Storage Options for the New Zealand Electricity Sector

Operational and Organisational Issues

Prepared for the

Ministry of Business, Innovation and Employment

by

E Grant Read

FINAL

V4.7

14 July 2022

DISCLAIMER

EGR Consulting Ltd

Executive Summary

1. This report has been prepared at the request of MBIE, as a contribution towards developing a comprehensive framework for understanding and assessing options for managing a large-scale storage facility, and integrating such a development into the New Zealand electricity market.
2. Much of the material presented here is new, and relates to possible future situations of which we have no direct experience, and have not yet been thoroughly researched. Accordingly, our goal has been merely to identify issues that are likely to need attention, and directions that are likely to be productive. So, all "conclusions" presented here should be regarded as tentative, and subject to further discussion, verification, and revision.
3. We focus mainly on the management and integration of pumped storage hydro because pumped storage hydro developments are being actively investigated, and because they raise more complex management issues than other options of which we are aware.
4. The report does not attempt to cover any issues relating to the organisational form, structure, ownership, or governance, of any of the hypothetical entities discussed in the context of the options explored. Nor is any consideration given to transitional or establishment issues of any kind.
5. The report does outline a wide range of structural options, though, focussing on the decision-making arrangements implicit in each, and on how a storage facility would interact with the spot/hedging market, and with market participants, the System Operator, and any host system manager.
6. It also provides a preliminary assessment of which of those options seems most promising for further development, while recognising that the preferred option will ultimately depend on the location, scale and nature of any development that might proceed.
7. Chapter 1 merely outlines the context of this report, the history of the theory underlying it, and the structure of the report itself.
8. Chapter 2 summarises the general principles of storage management, as developed in a much more comprehensive report recently released by MDAG, while highlighting aspects of particular relevance to the management of storage facilities, and surveying storage developments that might prove attractive in the projected environment.
9. Chapter 3 applies that theory to the management, analysis, and organisation of "stand-alone" storage facilities, including, but not limited to, batteries and pumped storage.
10. Chapter 4 extends that work to provide a more detailed framework for understanding the optimal operation of pumped storage hydro facilities embedded in host system catchments that may have their own storage /generation facilities, the opportunity costing of energy stored in such facilities, and the kind of situation in which conflicts might arise, and coordination might be required. The situation is potentially complex, and highly dependent on the configuration of the host system pumped storage, and the positioning of the pumped storage facility within that system. Broadly, though:
    - If the host system is operating on a run-of-river basis, then it has no control over flows, so the issue is not so much one of "coordination" but of controlling pumped storage operations to live within the available downstream flow freeboard/headroom.
    - In that case, incentives are broadly aligned, and the pumped storage manager can ignore upstream generation capacity, but ignoring downstream generation capacity could reduce national benefit, to a greater or lesser extent, depending on the relative release productivity ("head") and utilisable flow rates of the pumped storage and downstream generation stations. So, a coordinating agreement may be desirable, to extract the potential national benefit from the pumped storage development.
    - Downstream host system storage provides greater control of downstream flows, and leaves the operational objectives of pumped storage and host systems fairly well aligned, suggesting that a coordinating agreement is less likely to be required, or problematic.

- Upstream host system storage creates diametrically opposed operational objectives for the pumped storage and host systems, though, suggesting that a coordinating agreement may be both necessary to extract the potential national benefit from the pumped storage development, and problematic.

11. That chapter is complemented by four Appendices containing more detailed discussion of host system conflict and coordination issues.

12. Appendix A discusses examples where coordination of host system interactions might be more, or less, complex and /or critical:
    - A key conclusion is that, from a national benefit perspective, the economic priorities of pumped storage operations must dominate those of the host system, if the "head" of the pumped storage facility is greater than that of the combined upstream host system and/or its storage capacity is greater, and that significant impacts are likely, even in less extreme cases.
    - A national benefit optimisation is therefore likely to recommend operating the upstream host system in ways that extract maximum benefit from the pumped storage facility, at the expense of any direct profitability from upstream generation.
    - Accordingly, the theoretical national benefit contribution of enlarging the lower reservoir from which pumped storage draws and to which it discharges, to create more "buffer storage" is not so much to increase pumped storage flexibility, but to allow upstream generation to still contribute to national welfare by persisting with something more like its normal profit-driven operational pattern.
    - The optimal trade-off between the explicit cost of developing extra buffer storage and the implicit cost of reducing competition by requiring more coordinated operations will depend on the specifics of particular proposals.
    - The need for an agreement to deal with these interactions could become a key factor determining the likely desirability, and optimal design, of some potential developments, and that complexity could, in turn, be a key consideration limiting the range of organisational options that might be contemplated, in those cases.
    - But this may be only a minor issue for other potential developments, particularly if the flow freeboard/headroom and/or buffer storage can easily be made large enough, relative to the upstream host system reservoir storage.

13. Appendix B discusses the kind of agreement that could be used to resolve conflicts and coordinate operations, where necessary.

14. Appendix C discusses the water and/or energy trading concepts that might be employed in forming such agreements.

15. Appendix D discusses the interpretation of modelling results, and particularly optimised recommendations about buffer storage requirements, in light of the fact that optimisation models normally assume that perfect coordination can and will be achieved.

16. Chapter 5 then moves on to consider the various ways in which a pumped storage facility could be used to provide hedging to market participants, as well as hedge its own energy cost risk:
    - The facility would not be a net generator, but could buy night/summer energy contracts and sell day/winter energy contracts, reflecting its role in managing regular cycles. And it could also buy put options, and sell call options to create a loose "collar" on market prices, reflecting its role in limiting both short and long-term price volatility.
    - Allowing a single organisation of the size contemplated to trade in these markets would exacerbate any market power concerns, raising the prospect of more complex regulatory requirements.
    - The essence of what storage capacity can contribute to the system is really inter-temporal arbitrage, though, and "Financial Storage Rights" have been proposed to represent that capability. But those would effectively require parties to commit to "transferring energy" from one specific period to another, whereas a major role of the proposed a facility would be storing energy for use in crises, the timing of which can not be predicted.

- Thus. we suggest the natural form of hedging provided by the facility would really be “virtual capacity” contracts that would allow holders to make their own decisions about when the corresponding pumping/generation capacity should be called upon.
- Alternatively, participants might prefer what we call “tank option” arrangements, under which some other party takes responsibility for purchasing and storing energy, leaving them to purchase, and ultimately call on, that energy, within specified capacity limits, when they need to.

17. Chapter 6 then provides the core contribution of this report, which is to outline and assess a fairly wide range of alternative operational and organisational arrangements for a general large-scale pumped storage hydro facility in three broad groups:
    - “Unified” arrangements under which rules or models are used to guide and constrain the behaviour of an entity that would otherwise be in a position to exercise too much market power, or to disrupt the electricity market in potentially unpredictable ways.
    - “Diversified” arrangements designed to allow a variety of electricity market participants to effectively determine the dispatch of their allocated shares of the facility, in ways that do not increase the market power issues already inherent in the status quo, and may possibly reduce them.
    - “Hybrid/mixed” arrangements which retain the concept of “diversified” management of storage and generation, but could allow a lightly regulated “unified” approach to the management of all or part of the pumping capacity to build up stored energy that would then be available to be purchased and managed by the holders of simplified “tank options”.
18. Finally, Chapter 7 presents our conclusions. We do not think it appropriate to make a definitive recommendation with respect to the choice between unified and diversified/hybrid organisational arrangements at this stage, and realise that that decision may depend on wider issues then those considered here, but we have summarised our views on what we believe to be the most promising unified and diversified regimes, at this stage.
19. If a “unified” organisational approach is to be taken, with the pumped storage facility manager responsible for pumping/storage/generation/hedging strategy, we would favour an operating mandate based on the application of a national cost-benefit optimisation model, with parameters agreed by industry experts.
    - We consider that, while this might not provide so much short-term certainty around physical/market behaviour, it would actually provide greater “clarity of objective” than “rule-based” alternatives, and prove more predictable and durable, in the long run.
20. Alternatively, we suggest that serious consideration should be given to a “mixed” regime under which some “virtual capacity” rights are owned and managed by larger market participants, but the facility manager, or some other entity, retains and manages rights to some pumping capacity, and uses it to store energy for sale to smaller participants holding “tank” storage/generation capacity rights.
    - We believe that would allow market participants to acquire a virtual pumped storage facility, customised to complement their physical portfolio, and then to manage that virtual facility as part of their portfolio, and to support their offerings in more conventional hedge markets.
    - That approach would take market power away from the pumped storage facility, diversify decision-making with respect to storage strategy, and enhance competition in the broader market.
    - It remains to be seen whether planned developments would be large enough to make all aspects of that “mixed” package viable, or desirable, but we suggest that it provides a comprehensive enough template to allow research, development, and sectoral discussion to proceed, in the meantime.

Contents

Storage Options for the New Zealand Electricity Sector

Operational and Organisational Issues

1. Introduction

Context

This paper has been prepared at the request of the New Zealand Ministry of Business, Innovation and Employment (MBIE), with a view to informing discussion relating to the "NZ Battery" project. The motivations of that investigation, and key options being considered, are well known, and do not need rehearsing here. At the same time the Market Development Advisory Group (MDAG), has been conducting a parallel investigation into how the New Zealand electricity market might operate in the emerging environment without any substantial addition to national storage capacity, and released its issues discussion paper in February 2022.[1]

Although we did not participate in MDAG's investigation, we were asked to prepare a paper that has been released along with their report, on the application and implications of storage management theory, and particularly on the opportunity costing water, or other storables, in the context of the kind of market conditions MDAG expect to have developed by 2030, and then by 2050.[2] The discussion in that paper draws, in turn, on an earlier overview of electricity market economics, prepared for a generators' consortium, and released in 2018.[3] In particular, Appendix C of that earlier report outlined the theoretical concepts underlying the discussion in our paper for MDAG, and in the current report.

History

The conclusion of our report to MDAG was that, while the balance of generating capacity in the power system will be changing, the classical theory outlined in earlier papers has not changed. Indeed, much of it is rooted in analyses of the optimal management of electricity systems, going back as far as the work of Massé in the 1940's.[4] All of that earlier work was developed in a public sector context, assuming a framework of centralised national cost-benefit optimisation. However, the central concepts emerging from that work, over many decades, focus on economic interpretation of the "shadow prices",[5] calculated by the optimisation process as its measure of the marginal value of each resource utilised in the optimal solution.[6]

---

1. *Price discovery under 100% renewable electricity supply: Issues discussion paper*, MDAG report, February 2022, https://www.ea.govt.nz/assets/dms-assets/29/01-100-Renewable-Electricity-Supply-MDAG-Issues-Discussion-Paper-1341719-v2.4.pdf
2. *Opportunity Costing in the NZEM: Implications of Decarbonisation*, EGR Consulting Ltd report to MDAG, January 2022, https://www.ea.govt.nz/assets/dms-assets/29/05-Water-Values-under-100-Renewable-Electricity-Dr-Grant-Read1341584-v2.1.pdf
3. *An Economic Perspective on the New Zealand Electricity Market*, Prepared by EGR Consulting Ltd for a broad generator consortium, and submitted by Meridian in response to MBIE's *Electricity Price Review: First Report*, October 2018 https: //www.mbie.govt.nz/dmsdocument/4195-meridian-energy-electricity-price-review-first-report-submission
4. P Massé Les Réserves et la régulation de l'avenir dans la vie économique. Hermann & Cie. Paris, 1946.
5. Also known as "Lagrange multipliers", "dual prices", or "dual variables".
6. Here "resources" include anything for which a "constraint" is defined in the formulation of the optimisation problem, because it is recognised as being potentially in short supply. Those constraints may limit concrete observable "commodities". like "water in the reservoir", but also more abstract concepts like 'the ability to release extra water down the river at an environmentally acceptable rate".

At the time, that economic interpretation was partly a tool allowing “decomposition” of sectoral planning/operating optimisation problems that were too big to be solved by the technology available at the time, and partly a way of developing an intuitively comprehensible theoretical framework for understanding and managing these inherently complex processes. It has long been realised, though, that exactly the same mathematics, and conceptual understanding, would apply in a hypothetical perfectly competitive market context.

The major development of the last 30 years, then, has been to create actual markets based on these concepts, even when perfectly competitive assumptions can not be fully satisfied. Discussion of the strengths and potential weaknesses of that development lie outside our present scope. But the key point is that while allowances must be made for the various ways in which the New Zealand electricity market environment may differ from the perfectly competitive ideal, the same decades old theory still provides an intuitively comprehensible theoretical framework for understanding and managing storages, in the market environment as it is now, and as it will be in future.

Outline

Accordingly, the basic structure of this report is to:

- Outline, in Chapter 2, the basic theory of storage management in a national optimisation or perfectly competitive market environment, while highlighting aspects of particular relevance to the management of pumped storage hydro facilities;
- Address, in Chapter 3, the application of that theory to the management and analysis of “stand-alone” storage options that were not considered closely in our report to MDAG;
- Focus, in Chapter 4, specifically on the application of that theory to the management and analysis of hypothetical pumped storage hydro facilities embedded in host systems, because it seems to us that the integration of this technology into the market, and into an existing host system could pose unique and complex changes;
- Develop, in Chapter 5, a range of hedging strategies and instruments by which the pumped storage facility could meet the risk management needs of market participants and end-consumers, and of itself;
- Consider, in Chapter 6, a range of organisational options and operational regimes that have been proposed for managing pumped storage hydro, mainly to deal with the problem that a facility of the size currently being envisaged could definitely not operate on a “perfectly competitive” basis in the small New Zealand market; and finally
- Summarise our conclusions, in Chapter 7.

Scope

This report is preliminary, in the sense that it identifies quite a number of potentially important topics and issues that it does not attempt to resolve. It also focusses mainly on pumped storage hydro, because that is an important option that is reasonably well understood, having been long established elsewhere, and which clearly needs to be analysed. We believe the framework provided here can be simplified and/or extended to cover a wide range of other options though.

Much of the material presented here is new, and should be regarded as tentative, and subject to revision. No attempt has yet been made to formulate, solve, or analyse formal mathematical optimisation models as a means of validating the various propositions advanced in the text, or to apply them to any specific proposal. While some of the discussion could provide a basis for developing a quantitative analytical framework for pumped storage hydro, or other similar projects, no attempt has been made to develop such a framework or perform such analyses.

The report also makes no attempt to cover any issues relating to the organisational form, structure, ownership, or governance, of any of the hypothetical entities discussed in the context of the options explored. Nor is any consideration given to transitional or establishment issues of any kind, including the development of rules, charters, or statements of corporate intent, even at a high level.

Note on terminology

All references to electricity prices in this report should be interpreted as referring to half-hourly spot prices in a perfectly competitive wholesale market, and/or to the equivalent "shadow prices" that would theoretically emerge from an economic dispatch optimisation, in a centrally planned environment. While we do discuss organisational arrangements designed to limit the exercise of market power by the manager of any new pumped storage hydro development, we do not address the implications of potential discrepancies between real market prices and the market prices that would theoretically emerge from the perfectly competitive market implicitly assumed by most discussions in this report.

Many analysts and industry discussions, including our own, refer to a "Marginal Water Value" or "Marginal Value of Water" (MWV or MVW) as if it was defined directly in $/MWh terms. For a simple hydro system, we might also call this the Marginal Value of Storage (MVS), and our report to MDAG, used that term to refer to the marginal value of water that could be stored in a reservoir, as distinct from the "Marginal Cost of Release" (MCR), which will generally equal MVS, for a simple hydro system, but could be zero for flows that have to be spilled, even when MVS is strongly positive.

The advantage of defining these concepts in energy terms is that they can then be compared directly with electricity prices. For that reason, we will sometimes refer to the Marginal Cost of Generation (MCG), and Marginal Value of Pumping (MVP) for pumped hydro systems, even though those may not be simply set by a single marginal water/energy value, in that more complex setting. This report will also frequently refer to Marginal Value of Stored Energy (MVSE) in contexts where the term could apply in very similar ways to all storage technologies.

However, parts of this report focus more on understanding operations and valuations in river chains, and particularly in pumped storage hydro schemes. For those purposes, it will be necessary to differentiate between water value concepts more carefully than in our report to MDAG. So, we will use the term Marginal Water Value (MWV) to denote marginal water values defined in $/unit of water volume, that being the most likely marginal value to be reported by a formal mathematical optimisation detailed enough to be concerned with effects like the dependence of pumping/generation efficiency on flow rates, or head.[7] As discussed in Appendix C, though, what is really valued in these contexts is not the water, as such, but the potential energy stored by holding that water at a certain location in a catchment, from whence it will be released to generate while passing through one or more hydro power stations, before eventually reaching sea level.

Finally, in most contexts the relevant driver of managerial behaviour should be an expected value, which we will denote by EMWV/EMVSE, except when discussing values determined in a hypothetical deterministic context, or for a particular simulated scenario.

---

[7] The most obvious "unit of water volume" here would be a cubic meter, which corresponds to a metric tonne of water. But that would imply prices defined in $/Kilojoule, whereas the electricity sector generally prefers to define energy in kWh, or MWh. Because there are 3600 seconds in an hour, 1 kWh equals 3,600 Kilojoules, and 1 MWh equals 3,600,000 Kilojoules So that conversion factor would need to be applied, if we wanted to compare MWV directly with electricity prices quoted in $/MWh. Or, reversing that conversion, the "water volume unit" that, when raised by 1m, would store 1 MWh of potential energy, is actually 360,000 cubic meters, with the gravitational constant g approximated to 10. But, except in Appendix A, we will generally just refer to a "unit of water volume", or "water unit", because the precise units do not really matter for the purposes of the conceptual discussions in this report.

2. Managing Storage in the Future Electricity Sector

2.1. Introduction

The MDAG report explores and simulates what might be described as a “business as usual” scenario for electricity market operation, under which:

- All conventional fossil-fuelled thermal generation would be retired;[8]
- That capacity would be replaced by a mix of new geothermal, solar and wind plant, with substantial further capacity development to cover substantial load growth, as electricity takes over from fossil fuels, across the economy; and
- That geothermal/solar/wind capacity might be supplemented by some “green peaking thermal” capacity, perhaps fuelled by biodiesel or hydrogen; but
- Importantly, no new hydro capacity, including no pumped storage hydro, would be added; and
- The market would continue to operate as it does at present.

MDAG concluded that, while this rapid re-orientation would obviously imply major challenges for the sector, those challenges could be met by:

- Significant development of various forms of “Demand Side Management” (DSM); and
- Significant investment in battery capacity.

Still, hydro storages, and river chains, might need to operate quite differently than they have in the past, and it would be highly desirable to augment their capacity by expanding less conventional “storage” options, such as large-scale pumped storage hydro, or demand side flexibility, to assist with limiting:

- Cyclic day-night price differentials;
- Cyclic summer-winter price differentials; and
- Random price volatility, over time frames ranging from hours, to days, weeks, and years.

Our report to MDAG complemented that study, by explaining the theory of how any of these storage options should be operated, and valued, first from the perspective of national cost-benefit optimisation, which we believe to be theoretically equivalent to a perfectly competitive market, and then in a more realistic market environment. That report focussed mainly on management of conventional hydro storage, though, and only briefly touched on the optimal operation and valuation of less conventional storage options, such as demand side inventories, batteries, or pumped storage hydro.

Thus, the intent of this current report is to flesh out that discussion, particularly with respect to pumped storage hydro, with a view to developing an understanding of, and possible management regimes for the kind of large-scale developments currently being considered by MBIE. Before doing so, though, we should outline the general theory, and highlight some aspects of particular relevance to the management of less conventional large-scale storage options, including demand side flexibility and pumped storage hydro.

8 References to “thermal”, in this report should be taken as referring to “fossil-fuelled” thermal, unless explicitly stated otherwise.

2.2. Basic concepts

Chapter 2 of our report to MDAG discusses the basic theory of storage management, as it would apply to a single stand-alone stockpile of any kind, and as it has traditionally been applied to long term reservoir management in the New Zealand electricity sector.

The core concept is that storage capacity provides a physical means of "arbitraging" across time. That is, it allows units of the stored commodity to be carried forward from periods when they have a low marginal value, because they are in relatively good supply, to periods when they have a higher marginal value, because they are in relatively poor supply. In the New Zealand electricity sector, this has traditionally meant carrying water forward from summer to winter, and from night to day.

In that context, managers are constantly deciding how much water to release now, and how much to save for future use. For that purpose, it is useful to define the Marginal Value of Stored Energy, MVSE, as the "opportunity cost" of not having a unit of "potential energy" (in the form of water stored at height), available to use at some future date as a result of releasing it now.[9] We can think of the arbitrage process as proceeding by working through the units of water currently available, and identifying the best <u>remaining</u> opportunity to utilise each successive unit of water stored:

- At first, units would be assigned to the future periods in which the supply/demand balance was most critical, as determined by their having the highest prices, in either a national cost-benefit optimisation or perfectly competitive market context.
- But, as more and more water is put aside for release to support generation in those periods, the prices in those periods will gradually drop, and/or generation from our reservoir will reach its upper limit.
- Either way, later increments of stored water will be progressively assigned to other periods, in which they have progressively lower marginal values.
- Eventually, we will run out of water or, more likely, reach a point at which the next unit of water we could store would actually have no higher value than the last unit of water released to support current generation.

The "incremental arbitrage" process described here is just a way of illustrating the principles involved. Real optimisation models actually employ more direct methods to determine the "marginal value" of storing water for the future, and can equate it exactly to the marginal value of releasing water.[10] The key point is, though, that any optimisation algorithm will stop when it reaches the point at which it has identified the "marginal economic opportunity" for utilisation of stored water. And the MVSE is then set by the "opportunity cost" of not taking that opportunity.

Note that, even if there were no limits on storage capacity, and the future were completely known, the arbitrage process will not completely eliminate electricity price differences between periods, because:

- prices can not be lowered any further in periods with a relative excess of demand, once the upper limit on generation capacity is reached; and
- prices can not be raised any further in periods with a relative excess of supply, once the lower limit on generation capacity (possibly zero) is reached.

Also note that the biggest challenge in all this is that the future is not known, so there are many scenarios to consider, each with its own "marginal economic opportunity" and opportunity cost, MVSE. It can be shown, though, that ignoring the possibility of risk aversion, the optimal marginal water value, on which decisions should be based, is EMVSE: That is, the expected value of all those MVSE values.

Starting from Chapter 2 of our report to MDAG, we emphasise that much of what is commonly described as electricity price "volatility" is really predictable "cyclic variation". So, before tackling the mathematical and conceptual challenges of stochastic optimisation, it is helpful to first examine reservoir management from a "deterministic" perspective, focussed on maximising the value of reservoir storage

9 See *Note on Terminology*, in Chapter 1.

10 This assumes that the volume of water stored/released are continuous variables, with differentiable cost/benefit functions.

capacity to transfer water from low valued summer periods to high valued winter periods, and from low valued night-time periods to high-valued daytime periods.

"Deterministic cycling" in a single hydro reservoir[11]

Looking at reservoir management from a deterministic perspective, we argue that the common assertion that the marginal value of water must be zero when the reservoir is full, does not hold, if the reason the reservoir is full is because the manager (with perfect foresight) has decided it should be full so as to carry as much water as possible forward from a low value period to a higher valued one. In fact, in that idealised world:

- Stock levels should be expected to regularly cycle up, during periods of relative over-supply, and down during periods of relative under-supply, with the key cycles being daily, for small hydro reservoirs, and annual for large ones.[12]
- Since reservoir capacity is costly, we do not expect reservoirs to be built so large as to be able to fully arbitrage away all marginal value differences, in either daily or annual cycles.
- Thus, under deterministic assumptions, optimal reservoir management must involve holding the reservoir approximately full (apart from minor daily cycling) for some time at the end of each filling cycle, and approximately empty (apart from minor daily cycling) for some time, at the end of each emptying cycle.
- For an annual reservoir, we should expect to see this pattern when a deterministic analysis is performed on the assumption of expected inflows, and probably for most individual hydrological years, with the length of time optimally spent in the full and empty states depending on how different the supply/demand balance is in summer vs winter, in each scenario.
- From this perspective the marginal value of water/energy in storage (both MWV and MVSE) would be constant, at a lower level during the summer/night period, then rise over the period when the storage is full to a higher level, which is then maintained over the day/winter period, before falling again over the empty period.
- While at their storage bounds, reservoirs would effectively be operating in run-of-river mode, albeit with some daily cycling, but non-supply and spill are both very unlikely, with perfect foresight.
- And the day/night situation will be similar for most small reservoirs, on most days, although some mid-sized storages may not need to reach their storage limits on some days.

Managing stochasticity in a single hydro reservoir [13]

That perspective is clearly unrealistic, on its own, because there are many sources of true volatility in the sector, making a stochastic perspective important, too. From a "purely stochastic" perspective (i.e., assuming that flows and market requirements are purely random, with no predictable average patterns), the main function of any storage facility is to buffer the effects of random fluctuations in the supply/demand balance. From that stochastic perspective:

- Since reservoir capacity is costly, we do not expect reservoirs to be built so large as to be able to fully absorb all such fluctuations, so stock levels should be expected to reach both full and empty bounds. So, we may see spill, and may also see non-supply if (but only if) the reservoir is critical for national supply security.
- Storage deviations should be managed to revert, when possible, to levels far enough below the upper limit to allow high flows to be captured when they randomly occur and, more importantly (for a critical reservoir), far enough above the lower limit to allow non-supply to be avoided when flows randomly drop off.
- The Marginal Cost of Release (MCR) will obviously be zero, when spill becomes inevitable, and that may happen even before the reservoir is full.

---

11 The summary here is taken from the Executive Summary from our report to MDAG, with only minor changes.

12 We will acknowledge, but largely ignore, weekly cycles, because they do not add much insight.

13 The summary here is taken from the Executive Summary from our report to MDAG, with only minor changes.

- For a large annual reservoir “shortage” may mean national non-supply, and optimal reservoir management must favour high stock levels, with an inevitably increased probability of spill, over low stock levels, with an unacceptably high probability of expensive non-supply.
- For a small daily reservoir, though, “shortage” may just mean inability to take full advantage of a (probably moderate) unexpected intra-day price spike. So, it may well be optimal to use the water when prices are known to provide a reasonably valuable use for the water, rather than holding water back in case a higher spike occurs, but eventually only finding fewer valuable uses.

Synthesis[14]

Our report to MDAG discusses how the interplay between the deterministic and stochastic perspectives can imply a wide variety of outcomes, depending on the balance between energy capture, storage capacity, and utilisable release capacity, in different hydro systems. Thus, a wide variety of marginal water values should be expected, from different reservoirs, at the same time, and from the same reservoir at different times. In general, though, the Expected Marginal Value of Stored Energy (EMVSE) can be estimated by simulating realistic management (i.e., without perfect foresight) of a large number of hydrology sequences, and determining the (conditional) MVSE for each one, from some future "marginal economic opportunity" in that sequence.[15]

The marginal opportunities available vary greatly, depending on the reservoir, storage level, time of year or day, and scenario, but:

- There is no marginal economic opportunity available to release more water from a reservoir in periods when that reservoir should optimally be releasing at its maximum utilisable release rate, no matter how high the price may become. And there is no marginal economic opportunity available to release less in periods when that reservoir should optimally be releasing at minimum, no matter how low the price may become.[16] So, increased price volatility, implying prices further above/below the price levels at which maximum/minimum release becomes desirable in those periods, will have no impact at all on the marginal water value of that particular reservoir.
- There is no marginal economic opportunity available to store more water for release in periods beyond the next time when a reservoir is expected to be full in a particular scenario. And there is no economic opportunity available to store less water for release in periods beyond the next time when a reservoir is expected to be empty in a particular scenario, either.[17] So, we should expect to see significant cyclic variation in MVSE across seasons (for larger reservoirs), and across the day (for smaller storages).

As a result, water that can actually be held in storage might be given a high opportunity cost value (MVSE) if it seems at all likely to be required to avoid non-supply at some later date, even when excess (unstorable) water is being spilled from the same reservoir, at an implied MCR of zero.

Importantly, the need to hold storage levels away from bounds to deal with volatility reduces the effective capacity available to arbitrage between low and high valued periods. So, it actually increases the expected cyclic variation between day and night-time prices, and between summer and winter prices, thus increasing the importance of the fundamental insight derived from the deterministic analysis: Namely that the underlying MVSE must be rising while reservoirs are relatively full around autumn, and falling while reservoirs are relatively empty around spring.

---

[14] The summary here is taken from the executive summary from our report to MDAG, with only minor changes.

[15] This kind of “simulation” may not be explicit, but it is implicit, in some form, in all stochastic reservoir management optimisation packages.

[16] This discussion, like that in our report to MDAG, accounts for upper/lower flow limits, but ignores the possibility that limits may be imposed on the rate of flow change. If such limits are binding, they will act like upper bounds on flows in some periods, and lower bounds on flows in other periods.

[17] This discussion, like that in our report to MDAG, accounts for upper/lower storage limits, but ignores the possibility that limits may be imposed on the rate of storage volume change. If such limits are binding, they will act like upper bounds on storage volume in some periods, and lower bounds on storage volume in other periods.

Discounting, wastage, and head effects

Finally, we should mention three effects that were not discussed in our report to MDAG, but may have an impact on optimal operation of large-scale storage developments, including pumped storage hydro.

- First, stockpiled water, or energy of any kind, is implicitly subject to an "interest charge", just like any other commodity. That may seem strange, in cases were nothing was paid to fill the stockpile, but the arbitrage process described above always involves comparing the value of an opportunity to generate now, with the value of an opportunity to generate in some future period. So, a discount rate should really be applied, just as for any other such comparison.[18] As a result, marginal water values should not actually be constant, as in the simplified discussion above, but rising at the discount rate, over free trajectory arcs.
    - That should also induce a slight bias towards discharging water/energy sooner, if the alternative is leaving its inherent value idle (i.e., not "earning a return") over time. Again, this effect has generally been ignored in New Zealand discussions, because it will be cut off, and the cycle re-started, every time a storage capacity limit is reached.[19] But the effect would become more significant if water (or energy) were being stored for longer periods in larger scale storage facilities that approach capacity limits less often.
- Second, stockpiled water, or energy, is subject to wastage, just like any other commodity. The mathematical form of that wastage will differ between technologies. Batteries will slowly lose their charge, and water will slowly seep away and/or evaporate.[20] These effects are not generally significant in a daily cycling situation, and they have often been ignored as insignificant in longer term New Zealand reservoirs, too.[21] But it should be recognised that this form of "loss" has quite a different impact from the "round-trip loss factor ratios" discussed later.[22]
    - Basically, it induces a slight bias towards discharging water/energy sooner, if the alternative is letting it waste away, over time. It also implies that marginal water values must rise over time, as the slowly shrinking volume becomes more costly to provide, per unit, and scarcer. This effect is cut off, though, and the cycle re-started, every time a storage capacity limit is reached. It will become more significant if water (or energy) is being stored for longer periods, in larger scale storage facilities that approach capacity limits less often.
- Third, though, stockpiled water has one particular property not shared with other commodities, including "energy", namely the "head effect". The potential energy stored in water in a hydro reservoir is roughly proportional to its "head", that is the elevation of its surface above the point at which it will be discharged from any associated generation facility. So, the more we fill the lake

---

18 Mathematically, the issue is what difference adding a unit now will make to all the simulated trajectories from which EMVSE will be calculated. In some cases, the 'extra unit' may be held for a long time, and in others released fairly soon, thus letting the trajectory revert to the level it would have followed if there had been no extra unit.

When discounted back to the present, that marginal opportunity will set the current MVSE for that trajectory, and the average of all those MVSE values will form EMVSE. So, the higher the discount rate, and the further away the marginal release opportunities, the lower EMVS will be, and the less willing a rational manager will be to hold more energy in storage.

19 The extra unit will definitely have to be shed before the next time the storage trajectory reaches the full level, or otherwise it will be spilled at that time. And the extra unit will probably be used before the next time the storage trajectory reaches the empty level. Or, otherwise, it may be used to avert shortage at that time.

20 For hydro, wastage includes leakage, and some reservoirs become particularly leaky when certain geological levels are reached. In the limit, leakage may be great enough to match natural inflows, making it impossible to build storage up above that level, without pumping, and potentially even with pumping.

But evaporation is an issue, too. Although less severe in New Zealand than in warmer climates, it is not zero. The form of the effect depends on the shape of the reservoir bed. If the sides of the reservoir were vertical, then adding more water will not actually change the surface area, or the evaporation rates, so the incremental wastage due to adding one unit will actually be zero. Generally, the reservoir surface area will increase as incremental water is added, though, so there is some effect. But the rate of change, and hence the EMVSE impact, will vary non-linearly, depending on the surface area of the reservoir at each contour level.

21 With the notable exception of leakage at Waikaremoana.

22 This is just a particular application of "Hotelling's Rule", a long-established result in the economic literature. See: H. Hotelling: The Economics of Exhaustible Resources", *Journal of* Political Econ*omy*, Vol 39, Issue 2, p. 137–175, 1931.

the higher that surface level will be, and the more will be generated when each water unit is released. And that means that each unit of water stored in a reservoir contributes to increasing the productivity of other units being released, for as long as that unit is stored.[23]

- In principle, the effect of this is the opposite to that of wastage and discounting, inducing a bias towards discharging water later, so as to benefit from this head effect for longer. But the mathematical form of the effect is awkward (non-convex), and it only applies to generation facilities fed directly from reservoirs (rather than via an open river/channel). Also, it is only significant if the allowable range of surface level is significant. So, it has been ignored in most New Zealand discussions to date, because those conditions have seldom applied, but it could be material for large pumped hydro storage.

2.3. Traditional non-hydro storages

Thermal fuel stockpiles/trading

Sections 2.2.4 and 2.3.1 of our report to MDAG discuss the situation of thermal generators in New Zealand, and how "storage" in the thermal system has traditionally related to hydro storage. Since thermal generation is being phased out, there is no need to go into detail. But it is important to understand, in broad terms, how that sector has traditionally contributed to supplementing hydro storage, over various timeframes, and what capabilities might therefore need to be replaced.

In the case of gas, for example, there has always been long-term storage in gas fields which could be drawn down to a greater extent in dry years, to provide inter-annual storage not available in the New Zealand hydro system. And there has always been some very short-term storage in the form of overnight "line-pack" that could be drawn upon to allow extra gas-fired generation for a limited period, each day. Contractual arrangements may have complicated and obscured those realities, but the same opportunity costing logic applies to each of those limited storage resources, as to the hydro reservoir discussed above.

The coal supply system had its own explicit stockpiles, historically including pre-stripped open-cast coal seams, to which the same opportunity costing logic applied, but it also had the ability to replenish those stocks, via a supply chain that implied quite strict limitations on inter-temporal variability, each of which implied similar, but slightly different opportunity costing issues.

But the key point is that the "storage" capacity accessible through the thermal sector always went well beyond the obvious stockpiles being held, in three respects:

- First, it has also been possible, at least in principle, to purchase fuel that would have otherwise been used by other parties in the New Zealand economy. This may not seem like it is drawing on "storage", but in fact it often is. If those other parties are producing less of some product, it will sometimes be the case that consumption of that product actually falls, in the sense that end-users consume less at that time, and never make up the difference.[24] More often, though, the producer may draw down stockpiles of the product to meet customer demand, effectively forcing their customers to draw down their own stockpiles of that product, or of some end-product they would otherwise have produced from it. Or, it may be that end-consumes are not actually reducing ultimate consumption of whatever end-products are involved, but merely deferring it, which can be seen as a form of "storage" in itself.

---

[23] This effect applies irrespective of whether the water is drawn off the surface level, and thus "falls further" when released, or from a lower level where the energy is stored in the form of pressure from the weight of water above. The form of the effect depends on the shape of the reservoir bed, though. If the sides of the reservoir were vertical, then adding more water will raise head relatively quickly. Generally, though, the reservoir surface area will increase as incremental water is added, making it harder and harder to raise the surface level as the reservoir fills.

[24] This is more likely with services, like heating, than it is with commodities.

- Second, many New Zealand fuel users are directly involved in international networks trading the products they produce. So, when they reduce production of some product, New Zealand is either exporting less, or importing more, of that product, and the slack is taken up within that international trading network. Unless worldwide consumption actually falls, the New Zealand electricity sector is ultimately drawing on "storage" somewhere in that international trading system to deal with fluctuations in its supply/demand balance.
- Third, and more obviously, fuels like oil and coal are themselves internationally tradeable.[25] Thus an oil-fired generator in New Zealand does not actually need to hold a large stockpile locally, provided it can buy more from the international oil trading network, at short notice. When it does so, oil production must eventually increase somewhere, thus ultimately drawing down the planet's underground oil stockpile, but there is also a long chain of intervening oil storages, at ports and refineries, and in transit, that will be drawn down in the meantime.

Thus, it should be recognised that when we talk about "storage" in the thermal fuel supply sector, we are not just talking about the stocks held at New Zealand power stations, but storage in the whole vast trading network supplying and consuming those fuels, outside the New Zealand electricity sector. In principle, the same opportunity costing principles apply to all those storages as to the New Zealand hydro storages, but the application of those principles is mostly invisible to the New Zealand electricity sector, which must just accept whatever price deals can be negotiated with the fuel suppliers. And, whether the storage implications are visible or not, the end effect is that "tradability" of fuels has effectively been a substitute for fuel "storage", so far as the New Zealand electricity sector is concerned.

Demand side stockpiles/trading

While the above discussion of thermal fuel stockpiles and trading arrangements may seem increasingly irrelevant, it is important to understand that, when Section 2.3.1 of our report to MDAG talks about an increasing role for DSM, it is talking about utilising very much the same set of mechanisms, except on the demand side of the electricity sector, rather than the supply side. There, we discuss the distinctions between demand reduction; and demand deferral, and between implicit elastic demand responses to market prices, and explicitly contracted demand responses.

Overall, we conclude that there could be significant potential for demand-side response to enhance effective system storage capacity, and/or replace thermal flexibility, but also much development still to be done in that area. Options would include the aggregation of a great many small-scale (e.g., household) responses. But they could also involve a few large-scale responses, such as establishing a large-scale electrolysis plant to produce green hydrogen that could be traded internationally (e.g., as ammonia). Then, production could be reduced, with exports being reduced or imports increased, in "dry years", or perhaps when hydro storage falls to a certain guideline level, or projected electricity market prices rise high enough to make New Zealand electricity sales more profitable than supply to the electrolysis plant, over some planning horizon. [26]

Even the "reduction" options are implicitly utilising the storability and tradability of goods, feedstocks etc outside of the New Zealand electricity sector, to substitute for the storability and tradability traditionally called upon within the electricity sector, or in associated fuel sectors. In both cases, the net effect is to draw on the inherent diversity and flexibility of the international production/storage/trading system, which is vastly greater than that of the local electricity sector.

Although this report does not do so, we expect that investigations will eventually address the kinds of operational and organisational arrangements that would be required to access some of that potential response, particularly as a substitute for, or complement to, the kind of large-scale supply side storage options being considered. In fact, something like an explicit arrangement with a major exporter of an electricity-intensive product like aluminium or hydrogen, could be the nearest substitute available for the thermal sector storage/trading arrangements operating under the status quo. The operational/

---

25 In principle, gas is tradeable too, but large-scale LNG trading infrastructure was never developed in New Zealand.

26 Shorter term flexibility could also be feasible, and highly valuable, but the details are yet to be investigated.

organisational issues, though, are likely to centre on the extent to which such demand side response need to be contracted, and perhaps then explicitly managed like supply-side storages within the electricity sector, or left implicit, perhaps as elastic responses to electricity price levels.

2.4. Multiple independent storages

Section 3.2.1 of our report to MDAG discusses the extension of the above theory to cover management of multiple storages that are "independent", in the sense that they are not hydraulically connected within the same river catchment. Irrespective of whether those storages are hydro reservoirs, or something else, and no matter who manages them:

- The EMVSE in each of them must now depend on the whole vector of storage levels across all of them.
- Thus, in principle, there are now ***n*** EMVSE surfaces, each with ***n*** dimensions, for a system with ***n*** independent storages.
- But optimal arbitrage, either via centralised optimisation or perfectly competitive market interaction, will act to equate EMVSE values, as closely as it can within the limits implied by minimum/maximum limits on utilisable release from each storage, and the limits and losses imposed by the transmission system.

Fortunately, this last bullet implies that participants need not really concern themselves about the precise level of every storage in the system, but can make reasonable decisions based on considering the storage levels in a few key (aggregate) storages. It also means that any misalignment in EMVSE or storage levels, whether due to managerial misjudgements, or unexpected events, will be self-correcting over time.

This kind of equilibration between EMVSE levels in storages linked by a "lossy" transmission system provides an important analogy to the kind of equilibration we should expect to see, and plan to ensure, between the upper and lower storages in a pumped storage hydro facility, linked together by a "lossy" hydraulic system.

2.5. Linked storages

Section 3.3 of our report to MDAG discusses some issues of critical importance to understanding the economics of pumped storage hydro operations. Once two or more hydro stations are linked together into some kind of river chain, their EMVSE and optimal generation patterns are no longer independent.[27] Under moderate conditions, when no release/flow/generation/storage limits constrain operations, EMVSE levels might be quite closely aligned down the chain, but of course the potential energy stored in the water is being released as it passes down the chain. So, it becomes important to distinguish between EMVSE, defined in $/MWh terms, and the "Expected Marginal Water Value" (EMWV), in $/ unit of water volume:

- If a unit of water in an upstream storage can be released to instantly arrive in the next downstream storage, then its EMWV must be at least as high as a unit of water in that downstream storage, at that moment, even if that release is a spill that generates nothing; and

[27] Many models and discussions treat this short run river management problem as largely deterministic, but uncertainty is still a significant issue, so we will refer to EMWV/EMVSE, rather than MWV/MVSE.

- The optimal release/generation pattern for the intervening power station will be driven by the difference between the EMWV in the upstream storage, and the EMWV in the downstream storage.[28]

As discussed in that section, the picture can be significantly complicated by issues like flow limits, and delays, either of which can make it impossible to instantly transfer a marginal unit of water from an upstream to a downstream reservoir, thus making it possible for water that is already downstream to actually be more valuable than water still trapped upstream, at particular times. Luckily the interaction between the upper and lower reservoir in a pumped storage facility is typically more direct, but flow limits still apply, and the interactions between the pumped storage facility and the host system within which it is embedded must account for all the flow/storage limits and delay times applying to that system. The most fundamental difference, though, is that the "upper" reservoir is not actually "upstream" from the "lower" reservoir, or vice versa, in a system where water can circulate.

2.6. Sectoral overview

Although they were prepared independently, the general picture emerging from our conceptual report to MDAG aligns strongly with MDAG's own quantitative modelling work. Clearly:

- Managing "cyclic variation" in the supply/demand balance, over days, weeks, and seasons, has always been, and will always be, a critical issue in the New Zealand power system; and
- So has managing random fluctuations (mainly) caused by the impact of weather on demand and on the availability of energy on the supply side.[29]

Historically, three main types of "storage" have worked together in a loosely coordinated way, to manage both cyclic variability and random volatility:

- Conventional hydro system storage, as discussed above;
- Demand-side storage, such as the storage of energy in domestic hot water cylinders; and
- Thermal fuel storage, both explicit and implicit in the international trading system.

The emerging challenge is that:

- Adding more intermittent renewable capacity will increase the volatility of supply/demand balances, at least over short to medium term planning horizons;[30] while
- Removing thermal capacity will put heavy pressure on the ability of the other two storage options to manage both cyclic variation and volatility; and
- There is significant work yet to be done if the conventional DSM component is to increase to the levels assumed in MDAG's simulations.

MDAG argues that batteries will become an increasingly economic option for managing the day/night cycle, and intra-day volatility, but they will not be able to deal with longer term supply/demand balance issues, in the foreseeable future. So, the implications are that:

- Storage and flow levels in existing river chains will have to be varied more aggressively in order to manage supply/demand balance daily cycles, and short-term volatility; while
- The difference between summer and winter prices is expected to be much greater than in the past.

---

28 The pattern of MVSE values down the chain is trickier. If some station is so overloaded that incremental water must be spilled past it, the EMWV in its storage will equal than in the storage below, and those values will be monotonically non-increasing down the chain. But spilling the water reduces its potential energy without decreasing its value, so its EMVSE actually rises.

29 Traditionally this has just been the availability of water to "fuel" that generation, but increasingly also wind, and now solar.

30 The impact over some other planning horizon is less clear. Since hydro generation capacity is not anticipated to expand greatly, wet/dry year variability will gradually decrease, as a proportion of total New Zealand load. And, while increased solar generation will increase the importance of the winter/summer supply/demand imbalance cycle, it will reduce the day/night imbalance cycle.

- But traditional hydro storages will need to more consistently store as much water as possible through autumn into winter.
- So, there may be a disconnect between marginal water values and electricity prices over summer, when wind, solar, geothermal and uncontrollable hydro are expected to set quite low electricity prices, sometimes “spilling energy”, while marginal water values stay high, as major hydro storages release as little as possible in order to build stocks up to meet autumn targets; and
- There may also be a disconnect between marginal water values and electricity prices, at times over winter when hydro is generating at maximum capacity, and prices must rise high enough to induce significant demand-side responses, as the only way to manage supply/demand balance over critical periods.

Since hydro generation capacity is not anticipated to expand greatly, wet/dry year variability will gradually decrease, as a proportion of total New Zealand load. As a result, the most critical concern is predicted to gradually shift away from “dry years”, defined in terms of aggregate inflows into long-term storage reservoirs, towards extended periods of dull calm weather (dunkelflaute) conditions. In the New Zealand context, the situation will be most critical in winter when we must also add cold to the above dunkelflaute description and could experience extended periods of low tributary flows, even in “wet” years.

With traditional thermal capacity unavailable, batteries may well provide enough storage capacity to deal with this issue over one daily cycle, and expanding traditional DSM measures would help, possibly assisted by some “green thermal”. But the supply/demand balance is still expected to be very tight, and likely to become critical if those conditions continue over days and weeks, as analysis of historical weather patterns suggests they sometimes will.

So, while MDAG did not study large-scale storage options, it did highlight their potential value, in four particular roles:

- Replacing the “storage” traditionally available from the domestic/international thermal fuel production/trading system, in dealing with wet/dry year fluctuations over multi-year timeframes, which has always been a problem in the New Zealand power system and will remain so, for as long as the sector relies heavily on conventional hydro;
- Assisting conventional long-term (traditionally annual) hydro storages in dealing with the inter-seasonal supply/demand balance cycle, which is projected to imply a much stronger annual price cycle than in the past;
- Assisting batteries and conventional short-term hydro storages in dealing with intra-day supply/demand balance cycles and increased short-term volatility due to reliance on weather dependent renewables; and
- Assisting conventional mid/long-term hydro storage in dealing with extended dunkelflaute situations.

While MDAG concluded that the market could continue to produce acceptable, albeit more volatile, outcomes without significant storage augmentation, it clearly also recognised the potential contribution from two main longer-term “storage” options:

- Large-scale pumped storage hydro, or some similar technology that would allow energy to be stored within the electricity sector and/or “fuel supply” infrastructure, for eventual conversion into electricity, when required; or
- Large-scale demand side flexibility, e.g. from an export-oriented hydrogen plant, able to dial down its electricity consumption when the supply/demand balance becomes critical, in which case the “storage” being utilised is the ability of the international trading/storage system for the export commodity to take up the slack.

Accordingly, the current report is intended to complement our report to MDAG, by focussing on developing a conceptual framework for the analysis, operation, and organisation of such options, in a realistic New Zealand market context.

2.7. Non-traditional storage options

2.7.1. Introduction

Pumped storage hydro is obviously an important option, and a particular focus of current investigations, but it is far from the only option under consideration. The ultimate goal of this investigation is to develop a general framework for understanding how each of the options likely to be considered might be operated, organised, and integrated into both spot and hedge markets. It should be recognised, though, that each option has its own technical characteristics, likely operating mode, and organisational options and challenges. Thus, our discussion can not be entirely generic. Broadly speaking we see four types of options, and will discuss each separately below.

2.7.2. Passive supply side storage options

These options may be thought of as "incidentally embedded" in the supply side of the electricity sector, in the sense that there is some store of energy, in a tank or reservoir, closely linked to a generator. But, while it may have been supplied from a wider distribution network the fuel (or water), once delivered, is stored for the use of that particular power station, and will be managed for that purpose alone. This category would include traditional hydro reservoirs and fuel stockpiles, but also some potential green alternatives.[31]

The management of traditional hydro is well understood, and has already been discussed extensively in our report to MDAG. So, we see no need to discuss its operation any further in this report, and note that it already has established organisational structures, and is reasonably well integrated into both spot and hedge markets. We only note that the storage aspects of traditional hydro systems will become increasingly important in future years, and suggest that, if a market were to be developed for some form of storage-based hedges, as discussed in Chapter 5, serious consideration should be given to designing that market in a way that would allow traditional hydro to participate.

The management of traditional thermal will become increasingly irrelevant, but MDAG has identified a potential role for "green thermal", perhaps fuelled by hydrogen, or biofuels. The fuel supply chain would obviously be different, for each technology, but the overall structure of the biofuel supply chain does not seem very different from that for existing thermal capacity. We expect there would be stockpile of some kind, stored for station use, and a limited rate at which that stockpile could be replenished from local sources, and/or the ability to import. This seems very similar to the situation currently applying to coal, and we expect it would be managed in a very similar way to that discussed in Section 2.3.

Hydrogen could present some different issues, though, if it was locally produced by electrolysis. Obviously, it would not make sense to use electricity to produce hydrogen at times when prices were high enough to justify producing electricity from hydrogen. But one could imagine a closed cycle, in which hydrogen was produced at times when electricity prices are low, and then used to fuel generation when electricity prices are high, thus forming a kind of "battery", which could be analysed and managed as such under the "stand-alone" category below.

Alternatively, we could imagine hydrogen powered generation occurring in the context of a much wider trading network, in which hydrogen night be exported at some times, and imported at others. If there was no local hydrogen powered generation, this would be an example of "embedded" demand-side response, as discussed below. Local hydrogen powered generation enables a net switch from absorbing electricity to "charge" the hydrogen storage system, to "discharging" the hydrogen storage system to produce electricity. That would create a situation more akin to that of embedded pumped storage hydro, as discussed in Chapter 4, below. There may not be any physical linkage between the hydrogen-

31 Gas is an exception, in that it is supplied by a network, within which interactions quite similar to those discussed later for pumped storage could arise. But there seems little point discussing the intricacies of those arrangements, given that gas is supposed to be eventually retired, as a fuel, in New Zealand.

producing facility and the hydrogen-powered generation, either in the electricity network, or the hydrogen trading network. But the traded price of hydrogen should optimally be driving the behaviour of both facilities, in a similar way to host system EMWVs, impacting both operational and hedging strategies.

Geothermal may also be seen as a special case. Although generally thought of as “renewable”, most geothermal developments have significant incidental carbon emissions, to the extent that, unless a suitable carbon capture technology can be developed, future geothermal development seems likely to significantly curtailed by rising carbon prices. We understand, though, that there has been discussion of the possibility that geothermal plant could be developed, but only used in standby/backup role.

At one level, if we believe the carbon price will be set at an appropriate level, that that seems like a decision that could be left to the market, with potential investors making their own decisions as to whether it is worthwhile investing in plant that will only run when market prices are high enough to cover the carbon price, plus variable O&M. But there could be policy reasons for treating this kind of development as providing a form of storage capacity, perhaps operating under a stricter carbon limit than that implied by the carbon price.

Given the longevity of carbon in the atmosphere, it makes little difference when, in its lifetime, such plant may operate, so it probably makes most sense to think of it as having a lifetime carbon budget. In that case, it is the carbon budget, not the geothermal resource, that constitutes the “storage” to be managed. The facility manager should logically determine an opportunity cost to ration its stock of carbon release opportunities over its lifetime, and then add that opportunity cost to the carbon price when determining operational strategy. But note that discounting would play a very significant role over this kind of planning horizon, and the opportunity cost should not be constant, but rising at the discount rate, as in Hotelling’s rule, thus implying an increasingly restrained operational policy, as time progresses.[32] [33]

2.7.3. Stand-alone active storage options

The facilities we have in mind here are dedicated to storing electricity, in the sense that (unlike with traditional hydro or thermal) electricity is taken out of the electricity system to build up some store of potential energy, so that it can be later released back into the electricity system. In principle, this could be achieved using a very wide range of technologies to store the potential energy, including compressed air, and flywheels, but the two most relevant options in the current context would be batteries and pumped storage hydro.

Batteries are an inherently “stand-alone” technology. While some batteries might be associated with particular (typically intermittent) generators, their internal operation is in no way connected with the generation facility they may be associated with. The same could be true of a pumped storage hydro development that drew water from a source not otherwise associated with electricity generating infrastructure, and released to that source, or possibly to another similarly un-associated water body.

---

32 In other words, it should behave rather like the “Depletion Related Opportunity Cost” calculated for gas-field depletion. As in that case, though, that may mean that current opportunity costs are set by working back from some assumed future value to a possibly very low current value. And that would imply that the geothermal plant might expect to operate reasonably freely in its early years, adding only a small increment on to the carbon price, but become increasingly restrained as time progresses, and its “stock” of allowable carbon discharge is depleted. In fact, under various simplifying assumptions, Hotelling’s rule also implies that the usage rate should be decreasing at the discount rate, if the plant had an infinite life. We have not researched the issue to determine how that broad outline might be modified by the particular characteristics and maintenance cost structure of geothermal developments, or by factors such as the carbon price probably also rising over time, and dry year response requirements possibly falling over time.

33 The economic and scientific logic discussed here actually applies to all existing thermal and geothermal plant, with a critical issue then being the technical and economic life of that plant over an extended period of rising maintenance costs and decreasing utilisation. In all cases, the prospect of technical failure and/or obsolescence, not to mention annual maintenance costs, must increase the incentives implied by Hotelling’s rule to utilise available resources earlier, rather than later. And that will be particularly the case if doing so allows stocks to be built up in other, more durable, storage developments.

While the discussion of stand-alone storages in Chapter 3 focuses on batteries and pumped storage hydro, we believe it will largely apply to any other stand-alone technologies. The main difference is that some technologies are essentially self-contained black boxes (like batteries), while others store potential energy by drawing in some medium, such as water or air, from the environment, and holding it under pressure, or at elevation.

The organisational challenges may differ significantly, though, depending on the extent that the technology might be dispersed across multiple small-scale sites, vs focussed in one large-scale development. While our discussion in Chapter 6 is focussed on pumped storage hydro, we believe that it would largely also apply to other large-scale stand-alone developments, without the complexities potentially arising from "host system interactions". Dispersed smaller scale developments might benefit from centralised development of a market, for the kind of hedge products discussed in Chapter 5, but should not need the same kind of attention as large scale developments, for which market power is a significant issue.

2.7.4. Active storage options embedded in supply-side systems

The most obvious example here would be pumped storage hydro. Such developments might be retrospectively embedded in existing "host" hydro systems, but they could also be designed as part of an integrated catchment development that also included conventional hydro. The fundamental difference is, though that these are facilities specifically designed to provide "active" storage. That is, water would be actively pumped up to a dedicated "upper reservoir", at some cost, rather than just being passively "held" in some structure that essentially just delays its ultimate passage to the sea. We have devoted a whole chapter to the discussion of such developments, because we think the mix of (large-scale) pumped storage and conventional hydro in the same catchment poses unique operational and organisational challenges beyond those faced in the current market, or by the other options considered here.

2.7.5. Embedded demand side storage options.

A variety of demand-side "storage" options were briefly discussed by Section 4.3.3 of our report to MDAG, which suggested that they could play a significant role, with possibly greater potential than pumped storage hydro. There is a great diversity of possible developments, though:

- At one end of the spectrum, large-scale export industries could be deliberately developed in such a way as to ensure flexible demand reduction in response to short and/or long-term fluctuations in the electricity supply/demand balance. Production of hydrogen by electrolysis, or even development of more flexible smelting activities, would provide prime examples.
- At the other end of the spectrum, encouragement, and perhaps co-ordinated management, of distributed small-scale flexibility, including EV/solar battery charging, and even traditional hot water heating loads, could also make a significant contribution to the DSM component which MDAG considered to be critical to managing supply/demand balance, in the absence of traditional thermal capacity.

We describe all of these options as providing "embedded storage", because what is actually being stored is a product, such as hydrogen, or heated water, and the storage facility is embedded in some demand-side system somewhere on the planet, and not in the New Zealand electricity system, or necessarily even in New Zealand. We see this as being somewhat similar to the way pumped storage hydro is embedded in a host hydro system, in that:

- The medium of potential energy storage has its own value, as a product or feedstock in some process other than the energy storage process itself; so
- The effective cost of charging the storage will include the cost of taking that medium away from the production/storage process in which it would otherwise have been involved, at some time; and
- The effective value from discharging the storage will include the value of returning that medium to the production/storage process from which it was withdrawn at some earlier stage.

As a result, optimal charge/discharge decisions should not just be driven by electricity prices, but also by the fluctuating valuation of the storage medium in other markets. And that also means that the hedging these options can provide to electricity market participants might be affected by factors other than the electricity price at that moment.

Unlike pumped storage hydro, though, the other markets whose price movements may play an important role typically have no connection with the New Zealand electricity market.[34] The delay and cost involved in switching from charging to discharging, and the organisational issues will also differ widely between options. Thus, a detailed discussion can not be generic, and seems inappropriate, at this stage.

[34] Hydrogen being a possible exception, in that it might be both produced and consumed within the New Zealand electricity sector, at different times. And it might be both exported and imported at different times, too.

3. Optimal Stand-Alone Storage Operation

3.1. Introduction

This chapter discusses the idealised operation of a large stand-alone storage facility on the assumption that this could be achieved through centralised optimisation, or in a perfectly competitive market. While that assumption is obviously unrealistic it enables us to establish a basic theoretical reference point, to which later discussions can refer. Also, while this discussion focuses, at times, on the specifics of pumped storage hydro, the principles are broadly applicable to any "stand-alone" storage facility: that is, to storage facilities that are either fully "closed" or fully "open", in the senses discussed below.

The most obvious "closed system" would be a battery, and we discuss the optimal operation of batteries before moving on to consider pumped storage facilities that might either be fully closed or fully open, but have no interaction with any "host" hydro system associated with either the lower or upper reservoir. Studying these stand-alone systems allows us to understand the properties, and optimal management, of some of the simpler options under consideration. But it also provides a basis for understanding the much more complex case of storage facilities "embedded" in other systems, including the embedded pumped storage facilities discussed in Chapter 4 below.

3.1.1. Closed systems

It has often been said that pumped storage is just like a big battery, but that is an over-simplification for most real systems. It would be true, though, if the pumped storage facility was a stand-alone closed loop system, in which the same water was constantly re-cycled. That would be possible, if the combined rate of inflow to upper/lower storages at least matched the combined rate of leakage or evaporation.[35]

In addition to the obvious limits implied by the MW capacity of the pumps/generators, the effective cycling/storage capacity of the system would be set by the minimum of the upper and lower reservoir capacities. Even if the upper/lower reservoir could contain more than the other reservoir, the extra could never be run up/down into that other reservoir. The pumped storage hydro facility could operate "just like a battery" within those capacity limits, though. Conversely, the principles discussed here apply equally to any battery, or other "closed" storage system, which contains its own storage medium, rather than relying on potentially constrained supplies from some other system.

3.1.2. Open systems

Pumped storage hydro could also operate "just like a big battery", if the water was drawn from, and discharged to, the ocean, or some body of water large enough as to not be materially affected by the pumped storage hydro operation. In that case, its operation will not be limited by interactions with any other system, and its effective storage capacity is just the energy storage capacity of the upper reservoir of the pumped storage facility. And the same will be true of any other facility, like compressed air storage, drawing its storage medium from an essentially unlimited source.

[35] If water loss occurs from the upper storage, it may be seen as equivalent to the gradual "leakage" of energy from a battery, over time. This loss of water in storage is quite different from the "round-trip" loss of electro-mechanical energy due to inefficiencies in the pumping process, as discussed elsewhere. Such leakage does not fundamentally change the principles discussed here, but we will discuss it, along with "discounting", in Section 3.3.4 below.

3.2. Optimal battery operation

Section 4.3.2 of our report to MDAG briefly discusses the application of opportunity costing principles to batteries. Although similar to hydro reservoirs, they differ in three important respects. Ignoring losses incurred by storing energy over time:

- There is no “free” natural inflow to replenish them, or outflow to be maintained; but
- They can be deliberately recharged at a cost determined by electricity prices; although
- There is a “Round-Trip Loss Factor (RTLF) involved in each charge/discharge cycle.[36]

Looking at the situation from a deterministic perspective, the same arbitrage concepts apply as for conventional hydro reservoirs. As in that case, we can start by identifying the most profitable arbitrage opportunity, and then working down the list as opportunities are fully exploited. But we are now looking for profitable exchanges between periods for which the electricity price differential exceeds the RTLF. Given the limited storage capacity of batteries, we will restrict our search to be over a single daily cycle.[37] Then, we are limited by two factors: [38]

- The charge and discharge rates (which need not be the same) forcing us to move on to the next most profitable trade, once either limit is reached in one of the pair of periods we are currently arbitraging between; and
- The storage capacity of the battery.

If the daily price patten has a single (typically evening) peak, then the battery should have only one primary intra-day cycle, which should optimally look just like a miniature version of the inter-seasonal cycle discussed for long-term hydro reservoirs in Section 2.2 above. That is, the marginal value of stored energy (MVSE) should be:

- Constant over the period during which storage is “filling/charging”; then
- Equal to the rising electricity price, over the period during which storage is held “full”; then
- Constant at a higher level over the period during which storage is “emptying/discharging”, including the evening peak; and finally
- Equal to the falling electricity price, over the period during which storage is held “empty”.

Analysts who customarily work with batteries associated directly with solar installations may assume that the “filling” period will be “daytime”, and the “emptying” period, “night-time”. Such physical associations might sometimes be justified by physical transmission limits. But they are more often individually rational, but nationally sub-optimal, responses to imperfect market arrangements, under which the buy and sell prices of on-site electricity are unequal.[39] Since our assumed paradigm here is national cost-benefit optimisation, and our primary focus on large-scale grid-connected developments, such as pumped storage hydro, the proper comparison is with large-scale grid-connected batteries, buying and selling at spot prices. Such facilities should be charging when prices are lowest, and discharging when prices are highest, irrespective of the marginal source of the electricity they absorb, or displace.

If the system were to become solar dominated, then prices could be at their lowest across the middle of the day. On the other hand, wind, geothermal and run-of-river hydro also seem likely to make a significant contribution to charging such batteries, when loads are low overnight. So, a two-cycle operation may develop, with such sources being drawn on to build up storage before a morning-peak run-down, while solar contributes more to building up storage for an evening-peak run-down.

More generally, secondary cycling can be optimal at any time of day. The primary storage cycle described above is driven by broad price differentials between the filling and emptying periods (most

---

36 For example, RTLF =1.2 for a 20% round trip loss situation. Noting that this is a proportion, not an absolute value.

37 This can be generalized, for larger batteries, such as may be required to manage through several dull calm days.

38 Assuming that the battery is small enough that its charge/discharge rate has no impact on electricity market prices.

39 This is because buy-back rates are typically much lower than the per-unit prices paid by most consumers, mainly because the latter include a large component required to recover fixed costs, which do not actually vary on a per-unit basis.

likely day/night), and limited by battery storage capacity. But that analysis implies that MVSE is expected to be constant, at a relatively high level, over the emptying period, and at a relatively low level, over the filling period. Then, within each such period:

- The conditions and prices outside the period are irrelevant, because no more energy can be transferred into, or out of the period, in response to those conditions; and
- Storage capacity bounds are irrelevant, except inasmuch as they define the aggregate energy charge or discharge required across the period; but
- The MW charge/discharge limit becomes critical, because that determines the extent to which the aggregate energy charge/discharge requirement can be met by focussing solely on the period(s) of absolute lowest/highest price; and
- Irrespective of whether the overall goal for the period is to charge or discharge the battery, its intra-period MVSE should determine how it responds to intra-period electricity price variations, by:
    - Discharging when prices are above that MVSE.
    - Charging when prices are below that MVSE, adjusted (downwards) for RTLF.[40]

It may seem strange that a battery could behave in a potentially counter-cyclical way, but this results from the interplay between the electricity price pattern, its RTLF, and its two capacity limits:

- If it had unlimited MW charge/discharge capacity, and prices change monotonically from lowest to highest, and back again, then it should fully charge in the absolute lowest-priced market trading interval of the day, then hold that charge to be fully discharged in the absolute highest-priced market trading interval of the day.[41] And, in that case, the MVSE applying to its "filling period" would be just the very lowest price, the MVSE applying to its "emptying period" would be just the very highest price, and it should not respond to any price variations in between.
- More realistically, it will:
    - Charge at its maximum possible rate, over some set of market trading intervals within each filling period, and its MVSE for each such period will be set by the highest price at which it charges.
    - Discharge at its maximum possible rate, over some set of market trading intervals within each emptying period, and its MVSE for each such period will be set by the lowest price at which it discharges.
- If prices do cycle monotonically, then those highest/lowest priced trading intervals will all be grouped into two contiguous periods of time, thus defining the emptying/filling periods referred to above, and the battery should ignore the (monotonic) price variation within, or in between them.
- If prices do not cycle monotonically, though, there is no reason why the battery should not respond to price fluctuations, whether expected or unexpected (see below), by temporarily deviating from its overall charge or discharge goals for that period; and
- The greater its MW charge/discharge capacity, the greater the deviation can be, without jeopardising its ability to meet those goals by the end of the period.

The same may also be true, if prices do not rise/fall monotonically over the intervening periods, during which the battery storage is supposed to be full/empty, with MSVE falling/rising monotonically.[42]

Layering a stochastic perspective over the deterministic cycling discussed above, each battery manager should ideally be assessing an Expected MVSE (EMVSE) for energy stored in the battery, at any time, as a function of that battery's storage level, national storage levels, and load/weather/price forecasts. As for the large-scale hydro reservoirs discussed in Section 2.2, EMVSE will be the expected value, over a

---

[40] That is below MVSE/RTLF. For pumped storage we will refer to these critical price levels as the Marginal Cost of Generation (MCG), and Marginal Value of Pumping (MVP), but the terminology seems inappropriate for batteries.

[41] Assuming that the battery is not so large that its charge/discharge rate is actually limited by available supply/demand, or by local transmission/distribution capacity.

[42] Although the likelihood that this kind of deviation will be optimal for a battery is much lower than for a conventional hydro reservoir, because any short-term deviation will imply an actual cost, determined by the RTLF.

wide range of possible scenarios, of the best available marginal opportunities for using the energy stored in a particular battery. Thus, these EMVSE values will vary greatly, depending on the charge/discharge and storage capacity of each battery, and the charge/discharge opportunities available to it.

For each individual battery, though, we note that:

- There is no chance of forced spill becoming necessary before the battery is fully charged; and
- In physical energy terms, there is no point in spilling storable energy in any period, just to create room for the possibility that more storable energy will turn up later; and
- While it may be desirable to leave headroom so that the battery can quickly absorb a sudden supply surge, such as might occur if the windspeed or effective solar irradiation suddenly increases, that might be considered less an issue of "energy market" economics, than of ancillary service provision, for which the battery should be compensated separately.
- Still, the likelihood that a supply surge might last long enough to cause market prices to drop is not an issue for ancillary services, as currently defined, and should be factored into the manager's assessment of the distribution of prices likely to be available over the rest of the cycle. And the manager should become increasingly picky, as the battery approaches its fully charged state, aiming to charge in the lowest priced periods still expected to be available before the battery is full.

Conversely, at the other end of the charging cycle:

- There is probably little chance that a national crisis will be triggered by the fact of any individual battery being fully discharged;[43] and
- In physical energy terms,
- it may be desirable to leave some spare capacity charged, so that the battery can quickly respond to a sudden supply shortfall, such as might occur if the windspeed or effective irradiation suddenly decreased. Theoretically, though, that may also be considered less an issue of "energy market" economics, than of ancillary service provision, for which the battery should be compensated separately.
- Still, the likelihood that a supply drop-off might last long enough to cause market prices to drop is not an issue for ancillary services, as currently defined, and should be factored into the manager's assessment of the distribution of prices likely to be available over the rest of the cycle. And the manager should become increasingly picky, as the battery approaches its fully discharged state, aiming to discharge in the highest priced periods still expected to be available before the battery is empty.

3.3. Deterministic cycling of stand-alone pumped storage

While accounting for stochasticity is obviously important in reservoir management it is not the only, or even necessarily the main, reason storage facilities are built, or driver of operational strategy. Just as in our report to MDAG, we start by discussing the underlying pattern of "deterministic" supply/demand balance cycles, which storages must manage. These imply predictable daily (day/night) and seasonal (summer/winter) patterns to both optimal storage levels and MVSE values.[44] According to the MDAG study these patterns are expected to become much more marked as thermal generation capacity, which has traditionally played a major role in moderating those patterns, is withdrawn. Stochasticity then implies an overlay to these expected patterns, as discussed in the next section.

Section 3.2 treated battery storage like a "black box", with a single MVSE value for the whole facility. With pumped storage, though, we can clearly see that energy is stored by pumping water, and that the

---

43 Given the likely correlations, further consideration should be given to the difference between individual and collective perspectives, but the manager should really be accounting for those correlations, and the likely state of competing storages, when making price projections.

44 We will again ignore weekly patterns, because the principles involved are the same.

same unit of water has more stored energy when it is in the upper reservoir, than when it is in the lower reservoir. So, it becomes useful to focus more specifically on the MWV of water at various elevations in the system. In Chapter 4, the upper reservoir of the pumped storage facility will be the main reservoir specifically constructed to create the pumped storage facility, with water being pumped up to that reservoir from a lower "buffer storage" embedded in the "host system".[45] So, for simplicity, all of our discussions focus on the MWV in that upper reservoir.

- This discussion would apply directly to the "open-loop" system described above, in which the MWV for the lower "reservoir" would always be zero, because that water resource can never be constrained, or run out.
- It may also be applied directly to the "closed-loop" system described above, provided the upper storage is smaller than, or at least constrained to operate within tighter limits than, the downhill storage.[46] In that case, after an initial calibration, the limits on lower storage should never be binding, as discussed in Section 3.1.[47] So, the MWV in the lower storage should be constant over all time, as discussed in Section 2.2. Mathematically, it could be set to any arbitrary value, but it may as well be set to zero because it is only the difference between the MWV in this lower reservoir and that in the upper reservoir that determines optimal pumping/generation strategy, as discussed in Section 2.5.[48]
- The discussion still applies if the lower reservoir of the pumped storage facility is smaller than the upper reservoir, but the application is a little more subtle. In this case, after an initial calibration, the limits on upper storage should never be binding, and consequently the MWV in that upper storage should be constant over all time. But the effective storage capacity of the whole facility is now set by the capacity of the lower storage, and it is now the MWV in that storage that will change as soon as it reaches either capacity bound.
- Again, though, it is only the difference between the MWV in this lower reservoir and that in the upper reservoir that determines optimal pumping/generation strategy. So, we will again simplify the discussion by implicitly assuming that that difference is reflected entirely in the MWV of the upper reservoir of the pumped storage facility, while the MWV in the lower reservoir is implicitly assumed to stay constant, at zero.[49]

In summary, then, the marginal value of potential energy stored in the pumped storage facility is not necessarily the full value of energy stored in the upper reservoir of the pumped storage facility, but the difference between upper and lower MWV levels, just as the potential energy reflects the difference in elevation between the upper and lower reservoirs.[50]

---

[45] While the primary host system could be associated with the upper reservoir of the pumped storage facility, that seems very unlikely in New Zealand.

[46] Technically, the lower reservoir's MWV may be set to any arbitrary value, and held constant over the planning horizon. The upper reservoir MWV discussed in this section would then be additional to that lower reservoir's MWV, but this makes no difference to the optimal operating policy, which is driven by the difference in MWV values, as discussed in the next section.

[47] This is true, even under extreme stochasticity. There just can not be enough water in the upper reservoir of the pumped storage facility to move the storage level in the lower reservoir from empty to full.

[48] Note that this mathematical discussion of economic operation in a closed storage system ignores the intrinsic value of the storage medium, whether that be water, or air, or the chemical constituents of a rechargeable battery.

[49] Mathematically, if the two storage capacities are identical (most likely because the upper storage has been constrained to operate within limits exactly matching those in the lower storage), one storage will be exactly full when the other is exactly empty. So, the MWV could change in both simultaneously. But the MWV difference will still behave as discussed here.

[50] In fact, the situation within this closed storage system is not very different from that within a conventional battery, or more exactly cell. A detailed internal analysis could focus on whether it is the anode or the cathode capacity that is limiting cell capacity, and analyse the internal MVSE behaviour in much the same way as we have done here. But, for most purposes, economic analyses will just treat the cell (or indeed the battery) as a single unit, with a single MVSE. We are only discussing the MVSE and/or MWV of individual reservoirs here, in order to align our understanding of this situation with our discussion of marginal water values in upstream/downstream storages in Section 2.5, and of interactions with the host system in Chapter 4.

....*continued overleaf*

We focus first on "seasonal" cycling within an annual planning horizon, because that is most closely analogous to traditional long-term reservoir management in New Zealand. We then move on to consider both longer and shorter planning horizons.

3.3.1. Seasonal cycling

While our report to MDAG did not attempt any numerical analysis, it accompanied an MDAG study that projected the likelihood that there could be very large inter-seasonal price differentials, unless means were found to increase long-term storage capacity, on the supply side, or to find significant flexibility on the demand side of the New Zealand electricity sector. Since optimal national cost-benefit maximisation would be characterised by prices being as constant as possible across time, such price differentials would be indicative that costs were being imposed on the New Zealand economy and society, due to lack of storage capacity. Conversely, the inter-seasonal price differential provides a measure of the national benefit obtainable by increasing storage capacity to allow enough water to be carried forward to produce one more unit of electricity in winter, instead of summer.

Thus, we would expect any large storage facility being used to maximise the value of available energy by moving as much as possible from summer to winter, with the second order benefit of reducing inter-seasonal differentials.[51] In many respects a stand-alone storage facility should operate very similarly to any other hydro reservoir. Thus, the same arbitrage logic applies. But it is also like a large long-term battery, in that:

- The facility will absorb more energy than it returns to the system, and must be thought of as adding a net load to be covered by additional generation elsewhere in the system; and
- There would be a significant "dead band" within which prices could differ without triggering any pumped storage; so [52]
- Even an infinitely large pumped storage facility could only arbitrage away electricity price differences to within the proportion of energy it loses in each pump/generate cycle.

The MWV plays a critical role here, not just in determining the optimal rise and fall of storage across the seasons, but the level of acceptable price variation/volatility beyond which the facility will respond, at any point of time. Because the losses are proportional, the optimal strategy will be to respond to small price variations when storage is relatively high, and MWV is close to zero, but only to much larger variations when storage is relatively low, and MWV is high.[53]

And we can determine an MWV for a pumped storage facility, just as for any other storage. It's just that, if we think of its MWV still being the value of what is actually stored, and setting its MCR when

---

We also assume that pumping inefficiency is manifested by requiring more energy to achieve the same flow rate for pumping as for generation, because this simplifies of certain discussions, particularly in Appendix A. In reality, pumping inefficiency will most likely be manifested by requiring the same amount of energy to achieve a lower flow rate for pumping than for generation. But we do not believe that changes any of the conceptual conclusions in this report.

51 This being a second order benefit, in the technical sense, since pursuing this kind of development with the aim of reducing price differentials would be an "exercise of market power", and technically incompatible with the perfectly competitive assumptions made in this section.

52 There will also be a threshold time duration below which it is not worth switching from generation to pumping, due to costs, losses and constraints around the switching process itself. But we understand that, at least for pumped storage hydro, that threshold is low enough to be ignored in the context of inter-seasonal, and even daily, cycling.

53 So, if the round-trip loss was 20%, say, prices could differ by up to 20% before the facility would act to pump in the lower priced periods (thus tending to raise prices in those periods), or to generate in the higher price periods (thus tending to lower prices in those periods). Thus, if the MVSE for the storage was, say, 120 in some week:

- It must be optimally generating in any market interval in that week for which the electricity price lies at 120, or above, and may hold market prices at that level over a significant range of supply/demand balance variation.
- It must also be optimally pumping in any market interval in that week for which the electricity price lies at 100, or below, and may also hold market prices at that level over a significant range of supply/demand balance variation;
- But it should not respond at all to prices fluctuating across the range from 100 to 120;
- And it could not respond any more to control price fluctuations in the range below 100 (when it should be pumping at maximum rate), or above 120 (when it should be generating at maximum rate).

marginal, then we need to subtract the pumping losses when determining what it should be prepared to pay as a load in the market, when pumping.[54] In that respect, the situation of a stand-alone facility is essentially similar to that for a battery. Rather than referring to the Marginal Cost of "Release", we will refer to its Marginal Cost of Generation (MCG) which, for a simple stand-alone storage will actually be its MVSE (or EMVSE in a stochastic setting), and its Marginal Value of Pumping (MVP), which will be MVSE/EMVSE adjusted for losses.[55]

If all the larger reservoirs already existing in New Zealand were to be replaced by a single pumped storage facility, and the ratio of its upper storage reservoir capacity to its generation capacity was similar to that existing aggregate system, then we would not have actually added any storage to the national system, in this thought experiment, just made it more controllable.[56] So, if the rest of the system (including thermal generation) remained the same, the annual cycle of that pumped storage facility would also be similar to that of larger reservoirs under the status quo, and its MWV should behave in much the same way as for those reservoirs. From a deterministic perspective:[57]

- Storage should be full in autumn, carrying forward as much water as possible from the summer season, in which it has lower value, to the winter season, in which it has higher value.
- So MWV for the upper reservoir of the pumped storage facility must be rising when is full.
- In fact, the optimal deterministic policy is likely to involve holding that reservoir more-or-less "exactly full" for several weeks in autumn.[58] But, unlike the conventional reservoirs discussed in the MDAG report:
  - If there are no (appreciable) inflows arriving at that upper reservoir, we would not need to worry about "avoiding spill" from it, and it could be held exactly full during the autumn period.
  - Rather than risking spill, we would be just foregoing further opportunities to pump more water uphill during that period, while prices rise to be high enough to make release attractive.

Conversely, the optimal deterministic policy is likely to involve holding the upper reservoir of the pumped storage facility more-or-less "exactly empty" for several weeks in spring, with MWV falling, as it attempts to carry forward as little water as possible from the winter season, in which it has higher value, to the summer season, in which it has lower value. But the situation faced during that part of the cycle would be much more like that of the conventional reservoirs discussed in the MDAG report. Irrespective of any inflows arriving at the upper reservoir; the critical factor would still be the need to hold some precautionary storage to guard against the possibility of a national supply shortfall occurring before water is abundant enough, and prices low enough, to make pumping attractive again.

Also:

- MWV must be constant along each intra-season trajectory arc, i.e., whenever the reservoir is neither empty nor full, and
- This whole cycle must repeat, every year.

Comparing this cycle with what we might expect to see in a similarly proportioned conventional reservoir:

- The ability to control "inflows" via pumping should mean that a deterministic cycle can be more reliably maintained; while

---

54 Accounting for the basic efficiency loss on the uphill leg corresponds to the physical reality that the primary inefficiency occurs in the pumping process, and it implies this definition of MVSE, which aligns best with that employed for conventional hydro storage. The way in which efficiency ultimately depends on the loading levels of generating units on the downhill leg can then be accounted for in just the same way as for any conventional hydro generator.

55 Specifically, MVP= MVSE/RTL.

56 This comparison is only approximate, with the inflow via pumping seen as analogous to natural inflows.

57 As noted in Section 2.5, we ignore the possibility that the critical constraints in the host system could be limits on the rate of change of storage, rather than on storage volume, per se. As before, such constraints would act like upper limits in some periods, but like lower limits in other periods.

58 Although it may still cycle up and down from the storage limit, typically daily, as discussed below.

- The RTLF ratio means that the storage level should be optimally expected to stay constant, with the facility neither pumping nor generating, for significant periods of time, while the electricity price moves through the dead-band between MVP and MCG.

3.3.2. Daily/mixed cycling[59]

The situation is slightly more complicated if the pumped storage facility can operate on a daily cycle, as well as, or instead of, the seasonal cycle discussed in Section 3.3.1, with the details depending significantly on the storage capacity of both upper and lower storages. If both storages are large enough to operate on a seasonal cycle, then they will also have the physical capability to easily accommodate a daily cycle. But the question is the extent to which doing so would deliver economic benefits.

Such cycling may be seen as slowing long-term storage filling/emptying rates. But the upper reservoir should not be filling in hours when electricity prices are above its MVP, or emptying in hours when electricity prices are below its MCG. In part, the MVP and MCG are determined by MWV of the upper reservoir, and imposing a daily cycle on top of the seasonal cycle implies that (in this deterministic world) storage in that reservoir will:

- Fall, then rise back further to a new higher "peak" level (with the same MWV) in each day of its summer filling phase;
- Fall, then rise back to exactly the same peak storage level (with a slightly higher MWV) in each day of its autumn level-holding phase;[60]
- Fall, then partially rise back to a new lower peak level (with the same MWV) in each day of its winter emptying phase; and
- Fall, then rise back to exactly the same peak storage level (with a slightly lower MWV) in each day of its spring level-holding phase.[61]

But notice that the MWV, for a large upper reservoir, is basically driven only by long term factors.[62] So, if the lower reservoir capacity is not constraining, then:

- If the upper reservoir MWV is high, the upper storage must be relatively empty, and is prepared to pay a relatively high price to build up long-term storage. Thus, if electricity prices are low enough, it may optimally spend most of each day pumping, and little or no time generating.
- If the upper reservoir MWV is low, the upper storage must be relatively full, and is only prepared to pay a relatively low price to build up any more long-term storage. Thus, if electricity prices are high enough, it may optimally spend most of each day generating, and little or no time pumping.

Thus, such a facility could play a significant and valuable daily cycling role when its storage is at moderate levels, and/or daily price cycles are strong. And, because RTLF is proportional, it may cycle quite actively, in response to quite small absolute price differences, if prices are low and its own MWV is approaching zero. But we might describe this as "incidental daily cycling", because it may not be physically able to play that role much when its storage is at very high or low levels, and it may not be economic to do so if daily price cycles are low relative to its dead-band, which will be much wider when storage is low, and MVSE high.

At the opposite end of the spectrum, a closed pumped storage system in which both the lower and upper storages are small will obviously be less physically capable of cycling. It may not even be able to accommodate the daily cycle that would be optimal, given its RTLF, without reaching both upper and lower storage limits in at least one storage. Still, such a storage will cycle daily, irrespective of the

[59] As for other reservoirs, there is likely to be a weekly cycle interposed between the daily and seasonal cycles, but the principles discussed here can easily be generalised to cover any combination of regular cycles.

[60] In autumn, we can think of the peak being the full level, from which storage falls, and to which it returns, each day.

[61] In spring, we can think of the trough being the empty level, to which storage falls, and from which it returns, each day.

[62] For a small facility, the optimum is likely to be alternating between maximum generation and maximum pumping modes, whenever the day/night differential exceeds the stated thresholds, with that strategy perhaps being moderated switching from maximum pumping to maximum generation accounted for a significant proportion of the load.

longer-term situation, because it can not respond to that longer-term situation. Indeed, that is the role for which such facilities are frequently designed, elsewhere.

If electricity prices are low, on any day, then its MWV, MCG and MVP will all be low on that day, and the difference between MCG and MVP will also be low, so it may cycle strongly, in response to quite small price variations. If electricity prices are high, on any day, then its MWV, MCG and MVP will all be high on that day, and the difference between MCG and MVP will also be high, but it may still cycle strongly if price differences across the day are also large. At a more detailed level, it may also engage in secondary cycling and/or stochastic response, just like the battery discussed in Section 3.2

In the New Zealand context, though, it is quite likely that one storage will be much larger than the other. If the upper storage is large, then its MWV will be set by long-term considerations, including inter-seasonal, and possibly even inter-annual arbitrage. So, it will only change slowly, reflecting where its storage level sits, relative to its optimal long-term strategy. If the lower storage only has a daily capacity, though, its MWV will reflect conditions on the day, and should be expected to cycle between high and low values each day. Since the MCG/MVP for the pumped storage facility is determined by the differences between these MWVs, it will also cycle daily, implying that a daily pumping/generation cycle, with possible secondary cycling, is optimal for the facility, irrespective of whether the upper storage is in a long-term filling or emptying phase.

The assumption that the lower storage would be so small as to only permit daily cycling is obviously extreme, but broadly similar conclusions would apply to situations involving somewhat larger lower storages. And we believe those conclusions to be basically valid. There is a fundamental problem with the example discussed here, though, because no-one would construct a "closed" pumped storage scheme in which the two storages were so imbalanced. And we could not pursue a long-term management strategy to build up, or draw down, stock levels in the upper storage, unless the system was being replenished from some external source, and can discharge to some external sink. So, we will reconsider the situation described here in Chapter 4, as a type of "embedded" pumped storage system.

3.3.3. Additional storage capacity

The above discussion applies to a hypothetical situation, in which the existing inter-seasonal storage reservoirs were all replaced by a pumped storage facility, with a similar ratio of storage to MW capacity. But what if a significant addition to storage capacity is possible? If this new facility has the same MW capacity, then the fact that it has more storage capacity will seldom have any further impact on intra-day price differentials, once the storage capacity reaches the point where it is neither full nor empty within each daily cycle, implying a constant MWV/MVSE over that planning horizon.[63] The amplitude of the optimal annual storage cycle will obviously increase, though, to utilise all of the capacity available, by:

- Building up storage faster in summer, at a higher MVSE than in the more constrained case above;
- Holding storage full for a shorter time in autumn;
- Running down storage faster in winter, at a lower MVSE than in the more constrained case above; and
- Holding storage empty for a shorter time in spring.

As that (hypothetical) expansion process progressed,[64] the summer and winter MVSE values would be coming closer together until, eventually, they would be equal, once the storage capacity was large enough to fully accommodate the maximal annual arbitrage cycle for the assumed MW capacity and loss factor. At that point, a facility with the assumed MW capacity/loss characteristic could do no more to equalise summer/winter MVSE levels, or electricity prices. But that does not mean there will be no inter-seasonal (or indeed intra-day) price differentials, because:

63 Although some affect may be noticed if storage is close to limits, in the spring/autumn phases discussed above.

64 Noting that the "expansion" process referred to here should not be thought of as describing development over time. It is just a thought experiment, such as might be conducted to determine the optimal size of the facility, during the design phase.

- The MW capacity of the facility will still limit its ability to arbitrage perfectly. If that capacity is relatively small then, the facility might be pumping at maximum capacity in all periods when electricity is priced below its MVP, and generating at maximum capacity in all periods when electricity is priced above its MCG, but still having little impact on electricity prices in either season.
- Even a facility with infinite MW capacity would only be able to arbitrage price differences away down to its RTLF[65]. So, it would not actually be either pumping or generating at maximum capacity in all periods, but may sit idle for significant periods, in any season, and particularly during the transitional/shoulder autumn/spring seasons when prices are likely to spend more time, on average, at moderate levels.

3.3.4. Discounting, wastage, and head effects

The impact of discounting, wastage, and head effects was briefly discussed at the end of Section 2.2 which commented that these effects have normally been ignored in New Zealand discussions, because water was not stored long enough for the first two effects to have much impact, and because few power stations were directly connected to reservoirs subject to appreciable head variation. That may no longer be the case, though, if a large-scale storage development is pursued.

Storage trajectories in a large enough storage facility may not reach, or even approach, either bound for years at a time. So, an incremental storage unit may be carried through for many months or years, on some trajectories, before it reaches the period of optimal (discounted) marginal release "opportunity". And the longer water is to be held before release, the more evaporation there will be too, further reducing incentives to fill the reservoir. In the limit, a very large storage capacity might never be filled, because we simply can not find opportunities with a high enough discounted opportunity cost, anywhere in the future, to justify pumping to increase storage levels.[66]

On the other hand, for pumped storage hydro, the head effect will tend to offset discounting and wastage effects. And head effects are likely to be more significant for pumped storage than for conventional New Zealand hydro, because an upper reservoir with a significant storage capacity, must also operate over a significant range of storage levels.[67] The most obvious cost of a pumped storage facility is incurred in creating the civil/mechanical system to raise water to a significant height. But there is also another cost incurred in pumping enough water up to fill the upper reservoir up to the minimum operating level.

That cost may be quite significant if pumping must be done at times when electricity prices are not close to zero, and it effectively forms part of the capital cost of the plant, to which the discount rate applies. On the other hand, that cost may be quite low, and adding further water may be a relatively cheap way of increasing average head levels, and hence average productivity, if the reservoir has steep sides, the

---

65 That is, until the ratio between the higher and lower prices equals RTLF.

66 To put this in perspective, note that an incremental unit will not really be "needed" to avoid shortage until all the units already stored in the upper reservoir have already been exhausted. If the discount rate was 7% and 3% of stock was lost each year, the effective discount rate would be 10%. If we were thinking of running the facility according to the traditional 1:20 design dry year criterion, the marginal opportunity might be 20 years away, on average, suggesting that the pumped storage facility manager might only be prepared to pay around 15% of the eventual expected value of whatever shortage/DSM cost might be expected, on average, across a wide range of future years, for an incremental water unit. Or, if an incremental water unit was held, on average, for 10 years before being used, the pumped storage facility manager should only be prepared to pay around 39% of its eventual expected value in pumping costs. Those estimates are obviously extreme, and very rough, but they do suggest the possibility that the effect may be quite significant.

This becomes a design issue, though. It will not be optimal to build a reservoir so large that it is not expected to ever be filled. So, conversely, we expect that any reservoir that is actually built will also be filled, at least occasionally.

67 Assuming that, unlike most New Zealand power stations, a pumped storage hydro plant would be fed directly from the upper reservoir, with no intervening open channel.

discount rate is low and pumping energy can be obtained at time of very low electricity price.[68] If the reservoir is large enough that the optimal set of storage trajectories is seldom constrained, it could conceivably be optimal for them to lie at the top of the reservoir range, if the head effect is strong, and the pumping cost and/or discount rate are low enough, but at the bottom of the storage range, if the pumping cost and/or discount rate is higher.

3.4. Stochastic management of stand-alone pumped storage

Section 3.3 discussed the idealised operation of a large pumped-storage hydro facility as if it could be optimised and/or operated in a perfectly competitive manner, without considering any direct interaction with a host hydro system. That discussion covered the management of predictable daily/annual cycles, on the assumption that the upper storage capacity was no more than would be required to support full utilisation of the facility's MW pumping/generation capacity in arbitraging over an annual horizon. Here we retain the assumption of stand-alone operation in an optimised/perfectly competitive environment, but extend the discussion to consider the impact of uncertainty in two distinct timeframes:

- The intra-annual timeframe, within which management of stochastic fluctuations in the supply/demand balance must be integrated with the management of the underlying daily/annual cycles discussed above; and
- The inter-annual timeframe, over which extra storage capacity could be utilised to manage year-to-year supply/demand fluctuations, due to hydrological variations, etc.

3.4.1. Real-time stochasticity

Electricity systems use "ancillary services" to manage real-time stochasticity in the supply/demand balance. A pumped storage hydro facility could be used to provide ancillary services, including regulation and all kinds of contingency reserve, in a manner broadly comparable to provision of ancillary services from batteries and/or other short-term hydro facilities. Pumped storage hydro (or batteries) could also play a rather different role, though, by standing ready to cover sudden fluctuations in wind/solar generation over a time frame longer than that covered by traditional ancillary services in the New Zealand market. This may be thought of as super-imposing a stochastic element over the basic intra-day cycle discussed above. Specifically:

- If electricity prices (are expected to) lie within the "dead-band", the facility should be neither pumping nor generating, in which case it could commence generating to cover a sudden supply shortage, or commence pumping to absorb a sudden supply surge.
- If electricity prices (are expected to) lie above the "dead-band", the facility should be generating, in which case, it could:
    - Reduce generation, and maybe start pumping, to absorb a sudden supply surge, if generating at more than its minimum rate; or
    - Increase generation to cover a sudden supply shortage, if generating at less than its maximum rate.
- If electricity prices (are expected to) lie below the "dead-band", the facility should be pumping, in which case, it could:

---

[68] In the limit, if pumping only needs to occur at times when the electricity price is zero, the pumping efficiency is actually irrelevant. So, the marginal cost of storing energy is zero, too, as is the "holding cost" implied by the discount factor. In theory, electricity prices can actually become negative. Typically, this is a result of localised transmission "spring-washer" response to loop-flow constraints in the transmission network, or inflexible plant preferring to operate at a loss, rather than de-synchronise, during intervals when electricity is in excess supply. In theory, that could mean that (inefficient) pumping was actually profitable, in its own right. But we will largely ignore the possibility that such conditions could become common enough to have any material effect.

- o Reduce pumping, and maybe start generating, to absorb a sudden supply shortfall, if pumping at more than its minimum rate; or
- o Increase pumping to absorb a sudden supply surge, if pumping at less than its maximum rate.

Although those responses do not sound significantly different from those involved in providing traditional ancillary services, there is a difference in the time frames involved, and perhaps the granularity of the response. Partially loaded pumped storage hydro units might provide significant responses in traditional 6/60 second ancillary service timeframes, but that response might be maintained over longer time frames and/or more substantial responses could be available in slightly longer time frames:

- Given a few minutes notice, we understand that units could be stopped/started to increase/decrease generation/pumping rates over a wide range; and
- Given a few more minutes, we understand that the entire scheme could be switched from generation to pumping mode, giving a maximum total net response band of approximately twice the plant's MW capacity.[69]

These possibilities may, or may not, prove to be an important factor determining the economic viability of any particular facility, but we will ignore them here, because they have little direct connection with the opportunity costing of stored water. Increased generation/pumping obviously implies extra water taken from, or added to, storage, but this will only happen when such responses are actually triggered. And, when that happens, the extra will probably be accounted for as part of regular market operations, once the response extends beyond the end of the market interval in which it is triggered. Most of the time, though, providing this type of "extended ancillary service" implies being on standby to respond. So, it is actually a way of obtaining significant value from a facility without using appreciable quantities of stored "fuel", and thus has only minimal interaction with setting the marginal value of that fuel via opportunity costing.[70]

3.4.2. Daily/weekly/monthly stochasticity

As discussed above, "real-time" stochasticity, within each dispatch interval merges fairly seamlessly into management of intra-day stochasticity, across multiple dispatch intervals. Such stochasticity has always been a major feature of the NZEM, extending out to weekly, and even monthly stochasticity, due to strong correlations creating persistent wet/dry flow and/or high/low demand sequences over those time scales. And it will become increasingly important as wind and solar capacity increases, with strong correlations creating persistent dull/bright and/or calm/windy conditions across the country, adding to, and correlated with, traditional stochasticity over similar time scales.

A pumped storage hydro facility should optimally respond to such mid-term fluctuations by deviating from its target daily/annual filling/emptying cycles in a manner very similar to that of conventional hydro facilities. There are some important differences, though:

- First, the loss-induced dead-band still applies, so the facility should never be generating unless the market price exceeds its MCG, or pumping unless the market price is below its MVP. Thus, in the absence of thermal:
  - o The first line of response to short term supply/demand balance fluctuations should still be by varying the output of conventional hydro facilities, which can be done without incurring the losses involved in cycling batteries, or cycling water through the pumped storage hydro facility;

69 Exactly twice, if pumping and generation MW are identical, but possibly more, if pumping inefficiency is (partly) manifested by requiring more energy to achieve the same flow rate for pumping as for generation, as assumed in Appendix A.

70 Opportunity costing concepts are relevant, though, inasmuch as the prices set for ancillary services in a co-optimised market always include the opportunity cost of displacing generation. And that opportunity cost, itself, reflects the difference between the electricity price in that market trading interval, and the offer which, under perfectly competitive assumptions, would be determined by the EMVS.

- Short term demand response (e.g. deferred water heating, and in future deferred EV charging) should also play a significant role; and
- If batteries have lower round-trip losses than pumped storage hydro, they would provide the next line of response;[71] but then
- The pumped storage facility should respond when that combined conventional hydro/battery strategy proves unable to hold prices within the dead-band implied by its RTLF.[72]

- Second, the same loss-induced dead-band also limits the degree to which the EMVSE in the pumped storage facility can be aligned with EMVSE levels in other major reservoirs, in its own island, creating a situation very much like that already existing between South Island and North Island storages. So, while optimal generation/pumping strategy will constantly work to bring storage ratios back into the balanced zone, random variations in hydro inflows, and supply/demand balances, will constantly tend to force storage ratios out of the balanced zone.
- Third, in this stochastic environment, forecasting becomes a significant issue, which should theoretically add even more dimensions to each reservoir's EMVSE surfaces, and hence help determine optimal joint release/pumping strategies for them all.[73]
- Finally, though, the controllability of pumping means that:
  - Ignoring any significant inflows to the upper reservoir, or assuming them to be managed under a protocol whereby they are passed through to preserve a status quo flow regime, there is no need to worry about any probability of spill arising as a result of holding the upper reservoir full; and
  - Provided there is sufficient water available in the host hydro system, it would actually be much more feasible to drive the upper reservoir's storage levels toward a pre-determined target, such as "full by the end of autumn", than it is for a conventional reservoir, where filling is very much driven by whatever weather patterns occur, determining hydro inflows, and increasingly also hydro generation requirements as (uncontrollable) solar/wind replaces controllable thermal capacity.

That kind of target focussed strategy does not become desirable just because it is feasible, though. The reality is that weekly/monthly/seasonal supply/demand balance fluctuations eventually morph into annual fluctuations. And it is quite possible that it will not be optimal to force storage to meet, say, an end-of-autumn target, if the year is turning out to be one in which supply is lower than the average assumed in setting the target.

In such a year, more and more of the forward simulations from which the summer EMVSE is set will fall short of filling the upper reservoir in autumn (and probably conventional reservoirs, too), implying that their EMVSE levels will already reflect the supply/demand balance expected in the upcoming winter. So, the EMVSE will progressively transition to a winter level even though storage may never actually reach its full level. That means, though, that the full storage range will never be used in such years.

So, we can think of increasing the design capacity of the upper reservoir to be large enough to optimally manage all intra-annual cycling and stochastic variation requirements in some years, and then in many

---

71 The relevance of round-trip losses needs to be considered carefully. The MDAG study suggests that, at least on summer days, electricity prices could be near zero much of the time, and that may make losses (almost) irrelevant, much of the time. That would seem to be the case for both battery and pumped storage hydro in daily cycling mode, over summer, and for long-term pumped storage hydro filling over summer, as in the illustrative examples considered in Appendix A. If energy is truly in excess supply, then the fact that one storage technology absorbs more of it than anther does not seem like it should be a matter of concern. MDAG suggests that winter prices may be much higher, though, and that may make losses a critical factor for batteries, and also for pumped storage hydro in a daily cycling mode.

72 Or pumped storage hydro would respond before batteries, if its round-trip losses are less. In reality, a mixed response is likely to be optimal, because the effectiveness of various responses will depend on the location of the responding storage relative to the drivers of supply/demand imbalance. Since batteries will typically be sited much closer to loads than any particular planning horizon facility, they are more likely to be first responders, because their use does not incur intervening transmission loses, or encounter intervening transmission constraints.

73 Or, the effect may be approximated by reading each MVS surface at an "effective" storage vector, adjusted to account for flows etc. expected to occur over the near future.

years, and eventually in all years. And, while there may be no clear dividing line, at some point along that spectrum, our focus must shift from the primarily intra-annual analysis discussed above, to the broader inter-annual analysis discussed below.

3.4.3. Inter-annual stochasticity

If the storage capacity of a reservoir is large enough to accommodate the optimal annual arbitrage cycles of a significant proportion of the years in our simulation sample, we may consider that reservoir to be "inter-annual" rather than merely inter-seasonal, or "intra-annual". The upper reservoir of the mooted pumped storage hydro facility may fall into that category, but so may some existing reservoirs, if they are able to share more of their current intra-annual arbitrage role with a pumped storage facility.[74]

Either way, the effect would be to increase the overall ability of the system to absorb and manage year-to-year fluctuations in the supply/demand balance, such as those traditionally attributed to wet/dry year variation. In principle, we can think of building up storage in "wet" years so as to be able to cover deficits in "dry" years as being just another form of arbitrage. This kind of arbitrage will not imply a predictable regular cycle, though.

While our reservoir should be more likely to be full around autumn, and/or empty around spring, we can not expect it to regularly reach, or perhaps even approach, such limits. In fact, we should expect it to follow something more like the stochastic ideal, building up storage over wet/windy/sunny years so as to be ready to run it down to avoid deficits in dry/calm/dull years. That optimal operation could be determined dynamically, by re-running a stochastic optimisation every week, say. Or (in principle) a suitable set of EMVSE surfaces/guidelines could be pre-computed to guide decision-making over the whole feasible range of storage levels. Either way, the optimal strategy will surely be to build/hold relatively high stock levels during wet/normal years because:

- The national costs potentially incurred during non-supply events will be high;
- Once thermal capacity is removed, all storages will tend to fall much more quickly towards levels at which such costs would have to be incurred; and
- Holding high stock levels in the upper reservoir actually implies no risk of spillage, so the EMVSE of a full reservoir may remain at a much higher level than for a conventional reservoir.[75]

Thus, the stochastic equilibrium of simulated storage trajectories can be expected to show them all falling steeply over winter, but most returning to high levels by the end of the next summer. The fall of some will be steeper and deeper, though, perhaps even emptying the upper reservoir, with a slower multi-year re-fill phase to follow. And the motivation to fill the upper reservoir will be inhibited by discounting and wastage, as discussed in Section 3.3.4 above.

The EMVSE in a reservoir operated that way will be its MCG, i.e. the minimum price at which it is optimal for storage to be run down in order to support the supply/demand balance. But, after adjustment for losses it will also set MVP, the maximum price at which it is optimal for storage to be built up by paying to pump water uphill. So, these two must be in long-term equilibrium over intra-annual and inter-annual planning horizons. Clearly, EMVSE will rise whenever storage falls far enough below its target equilibrium range to justify an inter-annual effort to re-build stock levels. There is no reason to expect much inter-seasonal pattern to EMVSE, though. If the capacity is large enough to easily accommodate the range of trajectories needed to manage both intra-annual cycling, and the general run of variations on the inter-seasonal supply/demand balance pattern:

- Pumping would occur whenever electricity prices fell below MVP, mainly in summer; and
- Generation would occur whenever electricity prices rose above MCG, particularly in winter;

---

74 Note that this is not a function of absolute size. Historically, Hawea could be regarded as an inter-annual reservoir because, at that time, its operating rules allowed storage levels to be varied over a range that was quite large relative to its fairly modest inflow rate.

75 The fact that it never faces any risk of spill means that the EMVS can not to fall to zero unless the storage is absolutely full, but it could still be calculated from relatively low value opportunities available over a fairly short period, if storage is approaching the full level.

- But EMVSE itself would only vary in response to storage being relatively higher/lower than expected for the time of year, and not as a function of the time of year itself.

3.5. Achieving balance with other storages

3.5.1. Storage balancing in a multi-reservoir environment

The above discussion implicitly assumed that the entire national system was dependent on a single pumped storage hydro reservoir. But such a facility would really be additional to the current national portfolio of both short and long-term reservoirs. Thus, a national optimisation would provide an integrated operational plan, involving balanced cycling of all these storages, over both daily and annual planning horizons.

Based on the above discussion, a single EMVSE could be determined for an assumed national reservoir as a function of that reservoir's storage level. Section 3.2 of our report to MDAG discusses how to generalise that logic, but concludes that, if we are simultaneously optimising management of n reservoirs we need to think in terms of an n-dimensional EMVSE surface for each of those n reservoirs. That is, the marginal value of water stored in any particular reservoir now depends on the water storage level in, and opportunities available to, each reservoir. In fact, it may well depend more on storage levels in some other (larger) reservoirs than it does on its own.

If we simplify our analysis by just considering the interplay between our pumped storage facility and an aggregate national (conventional inter-seasonal) storage reservoir, the situation would be broadly analogous to the way in which SPECTRA always balanced South Island and North Island storages, under the status quo.[76] If there had been no inter-island loss factor, it would have computed a single "balance guideline" along which the North Island and South Island EMVSE levels were equal. As discussed in our report to MDAG, South Island and North Island hydro would have swapped merit order positions when that guideline was crossed, so:

- On one side of the guideline, South Island storage would be calculated to be relatively fuller than North Island storage, causing South Island release to be preferred over North Island release, and thus bringing the storage ratio back towards the balance line; while
- On the other side of the guideline, North Island storage would be calculated to be relatively fuller than South Island storage, causing North Island release to be preferred over South Island release, and thus bringing the storage ratio back towards the balance line.

In reality, there are losses on the inter-island link, and two guidelines were always calculated, one at which the South Island/North Island EMVSE ratio equalled the marginal loss factor, and one at which South Island/North Island EMVSE ratio equalled the inverse loss factor.[77] Then, optimal balancing was achieved, as above, except that there was a dead-band within which the losses were too high to make balancing worthwhile. Thus. the effect was to move the storage balance back towards the "more-or-less balanced" zone, whenever it moved outside that zone. But this discussion is not just about how

---

[76] The SPECTRA model was a development of PRISM, originally introduced by the Ministry of Energy c. 1984, and subsequently used for many years by ECNZ and a number of electricity sector participants. Its mathematical underpinnings are similar to those of SDDP, but instead of focussing on approximating the MWV surface within a region visited by some scenario set, it directly constructs an approximately optimal MWV surface in a way that is very efficient for 1-2 reservoirs, but computationally and conceptually challenging for higher-dimensional problems. See:
E.G. Read, J.G. Culy, T.S. Halliburton, and N.L. Winter: "A Simulation Model for Long-term Planning of the New Zealand Power System", in G.K. Rand (ed.) *Operational Research 1987*, North Holland, p.493-507

E.G. Read and M. Hindsberger "Constructive Dual DP for Reservoir Optimisation" in S. Rebennback, P.M. Pardalos, M.V.F. Pereira and N.A. Iliadis (eds) *Handbook on Power Systems Optimisation* Springer, 2010, Vol I p3-32

[77] Thus a 10% marginal loss factor in each direction would imply balancing guidelines at north/south EMVSE ratios of 0.9 and 1.11.

SPECTRA worked, it describes the way in which any optimisation would work, and any competitive market should work, to balance storage levels as closely as they can/should be balanced given the physical balancing mechanisms available.

In that case, inter-island transfer was the key balancing mechanism available, and indeed it still is, and will remain so. At the same time, any optimisation, or perfectly competitive market, should theoretically be able to achieve perfect EMVSE alignment (± intra-island losses) between conventional reservoirs with similar storage/flow capacity ratios within the South Island. Pumped storage hydro would add another balancing mechanism, very similar to inter-island transfer, though. The inter-island situation is more complex, because the South Island and North Island both have their own inflows, loads and constraints to be respected, but we can think of the entire South Island storage system as traditionally acting very much like a large "upper" reservoir, with limited and lossy pumping/generation capacity acting very much like limited and lossy HVDC capacity.[78] Thus:

- Rather than "pumping water uphill", the North Island has traditionally transferred energy to boost South Island storage by southward transfer displacing South Island hydro generation, at times when there has been a relative surplus in the North Island, and shortage in the South Island; then
- Rather than "running water downhill" to generate, the South Island has traditionally transferred energy to meet North Island loads (and/or preserve North Island storage), at times when there has been a relative surplus in the South Island, and relative shortage in the North Island; but
- Re-balancing has always been limited because energy could be "trapped" in particular storages by limits on both inter-island (or indeed intra-island) transfer capacity, and the release/generation capacity of individual hydro systems; in particular
- Once the reservoir with the lower ratio of storage capacity to inflows (i.e., the North Island aggregate reservoir in this case) is full, then further storage can only be accommodated in the comparatively larger reservoir (i.e., the South Island aggregate reservoir in this case) in an "unbalanced" fashion; while, in any case
- Round-trip losses of around 20% have traditionally meant that North Island and South Island marginal water values could differ by that much, at any time, before inter-island re-balancing action was justified; [79] and
- And, of course, the link capacity has limited the extent to which that balancing action could be effective. So, both losses and link capacity limit the extent to which:
  - Transfer to/from South Island storage could be used as a strategy to by-pass North Island storage limits in equating North Island marginal water values over time; and/or
  - Transfer to/from North Island storage could be used as a strategy to by-pass South Island storage limits in equating South Island marginal water values over time.[80]

Because the situation is so closely analogous, we can expect essentially similar outcomes from the optimal balancing of pumped storage hydro operation with conventional hydro reservoir management. Simplifying the situation down to consider the balance between a large pumped storage hydro facility and conventional South Island hydro storage:

- Centralised optimisation, or perfectly competitive market interaction, would try to ensure that the storage held in the upper reservoir was kept in balance with other South Island storage levels by:

[78] The loss situation is symmetrical, in this case, making the choice of which island is uphill vs downhill rather arbitrary. We have made the South Island "uphill" because it has a simpler pure hydro system, with more storage.

[79] The simple ±x% marginal loss factor typically assumed in SPECTRA was never a perfect model of losses for the original HVDC technology, and would be less so for modern technology. But we retain it here for illustrative purposes, and because it implies round-trip losses very similar to those for pumped storage hydro. And, while we have assumed all losses to occur during the uphill pumping phase, it is really the round-trip losses that count.

[80] No actual water is moved between islands, but the effect is as if incremental energy was passed from one island to the other, being effectively carried as extra water in South Island reservoirs over autumn, in order to "get around" the storage capacity limit preventing any more being carried forward in the North Island.

  - Preferring pumped storage generation to releasing water from inter-seasonal storages to support conventional generation, and/or reducing pumping, when the pumped/conventional storage ratio gets too high, as indicated by the pumped/conventional EMVSE ratio being less than 1/RTLF.
  - Preferring conventional hydro generation to pumped storage generation, and/or increasing pumping, when the pumped/conventional storage ratio gets too low, as indicated by the pumped/conventional EMVSE ratio being more than RTLF.

- But round-trip losses of around 20% will always mean that the pumped storage EMVSE can differ from other South Island marginal water values by that much, at any time, before re-balancing action would be justified, and that will also limit the extent to which transfer to/from the upper pumped storage reservoir can be used as a strategy to work around limits on conventional hydro storage, in equating South Island marginal water values over time.[81]
- And, re-balancing will also always be limited because energy may be "trapped" in particular storages by limits on both pumped storage pumping/generation capacity, and the release/generation capacity of individual conventional hydro systems.
- In particular, once the reservoir with the lower ratio of storage capacity to inflows (perhaps the aggregate South Island conventional storage reservoir, in this case) is full, then further storage can only be accommodated in the comparatively larger reservoir (perhaps the upper reservoir of the pumped storage facility, in this case) in an "unbalanced" fashion.

3.5.2. Capacity balancing in a multi-reservoir environment

Finally, if we reconsider the hypothetical process of increasing the design capacity discussed above, but now assuming that this pumped storage capacity is additional to the existing national storage portfolio, all operating in a balanced way, as above:

- Expanding the storage/pumping/generation capacity of the pumped storage facility would not only reduce summer/winter EMVSE ratios in that reservoir, but in all other inter-seasonal reservoirs;
- The fact that all of these reservoirs would then be working in concert to reduce summer/winter price differentials, means that we would more quickly reach the point beyond which there would be no value in further expanding the pumped storage reservoir capacity for cyclic inter-seasonal arbitrage purposes; although
- Increasing capacity beyond that point could further assist in dealing with stochastic elements of the situation, including wet/dry year variations.

At the same time, the marginal value of increasing any other storage capacity would also be falling, and could conceivably reach zero, for some existing reservoirs, indicating that they, too, had enough storage capacity to accommodate their optimal inter-seasonal storage cycles without being constrained by storage capacity, and thus could maintain EMVSE "as constant as possible", given changing expectations, over that annual cycle.[82] If so, the pumped storage hydro facility would have added more than enough storage/pumping/generation capacity to compensate for the loss of inter-seasonal thermal storage/trading capacity, leaving New Zealand with a more robust system, overall. But that outcome is frankly unlikely, and would not be optimal unless the marginal cost of building pumped storage capacity was very low.

[81] But note that, because this is a ratio test, quite a fine balance can be achieved when absolute electricity price or EMVSE levels are low. In the limit, it will be easy to justify pumping for quite marginal gains whenever electricity prices are (near) zero, or even negative, because the alternative to pumping is some form of (probable) spill.

[82] Note that arbitrage through the pumped storage hydro facility can only directly reduce the winter/summer EMVSE ratio in some other reservoir to, at most ± RTLF. But that does not preclude the possibility that, by reducing electricity price differentials to that level, the additional pumped storage capacity will reduce profitable arbitrage options down to a level at which conventional reservoirs with enough MW capacity can arbitrage away the remaining summer/winter differential in their own EMVSE.

4. Embedded Pumped Storage Operation

4.1. Introduction

Chapter 3 discussed the idealised operation of a large pumped-storage hydro facility as if it could be optimised and/or operated in a perfectly competitive manner, without considering any direct interaction with a host hydro system. Here we retain the assumption of operation in an optimised/perfectly competitive environment, but generalise the discussion to cover more realistic configurations in which the pumped storage hydro facility is embedded in a host catchment that is being managed to support its own conventional hydro scheme, potentially both upstream and downstream of the pumped storage facility's intake/discharge point.

Unlike batteries, pumped storage hydro is a very site-specific technology. So, we first describe the kind of pumped storage hydro facility we have in mind, and to which this theoretical discussion might ultimately be applied. But we then proceed to work through a series of simpler configurations to build up a picture of the various interactions involved, with a view to identifying the kinds of situations in which a co-operative working agreement might become important, and those in which it may not be necessary at all.[83] Finally, we briefly discuss the forms such an agreement might take, if it did prove to be necessary or desirable, and expand on that discussion in Appendix B.

4.2. Reference system

Rather than tackling analysis of any particular real system, it seems more helpful to develop some insights into how pumped hydro systems should operate, in general, and to identify areas and issues that should be examined more carefully when analysing any particular development proposal. Figure 1 shows a simplified, but reasonably general, topology for the generic system we will discuss.

For simplicity, we assume that pumping is used to raise water from a downhill "host system" to a dedicated upper reservoir, and generation then occurs when it is run down again, to the same location.[84] We will ignore the possibility that the upper reservoir of the pumped storage facility might also be embedded in an upper host catchment with enough natural inflow to significantly impact the pumped storage operation.[85] We will also ignore the possibility that the pumped storage facility may discharge to a different catchment, or sub-catchment.[86]

---

83 Appendix A takes a rather different tack, working through some highly stylised numerical examples, in order to get a feel for the relative importance of various capacity dimensions and siting considerations that may materially impact on the assessment and operation of potential developments, and perhaps limit the choice of operational and organisational regimes that might be appropriate, in such cases.

84 The opposite situation, in which water is run down to a dedicated reservoir, and then pumped up again, is technically possible, but seems unlikely in New Zealand.

85 The discussion here would thus strictly only apply if the arrangement was that natural inflows to the upper reservoir of the pumped storage facility were immediately passed through to become outflows in the upper catchment. But it would still apply, with minor modifications if some of those flows were captured in the upper reservoir of the pumped storage facility, and managed to better achieve the objectives of the pumped storage hydro facility and/or the upper catchment manager. Such management could also include a net transfer of water from the upper catchment to the lower "host system" catchment, or vice versa. But this would not materially alter our conclusions unless the flow volumes involved were significant, relative to those drawn from the lower catchment.

86 Pumping can actually be used in variety of situations, other than "cycling" to increase generation capacity at critical times, as described here. Appendix A notes the possibility that a (typically low head) pumping operation could be used to enhance storage capacity by transferring water from one reservoir to another reservoir with greater spare storage capacity, with the intent that it later be returned, or released into the same river system, as was once proposed between Wanaka and Hawea. Pumping could also be used to transfer water from one catchment to another, where it would have greater value because it can be stored for longer and/or will pass through more downstream generation capacity.

Thus, we will use “upstream” to mean that part of the (lower) host catchment from which water flows down to the pumped storage facility’s charge/discharge point[87], and “downstream” to mean that part of the (lower) host catchment to which water flows down from that point. We have also simplified the diagram by aggregating all upstream storage/generation capacity into a single representative reservoir/station, and similarly downstream. And we will initially ignore flow delays, which could actually be a significant factor modifying the conclusions we reach about the compatibility between the pattern of release timings desired by one manager vs another.[88]

Tributary flows are not shown on the figure, but can occur at any point in the host system. Intuitively, we might expect tributary flows upstream of the buffer storage location to reduce any problems with respect to water availability for pumping, and tributary flows downstream from that point to increase any problems with respect to downstream congestion. That is not necessarily so, though, because flow limits at any point on the river will have been set in relation to natural flows at the point, and will steadily increase as we move downstream, and more and more tributary flows are collected. So, tributary flows do not necessarily increase the average freeboard/headroom available for pumping/ generation, while the fact that they are uncontrolled can be expected to increase the frequency with which the limits are reached.

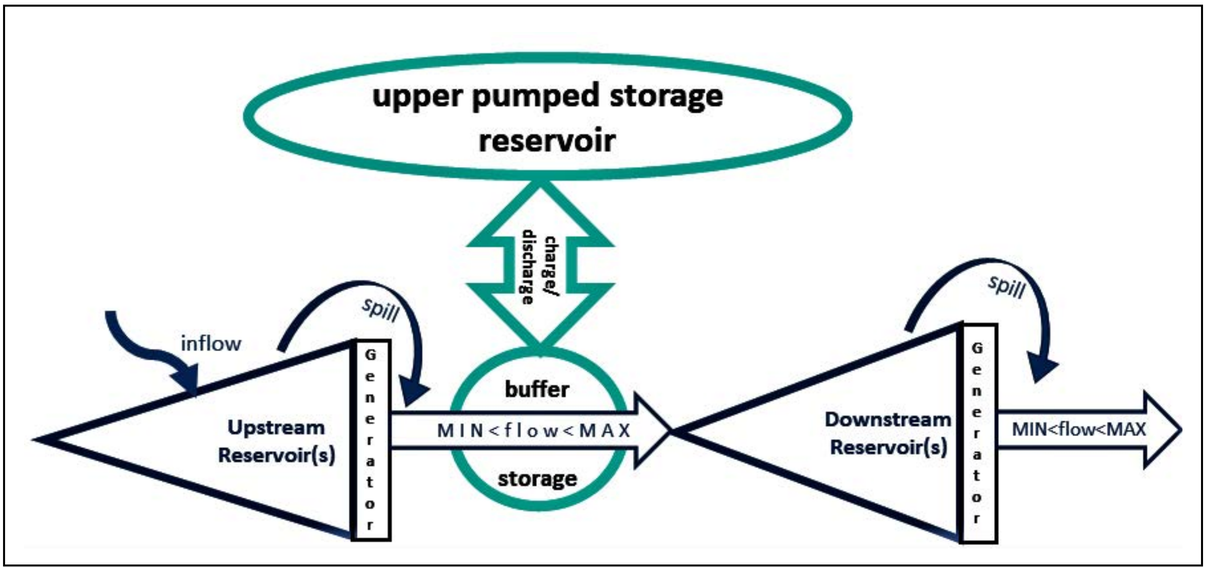

Figure 1: Pumped Storage Embedded in Host System

Our discussion will proceed by considering a range of alternative configurations created by removing various elements to form even simpler systems. Importantly, those various re-arrangements are not intended to suggest anything about the order in which facilities might be built.

While we distinguish between a pumped storage facility manager and a host system manager, their assumed relationship is essentially symmetrical, in physical, mathematical, and economic terms. Irrespective of the historical order of development, or organisational arrangements, the joint system, once built, will form an integrated whole, the operation of which should ideally be jointly optimised (directly, or indirectly in response to competitive market signals) to maximise collective benefit. Much of our discussion will focus on the reasons that may not happen, though, and the implications that may have for the national benefit actually delivered.

Thus, we are concerned to identify situations in which some agreement, or market mechanism, may need to be found to coordinate activity sufficiently to deliver something like the national economic potential of particular developments. Section 4.8 discuss a number of possible types of agreement that might be used to increase national benefit by improving coordination. But the kind of agreement that

87 Also referred to as the “buffer storage location”, irrespective of whether there is any significant buffer storage, or not.

88 For example, a 12-hour delay will typically make upstream/downstream release patterns that were perfectly in synch, perfectly out of synch, and vice versa.

might be reached between the parties may depend heavily on the fact that, in the New Zealand context, development options are likely to involve inserting a new pumped storage facility into a host system already operated by an incumbent manager.

Thus, historical precedence may play a very important in the real world, but it needs to be understood that it will play no role in any formal optimisation of operational strategy. So, we extend our discussion as far as to identify, and comment on, the areas in which tension is likely to develop, and mis-understandings and mis-valuations arise, between incumbent host system management, pumped hydro system management, and national benefit optimisation.

4.3. Host system status quo

Before discussing the management of pumped storage, we should perhaps consider how the host system would operate, in the absence of any pumped storage facility. With or without a buffer storage, the host system shown in Figure 1 is no different from any other hydro chain. Thus, the pattern of marginal water values will be as discussed in Section 2.5 above, basically decreasing down the chain, and perhaps responding dynamically to various stochastic events, but exhibiting some broadly predictable cyclic patterns over time. But the details depend heavily on the relative capacities of the storages involved. There may actually be several upstream or downstream reservoirs and/or generation stations in the host system but, accepting the simplified representation assumed in Figure 1:

- The marginal water value will generally be zero, at all times in the river reach below the (last) downstream reservoir.[89]
- The downstream storage capacity is likely to be low, in most New Zealand systems, and the downstream reservoir could well operate over a constrained daily cycle, with the EMVSE at night being lower than that in the day.
- The bulk of host system storage is most likely to be in, or close to, the furthest upstream host system reservoir, and that reservoir is likely to be able to manage storage over a weekly to annual planning horizon. So, its EMVSE will probably vary very little over most daily cycles.
- Its EMVSE may vary significantly over longer time cycles, though, and is likely to at least have distinct summer (low) and winter (high) values.
- If there are intermediate upstream reservoirs, they are likely to be smaller, and may well cycle daily.

As for any conventional (gravity driven) river chain, the basic logic is that upstream reservoirs directly influence the physical operations of downstream reservoirs, whereas downstream reservoirs signal back the economic implications to upstream reservoirs. Thus:

- Physical releases from the upstream reservoir will inevitably flow down to downstream reservoirs, and thus influence EMVSE values there by forcing physical schedule adjustments; but
- That impact should be accounted for by the upstream reservoir manager, who must recognise that the effective economic contribution made by any physical release, and hence the EMVSE of the upstream reservoir, includes the EMVSE of water delivered to the downstream reservoir.

Conversely:

- Physical releases from the downstream reservoir can not have any physical impact on upstream reservoir operations; and
- Upstream EMVSE values play no role in determining downstream EMVSE values, because the downstream water can not be returned to the upstream reservoir.

Thus, when we talk about downstream congestion forcing a reduction in upstream release rates, for example, we do not (generally) mean that water is physically prevented from leaving the upstream reservoir. We mean that the implied penalty for violating the downstream maximum flow limit is high

89 Ignoring the possibility that downstream flows could have a positive marginal economic value, if supporting uses outside the electricity sector, or possibly a negative marginal economic value, under flood conditions.

enough to make the downstream EMVSE effectively negative enough to make it uneconomic for the upstream reservoir to release any more.[90]

Of course, if a pumped storage development was small, then it might hardly disturb the host system environment, and we could thus analyse its optimal behaviour in terms of interaction with this status quo EMVSE pattern. But pumped storage developments large enough to be worth analysing, at the national level, will inevitably have major physical impacts on the host system, and radically alter the EMVSE pattern within that system. So, we must analyse mutual interactions between the EMVSE in the upper reservoir of the pumped storage facility, which must exist in equilibrium with upstream/ downstream host system EMVSE levels determined in this new joint system environment.

Unfortunately, though, the key function of pumped storage operations is to break down the strict upstream/downstream logic outlined above, with different upstream/downstream relationships applying in pumping mode than in generation mode. Thus, our intuitions derived from analysis of traditional non-pumping hydro systems are not necessarily correct. So, we will start from an analysis of the situation without any host system development, and systematically work through the implications of adding various elements to the host system configuration.

4.4. Run-of-river host system

4.4.1. Introduction

First, consider a simplified version of the system in Figure 1, in which the host system has no significant upstream or downstream storages, and hence no control over flows. Although that may mean that pumped storage operations are tightly constrained, it does not necessarily imply the need for a bilateral "coordination" agreement, because the host system manager can not take any actions to comply with such an agreement. Thus, the issue here is more one of the unilateral actions the pumped storage facility manager might take, and the conditions that might be imposed upon the pumped storage facility manager to limit such actions.

In fact, if there was no pumped storage facility, we would not expect any minimum or maximum flow limits to be set on any river reach either, because there is nothing the host system manager can do to comply with them. This changes once we add a pumped storage facility, though, because now flows can be controlled, and we should expect that limits will be set. Upstream flows can still not be controlled, but we have to allow that there could be simple upper/lower bounds on flows immediately downstream from the charge/discharge point.[91] In practice, the effective limits at that point may be set by minimum/maximum flows allowed at points further downstream, minus any intervening tributary flows, but that does not change the discussion here, in the situation where the host system manager can not control flows at any point in between.[92]

90 Note that real-world reservoir managers most likely see these issues in much more physical terms, but we are talking here about the optimal "shadow price" structure that would be calculated within an optimisation model of the situation and, when we say "negative enough", we mean "just negative enough", because that shadow price will be set to achieve the goal of limiting upstream release to the required level and no more, but also no less.

91 Noting that the absence of such limits under the status quo does not mean to say that they could not be imposed, as soon as a pumped storage development threatened to increase peak flows, and/or variability, above current levels, or decrease them below current levels.

92 For simplicity, we will ignore the likelihood that different limits may apply in recognised "flood" and/or "drought" conditions than in "normal" times.

4.4.2. Undeveloped host system

The simplest possible scenario, here, is that the host system is completely undeveloped, in the sense that there is no upstream/downstream or buffer storage able to control flows, and no upstream/downstream generation capacity, either. The pumped storage facility must operate entirely by drawing water from, and discharging to, a river reach in the host system, with no appreciable buffer storage at all.[93] As in Section 3.3, we assume that pumping inefficiency is manifested by requiring more energy to achieve the same flow rate for pumping as for generation, because this simplifies of certain discussions here, and particularly in Appendix A.[94]

This is not quite the same as the "stand-alone" scenario discussed in Chapter 3, though, because the natural river flow rate, and the upper/lower limits that might be set on flows downstream from the charge/discharge point, can constrain pumped storage operations.[95] If we ignore any natural flows into the upper reservoir of the pumped storage facility itself, then the effective limits on pumped storage operations will vary from day to day, as follows:

- The maximum rate at which pumping can occur, at any particular time, t, is not just limited by the pumped storage facility's pumping capacity, but by the flow freeboard, that is the difference between the flow available at that time from upstream, *HSF,*[96] and the minimum flow required immediately downstream from the buffer storage location (*HSFmin*):
  HSF(t) - HSFmin
- The maximum rate at which generation can occur, at any particular time, is not just limited by the scheme's generating capacity, but by the rate at which water can be returned to the river, given the headroom available at that time between the upstream flow and the maximum flow allowed immediately downstream from the buffer storage location (*HSFmax*):
  HSFmax - HSF(t)

These flow limits could be quite restrictive, in some cases, but they need not be:

- If the maximum pumping requirement is less than the freeboard between the natural flows and the minimum flow requirement, there may (almost) always be enough flow to pump from, even without a significant buffer storage.
- If the maximum generation discharge is less than the headroom between the natural flows and the maximum downstream flow limit, there may (almost) always be enough headroom to allow full generation, even without a significant buffer storage.

We may think of systems satisfying those criteria as "effectively open", in the sense that the pumped storage facility can operate in an unconstrained fashion, as if it was a "stand-alone" facility, with no host system constraints to consider. Or (at times) its operation may be significantly constrained, but there is nothing any party can do about it, so no "coordination" is required. The situation becomes more complicated, though, if the host system has itself been "developed", or is to be developed, for power generation.[97]

---

[93] That is, just enough to meet technical requirements, but not enough to allow average outflow from the buffer storage to differ from average inflow to the storage, over a period of an hour or more.

[94] In reality, pumping inefficiency will most likely be manifested by requiring the same amount of energy to achieve a lower flow rate for pumping than for generation. But we do not believe that changes any of the conceptual conclusions in this report.

[95] As noted in Section 2.5, we ignore the possibility that the critical constraints in the host system could be limits on the rate of change of flow, rather on flow rates, per se. As before, such constraints would act like upper limits in some periods, but like lower limits in other periods.

[96] This will be just the "natural flow", including any flow passing through upstream generation plant, and any upstream "tributary" flows, because neither is controllable.

[97] Similar, but different, considerations would apply if it was developed, in whole or in part, for other purposes, but we will focus solely on generation-oriented developments here.

4.4.3. Accounting for upstream generation capacity

Of all variations, upstream generation capacity is the easiest to account for, because it makes no difference to optimal, or even feasible, pumped storage operations.

The presence of a pumped storage facility does add considerable value to upstream host system water, which can now be stored and utilised to capture some potentially much higher value in some future period. And that may considerably enhance the value of upstream control structures that could capture and store that water, which is now more valuable.

But we are assuming there are no upstream control structures, in this scenario. And, no matter how much generation capacity there may be, upstream, it can not impact the flows available to the pumped storage facility unless it also has associated storage that can be used to modify the natural flow pattern. Nor can pumped storage operations have any physical impact or economic influence on upstream generation activity, for the same reason.

4.4.4. Accounting for downstream generation capacity

Downstream generation capacity is a different proposition, though. Physically, it makes no difference, inasmuch as the pumped storage facility still receives the same flow, and its discharge is only constrained by downstream flow limits, as above. But the presence of downstream generation capacity changes the economics of the situation, and may also imply tighter flow limits.

While the pumped storage manager could arguably ignore the downstream system, the host system manager definitely can not ignore the pumped storage facility:

- Adding the pumped storage facility means that flows in the entire catchment downstream of that facility are now controllable, and that potentially adds considerable value to the downstream host system assets.[98]
- Physical control is entirely in the hands of the pumped storage facility manager, though, and the question is whether they would, or should, manage the situation entirely for the benefit of that facility, or also account for the implications for the host system?
- Clearly, national benefit maximisation requires the latter, and it will presumably occur if the whole system is controlled by the same manager.
- But there is no mechanism to enforce that outcome if the two systems are under separate management.

The question is, though, whether that matters very much. And the answer depends on the degree of alignment between the natural objectives of the two systems, and the relative productivity of water passing through the two systems. Fortunately, in this very simple case, the objectives of the two managers are roughly aligned:

- The pumped storage facility will want to pump/generate when electricity prices are low/high; and
- That will reduce/increase flows flowing down to downstream generators, possibly to their minimum/maximum; and (ignoring delays)
- That will reduce/increase flows passing through those downstream generators, to produce power at exactly the times when they would have wanted to minimise/maximise those flows, if they could control those flows themselves.

As a result, downstream generators should be better off than they would have been, without the pumped storage facility, and the gain they make is a genuine improvement in national welfare. But the alignment of incentives is not perfect, and nor is the gain as great as it could be, and arguably should be, because the maximum downstream flow limit may be much higher than the maximum flow rate downstream generators can actually utilise for generation. And, since there is assumed to be no way to control

98 It also adds considerable value to the upstream host system water, which can now be stored and utilised to capture some potentially much higher value in some future period.

downstream flows, the potential energy/economic value of the excess flow will simply be lost as it spills past those downstream generators.

At this point, the order in which developments occur may make a real practical difference:

- If the host system and pumped storage facility were designed as an integrated catchment development, the downstream generators would presumably be designed to utilise something like the maximum discharge from the pumped storage facility, plus reasonably high natural flows, and the losses may not be too great.[99]
- If the pumped storage facility is retrofitted into an existing host system, though, the downstream generating stations would only have been designed to utilise the natural flows, without any additional flow from pumped storage arriving at the same time. So, the losses may amount to a significant proportion of the potential downstream energy value of the pumped storage discharge.
- The significance of that loss will depend heavily on the relative productivity of water passing through the two systems, though. Pumped storage facilities tend to be built to take advantage of high "head" situations, in which water can be pumped to a great height, and then produce a great deal of energy, per unit of water. It would not be surprising for the head of a pumped storage facility to be 10 times as much as that of a typical (more-or-less) run-of-river hydro station in the host system.[100] So,
  - If there is only one such station downstream, the economics of pumped storage operation should dominate that of the host system. Ignoring delays, power generated by both stations will receive the same high price at times when the pumped storage station is releasing strongly, and the total generated per unit of water released will fall from 110% of the pumped storage generation rate to 100%, when the downstream utilisation limit is reached.
  - That efficiency drop is probably much greater than any due to changes in unit loading in the facility itself. So, from a national benefit perspective, it would seem irrational and clearly sub-optimal for the pumped storage facility manager to optimise loadings to maximise efficiency of the pumped storage facility itself, while ignoring this downstream effect. But we can imagine the inefficiency being considered "acceptable", if there were larger issues at stake.
  - The inefficiency would be much more significant in other cases, though. It is really the total combined head of all downstream stations that should be traded off against that of the pumped storage facility, and it is not hard to imagine situations in which that combined head is maybe as great as that of the pumped storage facility. And that would mean a 50% efficiency loss, on marginal releases, once pumped storage discharge (plus natural flow) exceeds the downstream utilisation limit.
  - In all cases, accounting for this effect would probably not restrain pumped storage generation during periods when electricity prices are at their very highest, but it would tend to moderate releases to keep downstream flows within utilisable limits in shoulder periods, and so spread releases out over more of those periods.

Beyond that economic effect, the addition of pumped storage discharge on top of natural flows could potentially exceed the design limits of downstream spillways, when natural flows are already high. So, there could be times when the pumped storage facility, as the only controllable element in the catchment, would be required to throttle back its discharge to relieve pressure on downstream dams, irrespective of electricity prices. Indeed, we would not be surprised if it came under pressure to pump at such times, even if electricity prices are high.

So, even though the pumped storage facility may be physically able to ignore downstream conditions when determining its pumping/generation strategy, we expect it would actually be legally and morally required to account for and manage them. These kinds of restriction are not additional to those assumed above, though. They just mean that if *HSFmax* has to be set to keep downstream flows within limits implied by spillways designed to only accommodate natural flows, those limits are likely to constrain

---

[99] And, in that case, the entire system would probably be under integrated management anyway, so the pumped storage operations would be optimised to account for the economics of downstream generation.

[100] As assumed in the core example of Appendix A.

pumped storage operations more often, and perhaps much more often, in this retrofit situation than in a catchment developed as an integrated whole.[101]

In other words, the "effectively open system" scenario suggested above seems much less likely in this retrofit situation. Adding a buffer storage would help to de-link the two systems, and hopefully reduce the need for a coordinated optimisation of joint operating strategies, as discussed in Section 4.6 below. But the two systems may be tightly enough coupled, often enough, to make some kind of joint operating agreement desirable, if not necessary.

4.5. Accounting for downstream storage

With no pumped storage facility, storage in the downstream river chain can perform two basic functions:

- Storage (particularly near the top of the chain) can allow the host system manager to accumulate flows for release at times when electricity prices are higher, thus increasing the variability of the river flow, typically in a daily cycle; and
- Storage (typically further down the chain), can allow the host system manager to smooth out the fluctuating flow pattern so created, back towards a more natural or stable pattern, acceptable for release from the last hydro station in the chain.

At the same time. these storages, when full, serve to hold water at a higher elevation, thus increasing its head, and hence release productivity. In fact, a "fully developed" chain would consist of a series of storages, each releasing directly into the next, with no intervening river reaches. This maximises aggregate release productivity, because power is generated when water is released in a series of steps, the total head of which adds up to the entire elevation difference between the top and bottom of the chain.

And that has the side effect of effectively eliminating river reach flow limits between the top and bottom of the chain, too. Thus, the flow management problem may be transformed into a storage management problem, except downstream of the last dam in the chain.

As above, upstream generation is irrelevant if there is no upstream storage to control it, so the controllable system assumed in this section effectively starts with the pumped storage facility, and/or its buffer storage. And there are two possible configurations:

- If there is a river reach between the charge/discharge point and the (first) downstream storage, the pumped storage manager would be solely able to manage flows on that river reach, and presumably be responsible to keep them within limits, as above.[102]
- If the charge/discharge point is actually in the downstream storage shown in the figure, then there is no river reach, and hence no flow limit, to be concerned about immediately downstream from the pumped storage, and the downstream reservoir effectively is also the buffer storage. The downstream issue then becomes just one of managing the combined buffer storage/downstream storage to meet any flow limits further downstream.

Either way, the key points are that:

- The two managers would still have broadly similar incentives to maximise/minimise flows at the same time; and

---

[101] That also raises the prospect that, to be successful, a major pumped storage development might have to be complemented by significant modifications to downstream facilities. Such a re-development may well be economic, not just because it allows the pumped storage facility to be used more flexibly, but because the pumped storage facility can, and will, control flows in such a way as to make downstream generation capacity more profitable. But we understand that such possibilities lie outside the scope of present investigations.

[102] The situation beyond that is less clear. While the downstream host system reservoirs would be directly managing downstream flows, their ability to do so could be overwhelmed by unnatural flow fluctuations controlled by the pumped storage facility. So, the responsibility for keeping flows between limits may become an issue to be resolved between the pumped storage facility and the host system managers.

- Any intervening storage would assist in dealing with the coordination problems discussed in Section 4.4.

Notice that, even if the downstream storage might historically have been thought of as part of the host system, it is actually controlled equally by the two managers: One adding to, and the other subtracting from stock levels. And, thankfully, their incentives are complementary, in this case:

- Both will want to generate at peak times, so the pumped storage facility will be releasing water into the downstream reservoir at much the same time as the host system manager is releasing water from that reservoir. Thus, it may look like they have conflicting objectives: one trying to fill the reservoir, while the other tries to empty it.
- Actually, though, the pumped storage manager is not "trying" to do anything. The filling of the downstream storage is just an inadvertent side-effect of its generation, providing no reward to the pumped storage facility, but probably no penalty, either.
- The host system manager is not "trying to empty the reservoir" either, in fact they are trying to sustain maximum generation for as long as possible, over the peak, and every water unit arriving over that time helps towards that goal, by extending the period over which the storage empties.
- Thus, the impact of the pumped storage facility is to "flatten out" the storage trajectory, which is exactly what the arbitrage process discussed in Section 2.2 is trying to do. This also reduces the EMVSE differential between filling and emptying phases, and extends the unconstrained storage arcs over which those EMVSE levels can be held "as constant as possible", given changing expectations.
- If the maximum discharge rate from the pumped storage generator, plus the natural flow, is greater than the maximum utilisable flow rate in the downstream generator then storage will actually be rising during this period instead of falling.
- Similar effects will apply when electricity prices are low. The host system manager will want to minimise releases, but the pumped storage manager will not just minimise releases, but actually start pumping.
- If the pumped storage manager must maintain minimum flows on an intervening river reach, while the downstream manager can stop release completely,[103] then downstream storage will still build up at the minimum flow rate. If the pumped storage facility harvests water directly from the downstream storage, though, it is quite likely that downstream storage will actually fall over these low-priced periods, rather than rising as we would normally expect.
- But it will extend the period over which the downstream generator can hold release to a minimum, knowing that it does not need to have so much water stored in its own reservoir, because the water is instead being stored in the pumped storage reservoir, to be released at a time when both generators will be wanting to generate.
- Thus, the downstream storage trajectory may not only flatten out, but flip to follow a cycle that would be the exact opposite of that normally expected. If the pumped storage facility has high enough maximum charge/discharge flows, relative to the downstream facility's flow/storage capacity, the trajectory will be pushed against the upper/lower capacity limits of the downstream storage, and the situation may become much as discussed above, for the non-storage case, with downstream constraints limiting the national benefit from higher levels of pumped storage generation.

The balance of these physical effects will depend on the parameters defining each individual case, but the direction of the economic effect is clear:

- All of these effects, while unintended by the pumped storage manager, will increase the profitability of downstream generation, and that impact may be proportionately quite large, if the host system has little downstream storage; and
- The presence of the downstream reservoir extends the periods over which the pumped storage facility can operate in relatively unconstrained fashion, if it chooses to; but

[103] Because they discharge directly into the next storage.

- There is still likely to be a point beyond which further pumped storage release would force the downstream generator to spill water that it can not utilise.

So, the incentives are broadly aligned and, depending on its capacity, the downstream storage enables a greater or lesser degree of operational decoupling. But a cooperative agreement may still be worthwhile, if the real-world national benefit is to be brought up to the level that would be computed by an optimisation model.

Given the way an upstream pumped storage facility looks likely to boost the value of downstream generation assets, one would expect a rational host system manager to co-operate in negotiating an arrangement that could significantly boost the chances of such a development proceeding. But we do appreciate that there may also be national policy (e.g. competition) reasons for thinking that an agreement would be undesirable.

4.6. Accounting for upstream storage

4.6.1. Introduction

The situation with respect to downstream storage seems relatively benign, because the incentives of the two managers are broadly aligned with each other, and with the national interest. Thus, it should not be hard to agree on a value-enhancing coordination agreement, but it may not add enough value to be worthwhile, if it causes problems in other areas, such as by reducing competition. The situation with respect to upstream storage is more difficult, though, and more difficult to understand.

Basically, because pumping reverses the natural logic of catchment flows, it can also reverse traditional understandings of how a catchment should be operated, upstream from that point. And that also means that the strategies recommended by an optimisation model, and the marginal water values and valuations it will report, may be very different from strategies and valuations based on the assumed continuation of historical practices, and with intuitive understandings derived from historical experience.

Adding upstream storage clearly enhances the value of the host system. In principle, from a national benefit perspective, it must enhance the potential value of the pumped storage facility, too, because it introduces the ability to control upstream flows in such a way as to maximise the potential for both pumping and generation activity, at times when those are most valuable to the pumped storage facility, and the nation. Conversely, from a potential national benefit perspective, the presence of the pumped storage facility must increase, and probably greatly increase, the theoretical value of all upstream water, and control/storage capability, because that water can now be made much more productive, by pumping it up to a greater height, at low marginal cost.

Unfortunately, though, the examples in Appendix A show that upstream storage capacity also introduces the ability of the host system manager to control upstream flows in such a way as to maximise the value of host system production, by following a policy that, perhaps inadvertently but inevitably, has the side-effect of minimising the potential for both pumping and generation activity, at times when they would be most valuable to the pumped storage facility, and the nation.

Paradoxically, then, the presence of upstream storage could actually reduce the national benefit delivered by the pumped storage facility, unless the extra trouble and expense is taken to either establish a cooperative agreement to maximise the value extracted from the integrated catchment or build a buffer storage large enough to decouple the two systems. And, as we will see, a national benefit optimisation will only recommend building buffer storage as an adjunct to a cooperative agreement, not a substitute.

4.6.2. Conflicting objectives

The reason for potential conflict is easy to understand:

- The upstream host system manager and the pumped storage facility manager will, and should, both want to maximise releases when electricity prices are high, thus threatening congestion at downstream bottlenecks, and potentially forcing priority to be assigned to one or the other[104]; but
- The pumped storage facility manager will, and should, want to maximise pumping when electricity prices are low, thus requiring water to be available at times when the upstream host system manager will want to minimise releases.

If there is plenty of buffer storage, and/or flow headroom/freeboard, that potential conflict will be reduced, and may even disappear. The national benefit-maximising resolution of any remaining conflict is also very easy to understand, though:

- Any downstream congestion issues must be resolved in favour of prioritising generation from the upstream party with greater release productivity. And that will be the pumped storage facility, if its head is greater than that of the aggregate upstream host system.[105]
- Other things being equal, any upstream water availability issues must also be resolved in favour of retaining water in, or transferring water to, the reservoir with the greater net head, after accounting for efficiency losses. And the examples in Appendix A show that that will also be the pumped storage facility if its head exceeds that of the aggregate upstream host system by more than the efficiency losses.
- This is because, if the upstream reservoir release and pumping (and hence the losses) all occur at the same time, the return from generation just nets off the cost of pumping, and the net cost of transferring water by releasing it from the upstream host system reservoir, then pumping it up to the upper reservoir of the pumped storage facility can be very low, if the electricity price is low. The net value, once transferred, is proportional to the ratio of the pumped storage head to the upstream reservoir head, which would be a 10-fold value increase, in the examples studied.
- Other things will generally not be the same, but comparing the characteristics of possible pumped storage developments and host systems in New Zealand suggests that the effect of those other factors will most likely tip the balance even further in favour of prioritising pumped storage operations over traditional host system operations.
- In fact, Appendix A argues that, even if the pumped storage facility and upstream host system generator have the same head, but the upper pumped storage reservoir is larger, the national interest will be best served by sacrificing the potential returns from upstream generation in the relatively near future, in order to transfer the water to a location where it can be stored for use in (possibly much) later periods, when electricity prices are expected to be (possibly much) higher.
- And it is further argued that, even if the upper reservoir of the pumped storage facility and the upstream host system reservoir have the same head, the same generation capacity, and the same storage capacity, it will often be optimal to sacrifice the potential profitability of the upstream host system reservoir by releasing in low-priced periods, in order to transfer up to half the available flows into the upper reservoir of the pumped storage facility, which has no natural inflow of its own, so as to later be in a position to utilise both stations simultaneously in any future high-priced period.

Thus, in the absence of any buffer storage, the overall conclusion is that, in many realistic cases, national benefit optimisation will require that the role of the upstream catchment of the host system be fundamentally altered, to become primarily a system for collecting flows to be released at times of <u>low</u> electricity price, so that they can be pumped up to the upper reservoir of the pumped storage facility, thus maximising the value of the pumped storage facility, rather than returns from upstream host system generation itself.

---

[104] If both have the same release capacity, the effect could be to double downstream flows, so the change is not trivial, with the combined peak flow potentially also around twice the utilisable flow rate of downstream stations deigned to match release rates from the incumbent host system stations upstream.

[105] There could be a whole chain of upstream power stations involved, and their aggregate head becomes relevant, but the situation gets a bit more complicated if there is storage/buffer capacity between those upstream stations.

We will see below that buffer storage, including downstream host system storage, can improve the technical/economic situation, but it does not resolve the fundamental conflict, and could actually make it worse:

- In Figure 1, the charge/discharge point for the pumped storage facility is shown as being in a river reach, with a buffer storage, which the discussion so far has assumed to be minimal. We have already suggested that, if the downstream storage stretches further back up the river, the pumped storage facility will be able to charge/discharge directly from/to it, thus effectively becoming a buffer storage that can be used to decouple pumped storage and upstream operations, to some extent.
- But that could leave a river reach between the upstream and downstream host system reservoirs, forcing the upstream host system reservoir to release enough to maintain minimum flows, in that reach. And, with no minimum flows to maintain immediately downstream from its own charge/discharge point, those minimum flows will keep replenishing the downstream reservoir, from which pumping will occur.
- If the buffer/downstream storage extends back as far as the upstream dam, as is common on fully developed river reaches, that may increase the flexibility available to the pumped storage facility manager. But there is now no river reach on which flows must be maintained. So, the host system manager may be free to adopt a more extreme strategy of completely shutting down generation when prices are low, thus cutting off any replenishment of the downstream storage from which pumping occurs.

4.7. Adding buffer storage

It may seem odd to discuss buffer storage last, because it must be one of the first things that any pumped storage system designer will consider. Indeed, a minimal buffer storage will be required, just to allow technically feasible/efficient operation of any planned facility. But the issue here relates to building extra buffer storage, above that level, to enhance flexibility. And that raises the issue of exactly what "buffer storage" is supposed to achieve, given that it is really just another storage in a host system which already has "upstream" and "downstream" storage. Indeed the "buffer storage" may actually be an existing "downstream" storage in the host system.

In fact, buffer storage is arguably the last thing that should be considered, in a retrofit situation, because all of the host system storages are already in place. They may not be optimally configured for pumped storage purposes, but any national benefit optimisation will assume them to be available at no (avoidable) cost, and will optimise their use to maximise benefits from the new integrated system being planned.

Having done so, it may determine that extra buffer storage is justified, but it may not. It may be possible to utilise an existing storage as buffer storage for the pumped storage facility. Or the optimisation may be so successful at utilising other host system storages to marshal sufficient flows into the (minimal) buffer storage, at the right times, that no extra buffer storage appears to be required. If buffer storage is to be recommended, though, it seems important to understand exactly what that justification might be based on, and whether it would actually deliver the assumed benefits, in real life.

4.7.1. Basic operational limits

Because flows can be accumulated in the buffer storage, we need not be so concerned about instantaneous flows at any particular time, but must still account for the cumulative flow over any interval, as well as the state and capacity of the buffer storage. Specifically:

- Assuming the lower limit on buffer storage to be zero, the maximum amount that can be pumped, over an interval, *i*, of duration *Dur(i),* is limited by both *BSV(i),* the volume in the buffer storage at the start of *i*, and the flow situation in the host system:

$$BSV(i) + HSF(i)) - HSFmin \times Dur(i)$$

- The maximum amount that can be generated over interval *i* is limited by both the unused volume in the buffer storage, *BSVmax-BSV(i),* and the flow situation in the host system:
  BSVmax - BSV(i) + HSFmax x *Dur(i) - HSF(i)*

If there is enough flow headroom/freeboard these limits may not constrain pumped storage operations much at all, even without a buffer storage. But, while this kind of "replenishing" arrangement allows for more flexibility than one based solely on flows, it may still imply significant constraints on pumped storage operations. And that suggests that there may also be significant gains to be made by coordinating operations between the pumped storage facility and the host system, to maximise benefits to the nation, and to distribute some of those increased benefits in ways that improve returns to both parties.

That raises the question of how the buffer storage would actually operate because it is just as much part of the host system as of the pumped storage scheme, and its level will be determined by the cumulative actions of both managers. *HSF* in the limits above is no longer just the natural flow, but may be heavily modified by upstream release management. And the outflow will be determined by the upstream release, in combination with the pumping/generation level decided by the pumped storage facility manager.

It turns out that we can develop a narrative, and intuitive understanding, based on the assumption that it would be controlled by the host system manager. And we can develop an alternative narrative, and intuitive understanding, based on the assumption that it would be controlled by the pumped storage facility manager. But these two narratives turn out to be incompatible, and to have quite different implications for upstream/downstream and pumping/generation flows.

So, we must ask how these conflicting narratives can be reconciled: How would or should the buffer storage operate? and, for example, what would its marginal water value actually be? It turns out, once more, that a national benefit optimisation will determine a strategy, and implicitly determine corresponding marginal water values, that do not align with either perspective. And that observation has significant implications with respect to the way buffer storage development might be evaluated, and raises the spectre that its real-life operations might not actually deliver the calculated national benefit, without a coordination agreement.

First, though, we should re-visit the discussion of Section 3.3.2, which raised the prospect that a pumped storage facility in which the lower storage was significantly smaller than the upper storage could, and should, adopt what we might call a "mixed mode" cycling strategy, with the smaller lower storage inducing shorter term sub-cycles within the longer-term cycling strategy enabled by the larger upper storage.

4.7.2. Mixed mode cycling

Section 3.3.2 noted that a pumped storage facility with imbalanced storages only made sense if that system could be replenished from some external source, and discharge to some external sink. But, assuming a run-of-river host system so we can ignore any issues with respect to interactions with the host system manager, that is exactly the situation we have described here, with the buffer storage acting as the lower reservoir.[106]

In this context, the equations above define strict limits on what can be pumped, or generated, over any interval, and those must be adhered to, no matter how attractive pumping or generation may appear, when comparing electricity prices with the EMVSE in the upper storage. Those equations actually combine two kinds of limits:

- First, with no buffer storage, the average pumping rate could not exceed the freeboard of host system flows available, over and above minimum flow limits; and the average release rate could not exceed the headroom available, over and above the host system flows.

[106] In principle, all of these issues could be discussed in a more general context, which might not include any buffer storage, as such, but the extra complexity seems best avoided.

- Second, a buffer storage allows those instantaneous constraints to be replaced by looser limits on longer-term aggregates, but imposes its own limits on what can be pumped or generated over any time interval.[107]

Thus, if the buffer storage has only a daily capacity, the pumped storage facility can not physically pump or generate solidly for extended periods. As discussed in Section 3.3, the EMWV in a large upper storage will be (nearly) constant over any day, but the EMWV in the lower storage will cycle daily between higher and lower values, and so will the MVP/MCG for the facility as a whole.

The optimal EMWV pattern for the lower storage can be set relative to the upper storage EMVV, so as to achieve the maximum possible aggregate storage build-up/draw-down, or some more moderate amount, as required by the long-term strategy for the upper storage. But the pattern will also drive the facility to maximise generation, in those hours when electricity prices are greater than its MCG, and to maximise pumping in those hours when electricity prices are lower than its MCG.

Increasing the size of the buffer storage does not fundamentally change this situation, but does allow the pumped storage facility to pursue longer term goals more aggressively, over more sustained timeframes, at the possible cost of reduced contributions to shorter term cycling and/or stochastic responses. As discussed in Section 3.3, a pumped storage facility in which one or other storage faced a daily capacity limit might alternate between pumping and generating in response to price variations that exceeded its RTLF, in proportional terms, but were quite small in absolute terms, on a low-priced day. But a facility with two large storages might pump solidly all day, taking the opportunity to build up stocks for use in later days, when prices are higher.

In the limit, the "buffer storage" could be a large long-term reservoir in its own right, operating in a pattern that mirrored that of the upper storage, and hence in an opposite cycle to that of a conventional hydro reservoir, but exhibiting similar patterns of alternation between:

- Periods when storage is at a bound, with EMWV either steadily rising or falling, and incoming flows either passed through or pumped up; and
- Periods when storage is away from its bounds, with the facility probably alternating frequently between pumping and generation, to keep EMWV as constant as possible.

In other words, the pumped storage facility's operation would be primarily driven by longer-term priorities, while secondarily assisting with the management of shorter-term fluctuations in the supply/demand balance.

4.8. Managing host system interactions

We have discussed "host system interactions" in some detail, because it seems to us that managing these direct physical interactions may be the most difficult problem to resolve when it comes to organising and managing any "embedded" pumped storage hydro facility, and its interaction with the market, both short and long-term. Thus, in principle, if it is going to maximise national benefit:

- The volumes it can pump or generate are always strictly limited by the state of the shared buffer storage, and by upstream and downstream flow conditions in the host system; and
- The price at which it should be prepared to buy electricity for pumping, and the price at which it should be prepared to sell any electricity it generates, should not only depend on the MWV in its own upper storage reservoir, but on the difference between that MWV and that in the buffer storage, which will reflect marginal values throughout the host system; and
- Those host system MWVs should therefore also impact any hedge products it might offer, as discussed in the next chapter.

[107] Since the limits above apply over every interval, they basically replicate traditional limits, without introducing a specific storage variable, and storage balance equation.

But the obvious problem here is that storage/flows for hydro generation in most catchments are currently managed by independent power companies, so a major re-organisation of the sector might need to be negotiated in order to achieve that kind of “nationally optimal” outcome, while also preserving competitive balance. We can also imagine that an incumbent system manager might not welcome a development that threatened to take control of a catchment they have an established right to operate in.

If the host system is to be managed by an independent party, some of the information listed above would normally be regarded as private. Indeed, the MWV values referred to may not even be formally calculated. The same is true of other parties interacting via the market, under the status quo, and that does include some parties with upstream/downstream relationships. In those cases, though, the incentives of the various system managers involved do at least broadly align with the national interests, though, even in an imperfectly competitive market.

By way of contrast, it seems that that situation could change radically if a large pumped storage facility was added into any host system, and expected to operate independently of it. In fact, the direct interests of the host system manager, if measured solely in terms of profits from its own generation, seem likely to be diametrically opposed to the national interests, particularly upstream of the pumped storage facility. Conversely, even if the pumped storage facility manager set out to maximise national benefit, it could not so without having control over the host system, or agreement to secure a coordinated strategy, or mechanisms to align incentives.

None of that may matter, if there is actually little that can be done to influence flows, particularly in the upstream host system, either to further the host system manager’s interests or those of the nation. And the same might be true if downstream flow headroom/freeboard is always great enough to allow for largely decoupled operations. In some cases, it might also be cheap and easy to build a buffer storage large enough to allow largely decoupled operations that more-or-less maximised national benefit.

In other cases, though, we suggest that the viability of proposals to retrofit pumped storage into host systems already managed by incumbent generators could depend heavily on finding reasonably efficient ways to manage host system interactions.

Of course, the obvious resolution here is to integrate the management of the combined system under the control of a single organisation. That is what any national benefit optimisation will implicitly assume, as above. And it is also what we should expect to happen in any unconstrained market environment, given the potential upside of integrated operations to the incumbent host system manager, and the potential risks for any other party entering an environment controlled by the incumbent host system manager.

That “solution” also seems compatible with any of the organisational options discussed in Chapter 6. Of itself, though, it only increases concerns about market dominance. Those concerns might be partially addressed by restructuring, so that the entity controlling the integrated catchment system had no other generation assets. They might also be fully addressed by establishing integrated management, but then extending any of the regimes discussed in Chapter 6 to cover the whole catchment, and not just the pumped storage facility itself.

Arrangements that align incentives without requiring full integration should obviously be explored, though. Designing such arrangements lies outside our present scope, but Appendix B outlines several options for consideration, and briefly comments on their likely performance, and potential compatibility with the various organisational options discussed in Chapter 6.

Importantly, we expect that a credible study of the economics of any particular development proposal would have to employ an optimisation model to determine its optimal operating strategy, and economic contribution. But conventional optimisation models can only maximise a single objective, which in this case we expect to be national benefit.[108]

[108] Iterative modelling strategies could be developed to produce more realistic (sub-optimal) estimates, as in “gaming” models. But discussion of that kind of modelling development lies well beyond our scope, here.

The mathematics on which such models are based implements a rigorous economic logic that must, by design and mathematical necessity, implicitly ignore any distinction between assets controlled by one manager vs the other. So, it will come up with an operating strategy that maximises total national benefit from the system, assuming it to be an integrated whole, and value the development on that basis. In other words, the model must implicitly assume that, wherever ownership may lie, the parties involved will come up with a perfect coordination agreement.

So, in assessing the results of such studies, it should be recognised that the calculated benefits may not (fully) eventuate without a coordinating agreement. The extent of the discrepancy will vary from case to case, but seems likely to be greatest in cases involving significant upstream storage. Hence the discussion of possible arrangements to at least approximately align incentives with national benefit priorities in Appendix B, and on the implications for interpreting optimisation modelling results in Appendix D.

5. Hedge Market Interactions

5.1. Introduction

As will be seen, some of the organisational options discussed in Chapter 6 rely heavily on contracts that can be interpreted as "hedges", albeit perhaps of unconventional form. It should be stressed that, while these contracts may be described as "financial instruments", they are not necessarily designed to be traded at all frequently, or easily. The more complex "natural" forms discussed below are much more in the mould of a traditional lease: Simpler and more flexible than outright ownership, but giving the long-term holder benefits similar to outright ownership, and hence exposed to some of the complexities normally associated with ownership.

On the other hand, the art of mainstream hedge market design rests heavily on finding a hedge product form that a wide range of parties feel they can value, and hence are prepared to trade. And there are already several broad forms of hedge available in the New Zealand market:

- Strips of energy "futures", representing energy to be delivered to a specific location, at specified times of day (in broad base/peak blocks) in specified months, perhaps months or years ahead;
- Options on those energy hedges, and "swaptions", under which the holder may call on specified capacity, at any time over the contract period; and
- Locational swaps, known as "Financial Transmission rights" (FTRs) allowing the hedging of inter-nodal price differences, potentially also at specified times of day, in specified months, perhaps also months or years ahead.

So, the issue arises as to whether the manager of a pumped storage hydro facility should merely provide "natural form" contracts, or (also) interact with those established markets, and/or perhaps with some market trading a new class of product. Accordingly, we will start by discussing the "natural form" of contracts that could be supported by storage facilities, including batteries and pumped storage, then move on to discuss prospects for providing hedges that traders, and some market participants, might find more attractive, for various reasons.

5.2. Natural hedging for generators, loads, and traders

Our own thinking, over the years, has tended to focus on the fact that all hedges are still forms of contract. We think it important not to entirely lose sight of the fact that the original purpose of contracts was to induce parties to fulfil promises to deliver goods or services of value, on the one hand, and to pay for them on the other. Thus, whatever deals may be done by third parties in trading rooms, perhaps unconscious of the implicit physical realities, there must ultimately be parties able to (at least approximately) back these contracts with physical production, and parties ready to pay for physical delivery.

As a result, we have tended to favour "hedge products" that reflect, as closely as possible, the properties of the physical capacity ultimately underpinning the value of those products:

- For geothermal plant, that would be a base-load "strip" of energy futures, as might be traded in the current hedge market, and the same would be true of the "minimum running" component of hydro generation.
- Less predictable run-of-river hydro and wind generators might also back similar base-load products, if there is not much pattern to their expected output capabilities, albeit at some risk to themselves if unable to match the volumes sold, when conditions go against them on the hour.
- Solar plant seems unlikely to back that kind of product, though, because it knows it simply can not generate over half the 24-hour cycle (on average). Thus, it would presumably prefer to offer day-time hedges. It would also have to sculpt its monthly offerings over the year, offering more in summer than in winter.

None of those capacity types would seem like natural backers of call options, though, because their output is basically independent of market price, and not controllable to match option calls.[109] Thermal plant presents a range of more interesting cases. In theory, the natural form of contracting would reflect the cost and capability structure of thermal plant, which can be characterised as a fixed annual sum being paid to provide the physical option of running the plant to generate whenever electricity prices exceed its "Short Run Marginal Cost" (SRMC), including fuel and variable O&M.[110] In other words, it would ideally want to (only) offer call options to the market. In reality, though:

- A base-load thermal generator may be quite confident in offering base-load hedges, even years ahead, because it believes that it will (nearly) always be operating in that mode, since prices will seldom fall below its SRMC.
- A "shoulder" generator may think similarly, because it also considers its future generation schedule to be fairly certain, but restrict itself to only offering, say, day-time hedges in winter months.
- A "peak support" generator may be much less certain, particularly in the inherently volatile hydro-dominated NZEM, and would more naturally offer call options, or some combination of call options with sculpted energy hedges covering the generation levels it is fairly certain it will be called for.
- But true "extreme peaking" capacity, which may only be called on for a few hours in a typical year, and perhaps for a more extended period every decade or so, really can not offer anything but call options, because it has little idea as to when it will be generating.

The broad picture, then, is that thermal generators might effectively be selling across some combination of three different modes:

- Strips of energy futures matching generation levels they are fairly certain of achieving; plus
- "Natural Form" call options covering (ideally) their remaining available capacity; plus (or potentially minus)
- Supplementary spot trading, when high/low prices make it more economic to generate above/below contracted levels.

The above discussion omits an important factor, though, which is that the "natural form" of the contracts consumers ideally wish to buy is typically quite different from the natural form of contracts generators ideally wish to sell:

- Most small consumers will simply want to buy from a single retailer, on a fixed price/variable volume basis;
- Larger consumers may be willing and able to buy more on a spot basis but, if they think they know exactly what their load profile will be, they will also want to buy conventional hedge strips to approximately match that profile; while
- Traditionally, some very large consumers might consider investing in their own power stations, to provide a "physical hedge option".[111]

Such physical investments would, by definition, provide hedging in the "natural form" applicable to the technology employed. So, large consumers contemplating such investments should be trading that option off against the option of sticking with their core business, while buying hedges off the market. If hedge contracts were available in the natural form corresponding to that technology, that would effectively give them access to a "slice" of that type of capacity developed by some other party in the market, but probably at a lower cost, due to economies of scale and scope.

That mode of thinking might be expected to spread, in the scenarios discussed by MDAG. Businesses and households contemplating rooftop solar investments, and trying to match their load patterns to

---

[109] Ignoring the possibility that market prices may fall below production-dependent O&M costs.

[110] Although, as discussed in Section 2.3, much of the fuel cost may also be locked in, thus lowering the true SRMC to a level at which the thermal generator is effectively committed, in advance, to generate according to some profile, almost irrespective of market price, and may hence issue energy hedges accordingly.

[111] Many smaller entities have also made investments in backup capacity in the form of Diesel generators, although the motivations for that investment often relate more to dealing with localised transmission/distribution outages than with energy market pricing peaks.

output and/or buy batteries to facilitate that match, should logically consider buying virtual battery capacity instead.[112] And, if DSM is going to play as large a role as MDAG considers it may need to, consumers should not actually "know" their consumption levels far in advance, because they should be planning to respond flexibly to future market signals. So, they should be looking for flexible electricity supply arrangements, rather than sculpted strips of futures. They might achieve that by buying capped hedges, or by selling their own call options, reflecting their willingness to reduce demand, at prices above specified levels.

Still, the reality is that most consumers will not understand, or want to purchase, a whole portfolio of hedges, each in the natural form representing the technological capabilities of the ultimate underlying technology. So, the role of the retailer is to re-package what the generation sector can naturally offer into forms acceptable to the consumer.

The hedge market plays an important role in that re-packaging by allowing generators who wish to act as retailers to purchase hedge products to complement their own capacity, and stand-alone retailers to purchase a virtual portfolio of capacity to back their retail commitments. In the process, though, a third factor emerges to determine the shape of available hedging products: Namely the hedge form considered most natural by traders, who often come to the sector with experience of other markets, in which the underling natural forms of contracting may be quite different from those natural to the electricity sector.

Traders prefer, and indeed require, simplified products that may not match the technical characteristics of the underlying technology terribly well at all. And that tends to force market participants into dynamically adjusting market positions to reflect their ever-changing physical situation, rather than relying on the automatic adjustment implicit in more "natural" contract forms, such as call options.

5.3. Natural hedging for hydro systems

The above discussion directly applies to hydro plant that has no associated storage, and hence can only adopt a run-of-river operating mode. Such plant could naturally support base-load hedges, up to a generation level corresponding to its minimum flow level. Beyond that, it would face increasing risk if it were to offer energy hedges, or if it were to enter the retail market, because it would be facing the risk of having to buy power in from the spot market in order to meet its retail commitments. So, if it wanted to make commitments beyond that point, it would be looking to purchase hedges itself, to cover its own risk exposure.

A hydro generator with physical storage options available can provide much of its own hedging, by storing water arriving in periods when its commitments (and/or electricity prices) are low, so as to be able to generate in periods in which its commitments (and/or electricity prices) are higher. But its physical self-hedging capability will be limited, and it may well still seek supplementary hedging from the market. And/or it could use its physical hedging capacity to back hedges that would enable other market participants to cover some of their own risks. So, we turn to discuss the natural form of hedges backed by hydro generators with storage.

[112] Assuming the metering and network cost recovery arrangements could be worked out to create a level playing field, and that such hedge products were available, and comprehensible.

5.3.1. Virtual hydro system slices

Barroso et. al. discuss international examples involving "virtual model" representations of physical systems used as a basis for contracting in the hydro sector.[113] In particular, Hunt and Read proposed a "virtual reservoir" to allow competing parties to drive operational strategy for the Tasmanian hydro system by exercising "swing options", while retaining integrated management of the physical system.[114] Chapter 6 discusses the possibility of adopting that paradigm to allow "diversified" management of a pumped storage hydro facility in New Zealand. But the point, here, is that Hunt and Read proposed to express those "virtual reservoir" contracts in the form of "financial options" that would approximately reproduce the "physical hedging" made possible by the storage, in its "natural form".

The natural form of physical hedging provided by pumped storage hydro will be discussed further below, but the natural position of a traditional hydro plant that has associated storage is that:

- Storage levels are built up by uncertain inflows received into the reservoir;
- The reservoir manager assesses an EMVSE for stored water, using some form of opportunity costing methodology, as discussed in Chapter 3, and thus determines how much of that water should (in that manager's opinion) be optimally released for immediate use, and how much stored for future use;
- The manager makes market offers intended to produce something close to that "called for" release;
- The manager then implements the release determined by the market;
- Storage levels adjust to reflect the net impact of inflow minus release; and
- The process continues, with the manager being responsible to ensure that physical efficiency is maximised, and physical limits respected, while managing a continuously changing stream of inflows, and market conditions.

Accordingly, we suggested, in the Tasmanian proposal referred to above, that the "natural form" of a financial contract for hydro reservoir storage would be to:

- Form a mathematical representation of the system, with all of its various reservoir/tributary inflow streams, storage capacities, release/generation capacities, flow limits, conversion efficiencies, etc, such as would be employed by a (probably simplified) formal mathematical optimisation model of the system; then
- Sell/lease "slices" of that "virtual reservoir/system model" as hedges to be "managed" by the lessees, as they see fit; where
- Lessees "manage" their slices, not by exercising "dispatch rights", but by calling the corresponding MW under the virtual reservoir hedge contract, on the understanding that their share of the virtual storage capacity will:
  - Fall to reflect the volume released to support the dispatched generation, but also
  - Rise to reflect their share of incoming inflows.

Such "options" may seem complex, to those familiar with trading standardised options on generic trading platforms, but the intent is not to create a tradable product. It puts each option owner in exactly the position they would be in, if they owned a scaled down version of the system themselves, and used established techniques to manage it, but without any responsibilities for physical maintenance, dispatch, environmental compliance etc. Conversely, the physical system manager would be responsible for all those things. They would also interact with the spot market to determine physical generation schedules,

---

[113] L.A. Barroso, S. Granville, P.R. Jackson, M.V. Pereira & E.G. Read: *Overview of Virtual Models for Reservoir Management in Competitive Markets* Proceedings 4th IEEE/Cigré International Workshop on Hydro Scheduling in Competitive Markets. Bergen, Norway, 2012

[114] D. Hunt, E.G. Read, P.R. Jackson, L. Barroso and S. Granville: *Tasmanian Market Reform: Commentary on Panel's Preferred Options* Concept Consulting Report to Aurora Energy, Tasmania, Australia, 2012.

Summarised in Section III B of Barrosso et al, and also:
P.R. Jackson and E.G. Read: Financial Reservoir Models Supporting Competition in Integrated Hydro Systems *ORSNZ 2014.*

but they would do so with strong incentives to approximately match the aggregate "calls" from the various lessees which they must ultimately honour, in a financial sense.

Returning to the original discussion, then:

- A hydro generator could issue "slice" contracts in a form that directly provided the hedging capability of the underlying hydro system in its natural form to slice holders;
- Those slice holders could be end consumers who would call on their slices as required; or
- They could be other generators who would use the slices to complement their own physical capacity, in order to support their own retail/wholesale market offerings; and/or
- They could be intermediaries offering more conventional hedges to current markets.

But the key point is that it would be the slice holders, not the facility owner, who would be in a position to offer conventional hedges to the market, just like any market participant owning similar physical capacity under the status quo.

5.4. Hedging for batteries and storage facilities

5.4.1. Hedging in the new environment

In a traditional market environment, the physical storage capacity implicit in the thermal system has been implicit in the hedges backed by thermal plant, and the natural form of those hedges would really be as call options of some kind, including the swaption type arrangements currently in place. Since that type of option will not be available from thermal generators in the new market environment, we need to ask what will take its place. The answer must obviously relate to the types of physical capacity envisaged as growing in the new environment.

Based on the MDAG report, we might expect to see:

- Increasing volumes of base-load futures strips supported by geothermal, and perhaps also by diversified wind portfolios;
- Increasing volumes of sculpted futures strips, at least for summer day-time periods, supported by large scale solar developments;
- Possibly some call options with high strike prices supported by green thermal capacity;
- Possibly some call options with a range of strike prices, supported by demand reduction capacity; and
- Possibly some new kind of options supported by battery and/or other storage capacity, potentially including pumped storage hydro.

But what kind of hedging can "storage" actually support, and what is its "natural form"?

5.4.2. Active trading in conventional hedge markets

Unlike the hydro generators discussed above, storage facilities, on their own, are actually net loads, not net generators. So, we should not expect them to simply issue and support conventional hedges using strips of energy futures.

In principle, a storage facility could be said to only have the ability to sell hedges against price variation/volatility, as discussed in the sections below. In practice, we believe the situation might be seen differently, though, because the storage cycle for a large-scale storage facility could be even longer than the market trading cycle. Participants will be very much aware of the quantum of energy in stock, and may expect to see the manager of that stock active in the market:

- To buy in enough spot power to build stocks up, if stocks and prices are both low enough, while perhaps also covering its buying risk in the hedge market;

- To sell spot power, at the risk of running stocks down, if stocks and prices are both high enough, while perhaps also covering its selling risk and/or providing other market participants with the means to cover their buying risk, in the hedge market.

So, when various parties talk about the desirability of a pumped storage hydro facility supporting some kind of hedging, we suspect they may not be thinking in terms of such a facility issuing hedge instruments corresponding to its storage/generation/pumping capacity, but of that facility effectively selling rights to the actual stock held in that storage capacity. We believe the natural form of those rights would be the "tank options" discussed in Section 5.4.6 below. But perhaps market participants are expecting to see the manager of such a stock:

- Selling strips of cap options and/or winter/peak energy options;
- Buying strips of put options and/or summer/off-peak energy options; and possibly even
- Buying or selling energy strips, as for any other storage-based generator.

Such a strategy would seem most compatible with the kind of "unified" management structure discussed in Section 6.3. but there would surely be concern that a large facility dynamically exercising its own judgement and discretion with respect to the volume and pricing of such hedges might significantly impact that market, even if it does not set out to do so. So, we expect that rules would be required to govern any such operations in hedge markets, just as for the spot market. We turn to consider other hedge concepts that could be applied, though, including the hybrid arrangement discussed in Section 6.5, combining conventional market operations to buy pumping power, with the "tank option" concept of Section 5.4.6 below.

5.4.3. Collars

The key role of storage is to allow arbitrage between periods where price differences exceed a quite specific threshold, determined by its RTLF. So, in principle, a storage facility should be able to support some kind of "collar" option, protecting the holder against price fluctuations outside a specified band. The instrument that can be naturally supported is not a conventional "collar" though:

- A "collar" is normally constructed by selling a "call option" while buying a "put option" at a lower strike price, thus leaving the holder of the underlying commodity/share/instrument only exposed to price fluctuations between the put and call strike prices.
- In this case, it is the consumer who wants to be protected against volatility in the price of a commodity they do not hold, but believe they will need to buy, so they would be buying the collar: That is buying a call and selling a put.
- And we might think of a battery, say, as providing the physical support for a strip of such collars, with strike price buy/sell spreads determined by its RTLF.
- But even a battery of infinite capacity can only provide protection against proportional price volatility, and would not support an absolute spread between strike prices in some future period for which price levels are unknown.

In reality, a storage facility's ability to limit exposure to price volatility is not just limited by its MW charge/discharge capacity, but by its storage capacity. As discussed in Sections 3.2, 3.3.2 and 4.7.2 above, a battery, or a pumped storage hydro with either upper or lower storage small enough to reach its capacity bounds over a daily cycle, would find itself optimally performing two distinct types of arbitrage, and potentially providing two distinct types of hedging capacity over that cycle:

- Day to night hedging, limited by its storage capacity; and
- Intra-day and intra-night hedging, limited by its pumping/generation capacity.

In this intra-day setting, there could also be quite long periods between the "day" and "night" zones within which no hedging can be provided. Even if storage capacities are large enough not to constrain intra-day cycling, similar effects can be expected over longer planning horizons. And the real pumped storage situation will typically be much more complex, because the buffer storage will probably be embedded in a host system with its own generation capacity, implying a more complex set of constraints on its operation, and also a more complex set of marginal water values reflecting the value of water in host system operations.

The constraints involved will obviously vary from system to system, and it would be difficult to establish their pattern, or assess their impact, without applying a formal optimisation model. In general, it seems clear that the natural form of hedging available from a generic pumped storage hydro system would be more than just a simple "collar" derived from its RTLF. But some pumped storage facilities could have enough upper/lower storage capacity that such complexities can largely be ignored because they are only experienced in exceptional extreme circumstances.

But there seems to be a more fundamental issue here. As we understand it, a collar designed to hedge against price volatility for a single stock/commodity normally involves a put/call combination applying simultaneously, with either one or the other possibly called in a specified time period. That might seem applicable to a battery, selling future rights to its storage and charge/discharge capacity, as a means of hedging against daily price cycling and short-term volatility. But, when we look closely, the hedging provided by a storage facility is always between a commodity available at one price in one hour, and at another price in another hour, possibly on the same day, or possibly years ahead.

If we see the put side of a collar as representing pumping, there may be a long delay between the time at which the put holder might exercise their right to (implicitly) sell power for pumping to add water to the upper storage, and the (possibly distant) future price at which the facility might expect to be called upon to generate to cover the call side of the collar. And it is not clear how a collar can be designed to hedge against, say, inter-annual volatility alone, or intra-day volatility alone, or both together, without confusing the two, in some way.[115]

The RTLF does play an important role in that it determines a pricing dead-band within which the facility will not attempt to equalise prices between periods. But it should be recognised that efficiency losses are not the essential contribution made by any storage facility, any more than they are in the transmission system. The ideal storage would actually have no efficiency losses, and the collar discussed above would then collapse to a single price level. But pumped storage has no ability to guarantee any particular price level in future periods. Thus, we suggest that losses should be seen as they are in the transmission sector; that is, as a complication and limitation on the hedging that can be provided.

Rather than creating collars, designed to implicitly guarantee a fixed future price differential (+/- losses), we might think of the "collar" paradigm as providing a loose description of a trading strategy involving the purchase of storable energy, possibly via put options,[116] in some periods, in order to support the selling of call options, perhaps a few years ahead. A variation on that strategy might be adopted by the facility manager under some of the unified organisational regimes discussed later, in Section 6.3. First, though, we need to discuss financial instruments that do represent the essential function of storage, which is to allow energy to be transferred between periods.

5.4.4. Financial storage rights

We note that storage capacity is very much like transmission capacity, but operating over time, rather than space. Thus, a "Financial Storage Right" (FSR), can be defined in a very similar way to a "Financial Transmission Right" (FTR), and ownership of (physical or virtual) storage capacity would provide a basis for issuing such rights. We have studied the possible creation and trading of such rights in the arguably more complex context of a network-based water market, where water may be simultaneously

[115] We can imagine a strip of collars, but that would imply puts and calls being exercised in various periods, with no mechanism linking them, or even balancing the volumes exercised. If prices go high, then a great many calls may be exercised over an extended period. Or a great many puts may be exercised, if prices go low. There is no guarantee that the overall pattern would average out to be something that the storage facility could actually support with a realistically balanced pumping/storage/generation schedule.

[116] Although it does not need, and probably should not be taking on, an obligation to buy power in any particular period, if it wants to be free to take advantage of low power prices whenever they arise.

traded over both space and time.[117] Specifically, we concluded that the network/storage owner should be able to support a wide range of combined FDR/FSR rights to hedge the risks involved in such trading, subject to a "revenue adequacy" test, similar to that applied in clearing FTR markets.[118]

That paradigm implicitly assumes that there is enough water being traded in each specified location (the upper, and potentially lower, reservoir(s) of the pumped storage facility, in this case), with enough competition to set credible spot prices in both the initial and terminal FSR periods. It should also be recognised that while FDR/FTRs are described as "rights", the standard FTR form is arguably just as much an "obligation". While no party is physically forced to transmit power corresponding to the FTR, the FTR holder has effectively swapped a position at one node for a position at the other. And that swap only deals with their inter-locational risk exposure precisely if they actually do "send" that amount of power from the sending node to the receiving node, at that time.[119]

Similarly, an FSR only deals with inter-temporal risk precisely if the holder actually does store water from the sending period to the receiving period. That may suit a participant who knows they want to store water over the summer for use in the next winter, but what about the participant who wants to store water for precautionary use in some unknown future period, maybe days ahead, or maybe years ahead?

FTR "calls" can be defined to remove the obligation implicit in the standard form, and FSR calls can be defined similarly. But does this participant really want to guard against the risk of high marginal water values in the upper reservoir of the pumped storage facility? That may be helpful if the participant is focussed on managing seasonal/dry year risk, but EMVSE and EMWV will hardly shift at all over the planning horizon of concern to a participant wanting to manage the price/volume risk faced by a solar/wind portfolio. It is really the EMWV differential between the upper storage and the buffer storage that aligns with electricity prices, and we have seen that that differential can vary across each day, at least if one or other storage is small.

So, a set of hedges precisely matching the entire transaction could involve an FDR call representing transfer from the buffer storage to the upper storage, with a loss adjustment and a price term reflecting the cost of electricity purchased for pumping, plus an FSR representing storage in the upper storage for a specific period of time, plus another FDR call, representing transfer between the upper storage and the buffer storage, with a price/value term reflecting the electricity generated by that transfer.

Of course, that could all be bundled, but there is another problem here, in that the hypothetical owner can only "use" this FSR/FDR bundle if they actually have rights to water in the buffer storage in the first place. They will presumably be generating excess energy in some future period, and the transmission system will generally be able to transfer that energy to the pumped storage location to be used for pumping. But, in an embedded system, will there be water to pump? To complete this transaction, they also need to be able to buy water in the buffer storage, or hold some kind of right giving them access to a share of the water available in that storage at the particular time they want to pump it, as a way of storing energy. Once we include that factor, though, the required financial instrument starts to look very much like a "virtual system slice", as discussed in the next section.

We suggest that these FSR/FDR concepts should be investigated further, but we are not aware of any jurisdiction using rights of the forms we proposed.[120] In their simplest form, they could prove useful if participants can buy and sell energy stored in the facility, and then want to buy rights corresponding to storing that energy over an extended period, perhaps in the explicit form of water held in the upper reservoir of a pumped storage facility. It remains to be seen, though whether participants actually want to buy and sell rights to move water, or energy, from one specific period to another specific period.

[117] Mahakalanda, Read, Starkey, and Dye: Financial Hedging Instruments for Water Markets: *ORSNZ 2014*
Mahakalanda, Read, and Dye: Financial Rights and Obligations for Water Delivery in Mixed Use Catchments, *ORSNZ 2015*

[118] The distinction between FDRs (Financial Distribution Rights) and FTRs being essentially just that FDRs apply to "transmission" of water, not electricity.

[119] By selling it at the sending node, and buying it back at the receiving node, ignoring losses.

[120] Although we have not recently searched the literature to judge that question, either way.

Alternatively, they might want to hold rights representing storage capacity that can be managed to transfer stock from any period to any other period (within limits), as discussed in the next section.

5.4.5. Virtual storage system slices

Although the FSRs described in the previous section may sound very much like the virtual hydro system slice contracts discussed in Section 5.3.1, there are two fundamental differences:

- First, those hydro system slice contracts included proportional slices of hydro system inflows. Thus, they were not merely contracts for "capacity", but for energy as well.
- Second, those slice contracts allowed the holders to exercise their rights in a way corresponding to utilising the natural capacity of the reservoirs to hold water for as long or short a time as they ultimately decided they wanted to, rather than specifying the exact length of time when the storage capacity instrument is bought, as required by the FSR concept:
- So, if we wanted to apply that concept to the whole pumped storage facility, each slice could include a slice of the pumping/generation and upper reservoir storage capacity, along with:
- A slice of the buffer storage capacity, the inflows to that buffer storage capacity, and the headroom available for release in each period; or
- A slice of the entire host/pumped storage system, including all inflows, storage, generation and flow capacities.

With the second form of contract, the holder would exercise financial rights corresponding to determining a release/generation/pumping schedule for the entire host system, along with its embedded pumped storage facility. But, of course, that would only make sense if management of the two systems were integrated, or perhaps just aligned via one of the contractual mechanisms discussed in Section 4.6 above. [121]

Thus, the first form of contract would be more appropriate if there were no arrangements in place to secure flows to the buffer storage at times appropriate for pumping. That would allow right holders to exercise options equivalent to managing flows arriving at the buffer storage, through the pumped storage facility, and ultimately back to the buffer storage. But it would mean that all participants were reduced to exercising rights corresponding to harvesting flows arriving at (possibly inappropriate) times that they might predict, but could not control. Likewise for generation.

That might be quite acceptable, if the buffer storage is large enough to decouple the desired pumping/generation pattern from the host system manager's preferred inflow/outflow pattern. Or, it could be acceptable if the host system manager actually has so little control over the inflow/outflow pattern that the exercise of options by parties holding "slice rights" for various host system capacity elements can actually have no practical impact on operations. Even so, we can imagine that some participants might be daunted by the challenge of managing this degree of complexity, and suggest a simplified alternative in the next section.

---

[121] Note that unlike a natural hydro system, where inflows that arrive will naturally be stored for later use, irrespective of any rights being exercised, it might be thought that valuable water would be allowed to flow on past the buffer storage, if some participants choose not to exercise rights corresponding to pumping those flows, on the day. The same may be said with respect to the rights to use pumping capacity, or generation capacity, or host system flow capacity. That need not happen, though, because these are not physical dispatch rights:

- The pumped storage facility manager could be allowed to pump/generate using any excess flow capacity available, on the day, and keep the corresponding energy stored in the upper reservoir, or sold to the market, on their own behalf, contribute it to a collective pool, or make it available for purchase by other participants.
- Or, a market could be established to trade excess rights, of various kinds, on the day.

5.4.6. Tank options

Finally, we suggest that some participants might be more comfortable with a simplified form of virtual reservoir contract that we will call a "tank option". The simplification here is that we do away with the necessity to manage pumping, or any associated host system interactions, and assume that energy stored in the upper reservoir of the pumped storage facility is available for purchase.

We have previously described similar options as "a kind of swing option". According to Chen [2021][122], swing options are…*a type of contract used by investors in energy markets that lets the option holder buy a predetermined quantity of energy at a predetermined price while retaining a certain degree of flexibility in the amount purchased and the price paid.*

Noting that, while these contracts can contain a wide variety of conditions, Chen suggests that a typical swing option contract *delineates the least and most energy an option holder can buy (or "take") per day and per month…, how much that energy will cost (known as its strike price), and how many times during the month the option holder can change or "swing" the daily quantity of energy purchased.*

What we have in mind could be described as an extreme example of this option type, in which the take <u>rate</u> may actually be varied in each market trading interval, but the total take over the duration is fixed, in "take-or-pay" fashion. Alternatively, we could describe our "tank option" as a bundle of "American" style call options,[123] each for the specified MW capacity, and callable at any time up to the contract duration, but with the additional restriction that the maximum call across all options in the bundle can not exceed the MW capacity limit in any market trading interval.

Since neither description seems quite perfect, it seems best to provide our own definition of what we will call a "tank option", which works as follows:

- The duration of the contract may be anywhere from one day to several years, probably with rollover provisions.
- The maximum take <u>rate</u> corresponds to some proportion of the available generation capacity (preferably fixed MW, hopefully subject to only occasional scaling). [124]
- The minimum take <u>rate</u> is zero (corresponding to the facility being in an off, or pumping, state)
- The frequency with which the take rate can be varied is (probably) half hourly, and (preferably) subject to notification far enough ahead to allow for the reconciliation/offer process described in Section 6.4.2.
- The <u>total</u> energy available to be called over the contract duration is fixed, and paid for on a take-or-pay basis, but:[125]
    - o It can be topped up at prices that are either announced by the pumped storage facility manager or (preferably) determined by a regular auction process; and
    - o It can probably also be sold down, similarly; and
    - o It can either be sold, or rolled over, when the capacity contract is rolled over.
- If the "strike price" is thought of as the price at which energy is bought, as in Chen's description, then it can also be interpreted as the price paid <u>for</u> the contract quantity, because the total quantity bought at that price is fixed.
- But that means that the "strike price" is not actually relevant to the operation of the contract.

---

[122] See: https://www.investopedia.com/terms/s/swing-option.asp#. Chen suggests *"take-and-pay options", "variable base-load factor contracts,"* or *"swing contracts"* as alternative names. But note that this is not the same as a "swing option trading strategy", which we understand to be a trading strategy employed by option traders, not a type of option.

[123] Unlike the more familiar "European Calls", which are often traded in strips because each can only be called at a specified time, "American Calls" can only be called once, but at any time up to the expiration date.

[124] As discussed below, this might be sculpted on a daily/weekly/seasonal basis.

[125] In this respect, our tank option differs from Chen's description of a swing option, in which the total "take" over the contract duration would be variable, within specified limits., and paid for at an agreed strike price.

- Instead, the “optionality” provided by the contract is the freedom the holder has to choose when, and to what extent, they exercise their take rights.[126]
- Thus, the price paid under the contract would have two distinct components:
  - A payment for the “reserved” generation capacity, over the contract duration; plus
  - Payments made for stored energy, whenever that is topped up.

We note that, if we allow both energy and capacity components to be traded separately, this might no longer be described as creating/trading a single option of any specific recognised form. That does not invalidate its potential usefulness in this context, though. In fact, we note that buying a device of specified capacity to perform some tasks, and then “topping up the tank” as required is, if anything, more normal than buying a pre-filled single use device. Hence our reference to a “tank option”, which we see as analogous to buying a rechargeable battery, or refillable gas cylinder, or car, rather than a pre-charged throwaway, without having to deal with the technology by which the charge/gas/fuel is delivered to the point where it is available to re-charge the battery/cylinder/tank.

We also note that the concept of a fixed quantum of energy being available to any particular option holder over a contract duration effectively breaks down, if stored energy can be freely traded between capacity option holders. Rather than all tank option holders independently determining their own (possibly quite different) EMVSE, the prices paid in that stored energy market would effectively determine an equilibrium consensus EMVSE.

We do not see that as being at all undesirable, but note that a similar outcome might be achieved if similarly active trading were to be facilitated under the sliced virtual reservoir paradigm discussed earlier. What would still distinguish this “tank” option, though, is that tank option holders would not be responsible for bringing new energy to the stored energy market. Instead that energy would be pumped up by the facility manager, or perhaps by slice owners pumping extra energy, over and above their own requirements, to sell to the tank option holder market, depending on which of the organisational options discussed in Chapter 6 is preferred.

Finally, if we were to consider development of a new market actually trading some form of “storage option”, it would be highly desirable for the form of those options to be such as could be backed by all storage technologies, rather than being specially tailored to match some particular technology. So, one major advantage of these tank option arrangements is that they are less specific to the pumped storage hydro technology.

In order to fulfil that requirement, though, it might be necessary to forego the flexibility provided by allowing participants to specify offer curves, and revert to the simpler regime under which they merely specify the quantity called, and leave the facility manager to match the aggregate quantity called, as best they can, in the spot market. We would be cautious, though, because we see potential for simple quantity calls to create instability, and commercially disappointing outcomes in the spot market, with possibly unstable alternation between pumping and generation states, as discussed in Section 6.2.4.

Another possibility would be to recognise all rights of this form as giving their holders the right to make offers in the spot market, with the pumped storage facility manager merely acting as an agent, in this instance. But we suspect that the form of bid involved, the frequency with which bids can be changed, and gate closure requirements are all likely to be technology-specific, while the location will be specific to a particular facility. Thus, a more realistic goal might be to recognise that participants holding storage rights corresponding to particular facilities might then be in a position to offer rights of a similar, but more generic, form into a broader market.

[126] If a participant is managing this contract as a form of virtual storage, they may well want to specify their own “strike price” at which the contract is automatically called, or “exercised”, rather than making a specific call with respect to each trading interval. But that exercise price could change over the duration of the contract, depending on their assessment of the situation and remaining stock, and is not directly related to the take-or-pay price paid for the contracted energy quantity.

6. Organisational Arrangements

6.1. Introduction

The primary issues determining whether any particular storage development should proceed include its technical feasibility, the national and/or commercial costs and benefits it might deliver, its environmental impact, its ownership structure and the risks faced by its owners. All of those issues would need to be carefully assessed, both qualitatively and quantitatively, but that kind of assessment lies outside our present scope. Our focus is on how decisions might be made with respect to operational strategy, once a facility was built; how it might interact with markets; and what the combined effect might be on market outcomes, both short and long-term.

While much of the discussion in this chapter may apply to other storage options, it is also true that operational and organisational issues may differ widely. Some options may be large-scale developments, while others could be distributed throughout the network, some embedded and some stand-alone, etc. So, in the interests of clarity and concreteness, this discussion focusses on large-scale pumped storage hydro, and does not consider possible generalisations to cover other technologies.

Accordingly, this chapter provides a preliminary high-level overview of several organisational structures and/or arrangements that could be made to allow a large-scale pumped storage hydro development to achieve something like its full potential, in national cost-benefit terms. None of these proposals is perfect, so we list their pros and cons, and make some tentative suggestions about ways in which the negative factors might be mitigated. We do not attempt to discuss any details of the kind of institutional arrangements that might be required to make any of these proposals work, or of any other transitional arrangements that might be required.

This chapter is structured to lay out the spectrum of options available to deal with one key issue: Namely the potential of various organisational arrangements to limit the negative effects of market dominance and distortion of investment incentives. For this purpose, the options are dealt with in three broad groups:

- “Unified” regimes under which a single organisation is made responsible for operational decision-making, so the focus is on devising arrangements to limit and mitigate the market dominance of that single organisation; versus
- “Diversified” regimes under which a single organisation still operates the facility, but operational decision-making is driven by the competing demands of a number of parties, none of which is deemed to dominate the market, in its own right; and
- “Hybrid” regimes, which combine elements of both.

Not surprisingly, the general thrust of the suggested mitigations and resolutions is towards a hybrid regime that combines organisational and contractual mechanisms in such a way as to gain some of the benefits of diversified regimes, while retaining the benefits of a unified regime, as much as possible.

First, though, Section 6.2 identifies several other “key issues” that will eventually need to be addressed in the context of each organisational regime. While the main description and discussion of each regime largely ignores those issues, the section on each broad group of options includes a sub-section briefly outlining the way in which those key issues could be handled, under regimes of that type.

6.2. Key issues

6.2.1. Market impacts and economic efficiency

Productive efficiency

The distinction between "productive" and "allocative" efficiency is somewhat debatable, in the electricity sector. If a single organisation were running the sector, or even managing a portfolio of assets within it, we might see the coordination between capacity elements within the sector/portfolio as an internal matter, contributing to "productive efficiency". But these elements actually interact via a spot market, and we would typically think of pricing in that market as an "allocative" issue. Then, if we restrict attention to the management of a single hydro station, or catchment, calculation of its EMVSE vector is a critical issue, but its close link to spot prices makes it debateable whether it should be considered as a "productive" or "allocative" efficiency issue, too.

There are definitely productive efficiency issues when it comes to minimising the cost of operations and maintenance activity. And, of course, efficiency should be pursued in that area, but it is a relatively small component of the total hydro system cost, so pursuit of productive efficiency, so defined, should not be the primary consideration when assessing organisational options.

Market dominance and allocative efficiency

One obvious issue is that a single facility capable of creating a large supply/demand swing, over a short time interval, would also have considerable power to shift market prices in the comparatively small New Zealand market. While no-one would claim that the market currently operates under perfectly competitive conditions, it seems reasonable to rule out the option of having such a facility controlled by a single commercially motivated party, unless special organisational/contractual arrangements are made to curb its implicit market power. Thus, this chapter focusses on some particular arrangements that have been proposed, or might be considered to deal with that issue.

Investment incentives and dynamic efficiency

There are three distinct issues here:

Impact of spot market signals

First, one of the key concerns motivating several of the proposals in this chapter would seem to be concern that a large-scale pumped storage development, not driven directly by market forces, could distort the pattern of market prices. That could obviously have detrimental effects on the operating patterns of incumbent generators, and of new generating capacity, once built. But it could also have detrimental effects on the incentives for new plant to enter the market, and shift the mix of entry capacity away from the optimal long-run balance implied by the technology costs.

To be clear, though, the intended effect of investment in a facility of this type would be to lower peak/dry-year prices, and probably also to raise off-peak/wet-year prices. That would clearly reduce incentives for alternative investments in peak/storage capacity, and possibly raise incentives for baseload/intermittent generation, relative to their expected levels with no such investment.

If the facility could be expanded continuously, and indefinitely, the effect would be to maintain a price pattern in long-term equilibrium. But the lumpy nature of the investment options means that prices, and price differentials, should optimally be depressed for some years after the capacity is commissioned, thus deterring any broadly similar investments for an extended period, starting from the time when the large-scale investment was first mooted.

None of this is inherently inappropriate, though. If the pumped storage hydro facility is, itself, sub-optimal, that investment decision may be considered an instance of "dynamic inefficiency", in its own right. But the situation will not be improved by encouraging further peak/storage capacity until load growth, or capacity retirement, eventually brings the Price Duration Curve (PDC) back to an optimal long-run shape which can be supported by ongoing market investment. The concern would be if the dispatch/pricing strategy adopted by the facility, after it became operational, were pushing market prices away from their optimal (post-investment) levels, whether due to market power issues, or inappropriate rule-setting.

Impact of intervention

Second, though, many of the organisational arrangements discussed below could not proceed without Government involvement. Indeed, such involvement is already evident, with respect to investigating development options. And it would not be surprising to see that involvement follow through into decision-making and/or involvement with any actual development, and/or the rules under which it would operate. The details of how that might happen, and whether it should happen, lie beyond our current scope, but the key question, with respect to investment incentives, is whether that process is seen by (potential) market participants as a one-off response to a unique situation, occurring only once, at this particular point in history, or as a pattern for future involvement.

If market participants believe that capacity will continue to be built and/or operated on a non-commercial basis, they will see no point in pursuing investment in options that could possibly deal with the same issues on a commercial basis. There is a real danger that unrealistic expectations and irresponsible decision-making will be encouraged, if participants, and ultimately customers, come to rely upon such interventions to shield them from the consequences of their own actions, or inaction. And there is also a danger that, if the effect of such intervention in one investment class is to undercut the commercial viability of investment in other classes (e.g. by reducing the revenue component they should optimally derive from peak/shortage pricing) an initial intervention will create a "vicious spiral" in which more and more intervention is required.

Impact of restrictive rules

Finally, though some other jurisdictions have attempted to shield the market from the logical implications of similar interventions, by keeping "reserve" plant built in response to non-market incentives, out of the market, as much as possible, for example prohibiting operation of such plant at times when market prices would normally indicate that it optimally should be operating. Since that kind of policy creates its own obvious inefficiency, the pros and cons need to be considered carefully, and are discussed further in the context of the "rule-based" regimes under which such a policy could be implemented.

6.2.2. Hedge market interactions

The presence and operation of a large-scale storge facility would diminish the risks faced by all market participants by reducing supply/demand balance fluctuations, and market price volatility. Individual participants would still want to hedge the (probably quite substantial) remaining risk, though. Chapter 5 discussed a variety of ways in which the pumped storage facility manager could use the physical hedge capacity provided by that asset to back various kinds of financial instruments that would then be bought, and potentially traded, by market participants. Some of those options would be quite specifically applicable to particular organisational regimes, though. So, the discussion of each regime below comments on the kinds of hedge market interaction appropriate to that regime.

6.2.3. Host system interactions

Chapter 4 suggests that, while managing direct physical interactions with the host hydro system may be easy in some cases, it may present quite difficult problems in other cases, including those discussed in Appendix A. As noted in Section 4.6 above, the obvious solution would be to extend whatever regime is devised for managing the pumped storage facility to cover the entire host catchment. But storage/flows for hydro generation in most catchments are currently managed by independent power companies, so a major re-organisation of the sector might need to be negotiated in order to achieve that outcome. And that degree of integration would also add to the market dominance issues already inherent in very large-scale pumped storage development.

Appendix B therefore discusses a number of contractual arrangements that could be employed to achieve a sufficient degree of coordination, without complete integration. The issue is, though, that while full integration remains a technically feasible solution under any of the organisational regimes discussed here, various contractual options seem more compatible with some regimes than others. Conversely, the complexities potentially arising as a result of host system interactions would be more easily handled under some regimes than others. Thus, for each regime, we will comment briefly on how those interactions might best be handled, and what particular difficulties might arise.

6.2.4. Real-time issues

Technological complications

Much of our theoretical discussion, and some of the organisational options discussed in this chapter, implicitly assume that the pumped storage facility can be treated as providing a generation/pumping service with convex costs, over a continuous range of operating levels. That is consistent with the assumptions made about other generating facilities in the market design. But it is obviously recognised that accommodations must be made for the reality that facilities actually consist of individual machines that each must be started/stopped at particular times, and may have to follow a particular pattern of generation levels, and delays, while they move generation up into the range where they can actually operate with convex costs over a continuous range, as assumed.

Many other markets employ “integer” market-clearing optimisations to deal with such effects, particularly for large inflexible thermal plant, but that has not been considered necessary in New Zealand, to date. On the other hand, an integer formulation has been considered appropriate to deal with the various possible configurations of the inter-island HVDC link. So, either option could possibly be adopted to represent the unit commitment decisions for the proposed facility in market-clearing.

At this stage, the technological characteristics of the facilities under discussion do not seem fundamentally more difficult to accommodate than those of, say, large thermal units, though. So, most of our discussions assume that arrangements to deal with this issue could be made, in due time, without major market disruption.

The technological issues would need to be considered more carefully, though, when debating the workability of some of the organisational arrangements discussed in this chapter, where multiple decision-makers might implicitly be involved in decisions regarding the status of a single machine, as discussed in Section 6.4.2 below.

Ancillary service market interactions

Section 3.4.1 discussed the possibility that a large-scale pumped storage hydro facility could become a major provider of ancillary services, particularly regulation and contingency response. The ability to actually provide ancillary services will depend on unit commitment, though, and the most difficult issue would probably be the turnaround delay, when switching from pumping to generation, or vice versa. But the delay involved is certainly less than that to start up many larger thermal units in markets, worldwide, and the technical issues involved here seem not too different from those faced by existing units of similar size.

We note, though, that providing ancillary services of that kind involves withholding capacity that might otherwise have been used to generate, or possibly to pump. And that may be problematic, under some of the regimes discussed here, because the capacity being withheld might have been assigned, in some way, to multiple participants.

System interactions

One major issue, with facilities of the size being investigated, is that swinging from full generation to full pumping mode will have a substantial impact on the supply/demand balance, particularly in the relatively small South Island sub-system. This could greatly affect prices, implying that caution needs to be exercised with respect to the extent that the theoretical discussions of previous sections rely on perfectly competitive assumptions, irrespective of any speculations about the potential for deliberate price manipulation. A large facility could also challenge physical limits, at times, such as intra-island and inter-island transfer capacity.

Accordingly, we expect that moderation would need to be exercised with respect to the level of pumping/generation activity, and the rate at which it changes. In particular, responses that may seem optimal based on observed market prices will need to be moderated by consideration of the prices likely to eventuate if that response were actually to occur. It is very likely, for example, that prices will seem low enough to make pumping economic, so long as no pumping is actually occurring, but would jump up to levels high enough to make pumping uneconomic, if the pumped storage facility were to respond by switching to maximise pumping.

This is not an issue of poor market design, market failure, or market manipulation, but inherent in the underlying national cost-benefit optimisation. The optimal response, if the underlying cost functions actually were continuous and convex, would be to pump at a rate that leaves prices exactly at the level where the system, the nation, and the pumped storage hydro operator, are indifferent as to whether any further pumping is worthwhile. So, if the interaction is managed well, the implication should be that market prices are actually stabilised by the pumped storage facilities offer/bid, potentially over quite a wide supply/demand balance range.

There could be an issue, though, if the pumped storage facility was operating under simplistic rules specifying that maximal pumping or generation be triggered by a physical or price-based test. And there could also be an issue in the context of some of the “diversified” arrangements discussed below in dealing with the reality that market prices actually will move in response to the aggregate actions of individual participants, even though they each believe that prices will be unaffected by their individual actions.

6.3. Organisational arrangements for unified decision-making

6.3.1. Introduction

This section discusses options in which all operational decision-making occurs within a single organisation, but most likely bound by behavioural rules of some kind. We also cover two alternative cases, in which a single entity is allowed to operate freely to maximise its own profit, on the one hand, or constrained to follow the recommendations of a national benefit optimisation model, on the other.

By way of contrast, the next section discusses options under which, while one organisation is responsible for the maintenance of the facility, and for actual real-time control, contractual arrangements incentivise that organisation to adopt an operational/pricing strategy reflecting the aggregate demanded by individual market participants, each making their own decisions about how the portion of the facility to which they have been assigned "rights" should be operated.

6.3.2. Commercial operation

Description

In principle, a large-scale pumped storage hydro facility could be operated by a commercially focussed power company, just like any other plant in the sector. As such, it would operate essentially as described in Chapter 3 and/or 4, with no other specific restraints on its behaviour. That is, it would dynamically optimise its pumping/generating strategy in repose to market prices, and perhaps to shape market prices, with a view to maximising its own profit.

Discussion

Such a company could be owned by a major gentailer, or by a consortium of generators and/or gentailers. Of itself, though, it would be a net buyer from the market, and so unlikely to support a presence in the retail market. As a net buyer, we might think that its overall incentives would be to keep prices down, but (in the absence of any contractual commitments) it would only be motivated to do that when it was pumping. When generating, it would benefit from higher prices, thus giving it an overall incentive to exacerbate price differences, or at least not to pursue its arbitraging role to the point of eliminating the differentials that make that role profitable.

None of that seems problematic, if the facility is small, but it could be quite problematic if the facility was large enough to systematically influence pries, potentially to its own advantage.

Pros

- Simplicity;
- Internal efficiency; and
- Consistency with status quo arrangements.

Cons

- Market power/distortion; and
- Likely perception of increased risk to other market participants due to potential for exploitative, manipulative, and/or erratic behaviour.

Mitigation

The negative points noted above could theoretically be resolved by a joint ownership structure, but that raises the prospect of all generators colluding to manage sectoral outcomes by coordinating the behaviour of this facility, and potentially much else besides.

Public ownership might be suggested, but that model does not seem suited to "commercial" operations. While a public owner might pursue profit maximisation less aggressively, that suggests less focus on cost minimisation, and potentially greater focus on other objectives, at the expense of economic efficiency. So, even if that kind of ownership arrangement was preferred, we suggest that it might be better cast as a version of the "model-driven" arrangements suggested in Section 6.3.5 below.

Various forms of "contracting" might be proposed to mitigate a large facility's market power, but we must ask what form those contracts might have, and who they would be with:

- If the contracts were with the Government, or some regulatory agency, those contracts would presumably specify rules of conduct, and penalties for misconduct, in which case this would look like a particular instance of a "rule-based" regime, as discussed in Section 6.3.3 below.
- If the contracts were with market participants, behaviour would be determined, in some way, by those participants, in which case this would look like a particular instance of a "diversified" regime, as discussed in the Section 6.4 below, albeit perhaps with more traditional, less flexible, and in our view generally less desirable, contracting arrangements.

Dealing with key issues

With respect to the other "key issues" identified in Section 6.2, we note that:

- Technological, ancillary service, system, and hedge market interactions would be handled just as for any other large market participant.
- In principle, this regime might also deal with host system interactions perfectly, because there would be strong commercial incentives for the two parties to cooperate, if not amalgamate. But that kind of cooperation and/or merger would create an even larger entity, with greater market power, So, it would surely be prevented, thus leaving the parties with potentially quite limited coordination.
- Investment incentives would obviously be impacted by the presence of the facility, irrespective of regime, but the nature of that impact would depend critically on the decision-making process surrounding investment in the facility itself. In principle, a fully commercial development, driven entirely by market price signals, would not "distort" incentives at all. If the facility is large, it may suppress investment in other capacity for several years, but that is actually the correct national benefit optimising response, and would also occur under a centrally optimised regime. Essentially the same issues would arise for any other investment of similar scale.[127] If the investment is inconsistent with market price signals, though, its impact on investment signals will depend on the degree of inconsistency, and on whether that investment is seen as a one-off anomaly, or as a precedent for more similar intervention, as discussed in Section 6.2.1.[128]
- But the regime, on its own, clearly does not control market power, and this forms the central motivation for considering all the other regimes discussed below.

[127] The more concerning issue may be that optimal operation of an optimally sized facility, once built, would suppress its own profitability for such a long period that it would not actually enter, or be optimally scaled. And it is notable that no large-scale hydro generation developments have actually proceeded since market start. That partly reflects technology trends, but is also an obvious motivation for giving special consideration to investment in large scale storage facilities, at this time. Thus, we do not address it as an "issue" specific to any of the regimes discussed here, because it is inherent in the assumptions underlying the whole report.

[128] Again, noting the likelihood that a large-scale investment that is actually "optimal" will appear "more than economic" at prices projected assuming that it does not enter, but then appear "uneconomic", for many years after entry, at prices projected assuming that it does enter, and then operate in an economically optimal perfectly competitive manner.

Assessment

It is not impossible to imagine a large storage facility being operated as a "commercial" enterprise. But we believe that the market, and the public, would want to see that a "commercial operator" was constrained and/or incentivised either to operate in particular ways, or to cede effective control to other parties who would face enough competitive pressure to not exploit the potential market dominance that such a facility might have. Thus, this does not really qualify as a viable option, in its own right.

6.3.3. Rule-based regimes

The discussion in this section covers three broad types of "rule-based" regime, and briefly outlines their specific pros and cons, before concluding with an assessment of all three types.

Physical rules

Description

Rather than dynamically optimising its pumping/generating strategy in repose to market prices, and perhaps to shape market prices, a unified organisation operating a pumped storage facility could be given a set of well publicised rules under which specific physical responses (such as "maximise pumping" would be mandated in situations meeting specified physical criteria (such as "storage below some calculated guideline"), or it could develop and manage such a set of rules, in consultation with the sector, and the regulator.

Discussion

Possible rationales for this proposal might be to establish the facility in such a way that even the general public have a clear understanding of what it is supposed to do; and to prevent it from exploiting market power to influence prices in ways that enhance its own profits, but do not benefit the nation. A further objective could be to prevent the facility from taking on other roles, even if they do turn out, or appear to be, economic from a national cost benefit perspective.

Pros

- Clarity of purpose;
- Physical predictability; and
- Control over deliberate price manipulation.

Cons

- Complexity of developing rules for a wide range of plausible situations involving combinations of water levels in various reservoirs, weather forecasts, host flow conditions etc.;
- Obvious economic inefficiency if the facility is operating when the economic signals say it should not be, by pumping when the price is higher than MVP, or generating when the price is lower than the MCG;
- Less obvious economic inefficiency when the facility is not operating when the economic signals say it should be;
- Possibly more hidden inefficiency, if host system conditions impose significant constraints on operations, and the physical rules do not account for those conditions in ways yielding similar outcomes to the (hypothetical) optimal national benefit optimisation assumed in Chapters 3&4;
- No control over inadvertent price impacts, across any price range; and
- Very likely excessive discouragement of some investment types, and excessive encouragement of others, in ways that may not be easy to understand, or offset.

Mitigation

The following measures might be suggested to mitigate some of the negatives raised above, but see further discussion of the likely effectiveness of these measures, as they might apply to all of these rule-based options, under the "Assessment" heading, at the end of this Section, 6.3.3.

- Making the rules comprehensive enough to allow operation in a wide range of circumstances, including the status of storage and projected flows in the host system, and in other significant national reservoirs;
- Making the rules about physical operation subject to economic as well as physical tests, as discussed below;
- Optimisation of the rules, as discussed in Section 6.3.5 below; and
- Compensation to "correct" incentives for affected investment classes, although that might not be easy to compute or agree.

Price-based rules

Description

Instead of pre-defining physical tests and rules to determine operational strategy, we could possibly pre-define market price levels, or price ratios, at which the facility would respond in particular ways:

- Although round-trip losses imply a natural percentage buy/sell spread, a specific rule might be helpful in precluding certain kinds of gaming; and/or
- The facility could be bound to generate at whatever rate seems necessary to try to limit market price volatility; and/or
- The facility could be bound to generate at whatever rate seems necessary to try to keep the market price below some pre-announced cap level; and/or
- The facility could be bound to pump at whatever rate seems necessary to try to keep the market price above some pre-announced floor level; and/or
- The facility could be bound to make buy/sell offers at specified levels in spot and/or hedging markets.

Discussion

One rationale for this proposal could include limiting the risks faced by both producers and consumers, while preserving market-driven signals over most of the price range within which useful production/ consumption responses can be expected. Another could be providing other storage managers with clear reference points against which to opportunity cost their limited "fuel" stocks.

It is unclear what rules might be set across the, possibly wide, range of prices between the floor and the cap, though. As a net load, the pumped storage facility may be thought to have incentives to push overall price levels <u>down</u>, not up. Really, though, its profits would depend on price differentials, so a profit maximiser would want to see lower prices when the facility is pumping, but higher prices when it is generating. And it could achieve that by simply restraining both pumping and generation activities so that prices do not fall as much as they might when generating, or rise as much as they might while pumping. So, if the specified spread takes away its freedom to exercise that kind of market power, no other rule might be considered necessary.

Pros

- Clarity of purpose;
- Commercial predictability;
- Specific controls over both deliberate and incidental price manipulation; and
- Relative simplicity.

Cons

- Possibly inadequate control over price levels across the range between specified limits;
- Arguable economic inefficiency in over-riding market response at the extremes of the price distribution;
- Clear economic inefficiency if the configuration implies that (probably private) upstream/downstream host system EMVSE levels should pay a significant role in determining optimal dispatch, but they are not accounted for in the spot/energy market price response rules;[129] and
- Distortion of entry incentives, with peaking capacity perhaps unduly discouraged by intervention to set a cap on market prices, and/or baseload/intermittent capacity perhaps unduly encouraged by intervention to set a floor on market prices.

Mitigation

Much the same set of measures might be suggested as for the "physical" rule-based regime, to mitigate some of the negatives raised above, and there is further discussion of the likely effectiveness of those measures, as they might apply to all rule-based options, in Section 6.3.4, below. But measures that might be specifically considered when rules are driven by, or oriented towards, setting prices rather than physical triggers or outcomes, could include:

- Computing, declaring, and accounting for the nationally optimal upstream/downstream water values for the host system, jointly with those in the pumped storage facility itself;[130] and
- The encouragement of independent investment in renewables implied by supporting a price floor being balanced against the discouragement implied by attempting to maintain price caps.

Hybrid rules

Description

The facility's operation could be governed by a combination of physical and price-based rules. For example, a "forbidden zone" could be defined, with a rule specifying some combination of DSM response and pumped storage generation, when aggregate storage falls below that level, while leaving operations perhaps subject only to price-based rules when storage is above that level.

Discussion

The rationale for this kind of proposal could be to provide a higher level of "physical" assurance to consumers and/or their political representatives that "the lights will not go out" than might be provided by a purely price-based regime. In the past, hybrid arrangements of this kind have often been justified on the basis that the "costs of non-supply" are not well known. This lack of knowledge, in turn, has partly reflected the fact that mechanisms for implementing "Demand Side Management" of any kind have been poorly developed, and often imposed upon consumers in ways that have provided no way for the costs of that imposition to be assessed.

This is an area that is going to require significant attention, if New Zealand is to live comfortably with a power system relying heavily on intermittent renewables. And the MDAG study expresses the hope that, as traditional thermal support and reference points are withdrawn, clearer marginal cost indications may eventually be associated with "DSM guidelines" of this type. That development could hopefully see "physical" and "price driven" management philosophies converging. In the meantime, though, some

---

[129] As discussed in Section 4.5, the true national cost of delivering a unit of water to the upper reservoir of the pumped storage facility optimally must implicitly include the opportunity cost of inducing the upstream host system reservoir to release that water at a time that would not otherwise be optimal for it. And the true national cost of releasing a unit of water from the upper reservoir of the pumped storage facility could also include the opportunity cost of inducing the upstream host system reservoir not to release water at that time which would otherwise be optimal for it.

[130] Although it is hard to imagine that happening, unless their operational management is integrated.

parties may well derive additional assurance from the addition of physical guidelines, applying to situations where storage may be approaching extreme levels.

Pros

- Some degree of continuity with traditional practices
- Greater freedom for market forces to operate over a wide range of storage levels. than under a regime based entirely on physical rules; and
- Increased political/public assurance and physical predictability, relative to price-based rules.

Cons

- Increased complexity and potential confusion;
- Economic inefficiency due to likely inconsistencies between pricing guidelines, and the prices implied by physical guidelines and
- Risk of gaming to exploit any inconsistencies between physical and economic rule sets, including the boundary defining the regions within which applies.[131]

Mitigation

Particular care would need to be taken to try and align physical and pricing guidelines, so that they do not give incompatible signals. Apart from that, most of the measures applicable to either physical or price-based rules could also be applied to a hybrid rule set. But we now turn to discussion of the likely effectiveness of all those measures, as they might apply to any of these rule-based options.

Dealing with key issues

With respect to the other "key issues" identified in Section 6.2. we note that:

- These regimes all aim to control market power by restricting the range of actions available to the management of the regulated facility. This does not actually mean that their actions will not influence market prices, or that the pumped storage operator will not profit from them, but restricts their ability to prefer actions that increase their profits. The effectiveness of this kind of conduct regulation may be debated, and experience suggests that the apparent simplicity of whatever rules were proposed could prove deceptive, once implementation development started in earnest. But the kind of development path involved at least seems clear.
- All of these unified regimes are basically designed to internalise decision-making with respect to unit commitment, and loading levels, and the issues involved would be much the same as for other large-scale facilities under unified management.
- Technically, a rule-based regime does not make it any more difficult to manage interactions with the host system, but it does raise the prospect that explicit rules may need to be agreed, publicised, and monitored to manage those interactions. And other rules would probably need to include features reflecting the impact of costs (or potentially benefits) and constraints potentially arising from that interaction.
- Discussions about "rules" naturally tend to focus on rules governing physical operations and/or spot market interactions. Chapter 5 argues that hedge market interactions are important, too, and a rule-based operational regime raises the prospect that the dispatch/pricing actions dictated by the rule will be sub-optimal with respect to system/market conditions, at any particular time, and hence also with optimal hedging strategy. A storage entity trading hedges at uncontrolled market prices could find itself exposed to significant risk if simultaneously required to trade at specified prices in the spot market, particularly because all other traders would know those prices, and understand that the storage trader could not respond dynamically to any arbitrage strategy they might wish to pursue. So, we suspect that any rules determining spot market offers and bids would also explicitly

[131] By design, there will be inconsistencies to exploit, because physical guidelines will actually have no effect, unless they are sufficiently inconsistent with pricing guidelines and/or market prices to incentivise or force behaviour to differ from optimal responses to market signals.

or implicitly determine hedge market offers and bids. But that risks placing the storage trader in an even more extreme position, vis a vis market valuations, and/or forcing them to withdraw from the hedge market entirely. This aspect of the situation should be examined more carefully before any commitments are made in this area.

- As noted earlier, the existence of a pumped storage facility would obviously have a major impact on investment incentives that, of itself, is intentional and appropriate. Under a rule-based regime the specific rules under which a particular facility operates will determine whether the impact on the pattern of market prices will encourage or discourage particular types of entry, relative to their optimal levels. Rules could be devised or adjusted specifically to encourage or discourage particular technologies, and some may see the opportunity to tweak rules in pursuit of "policy goals" as a positive feature. From an economic perspective, though, it would be regarded as undesirable, unless a clear case can be made in favour of intervention to correct some identified "market failure". And pervasive distortions should be expected as an inadvertent by-product of policies intended to further other goals. Thus, we are more inclined to see this as a negative feature, but discuss the issues further in Section 6.3.4 below.

6.3.4. Assessment of rule-based regimes

Comparison with status quo

One common feature of all the rules-based options is that they imply a prominent role for some kind of regulatory body tasked with making, monitoring. and enforcing rules. Economic purists will definitely see that arrangement as "imperfect", but such imperfections often represent necessary compromises when the underlying situation itself is imperfect. And there clearly could not be a perfectly competitive market, if sectoral operations are physically dominated by a single facility, controlled by a single party, as assumed here.

Indeed, it should be recognised that the sector already operates under a hybrid regime, with particular actions triggered by physical measures of aggregate system storage levels. We have not reviewed the operations of that regime, and thus have no opinion as to whether it is actually necessary, or efficient, under the status quo. But the underlying motivation of ensuring supply security will surely become more important in coming years, and we must assume that any significant new storage would surely be included in the calculation of national supply security, and expected to cooperate with whatever actions were deemed to be necessary.

Conflicting motivations for rule-making

The range of actions available to pumped storage would be wider than for conventional hydro, with daily cycling and/or contingency response occurring irrespective of national storage system status. So, we could imagine new rules being developed to manage particular aspects of pumped storage operation across a wider range of situations. But the motivation for developing rules to deal with situations other than a national shortage of stored energy would surely be growing concern about the prospects that socio-economic damage might be caused by poorly coordinated management of situations such as the dunkelflaute, as identified by MDAG. This is a national electricity market issue, though, not specifically triggered by the mere presence of a major storage facility. In fact, the presence of such a facility may considerably reduce the perceived need for such rules.

So, while the form of some rules may be similar, the motivation behind many of the rules that could be contemplated for application to a large storage facility is actually national supply security, and nothing to do with the specifics of that facility, or control of market power. And that may be problematic, because it seems likely that those, really quite different and potentially conflicting, priorities will not be clarified, but confused. They may even be consciously or sub-consciously obfuscated, with the real motivation behind rule settings, or actions, being different from the announced intentions.

That is concerning, because positive rules of the type mainly discussed above, that is rules that do not prohibit some type of activity, such as collusion, but require some party to act in a particular way, do not actually reduce the ability of a large facility to influence market prices and outcomes. They just channel that power to influence market prices and outcomes in particular ways. And that effect will be reinforced, to the extent the rules provide "certainty". Accordingly, it should be recognised that applying such a regime to a major facility would not just be an incidental or inconsequential change, but a distinct move towards quite intrusive "market management" of a type that was rejected by the WEMS study that established the market, and has been largely eschewed to date.

Inflexibility

Given the potential (and intentional) importance of a large storage facility in determining market outcomes, the potential sub-optimality and inflexibility of the rules must also be of concern. The simpler the rules are, the more often they may need to be adapted to deal with circumstances that differ from those assumed in deriving them. And significant effort should, and hopefully would, be expended on optimising the rules, and then hopefully updating them regularly as circumstances change. In fact, the whole rationale for having such a facility may change if it turns out, as the sector evolves, that the problem originally identified is not actually the most important one, in future. In fact, MDAG predicts that priorities probably will shift away from "dry year firming", towards intra-day peaking, and dunkelflaute management.

In the limit, such updates could be performed quite frequently, and the regime could become very much like the "model-driven" regime described in Section 6.3.5, below. We will argue that that kind of regime offers its own kind of "certainty", but it clearly does not provide the simplistic form of "certainty" that may be envisaged by its advocates. To do that, rules and settings would have to remain constant for extended periods, and only change in a very measured and controlled way. That can be expected to decrease the average goodness of fit to particular circumstances, but that loss has to be balanced against the fact that every change increases the opportunity for "mission creep", influenced by conflicting objectives.

Long-term instability

In fact, while the general public may welcome the supposed "certainty" of rules imposed by "the Government", the sector may well see that as increasing long term risk, rather than reducing it. The problem is that, no matter how well intentioned the current government may be, and no matter how rigorously rules might be derived to meet that government's objectives, it is simply impossible for any current government to stop future governments from changing the rules and/or allowing a gradual process to erode and undermine their original intent.

Steps should obviously be taken to strengthen the independence of the entity responsible for tasks like rule-making and monitoring. But the suspicion must surely remain that the rules would actually be changed if, say, it turned out that those rules implied allowing prices to be high enough, or volatile enough, to create widespread unhappiness amongst voters. The longer history of the New Zealand electricity sector shows that past governments have been quite ready to subvert electricity sector economics in the name of managing other goals related to employment, inflation, revenue collection, or social objectives. So, it is not hard to imagine future Governments exerting and responding to similar influences, in a world where pressures of various kinds look likely to increase.

Prohibition of beneficial activity

This is particularly true if "clarity of purpose" is to be achieved by "negative" rules, precluding a facility from performing some roles that it is physically capable of doing, and which would actually deliver measurable benefits to the nation. For example, one could imagine a prohibition on daily cycling, as a way to boost private investment in batteries. Such a rule might actually be relatively easy to state, and enforce, but gives rise to two distinct causes of concern:

- First, we are aware of similar prohibitions being imposed on, say, extreme peaking reserve capacity, in overseas markets, for broadly simar reasons. Such a rule might be "optimal" in some second-best sense, when it serves to maintain old capacity that would otherwise be scrapped, in a backup role. When applied to a new development, though, the problem is that the net cost to the nation would rise significantly, and perhaps prohibitively, due to the cost of investing in large-scale batteries, as well as investing in a large-scale facility that could perform a similar role, with broadly similar efficiency, but is prohibited from doing so.[132] And consumers/taxpayers would ultimately have to pay those extra costs, one way of another.
- Second, though, the obvious artificiality of such a prohibition would surely leave a potential battery investor wondering whether the policy was durable. And factoring in the risk of a later policy reversal can only increase potential investors' assessed risk, and required rate of return, to the ultimate cost of consumers.

In other words, the apparent "predictability" offered by these rule-based regimes could be considered deceptive. Such rules make no difference to outcomes, unless they do imply incentives at odds with those emerging from the market. That tension between conflicting objectives implies an inherent instability. And the response of future governments to future public pressures is arguably even less predictable than that of a profit-maximising market participant, who can at least be relied upon to pursue that goal consistently over time, and can be consistently regulated accordingly.

Efficiency and Effectiveness

Finally, this kind of intrusive rule-making was rejected by the WEMS study because there was a broad consensus, internationally, that it would have a negative impact on productive efficiency (by dis-empowering and thus incentivising management), allocative efficiency (by forcing what are likely to be sub-optimal prices on the market), and dynamic efficiency (by limiting scope for innovation, and distorting investment signals). We will not attempt to re-hash those arguments here, but do not believe that consensus has shifted much, over the intervening years, or to assess the actual magnitude of any effects, in this particular context. Other things being equal, though, the "structural regulation" approach represented by the "diversified management" regimes discussed in Section 6.4 may be preferable.

6.3.5. A model-based regime

Previous sections have discussed the possibility that a single entity might determine and control pumping/generation/storage/hedging policy according to pre-determined rules of some kind. It may be argued that such a regime has the advantage of giving the market, and perhaps the populace, some degree of certainty. Put another way, it will definitely drive market outcomes in ways that are predictable, at least between rules re-sets.

---

[132] Note that a facility capable of managing inter-annual hydrological fluctuations, will also be capable of managing intra-day or inter-seasonal cycling, with the minimum timescale over which it might operate mainly limited by the speed with which pumping/generation rates can be adjusted, or switched. In fact, this is the only role of many pumped storage hydro facilities, overseas. We suggest that it should be considered the primary role here, too, in the sense that it will be the cheapest role to provide for, because it does not require construction of a large upper reservoir, or even a particularly large buffer storage. Such a project may not be worthwhile, in the New Zealand context, because it would involve a large investment in pumping/generation capacity for a facility that could only play a single role.

Once those costs and benefits have been established, though, the question becomes one of determining the incremental benefit of increasing storage capacity, so as to enable arbitrage over longer timeframes. That may well turn out to be the primary driver of any development, in the sense that it delivers the largest economic benefit. But the point is that it could well make economic sense to "prohibit" a facility from performing a longer-term arbitrage role, by simply not paying the cost of building storage capacity beyond some level. But, if that cost has been paid, it does not make economic sense to prohibit a facility from also performing a shorter-term arbitrage role.

That is, of itself, a kind of market "distortion", but it may be considered more desirable than distortions driven by commercial parties, because it is not directly driven by any advantage the owner may be seeking. That does not, of itself, necessarily mean that it delivers benefits to any party, though, or that the advantages it delivers to some parties outweigh the disadvantages to others, across the economy as a whole. So, we must ask:

- Who would determine the rules?
- What criteria would be applied in determining them?
- What analysis would be used to determine their "optimal" settings?
- How complex might they become? and
- How often might they be revised?

These last two questions are closely connected, because we have seen that the situation to be managed is quite complex, and will be constantly changing. So, there will be a trade-off between increasing the complexity of the rules so as to cope reasonably well with a variety of possible future circumstances, vs simply revising the rules frequently, as circumstances change.

Ultimately, if we follow the latter route to its logical conclusion, we could get to a point where the rules, or rule settings, were being revised every month, or perhaps every week. At that point, the regime might be better described as "model-based".

Description

Under this option:

- A formal reservoir management model, applying internationally accepted stochastic optimisation techniques,[133] would be run regularly to maximise a net national benefit "objective function", assuming the current/forecast status of the facility, and other key components of the national electricity system, forecast weather /load conditions, host system flows etc.
- The EMVSE and/or MCG/MVP determined by that optimisation would be published and held constant for some period, such as a week or month, or perhaps adjusted in some pre-announced fashion, as a function of a hydrology index.
- Buy/sell offers in spot/hedge markets would then be based on the announced MCG/MVP.

Discussion

The MCG/MVP pair referred to here should really be determined by the difference between the upper and lower reservoir EMWV levels. The upper reservoir EMWV level should really be very stable, unless that upper reservoir is close to its bounds, and might only need to be updated, say, quarterly. And the more reliance is placed on index-based updating, the less frequently the model would need to be re-run.

The EMWV of the lower storage/buffer pond may be a trickier issue, and Section 4.7.2 suggests that it could vary substantially, over a much shorter planning horizon, if capacity and/or flows are limited. But much of that variation could be in the form of predictable cyclic variation, and much may depend on easily measurable host system storage/flow levels. So, even if the MVS difference is not constant over a week or month, its behaviour could still be sufficiently predictable to form the basis for an approximately optimal bid/offer strategy that might not "fix" bids and offers over the period, but could allow them to be determined as a function of easily observable parameters.

At one level, this might be interpreted as setting "price-based rules" with very frequent revisions. At another level, though, the "rules" become rules about how the objective function and other inputs to the model should be determined, how it should be applied, and how bids and offers should be determined from the results. In other words, it addresses issues 2-5 above. The challenge then becomes to establish

[133] Such as the SDDP or DOASA models regularly used by the Electricity Authority and others, for various studies. See:
M.V.F. Pereira, and L.M.V.G. Pinto "Multi-stage stochastic optimization applied to energy planning" *Mathematical Programming* v52, pp. 359–375, 1991
. A.B. Philpott and Z. Guan. Models for estimating the performance of electricity markets with hydro-electric reservoir storage. Technical report, *Electric Power Optimization Centre*, 2013.

an entity, that is accepted across both demand and supply sides of the sector, and by New Zealand consumers, as having the requisite expertise, and credibility to establish and manage such a process.

Pros

- A clear national benefit objective;
- Flexibility to adapt to changing circumstances;
- Ability to account for much greater complexity than could be codified in the form of rules;
- Short term predictability in the form of a pre-announced offer/bid strategy; and
- Longer term predictability with respect to the process, and the way an optimisation model will respond to changing situations.

Cons

- Possible extra cost;
- Inability to prevent the optimisation from recommending operating strategies that meet its national benefit objective;[134]
- Reliance on a "black box" not easily understood by the general public;
- Potential for dispute over modelling methodology and assumptions; and (conversely)
- Reliance on a single decision-maker and/or model.

Mitigation

- Establishing a regime under which the model assumptions and methodologies are regularly subjected to critique by local, and perhaps international, experts, somewhat like the process currently applied to SPD, but with less emphasis on real-time computational robustness, and more on the robustness of assumptions about future conditions, and on the way those are accounted for by the particular stochastic optimisation methodology chosen.[135]
- Establishing a regime under which MVS or other relevant parameters are not just determined by a single model, but by consensus between competing models developed and/or run by independent parties.
- Allowing the responsible authority some leeway to adapt model recommendations to account for factors not readily represented in the optimisation model.

Dealing with key issues

With respect to the other "key issues" identified in Section 6.2, we note that:

- This regime aims to control market power by having the pumped storage facility's market interaction strategy determined by a national benefit optimisation, which should produce something akin to perfectly competitive policies.
- An optimisation model should also be much more capable than any simplified rule set, of accounting for, and perhaps managing, detailed host system interactions.
- The discussion above assume that the optimisation would also be used to drive a hedge market strategy, without specifying what kind of hedges might be involved. We consider that the "virtual storage" rights discussed in Section 5.4.5 (and/or perhaps the FSRs discussed in Section 5.4.4) form a natural basis for the alternative "diversified" regimes discussed in Section 6.3 below, while the "tank options" discussed in Section 5.4.6 form a natural basis for the "Hybrid" regime discussed in Section 6.5.2. Thus, while either could be issued under this regime, the primary strategy optimised

---

[134] It may seem odd that inability to prevent an "optimal" outcome should be listed as a negative feature, and clearly it is not, from a mathematical optimisation perspective. But discussion of some of the "rule-based" options suggests that the ability to adopt rules that over-ride short term optimality in order to achieve longer term "policy goals" might be considered a "positive" feature. And the point here is that the mathematical structure of optimisation models would generally preclude that, or at least require it to be done by an explicit adjustment, outside the optimisation model. See further discussion in the "Assessment" section below.

[135] Recognising that there is much more debate, nationally and internationally, about the ideal treatment of future uncertainty and risk, than there is about the deterministic techniques applicable to a single period optimisation like SPD.

under this regime (for the capacity not covered by virtual/tank options) would most likely relate to operations in conventional hedge markets, but with a strong emphasis on put and call options, rather than "energy" contracts, per se.

- It seems no more difficult to manage technological, ancillary service, and system interactions under this regime than under any of the rule-based options discussed previously, with the extension that an optimisation model could also be applied to set ancillary service market offerings.
- Investment incentives would obviously, and intentionally, be impacted by the presence of the facility irrespective of regime. If we believe that investors are likely to be sceptical of the long-term "certainty" provided by announced rules, though, they should be reassured by the predictability, and likely durability, of a regime under which the activities of a major system component are systematically and transparently optimised to a well-defined objective. This is, after all, similar to the way in which the market has been cleared for over 20 years now, by SPD.

Assessment

This regime is somewhat similar to that adopted by ECNZ prior to market start, under which an optimisation model of what was effectively the entire national system was run weekly, to set hourly North Island and South Island "spot prices" for the coming week, with electricity bought and sold in the form of hedges settled against those prices. Integral to that regime was a commitment to set prices "as if in a perfectly competitive market", so the model was run with a national benefit objective.[136] It is perhaps even more similar to the widespread practice, in other hydro-dominated power systems, of using a national benefit optimisation model to determine the operating strategy for the entire national system, including MVS for each reservoir.[137]

The proposition discussed here is different, in that it would focus on a particular facility, and only set strategy for that facility, with spot prices later determined dynamically by market interactions. And it would provide much more transparent signals in the form of announced MVS levels. That transparency would also set up a dynamic under which other participants would feel some degree of pressure to align their own behaviour with the announced "perfectly competitive" MVS values. And they would obviously tend to set their offers in relation to those announced for the pumped storage facility. So, a single model could have significant influence on market outcomes, over and above controlling a single facility, which will itself have significant influence.

It should be recognised, though, that the whole point of developing a major storage facility would be to "have significant influence on market outcomes", in ways that "maximise national benefit". So, while the scenario above may be criticised as facilitating a form of tacit collusion, it may also be welcomed by many parties.

We have listed: "Inability to prevent the optimisation from recommending operating strategies that meet its national benefit objective", as a negative. But that is only negative if the ability to prevent the facility pursuing operational strategies that appear to improve national benefit was seen as a positive reason for adopting a rule-based approach. We acknowledge the possibility that there could be situations in which the policies recommended by a (relatively) short term operational optimisation should be over-ridden in the interests of longer-term objectives, including investment signalling. Depending on the mathematical properties of the particular optimisation model, adjustments might also have to be made to inputs or outputs to account for factors such as national security risk.

Thus, we accept that it may be desirable to give the responsible authority some leeway to adapt model recommendations to account for factors not readily represented in the optimisation model. But, if such practices are to be adopted, we believe the adjustments, and the reasons for them, should be explicit, visible, and clearly costed, rather than being implicit in some vague, unquantified, backroom trade-off that might not even be well understood by those making it. Thus, we suggest that the discipline imposed by formal optimisation would keep the primary focus on electricity sector economics, help prevent

[136] In fact, SPECTRA, the model employed, had been developed by the Ministry of Energy, and could not be run with any other objective. And it was run by the same team, using essentially the same assumptions, so far as we are aware.

[137] Most, if not all Latin American markets operate in this way, using the SDDP optimisation package.

"mission creep", and make the regime more credible, and durable, in the long run. The short-term predictability provided by rigid rules that purport to over-rule the natural variability inherent in the sector, would be replaced by longer term assurance that a consistent objective would be pursued, and that policies would continuously adapt to changing circumstances in a logical, verifiable, and predictable fashion.

We also suggest that the concern about a model driving market outcomes is quite misplaced, if a comparison is being made with the rule-based regimes discussed in the previous section. Each of those regimes would drive market outcomes just as strongly, only in ways that are less well adapted to changing circumstances, and almost certainly less effective in maximising national benefits. Nor would any of those ostensibly simpler regimes necessarily be any less expensive to establish and maintain. Industry/expert "rules committees" may seem inexpensive, if relying heavily on seconded personnel, and/or consultants paid for by participants, but the true national cost is another matter.

By way of comparison, a committee responsible for rulemaking would still have to do, or at least commission, its own modelling. But modelling would be more likely conducted in an ad hoc non-transparent fashion, in that case. Thus, the basis and motivation for the eventual rules would be less clear, even to experts, and hence less open to challenge, and more open to manipulation or inadvertent "drift". There would be potential for dispute over modelling methodology and assumptions under any of those regimes, too. It's just that the disputes may be less well informed, and less tightly focussed on national benefit maximisation, under a rule-based regime.

In summary, then, we do not regard this "model-based" regime to be perfect, but consider it to be superior to any of the "rule-based" regimes discussed in the previous section, if a "unified" management regime is ultimately preferred. But we would suggest that, rather than relying on a single model, the process should be broadened to at least allow for "contestable advice" and "competition in the marketplace for ideas", to inform the "unified" model-based management approach discussed in this section. But the next section takes that concept much further, turning to consider the potential benefits of much more "diversified" approaches to decision-making.

6.4. Organisational arrangements for diversified decision-making

6.4.1. Introduction

The previous section discussed regimes under which all operational decision-making occurs within a single organisation, but probably bound by behavioural rules, or driven by an agreed optimisation model. By way of contrast, this section discusses options under which, while one organisation is responsible for the maintenance of the facility, and for actual real-time control, contractual arrangements incentivise that organisation to adopt an operational/pricing strategy reflecting the aggregate demanded by individual market participants, each making their own decisions about how the portion of the facility to which they have been assigned "rights" should be operated.

Motivations

The motivations for these proposals are only partly about limiting any potential "abuse of market power". In theory, market competition is not just supposed to "control prices", but to create drivers to improve efficiency by reducing costs, boosting innovation, etc. So, healthy competitive markets are supposed to go beyond merely "protecting consumers" to deliver positive net benefits to the nation as a whole.

Beyond that, again, fostering diversity not only creates the potential for greater innovation, but moderates aggregate national strategies, to produce more robust outcomes over the longer term. This report has frequently referred to "optimisation models" and, by definition, any such model will produce recommendations that are as good as they can possibly be, given the data supplied to that model, and the assumptions inherent in its mathematical structure. It should be recognised, though, that different analysts will think it wise to build different assumptions into their models and obtain different data estimates, while different decision-makers will place different interpretations on the results.

Some, particularly smaller, participants may consider that no formal optimisation is justified at all, with respect to the operation of the capacity elements they control, but they will still contribute their experienced judgements.[138] In the end, an aggregate market strategy emerges as the result of interactions between bids and offers made by all decision-makers. Conversely, all decision-makers will be continually comparing market outcomes with their own projections, and adjusting their assumptions to more closely align with a market consensus, largely determined by what other parties believe to be better assumptions.

The result will be a market outcome that no one participant or analyst thinks is "optimal". But there is no reason to think that any one participant, or analyst, has a monopoly on perceptions of the true current or future state of nature, technology, the New Zealand economy, or whatever. Thus, the consensus market outcome represents a compromise between all of the typically more extreme positions the various participants might have taken. It will be dominated by the strategies of larger participants, who can be expected to have analysed the issues more carefully, but still influenced by all. Accordingly, it should be expected to be more moderate and robust than any one participant's strategy, and to adjust in a more stable way over time, as the perceptions of various parties change.

Against this, there may be no party large enough to pursue some possible large-scale developments that would actually be optimal, from a national perspective. So, that is the central trade-off underlying any comparison between the diversified regimes advanced here and the unified decision-making regimes discussed in the previous section, as it is in the wider investigation to which this report contributes.

---

[138] In fact, the managers of non-storable renewable capacity may have little control, and hence little to gain from optimisation.

Mechanisms

We ignore the possibility of fully physical diversified arrangements, under which a different party might be responsible for the physical maintenance and control of a particular unit, but all share the same civil infrastructure, and perhaps the same control room.[139] Instead we focus on regimes under which contractual arrangements are used to apportion rights corresponding to key system capacity elements to market participants.

Conceivably, that could be done using some kind of physical contract, under which the contract holder would directly instruct the facility manager to pump or generate with particular units at particular times. But that seems problematic, because:

- It would be difficult to diversify contract holding beyond the level of each unit being wholly assigned to just one market participant;
- It would be not just inefficient, but physically impossible, to have one unit generating and another pumping at the same time, if they share a common tunnel in which water must flow in a consistent direction, either from, or to, the buffer storage;[140]
- The only way the facility manager, or the contract holder, could ensure that the market cleared precisely the specified pumping/generation volume would be to submit totally inflexible offers, which does not seem desirable from any perspective; and
- The facility manager (and/or possibly contract holder) would then have to work with the system operator to match actual unit commitment/dispatch as closely as possible to cleared quantities.

While it is possible to imagine a chain of contractual arrangements to achieve all that, the whole process seems unnecessarily complex. Significant transaction costs and inter-party tensions seem likely to arise at several points, all in an attempt to achieve a precise physical matching that really does not need to be achieved, and ultimately will not be achieved, not least because machines and systems are not 100% reliable. Thus, we have passed over that possibility for essentially the same reasons that physical contracting arrangements were rejected in favour of financial contracting around a spot market in the original WEMS market design: A choice that has since been endorsed by virtually universal acceptance around the globe.

Accordingly, the key mechanisms employed in all of the regimes discussed here are not physical contracts, but financial contracts of types that have already been discussed as possible forms of hedging in Sections 5.3.1 and 5.4.5, above.[141] Specifically:

- The “virtual capacity” regimes discussed in Section 6.4.2 all employ variations on the “virtual system” slice contracts of Section 5.4.5; while
- The hybrid “energy backed capacity” regime discussed in Section 6.5.2 employs the “tank option” contracts of Section 5.4.6.

---

[139] Such arrangements do occur overseas, typically when a dam straddles a national border, and the option could be added to our list. But it would seem particularly inappropriate in a situation where all units are likely to share a common penstock tunnel in which flows must either go in one direction or the other.

[140] We understand that some facilities may allow circulating flows in order to maximise reactive power production, at the cost of losing some “active” power. That is probably not attractive in the New Zealand context but, in any case, it would be an ancillary service market issue, not directly related to the (active) power market interactions we are discussing here.

[141] For a general discussion of regimes of this type, see the paper by Barroso et al, referenced there.

6.4.2. Virtual capacity rights

Proportionally sliced system rights

Description

Participants would be assigned proportional rights to control "dispatch" of all three capacity elements, namely pumping capacity, storage capacity, and generation capacity. So, a participant with a 10% slice could decide what they wanted to do with their 10% of the generation capacity, their 10% of storage capacity, and their 10% of pumping capacity, in each market interval. More exactly, they would specify:

- How much of their 10% of pumping/generation capacity they wanted to pump or generate; with
- The net effect to be reflected as a change in their stored water level; which
- They would be responsible to keep between 0% and 10% of the total storage capacity.

All participants would submit their requests/instructions to the facility operator, who would come up with a physically feasible schedule, and/or market offer/bid, after:

- Netting off any conflicting requests (where one party wants to pump, and another to generate); and
- Optimising unit commitment/loading schedules across the day.

There are at least three basic variations on this regime, though:

- The original version of this proposal, developed by the BPA for the Columbia River system in the USA, has the participants requesting a physical dispatch schedule, and the facility controller responsible to come up with an optimal feasible schedule that approximately matches aggregate participant requests. [142]
- The variation proposed by Hunt and Read for Tasmania has the participants exercising their rights to call "swing options", thus incentivising the facility controller to come up with an offer schedule intended to produce a market dispatch schedule that approximately matches the aggregate of the obligations implied by those participant option exercises. [143]
- The variation discussed here goes further, by having participants not just "exercise their rights" to specify a quantity they wish to call, but express bid/offer curves that will indirectly determine that quantity. The facility controller then has to come up with a market offer schedule that approximately matches the aggregate of those participant offers and bids, while respecting the facility's physical capabilities.[144]

Discussion

The regime developed by BPA leaves the central authority actually controlling the catchment to find the optimal feasible schedule that approximately matches aggregate participant "instructions". The situation faced here is a bit more complicated, though, because it is the market, not the catchment manager, that ultimately determines dispatch, and dispatch profitability. So, it would be incompatible with that paradigm for a participant holding capacity rights to the pumped storage capacity to believe that they, or the facility manager, had an absolute right to determine dispatch quantities. We suggest that it would be more consistent to think in terms of participants specifying "dispatch offer/bid curves", which the controller might rationalise, and pass through, in aggregated form.

---

[142] See Section II-C in Barroso et al, above.

[143] See Section III-B of Barroso et al, above.

[144] This option was also recognised as a possible variant of the Tasmanian proposal, too. Arguably, the offer management arrangement between the participants and the facility manager could be described as more akin to "brokerage" or "agency" than a "financial contract", but the terminology should not be seen as the critical issue in assessing the concept.

"Rationalising" would include resolving situations where one participant wants to generate at a price lower than that at which another participant wants to pump.[145] Such conflicting "instructions" imply that a loss-free deal can be done much more efficiently by simply transferring rights to utilise potential energy stored in the pumped storage facility.[146]

The participant offer proposals will also presumably assume whatever conversion efficiency is implied by their contractual models, and may need to be adjusted to reflect the real conversion efficiency curve, as applied by the pumped storage facility manager to the aggregated offers. Depending on market arrangements, the facility manager will also have to decide, as real time approaches, whether it will actually be pumping or generating in the upcoming market interval(s), and either make an aggregate generation offer, or an aggregate pumping bid, tailored to produce a physically feasible unit loading schedule.

Ultimately, though, the market clearing process will determine how much of the aggregate offer is accepted, and the facility manager will ultimately be paid for what is generated, and pay for what is pumped, just like any other participant.

We could imagine detailed rules being developed to apportion the proceeds received from the aggregate cleared/dispatched amongst participants by aligning their offer/bid curves with physical reality, in accordance with their contracts. An exact matching is not actually necessary, though. The capacity contract holders could just be paid out in accordance with the simplified representation of "capacity" in their contracts, with the quantity "called" probably determined by the intersection of their own proposed offer curve and the market price. Either way, participants would:

- Be paid for the generation deemed to correspond to their accepted generation steps, at the market price; or
- Pay for the aggregate pumping deemed to correspond to their accepted pumping steps, at the market price; and
- Have their energy storage position adjusted accordingly.

The key point is that each participant would adjust their offers, over time, to manage their own storage slice, as if the capacity described by their contract was a physical reservoir under their direct control. Thus, they would each determine their own MVS estimates. But note that, if those MVS estimates diverge, profitable trading opportunities will emerge to exchange energy stored in the pumped storage facility.[147] Thus, if trading is permitted, stored energy can be expected to change hands until a consensus MVS emerges, albeit with some participants who hold more extreme views having to withdraw from further trading because their proportional storage slices are either full or empty.[148]

The ability of the facility manager to meet participant expectations depends on the mathematical model those participants have, in their contracts, of the system "slice" under their control. Contracts should probably assume constant conversion efficiency, for simplicity, and because the real marginal efficiency, on the hour, will depend on the aggregate volume cleared for all participants, irrespective of how many steps might be accepted in any individual's offer curve.

---

145 Simply netting off the offers could be done in SPD, or in SPD data processing, but that would not achieve a value-enhancing loss-free transfer of stored energy rights.

146 Appendix C discuss the difference between trading "water" and trading potential energy, or more exactly financial rights corresponding to stored potential energy, as discussed here. The possibility of also transferring storage capacity rights will be discussed under a later option.

147 If they diverge by more than RTLF, this will become evident in the form of conflicting pumping/generating bids, as above. But if active trading occurs, participants' EMVSE estimates should align quite closely, making such conflicting bids unlikely.

148 A likely reason for any persistent divergence would be that one participant found their storage strategy more constrained by their storage capacity allocation than another. So, the possibility of also transferring storage capacity rights seems important, in order to give those who find that they want to store more, the ability to do so. See discussion of the "deconstructed" virtual capacity regime, below.

- If the contracts assume that pumping/generation capacity always operates at maximum efficiency, the facility manager will only sometimes meet expectations, but often fall below them, leaving it in deficit if it is expected to pay out to participants on the assumption that everything was dispatched, and produced, exactly as "instructed".
- But, if modest enough assumptions are made about system performance, the facility manager should be able to meet all expectations, and have something to spare, either for its own profit, or to be distributed back to participants as some kind of dividend.

In principle, we would recommend allowing the facility manager to retain any excess, because that would strongly incentivise the manager to maximise efficiency. And it would not give them any market power, because the residual involved would be a very small proportion of the capacity being managed. It would expose them to risk, though, and that may not be acceptable, depending on their financial structure and ownership status.

Pros

- Retains security of physical management, and efficiency of having a single party actually manage the facility;
- Allows the facility manager to interact with the market in an essentially traditional manner;
- Allows multiple parties to participate in setting pumping/generation strategy, and consensus MVS;
- Allows participants to manage their system slices in an essentially traditional manner;
- Ensures that the facility will be utilised to respond to supply/demand cycles and volatility over any planning horizons in which one or more participant believes that to be beneficial;
- Should be able to deal effectively with the market power issue, unless some parties obtain excessively large slices;
- Provides an initial and ongoing market valuation of the facility, if slices are tradeable; and
- Supports natural development of hedge markets, by strengthening the ability of participants to offer hedges.

Cons

- Innovative;
- May seem "artificial" to some;
- Some extra costs would be involved in interacting with multiple parties;
- Some large participants might obtain excessively large slices;
- Some small participants may not want to bother giving much consideration to their MVS estimates, or offer strategy;
- Different participants may want/need different ratios of storage capacity to pumping/generation capacity, rather than all accepting proportional slices; and
- Compartmentalisation of capacity, in the sense that a party holding capacity rights might not wish to fully utilise them at times when other parties might have a better use for them.

Mitigation

- Rules could be set to limit undue accumulation of "slice" rights, but note that, if the intent is to limit market power, it is really the overall market position of each party that needs to be considered, including whatever other storage assets they may control, or contract for.
- Aggregating agents could act on behalf of, say, a collection of smaller wind /solar operators who might not have the resources or experience to manage a system slice including storage capacity.
- Rather than making all "slices" proportional, storage capacity and pumping/generation capacity could be assigned separately, and hence in different proportions, as in the "deconstructed" capacity right regime discussed in the next sub-section.

Deconstructed capacity rights

One problem with the "slicing" approach discussed in the previous sub-section is that each participant will want/need a different ratio of storage capacity to pumping/generation capacity, and possibly a different balance between pumping and generation capacity, too. Thus, a major hydro generator with significant annual storage capacity may mainly be interested in topping up that capacity with inter-annual storage to cope with wet/dry year fluctuations. Depending on their particular inflow patterns, they may, or may not, value the ability to quickly build storage up. Conversely, solar/wind generators should have a strong interest in complementary short/mid-term pump/generation flexibility, but perhaps much less interest in inter-annual storage. So, the regime discussed here "deconstructs" the slices discussed above to allow pumping/generation and storage capacity to be allocated in differing proportions to multiple participants.

Description

This regime would operate exactly like the proportionally sliced version discussed above. As discussed there, participants could propose dispatch quantities, or exercise option quantities. But, if they specify bid/offer curves, as above, there is really no operational difference between the proposals. It is just that each participant would be effectively controlling a system with a different ratio of storage to pumping/generation capacity.

Discussion

This proposal has similar properties to the sliced approach, but a potentially important pro-competitive advantage in that capacity allocations could be more precisely tailored to complement each participant's other capacity, thus providing each with a more balanced portfolio, and hence enhancing competition. It would also allow some capacity components to be traded separately, to improve and maintain alignment, over time. For example, a participant might want to purchase more storage capacity, if they find that they consistently want to utilise a greater storage range, as discussed above.

We would not suggest complete "deconstruction", though. What we are suggesting is the creation of plausible virtual systems, with linked generation/storage /pumping capacity, so that each such virtual system can be managed by a single participant, as with the slices in the previous section. If a participant owned the rights to pure pumping capacity, they could only use it by pumping water up into storage capacity, the rights for which are owned by other parties, and so on. So, it may be that storage capacity, and pumping/generation capacity, are traded separately, but perhaps not pumping and generation capacity separately. Or, it may be more a case of allowing linkages to be specified in any trading environment, as has been done in fisheries quota auctions, for example.

And it is not clear that linkages actually do need to be specified, though, because participants will not want to own rights corresponding to unviable virtual systems, e.g. with pumping but no generation capacity. But limits might need to be set on the ability of any party to acquire too great a proportion of the rights to bottleneck capacity elements, perhaps giving them too much market power in the context of interactions between potentially modest numbers of participants involved in the pumped storage arrangements. More broadly, measures might be required to monitor, and perhaps constrain the overall positions of major market participants in the national market for storage and/or generation capacity.[149]

[149] We will not attempt to develop precise guidelines here, but note that, since withdrawal of thermal "storage" capacity will leave some large hydro operators holding a significantly higher proportion of the national storage capacity than they do at present, acquisitions that redress that imbalance should generally be seen as benign,

Pros (relative to sliced approach)

- Greater flexibility;
- Better utilisation of capacity elements such as storage, as participants who do not wish to use so much give up their excess to those who wish to use more (but see further discussion of compartmentalisation below); and
- Better balance, allowing a wider range of participants to acquire balanced asset portfolios, thus strengthening competition in the wider market.

Cons (relative to sliced approach)

- Slightly greater complexity; and
- Greater potential for inappropriately concentrated control of bottleneck resources.

Mitigation

- Education;
- Provision for aggregators, as above; and
- Rules limiting undue accumulation, as above.

Dealing with key issues in virtual capacity regimes

Because the processes involved in implementing these virtual capacity proposals are less conventional than for the unified organisational structures discussed in Section 6.3, it is less clear how some of the key issues identified in Section 6.2 should be dealt with. Thus, we provide a rather fuller discussion of those issues, in this case, then go on to discuss two extra issues that only arise in this diversified decision-making environment.

Investment incentives

These diversified proposals would not, of themselves, entirely remove the spectre of a future government making an investment that the sector considered to be inappropriate. But they would at least provide mechanisms to both measure and mitigate the extent of any mismatch, because:

- The actual operation of the facility, including water value assessment, would be entirely market driven, in an environment that could arguably be more competitive than the status quo, even after thermal generation exits; and
- That operation would also imply a market valuation for the underlying capacity components, so:
    - If the market valuation adds up to exceed the investment cost, the investment would be demonstrably economic, and the government's role merely seen to be smoothing the development of a facility that is collectively economic, but possibly of too large a scale to be contemplated by any particular participant, or left under the control of one participant; or
    - If the market valuation adds up to something less than the investment cost, at least the extent of the discrepancy would be clarified, providing a basis for decision-making with respect to the desirability of incurring those extra costs on behalf of the nation, perhaps in order to achieve policy objectives other than strict electricity sector economics and, of course, who should pay that extra cost.

Thus, we believe this regime would reduce the likelihood of inappropriate signalling with respect to the value of new investments. We also suggest that, because the regime seems highly compatible with other market arrangements, it promises to be more durable, and hence provide more credible signals to investors, than regimes relying on rules that may be "fixed" in the short term, but subject to potentially arbitrary change in future. The comparison with the 'model-based' regime of Section 6.3.5 seems less clear, though.

Hedge market interactions

The interaction between the pumped storage hydro facility and the hedge market would be quite different, under any of these diversified arrangements, than under the unified arrangements discussed earlier, because:

- These arrangements, of themselves, involve creating “financial contracts”, albeit perhaps of rather exotic types, that would provide the holders with hedging very similar to that available from physical investment.
- The existence and allocation of these capacity rights would at least partly meet the demand from market participants wanting to buy hedges, but the change in net hedging activity, relative to the status quo is unclear, since market volatility is expected to increase as thermal capacity is replaced by intermittent renewable capacity, thus increasing hedging requirements.
- But, under these regimes, it is not the pumped storage facility manager, but the holders of the capacity rights who would be in essentially the same position as the owners of equivalent physical assets to sell, back, and trade more conventional hedges in the relevant markets.

Technological, ancillary service, and system interactions

These diversified regimes still internalise decision-making with respect to unit commitment, and loading levels, and leave a single facility manager interacting with the market, and system operator, to achieve acceptable unit commitment and dispatch.

But the delegation of decision-making with respect to pumping/generation levels to multiple parties who each have only limited market power, raises an issue, in that none of those parties will be in a position to appreciate how the aggregation of their individual offers will actually impact the marginal efficiency of pumping/generation, the market supply/demand balance, or prices.

So, many participants could suddenly decide that prices were high enough to justify using their capacity share to generate, say, at the same time. But, if they insisted on that dispatch level, that could cause the market situation to suddenly switch from one of moderately short supply to one of significant excess supply, and possibly even spill, at prices that might make generation unattractive.

This is one reason why the suggested implementation does not allow right holders to “insist” on a particular dispatch level, though. If, instead, each right holder merely proposes an offer curve for their capacity, the resultant aggregate curve, after adjustment for unit efficiencies, will, if anything, be steeper than the curve that might be proposed by a facility manager.[150] As a result, abrupt switches in dispatch levels or market prices seem less likely under this regime than under a unified management regime. The issue then comes back to the way contracts are set up and settled between the participants and facility manager.

We might ask how the quantity actually cleared in the market, or dispatched, is to be distributed between the right holders. As proposed, though, the quantity “called” is taken to be the quantity indicated by the participant offer curve, so there is no direct connection to the dispatched quantity. As discussed earlier, we suggest that contract parameters should be set so as to ensure that the facility manager is at little risk of being unable to meet participant expectations. And we see no reason why participants would not soon learn to submit sufficiently nuanced offers to ensure that they definitely buy/sell an amount they are comfortable with, given the prices emerging from each market interval, and are more-or-less indifferent with respect to marginal sales that might, or might not, be cleared with the facility manager, through this process.

The same approach might be taken to interactions between capacity right holders and the ancillary service market. In principle, each participant could be allocated responsibility for decision-making about

[150] To achieve this, allowance might need to be made for an increased number of steps, or perhaps an offer curve for each machine.

the participation of its share of the facility's ancillary services capacity, just as for the other capacities discussed above. In which case they might submit combined energy/reserve bids/offers to the facility manager, who would then be responsible for rationalising any conflicts, and conforming aggregate offers with actual energy/reserve capacity, before passing them through.

Participants might not have a strong interest in the ancillary service market, though, or they might already have too strong an interest in some cases. Either way, the facility manager could be given responsibility to manage, and offer, all ancillary service capacity, perhaps distributing the profits from doing so back to the participants whose capacity was utilised. We wonder if that complexity is warranted, though.

Because ancillary service prices include the opportunity cost of not using capacity to pump/generate, the returns on each participant's capacity should be at least as great as they would have been from submitting its energy offer, without modification to allow ancillary service market participation. Accordingly, the facility controller could just pay out right holders on the basis of the simplified facility performance representation in their contracts, and retain any extra profit from ancillary service sales, along with the efficiency gains discussed earlier, as part of its funding, and as an incentive to maximise efficient utilisation in both markets.[151]

Host catchment interactions

Finally, Chapter 4 has argued that host system interactions may not actually be problematic, for many possible pumped storage developments. And, of course, the least problematic developments are the ones most likely to proceed. Still, if host system interactions are expected to be problematic for any particular development, they will be just as problematic under these diversified proposals as for the unified options discussed earlier, with the slight added complication that the host system manager would probably also be a capacity-owning participant, and thus have potentially different incentives from other participants.

Appendix B discusses a number of mechanisms that could be used to deal with problematic host system interactions, and all are potentially applicable in this situation, because the primary interaction would still between a single facility manager and the host system manager. Solutions involving integrating management of the pumped storage facility with that of the host catchment, either organisationally or contractually, become more complicated, though. In that case, it would be logical to assign participants a "slice" of the combined system, thus giving them responsibility for bidding in a slice of the host system capacity, along with that of the pumped storage facility.

We again note that parties managing pumping capacity must not only have cheap electricity, but also water available to pump. So, if host system water availability may be an issue, it would be logical for shares of pumping capacity to at least be linked to shares of water available in the buffer storage. Ideally, though, if arrangements with the host system allow for some degree of control over flows entering the buffer storage, pumping capacity should be linked to shares of upstream inflow/storage/release capacity, in a fully integrated model. Likewise, shares of generation capacity should logically be linked to shares of downstream storage/generation/flow capacity. So, the "deconstructed" regime might still offer participants shares of only three system components:

- Pumping capacity, matched to proportional shares of upstream inflow/storage/release capacity;
- Storage capacity in the upper reservoir of the pumped storage facility; and
- Generation capacity, matched to proportional shares of downstream tributary/storage/flow capacity.[152]

---

[151] Unless that gives the facility manager itself too much market power in the ancillary service market.

[152] Unpumped upstream flows obviously become competing downstream flows and, while our discussion has ignored tributary flows, these should logically be apportioned and allocated, too, as upstream sources of pumpable water, and as competitors for downstream flow capacity.

Of course, participants could avoid dealing with all that complication if they were allowed to buy pumped storage hydro capacity shares without also buying capacity shares in the host system. But the implication of that would be that other participants would be purchasing proportionately greater rights to capacity in the host system, thus potentially concentrating control of the host system within a small group of participants, who would then have significantly different incentives from participants controlling only capacity in the pumped storage hydro facility.[153]

It remains to be seen if managing a slice of the entire integrated catchment would be thought desirable by any party, but we note that an experienced hydro system operator would only be adding maybe three virtual reservoirs to their existing physical portfolio and, while managing a pumping bid may be new to them, it does not seem conceptually difficult.[154]

At the other extreme, small solar/wind participants, who arguably most need short-medium term (i.e., up to seasonal) backup may have little appetite for mastering the intricacies of what, for each of them, would be a very small slice of the integrated system. Hence the suggestion, above, that (one or more) aggregator/agents might be established, or expected to emerge, to manage the aggregated capacity shares of multiple participants of that type. Alternatively, such participants may wish to buy simpler "tank options" that they, or an agent, could call at short notice, such as would be available under the hybrid and mixed regimes discussed in Section 6.5 below, or perhaps to buy more conventional call options.

None of this may be required, though. There will be an issue if host system interactions need to be managed on a regular basis, for a particular development, but such tightly coupled developments seem least likely to proceed. In many cases, it may suffice to simply scale back pumping capacity allocations under rare drought conditions, or generation capacity allocations, under rare flood conditions.

Issues specific to virtual capacity regimes

Compartmentalisation

We have argued that this kind of regime puts right holders in essentially the same position they would be in if they were able to invest in a pumped storage hydro facility of their own. In fact, the position of each participant should be much better than they could have achieved with any physical pumped storage option likely to be available to them, because the deconstructed version of the proposal allows them to customise their virtual facility quite feely to match their requirements. It should also be much cheaper. That does not mean that it achieves the best possible utilisation of the facility's capabilities, though.

As discussed earlier, a large pumped storage hydro facility could play many diverse roles and, importantly, it can play different roles at different times, and/or more-or-less simultaneously. Thus, there is nothing physically preventing it from cycling daily, and responding to short term contingencies, while also gradually building up storage to cover seasonal or inter-annual supply/demand balance fluctuations. So, the question is whether diverse parties holding rights of the type discussed here would actually be able, or incentivised, to draw on that full range of capability.

At one level, if we think of the right holders as maximising expected returns, it seems that they should be so motivated. They may each want more or less long-term storage, depending on their differing views of the future, but they will still all be in the same position to respond to enhance their returns by responding to short term market signals as they arise.

153 In the limit, the host system might be controlled by just a single participant, thus re-creating, to some extent, the problem that slicing the combined system was designed to avoid.

154 By way of comparison, we understand that each participant in the Columbia River system scheme, from which this slicing proposal was ultimately derived, must determine their desired dispatch for their slice of the entire system, involving a great many more reservoirs than we are discussing here.

The same is still theoretically true if the right holders are market participants, with a portfolio of other assets and sectoral commitments. But, while those participants will be interested in maximising expected returns, they will also be focussed on the kind of business they are in, and on mitigating the kind of risks, whether short or long term, arising from that business. So, each will adopt a strategy designed to trade off some loss in expected returns, with some reduction in their risk profile.

If the right holders all have similar interests and portfolios, as balanced broad-spectrum gentailers, or stand-alone wind generators, or whatever, they will all face similar risk profiles, and will all want to utilise the same capacity elements at much the same time. They may not want to utilise those elements to quite the same extent, and so full utilisation will only occur if all agree it to be desirable. That will moderate the aggregate facility management strategy, but such moderation is supposed to be a strength of diversity in decision-making. And the deconstructed regime above allows them to trade capacity elements to form virtual systems that are more congenial to their interests and convictions, if they consistently find themselves valuing those elements more highly than others do.

The situation is a bit different if the right holders have quite diverse interests, though. As an extreme example, consider a solar generator wishing to act as a stand-alone retailer to domestic consumers. If they just want to cover their own risks:

- They would have little interest in inter-annual storage, but would want a moderate volume of long-term storage capacity within which to build up storage over summer, for winter use; and
- They would also always have an interest in pumping/generation capacity, backed by a small volume of storage capacity, to deal with day/night cycling and short-term irradiance/load volatility; but
- They would need less cycling capacity in winter than in summer.

Conversely, a hydro generator may be primarily interested in the storage capacity required to provide dry year winter peak cover, with the pumping capacity to fill it, and the generation capacity to utilise it, when the time comes. But there is no physical reason why the generation capacity reserved for that purpose could not also be used for daily cycling or wind/solar shortfall support, say, during summer/night periods, and even in winter, under most hydrological conditions.

So, the contracts should allow for reserved capacity to be sculpted over time, and/or released/traded on, say, a month-ahead basis. But note, again, that these are not “dispatch rights”, and the facility manager is not obliged to align dispatch precisely with aggregate contract calls. It is free to operate as it thinks best in the spot market, with the usual incentives to align behaviour approximately with the volumes implied by the exercised contracts. Depending on its financial structure, and incentives, it might be prepared to take some risk by double-booking capacity in periods where simultaneous calls seem very unlikely. Or, it could offer some capacity on non-exclusive second priority contracts at a discount.

It should also be remembered, though, that these contracts are not the only “hedging” option available to our hypothetical solar-based retailer. Ultimately, spot market transactions will cover any mis-match, but it could also buy hedges from other participants in conventional markets. It could negotiate interruptibility/backup arrangements with its customers and/or invest in batteries to cover its daily cycling requirements. Or it could merge its solar capacity into the broader portfolio of a more diversified gentailer. The pumped storage facility, and the capacity rights discussed above, would greatly improve its position, relative to the option of having to build its own storage facility, or survive in the spot market without one. But they do not purport to completely de-risk its entire operation.

Allocation of capacity rights

If capacity rights are to be created, they must also be allocated, and the most obvious way to do that would be to auction rights initially, and then allow trading on a regular or ad hoc basis, as has been assumed above. Theoretically an auction of long-term rights could be used to determine the market’s view of the economic value of any proposed development, before any commitment is made to construction. And successive auctions might also be desirable to ensure that no one party has the opportunity to buy up rights to too much bottleneck capacity before the market is fully informed. Realistically, though, the innovative nature of the rights involved means that participants may not find them easy to value, at first. And the rights would not actually have any day-to-day value until the facility was built, and filled to a reasonable level.

Even then, a facility that is optimally sized for its lifetime role is likely to seem oversized in its early years, meaning that capacity rights might not reach a realistic long-term value for many years after commissioning. Also, it could be considered desirable to reserve a pool of system capacity rights to support development of a market in simpler financial instruments, as discussed in Section 6.5 below, and/or entry by technologies that are, as yet, not well represented in the New Zealand electricity sector. So, the structuring of the allocation process is a topic that would need to be examined more carefully, if any specific development were thought likely to be economic.

Assessment of virtual capacity regimes

Of the two regimes above, the "deconstructed" version seems best suited to New Zealand circumstances, particularly given the wide variety of participants potentially wanting to acquire extra storage capacity of some kind, in the emerging market situation. The discussion above suggests potential drawbacks in two major areas:

- We have not attempted to cost any of these regimes, but introducing an extra layer of interactions between the facility manager and market participants must increase the complexity and direct cost of both implementation/development and ongoing operations; and
- Compartmentalising capacity in the proposed way must theoretically imply that the facility's full potential flexibility is not always fully utilised.

On the other hand, the true national cost of implementing/operating/regulating the unified regimes discussed in Section 6.3 would be greater than the direct costs might suggest, and the kind of simplified rules likely to be agreed would also mean that the facility's full potential flexibility would not always be fully utilised. Accordingly, a more detailed analysis would be required to determine the net additional cost, and possible loss of flexibility, involved.

That said, these proposals do seem capable of delivering the intended advantages of increased competition and decision-making diversity. Importantly, it seems that virtual capacity could be made available in a form that would allow a broad range of market participants to complement their physical capacity, so as to provide each with a more balanced portfolio, and hence enhance competition in the broader market. Thus, this kind of regime seems worthy of serious consideration.

6.5. Hybrid/mixed organisational arrangements

6.5.1. Introduction

We suspect that, in talking about "dispatch rights", some parties may not be thinking of any of the regimes discussed above, but perhaps thinking of a pumped storage facility primarily as a generator, complete with fuel supply, that buyers of "dispatch rights" might dispatch only in generation mode. Accordingly, this section considers a hybrid regime, designed to implement that paradigm, and then goes on to consider a "mixed" regime under which that hybrid would be developed within the overall context of a deconstructed capacity-right regime, of the type discussed above.

6.5.2. Energy-backed generation rights

This regime is built on the "tank option" concept introduced in Section 5.4.6, and may be termed a "hybrid" between "unified" and "diversified" regimes, because it envisages "unified" management of pumping to fill the collective tank with stored energy, but a "diversified" approach to utilising that energy for generation, with multiple parties filling and managing individual (virtual) "tanks".

Description

Under this regime:

- The pumped storage facility manager would act as the only decision-maker with respect to pumping strategy, the only user of pumping capacity, and the only buyer of energy for pumping;
- They would regularly build up a stockpile of potential energy, acting as a unified buyer, and perhaps incidentally setting a floor price in the electricity market, particularly in summer; but
- Individual participants would purchase rights to both generation and storage capacity, in whatever proportions they may choose; and
- They would also purchase rights to stored energy, as it becomes available; and then
- They would exercise their rights to generate, using their reserved capacity and stored energy, when they saw fit, as in the deconstructed version of the virtual capacity right regime described in Section 6.4.2 above.

Discussion

The intention here is to create a simplified capacity right regime in which right-holders do not need to involve themselves with the pumping side of the pumped storage operation. And that simplification may be particularly desirable if the upstream host system interactions involved in facilitating efficient pumping turn out to be complex. Section 5.4.6 describes this kind of arrangement as a "tank option", although noting that really there are two separate kinds of transaction here: One buying tank/generation capacity, and the other filling the tank on a more-or-less regular basis.

The discussion in that section was focused on meeting the needs of participants seeking relatively long-term hedging cover. But, as discussed in Section 4.7.2 above, a large pumped storage facility can also deal with daily cycling and/or short-term volatility while gradually building up or drawing down storage to cover seasonal or year-to-year hedge requirements. Thus, one could also envisage rights to dispatch that function, perhaps more like the "collars" discussed in Section 5.4.3, being auctioned quite independently of the tank options, and bought by participants who are mainly interested in managing short term cycles and volatility, and might otherwise invest in more battery capacity.

A similar function might be provided though, by simply adding an "autofill" service to the tank option arrangement. The participant would:

- Specify the level to which they want their tank capacity topped up, probably by a specified time each day,
- Agree to pay the daily energy price for that top-up, perhaps subject to a cap if prices are not pre-announced; and then
- Call on their stored energy to generate when they wish to; and
- Possibly make supplementary energy purchases, from time to time, especially if they have storage capacity over and above the autofill level.

Pros (relative to virtual capacity right regime)

- Simplicity, from a right holder perspective.

Cons (relative to virtual capacity right regime)

- Reduced control for right-holders;
- Reduced diversity in decision-making; and
- Potential market power issues arising from the single-buyer arrangement.

Mitigation

- Rules to limit market power of the pumping agency as a buyer in the electricity market;
- Rules to limit the markup over pumping costs when pricing stored energy; and/or

- Establishing a competitive market for energy stored in the upper reservoir.

Dealing with key issues

With respect to the "key issues" identified in Section 6.2, this regime has much in common with the virtual capacity right regimes of Section 6.4.2, but some additional issues do arise:

- The single buyer arrangement for pumping energy gives rise to concern that the buyer will exert monopsony power, by limiting its demand to keep electricity prices lower than their optimal perfectly competitive level at off-peak times. The extent of that market power is debatable because MDAG projects that the optimal price of electricity will already be very low at those times, and the pumped storage facility will not actually be the only buyer of electricity. Nevertheless, a rule could be formulated, with the general effect being to support some kind of floor price in the electricity market. That might be considered a relatively benign intervention, and possibly even beneficial.
- There is also a market power issue in that the single buyer would be a single seller in the market for stored energy. A "cost-plus" rule may be suggested, although that does mean that, if pumping is optimal, the pumped energy would be sold at a price different from, and typically lower than, its actual marginal value, as determined by the market participants seeking to buy it. And that suggests a problem with determining exactly who would be allowed to buy this under-priced energy. An alternative suggestion is to set up a competitive market for purchasing stored energy, and that seems desirable because there may also be parties wishing to sell energy already stored in their tanks. More generally, there could be competing suppliers of pumped energy to that market, as in the "mixed" regime discussed in the next sub-section.
- That implies that pumping would be a profit-making enterprise, though. The profits could be re-cycled to capacity right holders, or perhaps more widely in the sector, as a form of dividend. But we have previously suggested that the incentives of the facility manager could be improved by allowing it to retain the incremental profits made by increasing returns from generation and/or ancillary services above the level assumed in its contracts with right holders. So, the same principle could be applied here, if that was compatible with the way in which the facility manager was set up. Alternatively, that profit-making function could be assigned to some other entity, as suggested in the mixed regime discussed below.
- At a technical level, the facility manager's interactions with right holders, the market, and the system operator would proceed in very much the same way as above, with the simplification that there would be no need to interact with participants over their pumping dispatch request/bids. A potential conflict would arise if the facility manager wishes to pump at a price level for which some right-holders wish to generate, or vice versa. As above, though, the efficient way to resolve that conflict is not to pump, or to generate, but simply transfer energy stored in the upper reservoir.
- Host system interactions would still be managed between the host system manager and pumped storage facility manager, as above, with the simplification that the effect of the arrangements made on the pumping side would not need to be reflected in the rights held by market participants.
- Hedge market interactions would be simplified, inasmuch as right-holders would be holding a simpler product, and would not need to concern themselves about hedging their own future pumping demands. Whether the pumped storage manager feels the need to hedge its own pumping demand may depend on what rules might be set with respect to the price it needs to pay. But tank option holders would still be dispatching generation from their tanks in much the same way as from any physical reservoir, and that virtual capacity should still factor into their ability to support more conventional market hedges, as above.

6.5.3. Mixed capability allocation

To this point we have distinguished between the regimes in Section 6.4.2 that give participants some form of control over and/or benefits from particular system capacity elements, and the hybrid option discussed above, which only allowed and required participants to control and benefit from generation capacity backed by energy stored in the upper reservoir of the pumped storage facility. Earlier, Section

6.3 discussed the possibility that the pumped storage facility could be controlled by a single entity that would buy and sell in conventional hedge markets and/or issue non-conventional hedges.

One reason it is not obvious which of these regimes should be preferred is that different types of participants might prefer one option over the other. But it should be recognised that we do not really need to adopt a "one size fits all" solution. A large enough facility could directly or indirectly support a mixture of all these options at the same time.

One possibility would be that part of the capacity might be managed to support direct participation by the facility manager, or some other agent, in conventional hedging markets, while the remainder was managed under a policy effectively determined by participants exercising rights of various kinds. That mix might not appeal to purists because, if the centrally managed proportion is large enough, the "official" (probably model-based) policy would tend to drive, or at least provide the reference point for, the policies adopted by "independent" participants. But that kind of situation is not all that uncommon, internationally, and the discussion in Section 6.5.4 suggests that it might not be all that inappropriate, particularly as a fallback position if a sufficient volume of rights could not be sold at a reasonable price, perhaps in the early years of operation.

Here, though, we will focus on a simpler situation, involving a mix of just two of the regimes discussed above, with generation decisions fully decentralised, but determined under a mix of the virtual capacity and hybrid regimes described above.

Description

Under this regime:

- A proportion of the pumping capacity would be managed by the facility manager, or by an independent agency willing and able to take on the role of building storage up to a level capable of supporting energy-backed tank options backed by a corresponding proportion of storage/generation capacity, as in the hybrid regime discussed in Section 6.5.1 above.
- The remaining storage/pumping/generation capability would be sold to participants who are familiar with the management of complex hydro systems, and desirous of complementing their physical capacity with virtual capacity that could be managed as an integral part of that portfolio, as in Section 6.4.2 above.

Discussion

As described, the hybrid regime gave the facility manager a monopsony/monopoly role quite different from that of any other agent. In this wider context, though, that role could be assigned to another entity that would hold virtual capacity rights just like any other market participant, and manage pumping under those rights to build up stored energy for sale to market participants, at announced prices, or via auctions, or on "autofill", or under whatever arrangements might be mutually convenient. Such an agency could operate as a profit-making enterprise, and could eventually emerge spontaneously under the virtual capacity right regime above, as the proportion of smaller generators grows over time. Or, it could be deliberately established, to support entry by smaller market participants, by ensuring that their hedging needs of are met, from the start.

In this mixed environment, there is probably no need to distinguish between "virtual capacity rights" for storage and generation, and "tank option capacity" rights for storage and generation. The real difference is probably just that some parties would choose to hold their own pumping capacity rights, and manage their own pumping to fill their tanks, whereas others would rely on buying stored energy. And the stored energy they buy might have been pumped by an entity dedicated to providing that service, or be surplus to the requirements of others who may have obtained it by exercising their own pumping rights, or perhaps by buying it from another party earlier.

Pros (compared to one or other of its component regimes)

- More choice of right-holding options for market participants;
- Competition between these two modes of hedging provision;
- Competition in the markets for stored energy and for buying energy to pump; and
- Guaranteed availability of relatively simple tank options for smaller participants.

Cons (compared to either of its component regimes)

- Increased cost and complexity.

Dealing with key issues

If we think of this regime as operating in the overall context of a virtual capacity rights regime, then it would deal with the all the "key issues" identified in Section 6.2 just like that regime. But some comment may be helpful:

- If the agency creating tank options, and managing pumping to support those options was not the facility manager, then the role of the facility manager in interacting with right holders, the market, the system operator, and the host system manager would revert to that envisaged under the virtual capacity regime.
- One essential requirement is that the capacity covered by each of these arrangements must not only have access to electricity market "energy" to support pumping, but also access to pumpable water and to downstream flow headroom, if either is sometimes in short supply. We suggest that allocation in proportion to pumping/generation capacity rights would be appropriate, irrespective of whether pumping to provide stored energy to tank option holders was managed by the facility manager, or by some other agency.
- If the majority of participants strongly preferred the simpler tank option arrangements, an agency having the sole right to create those options could still have a significant degree of monopsony/monopoly power. But a competitive fringe of virtual right holders would still provide a useful discipline, and point of comparison, and there is actually no need to give any party sole rights, except perhaps in an initial establishment phase.
- Participants holding virtual capacity rights and participants holding tank options would both be managing the dispatch of those options as a complement to their physical portfolios, and could both independently engage in hedge market activities, backed by their entire physical/virtual portfolio.

6.5.4. Assessment of hybrid/mixed regimes

There is a sense in which "mixed" arrangements can always be recommended because, provided the arrangements are mutually compatible, they offer participants the choice of enjoying "the best of both worlds". The main question is generally whether the extra cost and complexity can really be justified, and that is largely a matter of scale. So:

- A small-scale pumped storage development should obviously operate under a unified management regime, probably integrated with its host system, and would not require any particular regulatory attention. Its owner would presumably be left to account for its capacity when determining its hedge market offerings.
- A development large enough to cause concern about its market power could remain under unified management, but operating under either a rules or model-based regime. A relatively simple tank-based regime should be seriously considered though, as a means of providing effective hedging to participants, while controlling the market power of the facility in the storage/generation markets. Regulatory attention might then be confined to establishing that structure, along with relatively simple rules governing the purchase of power for pumping, if that is of concern, and sale of stored energy.
- The creation of more complex "virtual capacity" rights may well be justified, though, if the virtual systems formed by those rights would be large enough to provide attractive complements to the

existing capacity of larger scale hydro operators, in particular. Depending on the proportions involved, that might suffice to eliminate any ongoing need to regulate pumping/storage activities, while still allowing some portion of the capacity to be used to support simpler tank options, as above.

One obvious objection to this mixed arrangement would be that the strategies adopted by whatever entity was made responsible for managing pumping to provide energy for sale to tank holders would inevitably influence the judgments made by independent operators in managing their portions of the capacity. But we do not see that as a fatal objection, because, even if that entity was publicly owned:

- First, the importance of the effect will depend on the proportions involved. In the limit, if we had, say a publicly owned entity managing 20% of the pumping capacity to fill the tanks of small participants while the other 80% was split between four major operators, then none of those five parties should really have any more control than any other, and the publicly owned entity would only have more influence inasmuch as it may be more transparent about its strategies and marginal water value calculations.
- Second, the existence of an entity publishing strategies and valuations consciously intended to reflect perfectly competitive market assessments and actions would provide a benchmark that many parties might welcome. There are many reasons why such a benchmark might itself be unintentionally biased, and it is not hard to imagine the responsible entity feeling pressure to intentionally bias its assessments in one direction or another. But it would at least provide a starting point for a sectoral/public debate on such matters. And it would then be up to other participants to adopt complementary strategies based on their own assessments of future conditions, and to explicitly consider, and potentially explain, why those assessments differed from the "benchmark".
- Third, the real-world choice may not be between a perfectly competitive market ideal and this mixed regime, but between this regime and a regime in which a single party unilaterally determines pumping/storage/generation/pricing/hedging policies covering 100% of capacity, and publishes assessments that drive market outcomes much more strongly than in this mixed model, without any moderation by, or cross-check against, diversified policies determined by experienced independent managers.
- Finally, we wonder if the lumpy nature of investment in a major facility might make this kind of mixed approach necessary, at least in the early years. If the facility is optimally sized for conditions in, say 2050, then it may be oversized for conditions in 2030.[155] Also, we expect that it will be several years before sufficient water has been stored to support the issuance of substantial quantities of energy-backed hedges. The financing/underwriting of that phase of operations may be challenging, but that lies beyond our present scope. The point is that the optimal balance of capacity allocated to each operational/hedging mode can be expected to evolve over time. In particular, we might reasonably expect incumbents to express limited initial demand for storage capacity that has yet to be built, and even for capacity that has been constructed, but still empty, and perhaps not fillable for several years.

[155] But not necessarily so. If the main concern is evolving from managing wet/dry year fluctuations to managing extended dunkelflaute conditions, large-scale pumping/generation capacity backed by moderate storage may gradually become more important than large scale long term "storage".

7. Conclusions

We can not make a universal recommendation with respect to operational or organisational arrangements for the wide variety of possible storage options that are being considered, or could be considered. The goal of this preliminary conceptual discussion has thus been just to canvas options and suggest the most promising directions for further development. We also appreciate that there will be many perspectives yet to be considered with respect to many of the matters we have raised. Still, we should perhaps summarise our own preliminary views on some key issues, at this stage:

First, at a high level:

- Any single large-scale development will, by intention, have a significant impact on both prices and physical/volume outcomes in the spot market, the hedge market, and the investment market.
- Care will be required to make arrangements to ensure that the intended beneficial impacts are not unduly undermined by potentially negative "second order effects".
- Such effects could arise if a commercial entity was allowed to control the facility in a way that advanced its own commercial interests, at the expense of the national interest.
- But they could equally arise as a result of blindly following rules that are well intentioned, but unable to fully reflect the nuances of the situation, or adapt to change.
- Accordingly, if a "unified" organisational approach is taken, we would favour an operating mandate based on the application of a national cost-benefit optimisation model, with parameters agreed by industry experts.
- The main attraction of rule-based alternatives seems to be that rules could be simplified in such a way as to make it clear, to all parties, that the facility was to be used only for certain purposes in certain circumstances.
- Such "clarity of purpose" may seem attractive, and there could be other policy objectives involved. But we believe the best economic use of resources, if investing in a flexible facility, will generally be to actually use that facility to capture all the benefits which can be captured by playing diverse market roles, over multiple time scales.
- We also note that optimisation modelling of the broad type suggested above would still be highly desirable when determining rule settings, and in updating them, if they are to stay attuned to changing circumstances. So, confining that analytical process to an ad hoc backroom role supporting periodic rule-making by committee does not necessarily reduce costs, and does necessarily reduce transparency.
- Thus, we consider that, while a model-based regime might not provide so much short-term certainty around physical/market behaviour, it would actually provide greater "clarity of objective", and prove to be more predictable and durable, in the long run.

Second, while a large storage facility can provide physical arbitrage between periods, it can also provide physical management of supply/demand balance volatility across all planning horizons, from real time out to multi-annual. Thus, its optimal operation should deliver national benefits by reducing the amplitude of predictable intra-day, intra-week, and intra-annual price cycles, and also reducing unpredictable price volatility across all planning horizons.

- All consumers, and market participants, can expect to benefit from the implied reduction in price variation and volatility.
- But the physical capability of the facility would also provide a basis for supporting explicit hedges that individual participants could purchase to reduce the risk inherent in their own portfolio positions, and thus ultimately lower the prices they can offer to consumers.
- The ideal form of such hedges deserves further study, but our initial thinking is that:
    - Allowing a single organisation of the size contemplated to trade conventional hedges in established markets would strongly drive prices, no matter what that party set out to achieve, and might have to be bound by complex rules, in addition to those applying in the spot market.
    - While it is true that the optimal operation of a storage facility will tend to discipline market prices to lie in a precise proportional band, it is not clear that "collars" can really represent the essence of what a storage facility should be delivering to the market.

- Arguably "Financial Storage Rights" come closer to that goal, but they would need to be defined relative to active "spot" markets in stored energy, and many parties will not want to be locked into "transferring energy" from one specific period to another.
- Thus, "virtual capacity" option contracts seem more promising, because they allow stored/storable energy to be bought and sold in any period, and hence stored for any length of time. They would involve the contract holder making decisions about when to call the pumping capacity corresponding to their options, though, as well as the generation capacity.
- Accordingly, we suggest that some participants might prefer simpler "tank option" arrangements, in which some other party takes responsibility for purchasing and storing energy, leaving them to purchase, and ultimately call on, that energy, within specified capacity limits, when they need to.

Importantly, if the "virtual capacity" concept is implemented in a way that allows contract holders to specify market offers for the corresponding pumping/generation capacity, and thus implicitly for energy storage strategy, we believe that this would operationalise a form of diversified decision-making with respect to storage management policy, under which:

- Efficient and effective unified management of the facility itself would be maintained; while
- Participants would be able to customise their holdings of pumping/generation/storage capacity, and rights to flows available from the host system, to complement their physical capacity portfolio in meeting their risk management requirements, and those of their customers; and
- Those participants could optimise utilisation of those rights, in a basically conventional manner, as an integrated adjunct to their physical portfolios.
- The more balanced generation entities so formed may have less need to trade hedging instruments, but it would be those entities, rather than the pumped storage facility, that should be trading in conventional hedge markets.
- Competition between those entities should be expected to control market power reasonably well in national spot and hedge markets; although
- Some limits might need to be place on accumulating bottleneck capacity rights.

In some cases, though, a potentially important matter of detail may have to be dealt with:

- Storage developments embedded in host systems may find their operations compromised by complex constraints, and subject to conflicting incentives with respect to the utilisation of joint resources, such as buffer stocks and upstream storage.
- Pumped storage developments embedded in host catchments should optimally be accounting for the fact that the marginal value of the water they wish to pump must somehow also reflect the value of water as a "feedstock" for the host system's production/generation process. Likewise, demand-side response options will, more-or-less by definition, be embedded in a network of processes focussed on processing raw materials to produce products, the value of which has no direct connection with electricity market economics.
- In both cases, the valuation of the primary product, or feedstock, may vary significantly due to factors not (directly) related to the pumped storage operations, and that may compromise the facility's ability to support simple dependable contributions to spot and hedging markets.
- The simplest approach to that situation would be to integrate management of the host system and storage facility, and that solution may emerge naturally, or perhaps unavoidably, particularly with respect to demand side developments.
- At the other end of the spectrum, we would not recommend simply building a facility and relying on one-to-one trading of either water or energy to coordinate operations efficiently, or trying to establish an acceptable operating agreement after substantial capital has been committed.
- But there are several types of agreement/contract options that could be entered into, prior to making major capital commitments, that should be able to deliver mutual benefits (if the project does indeed deliver a net national benefit), and could align incentives well enough to ensure a reasonably efficient operational coordination.
- If the "diversified" management/hedging regime discussed above were to be adopted, the conditions implied by those arrangements should be passed through, in some approximate form.

- There may be little or no need for such complexity, though, in cases where the host system has little ability to manage upstream storage in ways that would limit water availability for pumping in low-priced periods, or there is sufficient buffer storage between the two systems.

Ultimately, a mixed regime might prove optimal, under which some market participants hold virtual capacity rights for the entire facility, including pumping capacity, while others purchase simpler tank option capacity, with no pumping capacity component, with those “tanks” being filled with energy purchased from the facility manager, and/or virtual capacity right holders.

The viability of that mixed regime would depend on the size of the facility under consideration, and it would be premature to pass judgment on that, at this stage. But we believe the general paradigm is well enough developed, and sufficiently promising, to provide a genuine “structural” alternative to regulation of activity controlled by a unified management structure.

We understand that this “diversified” management regime might be considered “innovative”, even at an international level. So, it would and should take some time for various parties to assure themselves of its workability and viability. Some form of innovation may be required, though, to effectively deal with the new, and relatively unique, situation New Zealand now faces in the emerging 100% renewable market environment. And New Zealand can draw on experience with pioneering the implementation of broadly similar concepts before, in the early days of market development.

8. APPENDIX A: Illustrative Host System Interaction Examples

8.1. Introduction

Chapter 4 focuses on the issue of dealing with host system interactions, because we believe it may be a critical factor limiting the economic value of some pumped storage hydro proposals and, crucially for this report, the kinds of operational and organisational arrangements that need to be, or can be, considered in those cases. Since we understand that a variety of proposals may be considered by MBIE, we will not focus on any particular proposal, but briefly discuss a few simplified examples to illustrate the range of situations that could be encountered.

Our goal is not to provide a basis for calculations, but to illustrate the key considerations involved, as simply as possible. We will discuss the implications of re-sizing, re-positioning, or removing various elements in Figure 1 from Chapter 4, reproduced here for convenience.

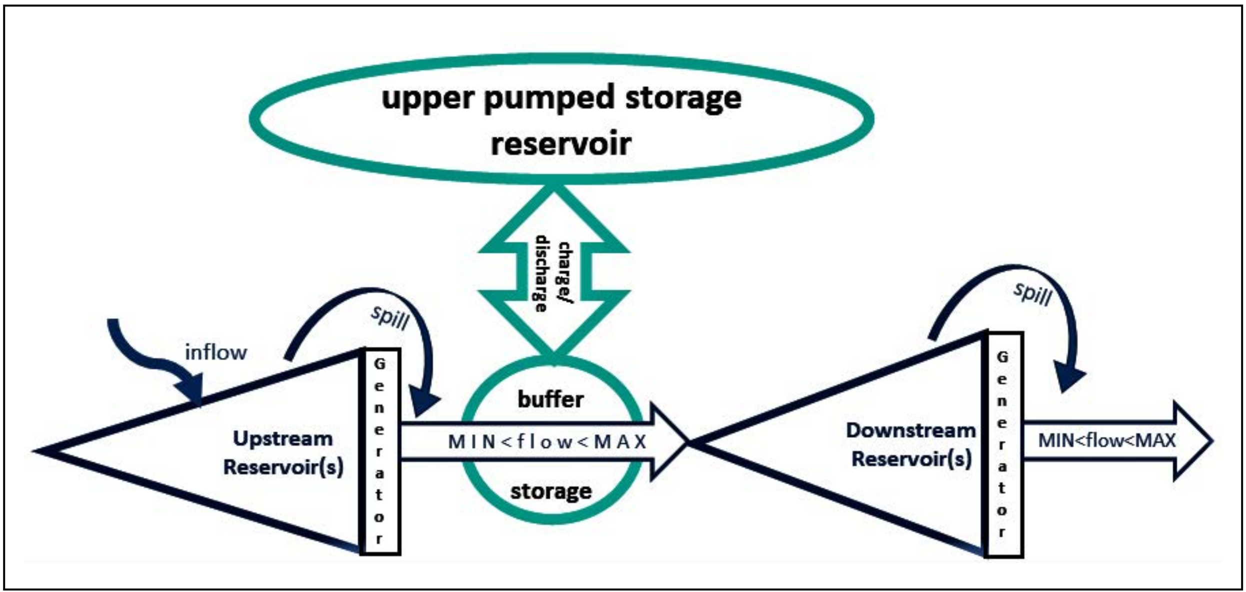

Figure 1: Pumped Storage Embedded in Host System (from Chapter 4)

Chapter 4 discussed the issues in generic terms, but we now focus on whether particular elements are small or large, in proportion to one another. Thus, our base example assumes a pumped storage facility that has:

- The same throughput and storage capacity as the upstream host system reservoir; with
- A head that is either the same as that for the power station associated with that upstream reservoir, or 10 times as great; and
- Capable of pumping at the same flow rate as when generating, with losses accounted for as extra energy requirements;[156] with

[156] In reality, pumping inefficiency will most likely be manifested by requiring the same amount of energy to achieve a lower flow rate for pumping than for generation. But this assumption simplifies our discussions here, and in Chapters 3&4. We do not believe that it changes any of the conceptual conclusions in this report.

- Pumping occurring when electricity prices are "zero", rather than "very low", so that the pumping "cost" term can be ignored entirely; and
- Releasing to the same point in the host system when generating.

As in the rest of this report, the discussion in this appendix mainly assumes a national benefit maximisation perspective, while noting where some of our conclusions about how various schemes should operate, from that perspective, might be in conflict with how they could operate, when managed in accordance with the perspectives of the pumped storage and host system owners.[157]

But the reader should be aware that this entire appendix is based on intuitive reasoning, rather than formal mathematical analysis or optimisation modelling. As such, it should be treated as advancing speculative hypotheses about the results to be expected if formal optimisation were to be applied to the various hypothetical system configurations discussed. Those hypotheses seem relevant to particular types of development, but it remains to be seen whether they have a significant impact on operational strategy, optimal configuration, valuation, or the marginal water value patterns and hedging strategy for any particular proposed development.

8.2. Developments implying strong interactions

8.2.1. Introduction

Since we can not supply a generic one-size-fits-all answer to the question of how important host system interactions may be, or what the optimal balance between various factors might be, our broad strategy here is to discuss the trade-offs arising in one base-case example, and then explore the sensitivity of those trade-offs to variation in a variety of parameters.

Confusingly, it will emerge that there are also several kinds of interaction to consider, and those different kinds of interaction may imply conclusions that seem contradictory, unless carefully distinguished. First, we have two basic kinds of interaction:

- Economic interactions, implying potentially strong national cost benefit reasons to favour modifying host system operational patterns to maximise benefits from a pumped storage facility, and potentially strong conflicts between those national economic priorities and the commercial benefits accruing directly to one party or the other.
- Physical interactions, which may align with the above, but probably more often limit those effects, because they limit the ability of the host system to actually respond to the economic priorities, no matter how strong those priorities may become.

As always, both kinds of interaction will be involved in any real situation, and a wide range of outcomes may emerge, but three generic kinds of situation can be identified:

- Situations in which the physical situation is controllable, but tight coordination of release/ generation/pumping schedules is required, in order to maximise economic benefits;
- Situations in which the physical situation is controllable, and tight coordination is either not required or easily achieved, but national cost-benefit economics imply the subordination of the normal pattern of host system operations, if the potential economic benefits of the pumped storage development are to be achieved; and
- Situations in which the physical situation is actually not controllable, so there is really no host system "operating policy" to be coordinated with, or subordinated, thus simplifying the managerial

[157] It may seem odd that such conflicting perspectives could persist, given the widely accepted understanding that (given the same assumptions about such things as resource constraints, discount rates, information, and transaction costs) perfectly competitive markets should produce the same outcomes as a national benefit optimisation. But the discrepancy can be attributed to the fact that, even if the electricity market is perfectly competitive, the hypothetical "market" for one-on-one water or energy "trading" between the pumped storage facility and host system certainly is not. Hence the discussion, in Section 4.8 and Appendix B, of alternative ways to organise that interaction.

and organisational situation, but implying that the potential economic benefits of the pumped storage development might not actually be achieved, at least to the extent ideally desired.

8.2.2. Adjunct host system

In the limit, if there was no generation capacity directly associated with any upstream or downstream infrastructure in the host system, its entire economic value would be as an adjunct to the pumped storage hydro facility. Ideally, that host system might just consist of a large reservoir with which the pumped storage hydro facility interacted directly. But that may not be physically possible. So, there could still be upstream and downstream reservoir systems, with the pumped storage facility drawing from, and discharging to, a buffer pond in a river reach somewhere in between, as in Figure 1.

The operational strategy of that upstream/ downstream "host" system would then be entirely driven by the value of ensuring that there is sufficient water available to allow maximum pumping when electricity prices are low, and the value of ensuring that there is sufficient downstream flow headroom available to allow maximum generation when electricity prices are high.

In other words, the upstream system operation would be the exact opposite of conventional hydro system operation, with releases <u>maximised</u> when electricity prices are <u>low</u>, and <u>minimised</u> when they are <u>high</u>. The downstream system might just let flows pass through, much of the time, and thus in synch with the pumped storage facility. But the reason for constructing downstream storage, if there is any, would be to boost downstream flows up to required minimum limits, by releasing more than it receives, when <u>low</u> electricity prices induce high pumping rates, and to keep downstream flows below maximum limits by releasing less than it receives when high electricity prices induce high generation rates.

That operational pattern might also be described as counter-cyclical, in the sense that it builds storage up/down at times when the pumped storage reservoir is falling/rising, but this is actually the way the last reservoir in a conventional hydro chain would be used to moderate flow variation downstream from that point. We will see, later, that this difference between the way in which pumped storage impacts upstream vs downstream operations largely determines the difficulty we should expect in aligning host system and pumped storage incentives.

The key point is that operation of a host system may be driven by the economic priorities, and physical requirements, of the pumped storage facility. Or, at the opposite end of the spectrum, the pumped storage facility could be so small, and unimportant, that the host system should just operate according to its own economic priorities and physical requirements, and largely ignore those of the pumped storage facility. Reality will generally lie between these two extremes, though, so the rest of this section explores the factors determining where the balance between these two polar opposites might actually lie for plausible system configurations.

8.2.3. Impact of head differences

If there is some upstream generation capacity in the host system, then a national cost benefit optimisation may still imply that the upstream host system's generation pattern be largely driven by the economics of pumped storage operations, depending on the relative productivity of release passing through those upstream stations, vs that in the pumped storage facility. Basically, the trade-off between maximising value derived directly from the pumped storage facility, and that derived directly from the host system, depends on two factors:

- The aggregate "head" of one system vs the other;[158] and
- The aggregate storage of one system vs the other.

[158] In this context "head" is basically just the height at which water is held, relative to some reference level, and what we are discussing here is really the head <u>difference</u> between the water level in a storage facility (e.g. in m above mean sea level), and the water level at which water is discharged from any associated power station. For a river chain, these values can be simply added to give "aggregate head", while ignoring any head "lost" over intervening river reaches.

Head is critical, because this essentially determines the potential energy stored in each unit of water, and is thus a major factor determining the marginal value of water, at any point in the system. For simplicity consider a hypothetical situation involving: [159]

- A host system with:
  - A large existing upstream storage reservoir from which water flows down through some "upstream" generating stations, each with a utilisable flow rate capacity of HSUF and a combined head of 100m;
  - No tributaries;
  - The same Minimum/Maximum flow limits (*HSFmin* and *HSFmax*) on all river reaches; and
  - No downstream storage; but
  - Some "downstream" generating stations with **twice** the utilisable flow rate capacity of those upstream (i.e. *2HSUF*), and another 100m head in aggregate head.[160]
- A proposed pumped storage facility drawing/discharging water directly from a river reach between the upstream and downstream stations, with:
  - The same (charge and discharge) flow rate capacity as each of the upstream host system stations;
  - A pumping efficiency factor (*EFF*) of 0.8[161]; and
  - A large upper storage reservoir; with
  - A 1000m head.
- But no appreciable buffer storage, and no host system river flow delays.

We will consider the impact of varying storage capacities later, but, in the absence of the pumped storage facility, a national cost-benefit optimisation, or perfectly competitive market, would clearly maximise host system flows when electricity prices are high, and minimise them when electricity prices are low. We will argue, though, that that optimal national cost-benefit strategy could radically change, once the pumped storage facility was built.

The whole purpose of such a facility is to pump significant volumes of water when electricity prices are low, and to release them when electricity prices are high. And, in the new market environment, prices are expected to be (close to) zero for much of the day, over extended periods (weeks to months) in summer.[162] So, the key issue will be water availability for pumping during those periods.

Assume (for simplicity) that all pumping can occur when electricity prices are actually zero. The marginal value of a unit of water stored in the upstream reservoir of the host system will typically be non-zero, because water can be stored in that reservoir for release in some later period. Let's say this is a large reservoir, starting at a level such that storage is not expected to reach its bounds within the next year. And let *Pannual* be the expected electricity price, across the marginal economic opportunities available in all scenarios over the next year.[163]

[159] As noted in Chapter 1, understanding operations and valuations in pumped storage hydro schemes requires us to discuss Marginal Water Value (MWV) defined in terms of $ per volume unit.

[160] This downstream capacity has (been engineered to have) enough processing capacity to generate from the combined release from the upstream host system reservoir and upper reservoir of the pumped storage facility. That is unlikely, unless the system was designed as an integrated whole, but the implications of varying that assumption will become evident below.

[161] Implying RTLF = 1/0.8 =1.25.

[162] We will ignore the possibility that it could actually be negative, under some circumstances.

[163] Noting that the best marginal opportunities available will depend on the utilisable flow capacity of the host system, and might be sooner or later, depending on the scenario considered, but the EMVSE is taken as an expected value over all those scenarios, as discussed in Section 2.2.

Then, letting *g* be the gravitational constant, and ignoring efficiency losses in generation, we can calculate the expected future marginal value, across all of those opportunities, of electricity generated by 1 cubic meter of water stored in this reservoir, and then falling 1m, as: [164]

*Vannual =3,600,000 * g * Pannual*

Since water released from this upstream host system reservoir will pass through stations with a combined head of 200m, we could say that the EMVSE of water in that reservoir would equal *200Vannual*, if we ignore the pumped storage facility. And we also know that the optimal release strategy would be to maintain minimum flows during these periods when electricity prices are less than *Pannual.* And, since the same minimum flow requirement applies right down the river, that leaves no spare water available for pumping in those periods.

But what value would an extra unit of water released from the upstream host system reservoir have, if it could be pumped from the river reach into the upper storage hydro reservoir?

- By generating or spilling when electricity prices are zero, the nation immediately loses *100Vannual* in potential future value; but
- With electricity prices at zero, pumping costs nothing (irrespective of efficiency); so
- That unit of water would then arrive in the upper reservoir at a net cost, to the nation, of *100Vannual.*

By assumption, though, this pumped storage facility has the same storage and release capacity as the host system's upstream reservoir. So, if we assume that it starts from the same relative storage level then, over the next year, water in that reservoir has marginal opportunities in exactly the same periods of each scenario, as it had when stored in the upstream host system reservoir. So, it faces the same *Vannual.* Critically, though:

- When the water is released from this reservoir, it will now fall through 1000m, thus generating electricity worth *1000Vannual*, when it is eventually run down; and then
- That same water will also pass though the downstream generating stations, where it will realise another *100Vannual* of value, just as if it had been released from the upstream reservoir in the host system, at the same times, in each scenario; so
- The gross value of power generated from that unit of water will be *1100Vannual*, on average; and
- The net gain from transferring it to the upper pumped storage reservoir would be *900Vannual*.

Clearly, then, water collected in the upstream host reservoir should be transferred to the pumped storage reservoir whenever electricity prices are low, and possibly even when they are moderately high. To be exact, pumping should occur whenever electricity prices are less than 90% of *Pannual/EFF*, if we assume, as above, that *Vannual* is the same for both reservoirs.[165]

If we start from a position of "equal headroom", though, and pursue a policy of transferring water from one reservoir to the other whenever electricity prices are low enough, headroom may not remain equal

---

[164] The units really do not really matter for this conceptual discussion, provided they are the same for all reservoirs. The key point is that *Vannual* is a measure of value per unit of stored potential energy, and we need to multiply it by the head to get the value per unit of water stored at that head level, i.e. EMVSE.

- To be exact, though, if water volume is measured in cubic meters, which weigh approximately 1000kg, the potential energy stored by raising 1 cubic meter of water by 1m is 9.8 Kilojoules, with the gravitational constant *g*=9.8, or approximately 10 Kilojoules, with *g* rounded to 10.
- The electricity sector generally prefers to define energy in kWh or MWh. though. Because there are 3600 seconds in an hour, 1 kWh equals 3,600 Kilojoules, and 1 MWh equals 3,600,000 Kilojoules So that conversion factor needs to be applied, if we want to compare MWV directly with electricity prices quoted in $/MWh.
- Reversing that conversion, the "water volume unit" that would store 1 MWh of potential energy, when raised by 1m is actually 367,347 cubic meters, with *g*=9.8; or approximately 360,000 cubic meters, with *g* rounded to 10.

[165] Technically, adopting that strategy actually changes some of the assumptions on which the argument was based. The most economic marginal opportunity available for utilisation of water in the upstream host system reservoir will generally not now be eventual release for generation in the upstream/downstream host system stations, but retention for eventual pumping into the upper reservoir of the pumped storage facility. So, its EMVSE would rise to reflect that. But that does not change the conclusion that its operating strategy should optimally be driven by the economics of pumped storage, not vv.

for long. If there were no inflows, stocks in the upper reservoir of the pumped storage facility should rise, while stocks in the upstream host system reservoir should fall, thus raising the likelihood of one becoming full, and of the other becoming empty, within a planning horizon less than the year assumed when calculating *Vannual.*

On the other hand, there will actually be inflows re-stocking the upstream host system reservoir, and none re-stocking the upper reservoir of the pumped storage facility, so pumping must continue, indefinitely. And the end result is that, so long as pumping is free, an equilibrium should be maintained under which the EMVSE values in the two reservoirs are as equal as possible. Or, more realistically, if pumping can all occur in periods when the electricity price is quite low, an equilibrium should be maintained under which the EMVSE values in the two reservoirs only differ by the highest (but still quite low) net marginal transfer cost incurred across those periods.[166]

But the key point is that, with a head difference of this magnitude, the entire host system should still be seen as largely an adjunct to the pumped storage facility, with its optimal operational strategy largely geared towards facilitating the operational strategy of the pumped storage facility, and not vice versa.

8.2.4. Impact of downstream flow limits

Note, though, that in the example above, full downstream utilisation of the released flows was only possible because we effectively assumed that the downstream stations could utilise the maximum combined simultaneous release from upstream and pumped storage reservoirs. It may well be optimal for the downstream station to have that much capacity, if developing the entire catchment from scratch. If pumped storage hydro is being added to an existing host system, though, the cost and value of expanding downstream capacity to that level would have to be considered as a design option associated with the new development.

Otherwise, if the downstream flow capacity was just designed to accommodate maximum flows from the upstream station in the host system, it would only be able to utilise that flow rate from either reservoir, but not both simultaneously. If the maximum downstream utilisable flow rate was treated as a hard limit, the priority should clearly be given to releases from the pumped storage facility, because (once downstream generation is accounted for) each unit of water released there will produce 5.5 (i.e. 1100/200) times as much energy in these critical high-priced periods as water released from the upstream host system reservoir.

The situation will not be that constrained, so long as it is possible to spill past the downstream power station. Once that starts happening, the marginal benefit of release from the upper reservoir of the pumped storage facility will fall by $1/11^{th}$ (100/1100), so it should not release any more until electricity prices rise to compensate. But the marginal benefit of release from the upstream host system reservoir will fall by 1/2 (100/200) so the balance will shift further in favour of prioritising release from the upper reservoir of the pumped storage facility over that from the upstream host system reservoir.

That prioritisation is not absolute, though, and release from the upstream host system reservoir might take priority over that from the pumped storage facility, if its EMWV is much lower than that of the upper reservoir of the pumped storage facility. And/or, if its EMWV is low enough, the upstream reservoir could keep releasing at maximum rates, even if spill is occurring downstream, and there could be spill right down the river chain, without necessarily precluding the pumped storage facility from releasing at its maximum rate, too.

Eventually, though, the combined downstream flows could reach the maximum allowable flow limit on the river. And we would then want to give absolute priority to releases from the upper pumped storage reservoir by backing off upstream host system reservoir release, because each unit of generation foregone upstream allows 10 times as much generation from the pumped storage facility. Critically, though it will be physically impossible to back off host system release once the upstream host system reservoir becomes so full that it has to spill. So, in reality, physics will prevail over economics, and

[166] Noting that, when the electricity price equals *MCEP,* the net transfer cost is: *MCEP*(1000/EFF -100).*

absolute priority will have to be given to upstream host system reservoir release, even though incremental release there may produce absolutely nothing, at either upstream or downstream host system stations, and is blocking highly productive release from the pumped storage facility.

This discussion suggests that downstream congestion could be a major issue constraining pumped storage development, under flood conditions, particularly in catchments where there is limited upstream reservoir capacity to keep flood flows within acceptable downstream limits. But it also suggests a further reason to transfer water from the upstream host system reservoir to the upper reservoir of the pumped storage facility, so as to avoid it being "trapped" in a location where it is ultimately unable to be utilised.[167]

In other words, our comparison based on the assumption that *Vannual* would be the same for both reservoirs because they had the same storage capacity, actually over-stated the value of leaving water in the upstream host system reservoir, because the "marginal opportunities" available for utilisation of water in the two reservoirs would actually not be the same. Suppose that *Vannual* for the upper reservoir of the pumped storage facility is calculated on the assumption that it would be generating in the highest-priced 10% of hours.[168] But there is only enough downstream capacity to allow full utilisation of the flows from one reservoir, or the other, and the logic above implies that pumped storage releases must be prioritised whenever possible. So, *Vannual* for the upstream host system reservoir would have to be calculated on the assumption that it would only be generating in the <u>next</u> highest-priced 10% of hours.[169]

All of this discussion serves to reinforce the same conclusion, though. Clearly, water collected in the upstream host reservoir should be transferred to the pumped storage reservoir whenever electricity prices are low, and possibly even when they are moderately high.[170] So, if the ratio of aggregate system heads is high, e.g. 1100:200 = 5.5:1 as assumed here, the entire host system should probably still be seen as largely an adjunct to the pumped storage facility, with its operational strategy geared towards facilitating the operational strategy of that facility, and not vice versa.

---

167 The above discussion might seem to imply the opposite, since it was the water in the upper pumped storage reservoir that seemed to be "trapped". But, if that water had not been in the upper reservoir of the pumped storage facility, it would still have been in the upstream host system reservoir. That incremental water would then have to be spilled from that reservoir, with its economic value completely lost. So, it is "trapped in a location where it is ultimately unable to be utilised". Conversely, while the water in the upper reservoir of the pumped storage facility is temporarily blocked from release, at least it is saved for productive use later on.

Further, it should be recognised that each extra water unit transferred to that upper reservoir is a unit that will not need to be spilled, and hence will not be blocking pumped storage release in that period. In other words, incremental units transferred could actually be released to generate from the pumped storage facility, rather than being spilled, while leaving the downstream flow situation exactly the same.

If that downstream flow would have been acceptable when fed by unproductive upstream reservoir spill, economic logic would argue that it should still be acceptable when fed by highly productive pumped storage reservoir release. So, that release should not actually be blocked. Or, if the environmental rules imply that it is blocked, the effect of that blocking will be to reduce downstream flows below the level they would have been, had the water spilled from the upstream reservoir, instead. And economic logic would argue that should only happen if benefits from that flow reduction exceed the potential benefits from increased generation in that period. Real-world environmental outcomes might not always align well with economic logic, though.

168 Noting that the hours in which generation is maximised do not count when assessing EMWV, or *Vannual*, because there is no "marginal opportunity" available there. So, the relevant price is really the cut-off price for this "top 10%" sample, in each scenario considered in calculating EMWV.

169 With the cut-off now at 20%.

170 To be exact, whenever they are less than 90% of *Pannual,* times the pumping efficiency factor, under these assumptions, where *Vannual* stays the same for both reservoirs. In reality, though, an equilibrium would be found as pumping gradually increases the EMWV in the upstream host system reservoir, and decreases the EMWV in the upper pumped storage reservoir.

Impact of storage capacity differences

From an economic perspective, the conclusion that host system operational strategy should be subordinated to the priorities of servicing an embedded pumped storage facility would be reinforced if the upstream storage capacity in the host system was significantly less than that in pumped storage facility. For any scenario, the MVSE in the upstream host reservoir will be determined by assessing the expected opportunity cost over the marginal economic opportunities available between the present and the next time that reservoir might expect to be empty or full.[171] Thus if the storage capacity of that upstream reservoir is sufficiently limited, its MVSE might be calculated over a week, at *200Vweekly* in this case, and that short term MVSE may often be near zero in the kind of circumstances envisaged by MDAG, under which electricity prices could be near zero for extended periods in summer.

By way of contrast, the MVSE in a large enough upper pumped storage reservoir could be calculated over a planning horizon of years, at *1100Vmulti*. Because there are more economic opportunities available over longer planning horizons, we typically expect *Vmulti* > *Vannual* > *Vweekly,* and the differences will often be quite large.[172] So, this gives us additional reasons to transfer water to the upper pumped storage hydro reservoir:

- First, we could characterise the short-term gains from transferring water from a reservoir supplying a station with 100m head to one with 1000m head as: *900Vweekly*
- Second, we could characterise the additional gains from transferring that water from a weekly to a multi-annual reservoir as: *1000(Vmulti – Vweekly)*
- Finally, we could characterise the further gains from transferring the feedstock for the stations further downstream from a weekly to a multi-annual reservoir as: *100(Vmulti – Vweekly)*
- So, the total net gain is now: *1100Vmulti – 200Vweekly.*

Obviously, that would imply even stronger economic incentives to prioritise supporting pumped storage operations over maximising the direct profitability of the host system. On the other hand, reducing host system storage capacity reduces the physical ability of the host system to respond to these strong economic incentives. So, there may be strong interactions, in the sense that host system conditions tightly constrain pumped storage operations, but less “need” for close coordination, because the host system has little ability to control flows in ways that maximise gains from the pumped storage facility.[173]

At the other end of the spectrum, the presence of ample upstream /downstream storage in the host system means that that system has the physical ability to support flexible benefit-maximising pumped storage operations, even without the large buffer storage discussed in Section 4.7. Increasing host system storage capacity also enables water to be held longer for use by the host system in more profitable future periods, and thus somewhat reduces the gains from transferring water to the upper pumped storage reservoir.

In summary, though, we should generalise the previous conclusions by stating that, if the pumped storage head and/or storage capacity are greater than those of the upstream host system reservoir, then

---

[171] As above, “the next time” might be sooner or later, depending on the scenario considered, but the MWV is taken as an expected value over all those scenarios, as discussed in Section 2.3.3 of our report to MDAG.

[172] It might be thought that *Vweek* could become very high if the smaller reservoir were going to run out of water within a week. But emptying a small reservoir does not create a national supply crisis, so its EMWV will never go high for that reason. The argument is actually the reverse: If there is a looming national crisis, it becomes increasingly desirable to transfer water to the upper pumped storage reservoir, where it has a much higher potential energy value, although it may become increasingly expensive, too.

On the other hand, it is true that the EMWV in the upstream host system reservoir could become very high if stocks there fall to a level where the ability to maintain minimum flows in the river reach between the host system reservoir and the buffer storage is threatened, because that can not be done by releasing from the pumped storage facility. But EMWV would not be driven by *Vweek*, in that case, but by physical feasibility. If the binding flow limit is upstream from the buffer storage, the consistent response is to adopt a minimum release strategy at the upstream reservoir, meaning that there is no more water available to be pumped up to the upper reservoir at all. But, if the binding flow limit is downstream from the buffer storage, it can equally well be met by release from the pumped storage reservoir, so there is no reason for the MWV to be higher in one reservoir and not the other.

[173] In the limit, there may be no control at all in the upstream host system catchment, so there can be no host system “policy” to influence or adjust at all. See discussion of “run-of-river” systems in Section 4.4.

the upstream host system should primarily be managed to complement the pumped storage operation, rather than the direct benefits that would normally accrue to the host system, from its own generation activities.

8.2.5. Impact of flow capacity differences

Finally, flow capacity differences are a significant factor, too. If we changed the above example so that, apart from its ability to pump, the pumped storage facility was identical to the upstream host system reservoir, in storage, flow capacity, and head, the optimal operating strategy would be more balanced. Water in the two reservoirs would now face the same future opportunities, generating at the same rate, and deliver the same downstream benefits, too.

There would still be two strong reasons to transfer water from the upstream host system reservoir to the upper pumped storage reservoir, though. If the electricity price is (close to) zero, pumping is (nearly) free, and we should do it, even though we gain nothing from the upstream reservoir release, because, by transferring water across in this way:

- We can double the usable storage capacity of the host system; and
- We can also double the utilisable release rate, and hence the MW capacity of the host system.[174]

These two factors are both important, because they both increase the EMWV of both reservoirs. Looking at the system as a whole:

- Doubling the maximum release/generation rate means that the same amount of water can be conserved and focussed on generation in the most valuable half of the periods over which it would otherwise have to be released, thus (probably significantly) increasing its value;[175] while
- Doubling the storage capacity means that the same amount of water can be distributed and balanced between the two reservoirs in such a way as to (probably significantly) increase the length of time before storage limits can be expected to constrain the optimal arbitrage strategy in either of them, and thus perhaps allowing *Vmonthly* to replace *Vweekly*, or *Vmulti* to replace *Vannual,* in determining the EMWV for both.

While this might be described as a "balanced" approach to jointly managing the two systems, the general effect is still to sacrifice the traditional management priorities of the host system, by inducing the upstream reservoir to release at times of low electricity prices, in order to optimise the joint operations of the new combined system. In fact, this is necessary, because the system is fundamentally imbalanced, in that all inflows are still entering via the upstream host system reservoir.

Accordingly, half of the flows would need to be transferred in this way, just to maintain a balanced storage position, if the potential of the combined system is to be maximised. And the best time to do that transfer is still when electricity prices are lowest, because the combined system must buy electricity from the market to make up the efficiency loss involved in the transfer, in the period when it occurs.[176] Thus, the pattern of upstream host system operations should still be exactly the opposite of what would normally be expected to maximise its own returns.[177]

The value of transfer to the upper reservoir of the pumped storage facility would be further increased, though, if that facility also had greater utilisable flow capacity than the upstream host system reservoir. If we were to double the utilisable release capacity of the upper pumped storage reservoir in the above example, transferring water would triple MW capacity, rather than doubling it, and further increase the

---

[174] With the caveat that the combined release might not be fully utilisable if the downstream stations can not process the extra flow, and there is no intervening storage to buffer it.

[175] Targeting the top 10% of prices, rather than the top 20% in the example case discussed in an earlier footnote.

[176] The transfer cost is no longer MCEP*(1000/EFF -100), but MCEP*(100/EFF -100) = 100MCEP*(1-EFF)/EFF which is positive, because EFF<1, and still proportional to MCEP.

[177] The incentives for prioritising the transfer of water across to the pumped storage facility must obviously reduce as the head of the pumped storage facility reduces relative to that of the upstream host system reservoir, but that case is discussed in the next section.

system's ability to focus release from that reservoir, but not the other one, on the periods when electricity prices were at their highest. As a result, the value of keeping water in that upper pumped storage reservoir would rise, when calculated over weekly, monthly or whatever periods, and storage in that reservoir could be more flexibly managed, too. Thus, the optimal joint strategy would be re-balanced to hold more water in this more flexible reservoir, and less in the traditional upstream host system reservoir. And that implies further increasing releases from that latter reservoir at times when pumping is cheap, because electricity prices are low.

So, we should further generalise our previous conclusions by stating that, if the pumped storage head and/or storage capacity and/or flow capacity, are greater than those of the upstream host system reservoir, then the upstream host system should primarily be managed to complement the pumped storage operation, rather than the direct benefits that would normally accrue to the host system from its own generation activities.

While physical coordination may be easily achieved in these cases, economic prioritisation is still very much an issue, and the direct economic incentives of the upstream host system manager, in terms of returns from their own generation, seem diametrically opposed to those of the pumped storage facility manager. So, if the national interest is best served by prioritising utilisation of the pumped storage facility, there would seem to be a need to find some mechanism to align the host system manager's incentives with the objective, in these cases. On the other hand, it should be recognised that some of the assumptions made above may not hold, so the next section explores cases where the host system manager's incentives may be more naturally aligned with the national interest, or have less potential for negative impact, and/or management of the two systems may be more easily decoupled.

8.3. Developments requiring weaker interactions

8.3.1. Introduction

All of the examples in the previous section involved situations in which the addition of a pumped storage hydro facility to a host system implied that optimal operation of the joint system, from a national cost benefit perspective, required prioritising the pursuit of broader benefits available from the wider joint system, particularly by transferring water into the more flexible/capable upper reservoir of the pumped storage facility, over the pursuit of potential benefits from traditional (upstream) host system operation. That prioritisation was optimal even if the pumped storage facility was identical, in most respects, to upstream capacity in the host system, and identically situated with respect to downstream capacity in that system.

This prioritisation effect looked to be potentially extreme in cases where the pumped storage facility was "larger" than the host system facilities, in some dimension, such as head, storage capacity, or utilisable flow capacity. Thus, we might expect the opposite to be true in cases where the pumped storage hydro facility is "smaller" than the host system facilities, in some dimension. But we also note some situations in which the lack of capacity of various kinds in the host system actually reduces the importance of "coordination", because it becomes physically impossible to affect pumped storage operations, for better or worse, no matter how desirable it might seem, from an economic perspective.

8.3.2. Smaller upper pumped storage reservoir capacity

First, suppose the host system was exactly as described in Section 5.4.1 above, with the pumped storage facility being as described, with a 1000m head, but with an upper reservoir that could only store enough water to sustain a daily cycling operation. Then, the joint operating strategy might still be dominated by the economics of pumped storage operations, but the balance now depends much more on details of the system configuration.

If the pumped storage facility was drawing/discharging directly from/to a river reach, with no appreciable buffer storage, as described above, the optimal joint operating strategy may not be much different from that discussed above. If there is no delay between water being released from the upstream host system reservoir and being available for pumping, then the upstream host system reservoir would still need to release at times when cheap electricity makes pumping attractive, even though electricity generation is unattractive at those times. The difference would just be that, once the water was delivered to the upper reservoir of the pumped storage facility, it would only be valued at an MVSE of *1100Vdaily* instead of *1100Vannual.* So, there would be times when the value lost by release from the upstream host system reservoir (*200Vannual, 200Vweekly*, or whatever, depending on how large that reservoir is, and how close its storage is to a limit) would be greater than *1100Vdaily.*

But Sections 8.3.6 and 8.3.9 point out that a significant natural flow delay, or modest buffer storage investment, could easily decouple the situation, at least to the extent of allowing the host system and pumped storage facility to operate independently over daily cycles.

8.3.3. Lower pumped storage flow requirement

Obviously, a pumped storage development with low flow capacity could just have low capacity all round, in which case it may be quite profitable, in its own right, but have little impact on the host system, or the electricity system. But, if a pumped storage facility has a high enough head, it could actually have a significant generation capacity, and even a significant energy storage capacity, without necessarily taking much flow from, or releasing much flow to, the host system. For example, if we doubled the head of the pumped storage facility, and halved its pumping/generation flow rate, in our original example:

- Its energy characteristics would remain the same.
- The incentive to sacrifice value by releasing water from the upstream host system reservoir so that it can be available to be pumped up to the upper pumped storage reservoir would be even stronger.
- But the actual impact on the host system's operating strategy would be halved, because the volumes involved would be halved.

So, in the limit, interactions between a very high head pumped storage facility and the host system might actually become quite simple, because the required actions are obvious, but have limited impact.

8.3.4. Lower pumped storage head

We have already noted that optimally accommodating a pumped storage hydro development with exactly the same storage capacity, flow capacity, and head as the upper reservoir of the pumped storage facility could imply a quite significant change to the host system operating strategy. The situation must obviously change, though, as the head of the pumped storage facility reduces, relative to that available in the host system.

If the host system was exactly as in our original example, and the pumped storage facility exactly the same, except with a head of 10m, instead of 1000m, then it would never be optimal to release water valued at 200*Vannual,* when electricity prices are zero, just to pump it up (admittedly for free) to the upper reservoir of the pumped storage facility where it would only be worth *110Vannual* once stored. On the other hand, with such a low head, it could be worthwhile to pump, even at times when the upstream host system station finds it profitable to operate, because:

- The per unit net cost of doing so would only be something like 12% of the profit being made by the upstream host system station generating at that time;[178] and
- The EMVSE in the upstream host system reservoir will eventually fall to a very low level (2*00Vmonthly,* 2*00Vweekly, 200Vdaily, 200Vhourly, then zero)* if it keeps filling due to high inflows, but

[178] If pumping was lossless, the penalty would be 10%, just based on the head ratio, but that rises to 12% if round-trip pumping losses are 20%.

- The EMVSE in the upper reservoir of the pumped storage facility will always be at least *110Vannual,* unless or until its storage level is also raised, by pumping.[179]
- So, the presence of the pumped storage facility does still affect the optimal operation of the host system. The difference is, though, that this low head pumped storage facility is primarily being used to augment the storage capacity of the host system. So, the operating strategy of the pumped storage facility is being subordinated to that of the host system, not vice versa.
- In the limit, if a “pumped storage facility” in that position had no head at all, it would just be an additional conventional storage, simply improving control over flows to the downstream stations.
- Or, if a low head pumped storage facility could be embedded above the upstream stations, it would be effectively controlling storage which is twice as valuable because the potential energy realised by its eventual release reflects a head of 200m, rather than just 100m. Low cost/benefit pumping/generation operations between the upper reservoir of the pumped storage facility, and the upstream host system reservoir would then mainly be driven by a desire to balance storage between them, so as to reduce spill.[180]

8.3.5. Looser flow limits

Figure 1 assumes that there could be both upper and lower flow limits on every river reach, and our discussion suggests that pumped storage operations could be constrained, and/or coordination required, if any one of those limits becomes binding. Specifically, if we take “host system flows” to mean the flows we would expect in the host system if the host system manager optimised operations in the absence of any pumped storage facility, then the issue relates to:

- The freeboard between host system flows from upstream (including any tributaries),[181] which limits pumping capacity; and
- The headroom between flows from upstream (including any tributaries), and maximum flow limits downstream of the pumped storage facility, which limits generation capacity.

So, relaxing those limits could obviously increase the value of any pumped storage development, and reduce the need for coordination, The significance of these limits, and even their existence, depends on the configuration of the pumped storage facility, though, relative to various host system elements.

If there is no downstream storage in the host system, so the pumped storage facility actually charges/discharges downstream of all host system facilities in Figure 1, then:[182]

- The upstream host system will presumably be responsible for maintaining flows between limits under the status quo.
- If the host system remains responsible for keeping its release within those limits, the pumped storage facility (with any associated buffer storage) must effectively be responsible to work within the implied freeboard/headroom restrictions.
- If the optimal host system release pattern is normally well within those limits, the pumped storage facility may normally be able to pump or generate as much as it likes, without any need to coordinate.
- The absence of any coordination agreement would presumably mean that host system release would always take priority, but the occasional loss of national benefit may be acceptable.
- We should caution, though, that if a buffer storage is built immediately below the last host system dam, there may be no river reach on which the host system would be responsible to maintain minimum flows, thus potentially allowing it to reduce off-peak release all the way to zero, rather

---

179 Because, by assumption, it can store water for at least a year and, when it does release that water, it will also pass through the 100m station downstream.

180 Such a development could be much more valuable, though, if combined with a high-head pumped storage development downstream, because it could then control release of water, through the upstream host system reservoir, to the upper reservoir of the downstream pumped storage facility where, in our example, it would ultimately be worth 1100V, not 200V.

181 Host system generation capacity does not really affect these flow limits, if it has no associated storage capacity.

182 Host system generation capacity does not really affect these flow limits, either, if it has no associated storage capacity.

than to the minimum flow levels applying in the status quo, perhaps significantly reducing flows available for pumping at times when pumping would be most valuable.

- Also, there may be cases where no maximum flow limit has yet been specified, because the host system has been deemed incapable of releasing at an unacceptably high rate, but that does not mean to say that a limit would not be imposed if a pumped storage facility threatened to substantially increase peak flows.

If there is no upstream storage in the host system, so the pumped storage facility charges/discharges upstream of all host system facilities in Figure 1, then:

- Neither the host system manager nor the pumped storage facility can affect any upstream flows, so the flow arriving at the buffer storage location will just be the natural flow;
- The pumped storage facility would presumably become responsible for maintaining flows between limits on the river reach immediately downstream from the buffer storage location; so
- There is no point in "coordinating" with respect to upstream flows, because the host system has no control over them, but the looser those limits are, the better, so far as pumped storage operations go; however
- There will be no such river reach if the pumped storage draws water from, and releases water to, an existing host system reservoir; and
- The responsibility for keeping flows within limits further downstream may become a joint issue with respect to which coordination would be required.

If there is both upstream and downstream storage in the host system, as in Figure 1, then:

- Looser limits on flows in the river reach immediately downstream from the buffer storage location would allow more flexible operation of the pumped storage facility, and reduce any need for coordination; and
- There would be no explicit limit there at all, if the downstream storage extends back as far as the buffer storage location, with no intervening river reach; but
- Looser limits on flows in the river reach immediately upstream from the buffer storage location would allow more flexible operation of the upstream host system; and
- Given the conflicting objectives discussed in Section 8.2 that should be expected to further constrain pumped storage facility operations, and increase the desirability of a coordination agreement.
- That situation would be exacerbated if the downstream storage and/or buffer storage extend back as far as the buffer storage location, with no intervening river reach, and hence no flow limits to restrain the upstream host system release pattern, as discussed above.
- The responsibility for keeping flows within limits further downstream may also become a joint issue with respect to which coordination would be required.

8.3.6. Convenient delay times

The whole situation could change dramatically with minor re-configuration of the host system. First, if there happened to be a 12-hour delay between water being released from the upstream host system reservoir and being available for pumping, with no delay between water being released from the upper pumped storage reservoir, and being available at the downstream host system reservoir, then a regular daily cycle could be maintained, in which water was:

- Released from the upstream host system reservoir during hours when electricity prices were at cyclic highs;
- Pumped up to the upper reservoir of the pumped storage facility, 12 hours later, during hours when electricity prices were at cyclic lows;
- Released from the upper reservoir of the pumped storage facility, 12 hours later again, during hours when electricity prices were again at cyclic highs; and
- Passed through the downstream host system station immediately afterwards, during hours when electricity prices were still at cyclic highs.

Such a convenient delay seems unlikely to occur naturally, and it might not help the pumped storage much in responding to random fluctuations in the supply/demand balance due to changing weather conditions. But it does illustrate the potential significance of delays, which were implicitly ignored in the previous discussion. And it thus also underlines the potential value of making a buffer storage at least large enough to create an artificial 12-hour delay in the host system flows, as in Section 8.3.9.

8.3.7. Less upstream/downstream storage

It might be thought that decreasing upstream storage makes the job of managing interactions with the host system more difficult, as operations might need to be coordinated much closer to real time. But there may be less point in trying to coordinate, because the host system itself has little control, to achieve any objective, whether that be to benefit its own operations or those of the pumped storage facility. In the limit, if the upstream catchment is operating on a pure run-of-river basis, there is no upstream "activity" to coordinate, and the pumped storage facility must just do its best to accumulate and store uncontrolled flows in its own buffer storage, to be pumped up at the best available opportunity.

It should also be recognised that, while upstream storage could, and should, theoretically be used to massage natural flows into a pattern more compatible with maximising the national benefit from pumped storage operations, the natural incentives of the host system manager are to massage flow patterns in the opposite direction. Thus, the less upstream storage there is, the less inappropriate the flow pattern arriving at the buffer storage location will be, and the less need for any coordination agreement to make that pattern more appropriate.

Similarly, lack of control in the upstream host system catchment will limit the ability of the host system manager to (inadvertently) exacerbate any downstream choking effects by adopting the normally optimal policy of releasing more when prices are high and pumped storage generation should be high, too. So, that would reduce the negative motivation for a coordination agreement to restrain that effect.

On the other hand, lack of control in the upstream host system catchment will also limit the ability of the host system manager to (deliberately) reduce any downstream choking effects by releasing less at times when prices are high and pumped storage generation should be high. So, that would also reduce the positive motivation for a coordination agreement to incentivise the upstream host system manager to exploit that ability in ways that enhance national benefit.

In the limit, if upstream flows can not be controlled, there is no point to a coordination agreement, because the host system manager will not be able to stop excess water flowing downstream at times when the pumped storage facility should ideally be generating. To the extent that flow levels lie within the capacity of upstream and downstream generation facilities, the host system manager will derive some benefit from them. But, if the combined flow reaches the allowable maximum, the uncontrolled flow will have to be prioritised, and any benefit derived from that uncontrolled flow will come at the expense of restricting much more productive releases from the pumped storage hydro facility.

Downstream storage capacity also plays a role here. A combination of low upstream and downstream storage capacity must make the whole host system catchment less controllable, and hence reduce the potential gains from "coordination" as above. So, we may conclude that a situation with little host system control could be one requiring "weaker interaction".

That conclusion may seem strange, given that Section 8.2.5 argued that the economic incentives to transfer water to the upper reservoir of the pumped storage facility will be stronger if that reservoir is large, relative to the upstream host system reservoir. The conclusions are not inconsistent, though, because the economic incentives referred to there are from a national benefit perspective. As the physical constraints bind more tightly, the economic incentives do become stronger, but they shift from being incentives to operate in more economic ways to become incentives to re-shape the system so as to be able to operate in more economic ways, in this case by increasing the extent to which upper host system flows can be controlled. But that extra control is worse than useless, from a national perspective, unless the host system manager's incentives are aligned with national benefit.

8.3.8. More downstream storage

The previous section argued that having less downstream storage could create a situation in which there was no point pursuing any coordination agreement, because the host system manager can not control flows. But that is not the same as saying that coordination would not be desirable, if it were possible.

On the other hand, if downstream storage is large enough, and flow limits loose enough on any intervening river reach, the combined outflow from the upstream host system reservoir and the pumped storage hydro facility could be held in that reservoir, to be eventually passed through at the maximum allowable flow rate over an extended period.

Still better, if there is no intervening river reach, the downstream storage effectively becomes a buffer storage for the pumped storage facility from which it can draw, even when incoming upstream flows are low, and to which it can release, even when incoming upstream flows are high.

So, larger downstream storage capacity can increase pumped storage flexibility, to the point where there may be no need for, or significant benefit from, any coordinating agreement.

8.3.9. Larger host buffer storage capacity

Finally, the "obvious solution" to all of the potential problems discussed above may seem to be having a large enough buffer storage to decouple host system and pumped storage operations. For example, with no delays, and a 12-hour buffer storage in the host system, the pumped storage facility could operate on a daily cycle, and also probably be ready to provide shorter-term dynamic response strategies across most of that cycle, by just:

- Starting with the buffer storage empty, at the beginning of the daily 12-hour period of highest electricity prices;
- Releasing water from the upper reservoir of the pumped storage facility to the buffer storage over the approximately12-hour period when electricity prices were at cyclic highs; while
- Allowing releases from the upstream host system reservoir to also build up the buffer storage over that same 12 hours (say) when that station also finds release most profitable;[183] but
- Allowing enough to flow on down to the downstream host system station to allow it to also generate at its maximum rate, over the same period; thus
- Maximising generation, and returns from all three stations, simultaneously, over the period of highest electricity prices; but then
- Minimising release from the upstream host system storage, and from the buffer storage, and hence generation from the downstream station, over the intervening hours when electricity prices are at their cyclic lows; while also
- Pumping the water accumulated in the buffer storage over the higher priced section of the daily cycle up to the upper pumped storage reservoir over that same period of lower prices.[184]

The details of strategy optimisation would obviously vary greatly from system to system, and from day to day. As discussed in Section 3.3.2, increasing buffer storage capacity can not decrease the system's physical ability to perform this kind of daily cycling role, but it may open up longer term arbitrage opportunities that over-ride it. The above example highlights some important points, though:

[183] The balance here will differ from system to system. But, if there are no tributary flows, and the upstream, downstream and pumped storage stations all have the same utilisable release capacity, then water will be arriving in the buffer storage at twice the rate required to sustain downstream generation over this period. So, half can be saved for later pumping, and half passed through, and it becomes arbitrary whether we see the half flowing downstream as being released from the upper pumped storage reservoir, as described here, or just passed through from upstream host system reservoir, with the pumped storage facility simply cycling "its own water" to and from the buffer storage.

[184] Noting that, actually, there should be a dead-band period between these phases, while electricity prices traverse the range between being high enough to justify generation, and low enough to justify pumping.

- A pumped storage facility with this configuration might operate effectively, on a daily cycling basis, without requiring any significant change to traditional operating strategies in the host system.
- Such a facility would not need to coordinate tightly with the host system because, in the absence of delays, they both have incentives to respond to the same signals in much the same timeframe.[185]
- They would not normally need to interact much over longer time frames either, because the pumped storage facility is assumed to have no long-term storage, and hence no ability to vary its operational cycle to manage events beyond the daily planning horizon, or need to forecast them.[186]
- In other words, optimality would not require the operational strategy of the host system to be subordinated to that of the pumped storage facility.

Importantly, though, none of that would be true once the storage capacity of the upper reservoir of the pumped storage facility got significantly bigger than that of the buffer storage. Arguably, the buffer storage should ideally have at least as much capacity as the upper reservoir of the pumped storage facility, if operational strategy is to be completely decoupled, as in the "effectively open" system described in Section 4.4.2. But a buffer storage large enough to at least "undo" the impact that conventional pro-cyclical upstream host system operations will have on natural flows, should serve to decouple the two systems to the point where a coordination agreement might be considered unnecessary.

As discussed in Section 4.7, though, the optimal outcome here depends critically on a careful comparison between the direct costs of building a buffer storage large enough to decouple operations, and the perhaps largely indirect costs (e.g. in terms of competitive balance) of a coordination agreement which would effectively harness upstream/downstream host system capacity to do a similar job. But complete decoupling may never be realistic.

A large upper storage reservoir would be required if the pumped storage facility is supposed to help manage longer term situations, such as dunkelflaute crises, or seasonal supply/demand imbalances, and much larger again to handle wet/dry year fluctuations. And that suggests that a correspondingly large storage capacity could also be desirable, either in the buffer storage or upstream in the host system, to accumulate net excess flows over a similarly extended period of relatively higher electricity prices (e.g. winter), ready to be pumped up in the corresponding extended periods of relatively lower electricity prices (e.g. summer).

Realistically, the potential for developing such large buffer storages may be severely limited by cost, practicality, and socio-environmental acceptability. And, if buffer storage capacity is difficult and/or expensive to develop, some form of coordinating agreement may be required, as discussed in Appendix B, and that may impact on the kind of organisational arrangements we can apply to efficiently integrate pumped storage developments into both spot and hedge markets.

Theoretically, increasing either upstream or downstream storage capacity in the host system could act as a reasonably effective substitute. But the point made by Section 4.7 is that, unless incentives are reasonably aligned, the nation may face the cost of building additional buffer storage, over and above that required to manage natural flow variations, just to counteract the impact that managing upstream storages to maximise host system benefits would have on the natural flow pattern. So, the greater the investment there has been, or may be, in upstream host system storage, the greater the potential requirement for investment in offsetting buffer storage. While our focus here is not on the economics or design of pumped storage developments, Appendix D discusses why this implies a need for caution in interpreting the recommendations of optimisation models with respect to buffer storage capacity.

[185] Ignoring variations in the daily routine due to maintenance etc.

[186] Again, ignoring the need to plan maintenance etc.

9. Appendix B: Managing Host System Interactions

9.1. Introduction

Chapter 4 discussed "host system interactions" in some detail, because it seems to us that managing these direct physical interactions may be the most difficult problem to resolve when it comes to organising and managing any "embedded" pumped storage hydro facility, and its interaction with the market, both short and long-term. Thus, in principle, if it is going to maximise national benefit:

- The volumes it can pump or generate are always strictly limited by the state of the shared buffer storage, and by upstream and downstream flow conditions in the host system;
- The price at which it should be prepared to buy electricity for pumping should not only depend on the EMWV in its own upper storage reservoir, but on the difference between that EMWV and that in the buffer storage, which will reflect marginal values throughout the host system;
- The price at which it should be prepared to sell any electricity it generates should not only depend on its own EMWV, but on the difference between that EMWV and that in the buffer storage, which will reflect marginal values throughout the host system; and
- All of these conditions thus also apply to any hedge products it might offer.

The obvious problem here, though, is that storage/flows for hydro generation in most catchments are currently managed by independent power companies, and some of the information listed above would normally be regarded as private to the host system manager. In fact, the EMWV values referred to may not even be formally calculated. So, there is an obvious temptation just to ignore these potential interactions, and hope that the incentives of all managers involved will broadly align with the national interests.

That may actually be the case, for upstream/downstream managers of conventional hydro, under the status quo. We have argued that the situation could change radically, though, if a large pumped storage facility was added into any host system, and expected to operate independently of it. It seems that the implied reversal of upstream/downstream relationships, can mean that the direct interests of the upstream host system manager, if measured solely in terms of profits from its own generation, can end up being diametrically opposed to the national interests. Conversely, even if the pumped storage facility manager sets out to maximise national benefit, it could not do so without having control over the host system, or some agreement to secure a coordinated strategy.

None of that may matter, if there is actually little that can be done to influence flows in the upstream host system, in particular, either to further the incumbent manager's interests or those of the nation. And the same might be true if downstream flow headroom/freeboard is always great enough to allow for largely decoupled operations. In other cases, though, we suggest that the viability of proposals to retrofit pumped storage into host systems already managed by incumbent generators could depend heavily on finding reasonably efficient ways to manage host system interactions. Accordingly, this appendix discusses various kinds of agreement or arrangement that might be used to achieve that goal.

Clearly, no commitment could be made to proceed with a major development without establishing some kind of agreement with the host system manager. Indeed, the construction process is likely to involve such large disruptions to normal operations that a close working arrangement would surely have to be established, and maintained over several years, at least. So, the direct costs of also establishing a longer-term agreement seem unlikely to be a major obstacle.

While an agreement that radically altered established operational patterns is not something that could, should, or hopefully would be simply imposed upon an incumbent system manager, our analysis also suggests that the host system owner would be one of the largest potential beneficiaries of any development. Conversely, a developer who proceeded to commit capital without entering into a long-term agreement could face the prospect of the host system manager actually capturing a large proportion of the gross benefit, as well.[187] So, if a development is actually economic, rational profit/benefit maximising parties should be able to reach an agreement as to how the net benefits of that development are to be allocated between them.

The nature of that agreement is another matter, though, and it could have a major impact on operational outcomes, and hence on the aggregate quantum of benefit available to the two parties, and the nation as a whole. While this section makes no recommendations, it does outline several options for consideration, and briefly comments on their potential compatibility with the various organisational options discussed in Chapter 6. Given the incentives outlined above, we think the real challenge here is to find a form of agreement that preserves, and hopefully enhances, the competitive balance of the national (and/or island) market.

9.2. Organisational integration

The most obvious way to manage host system interactions would be to merge the organisation managing the pumped storage facility with the organisation managing the host hydro system, in a single integrated structure, the operational mandate for which could just be an extension of any one of the "unified" or "diversified" options discussed in Chapter 6.

If the pumped storage facility was a relatively minor addition to the host system capacity, then it would be natural to allow the host system manager to develop and operate the facility as an integrated part of the host system. Such a small scheme might not radically alter the way in which the host system operated. But a proposal on that scale would not give rise to concerns about the ability of the market to develop or accommodate it, in the first place.

On the other hand, we understand that the reason MBIE is involved with this investigation is partly because the proposals under consideration are large enough to raise significant concerns about market dominance. And any scheme large enough to "dominate" the national market will obviously also dominate the operational economics of its particular host system. Thus, the host system could become very much an adjunct to the pumped storage hydro facility, not the other way around.

And we note that the economic "dominance" discussed here does not just relate to short/mid-term operational strategy, but to longer term development options, too. In our example, the presence of a pumped storage facility downstream makes all water captured and managed in the upstream host system catchment worth at least 5 times as much as it would have been without that development, thus radically altering incentives for enhancing upstream capture/control facilities and policies. Likewise, the presence of a pumped storage facility able to control flows to be released only at times when electricity prices are high, makes all actual or potential downstream storage/generation capacity much more valuable, too.

Thus, while we understand that MBIE may have no mandate to consider catchment-wide development plans, it seems obvious that, if we were designing the system from scratch, a national cost/benefit assessment would have to consider development of the host catchment as a whole. And it would most likely establish an integrated management structure for the catchment, too. Even in a situation where many costs are already sunk, and many constraining factors, such as river and lake-side developments are already in place, proper national cost-benefit analysis of a retrofit proposal should still really consider potential developments over the catchment as a whole. We believe that would be the natural outcome of a market-driven development, with no government involvement, too.

The most obvious way to achieve that kind of integrated approach to assessment, development, and operational strategy would be for the incumbent host system manager to develop, own and control the integrated system. Or, if the proposed development was too large for the incumbent to develop, on its own, arrangements would surely be made for the incumbent to at least be involved in planning and

---

[187] Since the host system manager could just carry on operations as normal, making the negotiating position of a pumped storage facility desirous of seeing changes to that traditional flow pattern extremely weak, once their capital was committed.

operating the pumped storage facility, and incentivised to consider complementary adjustments to the host system facilities.

In the New Zealand context, though, that arrangement could create a dominant market participant, even if the incumbent manager had no generation assets outside the host system. Section 6.3 discusses several mechanisms that could be used to control the market power of a large pumped storage facility, if it were managed by a single entity. In theory, the rule-based options could be applied to any entity, even one controlled by a commercial profit maximiser with other interests in the sector. In theory, they could also be extended to cover the entire host system.

Such rules could become very complex, though, and we wonder about the wisdom, and ultimate efficiency implications, of giving a commercial party that actually owns at least part of the system, the task of operating a system according to rules that it will inevitably believe to be sub-optimal, to a greater or lesser degree. The rules may imply that it foregoes profits that it believes it should be entitled to, but equally they may legitimise profits that it would not otherwise be entitled to. Neither circumstance seems likely to incentivise particularly efficient behaviour.

Furthermore, any incumbent host system owner would probably also own generation assets outside the host system, thus further increasing its potential dominance. And it seems unlikely the market would accept the credibility of a competing commercial party supposedly operating only part of its portfolio to meet some kind of national benefit objective, without somehow privileging the rest of its portfolio. On the other hand, the rules would become even more complex, if extended over a whole diversified portfolio.

The major alternative to applying a rule-based approach to a single entity managing the pumped storage facility, or possibly the integrated catchment, would be to leave a single entity responsible for the physical maintenance and management of the facility, or integrated system, but facilitate diversified competitive management of its operational strategy, using one of the virtual capacity right approaches proposed in Section 6.4. In principle, that regime could be applied to a facility, or system, that was part of a wider portfolio. In fact, the market is already familiar with arrangements under which multiple participants hold options that, when exercised, effectively determine the dispatch of some competing party's plant.[188] But the arrangements discussed in Section 6.4 seem more likely to be acceptable if the party actually operating the dominant facility/system has clearly defined responsibilities and incentives relating to that facility/system alone, and not confused, and possibly conflicted, with incentives arising from its wider portfolio interests.

So, we suspect that the host system would need to be separated from the incumbent's existing portfolio in order to form a stand-alone entity owning, or at least managing, the integrated system in the catchment hosting the pumped storage facility, and nothing outside of that catchment. The acceptability of that kind of restructuring would be a matter for negotiation with whatever incumbent host system owner might be involved, but it would clearly be a major issue. So, alternatives to organisational integration certainly should be explored, as discussed below.

[188] This is explicit in the case of a "swaption" but actually also implicit in any kind of hedge.

9.3. Water management agreements

In many jurisdictions parties interacting in a catchment are required to enter into a comprehensive Water Management Agreement (WMA). However, we are aware that two situations already exist, within the New Zealand hydro system, where an upstream and downstream party are supposed to operate independently, in the same catchment, without any real WMA.[189]

In principle, such arrangements must increase the risk of the downstream party, and can not produce "optimal" coordination. We understand, though, that the inefficiencies implied by those particular arrangements may have been deemed acceptable on the basis that the upstream party must eventually release whatever water they hold, and that the timing of that release should not matter too much, provided the downstream operator receives that water directly into a large enough long-term storage. It was also believed that the incentives of the two parties would roughly align, inasmuch as both would want to hold water back when prices are low, and maximise release when prices are high, thus establishing more-or-less compatible flows down the entire river chain. But it should be recognised that the pumped storage situation envisaged here would be very different, in that:

- It would involve a dynamic interchange of water, in both directions, between the pumped storage and host system operators, in a situation where coordinated timing could be quite critical to both parties, and to the national electricity system.
- The planning horizons over which the two operators are expected to optimise may be very different, with water taken from the host system possibly not being released back to that system until several years later.
- Even though the economics of the situation are likely to imply that the pumped storage facility should be given almost absolute priority with respect to river flow patterns, it is entirely dependent on flows which the host system physically controls.
- The incentives of the two operators are not at all aligned, but diametrically opposed, with the pumped storage facility manager needing the host system manager to generate (or at least release) at exactly the time when that party will least want to release, and to refrain from releasing at exactly the time when that party will most want to release.

Thus, we do not believe these existing (non-)arrangements should be seen as setting a precedent applicable to the situation of embedded pumped storage.

At a minimum, both parties will presumably want to avoid violating maximum and minimum flow/storage limits, in storages and river reaches directly affected by their pumping/generation decisions, and for which they might jointly be held responsible. So, we imagine that an agreement would at least be required with respect to the physical management of that part of the river system.

Any such agreement would obviously have upstream/downstream implications, even if the rest of the catchment was not covered by it. But it could just relate to management of extreme conditions, and would not necessarily cover aspects of system management that are not driven by physical/legal/socio-environmental constraints, but by economics.

It could potentially be extended to cover "economic" management of the whole system, or just specify rules approximately reflecting economic priorities in relation to flows to/from the buffer storage, where they both interact. But those aspects might be better covered by other trading/ contractual arrangements, as discussed below, in addition to a limited physical WMA.

[189] I.e., the Rangipo scheme, above Lake Taupo in the Waikato catchment, and the Tekapo scheme, above Lake Pukaki in the Waitaki catchment. In the latter case, we understand that there may be an agreement for the upstream party to cooperate in maintaining minimum flows, in extreme situations, but are not aware of any WMA operating when flows are at normal levels.

9.4. Water/energy trading

The concept of "water trading" may be considered controversial, if it is taken to imply that some party ultimately "owns water". Appendix C argues that what would actually be traded here are rights to the potential energy implied by water being held at various elevations, and thus not substantially different from, say, rights to chemical potential in a battery.

But the key point is that, however it is construed, we do not believe that an unstructured one-on-one "market", in which the pumped storage facility manager constantly trades or negotiates with the upstream and/or downstream host system operators over water flow rates, can be made "workably competitive". Nor do we believe that this kind of continual one-on-one interaction would fully dispel concerns over the possibility of collusion. It may actually be considered likely to increase the prospect of that happening.

The demand side of this hypothetical market could be made competitive, if the diversified arrangements discussed in Section 6.4.2 went as far as requiring pumping capacity right holders to purchase water in the buffer storage for pumping, or to trade between themselves for variations around a proportional allocation. But the supply side would then be a monopoly, able to set any price it liked, up to the level at which it essentially captured all the gains that those competing buyers, and hence the pumped storage facility, might expect to make from raising the water to the upper reservoir of the pumped storage facility. [190]

We are aware of the theoretical argument that the threat and counter-threat involved in this kind of situation could lead the parties to agree on a stable deal, which could then claim to be a "market solution". Our concern here, though, is really about what form of "deal" might work in this situation, and how well that might serve the national interest, rather than about the haggling process leading up to it. So, we turn to consider the kinds of "deal" that might be considered, while noting that it need not necessarily be constructed in terms of buying and selling "water", or even take the form of a "water" management agreement.

9.5. Contractual mechanisms

Section 6.3 discusses several kinds of contractual mechanism that could be used to diversify decision-making with respect to pumping/generation strategy, and very similar mechanisms could be proposed to manage interactions between the pumped storage facility and the host system.

In this case, physical management of specific water flows is actually the critical issue, so a physical dispatch interpretation might seem necessary. But some degree of mismatch between reality and expectations of physical water availability for pumping is not actually going to crash either the pumped storage facility or the national power system. It just means that the pumped storage facility will buy in less/more low-priced/surplus power than expected, and add less/more water to its storage than expected.

Thus, the two parties could agree to a contract that allowed the pumped storage facility manager to treat a virtual model of the host system as if it were a part of the system under its physical control, thus forming the basis of a "financial" contract giving sufficient certainty for the pumped storage facility manager to make corresponding market offers and/or incorporate that capacity into any virtual slices it

---

[190] We focus on the upstream/pumping interaction, but note that there is also a downstream interaction, in a situation where there is no physical mechanism stopping both the pumped storage facility manager and the upstream host system manager releasing simultaneously, possibly leaving the downstream manager no option but to spill any excess. At that point the downstream MWV should normally be zero, but it should logically be negative if the flows are high enough to threaten flooding downstream. And it may need to be negative before either party will reduce generation.

That may be seen as equivalent to imposing a penalty on release, and if it has a lower head, the upstream host system should be expected to respond by limiting release first, thus limiting the penalty payable by the pumped storage facility manager. It seems distinctly possible, though, that the upstream and downstream managers (who are probably the same party) could "collude" to extract a higher penalty out of the pumped storage facility manager. Indeed, the downstream manager could do that unilaterally, if negative prices can be set when there is no downstream flooding.

offers, and/or back corresponding hedges. Theoretically, that would leave the host system manager responsible for physically managing the host system, and incentivised to manage that system so as to at least deliver the performance implied by the virtual model, as and when called by the pumped storage facility manager, while potentially also benefitting from any performance delivered in excess of that standard.

The fundamental problem with using financial contracts in this context, though, is that they are normally settled against a spot price, set by a hopefully competitive market. In this case, the underlying market would be the one-on-one market discussed above. If the price for extra supply in that market is just whatever the upstream manager (who physically controls the inflow) says it is, then the issue is just whether the pumped storage facility manager is willing to pay that price.

We suspect that, if the pumped storage facility manager was responsible for securing water supply for the whole facility, they might really be more comfortable with a physical contract, specifying quantities to be delivered, perhaps contingent on a wide range of physical and market conditions, but at prices agreed in advance, and/or or linked to prices set in the more competitive electricity market.

Physical contracts could take on a wide variety of forms, but we do not think it would make sense to have multiple pumping capacity right holders competing to purchase water in the buffer storage, if the supply side was a monopoly, unless some rule or contract required that monopoly to make an acceptable amount available at a reasonable cap price. The lack of supply side competition could be dealt with by extending the virtual capacity concept all the way through the upstream host system catchment, as suggested in Section 6.4. But, of course, that would only be possible if a deal had been done to secure rights to manage the upstream host system in the first place, via some form of contract, or version of the integrated arrangements discussed above.

9.6. Energy swaps

Finally, we could extend the concept of trading "potential energy" to the point where there is no explicit reference to water, at all.

One possible deal would be for the pumped storage facility to buy the energy generated by releasing from the upstream host system reservoir at a price reflecting some compromise between the market price at the time, and the (presumably higher) long-term value of that amount of energy stored in the upper reservoir of the pumped storage facility. That arrangement would break down, though, if the desired upstream release level exceeded the generation capacity associated with the upstream host system reservoir. When prices are low, such spill may actually add almost as much economic value to the system as generation does, inasmuch as it makes water available for pumping, but there would be no incremental unit of generation to be bought from the host system manager, and hence no way to incentivise the release, under this form of "energy swap" proposal.

But note that the "water management agreement" proposed by Read for the Waitaki catchment was actually an energy swap in the electricity market, reflecting the portion of the potential energy released from the upstream reservoir (Tekapo) which was actually being (implicitly) held there on behalf of the downstream system manager.[191] The situation here is not identical, and further study would be required before making any definite proposal, but a reasonable deal might be to simply swap the energy content of the water released from the upstream host system reservoir for the same amount of energy stored in the upper reservoir of the pumped storage facility, without reference to any market price or marginal water value.[192]

[191] See Section III-C in Barroso et al, above, or:
E.G. Read *Energy Storage Management for Upstream/Downstream Reservoir Operators* Presented to EPOC Winter Worksop, Auckland, New Zealand, 2010.

[192] Except possibly a transfer fee to cover the actual cost of pumping their swapped share of the energy.

That would then allow the upstream host system manager to sell that stored energy (e.g. to the tank holders discussed in Section 5.4.6), or to manage it as part of their own portfolio, under one of the "diversified" management arrangements discussed in Section 6.4.

This seems like a very generous deal, for the host system manager, because it leaves them with energy that can be stored for much longer, and hence should be much more valuable. It also shifts the timing of generation from its downstream capacity into a future higher valued period. Perhaps it is too generous, especially in periods when the host system manager would want to release water anyway. But it also allows the pumped storage facility manager, or pumping capacity right holders to retain most of the gains from pumping.[193]

The details would need to be considered carefully, especially for cases in which various limits might bind. And we would need to resolve whether it should take the form of an option to be exercised by the host system manager, the pumped storage facility, or the pumping right capacity holders. Careful attention would need to be paid to the potential for gaming around situations in which the party holding the right to trigger that agreement might gain an extra benefit for an action they would have taken anyway, suggesting that the arrangement might perhaps be better triggered by market conditions, such as low electricity prices.

Be all that as it may, the potential advantage of this kind of swap is that it would obviate the necessity to consider any kind of "water trading". And that might also eliminate the need for constant communication between the parties. Also, if such a deal can be expressed in a clean mathematical form, it could also be incorporated into each party's optimisation models, just like the virtual facility contracts participants might hold and exercise under the diversified arrangements discussed in Section 6.3, and reflected in any pumped storage system capacity slices, or hedges, supported by those virtual system models.

[193] 90%, if pumping occurs at times of zero electricity price, and the head ratio is 10:1, as in the examples in Appendix A.

10. Appendix C: Energy Trading vs Water Trading

At various points we have referred to "trading" as a paradigm to explain the economics of optimal storage management. Chapter 2, and our report to MDAG, both use the "arbitrage" concept extensively to explain the definition, derivation, and optimal behaviour of marginal water values, and why those values drive and characterise optimal reservoir management practices, from a national cost-benefit perspective. But that discussion was not meant to imply, or propose, an actual water trading market, within which such notional "arbitrage" might occur. Indeed, the entire discussion was based on work done in the late seventies and early eighties, in a non-market public service context.

On the other hand, Section 5.4.4 discusses results from more recent work which focussed explicitly on the possibility of trading of water in a market, in the context of catchments where benefits could be derived from both consumptive and non-consumptive "use", with the latter being characterised by situations in which merely moving water from one place to another, or allowing it to move, produces electricity, or enhances environmental values.

The situation being considered here is rather different, though. It has been suggested that some kind of trading could be desirable between participants holding rights related to the pumped storage scheme. What would be traded is not actually water, though, but abstract representations reflecting the value deemed to be inherent in the right to utilise certain elements of system capacity, and/or the potential energy stored in the upper reservoir of the pumped storage facility.

In that context, no value is actually being attributed to "water", as such. It is merely the medium by which potential energy is stored. If the pumped storage facility was a closed loop system, exactly the same mathematics, and economics, would apply to potential energy stored by chemical reactions in a battery, or by the pressurisation of air, or kinetic energy, in other forms of storage.

The situation is a little more complex when the pumped storage facility is embedded in a host system, but not fundamentally different. Even though the sector constantly talks about the "marginal value of water", the units in which that value is often expressed are not $ per unit of water volume, but $/MWh, and what is really being valued is the potential energy stored by having water held at some elevation above sea level. The water clearly has its own intrinsic value, from various perspectives, some potentially quantifiable and others not. But none of that is reflected in the marginal water values calculated by electricity sector reservoir managers. Hence, in part, our preference for the more general term "Marginal Value of Stored Energy", MVSE, which is equally applicable to other forms of stored potential energy.

So, how should we describe the key interaction with the host system, in which a unit of water is released from the upstream host system reservoir, and pumped up to the upper reservoir of the pumped storage facility? In our basic example, so long as the release remains within the generation capacity of the station(s) associated with the upstream host system reservoir then, in energy terms:

- A unit of potential energy is removed from the upstream host system reservoir;
- That unit of energy (minus some small efficiency losses in generation) is contributed to the electricity system; but
- Maybe 12 units of electrical energy (assuming a 10:1 head ratio, and 20% efficiency loss in pumping) are then more-or-less simultaneously taken from the electricity system for pumping; and
- 10 units of potential energy are ultimately delivered to the upper reservoir of the pumped storage facility.

Thus, the net effect is to use 11 (i.e. 12-1) units of electrical energy to store an extra 9 units (i.e. 10-1) of potential energy in a more secure long-term location within the catchment hydro system. And that will be worthwhile in national cost benefit terms, so long as the long-term MVSE of 10 units energy stored in the upper reservoir of the pumped storage facility minus that of the 1 unit no longer stored in the shorter-term upstream host system reservoir, exceeds the cost of the 11 units of electrical energy consumed in getting it there, which it often will, when electricity prices are low.

11. Appendix D: Coordination Agreements and Optimisation

If coordination between the pumped storage facility and host system is a significant issue, then there is obviously a case for constructing buffer storage to decouple pumped storage and host system operations, as in Section 4.7. But how large should that buffer storage be? If the host system already has significant storage, a national benefit optimisation may well determine that no extra buffer storage is actually required. Instead, it may recommend that existing storage should be utilised to marshal sufficient flows into the (minimal) buffer storage, at the right times to maximise national benefit from the pumped storage facility. And that would have to be the prime consideration for many of the examples discussed in Appendix A. But will real life reflect the optimisation's projections of what should happen?

It is important to understand that, if a national benefit optimisation does recommend buffer storage development it will not be as an alternative to forming a coordination agreement, but in addition to such an agreement. To see this, consider the following:

- As explained previously, a national benefit optimisation will always implicitly assume that a perfect coordination agreement is in place; so
- A national benefit optimisation will always assume that, under that agreement, the operation of the upstream host system can be taken over to maximise the benefits from the pumped storage facility, at the expense of any direct benefit from the upstream system itself, if that improves national welfare; and
- As explained in Appendix A, that seems very likely to be the case, for situations in which the head of the pumped storage facility exceeds the combined head of upstream generation facilities and/or the storage capacity of the upper reservoir of the pumped storage facility exceeds the storage capacity of the upstream host system reservoirs.

For example, if flow constraints force a choice to be made between upstream host system generation and pumped storage generation, the facilitation of pumped storage generation will strongly dominate if, as assumed in those examples, the pumped storage head is ten times that of the upstream host system generation capacity.

So, if there is no buffer storage, the implicit base-line assumption in any national benefit optimisation must surely be that the upstream host system reservoirs will be operated in such a way as to form a substitute for buffer storage, as far as that may be possible, or necessary, to enable unconstrained utilisation of the pumped storage facility:

- In some cases, there may already be enough storage in the host system and/or enough flow freeboard/headroom that there is deemed to be no need for extra buffer storage, but that judgment will be made under the implicit assumption of a perfect coordination agreement.
- If the available host system storage turns out to be inadequate for that task, the optimisation will perhaps recommend more buffer storage to further reduce constraints on pumped storage operations, while still assuming that upstream host system storage is primarily used to enhance the value extracted from the pumped storage facility, under the implicitly assumed coordination agreement.
- If it does recommend buffer storage beyond that level, though, the reason is not to increase flexibility of pumped storage operations, or to enable decoupled operation, but to increase the flexibility of upstream host system operations, so that upstream stations might add their full capacity generation to that of the pumped storage facility, when electricity prices are high, rather than only generating when prices are low, so as to maximise pumpable flows at such times.

If a large buffer storage is technically feasible, and relatively cheap to develop, the incremental cost might be justified by the incremental benefits of increasing upstream host system flexibility. If the incremental cost is low enough, the buffer storage might even be built large enough to allow optimal upstream operational patterns to look very much like they would in a conventional hydro system, with no pumped storage facility embedded. But the incremental benefits involved must be an order of magnitude less than those from pumped storage generation, if the upstream head is only 1/10th of the pumped storage facility head, as in the Appendix A examples.

Even if the marginal cost of buffer storage is assumed to be zero, though, an optimisation will probably never recommend a buffer storage large enough to allow decoupling to the extent that a coordination agreement was redundant: The reason being that it will not, and can not, even calculate what a non-coordinated schedule might look like. Accordingly, we suggest that the buffer storage recommended by a national benefit optimisation will always be in addition to, not instead of, its implicitly assumed coordination agreement.

The cost of developing buffer storage beyond that level may be justified, if those costs are lower than the estimated losses due to factors not considered by the optimisation. If the objective is to decouple operations, it is worth noting that that objective should be attainable if the buffer storage is made large enough to "undo" the impact of upstream operations on the natural flow. So:

- If there is no upstream storage, the upstream manager can not change the natural flow patterns arriving at the buffer storage location, and the incremental cost of buffer storage capacity should be traded off against the incremental benefit of bringing those patterns into better conformity with the optimal pumped storage charge/discharge policy, primarily driven by electricity prices.
- If there is upstream storage, a national benefit optimisation will expect it to operate in ways that would tend to bring the flow patterns arriving at the buffer storage location into better conformity with the optimal pumped storage charge/discharge policy, and that would reduce the need to develop buffer storage there.
- If that upstream storage is operated independently, though, it should be expected to operate in ways that would tend to make the flow patterns arriving at the buffer storage location less compatible with the optimal pumped storage charge/discharge policy, and that would increase the need to develop buffer storage there.
- Thus, the more upstream storage capacity there is, the more flow patterns arriving at the buffer storage location should be expected to deviate from pumped storage requirements, and the more buffer storage will be required to bring the two into alignment.
- And that suggests that the national economic cost of not aligning incentives, whether through an agreement or other means, could be measured by the difference between the cost of the reduced buffer storage needed in the coordinated case, and the increased buffer storage development cost required to undo the impact of upstream system operations being optimised without regard to the economics of pumped storage operations.

The issues here could, and probably should, be investigated by carefully comparing optimisation runs with, and without, particular elements, or perhaps with strategies for some facilities optimised from one manager's perspective held constant, when optimising other strategies from a different perspective. But the details of modelling strategies go well beyond our current brief. What should be borne in mind is that:

- The operational patterns and asset/water values calculated by any optimisation will reflect the logic of national benefit maximisation, and may be quite inconsistent with intuitive perspectives on the situation implicitly adopting the perspective of either manager.
- The recommended developments may not perform as expected, unless a coordination agreement is actually implemented, along the lines discussed in Appendix B.