone.

96. Consistent with this view, Levi and Welch (2020) consider the practice of using comparator firm betas instead of own-firm betas and conclude on a pessimistic note. First (p427), "(h)istorical industry averages have almost no predictive power and should never be used." Second, although smaller, more carefully targeted, groups of industry peers are somewhat better predictors of realised beta (although still inferior to own-beta), this only applies if the peers are of similar size. Thus, in choosing comparator firms, it is important to not only match business activity, but also size.[13]

[13] By contrast, grouping on book-to-market is not successful. Firm size might therefore be seen as

97. Overall, the problems created by the use of foreign comparators in estimating beta suggest it would be preferable, wherever possible, to rely primarily, if not solely, on data from local firms.[14]

4.3.3 International practice

The AER employs a domestic index, as do regulators in NZ, the US and UK. However, the Netherlands uses a Eurozone index. With respect to the use of foreign comparators, AER employs only local firms in estimating beta. This contrasts with the practice among regulators in the Netherlands, Italy, NZ and the US (FERC) who all make some use of foreign comparators.

4.4 "Low-beta bias"

98. A commonly-observed phenomenon is that the relationship between actual securities returns and beta is weaker than implied by the CAPM. Specifically, the slope of the relationship is flatter than the CAPM predicts and the intercept is higher. As a result, average realised returns on securities with beta less than one are higher than they should be according to the CAPM and so the CAPM under-estimates true average returns. This is sometimes referred to as "low-beta bias" (Frontier, 2017).

99. Assuming this phenomenon is based on sufficiently data to be statistically robust (which is typically the case), there are two possible explanations:

an important criterion for AER to consider in choosing comparator firms (if own-firm betas are not available) but, to repeat, the analysis of Levi and Welch comes from US data, so it would first be necessary to determine if the same result holds in Australia.

[14] A side issue that arises with the use of foreign comparators subject to different regulatory regimes is whether or not the type of regulation (e.g., rate of return versus incentive-based) affects beta, and therefore whether foreign comparator betas require adjustment to take account of this. We discuss this point in greater detail in section 4.9.

- The CAPM is wrong;
- The CAPM is right, but markets are inefficient, i.e., the observed relationship reflects bias in investor expectations and hence mis-pricing.

100. If the CAPM is wrong, then the appropriate solution is to adopt another model (e.g., Black CAPM, Fama-French 3-Factor) that better fits the data. In the absence of any such model (or at least agreement about such a model), we understand that one suggestion is that the equity beta be raised to "offset" the observed bias. Because the problem lies not in the estimation of beta per se, we are sceptical about such an approach. First, a flatter slope actually implies a lower market risk premium and a higher riskless rate, not a higher beta. Second, any such adjustment must either be ad-hoc with no theoretical justification, or it must be reverse-engineered to produce the expected return generated by another model, in which case it is more sensible and transparent just to use that model directly.

101. If the bias is due to market inefficiency, then, as already noted in para 42, Stein (1996) has shown that the optimal approach for investors with a long-run investment horizon (which is presumably the case for regulators) is to use the standard CAPM without any adjustment to the beta. Intuitively,

> (t)his is because β – if calculated properly – may continue to be a reasonable measure of the fundamental economic risk of an asset, even if it has little or no predictive power for stock returns. (Stein, p432)

102. Overall, regardless of whether the observed low-beta bias is due to CAPM deficiencies or biased market expectations, we see no reason to adjust the estimate of beta. Doing so increases complexity, reduces transparency, and lowers accuracy. If the low-beta bias is a model problem, then the consistent approach is to adopt a superior model. If the bias is a market inefficiency problem, then the consistent approach is to stick with the standard beta estimate.

4.5 Non-synchronous regulatory and beta periods

103. An issue we have been asked to comment on is the AER legal requirement to set a beta once every four years that then applies to all new revenue determinations (which last five years) during that period. As a result, a regulated firm can be operating under a beta that is, according to the AER's own determination, out of date for much of the regulatory period. For example, suppose AER sets beta equal to β_t in year t and then subsequently revises this to β_{t+4} four years later. Then any firm that has a revenue determination in, for example, year $t+3$ will have its 5-year determination based on β_t despite the fact that β_{t+4} is the appropriate beta for four of the five year period. For brevity, we refer to this as non-synchroneity.

104. Three basic observations are in order. First, if the beta of Australian regulated firms is, or is approximately, a constant (see para 69), then $\beta_t \approx \beta_{t+4}$ and non-synchroneity is a non-problem.

105. Second, if beta is time-varying, then, so long as each four-yearly determination provides a full-information forecast of beta (which must be true by assumption), beta will only change in year $t+4$ due to the arrival of new information that was unknown and unpredictable in year t, i.e., the *expected* change equals zero. Thus, at least in the "long run", consumers are not disadvantaged by non-synchroneity and regulated firms are not incentivised to under-invest (as might be the case if beta was expected to increase between year t and $t+4$).

106. Third, a stable and transparent regulatory system would be expected to produce only small changes in beta between between year t and $t+4$.[15] Moreover, given the uncertainties involved in estimating beta and the associated (although not

[15] Technological advancements could of course markedly increase competition for what were previous natural monopoly activities, which could certainly have a larger effect on beta, but then those activities would no longer need to be regulated and the issue becomes redundant.

usually estimated or reported) high standard errors, any such small changes are almost certain to lie within the margin of error and so implementing them *within* a determination period seems likely to represent superfluous precision.

107. We understand that one suggested solution to the non-synchroneity problem is that a "formula" be created to bring the beta within a revenue determination period up to date following the estimation of a new beta. While it is difficult to comment sensibly on this without detailed information on the proposed formula, we are sceptical about the value of such an approach. Any such formula seems likely to be over-prescribed, increase complexity, and be of unknown accuracy. Moreover, the potential for beta to change within a revenue determination period seems inconsistent with the intention of such determinations being, at least in part, to give some degree of certainty to all involved.

108. Overall, while non-synchroneity is not ideal (unless beta is a constant), we are doubtful that the benefits of any formulaic solution will exceed the costs.

4.6 Adjustment for specific risks

109. An argument that is sometimes made in regulatory hearings, and which CRG has asked us to comment on, is that regulated firms should be allowed a higher beta as compensation for certain specific, or idiosyncratic, risks, e.g., technological changes that result in stranded assets.

110. In the CAPM world, only systematic, or market, risk is priced and so no allowance for specific risks is admissible. As Stulz (1999) points out, this means the cost of equity for any project is the same regardless of the firm that undertakes it. But not all firms have equal financial capacity and, intuitively, there would seem to be some difficulty with the proposition that a project which could wipe the firm out

if it goes badly can be of the same (or even less) risk than another project that has no significant implications for the firms financial health.

111. The missing link here is that, when dealing with company investment decisions, a blind application of the CAPM assumes that the firm's ability to fund investments is independent of its total (including idiosyncratic) risk. Consequently, a firm that approaches financial distress can always obtain fair-priced financing and thus avoid a distress situation. But as Stulz (p9) points out "In the real world, such costless recapitalization is a dream." Increases in total risk can reduce a company's access to capital, or increase the price it must pay for access, because of customer and supplier doubts about future viability, or because of agency problems that incentivise "bet the house" or empire-building behaviour.

112. By making future funding more expensive, higher total risk can result in the firm being unable to take advantage of valuable investment opportunities that it could have exploited if its financial position were stronger. The greater probability of this occurring in the future reduces the value of those opportunities now, i.e., the value of its investment opportunities, or real options, fall. Therefore, investment in a project that materially raises the total risk of the firm changes firm value not only by the static net present value (NPV), but also by the fall in value of future investment opportunities. As a result, investment in the project is justified if and only if the NPV exceeds the fall in value of the future investment opportunities. In effect, companies impute a higher cost of capital, and thus require a higher expected return, than would be implied by their systematic risk alone.

113. Given the current state of knowledge about this issue, we are agnostic as to whether or not such an adjustment should be allowed for regulated firms. While there is evidence that companies systematically set investment hurdle rates well in excess of their CAPM cost of capital (Jagannathan et al., 2016), with the dif-

ference increasing in idiosyncratic risk, the absence of any good theory capable of quantifying such a premium makes it difficult to rule out other possibilities.

114. One thing though is certain: this is not a beta issue. If an adjustment for specific risk is to be made, then this needs to be explicitly modelled and quantified and then separately allowed for. Adding an idiosyncratic risk increment to the best estimate of a firm's systematic risk, thus assuming that both risks face the same price, is ad hoc and incoherent. Proceeding in this way would potentially open up all kinds of inconsistent adjustments to beta estimation that we do not think should be encouraged.

4.6.1 International practice

115. As far as we are aware, no regulator allows any explicit adjustment of this kind. The NZCC does, however, set its allowed return at the 67th (rather than 50th) percentile of the estimated cost of capital distribution, which has a broadly similar effect to an allowance for idiosyncratic risk. Two things can be said about such a policy. First, the magnitude of the adjustment reflects the precision of the estimation process rather than the risk of the regulated assets; there is no reason why the two should be related. Second, such an allowance nevertheless seems preferable to adjusting beta.

4.7 Statistical adjustments

4.7.1 Shrinkage

116. All beta estimates are subject to sampling error. As a result, an estimate based on past data (one sample) often works less well when applied to future data (another sample). With unlimited data and a time-invariant beta, this should not present any problems. But in more constrained circumstances, problems can arise. For

example, suppose we wish to estimate the "true" batting average of test cricketer X. And suppose also that after two matches cricketer X is averaging 120. Because test batting averages over 50 are rare, we can be confident that X's true average is not 120. Yet 120 is what our past data on its own suggests.

117. When estimating beta, mean-reversion is, as we have frequently noted, sometimes thought to be present. In such a case, estimates of the current beta, which have to be obtained from relatively short data periods, may be sensitive to outliers and other sample peculiarities, thus rendering them unreliable.[16]

118. In such circumstances, estimation can often be improved by a statistical technique known as shrinkage (Copas, 1983). The idea is that the estimate based on past data should be combined with some prior estimate based on the underlying population to produce a superior estimate. For example, suppose the overall average of top-6 batsmen in post-WWII test cricket is 35. Then shrinkage involves using a weighted sum of 120 (the historical data estimate) and 35 (the prior) to arrive at a more realistic estimate of X's true batting ability.

119. When applied to beta, the prior is usually set to 1.0 since that reflects, by definition, the beta of the average firm. The most widely used beta shrinkage adjustment is that of Vasicek (1973) who weights historical estimates based on their standard error relative to the standard error of the of the overall market, the idea being that estimates with lower standard errors are given a higher weighting relative to the prior of 1.0.

120. Using US data, Levi and Welch (2017) find that the Vasicek adjustment improves the predictive accuracy of standard OLS beta estimates. Perhaps more interestingly, they also report that additional improvement is possible by using so-called double shrinkage. Specifically, for the purpose of forecasting future beta, they

[16] Note this assumes the regulator wishes to estimate current, as opposed to long-run, beta.

recommend using:

$$\hat{\beta} = 0.75\hat{\beta}^v + 0.25\hat{\beta}^p \tag{9}$$

where β^v is the Vasicek beta estimated from equation (8) and β^p is the prior beta which is categorised according to firm size. That is, the most accurate method for forecasting beta is via a weighted average of the (already-shrunk) Vasicek beta and a target beta based on firm size.[17]

121. Using data from 48 countries, Hollstein (2020) also reports that the Levi and Welch double-shrinkage beta is one of the best approaches for predicting future realised beta.

122. Whether shrinkage should be applied to regulated firm betas depends on the statistical properties of those firms' betas. If beta is a constant, then shrinkage is unnecessary.[18] If beta is mean-reverting, then the value of shrinkage again depends on the reason for reversion (see paras 66-67). If it reflects rational time variation in risk, then the regulator's objective is to estimate the current (i.e., up-to-date) beta and the results of Levi and Welch (2017) and Hollstein (2020) suggest shrinkage is required. If it reflects under- and over-pricing by investors, then the model of Stein (1996) implies the objective should be to predict "long-run" beta and shrinkage is less likely to be necessary.

123. A final issue arising from shrinkage use is a possible inconsistency when using foreign comparators. In such circumstances, it is often tempting to use the beta estimates provided by commercial operators, e.g., Bloomberg, Value Line. Some operators provide raw betas, others shrink their betas, while Bloomberg offers both (Levi and Welch, 2017). Obviously, it is important to ensure that any betas

[17] Levi and Welch recommend $\beta^p = 0.5$ (0.7) (0.9) according to whether the firm is in the smallest (middle) (largest) tercile.

[18] In equation (8), the parameter γ_1 is approximately equal to 1 minus the rate of mean reversion. So in the case where reversion is instantaneous ($\gamma_1 = 0$), beta is a constant.

sourced in this way are consistent with those estimated directly from domestic firms, e.g., if no shrinkage is applied to domestic beta estimates, then these should not be combined with a shrunk foreign beta.

124. Overall, any accuracy benefits from the use of shrinkage would need to be balanced against the greater complexity and lower transparency. Also, consistency would suggest shrinkage not be used if other choices indicate a belief that beta is constant, e.g., a long estimation period.

4.7.2 Thin trading

125. Infrequent, or thin, trading can affect beta estimates in two ways. First, when a security is thinly traded, or illiquid, the increased presence of zero returns (particularly in daily data) induces a downward bias in the estimate of that security's beta.[19] Second, if thinly-traded securities are present in a market, the downward bias in their beta estimates must be associated with an upward bias in heavily-traded securities on average (since the average beta equals 1).

126. Thus, thin trading can cause the beta estimate for a regulated entity to be under- or over-estimated. If the firm's securities are themselves illiquid, then its beta estimate is likely to be too low; if the firm's securities are illiquid in a market where other securities are thinly traded, then its beta estimate could be too high. These problems have encouraged researchers and practitioners to employ adjustments that mitigate the effects of thin trading in daily data.

127. Although several such adjustments exist, the most popular is that of Dimson (1979) which adds lagged and lead returns to the estimation process. However, recent research provides little support for the use of this adjustment: both Levi and Welch (2017) and Hollstein (2020) conclude that the Dimson beta is a less

[19] For a simple numerical illustration of this phenomenon, see P&S Group (2016).

accurate predictor of realised beta than the OLS beta and a considerably less accurate predictor than the Vasicek beta. One possible explanation for this is that the markets studied by these authors are strongly liquid, so the adjustment just incorporates superfluous information and thus results in a less accurate measure of beta. This suggests that, at least in modern well-developed capital markets (such as Australia), the upward-bias phenomenon is unlikely to be a phenomenon.

128. This leaves the possibility of downward bias if the regulated Australian entities are themselves illiquid. Determining whether or not this is the case obviously begs the question of what level of trading infrequency is needed for a security to be "thinly traded", in the sense of inducing a material downward bias in the beta estimate. Although we know of no hard and fast rule on this, or even a rule of thumb, an extension of the P&S Group (2016) simulation analysis would allow an approximate threshold to be determined. This could then be compared to the actual situation of regulated Australian securities in order to determine whether a potential problem exists. For example, suppose the analysis indicated that material bias in beta estimates was likely to occur only if the number of non-trading days for a security exceeded 5%. If the actual number of non-trading days for an Australian regulated firm were then observed to be 2%, thin trading concerns could justifiably be ignored.

4.7.3 International practice

129. AER eschews the use of both the Vasicek and Dimson adjustments. According to Anthony et al. (2020), only the Netherlands employs both adjustments, while the US energy regulator uses the simplified shrinkage adjustment of Blume (1971). None has yet made use of the Levi and Welch (2017) double-shrinkage adjustment.

4.8 Indirect beta estimation

130. Given the difficulties involved in estimating beta directly from returns data, CRG has asked us to comment on the possibility of indirect approaches, using variables that are observable (or more easily estimated) to predict beta.

131. We are not aware of any existing work in this regard, but can warily suggest a simple approach that utilises the Gordon Growth Model of stock pricing (see Brealey et al., 2020, p210). As we show in Appendix C, this implies the regression equation:

$$\beta_{it} = a + bDY_{it} + c\bar{g}_i + dR_{ft} + \zeta_{it} \tag{10}$$

where DY_{it} is the time t dividend-price ratio for asset i, $\bar{g}_i$ is the constant expected growth rate in dividends for asset i, and ζ is an error term.

132. Equation (10) could be estimated using, for example, all stocks listed in the ASX200 in order to obtain robust estimates of (a, b, c, d). Together with observed values of DY_i and R_f, and consensus analyst forecasts of $\bar{g}_i$, these could then be substituted back into (10) to obtain an estimate of beta for each firm i in which regulators are interested.[20]

133. Of course, such an approach may just substitute one set of problems for another and so be of little incremental value. Before it could be confidently applied, much testing would be necessary. Overall, we are not optimistic about the potential of this or similar indirect approaches.

[20] There is an element of circularity here in that the dependent variable β_{it} must first be directly estimated from returns data, but this is unavoidable.

4.9 Beta endogeneity

134. We have also been asked by CRG to comment on possible circularity in the estimation of beta and, more generally, beta endogeneity.

135. The regulatory objective in setting beta is to quantify the systematic risk applicable to the regulated assets during the regulatory period, i.e., *after taking into account any effect of the regulation on the assets' systematic risk*. The standard approach that estimates beta from data reflecting regulatory settings does precisely this. By contrast, the alternative approach of attempting to estimate beta from data "purged" of any regulatory influences would systematically overestimate beta (assuming that regulation tends to reduce systematic risk).

136. This matter potentially becomes more complicated if the type of regulation (e.g., price-cap versus rate-of-return) changes between regulatory periods. In this case, the problem is (again) not one of circularity, but simply one of having to use data from one regulatory system to estimate beta for another regulatory system.[21] Of course, whether this is a problem in practice depends on the existence of a systematic relationship between beta and regulation type.

137. Evidence on the relationship between beta and regulation type is decidedly mixed. Early work by Alexander (1996) based on five years of data from the US and UK found that the beta estimates for price-capped UK firms average almost double those of rate-of-return regulated US firms and that this difference is strongly statistically significant. However, Buckland and Fraser (2001) cast doubt on this conclusion by showing that the UK beta estimates were artificially inflated by two events that were independent of the regulatory system.

138. Subsequent evidence has not been very supportive of the view that beta is affected

[21] Obviously, this is also an issue when using foreign comparators from different regulatory systems.

by regulation type.[22] The most comprehensive study is that of Gaggero (2010) who compares beta estimates across a wide range of countries, industries and regulatory systems. Differences are small and statistically insignificant, resulting in Gaggero (p233) concluding that "...the methods of regulation have no impact on the level of systematic risk to which regulated firms are exposed." Using US data only, Allen Consulting Group (2008) and CEG (2013) arrive at the same position. Based on similar analysis, the NZ Commerce Commission (NZCC, 2010) also concludes that beta is materially independent of regulation type.

139. Absence of evidence is not the same as evidence of absence, particularly when, as is the case with beta estimates, the underlying data are noisy. Thus, the possibility of beta being sensitive to regulation type cannot be dismissed with 100% confidence. But the current body of evidence is more supportive of the view that beta is independent of the form of regulatory control.[23]

140. One possible explanation for this finding is that company betas primarily reflect their discount rate sub-betas, while differences in regulation type affect only cash flow sub-betas (see section 2.4). However, in the absence of any supporting analysis, this conjecture is speculative.

5 Correspondence with the AER approach

141. The method by which the AER determines beta is specified in its Rate of Return Instrument. The Rate of Return instrument is binding on the AER, and a new rate of return instrument is published every four years. Table 1 summarises the application of our section 3.3 criteria and section 4 analysis to the current instrument rate of return instrument, which was published in 2018.

[22] For a detailed summary, see Lally (2016).

[23] This conclusion partly rehabilitates the case for foreign comparators, but all the other problems remain — see section 4.3.2.

Table 1: Applying the criteria to the AER rate of return instrument

Issues in determining beta	**AER approach (2018 RoR Instrument)**	**Application of criteria**
Excess or raw returns	AER appears to use raw returns.	*Accuracy* unlikely to be material issue given use of weekly return data. *Consistency* with CAPM suggests a possible preference for raw returns. *Simplicity* slightly favours raw returns.
Frequency of data	AER favours use of weekly returns	*Accuracy* best served by high-frequency daily or weekly data. *Consistency* with CAPM favours high-frequency data as eliminates choice between raw and excess returns, and return period. *Simplicity* favours use of high-frequency data as avoids need to decide return period (e.g., month end or mid-month).
Estimation period	AER considers the longest-term data is most reflective of the equity beta value.	*Accuracy* will depend on whether beta is constant or mean-reverting. *Consistency* with CAPM will depend on the underlying cause of any mean reversion. *Simplicity* slightly favours short return period for data availability.
Beta constant or time-varying	AER gives most weight to longest estimation period, indicating a view that beta is approximately constant.	*Accuracy*: whether beta of regulated Australian network firms vary over time is an empirical question. *Consistency* with CAPM will depend on whether beta is constant or varies over time, but acting as if beta is constant reasonable if no strong evidence of mean reversion. *Simplicity* favours assuming beta is constant until evidence proves otherwise.
Domestic-world considerations	AER uses domestic data, with international data used as check.	*Accuracy*: lack of empirical support that a world market index/foreign comparators improves accuracy of beta estimates. *Consistency* with imperfect integration of global capital markets favours domestic index/comparators. *Simplicity* favours domestic index/comparators.
Low beta bias	AER retains the Sharpe-Lintner CAPM.	*Accuracy*: ad hoc adjustments would reduce accuracy. *Consistency*: if a superior model exists, then consistent approach is to adopt that model (not make ad hoc adjustments to Sharpe-Lintner CAPM). *Simplicity*: ad hoc adjustments would increase complexity and lower transparency.

Applying the criteria to the AER rate of return instrument cont...

Non-synchronous beta periods	AER sets beta once every four years and applies that to all new revenue determinations during that period.	*Accuracy*: if beta is constant no issue arises; if beta is time varying, changes are likely within margin of error. *Consistency:* an adjustment to beta within a revenue determination seems inconsistent with objective of certainty for all parties. *Simplicity*: adjustments would increase complexity.
Adjustment for specific risks	AER does not adjust beta for specific or idiosyncratic risks	*Accuracy*: adjusting beta for idiosyncratic risk would not improve accuracy of beta estimate. *Consistency*: adding idiosyncratic risk increment to best estimate of systematic risk would be incoherent. *Simplicity*: adjustments would increase complexity.
Statistical adjustment	AER does not adjust beta for shrinkage or thin trading.	*Accuracy*: whether statistical adjustments would increase accuracy depends on the statistical properties of regulated firm's beta. *Consistency*: cannot be determined absent information on the statistical properties of beta. *Simplicity*: adjustments would increase complexity.
Indirect beta estimation	AER does not estimate beta by indirect approaches.	*Accuracy*: an indirect approach may just substitute one set of problems for another. *Consistency*: may not be consistent with other elements of the CAPM. *Simplicity*: variables may be more observable or more easily estimated but fit with CAPM may be complex.
Beta endogeneity	AER considers that evidence is not persuasive that differences in regulatory frameworks warrant different betas.	*Accuracy*: Evidence on relationship between beta and regulation is mixed. *Consistency*: potential to introduce inconsistency where regulation impacts on cashflows but not discount rate sub-beta. *Simplicity*: attempting to purge data of regulatory influences would add complexity.

6 Conclusion

142. Beta estimation is an inexact science subject to considerable uncertainty. In our view, only a few statements can be made with any real confidence:

- Daily (or possibly weekly) data are preferred to lower-frequency data.
- A domestic index proxy is preferred to a world index.
- The so-called "low-beta bias" should not be accommodated by adjusting beta.
- Significant company-specific risks should not be accommodated by adjusting beta.

143. Ambiguous statements are just as common:

- The choice between raw and excess returns is immaterial if daily (and possibly weekly) returns are used, but not otherwise.
- If beta is believed to be a constant, then the estimation period should be set as long as possible; if beta is believed to be time-varying (mean-reverting), then the choice is more complicated and depends on the reason for reversion.
- The use of foreign comparators may be necessary in some circumstances, but doing so introduces a number of estimation problems.
- Beta shrinkage may be useful if regulated firm betas are believed to be mean-reverting, but not otherwise.
- Because they avoid the above problems, indirect approaches to estimating beta are attractive in principle, but are likely to create other, potentially bigger, problems.
- Existing evidence does not support the view that beta is affected by regulation type, but nor does it definitively rule this view out.

References

Allen Consulting Group, 2008. Beta for regulated electricity transmission and distribution. Available at https://www.aer.gov.au/system/files/Joint%20Industry%20Association%20-%20Appendix%20H%20-%20ACG%20-%20Beta%20for%20regulated%20electricity%20transmission%20and%20distribution.pdf.

Anthony, J., T. Brown, L. Figurelli, D. Harris, N. Nguyen and B. Villadsen, 2020. A review of international approaches to regulated rates of return. Brattle Group. Available at https://www.aer.gov.au/system/files/Report%20to%20the%20AER%20-%20A%20Review%20of%20International%20Approaches%20to%20Regulated%20Rates%20of%20Return%20-%2030%20June%202020.pdf

Blume, M., 1971. On the assessment of risk. Journal of Finance 26, 1-10.

Blume, M., 1975. Betas and their regression tendencies. Journal of Finance 30, 785-795.

Boyle, G., L. Evans and G. Guthrie, 2006. Estimating the WACC in a regulatory setting. NZ Institute for the Study of Competition and Regulation. Available at https://ir.canterbury.ac.nz/handle/10092/2437

Boyle, G., 2009. Capital market integration: a review of the issues and an assessment of New Zealand's position. Ministry of Economic Development and the Capital Market Development Taskforce. Available at https://researcharchive.vuw.ac.nz/xmlui/handle/10063/4034.

Brealey, R. and S. Myers, 1996. Principles of Corporate Finance (5th ed.). New York: McGraw-Hill.

Brealey, R., S. Myers and A. Marcus, 2020. Fundamentals of Corporate Finance (10th ed.). New York: McGraw-Hill.

Brown, P. and T. Walter, 2013. The CAPM: theoretical validity, empirical intractability and practical applications. Abacus 49, 44-50.

Bruner, R., W. Lia, M. Kritzman, S. Myrgrenc and S. Page, 2008. Market integration in developed and emerging markets: Evidence from the CAPM. Emerging Markets Review 9, 89-103.

Campbell, J., 1991. A variance decomposition for stock returns. Economic Journal 101, 157-79.

Campbell, J., A. Lo and A. MacKinlay, 1997. The Econometrics of Financial Markets. Princeton University Press.

Campbell, J. and J. Mei, 1993. Where do betas come from? Asset price dynamics and the sources of systematic risk. Review of Financial Studies 6, 567-592.

Campbell, J. and R. Shiller, 1988. The dividend-price ratio and expectations of future dividends and discount factors. Review of Financial Studies 1, 195-228.

CEG, 2013. Information on equity beta from US companies. Available at https://www.erawa.com.au/cproot/11653/2/Energy%20Networks%20Association%20-%20Draft%20Rate%20of%20Return%20Guidelines%20-%20Report%209%20-%20US%20Beta%20Comparables.pdf

Cochrane, J. 2005. Asset Pricing (2nd ed.). Princeton, NJ: Princeton University Press.

Cochrane, J., 2008. The dog that did not bark: a defense of return predictability. Review of Financial Studies 21, 1533-1575.

Cochrane, J., 2011. Discount rates. Journal of Finance 66, 1047-1108.

Copas, J., 1983. Regression, prediction and shrinkage. Journal of the Royal Statistical Society: Series B (Methodological) 45, 311-335.

Cornell, B., 1999. Risk, duration, and capital budgeting: new evidence on some old questions, Journal of Business 72, 183-200.

Dimson, E., 1979. Risk measurement when shares are subject to infrequent trading. Journal of Financial Economics 7, 197-226.

European Central Bank, 2020. ECB pushes for EU capital markets integration and development. Available at https://www.ecb.europa.eu/press/pr/date/2020/html/ecb.pr200303~5eaf4c119d.en.html

Fama, E. and K. French, 2004. The capital asset pricing model: theory and evidence. Journal of Economic Perspectives 18, 25-46.

Friedman, M., 1953. Essays in Positive Economics. Chicago: University of Chicago Press.

Frontier Economics, 2017. Low-beta bias: a report prepared for Actewagl Distribution. Available at https://www.aer.gov.au/system/files/Evoenergy%20-%20Rate%20of%20return%20-%20Appendix%208.3%20-%20Frontier%20-%20Low-beta%20bias%20-%20December%202017_Public.pdf

Gaggero, A., 2010. Regulation and risk: a cross-country survey of regulated companies. Bulletin of Economic Research 64, 226-238.

Grinblatt, M., and S. Titman, 2002. Financial Markets and Corporate Strategy (2nd ed.), McGraw-Hill Irwin, New York, NY.

Harris, R., F. Marston, D. Mishra and T. O'Brien, 2003. Ex ante cost of equity estimates of S&P 500 firms: the choice between global and domestic CAPM. Financial Management 32, 51-66

Henry, O., 2009. Estimating β. ACCC. Available at https://www.aer.gov.au/system/files/Attachment%20C%20-%20Henry%20-%20Estimating%20beta_0.pdf

Hollstein, F., 2020. Estimating beta: the international evidence. Journal of Banking and Finance 121, 1-20.

Hollstein, F., M. Prokopczuk, and C. Wese Simen, 2019. Estimating beta: forecast adjustments and the impact of stock characteristics for a broad cross-section. Journal of Financial Markets 44, 91-118.

Ibbotson, R., R. Carr and A. Robinson, 1982. International equity and bond returns. Financial Analysts Journal 38, 61–83.

Jacobs, M. and A. Shivdasani, 2012. Do you know your cost of capital? Harvard Business Review 90, 118-124.

Jagannathana,R., D. Matsaa, I. Meiere, and V. Tarhan, 2016. Why do firms use high discount rates? Journal of Financial Economics 120, 445-463

Koedijk, K., C. Kool, P. Schotman, and M. Van Dijk, 2002. The cost of capital in international financial markets: Local or global? Journal of International Money and Finance, 21, 905-929.

Kwiatkowski, D., P. Phillips, P. Schmidt and Y. Shin, 1992. Testing the null hypothesis of stationarity against the alternative of a unit root. Journal of Econometrics 54, 159-178.

Lally, M., 2005. The weighted average cost of capital for electricity lines businesses. Available at `https://comcom.govt.nz/__data/assets/pdf_file/0027/88362/The-Weighted-Average-Cost-of-Capital-for-Electricity-Lines-Businesses.pdf`

Lally, M., 2016. Review of WACC issues. Available at `https://comcom.govt.nz/__data/assets/pdf_file/0016/61225/Dr-Lallys-expert-advice-on-asset-beta-adjustments-and-Blacks-simple-discounting-rule-25-February-2016.pdf`

Lehmann, B., 1990. Fads, martingales, and market efficiency, Quarterly Journal of Economics 105, 1-28.

Levi, Y., and I. Welch, 2017. Best practice for cost-of-capital estimates. Journal of Financial and Quantitative Analysis 52, 427-463.

Lintner, J., 1965. The valuation of risk assets and the selection of risky investments in stock portfolios and capital budgets. Review of Economics and Statistics 47, 13-37.

McKenzie, M. and G. Partington, 2012. Report to the AER: Estimation of the equity beta (conceptual and econometric issues) for a gas regulatory process in 2012. Available from the AER website at https://www.aer.gov.au/system/files/RBP%20gas%20transmission%202012%20-%20Equity%20Beta%20report%20-%20McKenzie%20and%20Partington%20%28Public%29%20-%203%20April%202012_0_0.pdf

Morse, D., 1984. An econometric analysis of the choice of daily versus monthly returns in tests of information content. Journal of Accounting Research 22, 605-623.

NZCC, 2010. Input methodologies (electricity distribution and gas pipeline services) reasons paper. Available at https://comcom.govt.nz/__data/assets/pdf_file/0026/95570/Input-Methodologies-EDBGPB-Extracts-from-Chapter-6-and-Appendix-E-December-2010.pdf.

Orlowski, L., 2020. Capital markets integration and economic growth in the European Union. Journal of Policy Modeling 42, 893-902

Oxera, 2015. Turning down the volume: estimating cost of capital in the presence of noise. Available at https://www.oxera.com/insights/agenda/articles/turning-down-the-volume-estimating-cost-of-capital-in-the-presence-of-noise/, retrieved 5 May 2021.

P&S Group, 2016. Illiquid markets and thin-trading effect on beta estimates. Available at https://p-s.com/news/illiquid-markets-and-thin-trading-effect-on-beta-estimates/

Robertson, D., 2020. Estimating beta. University of Cambridge. Available at https://www.ofgem.gov.uk/ofgem-publications/145143

Sharpe, W., 1964. Capital asset prices: a theory of market equilibrium under conditions of risk. Journal of Finance 19, 425-442.

Stein, J., 1996. Rational capital budgeting in an irrational world. Journal of Business 69, 429-455.

Stulz,R., 1999. Whats wrong with modern capital budgeting? Financial Practice and Education 9, 7-11.

Vasicek, O., 1973. A note on using cross-sectional information in Bayesian estimation of security betas. Journal of Finance 28, 1233-1239.

Wright, S. and R. Mason, 2020. Comments prepared for Ofwat on the CMAs Provisional Findings Anglian Water Services Limited, Bristol Water plc, Northumbrian Water Limited and Yorkshire Water Services Limited price determinations: Cost of capital considerations. Available at https://www.ofwat.gov.uk/wp-content/uploads/2020/10/Wright-and-Mason-Comments-prepared-for-Ofwat-on-the-CMAs-provisional-findings-cost-of-capital-considerations.pdf

Appendix A Comparison with the average firm

Equation (5) can also shed some light on the conceptual question of how a regulated firm's β compares to that of the average firm (for which $\beta \equiv 1$). Recognising that β_r is firm-independent, the regulated firm version of (5) is:

$$\beta_i^r = \beta_{ci}^r - \beta_{ki}^r \tag{11}$$

where the superscript r indicates a firm subject to regulation. For the average firm:

$$1 = \bar{\beta}_i = \bar{\beta}_{ci} - \bar{\beta}_{ki}$$

where a bar ($^-$) indicates the average firm. Combining these two equations yields:

$$1 - \beta_i^r = \underset{(+)}{(\bar{\beta}_{ci} - \beta_{ci}^r)} + \underset{(?)}{(\beta_{ki}^r - \bar{\beta}_{ki})} \tag{12}$$

Equation (12) reveals that the excess β of the average firm over that of the regulated firm equals (i) the difference between the average and regulated cash flow betas plus (ii) the negative of the difference between the average and regulated risk premium betas. Because regulation tends to insulate firm cash flows from market returns, the first term is likely to be positive. However, the second term is, without further information, of indeterminate sign. If the risk premium beta of the regulated firm is approximately equal to that of the average firm, then the second term in (12) is approximately zero. As a result, the excess β of the average firm over that of the regulated firm is likely to be positive, supporting the view that the regulated firm β is less than 1.0.

Appendix B Estimation error using comparator firms

Suppose there are N firms in the comparator firm sample, and that the true beta of firm j is

$$\beta_j = \beta + \theta_j, \quad j = 1, \ldots, N,$$

where each θ_j is independently distributed with mean zero and standard deviation σ.[24] The cross-sectional variation in asset betas is due to intrinsic differences between the firms, e.g., different businesses, operations and regulation. Suppose that we have an estimate of each firm's beta. In particular, for firm j we can observe

$$\hat{\beta}_j = \beta_j + \varepsilon_j, \quad j = 1, \ldots, N,$$

where each ε_j is distributed with mean zero, variance ϕ^2, and correlation $\mathrm{Cor}[\varepsilon_j, \varepsilon_k] = \rho \geq 0$ when $j \neq k$. It follows that $\hat{\beta}_j - \beta_j$ has mean 0 and standard deviation ϕ. Thus, ϕ measures the precision of each firm's beta estimate, information that is readily obtained during the estimation process.

The average of the N estimates is

$$\bar{\beta} \equiv \frac{1}{N} \sum_{j=1}^{N} \hat{\beta}_j$$

The comparator approach uses this "average of the estimates" to estimate the beta $\beta_i \equiv \beta + \theta_i$ for a single firm, labelled i, using estimated asset betas from a set of different firms, labelled $1, \ldots, N$. Since

$$\bar{\beta} - \beta_i = (\bar{\beta} - \beta) + (\beta - \beta_i) = \frac{1}{N} \sum_{j=1}^{N} (\varepsilon_j + \theta_j) - \theta_i,$$

it follows that $\bar{\beta} - \beta_i$ has mean 0 and variance

$$\sigma^2 + \frac{\sigma^2}{N} + \frac{\phi^2}{N}(1 + \rho(N-1))$$

which is strictly decreasing in N. As N gets very large (i.e., a lot of comparator firms are used), the variance approaches

$$\sigma^2 + \rho\phi^2$$

By contrast, the variance of the estimate obtained from firm i data alone is ϕ^2 (i.e., sampling error). So the standard error of the "average of the estimates" beta will exceed the standard error from the single-firm beta if

$$\frac{\sigma}{\phi} > \sqrt{1 - \rho}$$

[24] This section is adapted from Boyle et al. (2006).

Even in the extreme, and unlikely, case that $\rho = 0$, this requires only that the error arising from intrinsic variation exceed the error arising from sampling variation. More generally, the requirement is only that the former not be too much smaller than the latter.[25] Given the potentially wide array of differences between domestic and foreign regulated firms, this condition seems eminently plausible, and certainly cannot be ruled out *a priori*.

Appendix C Indirect beta estimation using the Gordon Growth Model

The Gordon Growth Model states:

$$E[R_{it}] = DY_{it}(1 + \bar{g}_i) + \bar{g}_i$$

where DY_{it} is the time t dividend-price ratio for asset i and $\bar{g}_i$ is the constant expected growth rate in dividends for asset i. Combining this with equation (3) for the CAPM and solving for beta yields:

$$\beta_{it} = \gamma\{DY_{it}(1 + \bar{g}_i) + \bar{g}_i - R_{ft}\} \tag{13}$$

where γ is the reciprocal of the market risk premium. Equation (13) reveals that beta can be expressed as a linear function of the riskless interest rate (observable), the dividend-price ratio (observable), and the long-term growth rate in dividends (must be estimated).

Because of the duration effect discussed in section 2.4, the factor of proportionality γ is unlikely to be the same for all three variables, i.e., the duration effect would suggest

[25] For small N, the condition becomes

$$\frac{\sigma}{\phi} > \alpha\sqrt{1-\rho}, \qquad \alpha = \sqrt{\frac{N-1}{N+1}} < 1$$

which is even more easily satisfied.

that lower DY and higher $\bar{g}$ are associated with higher beta regardless of the Gordon Model. The different loadings for each variable could be estimated via the regression equation:

$$\beta_{it} = a + bDY_{it} + c\bar{g}_i + dR_{ft} + \zeta_{it}$$

where ζ is an error term.