RBB Economics

Response to the ACCC's analysis of the proposed Viterra port loading capacity auction mechanism

RBB Economics, 23 March 2012

1 Introduction and summary of main points

RBB Economics has been asked by Viterra to comment on an economic report produced by Darryl Biggar from the ACCC entitled "Analysis of the proposed Viterra port capacity auction mechanism" (referred to hereafter as "the ACCC's economic report"). The ACCC's economic report assessed the proposed grain port loading capacity auction mechanism and concluded that this mechanism has certain flaws. Specifically, the ACCC's economic report concluded that:

- The auction mechanism may not, in certain circumstances, come to an outcome at all. In other circumstances, the auction will reach a conclusion but the outcome of the auction will not reflect an efficient allocation of scarce port capacity (that is, a substantial proportion of port capacity will not be allocated at the auction despite high initial demand for that capacity at the start of the auction).

- The possibility of resorting to the first-in-first-served (FIFS) mechanism (which can be used for capacity that has not been allocated through an auction process) alters the incentives on exporters to participate in the auction process, particularly for slots with a high (effective) price and particularly for an exporter which can be reasonably sure of obtaining the allocation that it desires in the FIFS mechanism.

The ACCC's economic report could not – and should not – be used by the ACCC to support a conclusion that Viterra's proposed auction will not reach an efficient conclusion.

Firstly, the ACCC's economic report does not comment on how likely an adverse outcome actually is and more worryingly, presents a distorted view in the paper by only presenting

scenarios where the auction leads to less efficient outcomes and failing to report scenarios where the auction performs well. And although the ACCC's economic analysis gives the impression that the auction price will always increase by precisely the same amount as the rebate, this is not always (or even generally) the case. Indeed the ACCC modelled (but did not report) a scenario where the price and rebate did not move in lockstep and the auction worked well. In the three scenarios that we present using Viterra's data, the auction price does not increase by precisely the same amount as the rebate and the auction achieves an efficient outcome.

Secondly, the ACCC's economic report relies on assumptions that bear no resemblance to reality. The demand functions and supply constraints are completely arbitrary and lack any objective justification. In particular, the demand functions used in the model are static and do not reflect the substitutability that currently exists in the market. Moreover, this substitutability will become more important when exporters are faced with a market based mechanism. The model is, therefore, based on unrealistic, unsupported, and unjustified assumptions.

Finally, we have issues with the approach that the ACCC has taken on a number of other issues. The economic report argues that firms with "deep pockets" will be more successful and as a result capacity will be allocated on the basis of financial strength rather than willingness to pay. Such a finding is completely at odds with the generally held economic view that efficient capital markets will overcome issues of financial inequality. The willingness to pay does not rely on "deep pockets", but on revenues a grain exporter generates by exporting.[1] A more efficient firm will have a higher willingness to pay. That such firms are likely to have "deep pockets" is a result of its efficiency. The economic report also expresses a concern that the auction may be "gamed" by participants without defining the term or drawing a distinction between legitimate gaming and anti-competitive gaming.

Although we do not believe that the results of a simultaneous multiple round ascending bid clock auction can be modelled in a simple way, we have attempted to re-run the ACCC's model using assumptions that are grounded in reality and based on real world data from Viterra. When we do this we find that the assumptions used by the ACCC do not hold and the auction reaches an efficient outcome.

Our conclusion is that the ACCC's economic report simply cannot support a finding that the auction mechanism proposed by Viterra will not lead to an efficient outcome. And if the ACCC's view is that the auction (based on arbitrary, unrealistic and theoretical assumptions) may in some circumstances reach inefficient outcomes, then the onus should be on the ACCC to demonstrate how often it believes these outcomes are likely to be reached and the sensitivity of the assumptions used to derive those results.

---

[1] For example: there are two exporters, exporter A and B, both of which can export their grain at the market price of $100 per tonne. However, exporter A's costs are $70 and exporter B's costs are $80. This means that at most, exporter A's willingness to pay is $30 and B's willingness to pay is $20. The fact that B's willingness to pay is lower, reflects its efficiency (or lack thereof), not because A may have "deep pockets".

We have undertaken sensitivity analysis around the assumptions used by the ACCC by using information from Viterra's actual operations. In particular, we find that the concerns identified in the ACCC's economic report quickly dissolve when we use demand and supply assumptions that reflect the substitution options available to exporters.

The remainder of this note is structured as follows. Section 2 sets out our concerns with the process used by the ACCC to apply economic analysis and our concerns with the way that the results have been presented. Section 3 explains why we believe that the overtly theoretical approach taken by the ACCC limits the ability of the ACCC to draw conclusions about the performance of Viterra's proposed auction. Section 4 shows that even if we retain the ACCC's methodology (which we reject as insufficiently robust to generate results that can inform on the performance of Viterra's auctions) and use assumptions based on Viterra's actual data, then the model paints a very different picture. Section 5 presents examples where we believe that the ACCC's economic report is based on unsound economics. Section 6 concludes.

2 The ACCC's economic report presents a biased and pessimistic view of the proposed auction

The conclusion of the ACCC's economic report is that the auction mechanism proposed by Viterra may not, in certain circumstances, come to an outcome at all. The ACCC's economic report also argues that in other circumstances, the auction will reach a conclusion but the outcome of the auction will not reflect an efficient allocation of scarce port capacity (that is, valuable port capacity will not be allocated at the auction).

In this section we discuss two concerns about the weight that the ACCC may place on these findings. In particular, we argue that the ACCC should place little weight on the economic analysis. Firstly, because the paper only presents two scenarios, both of which suggest that the model may not lead to an efficient allocation of resources. And secondly, because it does not offer a view on how likely the circumstances it discusses in the report are to arise. In our view, this leaves the reader (and, importantly, the decision maker at the ACCC) with the impression that the auction mechanism proposed by Viterra will inevitably lead to an inefficient allocation of capacity. Our modelling suggests this is not the case.

We note that during a meeting held on 21 March 2012 with Viterra and its advisors (including RBB Economics) the ACCC stated that they were not in a position to know how frequently the adverse outcomes identified in the ACCC's economic report were likely to arise, nor whether the assumptions used to generate the results in the ACCC's economic report were grounded in reality. Yet our understanding based on the 21 March 2012 meeting was the ACCC was placing weight on the ACCC's economic report in order to reinforce concerns that it had around the proposed Viterra port capacity auction mechanism. This is a concern for two reasons:

First, we believe that the results presented in the ACCC's economic report are a distorted view of the actual modelling undertaking. Our review of the actual model developed by the ACCC shows that the ACCC prepared a scenario where the auction works well and clears. Yet this

scenario was not included in the final report produced by the ACCC, which only included two examples where the ACCC believes that the auction did not work well. We would be concerned if this biased and selective presentation of results influenced the view of the ACCC on the proposed Viterra port loading capacity auction mechanism.

Second, the ACCC's economic report provides no guidance as to how likely the adverse outcomes it identifies are to arise. It is not clear, for example, whether the ACCC believes that the outcomes it identifies represent a special case which may only arise in a very limited number of cases, or whether the auction design is so fundamentally flawed that it is unlikely to generate an efficient outcome in most cases.

Third, although the ACCC's economic report gives the impression that the auction price will always increase by precisely the same amount as the rebate, this is not always (or even generally) the case. Indeed the ACCC modelled (but did not report) a scenario where the price and rebate did not move in lockstep and the auction worked well.

As a result, we do not believe that the ACCC's economic report provides any meaningful insights about the ability of the auction mechanism proposed by Viterra to reach an efficient outcome.

3 The ACCC's economic report is not able to support the conclusion that the proposed auction is inefficient

The model used in the ACCC's economic report is a simplified and theoretical representation of the auction mechanism which is a revenue-neutral, simultaneous, multiple-round, ascending bid clock auction. It includes a number of (unrealistic) assumptions, such as including demand functions for ports which are completely independent of each other. More worryingly, it includes demand functions and supply constraints that are completely theoretical and divorced from reality. The ACCC during the 21 March 2012 meeting welcomed data from Viterra to apply in the model, but we are alarmed that the ACCC may be placing weight on a simplified and theoretical model without attempting to test that model and the results using empirical evidence from Viterra.

In this section, we show that the assumed demand and supply functions as well as the capacities in the ACCC's economic report are not supported by the actual data provided by Viterra. More importantly, the unfavourable auction outcomes presented in the ACCC's economic report are driven to a large extent by the relatively extreme demand conditions derived by assumptions across four slots. We show that if the excess demand conditions across the slots are smoothed out to reflect the conditions actually faced by Viterra, the auction model does generate efficient market outcomes.

3.1 The assumptions used in the ACCC's economic report are arbitrary

There are a number of problems with the assumptions made in the ACCC's economic report.

First, the ACCC assumes the price at a given port and at a given time is completely independent of the prices at other ports during the same time periods and/or different time periods. This is unrealistic as in reality exporters would substitute to various degrees between the different slots under auction.

Second, all demand functions and capacities are "made-up" and the ACCC's economic report fails to check the extent to which those assumptions reflect the realistic market conditions in South Australia. More importantly, our analysis shows that the capacity assumptions and demand conditions assumed in the ACCC's economic report are not supported by the actual demand and capacity data in South Australia provided by Viterra.

- **The capacities assumed in the ACCC's economic report fall into a very large range which is not supported by the Viterra's data.** The ACCC's economic report assumes four slots with capacities ranging from 10 to 1,000. The largest slot (1,000) is, therefore 9,900 per cent larger than the smallest slot (10). However, the actual capacities provided by Viterra are not that extreme and generally range between 50,000 tonnes to 276,000 tonnes. In other words, the largest slot is only 450 per cent larger than the smallest slot.[2]

- **It is incorrect to assume that the majority of slots are subject to excess demand**. The ACCC's economic report assumes that out of all four slots, there is excess supply at the price of zero for slot D only (*i.e.* if the price is set to be zero; the volume demanded by the market is less than the capacity of the port at the given time). The rest of the slots are subject to excess demand at the price of zero. However, Viterra's data reveals that during the October 2011-September 2012 season (which was one of the highest demand seasons in South Australia's history), there are a similar number of ports with excess supply as excess demand. See Section 4.1.

3.2 The ACCC's assumptions are unrepresentative

The assumptions in the ACCC's economic report rely on an unrepresentative set of assumptions of market demand relative to capacity. They do not represent actual data across ports during different time periods in South Australia. To describe this relative market condition, we calculate the excess demand as a percentage of capacity using the parameters the ACCC

[2] This is based on the theoretical capacity, rather than actual capacity "in-use" which has a smaller range of 46,000 tonnes to 233,000 tonnes for the 2011-12 season.

assumes for the scenarios it presents in Tables 1 and 2 of its economic report (reproduced below in Table 1).

Table 1: Demand and capacity assumptions of the ACCC's economic report

	Slot A	Slot B	Slot C	Slot D
Scenario 1: Demand function	***Q = 133.3-6.667P***	***Q=30-0.15P***	***Q=2666.67-66.67P***	***Q=138.88-2.78P***
Intercept	133.3	30	2666.67	138.88
Price coefficient	6.667	0.15	66.67	2.78
Capacity	**50**	**10**	**1000**	**200**
Excess demand as % of capacity	**167%**	**200%**	**167%**	**-31%**
Scenario 2: Demand function	***Q = 133.3-6.667P***	***Q=30-0.15P***	***Q=2666.67-66.67P***	***Q=41.67-0.833P***
Intercept	133.3	30	2666.67	41.67
Price coefficient	6.667	0.15	66.67	0.833
Capacity	**50**	**10**	**1000**	**200**
Excess demand as % of capacity	**167%**	**200%**	**167%**	**-79%**

Source: the ACCC's economic report; excess demand as a percentage of capacity is calculated by RBB using the ACCC data and the formula is (intercept-capacity)/capacity.

Table 1 shows that the excess demand divided by capacity across the four slots used in the ACCC's economic report ranges from -31 per cent to 200 per cent of the capacity under Scenario 1, and from -79 per cent to 200 per cent of the capacity under Scenario 2. In contrast, as shown in Figure 1 below, the majority of the slots during the harvest season in 2011-12 in South Australia are subject to excess demand (or excess capacity if negative) greater than -76.5 per cent and less than 125.3 per cent.

Figure 1: The frequency of excess demand / capacity across all slots during the harvest season in 2011-12 with excess demand shown on the x-axis

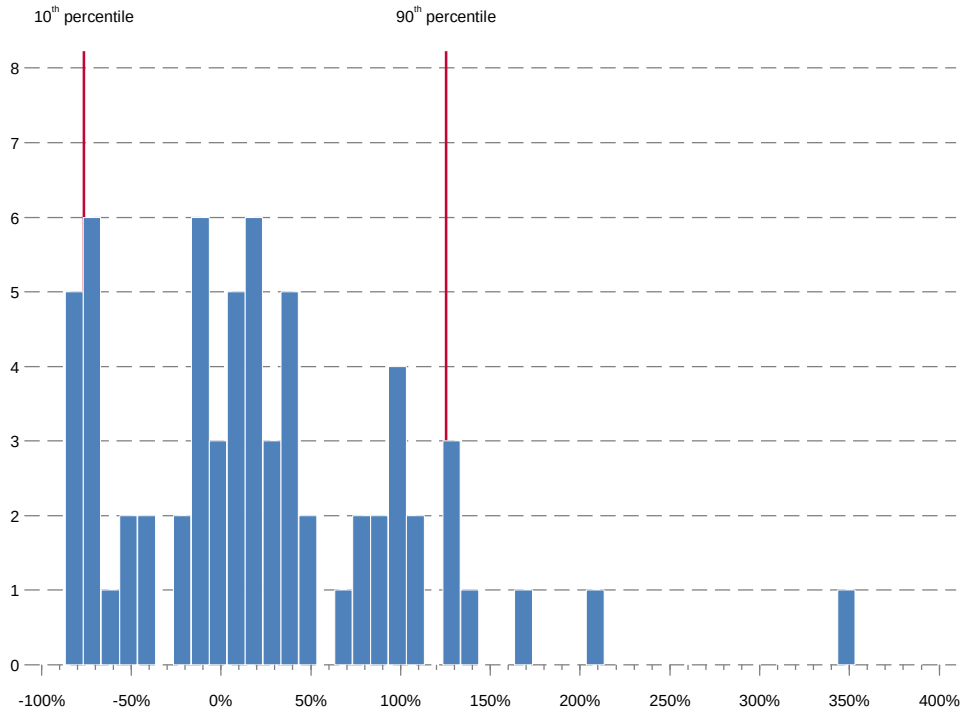

Source: RBB analysis of the Viterra demand and capacity data for the period Oct 2011 to September 2012. The ports included are Port Giles, Port Adelaide – Inner Harbor, Port Lincoln, Outer Harbor, Thevenard and Wallaroo.

Figure 1 illustrates that the excess demands for slots A, B and C in the ACCC's economic report all lie outside of the relevant range we see using data from the harvest season in 2011/12 in South Australia. This leads us to believe that the assumptions made in the ACCC's economic report on demand function and capacities do not represent the market reality. And it is these extreme assumptions that contribute in large part to the inefficient simulation outcomes in both Scenario 1 and Scenario 2 in the ACCC's economic report.

4 The use of data from Viterra provides a different view to that produced in the ACCC's economic report

As discussed in section 3, the ACCC's economic report is based on arbitrary, unsupported and undocumented assumptions, which means that the model is overtly theoretical. This section attempts to test some of the assumptions used in the model with data from Viterra in order to assess whether the results generated by the model apply when using empirical data.

4.1 Slots are not always oversubscribed

First, some of the assumptions made in the ACCC's economic report do not correspond to reality. One of the scenarios presented in the paper considers what happens when at some point in the auction process every slot is oversubscribed (page 3). We have collected data from Viterra for the 2011-12 season which represented one of the largest, most capacity-constrained shipping periods in South Australia's history.

Figure 2 shows that even during that period, there was always at least one slot that was under-subscribed. It should be noted that this is conservative as it does not include cases where there are currently no bookings for 2011-12.[3]

Figure 2: Number of ports with available capacity during different months (October 2011 – September 2012)

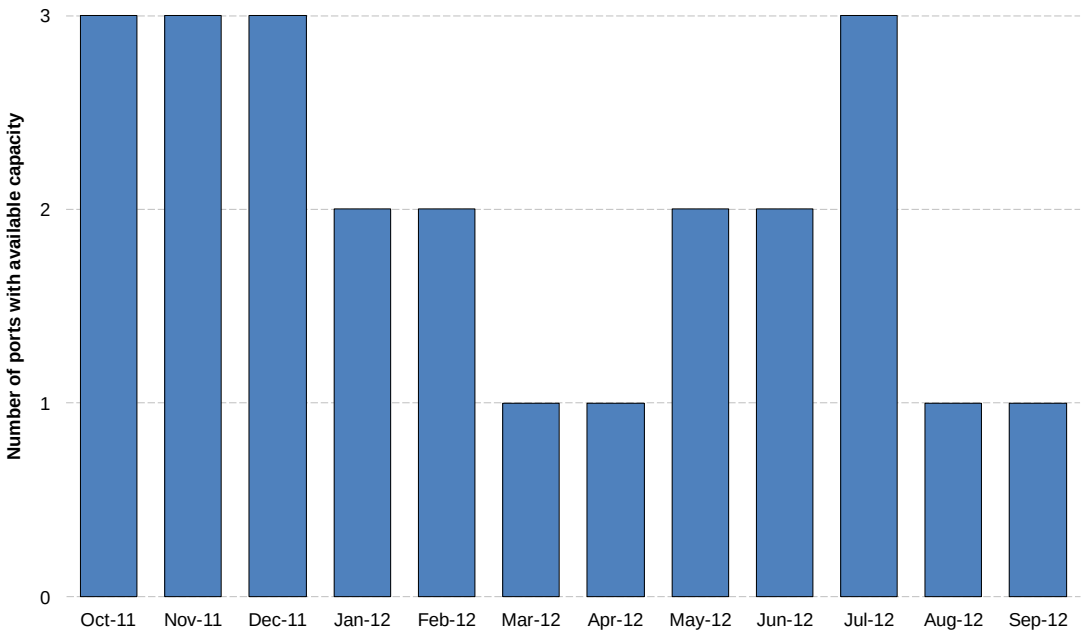

Source: Viterra data, RBB analysis

Actual data from Viterra, therefore, suggests that it is extremely unlikely that all slots will be over-subscribed. The implication of this is that the scenario presented in Table 1 of the ACCC's economic report, namely that Viterra's auction mechanism may not come to an outcome at all, is at best a extreme theoretical assumption and not an outcome that can be considered to be likely to occur as a result of the implementation of the auction mechanism proposed by Viterra.

[3] There are no bookings yet for July, August and September 2012 for Thevenard for August and September 2012 for Port Giles and for September 2012 for Wallaroo

Another serious concern that we have with the ACCC's economic report is that the demand functions are not only groundless, but assumed to be completely independent of each other. We have found that the adverse outcomes generated in the ACCC's economic model are largely the result of inelastic and independent demand assumed by the ACCC. What this means is that the ACCC assumes that exporters are extremely inflexible with regard to the port that they want to use and will not respond to relative price changes through the auction premium. This lack of flexibility is assumed rather than determined by analysis of actual data.

4.2 Results of an auction simulation using Viterra data

The annex presents an auction simulation using Viterra's data. It simulates six slots, each representing one of the grain ports operated by Viterra (Port Giles, Port Adelaide – Inner Harbor, Port Lincoln, Port Adelaide – Outer Harbor, Thevenard and Wallaroo). The model uses real-world capacity and demand intercepts using data from Viterra, and forming reasonable demand slopes using Viterra's knowledge of each port and the willingness of clients to move bookings between ports. The model finds that the auction clears effectively and in low, medium and high demand conditions.

In reality, Viterra proposes to run separate auctions for the harvest shipping period (1 October – 31 January) and the non-harvest shipping period (1 February-30 September). Clients can enter a quantity bid (in tonnes) for any half-monthly period at any port once the auction begins. For the harvest shipping period this represents eight half-monthly periods across six grain ports, which is 48 slots.

We model the general features of the auction using six representative slots, one for each of the grain ports that Viterra operates. We introduce complexity by varying the demand assumptions to represent a low demand case where there is low demand for most slots, a central demand case where there is high demand for some slots and high demand case where there is high demand for all slots.

As a first modelling step, we set the capacity and demand intercepts for the model. The capacity is invariant across each month so we simply use the capacity in-use figure for each port. We checked the variability of demand across each month and use the first quartile of demand for a particular port for the low demand case, the median demand for the central case and the third quartile for the high demand case. Table 2 below displays the capacity in-use, and the lower quartile, median and upper quartile of demand for each port across all months in 2011-12. We later use the lower quartile demands as our low-demand scenario, the median demands as our central case and the upper quartile demands as our high-demand scenario.

Table 2: Demand and capacity in-use across South Australian grain ports operated by Viterra for the 2011-12 season

Port	Capacity in-use (1000 tonnes)	First quartile demand (1000 tonnes)	Median demand (1000 tonnes)	Third quartile demand (1000 tonnes)
Port Giles	100	84.0	101.9	120.0
Port Adelaide – Inner Harbor	46	47.5	61.0	92.5
Port Lincoln	233	233.0	298.7	425.0
Port Adelaide – Outer Harbor	186	182.5	310.0	426.5
Thevenard[4]	85	20.0	20.0	36.0
Wallaroo	100	52.0	117.5	145.0

Source: RBB, Viterra data

Given that there is no auction data currently available for South Australia, which would illustrate how demand varies as the price varies, we assume demand slopes based on Viterra's assessment of the willingness of clients to move demand between ports (or different time periods).

The demand functions that we have used in our model simulation are based on assumptions. But unlike the approach in the ACCC's economic report, our assumptions are clearly documented and informed by evidence on how the grain export market in South Australia actually operates. We recognise that neither our approach nor that used by the ACCC can predict likely outcomes with complete accuracy. This can only be achieved by using real world data with the experience of one or more auctions in South Australia.

We ran our six-slot model for each of the demand scenarios described in Table 2 above. We describe the results of the model using these realistic and representative capacity and demand assumptions below.

For the lower quartile case, demand is below or close to capacity at every slot. This leads the auction to clear almost immediately. The auction allocates 83.3 per cent of capacity, following an initial demand of 82.5 per cent of capacity. The slight increase in demand is due to the change in relative prices for each port that is caused by the auction mechanism. For the median case, the auction clears at round 92 with 92.5 per cent of capacity allocated through the auction. This was following initial demand exceeding capacity by 159,000 tonnes or 21.2 per cent. For the high demand case, the auction clears after 138 rounds, with 95.4 per cent of

[4] For Thevenard, there were only nine observations as no bookings had been made for the last three months of the season, from July to September 2012.

capacity allocated. This was with initial demand exceeding capacity by 495,000 tonnes or 66.0 per cent. However, it should be noted that the number of rounds is conservative because it uses independent demand functions. In reality, the response of exporters moving to slots that have spare capacity or a more favourable price would lead the auction to clear more quickly.

In summary, the auction model clears in all three cases. With low initial demand the auction allocates more capacity than initially demanded. The auction allocates 92.5 per cent of capacity in the medium demand case and 95.4 per cent of capacity in the high demand case, where initial demand exceeded capacity. The model suggests the auction will work effectively in practice even with high demand for capacity.

4.3 The ACCC has not taken dynamic factors into account

The real issue, however, is how exporters will respond to the price signals that will be generated through the auction. Our understanding is that although exporters may have a strong preference for some ports (perhaps because the port in question may be a deep water port), they can and do switch their demand across ports. There are two important implications of this:

- First, the demand function will be far more elastic in practice than is assumed in the ACCC's economic report; and

- Second, the demand functions will be inter-dependent rather than independent.

The second point is crucial. With interlinked demand functions, exporters would explicitly move their demand based on the capacity available at each slot and the price of each slot. With independent demand functions the only way demand for one slot affects demand at another slot is indirect and takes place through the impact of demand at a particular slot on the rebate. This increased demand flexibility would require less adjustment in each individual price in order to ensure each slot clears and so would lead the auction to clear sooner.

What is lacking from the ACCC's economic report – and from our response to that analysis – is a view based on actual behaviour of how exporters will respond to actual price changes and how they will switch their demand across different ports in South Australia. This is the crucial piece of information and has been dealt with by assumption. Crucially, this demand response is likely to be different from what occurs in Western Australia, given the different port characteristics and different geographic dispersion of grains zones and ports.

What we do know, however, is that – despite the effect of the auction rebate premium - the auction is a market-based mechanism and exporters will react to the price signals in South Australia in ways that we cannot predict. This is - and increasingly will become - a dynamic market and the inter-dependencies between ports based on relative price changes must be taken into account.

The ACCC has not taken these dynamic factors into account in its analysis and we have only done so on the basis of assumptions. In our view, this means that the results that the ACCC has generated need to be treated with caution.

5 Additional comments around ACCC's economic reasoning

There are two additional areas where we disagree with the economic reasoning of the ACCC.

The first of these relates to the impact of "deep pockets" on the outcome of an auction. The ACCC argues that as the auction progresses and the auction price increases, a higher financial outlay will be required by the exporter. This higher financial outlay also increases the risk that the exporter does not use some of the capacity that it purchases at auction. The ACCC then argues that "in practice" as the auction price increases, some exporters may be forced to drop out due to less "deep pockets", which it defines as the ability to finance the auction price, or because they are less able to bear the risk of shipping less than the capacity purchased at the auction.

The ACCC then argues that "eventually, perhaps" sufficient exporters will drop out so that demand for slot capacity drops below the available capacity and the auction mechanism will come to a conclusion. Our concern, however, is not with the lack of confidence with which the ACCC appears to hold this view, but with its conclusion that "at this point, the auction mechanism is not necessarily allocating capacity to those who value it most highly, but to those with the deepest pockets or the greatest ability to bear risk".

The ACCC's conclusion is at odds with the generally accepted view that with efficient capital markets the bidder with best business plan, producing the highest expected profits, will get best financial backing and will be able to place winning bids. The ACCC's concern about the way capacity is allocated, therefore, is based on an assumption that capital markets are inefficient. In the absence of such an assumption, the "deep pockets" or increased risk appetite of some firms will be irrelevant to the ultimate result of an auction and capacity will be allocated to those who value it most highly.

The second area where we disagree with the economic reasoning of the ACCC is related to the discussion of the FIFS mechanism. Under the mechanism proposed by Viterra, capacity which is not allocated through the auction mechanism is made available to exporters through a FIFS mechanism. The ACCC's concern is that the presence of the FIFS mechanism can give rise to strategic behaviour and inefficient outcomes in the auction. The strategic behaviour it refers to is the potential for exporters to "game" the auction.

As with the concerns over the auction mechanism discussed in section 2, the ACCC raises concerns without providing an opinion on how likely it is that those concerns will arise. In fact,

the ACCC's economic report sets out three conditions for the gaming strategy that it identifies to arise. These are:

- The player's bid must be large relative to the excess demand for a slot;
- The player must be reasonably certain of the auction coming to an end; and
- The player needs to have a strong assurance of being able to acquire any capacity that is foregone in the auction mechanism through the FIFS process.

But how likely is it that those three conditions will be met? Our understanding is that the rules around the FIFS mechanism developed by Viterra restrict the ability of exporters to pick up large quantities of capacity through FIFS. This suggests that the third condition will, in practice, be unlikely to hold. And even if these conditions met, it is not at all clear whether the "gaming" or strategic behaviour identified by the ACCC is something that they need to be concerned about. Exporters have the choice of going to auction and/or using the FIFS mechanism. Depending on the risk appetite of an exporter, they may decide to withdraw from an auction and take their chances on the FIFS mechanism. They may succeed in securing capacity (even scarce capacity) through the FIFS mechanism. They may also miss out entirely if other exporters secure capacity through the auction system and pay a higher price (which reflects the value they place on the capacity). It is not clear why such an outcome – under either scenario – is a source for concern to the ACCC.

It may be that the ACCC needs to be clear about what it defines as "gaming" or strategic behaviour. In particular, they need to be clear about how they distinguish between "legitimate" gaming and "illegitimate" gaming. Is the ACCC, for example, concerned that there is fraud taking place in the auction? Or is the ACCC suggesting that there has been manipulation enabled by an abuse of market power? Either of those situations may require further investigation by the ACCC. But it is not at all clear why the strategic behaviour identified by the ACCC is necessarily a cause for concern.

6 Conclusion

Our conclusion is that the ACCC's economic report simply cannot support a finding that the auction mechanism proposed by Viterra will not lead to an efficient outcome. And if the ACCC's view is that the auction (based on arbitrary, unrealistic and theoretical assumptions) may in some circumstances reach inefficient outcomes, then the onus should be on the ACCC to demonstrate (with data) how often it believes these outcomes are likely to be reached and the sensitivity of the assumptions used to derive those results.

We have undertaken sensitivity analysis around the arbitrary and unrealistic assumptions used by the ACCC by using information from Viterra's actual operations. In particular, we find that the concerns identified in the ACCC's economic report quickly dissolve when we use demand and supply assumptions that reflect the substitution options available to exporters.

Annex A Auction simulation using Viterra data

This annex summarises RBB's auction simulation using Viterra data. This annex is structured as follows.

- Section 1 describes the data provided by Viterra;
- Section 2 describes how we obtain capacity and demand intercepts;
- Section 3 describes how we select demand slopes; and
- Section 4 describes the outcome of the model for low, central and high demand cases.

This model suggests the Viterra auction will work well even for high demand cases.

A.1 Data provided by Viterra

Viterra provided monthly demand and capacity data for the 2011-12 season (*i.e.* from October 2011 to September 2012) for each of the six grain ports it operates (Port Giles, Port Adelaide – Inner Harbor, Port Lincoln, Port Adelaide – Outer Harbor, Thevenard and Wallaroo). The data is monthly as that is how Viterra currently allocates capacity.

The demand data shows the total quantity in tonnes of all booking requests for each of Viterra's six grain ports for the 2011-12 season. The demand number captures periods of excess demand at a particular port at a particular month, because the original bookings are unconstrained by the actual capacity of each port for each month. Since these auction bookings were made without taking into account an auction fee (*i.e.* assuming a zero auction fee), this demand data allows us to set the intercept of the demand function (demand at zero price) in our modelling. We understand that this has been a season with exceptionally high demand so using this data to derive demand assumptions is conservative.

Viterra also provided capacity data for the 2011-12 season across each of its six grain ports. The capacity does not vary from month to month. Viterra provided both nominal capacity and the "in-use" capacity. This latter number is a more realistic capacity measure as it takes into account factors such as delays, periods of maintenance shut-down and other random factors that mean that capacity actually available is less than the theoretical maximum capacity for a port for a particular month. We use the "in-use" capacity measure for our modelling because it is both more conservative and represents the measure Viterra actually uses to allocate capacity.

A.2 Choosing a model

In reality, Viterra proposes to run separate auctions for the harvest shipping period (1 October – 31 January) and the non-harvest shipping period (1 February-30 September). Clients can enter a quantity bid (in tonnes) for any half-monthly period at any port once the auction begins. For the harvest shipping period this represents eight half-monthly periods across six grain ports, which is 48 slots.

We model the general features of the auction using six representative slots, one for each of the grain ports that Viterra operates. We introduce complexity by varying the demand assumptions to represent a low demand case where there is low demand for most slots, a central demand case where there is high demand for some slots and high demand case where there is high demand for all slots.

As a simplifying assumption we use independent demand functions. This is conservative. With interlinked demand functions, clients would explicitly move their demand based on the capacity available at each slot and the price of each slot. With independent demand functions the only way demand for one slot affects demand at another slot is indirect and takes place through the impact of demand at a particular slot on the rebate. This increased demand flexibility would require less adjustment in each individual price in order to ensure each slot clears and so would lead the auction to clear sooner.

A.3 Finding capacity and demand intercepts

As a first modelling step, we set the capacity and demand intercepts for the model. The capacity is invariant across each month so we simply use the capacity in-use figure for each port. We checked the variability of demand across each month and use the first quartile of demand for a particular port for the low demand case, the median demand for the central case and the third quartile for the high demand case. These numbers are displayed in Table 2 below.

Table 3: Demand and capacity (in '000 tonnes) in-use across South Australian grain ports operated by Viterra for the 2011-12 season

Port	Capacity-in-use	1st quartile demand	Median demand	3rd quartile demand
Port Giles	100	84.0	101.9	120.0
Port Adelaide – Inner Harbor	46	47.5	61.0	92.5
Port Lincoln	233	233.0	298.7	425.0
Port Adelaide – Outer Harbor	186	182.5	310.0	426.5
Thevenard[5]	85	20.0	20.0	36.0
Wallaroo	100	52.0	117.5	145.0

Source: RBB, Viterra data

Table 2 shows that the in-use capacity varies from 46,000 tonnes at the Port Adelaide – Inner Harbor to 233,000 tonnes at Port Lincoln. For the lower quartile of demand for a particular port across all months in the 2011-12 season, demand varies from 20,000 tonnes at Thevenard to 233,000 tonnes at Port Lincoln. These demands are below or close to the respective capacity-in-use at each port. For the median, demand varies from 20,000 tonnes at Thevenard to 298,700 tonnes at Port Lincoln. For the upper quartile, demand varies from 36,000 tonnes at Thevenard (still below the available capacity-in-use) to 426,500 tonnes at Port Adelaide – Outer Harbor (showing excess demand of 240,500 tonnes or 129.3 per cent).

A.4 Selecting demand slopes

Given that there is no auction data currently available for South Australia, which would illustrate how demand varies as the price varies, we assume demand slopes based on Viterra's assessment of the willingness of clients to move demand between ports (or different time periods).

Viterra find that customers are generally willing to move demand from one port to another. In periods of excess demand, Viterra has often suggested this to clients and moved bookings in this way. The large ports, Port Lincoln and Port Adelaide – Outer Harbor, are also the ones that clients are most keen to use because they have the required 13-14m depth to accommodate large Panamax vessels and can ship "just-in-time" and at high volume. However, even for these ports customers have been willing to move bookings to another port. Port Adelaide – Inner Harbor, in particular, can be used as an overflow for Port-Adelaide – Outer Harbor. As these two ports have the same supply chain they can effectively be used as one port with two berths. During a shut-down at Outer Harbor during November 2011, Vittera moved bookings to Inner Harbor. Thevenard is a less attractive port as less pilots, needed to direct ships into and out of

[5] For Thevenard, there were only nine observations as no bookings had been made for the last three months of the season, from July to September 2012.

port, are available, the channel is small and only 8.7m deep and the ships serving this port tend to be older.

The option to move the month of a booking has only recently become available. There is some willingness amongst customers to move bookings between time periods. During the peak shipping period, which typically lasts from January to April, and to some extent to June, customers may be willing to move their booking subject to it being within this general timeframe.

Using the information supplied by Viterra, we generated the following representative demand curves. These are displayed using the median intercept as described above.

Figure 3: Port Giles representative demand curve

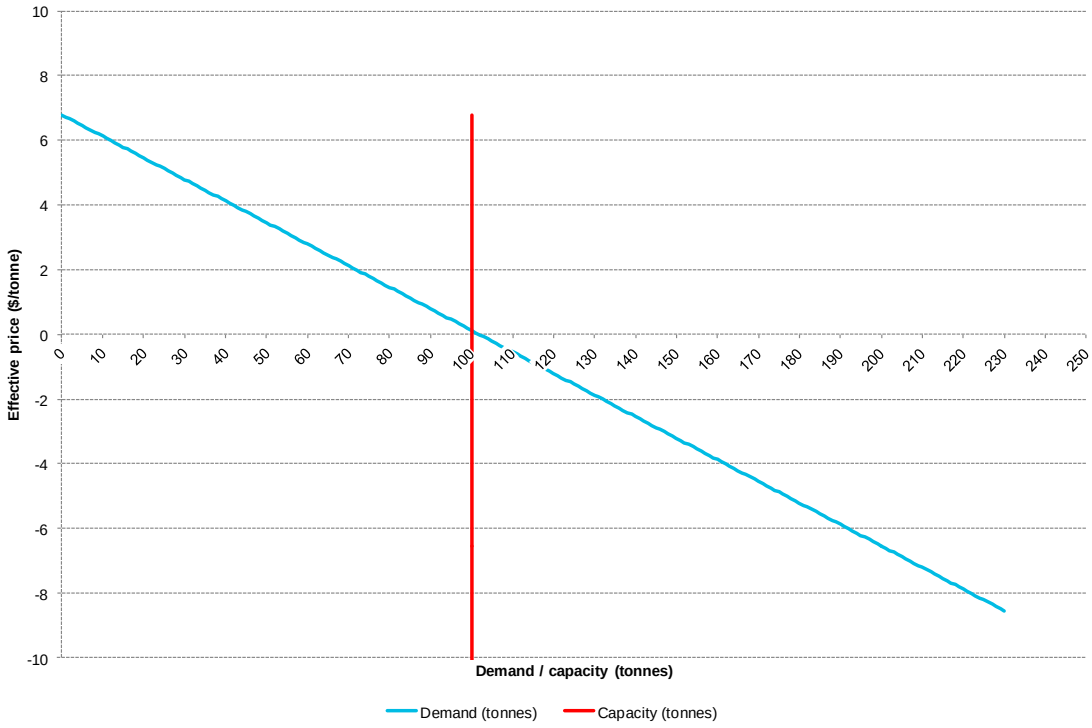

Source: Viterra data, RBB analysis

Figure 4: Port Adelaide – Inner Harbor representative demand curve

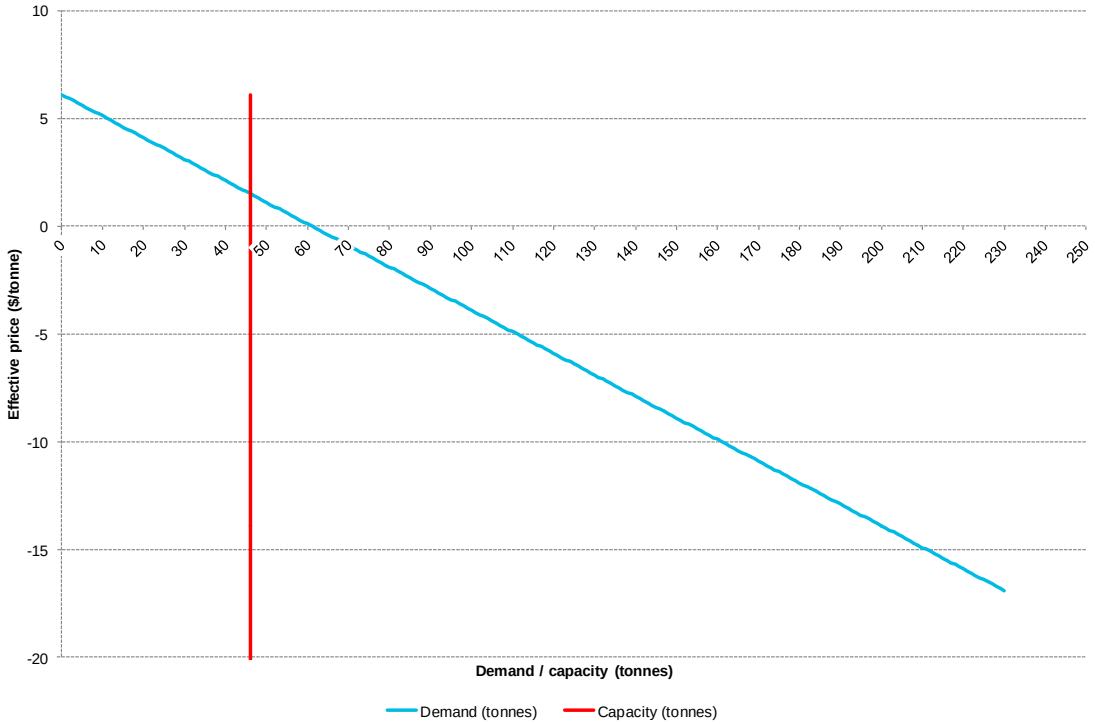

Source: Viterra data, RBB analysis

Figure 5: Port Lincoln representative demand curve

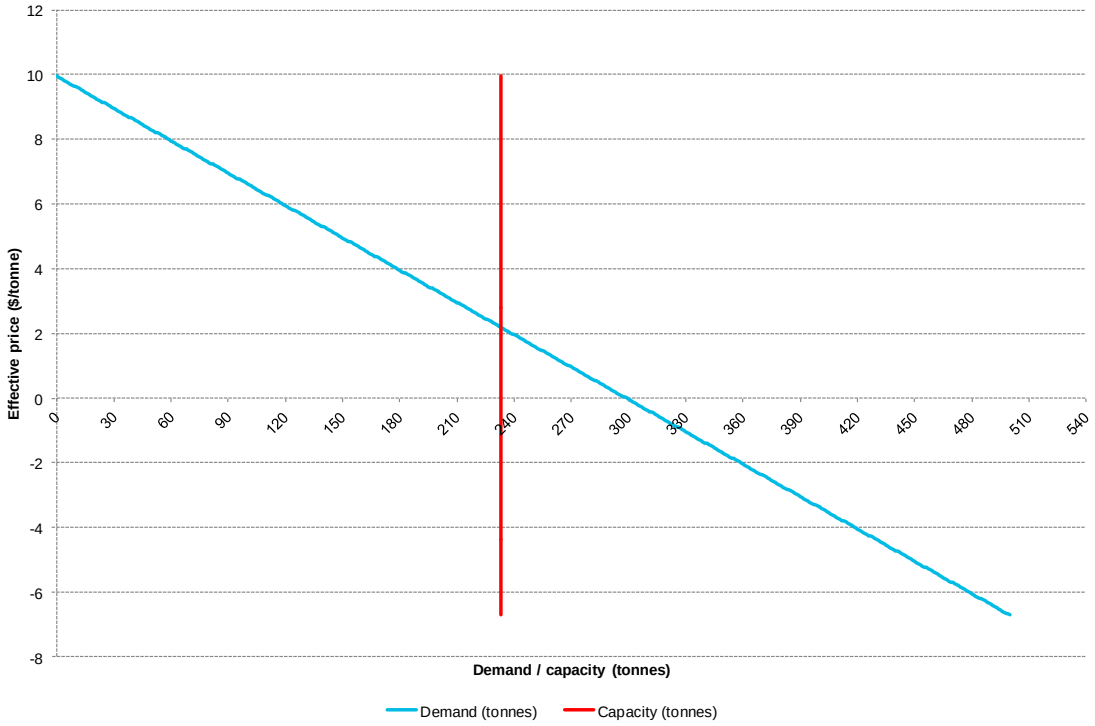

Source: Viterra data, RBB analysis

Figure 6: Port Adelaide – Outer Harbor representative demand curve

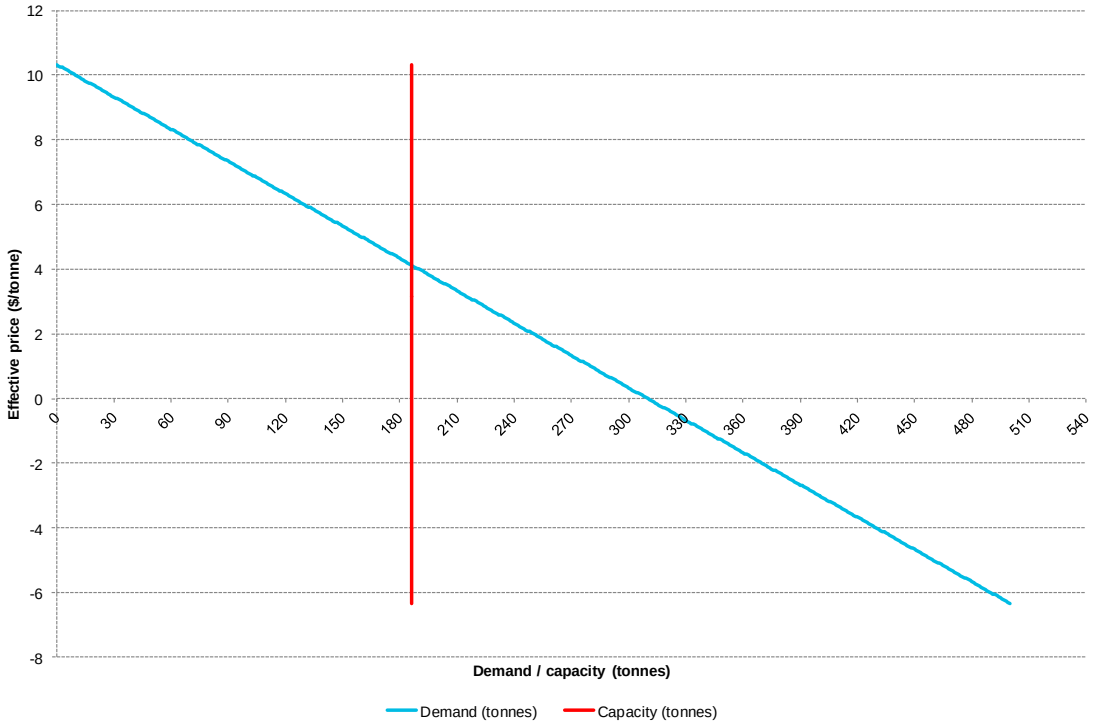

Source: Viterra data, RBB analysis

Figure 7: Thevenard – Outer Harbor representative demand curve

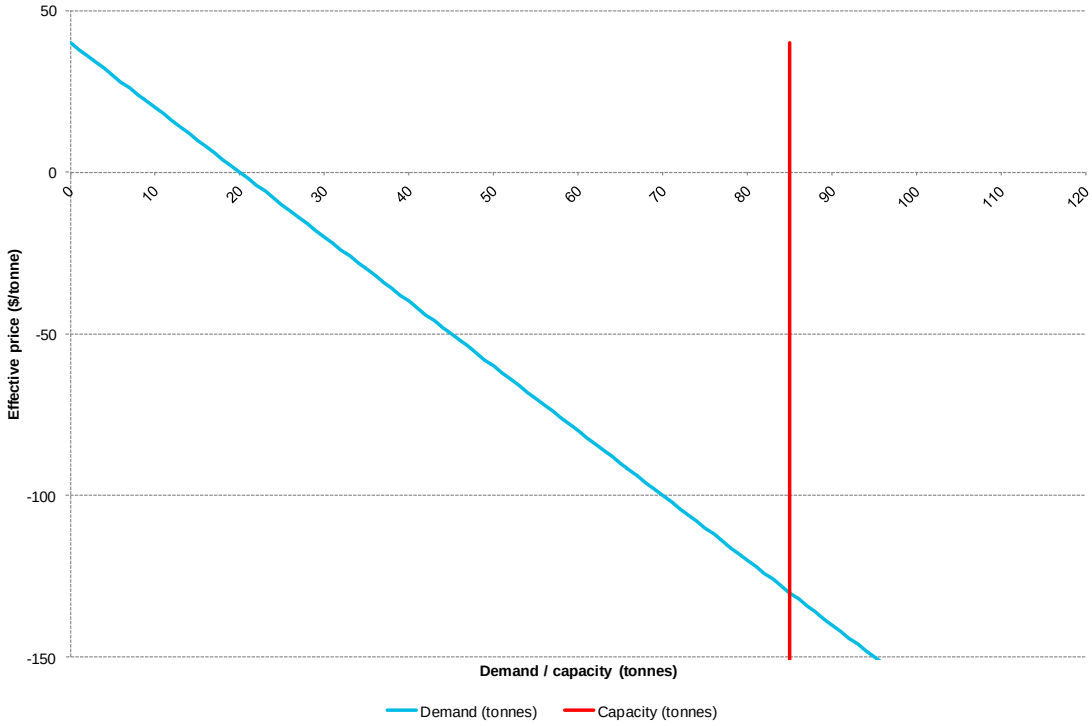

Source: Viterra data, RBB analysis

Figure 8: Wallaroo representative demand curve

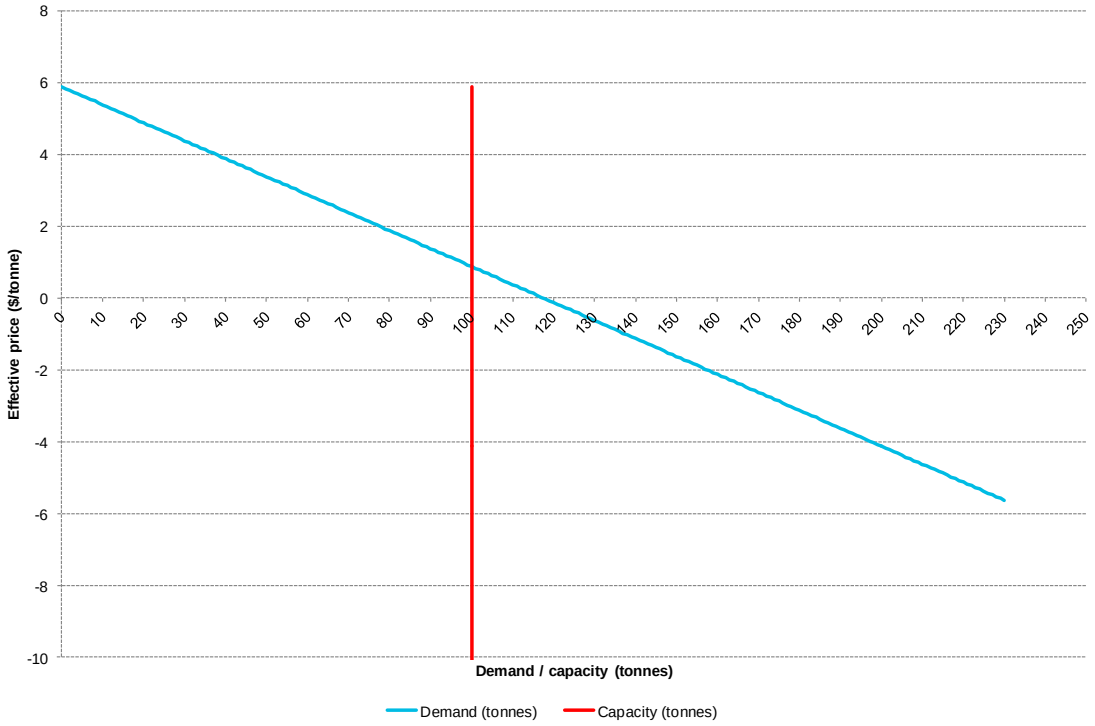

Source: Viterra data, RBB analysis

A.5 Outcome of model under low, central and high demand scenarios

We ran our six-slot model for each of the demand scenarios described in Table 2 above. We describe the results of the model using these realistic and representative capacity and demand assumptions below.

For the lower quartile case, demand is below or close to capacity at every slot. This leads the auction to clear almost immediately. The auction allocates 83.3 per cent of capacity, following an initial demand of 82.5 per cent of capacity. The slight increase in demand is due to the change in relative prices for each port that is caused by the auction mechanism.

Figure 9: Simulated auction results with low demand

Slot	Giles			Inner Harbor			Lincoln			Outer Harbor			Thevenard			Wallaroo		
Demand	*Q =84-15P*			*Q =47.5-5P*			*Q =233-30P*			*Q =182.5-30P*			*Q =20-0.5P*			*Q =52-20P*		
Capacity	100			46			233			186			85			100		
Round	Price	Eff Price	Rebate Demand	Price	Eff Price	Rebate Demand	Price	Eff Price	Rebate Demand	Price	Eff Price	Rebate Demand	Price	Eff Price	Rebate Demand	Price	Eff Price	Rebate Demand
0	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0
	0.00		84	0.00		48	0.00		233	0.00		183	0.00		20	0.00		52
1	0.00	0.00	0	0.50	0.00	0	0.00	0.04	0.04	0.00	0.00	0	0.00	0.04	0.04	0.00	0.04	0.04
	-0.04		85	0.46		45	-0.04		234	-0.04		184	-0.04		20	-0.04		53
2	0.00	0.00	0	0.50	0.00	0	0.50	0.22	0.22	0.00	0.00	0	0.00	0.22	0.22	0.00	0.22	0.22
	-0.22		87	0.28		46	0.28		225	-0.22		189	-0.22		20	-0.22		56
3	0.00	0.00	0	1.00	0.00	0	0.50	0.40	0.40	0.50	0.00	0	0.00	0.40	0.40	0.00	0.40	0.40
	-0.40		90	0.60		45	0.10		230	0.10		180	-0.40		20	-0.40		60

Source: Viterra data, RBB analysis

For the median case, the auction clears at round 92 with 92.5 per cent of capacity allocated through the auction. This was following initial demand exceeding capacity by 159,000 tonnes or 21.2 per cent.

Figure 10: Simulated auction results with central demand (last five rounds shown only)

Slot	Giles			Inner Harbor			Lincoln			Outer Harbor			Thevenard			Wallaroo		
Demand	*Q =101.875-15P*			*Q =61-5P*			*Q =298.725-30P*			*Q =310-30P*			*Q =20-0.5P*			*Q =117.5-20P*		
Capacity	100			46			233			186			85			100		
Round	Price	Eff Price	Rebate Demand	Price	Eff Price	Rebate Demand	Price	Eff Price	Rebate Demand	Price	Eff Price	Rebate Demand	Price	Eff Price	Rebate Demand	Price	Eff Price	Rebate Demand
88	39.00	0.21	39 99	42.00	3.21	39 45	41.00	2.21	38.79 232	43.00	4.21	39 184	0.00	-38.79	38.79 39	39.50	0.71	38.79 103
89	39.00	0.15	39 100	42.00	3.15	39 45	41.00	2.15	38.85 234	43.00	4.15	39 186	0.00	-38.85	38.85 39	40.00	1.15	38.85 95
90	39.00	-0.03	39 102	42.00	2.97	39 46	41.50	2.47	39.03 224	43.00	3.97	39 191	0.00	-39.03	39.03 40	40.00	0.97	39.03 98
91	39.50	0.26	39 98	42.50	3.26	39 45	41.50	2.26	39.24 231	43.50	4.26	39 182	0.00	-39.24	39.24 40	40.00	0.76	39.24 102
92	39.50	0.20	39 99	42.50	3.20	39 45	41.50	2.20	39.30 233	43.50	4.20	39 184	0.00	-39.30	39.30 40	40.50	1.20	39.30 93

Source: Viterra data, RBB analysis

For the high demand case, the auction clears after 138 rounds, with 95.4 per cent of capacity allocated. This was with initial demand exceeding capacity by 495,000 tonnes or 66.0 per cent.

Figure 11: Simulated auction results with high demand (last five rounds shown only)

Slot	Giles			Inner Harbor			Lincoln			Outer Harbor			Thevenard			Wallaroo		
Demand	*Q =120-15P*			*Q =92.5-5P*			*Q =424.955-30P*			*Q =426.5-30P*			*Q =36-0.5P*			*Q =145-20P*		
Capacity	100			46			233			186			85			100		
Round	Price	Eff Price	Rebate Demand	Price	Eff Price	Rebate Demand	Price	Eff Price	Rebate Demand	Price	Eff Price	Rebate Demand	Price	Eff Price	Rebate Demand	Price	Eff Price	Rebate Demand
134	59.50	1.58	58 96	67.00	9.08	58 47	64.50	6.58	57.92 227	66.00	8.08	58 184	0.00	-57.92	57.92 65	60.00	2.08	57.92 103
135	59.50	1.48	58 98	67.50	9.48	58 45	64.50	6.48	58.02 231	66.00	7.98	58 187	0.00	-58.02	58.02 65	60.50	2.48	58.02 95
136	59.50	1.33	58 100	67.50	9.33	58 46	64.50	6.33	58.17 235	66.50	8.33	58 177	0.00	-58.17	58.17 65	60.50	2.33	58.17 98
137	60.00	1.67	58 95	67.50	9.17	58 47	65.00	6.67	58.33 225	66.50	8.17	58 181	0.00	-58.33	58.33 65	60.50	2.17	58.33 102
138	60.00	1.60	58 96	68.00	9.60	58 45	65.00	6.60	58.40 227	66.50	8.10	58 184	0.00	-58.40	58.40 65	61.00	2.60	58.40 93

Source: Viterra data, RBB analysis

In summary, the auction model clears in all three cases. With low initial demand the auction allocates more capacity than initial demanded. The auction allocates 92.5 per cent of capacity in the medium demand case and 95.4 per cent of capacity in the high demand case, where initial demand exceeded capacity. The model suggests the auction will work effectively even with high demand for capacity.