



Effect of Single Commodity Risk on ARTC

12 April 2016

1 Introduction

Following the meeting with the ACCC on 24 March 2016, we have been asked by the Hunter Valley Rail Access Task Force (HRATF) to consider how the fact that the ARTC Hunter Valley network is dependent on a single commodity—coal—should be reflected in the remuneration received by ARTC.

Since ARTC's access prices are not dependent on coal prices, ARTC's exposure to single commodity risk is limited to a catastrophic situation where the coal price falls to sufficiently low levels for a sustained period of time to induce the closure of such a significant number of mines that the remaining volume on the network cannot sustain ARTC's full revenue requirement.


This note uses evidence from the research undertaken by the commodity market specialists Wood Mackenzie to assess the extent of such catastrophic single commodity risk. We then consider how investors in ARTC should be remunerated for such risk.

2 Relevant Single Commodity Risk

Despite the general economic slowdown and growth in other forms of energy (including renewables), global coal generation capacity has continued to grow in recent years. The most recent picture—derived from an essentially anti-coal report prepared by the Sierra Club in the US—is presented in the figure below.

While many of the pre-announced coal fired plants have been shelved over the last few years, many new projects are proceeding. The Sierra Club's Global Coal Plant Tracker estimates that of the projects announced over the past 5 years, approximately half have been shelved. The same source estimates that if all currently announced and pre-planned projects are cancelled, global coal generation capacity in 2030 will still be 505,000 MW higher than it is today, due to the projects already under construction.

Table 2.1: Coal Powered Generation Capacity Additions Net of Retirements (MW)



Sources: 1982–2009, Platts WEPP, December 2015; 2010–2015, Global Coal Plant Tracker, January 2016, and Sierra Club. Figures for 2015 preliminary.

Source: As presented in Boom and Bust 2016: Tracking the Global Coal Plant Pipeline

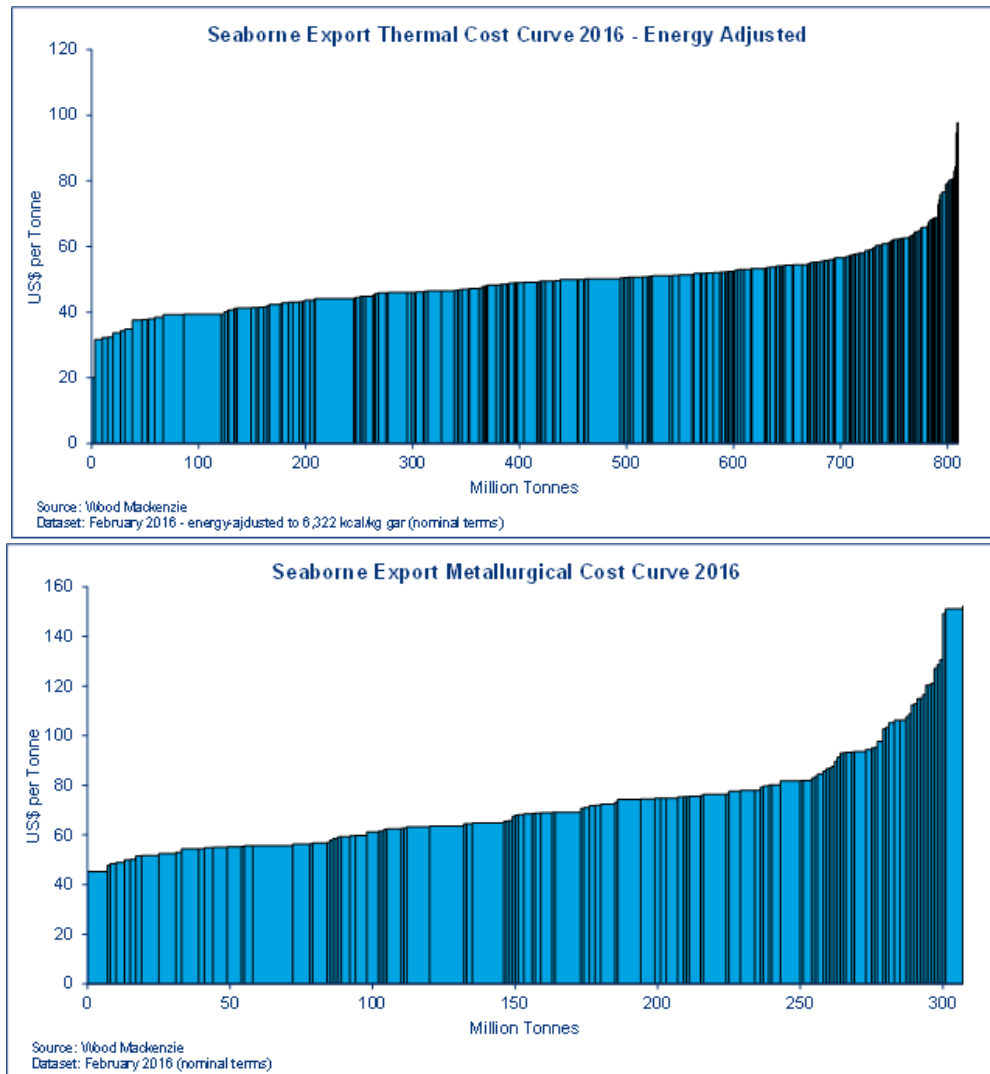
Given the current pattern of investment, global demand for thermal coal can be expected to continue growing (albeit at slower rates) over the expected mine lives in the Hunter Valley. Global seaborne trade—in which Hunter Valley coal plays an important role—will continue for the foreseeable future. In this environment, over the medium term, the coal price can be expected to be set by the long-run marginal cost (LRMC) of the marginal mines required to meet such growing demand.

At any point in time, however, prices will deviate from LRMC depending on the current supply/demand balance. For example, a shortage of production capacity will lead to prices that are higher than LRMC, and excess capacity will drive prices below LRMC towards the short-run marginal cost (SRMC) of the marginal mine.

At the most pessimistic, if global demand for coal remains flat and no new mines have to be developed, global prices will be set by the cost of the incremental production from the marginal mine. Since most of the existing mining capacity will continue to be needed, such marginal mine can be expected to be at the tail end of the current global cost curve.

In this context, the relevant question to consider is where the Hunter Valley mines sit on the global cost curve. The charts below show the thermal and metallurgical coal global cost curves developed by Wood Mackenzie, a specialist consultancy. Thermal coal cost curves have been adjusted by the energy content of the coal. The curves identify the relative positions of the Hunter Valley mines.

Figure 2.1: Thermal and Metallurgical Coal Cost Curves



Source: Wood Mackenzie

Note: mine identification removed for publication

The Wood Mackenzie cost curve needs to be interpreted with caution. It shows the incremental cost of additional production, which of course can vary over the life of the mine, and would tend to decline with scale of production. For example, while the Aустar mine appears at the top end of the thermal cost curve, this is because of its current state of development (preparation of the long wall). Its incremental cost would decline substantially once full production begins.

In addition, the incremental production cost as measured by Wood Mackenzie tends to be higher than what an economist would regard as the SRMC of the mine. Given that there are

also incremental costs involved in putting a mine into care and maintenance, it is likely that mines would continue operating throughout a commodity price cycle even if the price was below the measured incremental cost for a period of time.

The bulk of Hunter Valley thermal coal reserves are located towards the lower end of the global cost curve for seaborne trade.

The table below shows the proportion of Hunter Valley production capacity in each cost quartile. As the table shows, if the global seaborne thermal coal trade were to fall by a half, 81.6 percent of Hunter Valley capacity can be expected to continue operating. Similarly, if the global seaborne metallurgical coal trade were to fall by half, 96.2 percent of Hunter Valley capacity can be expected to continue operating.

Table 2.2: Proportion of Incremental Annual Production Capacity in each Quartile of the Cost Curve

Hunter Valley - Seaborne Export Thermal Coal

Quartile	Incr. Production (Mt)	Share Of Total (%)
1	93.5	63.4%
2	26.8	18.2%
3	21.0	14.2%
4	6.2	4.2%
	147.6	100.0%

Hunter Valley - Seaborne Export Metallurgical Coal

Quartile	Incr. Production (Mt)	Share Of Total (%)
1	15.1	72.1%
2	5.3	25.1%
3	0.5	2.4%
4	0.1	0.4%
	20.9	100.0%

Source: Wood Mackenzie

The cost curve analysis is consistent with the actual recent market experience. For example, [] which has been put into care and maintenance earlier this year is located, as expected, at the tail end of the cost curve. Of course, in the short term, various other factors in addition to a mine's position on the global cost curve can play a role, including the financial health and risk appetite of the current owner. However, the key conclusion is that a catastrophic decline in coal volume from the Hunter Valley of the scale necessary to undermine the ability of the remaining users to cover the full economic costs of ARTC would require more than a sustained softness in coal prices and lack of growth in demand. In fact, it would require a fundamental collapse in the demand for seaborne coal, with volumes falling by more than half.

Such a collapse in demand could not be driven by cyclical factors. Rather, it would require a fundamental shift in the global use of primary energy, with rapid substitution to alternative energy sources. As our analysis above shows, at the current already very high rates of growth in renewable energy, total global coal generation capacity is still likely to continue to grow. In other words, a significant decline in the global demand for coal over the next 20 years would require the LRMC of renewables to fall below the SRMC of coal-powered generation so that it is economic to retire the existing coal capacity before the end of its economic life. In addition, the renewable sector would need to be physically able to install sufficient additional capacity to replace early retirement of coal.

3 Finance Theory

As explained above, the exposure that ARTC has to single commodity risk is limited to a very low probability event in which the price of coal decreases to a point where long-term mine closures make it impossible for the remaining lives to cover the total revenue requirement. So, what guidance can we have from the corporate finance literature on how that risk should be reflected in the calculation of WACC?

The usual tool for the calculation of WACC—the Capital Asset Pricing Model (CAPM)—deals with the opportunity cost of capital within a diversified portfolio. Investors can eliminate company-specific risk by properly diversifying their portfolios. Investor is rewarded with higher returns for bearing systematic market-related risk.

While there are several competing risk and return models in finance, they all share some common views about risk. First, they all define risk in terms of variance in actual returns around an expected return; thus, an investment is riskless when actual returns are always equal to the expected return. Second, they argue that risk has to be measured from the perspective of the marginal investor in an asset, and that this marginal investor is well diversified. Therefore, the argument goes, it is only the risk that an investment adds on to a diversified portfolio that should be measured and compensated. In fact, it is this view of risk that leads us to break the risk in any investment into two components. There is a firm-specific component that measures risk that relates only to that investment or to a few investments like it, and a market component that contains risk that affects a large subset or all investments. It is the latter risk that is not diversifiable and should be rewarded.¹

3.1 Asset Beta

It is useful to recap that beta measures the tendency of the return of a security to move in parallel with the return of the equity market as a whole. If ARTC and other rail track providers in Australia were listed companies, estimates of their betas from the observed data would enable us to establish the correlation between their returns and the total market returns. Over the years, we would have expected that rail track returns may have been affected by fundamental shifts in demand of the type that could occur if coal volumes collapse: fundamental changes in how both passengers and goods move around Australia. Hence, if we were able to measure ARTC beta from direct market observations, we would likely correctly

¹ Damodaran, A. *Equity Risk Premiums (ERP): Determinants, Estimation and Implications—the 2015 Edition*, Updated March 2015, Retrieved from <http://pages.stern.nyu.edu/~adamodar/>,

measure the effects of periodic dramatic changes in dominant commodities on the variability of returns relative to the market.

Conceptually, there would be no reason to expect that past dramatic changes in how people and goods are transported would have been correlated with the systematic market risks. In fact, it is quite likely that past declines in rail utilization would have occurred during periods of strong economic growth. Similarly, in the future, the factors that may drive a hypothetical catastrophic decline in coal volumes—such as a dramatic technological breakthrough reducing the costs and increasing potential speed of deployment of alternative energy—would likely be correlated with strong economic performance.

Since ARTC's actual beta cannot be observed, we need to consider whether our estimates of beta—Castalia has suggested that ARTC's asset beta should be around 0.4—are consistent with the effects of single commodity risk on the variability of ARTC returns relative to the market. In the absence of direct observation, betas for entities such as ARTC are usually estimated by reference to:

- Similar businesses with observable betas (such as US railroads)
- Businesses with similar profiles of cash flows (such as regulated energy companies).

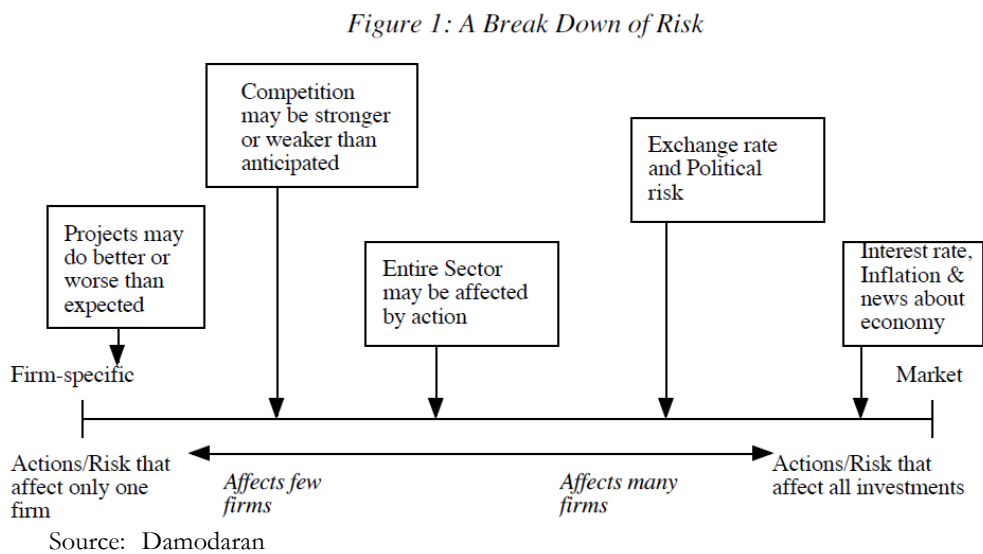
Since there have been dramatic shifts in commodity demands and transportation needs over time, any empirically observed proxies for the estimation of ARTC beta would have already captured the effects of single commodity risk on both the systematic market risk and on the covariance of rail track returns with that risk. In other words, it is unlikely that any adjustment would be needed to those empirical observations to proxy beta.

Similarly, there is no obvious reason why we would expect that the covariance between market returns and their own returns expected for other regulated entities would need to be adjusted for single commodity risk. All regulated entities potentially face the risk of low probability, high impact shifts in demand. For example, regulated electricity networks are highly alive to the risk posed to their business model by distributed generation.

Overall, in the corporate finance literature, there is strong academic consensus that the kind of rare event risk should not, as a rule, be expected to affect the beta.

Figure 3.1 shows examples of risks that affect firms and those that affect the market. The risks that affect the firm (left side of the figure) are diversifiable. According to finance theory, these risks should not be included in the calculations of the cost of equity because marginal investors will not be exposed to them. The rare event of mines shutting down is sector specific, and as such is diversifiable. An investor can hold shares in different sectors and eliminate this risk.

Figure 3.1: A Break-Down of Risk



A highly regarded Australian practitioner emphasizes that:

Businesses may face exposure to other business risks arising from low-probability events that may have a large impact on the cash flows of the business. Examples of such events may include: major or catastrophic losses arising from natural disasters (e.g. bush fires, storms, floods) or damage caused by third parties (e.g. terrorist attacks); fraud losses; and losses arising from liability for damages to third parties (e.g. liability for damage to the environment, or arising through litigation). Most of these types of risks should be considered business or industry specific and, therefore, largely diversifiable from the perspective of the marginal investor.²

3.2 Market Risk Premium

In an economy such as Australia's, with its dependence on primary commodities, significant shifts in commodity demand and prices are likely to have an effect on the non-diversifiable (systematic) portion of risk. The key question, then is whether a potential future move away from coal is likely to be already taken into account in the measured market risk premium (the return over the risk free rate that investors demand).

While the theory is concerned with the premium for future systematic risk, in practice we measure the market risk premium using historical data. So, we need to as whether such historical measurements may need to be adjusted for single commodity risk.

All companies and sectors, and in fact, the market itself, have faced catastrophic risks in the past—various low probability events, such as new inventions, rapid technological and demand shifts, wars and natural disasters. Financial theory suggests that as long as the period over which MRP is calculated includes such rare events, then the observed MRP is a fair measure

² Frontier Economics, a report prepared for the AER, *assessing risk when determining the appropriate rate of return for regulated energy networks in Australia*, July 2013

of the market's willingness to accept the non-diversifiable risk of such low probability high impact events.

There are relatively few studies of the effect of catastrophic events on systematic risk. Kadan Liu and Liu³ consider systematic risk (contribution of an asset to the total risk of the portfolio) in a setting that captures rare disasters. They conclude that the theoretical approach used in the Capital Asset Pricing Model can be generalized to include a such rare disasters.

Other studies have tried to measure the empirical effects of rare disasters. For example:

In Barro (2006) ... rare disasters are identified as large, instantaneous, and permanent drops in output. He calibrates the frequency and permanent impact of disasters to match large peak-to-trough drops in real per-capita GDP in a long-term panel dataset for 35 countries and shows that his model is able to match the observed equity premium... More recently, Barro and Ursua (2008) have gathered a long-term data set for personal consumer expenditure in over 20 countries and shown that the same conclusions hold using these data. A growing literature has adopted this model and calibration of permanent, instantaneous disasters (e.g., Wachter, 2008; Gabaix, 2008; Farhi and Gabaix, 2008; Burnside, et al., 2008; Guo, 2007; and Gourio, 2010).⁴

In other words, historical estimates of MRP do not appear to vary depending on whether rare events are specifically considered. This suggests that the existing estimates of MRP already factor in the effect of catastrophic events. As a practical point, for example, Damodaran argues that catastrophic events are one of the determinants of risk premiums⁵, yet when estimating cost of equity, he does not include an extra premium for catastrophic risk, thus assuming that this risk is already compensated in the risk premium.

More generally, market risk premiums calculated using historical averages should be expected to include the effect of periodic of catastrophic events. The historical approach is the approach used by most practitioners.

Most analysts, investors and managers, when asked to estimate risk premiums, look at historical data. In fact, the most widely used approach to estimating equity risk premiums is the historical premium approach, where the actual returns earned on stocks over a long time period is estimated, and compared to the actual returns earned on a default-free (usually government security).

In Australia, in estimating market risk premiums, regulators have used historical information for periods which include catastrophic events. For example, in the NSW State Water decision, the ACCC used the historical excess returns estimated by the AER (see Table 3.1 below), and market surveys.⁶ The historical excess returns use information from years going back to 1883. This would have undoubtedly included various significant rate events.

³ Kadan, Ohad and Liu, Fang and Liu, Suying, Generalized Systematic Risk (October 29, 2014). Available at SSRN: <http://ssrn.com/abstract=2444039> or <http://dx.doi.org/10.2139/ssrn.2444039>

⁴ Barro, R., E. Nakamura, J. Steinsson and J. Ursua, 2009, Crises and Recoveries in an Empirical Model of Consumption Disasters, Working Paper, http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1594554

⁵ Damodaran, A. *Equity Risk Premiums (ERP): Determinants, Estimation and Implications—the 2015 Edition*, Updated March 2015, Retrieved from <http://pages.stern.nyu.edu/~adamodar/>

⁶ ACCC, *Attachments to ACCC Final Decision on State Water Pricing Application 2014–15 — 2016–17*, June 2014, Page 88

Table 3.1: AER’s Estimates of Historical Excess Returns

Table D.1 Historical excess return estimates—assuming a use rate of distributed imputation credits of 0.7 (per cent)

Sampling period	Arithmetic mean	Geometric mean
1883–2012	6.3	4.8
1937–2012	5.9	3.9
1958–2012	6.4	3.8
1980–2012	6.3	3.6
1988–2012	5.7	3.6

Source: NERA, AER analysis.

Source: AER, *Explanatory Statement Rate of Return Guideline (Appendices)*, December 013

In conclusion, it appears reasonable to expect that 6.0 percent MRP already takes into account the reward that investors require for periodic rare events (as well as other systematic risks) and that no further risk premium is needed.

4 Conclusions

On the face of it, the current state of the coal industry suggests that ARTC may be facing heightened risks. This impression may be reinforced to some extent by the recent announcement by Moody’s of the review of the ARTC credit ratings for possible downgrade. ARTC have pointed to sector risks as a justification for higher beta.

The analysis in this paper shows that broad volatility in the coal sector and the financial strain faced by the existing mine owners are not relevant. Rather, the relevant risk for ARTC is of the coal price persistently settling at the level that leads to such a loss of volume that ARTC’s revenue requirement cannot be sustained. Moody’s announcement specifically describes the risk of such a downward spiral as the key issue for the ratings review.

Our analysis suggests that the probability of such an event is very low. Moody’s credit rating confirms this. Even a two to three notch rating downgrade compared to the current position would still keep the ARTC rating at above the regulatory benchmark rating used for regulated power companies (after also removing the uplift for Government-related issuers).

More importantly, a change in the credit rating, even if it were significant, does not mean that the approach to the cost of equity should be similarly adjusted. Unlike the cost of debt, where the risk premium regularly changes, the cost of equity is usually calculated using long-term estimates of MRP and beta. This is because there is no reliable way to observe forward looking values for these variables.

There is no theoretical or practical reason to expect that the current estimates of beta and MRP should be changed as a result of single commodity risk. The existing estimates already capture all relevant risk information, including the risk of rare events.