A Market-based Regional Approach to Valuing and Reducing GHG Emissions from Power Sector

An ISO-administered carbon price as a compliance option for EPA's Existing Source Rule

A DISCUSSION PAPER PREPARED FOR

Great River Energy

PREPARED BY

Judy Chang
Jurgen Weiss, PhD.
Yingxia Yang, PhD.

April 2014

THE Brattle GROUP

This report was prepared for Great River Energy. All results and any errors are the responsibility of the authors and do not represent the opinion of The Brattle Group, Inc. or its clients.

Acknowledgement: We acknowledge the valuable contributions of many individuals to this report and to the underlying analysis, including members of The Brattle Group for peer review.

This discussion paper is intended to outline an attractive regional approach to complying with anticipated regulation in Section 111(d) of the Clean Air Act. It is not intended to preclude other potential regional approached to achieve the same goals. We encourage and invite comments and will incorporate those in future improved versions.

Table of Contents

I. Summary of Proposed Approach

This discussion paper outlines a proposal for a regional, market-based approach as a compliance option for the EPA's forthcoming rule under Section 111(d) of the Clean Air Act to regulate greenhouse gas emissions from existing power generating sources ("Existing Source Rule").

The foundation of the proposed approach includes the following features:

- States in a region translate the required greenhouse gas emissions reductions promulgated by the EPA in its Existing Source Rule into targets for the relevant regional power markets.
- An independent party facilitates the states that participate in the approach to translate the emissions target into a regional carbon price.
- The ISO administers the wholesale electricity market with an added carbon price.
- Generators pay for carbon emissions based on the regional carbon price and their carbon emissions rate.
- Generators comply with Section 111(d) requirements through participation in the regional market with a carbon price that achieves the regional carbon emissions reduction target.
- The ISO collects the carbon revenues from the generators.
- The ISO refunds the carbon revenues to load-serving entities based on load share ratios or transition towards such an approach over time, should immediate full refunding to LSEs based on load-share alone be deemed too disruptive.
- The carbon price will be adjusted by the independent party based on the actual emissions compliance level relative to the targets.

The proposed approach would allow states to include in their State Implementation Plans (SIPs) used to meet Section 111(d) a commitment to meeting the regional targets as a way to comply with the greenhouse gas emission reduction targets for generation resources that participate in the proposed approach. To gain the EPA's acceptance of this approach, and recognizing that attaining agreement among multiple states on appropriate carbon prices will not be simple, we propose that an independent entity facilitate the participating states' efforts to agree on an initial set of carbon prices estimated to be sufficient to achieve required emissions reductions as well as on a mechanism for adjusting carbon prices should the actual path of emissions diverge from the path expected from the initial price path.

The independent entity could be the ISO itself, or some other entity set up by participating states. In either case, we anticipate that the ISO would be closely involved in the analyses underlying the determination of an initial carbon price path. By staying involved, important system cost and reliability issues would be considered in setting both an initial carbon price path and in adjusting this path if necessary over time.

After an initial carbon price path has been agreed upon, the ISO would use these carbon prices in its day ahead and real time market decisions, and thereby implement a regional carbon emissions constraint in addition to its other functions, namely to ensure reliable power supply at least cost to consumers. In practice, the ISO would charge each participating generator for CO_2 emissions at a rate equal to the carbon price (in $/ton) multiplied by each generating source's emission rate (ton/MWh). By doing so, the CO_2 emissions become an additional variable cost for CO_2-emitting generators. The ISO charges generators for carbon emissions and generators would therefore include these costs to their offers.

The ISO would then use the offers, including the carbon price component, to clear the markets and dispatch the system just as they do today. All else being equal, the normal system dispatch will tend to increase the dispatch of lower carbon-emitting resources relative to what would have been dispatched absent the carbon price. We anticipate that the wholesale market clearing prices will increase with the carbon price in place and, in the near-term, increase the energy margins (energy price minus generators' variable costs including carbon costs) of low-emitting resources relative to higher carbon-emitting resources. The effect will be to lower emissions in the short run due to re-dispatch, and, over time, alter the generation mix through generator entry and retirement to efficiently accomplish the intent of the EPA.

The increase in the market price will be paid by all buyers of electricity (notably all load serving entities, and ultimately end-users), and received by all dispatched generators. Carbon-emitting generation resources will effectively be paying the ISO for their emissions. This means that for every unit of energy production, generators would receive the market clearing prices (locational prices) less the carbon price (based on the amount of CO_2 emission from each generator). For example, if the market clearing price is $50/MWh and a generator emits 1 ton per MWh and the carbon price is $10/ton, then the generator would receive the market clearing price of $50/MWh less $10/MWh, and the $10/MWh would be collected by the ISO.

These payments result in carbon revenue collections by the ISO. We suggest that the most efficient, simplest and fairest long-term approach to using these revenues is to refund them to load serving entities (LSEs) in an incentive-neutral fashion, most easily in proportion to their load-share ratio. This approach is efficient in that it does not distort the price signal to generation and is the same way market benefits are currently shared across the market. It is fair in that ultimately end-users pay for the costs of lowering carbon emissions through their bills, but these costs would be partially offset by these payments, plus they will further benefit by new entry of cleaner generation that will tend to decrease the market clearing price.[1]

We recognize that our proposed treatment of these revenues will impact the distribution of costs for lowering emissions between utilities (and merchant generators). However, this approach does not inherently distort the distribution of costs that would have been experienced under a state-by-state or source by source approach.

The key advantages of the proposed approach when compared to state-by-state or plant-by-plant implementation of rules on each existing power generation source include:

- It makes use of an existing market system that is well equipped to determine least-cost solutions under constraints. The ISOs already implement least-cost dispatch subject to reliability constraints. Hence, the proposed solution is likely efficient, and in particular, is more efficient than direct control of individual generation sources;
- It takes into account a number of factors impacting the total cost of reducing CO_2 emissions from existing sources, such as constraints on the operation of plants, network congestion issues, the timing and magnitude of previously made investments, etc. Perhaps most importantly, the approach can explicitly reflect reliability contributions of plants.[2]

---

1 There is an important and complicated question related to the impact 100% refunding to LSEs would have on merchant generation and whether any reductions in profits that may result from having to pay for carbon without receiving any refunds would constitute an unfair treatment. An important question that should be asked is what the EPA's authority might be to impose direct costs on these same merchant generators in the absence of the proposed regional ISO-based carbon pricing system. Another important and perhaps even harder to answer question is what investors in merchant generation, at the time of acquisition, would or should have anticipated future regulation of emissions from such generation.

2 Robust mid-term to long-term capacity / reliability constructs may help plants remain economical even if they are required to pay for carbon emissions and/or are dispatched less under our proposed carbon pricing approach, as such plants, if required for reliability purposes, need to be compensated sufficiently through such capacity/reliability mechanisms. This would be true whether or not any plant is under vertically integrated utility or merchant ownership.

- It minimizes the impact on rate payers by fully refunding carbon revenues to end-use customers in a way that maintains appropriate conservation incentives;
- It implements a market-based approach through an independent entity that is perceived as being fair and not profit seeking and thus able to provide a predictable and stable system of carbon prices, which in turn reduces uncertainty for all market participants, with corresponding benefits in terms of lower risk premiums for necessary investments.
- It leaves investment decisions in the hands of market participants, with the likely result of significantly lower total cost for reducing emissions relative to a system of direct control of existing sources.

While the concept is appealing in theory, a number of important design issues will need to be developed. Perhaps most importantly, determining the initial path of carbon prices will be critical, both in terms of assuring the EPA that such a path achieves compliance with the Existing Source Rule, and in terms of achieving an equitable sharing of the costs of compliance between all types of generators and load. It is likely that accomplishing this will require a widely accepted modeling platform and a broad stakeholder process. The same holds true for the conditions under which the carbon price path will be altered and, if so, exactly how to adjust the path.

Other important issues include concerns about: (a) the treatment of entities entering or leaving the proposed approach while continuing to participate in the ISO-based wholesale electricity market, (b) the treatment of entities in a given state that are not a part of the ISO but sell into the ISO-based wholesale market, and (c) various other related "seam-related" issues. While our intention is to develop an approach that captures the efficiency of a regional market-based approach by encouraging broad participation, we prefer to develop an approach that can withstand the disruptions of entities' decisions to participate or not over time and the treatment of entities that choose to comply using alternative approaches.

The rest of this discussion paper describes these and related issues in an overview and sketches some potential solutions. Further, in an appendix, we also discuss why it is important to charge emitting sources rather than, as has been suggested by some, simply applying a "shadow price of carbon" to dispatch decisions.

II. Considerations in Determining the Initial Carbon Price Path

The carbon price path is the central element of the proposed approach. The price path will provide the economic signal for using low emitting resources through re-dispatch and through the entry and retirement of various generation technologies over time. High carbon prices provide a stronger cost-advantage for lower emitting resources and would thus lead to faster emissions reductions from the system than low carbon prices, but fast retirement may also lead to reliability risks if some high carbon emitting generators retire too quickly.

Setting the carbon prices would necessarily incorporate short-term and long-term operational and reliability concerns, while ensuring that the prices are sufficiently high enough to achieve the required emissions reductions over time. Given that the ISOs have the best insights into the resources needed to ensure reliable supply of electricity, the ISO would need to be involved, at minimum, to assist in the analytical efforts. Depending on the level of emissions reductions targeted (and presumably acceptable by the EPA), multiple price paths may exist to achieve the same overall reduction goal. This is particularly true if the existing source rule is translated into a mass-based emissions requirement by a certain year in the future (and therefore not being concerned about emissions prior to that year, or cumulative emissions over the entire time frame). If that is the case, there could be substantial latitude in developing the price paths that achieve the ultimately required emissions reductions. Given the flexibility allowed, in addition to reliability considerations, the carbon price path could also take into account the economic conditions of individual existing generators (for example, whether or not investments have recently been made and are not fully amortized), and the potential impacts on utilities and/or participating states.

A. Different Carbon Price Paths

It is also possible that the EPA might require that emissions be reduced in stages over time, removing some of the discretion over the emissions quantity in the intermediate years. In either case, emission reduction targets can be translated into various carbon price paths, each one expected to achieve the required emissions reductions over time. For example, higher initial carbon prices would create stronger near-term cost increases for higher emitting sources, leading to higher emissions reductions and likely a faster pace of retirements. As a consequence, gradual increases in the carbon price over time may only need to be moderate, if at all, to meet emission targets. Alternatively, a relatively low initial carbon price could be followed by somewhat larger annual increases. It may even be possible not to introduce a carbon price immediately, but rather, phase in a carbon price at a later date closer to the ultimate target year (e.g., starting in 2025 for x% emission reduction by 2030). Under a later start date, it is likely that ultimate carbon prices required to meet the emissions reductions requirement will be higher than in the other cases, but these costs could be anticipated and would impact retirements and new market entrants in advance of the carbon prices. Figure 1 below illustrates some of the potential carbon price paths that may meet an identical ultimate emissions reduction goal.

Figure 1: Illustrative Examples of Carbon Price Paths

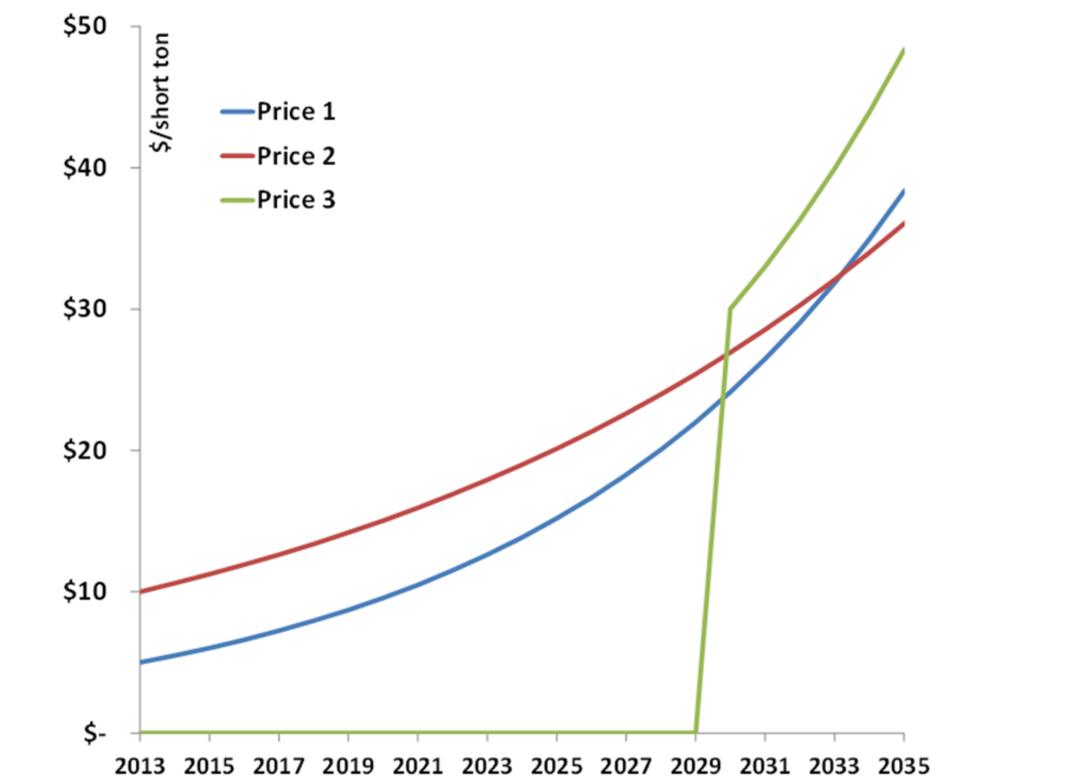

B. Carbon Price Adjustments

Since any initial carbon price path's ability to lead to the targeted (and required) emissions reductions must rely on market simulations that estimate the evolution of the power system over time, including its response to any chosen carbon price path, actual emissions will certainly deviate from simulated forecasts. This means that actual emissions will evolve on a trajectory that is different from the forecast under an initial carbon price path, with actual emissions being higher or lower than initially expected through the forecast simulations.

For the EPA (and market participants) to be confident that the required emissions reductions will be achieved, adjustments to the carbon price path in response to observed differences between actual and expected emissions will be necessary. If stakeholders can reach an upfront agreement, an automatic or formulaic adjustment, administered by an independent party, could be a part of the overall proposal and subject to approval by the EPA as part of the participating states' SIP. Another important element of the price adjustment formula is to create long-term visibility about the carbon price path and to minimize carbon price volatility as much as possible to help create certainty for investment decisions (retirements and new investments), which in turn will likely lead to lower financing costs and consequently to lower overall compliance costs.

It is important to note that the carbon price is not intended to represent a "social cost of carbon." Rather, the price simply represents the economic price signal required to achieve the necessary carbon emissions reductions implementing Section 111(d), over time.

There are many potential ways to adjust the carbon price paths over time. To reduce carbon price volatility, adjustments should be made not of single annual carbon prices, but of the entire carbon price path, and such adjustments should not be done too frequently or in response to very short-term deviations of actual from expected emissions. For example, an adjustment mechanism could be: "adjust the carbon price path once a year and only if actual cumulative emissions to date deviate more than x% from expected cumulative emissions." For each adjustment, both the level and the rate of change of carbon prices could be altered. As with the setting of the initial carbon price path, the choice will depend on numerous factors related to the operations of the power system and the expected impact of a change in carbon prices on costs and reliability of the system.

III. The Approach Functions Even if Not All States in the ISO Participate

Though the desired and the most efficient outcome is if all states within an ISO participate, the approach must be able to withstand changing participation. While it would be less efficient if some states with generation resources in a given ISO do not participate, the proposed approach can still continue to operate if some states decide to meet the emissions reductions obligation of their existing sources by some other means.

There are several possible ways to deal with plants in "non-participating" states in an ISO-administered carbon pricing system. First, we assume that all states will have to comply with Section 111(d) of the CAA by some means and, if the ISO-administered carbon pricing approach is not the choice, another method of compliance is in place through the states' SIP. Given such an assumption, the states that do not participate are incurring some other compliance costs and therefore should not be subjected to paying a carbon price when selling into the ISO-based wholesale market. (The net impact of a SIP in a non-participating state would be that the ISO-administered carbon pricing system is NOT needed to achieve the required emissions reductions for that particular state.) Under such a scenario, if the ISO does not impose a carbon price on plants from non-participating states, such plants will only include a carbon component in their bids if the measures implemented by such a state's SIP result in higher variable costs related to achieving the required state-wide emissions reductions. The ISO will simply dispatch such plants based on the offers received (as these offers will already incorporate the higher variable costs). In cases where a non-participating state's SIP does not impose measures that increase the variable cost of such plants, these plants will tend to be dispatched more relative to essentially identical ones (in terms of cost and fuel) participating in the ISO-administered carbon pricing system. To the extent that the locational marginal price ("LMP") at their location is increased by the ISO-administrated carbon pricing system, they may also receive higher energy prices (and therefore higher energy margins) than they would receive in the absence of an ISO-based carbon price.

We believe that there are two options for compensating plants from non-participating states: they can simply be paid the LMP, or they can be paid the greater of their bid or the LMP without the carbon price component, since for any bid-based central dispatch system to be efficient, bidders must be assured to receive at least their bid. As Figure 2 below illustrates, if plants from non-participating states receive only the LMP minus its carbon price component, they may be paid less than their bids when they are dispatched. Thus, it may be important to pay the greater of the generator's bid and the LMP without the carbon price component. However, to use such a mechanism, the strategy of the non-participating generator may be to increase its bids up to close to that of the LMP with the carbon price component but low enough to be dispatched. If this is the dominant strategy, the system is similar to simply paying the non-participating generator the LMP including the carbon price component.

In principle, there are three options for demand from non-participating states to take part in the ISO market. The first two options are that they can be required to pay the LMP (or applicable zonal price) and receive a carbon refund just like demand located in participating states, or they can pay the LMP (or zonal price) net of the carbon price component but not receive a carbon refund. (It is likely that, depending on how prices paid by demand are calculated relative to LMPs, the two options could result in the same net payments by load.) In either case, such payments will be somewhat higher than if there were no ISO-based carbon price system, since even with a full refund of carbon revenues to load as proposed, production costs (and therefore average prices net of the carbon component) will increase due to the re-dispatch towards more expensive (but less carbon intensive) resources that will be the result of the proposed approach. The third option is to charge non-participating load the LMP with no refund. This is an interesting option because it may provide incentives for states to participate.

On balance, generators from non-participating states will receive somewhat higher revenues (and higher margins) when compared to participating generators (but on average will likely incur greater compliance costs through non-ISO compliance at the state or individual source level), whereas demand will incur somewhat higher costs. For utilities with both sources and demand in the ISO, the net effect of the two will depend on the specific mix of resources and demand. Given the likely significant efficiency advantage of the proposed approach over non-regional alternatives, we also do not believe that any incremental benefits for non-participating states would actually create an incentive NOT to participate in the ISO-administered carbon pricing system proposed. However, we need to explore this further.

Figure 2: Pricing Options When Some Plants Are From Non-Participating States

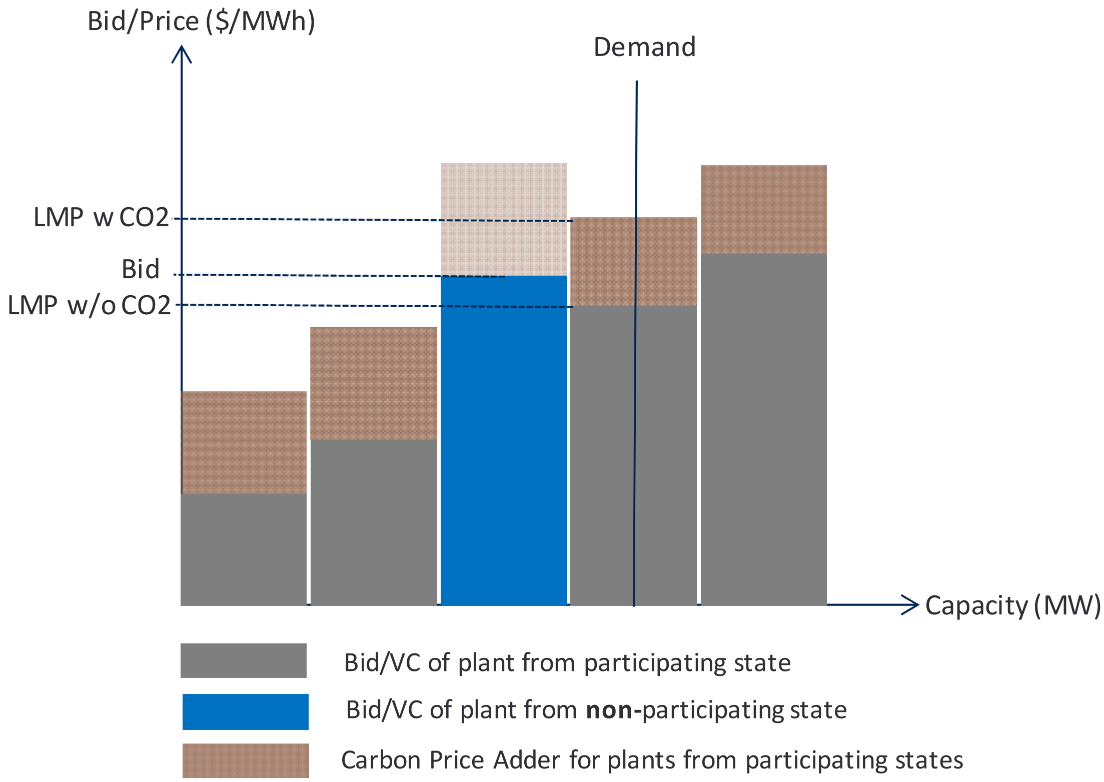

In summary, we believe that an approach along the lines outlined above assures the feasibility and viability of the proposed ISO-administered carbon pricing system even if some states with generation and demand in the ISO decide not to participate in such a system. We further believe that such approaches would not create sufficient incentives for individual states to *not* participate, but rather that offering participation in the approach would be very attractive, although we cannot exclude the possibility that some individual states will come to a different conclusion.

IV.The Relationship between 111(b) and 111(d) Compliance

Another important issue is whether it might be unfair to charge new sources compliant under Section 111(b) a carbon price, if they are part of an ISO implementing the proposed approach. We believe that the answer is no.

First, we believe it is important that new sources be subject to the same carbon pricing regime as existing sources. This ensures efficient dispatch and minimizes the cost of emissions reductions, from which all consumers benefit. Unlike with non-participating states, new sources are not by definition carbon-free (in the sense of emissions reductions being achieved outside the ISO system). They will have lower emissions than at least some existing sources, at least until these existing sources have retired. However, they may have higher emissions and/or higher cost, with or without consideration of the marginal cost of an incremental unit of emissions reductions (embodied in the carbon price), than other existing resources. A simple illustration demonstrates this point: if new and existing sources are treated differently, then a generator constructed the day before 111(b)/111(d) go into effect will be treated differently from a generator being constructed the day after. If the former were subject to a carbon payment while the latter is not, the latter would be dispatched before the former as long as the variable cost without carbon of the new source is lower than the variable cost with carbon of the existing source. As a consequence, a less efficient new combined cycle gas plant might be dispatched before a more efficient combined cycle gas plant, a dispatch that will lead to higher total cost of meeting demand, including higher costs of meeting emission reduction requirements.

Perhaps most important is a recognition that in a regional approach that imposes either a cap or a price on power plants participating in a jointly dispatched electric system, overall cost minimization (and hence efficiency) implies that all generators, whether new or existing, have to be treated equally. In a price-based system, this implies charging every generator a carbon price. In a cap-and-trade system, it implies requiring every generator to hold allowances equal to emissions.[3]

With this recognition, there are two issues that require a solution: whether, and if so, how, to reflect new source additions in the regional cap/target/price; and whether or not some form of compensation is required for new sources subject to a regional implementation of Section 111(d).

With respect to the former, we suggest that a solution that would increase the cap over time to account for load growth may be desirable and in line with the emission intensity concept traditionally associated with emissions regulation of both new and existing sources. While it would be possible to agree to an overall absolute cap independent of whether or not load growth over time will require significant additional capacity, doing so might, at least in the presence of

---

[3] In this context, it is important to note that under the existing SO_2 trading program under the CAA, new sources of emissions are required to purchase and hold allowances to cover their emissions.

actual load growth, ultimately require new sources to have emissions intensity below those required under a new source rule. Therefore, absent using the cap to implement some other policy requirement to limit absolute emissions, there would seem to be at least some risk that an absolute emissions cap covering both existing and new generation would not be compatible with the Clean Air Act. It could however be argued that, by definition, existing sources serve existing load and existing load only. Emissions from existing sources could therefore be legitimately capped at the regional level, implying lower total emissions over time to meet existing levels of demand. Incremental demand, on the other hand, would need to be met with new resources, which should be subject to the new source rule. Assuming an allowed maximum emissions intensity of new sources, we therefore propose to impute the additional emissions that would result from serving incremental demand from new sources meeting the Section 111(b) requirement, and to adjusting a regional cap over time as a function of the additional emissions thus be calculated. The net result of such an approach is illustrated in Figure 3 below.

A second question concerns the issue of whether or not new sources would or should have to be compensated for being required to participate in an (ISO-based) regional compliance system designed primarily to deal with existing sources, even they are already compliant with Section 111(b). Our basic conclusion is that, beyond the cap adjustment discussed above, no further compensation is necessary. The argument for this proposal is two-fold:

First, new generation enters the market if and when such entry is expected to be profitable. Once both Section 111(b) and 111(d) are implemented through SIPs, the resulting changes to the market environment will be known to potential entrants. No new plant will enter if it doesn't assume to be able to operate profitably. In pure economic theory, over the long-term, just enough plants will enter the market so that new plants exactly break even in economic terms, including an appropriate compensation for investors. This is true whatever the market environment, and therefore whether or not new sources will have to pay a carbon price or hold allowances in addition to meeting emissions standards under Section 111(b). More to the point, if for some reason entry by new plants were more profitable when the new plants do not need to pay a carbon price or hold allowances, one would expect more plants to enter the market until any such additional profits were competed away. Hence, at least conceptually, new plants cannot be harmed by being required to fully participate in a proposed regional system since they will, on average and in expectation, just break even in either case.

Figure 3

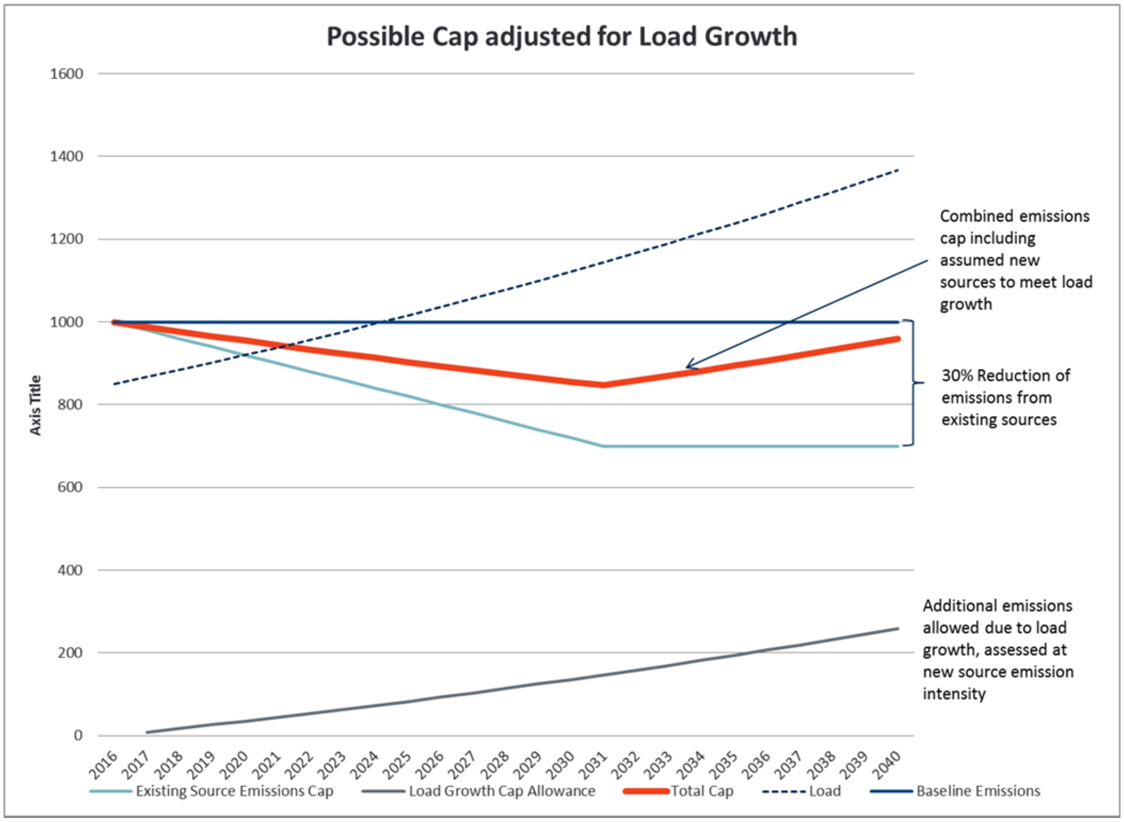

Source: Brattle Analysis

Second, and perhaps more practical, it is actually quite likely that new sources will benefit more from being part of a carbon-pricing or other market-based implementation of Section 111(d) than being part of a system without explicit or implicit carbon pricing, and therefore there is no need to compensate them for being worse off. The intuition for this is that new sources will likely be more efficient than existing sources of the same technology. They will also, on average, have lower CO_2 emissions per unit of output. Both factors will tend to favor them in the dispatch relative to existing units **once carbon prices are added**. This in turn means that they will: a) be dispatched more often and b) benefit from a market price that increases more than their carbon payments, as illustrated above. Hence, if anything, a regional carbon cap, implemented through either allowances or a carbon price applied to all generators, will likely create additional incentives for new plants to enter the market without negative financial impact on such new entrants. Consequently, there is no reason to provide additional compensation for such plants.

V. Imports and Exports

One of the important questions associated with any method of compliance is how to treat imports from and exports to exporting region. It appears that the SIPs will need to cover generating plants physically located in the complying states. Because states have permitting authority over the plants by geographic location, it would seem that the emissions should be dealt with at the exporting region. (However, it may be possible for an exporting resource to cede authority to the destination region if the destination region a compliant region operated by another ISO.) One would also need to ensure that self-scheduled resources would be subject to some form of control as well. In addition, how resources with legacy contracts would comply will need to be resolved. The legacy contract question is likely important since, if there is a loophole, existing generators may have incentives to sign long-term contracts before the approach is implemented.

There are a few important principles that need to guide the treatment of imports and exports. In the absence of transmission constraints, economic imports and exports should equalize prices across regions, i.e., imports and exports between two regions should lead to the two regions having a dispatch that is optimal for the combined region. A second principle is that carbon reductions should occur where the value of carbon reductions is largest, which in turn is reflected in a higher carbon price. Ultimately, power is generated to meet end user demand and the desire is to reduce the carbon emissions associated with meeting end user demand. A higher carbon price indicates a higher value and higher (expected) cost of lowering CO_2 emissions, so all else equal, imports of reduced emissions (in the form of power imports with lower than average emissions) would be desired by all regions and should be encouraged.

A simple conceptual model helps understand how these ideas should be reflected in dealing with imports/exports in an ISO-administered carbon pricing system.

Assume there are two adjacent regions. Region 1 has a relatively low cost generating resource mix composed of low cost (low-emitting) coal plants and mostly gas CTs, while Region 2 is composed of older and less efficient coal plants and a number of relatively efficient gas CCs. Average energy prices in Region 1 (absent carbon) would be relatively low, while those in Region 2 would be relatively high.

Absent imports/exports, the price in Region 1 would be lower than in Region 2. With imports/exports, one should therefore expect exports from Region 1 into Region 2.

Given that Region 1 has a large variable cost spread between high and low emitting resources, the carbon price that would be needed for a reduction in emissions is relatively high. By contrast, the cost to achieve emissions reductions in Region 2 would relatively modest, given the lower variable cost spread between the less efficient coal plants and the gas CCs. Consequently, achieving the same emissions reductions in Region 2 could be accomplished at a lower carbon price.

The following example shows this in a simple graph. Assume Region 1 has generators A, B, C and D and Region 2 has generators E and F. If the two regions are disconnected from each other, in the absence of carbon prices, generator C sets the market price at $45/MWh in Region 1, and generator E sets the market price in Region 2 at $25/MWh. Since market prices are higher in Region 1 than in Region 2 (due to a higher priced gas fired CTs setting the market price), an efficient export would sell power from generator E in Region 2 to Region 1, which would displace some of the generation from generator C (which has a higher variable cost than generator E). The resulting generation pattern is identical to what would happen if the two regions were dispatched jointly. This is illustrated in Figure 4 below.

Figure 4: Efficient Trade Without a Carbon Price

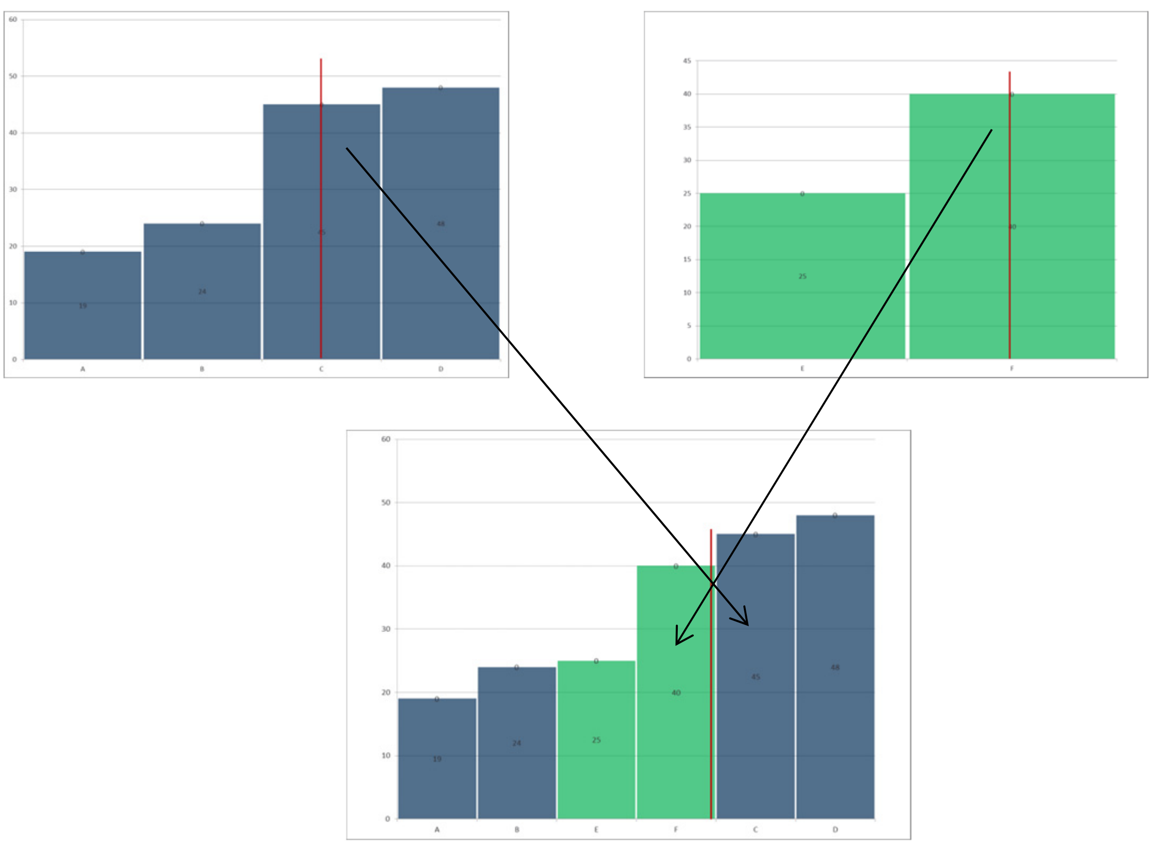

Now assume that in Region 1, it is possible to reduce carbon emissions to the required levels at a carbon price of $60/ton, but that the required emissions reductions in Region 2 require a carbon price of $30/ton. These carbon prices will increase market prices in different ways in the two markets.

Figure 5: Joint Dispatch of Regions 1 and 2 with Exporting Region Carbon Prices Applied to Exports

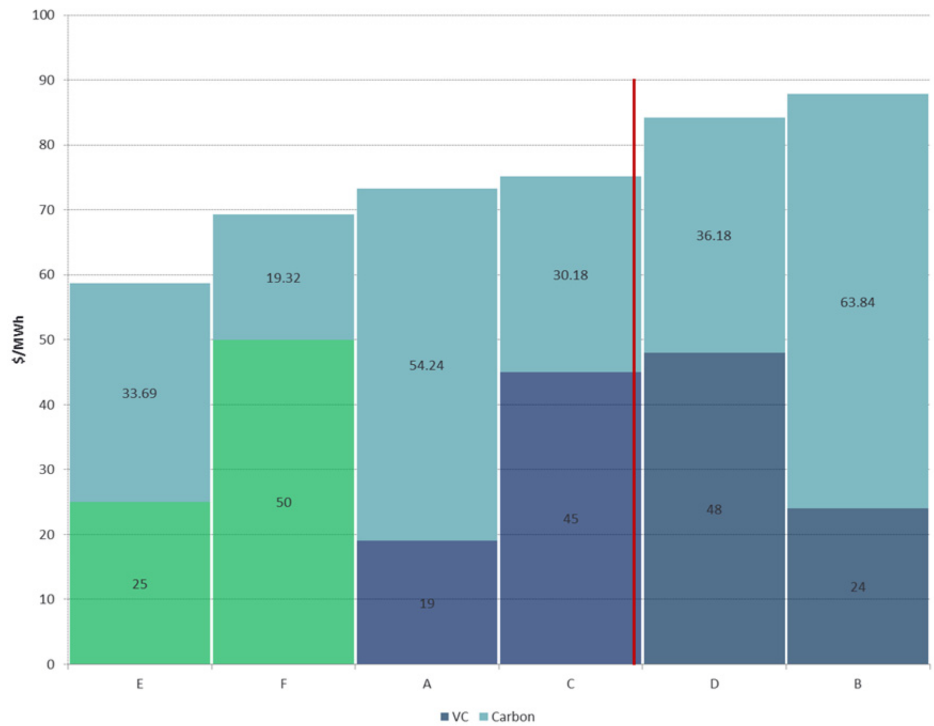

Figure 5 illustrates why charging exports the carbon price of the exporting region rather than the importing region leads to inefficient outcomes. In a joint dispatch of Regions 1 and 2, where the carbon prices of Region 2 are applied to the generators located in Region 2 ($30/ton rather than the $60/ton in Region 1), generators E and F, both located in Region 2, get fully dispatched. This is the same as saying that it would be economical, given the incentives, for generators E and F to export to Region 1. They become the lowest cost resources, even though both their non-carbon variable costs and their emissions intensity are higher than those of the corresponding Region 1's generators B and D, which get pushed out of the dispatch.

Figure 6 below shows the outcome if generators E and F were charged the carbon price in the importing region (Region 1), should they export power into Region 1.

Figure 6: Joint Dispatch of Regions 1 and 2 with Importing State Carbon Prices for Exports

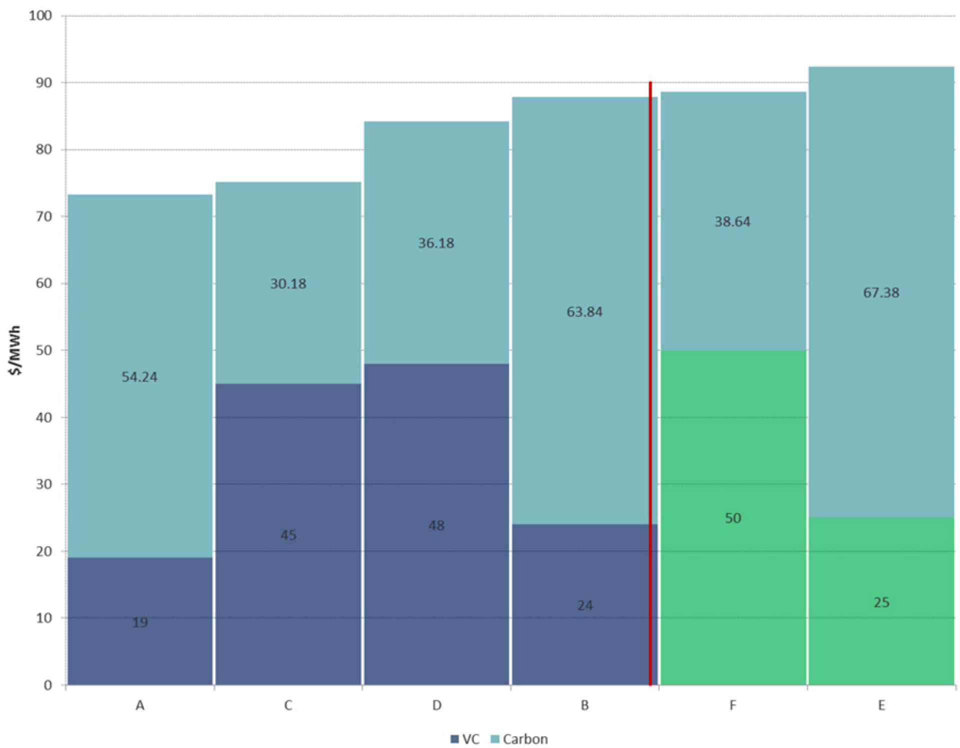

As seen above, in the case the resulting joint dispatch is efficient – lower variable cost and lower emissions intensive generation is dispatched to meet demand and there is no rationale for exporting power from Region 2 to Region 1.

Consequently, as can be seen in the graphs above, for incentives for exports and imports to be efficient, the carbon price of the importing region has to be the basis for charging for CO_2 emissions. This rule is in essence equivalent to the "first buyer" approach used in California under AB 32; and we therefore suggest that the methods used under AB 32 could provide a good basis for dealing with imports and exports under our proposed approach. Since unlike in California both importing and exporting regions will have to comply with Section 111(d) of the CAA, this also implies that the exporting region should not charge for CO_2 emissions.

An issue not addressed by the example above is how to calculate emissions if the imported power cannot be tied to a specific generation resource. This issue has also been addressed in AB 32's approach to charge the first deliverer of imported electricity for associated emissions.[4] Using the first deliverer creates regulatory oversight by in-state regulators. Also, the first deliverer approach allows clear identification of the importing entity (NERC E-tag, facility operator or scheduler). Mandatory Reporting Regulation 95111 sets default emission rates for unspecified imports, i.e., imports that cannot be traced specifically to a certain power generation facility. Using such default emissions factors allows treatment of non-unit specific exports, such as system sales. California recognized that resource shuffling, i.e., assigning imports to specific, low-emitting resources even though those resources may not actually be delivering the power, could create serious complexities. These issues arise in particular if imports are from regions without corresponding CO_2 regulations, which would not be the case under 111(d). However, the same issue could apply, to a lesser degree, in adjacent regions with different compliance approaches (or carbon prices) to Section 111(d). We propose to base the treatment of imports from adjacent regions on the approach developed in California, incorporating the experience gathered to-date, in particular as concerns default emissions rates from unspecified sources and related incentives for resource shuffling.[5]

VI.Shadow Price Approach vs. Actual Carbon Payments

While appealing in concept, a system of shadow prices only would likely not lead to the desired emissions reductions, since suppliers would have an incentive to distort their bids to the ISO. Current wholesale markets are based on the notion that all participating generations generally have no incentive to make offers to the market that deviate from their variable costs, a pre-condition for markets to deliver socially optimal outcomes. When bids represent variable (or marginal) costs, the ISO can dispatch the system to minimize total costs to consumers. In all ISO markets, there is a single (hourly) set of market prices (which may differ by location), and all

---

4 For a discussion, see http://www.arb.ca.gov/cc/capandtrade/meetings/050412/may4electricityppt.pdf.

5 There is some indication that existing default emissions rates may lead to poor incentives. See, for example, http://www.cacurrent.com/storyDisplay.php?sid=7214 for a discussion. At present, the default emission rate is equivalent to emissions from a relatively efficient gas-fired power plant. However, when imports are from systems with significant coal-fired generation, this can underestimate the emissions "contained" in unspecified imports. In response to this and related issues, CARB is in the process of revising some of its rules. For a summary, see http://www.vnf.com/1914.

market participants pay/receive that same market price. If an emitter has to pay for carbon emissions, then marginal costs increase correspondingly and a generator is only dispatched if its costs (including carbon) are lower than those of other generators that might be dispatched to meet demand. The right carbon prices will lead to switches in the dispatch, for example, by increasing the dispatch of natural gas fired generation relative to coal-fired generation.

A simple example illustrates how, without charging generators for their emissions, this switching would likely not take place. Let's assume there are only two generators, a coal- and a gas-fired generation, both of equal size. The variable cost of the coal-fired generator is lower than the variable cost of the gas-fired generator without a carbon price, but the same coal generator's variable cost would be higher than the gas-fired generator with a carbon price. Finally, let's assume that demand is such that 100% of the production from one of the generators will be needed, but only 50% of the other.

Before implementing any GHG emissions system, the coal-fired generator would be used 100% and the gas-fired generator 50%. The more expensive gas-fired generator's marginal cost (no carbon cost) would set the market price. The gas-fired generator would just earn its marginal cost, the coal-fired generator would earn its variable cost plus an additional energy margin (based on the clearing price set by the gas generator). With a carbon value paid by all emitters, the role of the two generators would be reversed. The coal-fired generator would set the market price and earn just enough revenues to cover marginal fuel and carbon costs. The lower-carbon gas-fired generator would earn an additional margin. This is the efficient outcome and would reduce carbon emissions, by reducing the generation from higher-carbon intense coal generation in favor of lower-carbon intense gas generation.

Now assume that the ISO does not charge for carbon, but uses the same carbon value to guide dispatch decisions. Assuming both generators bid their variable cost, the outcome would be the same in terms of dispatch (gas=100%, coal=50%), but not in terms of prices and margins. Now, the price would still be set by the most expensive generator on the system (which would be gas), so the gas-fired generator will earn the same margin (none, since price is equal to its marginal cost, even though it produces more), and coal-fired margins would be reduced as a consequence of it being only partially dispatched. This still looks appealing.

However, the coal-fired generator can do better. Assume that the carbon value is $10/ton, which, at an emissions rate of 1 ton/MWh for coal and 0.5 tons/MWh for gas, would translate into a cost disadvantage of coal relative to gas of $5/MWh. Since the market price does not include the cost of carbon and is set by the most expensive generator running, the coal-fired generator has an incentive to reduce its bid price by $5.01/MWh (tell the ISO that its marginal cost is really a bit more than $5/MWh lower than it really is). If it does so, the ISO will conclude that even with the carbon value used as a shadow price in dispatch the coal-fired unit seems cheaper than the gas-fired unit, so it will dispatch the coal-fired unit at 100% and the gas-fired unit at 50%. The gas-fired unit is still the most expensive unit and sets the market price. In short, the situation is identical to the pre-carbon regulation system. The coal-fired generator can underbid and essentially undo the effect of the carbon value.

VII. Conclusions

In this discussion paper, we have outlined a basic regional approach as a compliance option under Section 111(d) of the CAA. The simple idea is to determine a regional carbon price path that, if used in dispatching regional resources, would lead to regional emissions reductions sufficient to meet the emissions reductions requirements under the CAA, and to use this path as an alternative compliance mechanism. We believe that since it is a market rather than a plant-by-plant or state-by-state approach, a regional carbon price (path) will be more efficient than other compliance approaches, and hence, lower the average cost to consumers of meeting the requirements of Section 111(d). We also believe that refunding the carbon revenues that would be collected from CO_2 emitting generators (carbon price multiplied by generation) to load serving entities will be the fairest and simplest approach, further mitigating any emissions-reductions cost to ratepayers. We are aware of the possibility that using this approach may, in some cases, lead to the potential for sudden rate increases (or other financial impacts on any participating state's energy sector). We propose to assess the potential magnitude of any such impacts prior to deviating from our proposal of full refunding to LSEs at the outset, but in all cases suggest that full refunding to LSEs be both the goal and at the very least the end result of a transition should the potential be found to be significant..

We have outlined various approaches to deal with the kinds of issues that arise when different regions or states use different market mechanisms, for power generation and or for meeting regulatory requirements. Dealing with these "seams" issues will be important, especially since not all states will have all power generation participate in ISOs. The proposed approaches are likely one, but not necessarily the only feasible mechanisms for dealing with seams (and related) issues. We encourage readers to critically evaluate our proposals and, where necessary, suggest flaws, alternatives or improvements. Ultimately, detailed implementation of the approach we propose will require effort well beyond the current development of basic principles of the regional approach we propose.

CAMBRIDGE
NEW YORK
SAN FRANCISCO
WASHINGTON
LONDON
MADRID
ROME

THE Brattle GROUP