PJM Credit and Clearing Analysis Project
Findings and Recommendations

June 2008

MARKET REFORM

Table of Contents

1 STUDY BACKGROUND AND PURPOSE

PJM operates real-time and day-ahead-hourly cash markets[1] in electrical energy, as well as cash markets in related products (e.g. regulation, synchronized reserve). It also operates forward markets in certain products related to its cash markets, notably in Financial Transmission Rights (FTRs). One of PJM's key responsibilities is to administer the settlement of these markets, and to preserve the financial integrity of the settlement scheme through effective credit risk management. As such, PJM acts as the clearing house for all trade in its markets, in function if not in name.

PJM carries out these functions on behalf of all the participants in its markets, who ultimately bear the risk of default through the socialized guarantee they provide. "Because participants do not have any knowledge of other parties' spot exposures, and no way to manage these risks bilaterally, potential losses are both unpredictable and unhedgeable. As a result, participants are totally dependent upon the credit practices adopted by the (Market Operator)."[2]

When compared to other organized (i.e. centrally-operated) commodity markets, the credit practices utilized in most electricity cash markets – and it should be stressed that PJM is by no means unusual here – are deficient. In most cases the policies and processes currently in place – with long settlement timeframes and liberal unsecured credit – are an artifact of historical practice, having evolved in a piecemeal and often perfunctory manner from the days of wholesale settlement between regulated monopoly utilities. In a competitive wholesale market, with a diverse set of participants and where price is free to float, the risk profile is decidedly different, and must be managed accordingly.

The aim of this exercise was to examine opportunities to improve how PJM manages credit risk. In this context, the term 'improve' was tightly interpreted to mean:

- reduce the probability of participant default
- reduce the magnitude of default
- reduce the unpredictability (in timing and quantity) of charges received as a result of default.

This involved examination of issues of structure, policy, process and assessment methodology, many of which are intertwined.

This report is set out as follows:

- Section 2 provides a summary of findings and recommendations contained in the document.
- Section 3 addresses a select number of fundamental structure and policy issues.
- Section 4 examines key processes which should be present in a robust, centrally-managed, credit risk management mechanism, for both cash and forward markets.
- Section 5 examines opportunities for improving the management of defaults within PJM's markets, if and when they occur.
- Sections 6 and 7 discuss process and calculation methodology options for risk assessment in PJM's forward and cash markets respectively.
- Section 8 discusses options for the utilization of an external clearing house, to replace PJM as the entity responsible for all credit functions and a number of settlement functions.

[1] The term 'spot market' is also frequently used, though strictly speaking the only market that is 'spot', resulting in actual physical delivery, is the real-time market.
[2] Todd W. Bessemer and Francis X. Shields, *Spot Market Clearing: Solving the Electricity Credit Malaise*, Public Utilities Fortnightly, May 2005.

2 SUMMARY OF FINDINGS & RECOMMENDATIONS

Contained within the narrative of this report are a number of findings, recommendations and issues for further investigation, the most important of which are summarized below, by topic (and not necessarily in the same order as the remainder of this report). For those not well versed in these topics, however, many of these findings and recommendations may be difficult to follow in isolation, and it is highly recommended that they be read in the context of the discussion accompanying them, in the main body of this document.

Central Counter-Party

- ***Recommendation***: PJM should utilize a central counter-party structure for settlement and credit management.

Market Participation

- ***Recommendation***: PJM should establish minimum criteria for participation in its markets, to be assessed on initial registration. PJM should have a right to verify ongoing compliance with these criteria, at its discretion.

Credit Cover/Collateral – Policy

- ***Recommendation***: All positions should be 'fully collateralized', to cover potential exposures to a high degree of statistical confidence. Unsecured credit should not be extended to any participant.

Liquidation

- ***Recommendation***: In the event of participant default, PJM should have the right to liquidate any forward contracts not yet in delivery.

FTR Market Structure

- ***Recommendation***: To support liquidation, as the delivery period approaches, longer-term contracts (e.g. an annual FTR) should decompose into strips of the next most granular contract (e.g. a strip of monthly FTRs).
- ***Investigate***: the feasibility of shorter auction periods, and FTR products of shorter tenor.
- ***Investigate***: whether a reduced auction that only eliminates FTR options is feasible to run on a daily basis, and confirm it is viable for a weekly basis.
- ***Recommendation***: Pending the investigation into performance of reduced auction models, move to the weekly auctioning of FTRs. Re-examine the feasibility of daily auctions within two years of this move.

Reassignment of Settlement Obligation (eSchedules)

- ***Recommendation***: In the event of default, PJM should have the right to immediately terminate the forward portion of any eSchedules which serve to increase the defaulting participant's settlement liabilities. Any eSchedule commencing in the future should be immediately rejected.

- ***Recommendation***: Perform credit validation on any transaction to cancel an eSchedule. If the validation causes either party's exposure to exceed their credit cover, reject the transaction.

Trading/Submission Limits

- ***Recommendation***: Apply credit risk limits to all day-ahead market trading, FTR trading, eSchedule submissions and any other 'controllable' market activity.
- ***Recommendation***: Ideally, use trading limits rather than bid limits. Where this is not practicable, institute a process to assess required collateral soon after an auction, and release unused collateral (or provide a streamlined process by which participants can request its release).

Forward Market Settlement

- ***Finding***: Variation Margining is of limited utility if auctions only occur infrequently (e.g. once a month), as circumstances can alter significantly between opportunities for incremental price adjustment.
- ***Recommendation***: The settlement and billing of variation margin should be de-coupled from the cash market billing cycle, with bills issued as soon as variation amounts are calculated, and payable within one (1) business day, or before the next auction, whichever is sooner.
- ***Finding***: If Variation Margining is implemented, an imbalance account, held by PJM, will be required to manage overs/unders until the FTR is delivered in the Day-Ahead Market.

Cash Market Settlement

- ***Recommendation***: Reduce the cash market billing period to a duration of one (1) week or less.
- ***Recommendation***: Reduce the lag between the end of the billing period and payment to five (5) calendar days or less.

Credit Cover/Collateral – General

- ***Recommendation***: Set the Minimum Credit Cover Requirement for a participant equal to Actual Exposures + Resettlement Exposures + historically-predictable Future Potential Exposures + any controllable Future Potential Exposures already accepted (e.g. future-dated component of submitted eSchedules).
- ***Recommendation***: Set the Working Credit Cover Requirement equal to 110% of the Minimum Credit Cover Requirement. This is the level to which credit cover must be restored in the event of a margin call, and ensures that uncommitted credit cover is available for additional controllable exposures.
- ***Recommendation***: Validate any new controllable exposures, using bid/trading limits, to ensure they do not exceed the participant's uncommitted credit cover.
- ***Recommendation***: Provide each participant with an estimate of their Maximum Predicted Exposure, as part of the daily credit process.
- ***Recommendation***: For a terminating participant, continue to hold credit cover for resettlement exposures, until the resettlement window for their trading activity has closed.

Credit Cover/Collateral – Forward Markets

- ***Recommendation***: Assuming the use of both variation margining and liquidation, set initial margin requirements to cover the x% worst move over two auction periods.
- ***Recommendation***: Simulation should be used for calculating the potential exposure (and hence initial margin) of FTR positions, provided performance issues can be addressed satisfactorily.
- ***Recommendation***: Determine delivery margin requirements for FTRs based on the auction value of the FTR portfolio, as marked-to-market, minus the x% worst-case payout of the FTR portfolio (which could be negative).
- ***Recommendation***: The Simulated Scenario approach should be used for determining the Worst Projected FTR Payout component of delivery margin requirements - predicated upon a simulation approach also being followed for initial margin calculation.
- ***Recommendation***: Require 50% of estimated delivery margin requirements to be posted prior to the penultimate auction before delivery, with the remaining 50% posted prior to the final auction before delivery, and true-up after the final auction based on the final settlement price. Initial margin held against these positions should offset delivery margin requirements.
- ***Recommendation***: Adjust delivery margin requirements during the delivery period, to take account of those days that are 'delivered', and for which the FTR Payout (and Auction Payment) form part of the Actual Exposures in the Cash Market.

Credit Cover/Collateral – Cash Markets

- ***Recommendation***: Calculate Future Potential Exposure as the worst contiguous n-day exposure that the participant might incur over the future exposure period (n), to x% likelihood.
- ***Recommendation***: Perform an indicative settlement calculation the morning after the trading day, using approximated data as required.
- ***Recommendation***: Re-evaluate participant exposures, and associated credit cover requirements every business day, and possibly every calendar day.
- ***Recommendation***: Require the transfer of load responsibility to a Provider of Last Resort to be completed within a maximum of two (2) business days.
- ***Recommendation***: Calculate resettlement exposures based on historical trends for the participant, where available, or for the market otherwise.
- ***Investigate***: the possibility of shortening the resettlement window for some or all charges and payments.

Triggering a Default

- ***Recommendation***: Any failure to pay, or to post required credit cover, should result in an immediate 'margin call'. If the margin call is not satisfied within 24 hours, default proceedings should commence.
- ***Recommendation***: Provide formal warnings to a participant when its aggregate exposure reaches 80% of credit limit, and again when it reaches 90%.

Collateral Standards

- ***Recommendation***: PJM should continue to employ standards which all LCs must comply with, and stipulate them in a pro-forma LC. These standards should include that the LC is payable on demand, and payable the same day.

- ***Recommendation***: PJM should continue to maintain minimum financial standards (including required credit ratings) for LC providers. A list of accredited LC providers should be maintained and regularly reviewed, with participants able to petition to have additional providers added.

Default Structures

- ***Recommendation***: PJM should maintain a line-of-credit, or similar facility, to provide bridging finance for near-term shortfalls (to a reasonable limit).
- ***Recommendation***: PJM should explore the creation of an escalating guarantee structure, including residual funds held by PJM, a guarantee fund contributed to by the members, and default insurance. Socialization of default to members should be the last step in this structure.

3 MARKET-WIDE STRUCTURE AND POLICY ISSUES

3.1 Structural Issue – PJM as Central Counter-Party

One of the key settlement and credit issues presently facing PJM is whether it should become the central counter-party to all trade. PJM has always stated the position that it does not take title to the electricity that trades through its markets. PJM is also in the practice, however, of netting participant positions. For example, if a participant is owed $8,000 for its energy sales in the day-ahead market, and owes $10,000 for its purchases, it is considered to owe a net $2,000 to PJM. The participant's credit cover requirement (collateral, plus unsecured credit if applicable) is calculated based upon the projected risk of its *net* position.

Recent legal advice[3] raises serious doubts concerning whether, under current structures, such netting could be subject to challenge in a participant bankruptcy proceeding (the issue principally relates to whether 'mutuality' can exist between the defaulting party and PJM, if PJM is not a party in the transaction[4]). Ironically, bankruptcy is precisely when credit arrangements need to hold, as it and participant default generally go hand-in-hand. If netting was disallowed, this would mean, using the previous example, that in a bankruptcy proceeding PJM would need to payout the $8000 it owes, and would have a bankruptcy claim against the $10,000. The collateral required to secure the credit risk of these gross purchases would be substantially greater than that for the net.

Other than sticking with the status quo, the only other viable option is to utilize a central, or 'universal', counter-party (CCP) structure. Under this approach, all transactions are executed with a central entity, which becomes the buyer to all sellers and seller to all buyers. As the CCP holds one side of every contract, the required 'mutuality' exists, and it is able to facilitate multilateral netting of participant positions. While the CCP does take title, this does not present any inherent issues of conflict, as it always maintains a net zero position, and therefore has no incentive to act in anything other than an independent and price-neutral fashion. The multi-lateral nature of this structure is well-suited to PJM's markets, many of which execute through a multi-lateral auction process, where it is not possible to associate a single buyer with a single seller and establish a bilateral credit relationship, even if minded to do so.[5]

None of this is meant to imply that such a move is trivial. Advice from PJM counsel indicates that issues related to CFTC registration (or obtaining a 'no action' letter), tax treatment and financial statements may need to be addressed. Some participants have also flagged potential tax implications. On the other hand, these structures are hardly new. The use of a central counter-party is standard practice in all major commodity futures markets (including those in energy). A number of electricity cash markets around the world also use this model (e.g. Elexon, Nord Pool).

Used effectively, a central counter-party mechanism can serve to reduce risk and create capital efficiencies. Key to this effectiveness, though, is the application of sound risk management policies and processes. As written by Mark Twain in *Pudd'n'head Wilson*, "The wise man saith, 'Put all your eggs in the one basket and -- watch that basket!'"

***Recommendation**: PJM should utilize a central counter-party structure for settlement and credit management.*

[3] Memorandum dated March 17, 2008, from Novikoff et.al. (of Wachtell, Lipton, Rosen & Katz) to Vincent Duane, PJM General Counsel.
[4] This paper will not seek to provide further explanation of the rationale for this issue, as this has been set out elsewhere in greater detail and precision by legal advisors expert in this field. See Novikoff, *op.cit*.
[5] Ironically, most futures are traded in a continuous auction process, where it is relatively easy to identify a single buyer and seller for a given transaction, but are cleared using a central counter-party mechanism because of the credit and settlement benefits it provides.

3.2 Policy Issue – Collateralization

Probably the most fundamental policy question that needs to be addressed in any credit risk management scheme is whether to extend unsecured credit.

In a 'fully collateralized' regime, all positions must be secured by collateral in accordance with the risk they present. The intent is to ensure that potential exposures – determined to a given (high) degree of statistical confidence – are covered by cash or a liquid instrument rapidly redeemable for cash, such as treasury bills or an appropriately formulated letter-of-credit (LC).

In a regime that permits unsecured credit, certain participants are extended credit without the need to post collateral, generally based upon an assessment of their financial strength[6]. Parent guarantees might be utilized, but these are simply a displacement of risk from a subsidiary onto a larger (and presumably financially stronger) parent entity. The extension of unsecured credit to a participant involves an assumption, on behalf of the market, that the entity will not fail. Managing this risk is inherently more subjective than managing the risk of a position. Credit assessments performed by the market operator are unaware of the participant's transactions with other entities and exposures in other markets. Substituting for this information are bona fides provided by the company itself, or by external parties such as rating agencies. While the failure of an investment-grade company is an unlikely event, it does happen[7], and is certainly an event within reasonable statistical bounds; remember that the goal of risk management is to guard against what might happen, not to substitute the belief that nothing will. Further, when such events do occur, they usually happen quickly. Calls for collateral at such a juncture are often too-little, too-late.

Most electricity market operators are insulated from default[8] by a guarantee that allows them to socialize any shortfall. The providers of this guarantee – the remaining market participants[9] – ultimately bear this credit risk. The multilateral nature of these markets (coupled with the confidentiality of participant data) makes it virtually impossible for market participants to do anything of their own accord to predict this risk, or hedge against it.[10]

Clearly, full collateralization reduces market-wide credit risk, and hence directly decreases each participant's individual credit exposure as a member of the socialized risk pool. Considered solely from a credit risk point-of-view, the case for eliminating unsecured credit is compelling. However, full collateralization would require collateral postings from a number of participants who are in the habit of posting little or no collateral at present. Under the current settlement timeframe, these requirements would be substantial – equivalent to around 60 days' trading activity in PJM at present (with monthly billing, payment ~22 days in arrears, 3 days to remedy, etc.) – something that sub-investment-grade entities in these markets are already acutely aware of. Obviously this collateral has a cost – though ironically, if settlement timeframe is accelerated (a topic discussed in later sections), the market-wide cost of collateral would be substantially reduced from present levels, even though it would increase for some individuals.

[6] Though more unusual arrangements have existed. For example, in the initial top-up/spill arrangements for the Republic of Ireland, now replaced by the Single Electricity Market (SEM), all participants were extended unsecured credit beyond a fixed €20,000 collateral requirement, regardless of financial capacity.
[7] As seen with Barings, Enron, utilities in California, etc. While some, such as Enron, had been de-rated before their collapse, they were investment grade until very shortly before the event.
[8] Market operators still have strong incentives to prevent default, as repeated instances of credit failure would undermine the integrity of their markets.
[9] Or some subset, depending on the rules of the individual market.
[10] In order to cover itself a participant would need to know all the other parties that traded and their quantities awarded (which it won't always be entitled to know), or have some projection of this for the cash markets (which won't have transacted yet). It would also need to know its own share of overall trade (unknown ahead of the event), and each participant's postings. Assuming it could somehow come up with an estimate of its socialized share of potential losses, it would then need to take out credit insurance or credit derivatives against each of these parties (which it won't be able to obtain for most of them).

In the financial markets it is established practice that clearing houses, as the backstop to the entire system, do not extend unsecured credit to anyone, ensuring that the collapse of any one participant does not result in a cascading credit default, a massive loss of liquidity, or otherwise compromise the integrity of the market as a whole.

In arguing against full collateralization it could be claimed that PJM fulfils a public interest role that normal clearing houses do not. This argument has some merit, though it would be difficult to argue that the role NYMEX plays in the oil and natural gas markets does not also serve a similar public interest function[11]. It can also be fairly argued that PJM has a monopoly in many of the markets it operates, whereas most clearing houses do not, and therefore participants cannot opt out of the full collateralization regime. The converse argument, however, is equally compelling – that, because PJM is a monopoly, participants cannot opt out of the socialized default structure and the risk that entails, and therefore PJM has a duty to put in place appropriate protective structures. As a corollary, it should be remembered that PJM already forces many participants to be fully collateralized; it just doesn't force them all at present.

Much of this debate, then, comes back to PJM's role. Current credit practices are an outgrowth of the old 'power pooling' days, when PJM was an agent facilitating wholesale settlement for a club of monopoly utilities. Nowadays, though, PJM is the de-facto clearing house of a complex and competitive commodities marketplace, with volatile prices and a broad participant base. It is logical that it would seek to apply the disciplines of a proper clearing house. PJM must decide, however, whether the associated market-wide benefits countervail the cost of additional collateral for a subset of participants.

***Recommendation**: All positions should be 'fully collateralized', to cover potential exposures to a high degree of statistical confidence. Unsecured credit should not be extended to any participant.*

3.3 Policy Issue – Participation

3.3.1 Credit as a Barrier-To-Entry

One of the arguments frequently advanced against full collateralization (or sometimes, the requirement to post any meaningful level of collateral) is that it would constitute a barrier-to-entry for new participants, in particular smaller players. This, however, tends to be something of a red-herring. In and of itself, substantial collateralization is not a significant barrier to the participation of smaller players, provided reasonable processes and an appropriate settlement timeframe are in place.

This can be seen on futures exchange such as NYMEX and CBOT, which have a large and active community of 'market locals', most of whom are thinly capitalized and have no credit standing (and who would almost certainly be considered sub-investment-grade if they did). In fact, the use of full collateralization in these markets effectively levels the playing field, so that all entities can participate on an equal basis regardless of their credit standing or financial strength. Similarly for PJM, many new entrants, particularly the smaller ones, are presently posting full collateral to cover their potential exposures – all ~60 days of it. Acceleration of settlement timeframe would actually reduce the collateral requirements for these entities.

[11] This is one of the reasons why 'contract markets' are more strictly regulated by the CFTC than others facilitating trade in near-identical products.

3.3.2 Understanding of Risk as a Requirement for Participation

In order to preserve the integrity of the market, however, some restrictions on participation are required. An essential hurdle should be that an organization has the wherewithal to understand the trading risks they are taking, and therefore the credit risks they are imposing upon the market. The objective is not to protect participants from the consequences of their own actions. "Poor performers go broke. This is a Darwinian consequence of the free market"[12], and it is not the role of the market operator to prevent this. It is desirable, however, for the market to protect participants, to the extent practicable, from adverse consequences resulting from the failure of others, especially when such outcomes are unpredictable and unhedgeable.

A number of parties have suggested the Eligible Commercial Entity criteria defined by the Commodity Futures Trading Commission (CFTC) as a useful requirement for participation. These state that an Eligible Commercial Entity must be:

> "An eligible contract participant or other entity approved by the CFTC that has a demonstrable ability to make or take delivery of an underlying commodity of a contract; incurs risks related to the commodity; or is a dealer that regularly provides risk management, hedging services, or market-making activities to entities trading commodities or derivative agreements, contracts, or transactions in commodities."[13]

This is not an unreasonable set of criteria, and needn't form a significant barrier to entry. For example, market locals on NYMEX are considered Eligible Commercial Entities for the purposes of trading on the Intercontinental Exchange (ICE).[14] Nevertheless, such criteria are suggested as a guide, not a prescription, for PJM's markets.

***Recommendation**: PJM should establish minimum criteria for participation in its markets, to be assessed on initial registration. PJM should have a right to verify ongoing compliance with these criteria, at its discretion.*

Failure to maintain adequate standards for participation can erode confidence in a market, and in an extension of Gresham's Law, itself drive away participants and liquidity. As one industry player stated recently: "You have the stupid money coming into the market now ... and I think the smart money is beginning to get a little frightened about what the stupid money will do."[15] Effective standards, which preserve the integrity of the market, actually serve to increase participation.

[12] Todd W. Bessemer and Francis X. Shields, *In Defense of Markets*, Public Utilities Fortnightly, June 2007.
[13] http://www.cftc.gov/educationcenter/glossary/glossary_e.html
[14] The S-1/A SEC Filing, filed by Intercontinental Exchange on October 27, 2005, states: "We have also obtained orders from the CFTC permitting us to treat floor brokers and floor traders on U.S. exchanges and ICE Futures as eligible commercial entities, subject to their meeting certain requirements."
[15] Comments by Jeffrey M. Christian, Managing Director of the CPM Group, as reported in *Commodities: Latest Boom, Plentiful Risk*, New York Times, March 20, 2008.

4 RISK MANAGEMENT PROCESSES

Clearing houses preserve their financial integrity through a system of rigorous processes and robust guarantees. This is summed up well by the Bank for International Settlements (BIS) and the International Organization of Securities Commissions (IOSCO) in their *Recommendations for Central Counterparties*[16]:

> "A CCP (Central Counter-Party) has the potential to reduce significantly risks to market participants by imposing more robust risk controls on all participants and, in many cases, by achieving multilateral netting of trades. It also tends to enhance the liquidity of the markets that it serves, not only because it tends to reduce risks to participants but also, in many cases, because it facilitates anonymous trading. However, a CCP also concentrates risks and responsibility for risk management in the CCP. Consequently the effectiveness of a CCP's risk controls and the adequacy of its financial resources are critical aspects of the infrastructure of the markets it serves."[17]

This section examines key processes which should be present in a robust, centrally-managed, credit risk management mechanism, for both cash and forward markets. Application of these processes to PJM's markets is discussed in general. Specific methodology options for risk assessment in the various PJM markets are addressed in Sections 6 and 7.

4.1 Credit Cover – Introduction

For the purposes of this discussion, credit cover equals the financial security (i.e. collateral) posted by a participant, plus any unsecured credit allowance it is granted[18].

Credit cover requirements for a participant should be determined based upon the risk that its positions, in aggregate, represent. Stated at the highest level, this is the exposure which a participant might incur by the time a default is detected and resolved (specifics differ between cash and forward markets). This exposure varies based on actual market outcomes and cannot be precisely known before-the-event. Therefore, in order to have a high degree of confidence that coverage is sufficient, it should be based on potential exposures, calculated to a high statistical likelihood ('99% adverse move' is a measure used by many futures clearing houses, such as NYMEX[19]).

Key to credit cover requirements is the time period of any exposure. i.e. a 30-day exposure obviously requires greater collateral to cover than a 7-day exposure. This creates an intrinsic linkage to the processes for settlement and the closure of positions. Also key are trading limits, which serve to limit some exposures. It is useful to address these topics before returning to the discussion of credit cover.

[16] In this context, when BIS/IOSCO refer to a 'central counter-party' they are discussing organizations which provide centralized market settlement and risk management.
[17] Committee on Payment and Settlement Systems, *Recommendations for Central Counterparties*, Bank for International Settlements, November 2004.
[18] This analysis is irrespective of earlier recommendations that unsecured credit not be offered.
[19] This is also consistent with capital adequacy standards, such as those that banks are obliged to follow, that require the maintenance of risk capital (that is otherwise unencumbered) equal to 99% potential exposure.

4.2 Closure of Positions

4.2.1 Forward Markets

It is standard practice in most forward markets to liquidate the outstanding positions of participants in default. This serves two purposes. Firstly, the residual value of the liquidated portfolio (if any) becomes available to satisfy outstanding obligations. Secondly, and more importantly from a risk management perspective, it serves to realize, and therefore limit, the defaulting participant's losses.

This does not imply that such an action will achieve the best possible value for the portfolio. It is possible that the value of the portfolio could improve if held, but equally possible that it could get worse. Realizing losses eliminates this uncertainty. Without liquidation, positions must be held until they naturally expire – through cash settlement or delivery – with the default guarantees of the clearing house continuing to absorb any ongoing losses. This is the unfortunate situation in PJM's FTR market at present.

> ***Recommendation****: In the event of participant default, PJM should have the right to liquidate any forward contracts not yet in delivery.*

Applying a liquidation mechanism to PJM's forward markets does pose some challenges. Liquidation can only occur if there is a market to sell the positions into. However, in PJM's FTR market at present:

- FTR auctions are only run on a monthly basis: meaning that substantial time can elapse between the incidence of default and a position being liquidated.
- The minimum tenor of FTR products is monthly: meaning that there is no way to liquidate the residual component of the month that is 'in delivery'.

Potential solutions are available, such as more frequent auctions and products with shorter tenors. Various options, with associated pros and cons, are discussed in greater detail in Section 6.1.

> ***Recommendation****: To support liquidation, as the delivery period approaches, longer-term contracts (e.g. an annual FTR) should decompose into strips of the next most granular contract (e.g. a strip of monthly FTRs).*
>
> ***Investigate****: the feasibility of shorter auction periods, and FTR products of shorter tenor.*

Liquidation also poses a liquidity risk, that dumping a large position into the market will itself result in an adverse price move, or not get executed at all (this problem applies to any forward market, not just FTRs). On the other hand, it is generally undesirable to have the market operator making commercial judgments about the best time to trade a portfolio into the market, even if only for liquidation purposes.

4.2.2 Cash Markets

There is no liquidation of positions in a cash market, as a matter of definition. A market is deemed a 'cash market' if it directly results in delivery (i.e. the 'spot' market), or trades sufficiently close to the time of delivery that the product has been (or is being) delivered/consumed by the time a default could be detected. This leaves only a settlement liability, which is closed out by cash payment.

4.2.3 Reassignment of Settlement Obligations

A number of electricity markets provide their participants with the ability to reassign settlement obligations between each other, often on a defined-MWh per trading period, or absolute-dollar, basis. In PJM, this is referred to as the 'eSchedule' mechanism. From a credit perspective, eSchedules[20] are unusual. They can be submitted months in advance, giving them a forward 'feel', as they imply financial obligations into the future, but ultimately they are not forward contracts, and result in settlement liabilities in the cash market.

The typical use of eSchedules is as a convenience to participants – allowing them to offset obligations in the cash markets, based on bilateral arrangements that exist between them outside of these markets. For example, participants who have a forward contract with each other may use eSchedules to settle the 'floating' component of this[21], in order to reduce their cashflows in the cash market. However, what happens in the event of default by the party to whom obligation is assigned?

eSchedules may span many months. For the portion of the eSchedule in the past there is little that can be done in the event of default – other than retrospective revision of settlements, which might appear arbitrary and undermine market confidence. However, the eSchedule might also extend for some time into the future, posing enormous risk and significant opportunity to game the credit system. [22]

> ***Recommendation****: In the event of default, PJM should have the right to immediately terminate the forward portion of any eSchedules which serve to increase the defaulting participant's settlement liabilities. Any eSchedule commencing in the future should be immediately rejected.*[23]

Ultimately, the two parties to the eSchedule choose to enter into a forward arrangement with each other, which includes the associated credit risk. It is not the job of the eSchedules mechanism to absolve them of this risk. To do so would be tantamount to clearing forward contracts; an activity well outside the scope of the cash market, and requiring the disciplines of a futures clearing house.[24]

Issue: eSchedules Representing 'Real' Physical Obligations

One slight variation on the above concerns eSchedules that represent an underlying physical reality, rather than a bilateral arrangement between participants. The classic example is when Party B has load within Party A's service area, and an eSchedule is submitted to reflect B's responsibility for the load. If B defaults, it would be necessary (due to the 'obligation to serve') to allow such eSchedules to continue to be submitted until its load is moved off to alternate retailers. It is understood that PJM's eSchedule system supports this differentiation between types of eSchedule[25].

[20] For the sake of convenience the term 'eSchedule' is used throughout this discussion, though the arguments presented apply broadly to the reassignment of settlement obligations, not just PJM's specific implementation.
[21] The 'fixed' component of the 'fixed-for-floating swap' would be settled separately between the two parties.
[22] For example, assume Participant P does business with a distressed counter-party, Q. A clearing house would require sufficient collateral to cover the position. Q doesn't have the funds, though. Instead, the parties initiate an eSchedule (which presently requires no collateral – a topic to be addressed later). The inevitable occurs, and Q defaults. Now, rather than bearing the entire default, P gets paid in full, minus his small socialized share of the default, which isn't a great impost as he got the original contract cheap anyway (remember that Q was in distress). Meanwhile, the other participants end up 'carrying the can' for a deal P only did because he knew he wouldn't have to bear the adverse consequences.
[23] Another way to think about this is that each day of the eSchedule is 'peeled off' and applied as that day trades. Only then is that day of the eSchedule accepted by PJM for settlement.
[24] If the CFTC comes to the conclusion that PJM is offering credit protection over a forward instrument, there is a real possibility that it might seek to assert jurisdiction, and force the application of such disciplines.
[25] Even if this were not the case, this is a change that could be made if necessary.

4.3 Settlement/Variation Margin

Settlement processes serve to reduce a participant's potential exposure, through the payment of monies that offset this exposure. The more frequently such settlement occurs, and the shorter the lag between liabilities being incurred and being settled, the smaller the overall exposure at any given time. This not only reduces the collateral required to cover this exposure to a given probability, it also lessens the size of any exposure that falls outside this probability range – and which, in the event of default, would need to be socialized, or draw upon other financial resources maintained by the clearing house/mechanism[26].

However, while the same principles concerning frequency and time lag apply, the settlement processes used by cash and forward markets in electricity tend to differ substantially:

4.3.1 Forward Markets

The principal settlement process used in forward markets is 'variation margining'. This process does not concern settlement for actual delivery of product[27], but is a periodic process to accrue incremental gains and losses, based upon price movements (i.e. variations) that occur as a result of ongoing trading. For this reason, the process is also known as 'mark-to-market', and is a common feature of all major futures markets.

To give a simple example, consider how this works for a basic futures contract (this example is also shown diagrammatically in Figure 1):

- Party A buys Natural Gas at ABC Hub for October 2008 delivery at $9.86/MMBtu. A is deemed to have a 'long' position (assume he started from zero). Later on the same day, Party B sells Henry Hub Natural Gas for October 2008, at $9.92/MMBtu[28]. B has a 'short' position.
- At the end of the trading day, the final settlement price for the October 2008 contract is $10.06/MMBtu.
- Variation margin is calculated overnight and monies move the following day. A receives the differential of $0.20/MMBtu, between the price it traded-in at and final settlement. B must pay the differential of $0.14/MMBtu, between the price it traded-in at and final settlement.
- The following trading day, neither A nor B trade. The settlement price for the day falls to $9.60/MMBtu.
- A must now pay $0.46/MMBtu, the difference between the previous day's mark and today's. B receives $0.46/MMBtu.
- The following day B trades out of its short position, by buying at $9.66. A holds its position. The final settlement price is $9.80.
- B must pay $0.06/MMBtu, being the difference between the previous day's settlement price and the price it traded-out at. A receives $0.20/MMBtu, being the difference between the two day's marks.
- In the final analysis, it can be seen that the variation margin for A (+0.20, -0.46, +0.20) equals the difference between final settlement price and purchase price (9.80 - 9.86). Similarly, variation margin for B (-0.14, +0.46, -0.06) equals the difference between the price at which it initially sold, and the price at which it bought back in (-9.66 + 9.92).

[26] This topic is discussed further in Section 5 on Default Management Structures and Processes.
[27] Forward markets which involve physical delivery of product also have processes for settlement associated with delivery. This is more akin to PJM's cash market settlement (a topic discussed in the following section), and not relevant in the context of PJM's forward markets.
[28] Obviously to someone other than A. Assume there is a diversity of buyers and sellers.

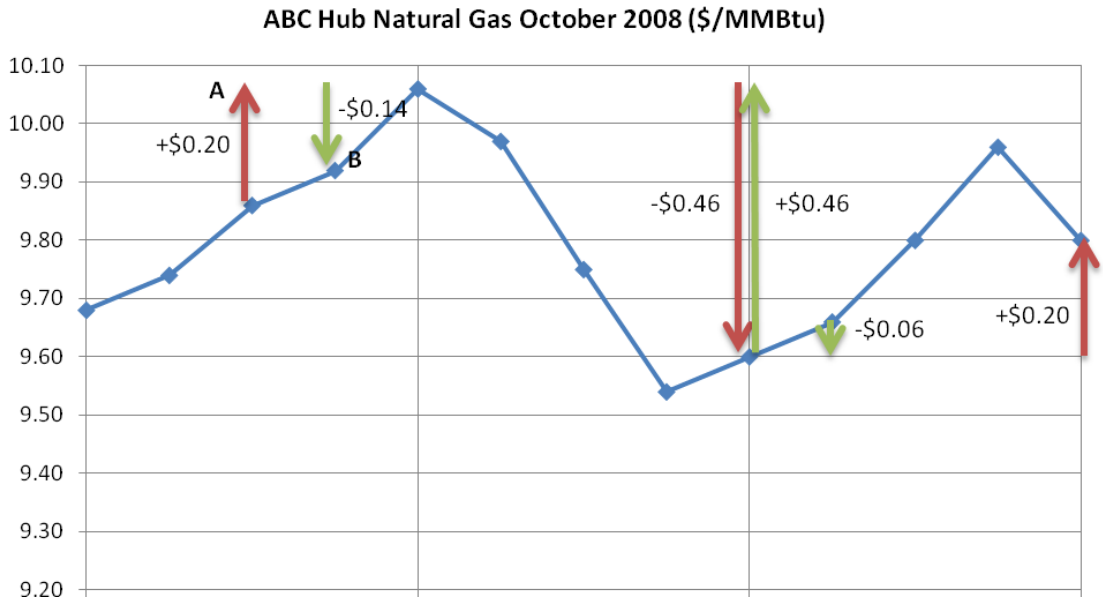

Figure 1 – Example Variation Margin Payments

There are a number of complications that must be considered in applying variation margining to FTRs, however.

4.3.1.1 Frequency of Mark-to-Market

It is only possible to mark-to-market as frequently as the market determines a price, which for FTRs occurs as part of the auction process. Currently this occurs monthly, with significant price movements possible during the month.

> ***Finding**: Variation Margining is of limited utility if auctions only occur infrequently (e.g. once a month), as circumstances can alter significantly between opportunities for incremental price adjustment.*

This issue could be addressed, at least in part, through more frequent FTR auctions, enabling price movements to be absorbed in a more incremental fashion.

4.3.1.2 Payment Lag

Currently payment in all PJM markets occurs as part of the standard monthly billing process, with payment 20+ days in arrears. There is little value in marking-to-market, however, if weeks then elapse before payment is made. There is no particular need, though, for settlement of variation margin to coincide with other payments. Settlement of variation margin on FTR positions could take place promptly after the auction (in most futures markets, for example, variation margin payments take place on a next-business-day, 'T+1', basis).

> ***Recommendation**: If Variation Margining is implemented, settle as soon as practicable after the auction, which need not coincide with cash market settlement.*

4.3.1.3 Payment Imbalances

Variation margining traditionally works on the basis of zero sum gain – for every change in price that results in payments to one set of parties, payments are required from other parties which, in aggregate, are exactly equal.[29] For FTRs, however, the net value of all outstanding contracts for a given period may vary from auction to auction. e.g. the total value of all October 2008 FTRs following the April auction may differ from the value after the March auction.

Under variation margining this would result in a settlement imbalance between periods – which would eventually be trued-up when the contract is 'delivered' in the day-ahead market. In the interim, a mechanism is need to manage these imbalances.

***Finding**: If Variation Margining is implemented, an imbalance account, held by PJM, will be required to manage overs/unders until the FTR is delivered in the DAM.*

The need for an imbalance account is partially offset if the holders of Auction Revenue Rights (ARRs) are margined, as these represent the 'short' side of the transaction for FTRs sold in the annual auction, balancing FTR holders' long positions. However, additional FTRs may be sold in the intervening auctions. For these, the short side of the transaction is the day-ahead market itself, or more precisely, the pool of congestion revenues from the day-ahead market for that month. Until the month rolls around, and its congestion revenues are available, the imbalance account would need to absorb any differential.

4.3.1.4 Margining of FTR Options and Handling of Option 'Premium'

As opposed to FTR obligations, a price is not established for every possible FTR option in each auction. In addition, one of the alternatives for conducting more frequent FTR auctions is to include the trading of options in only some auctions, due to issues with solution time. This raises a question. If a price is not established for an FTR option on a given path in each auction, how can it be marked-to-market?

The price paid for an FTR option can be thought of as having two components; the expected value of the path (the FTR obligation price), and the price paid for optionality (i.e. the protection from down-side risk). The latter of these is analogous to the option 'premium'. In each auction, there will be a price differential for each path, making it possible to mark the expected value component of the contract to market (subject to a limit that it not be marked below zero). It is necessary, though, to determine how the option premium will be settled – at the time the option is bought (standard practice in futures markets), at delivery (similar to current practice) or at some time in the interim[30].

***Recommendation**: Settle the option premium at the time the FTR option is bought. Mark the option to market after each auction based on the expected value component.*

The issues described above do not necessarily constitute an exhaustive set, nor do the potential solutions. The pros and cons of various solution options are discussed more comprehensively in Section 6.3.

[29] This works simply enough for standard futures, where for every contract traded there is a buyer and a seller, and if the position increases in value for one, it will decrease for the other.
[30] For example, the entire option could be marked-to-market at the next auction against the expected value of the path. This would result in the option premium being implicitly settled at the time of the first mark-to-market.

4.3.2 Cash Markets

Settlement in the cash market involves payment for actual product delivered, or traded very close to the event. In PJM, this includes the Real-Time Market (RTM), Day-Ahead Market (DAM), Regulation Market and Synchronized Reserve Market, as well as settlement of payments and charges for operating reserve, FTR payouts, eSchedules for the current billing period, and various non-market-traded amounts.[31]

Currently PJM operates a regime of monthly billing, with payment 20-22 calendar days in arrears, resulting (once other factors are considered) in exposures equaling around two months' activity. The settlement process can reduce these exposures through:

4.3.2.1 Shorter Billing Periods

Shorter billing periods result in more frequent settlement.

***Recommendation**: Reduce the cash market billing period to a duration of one (1) week or less.*

The majority of payments and charges in electricity cash markets (with PJM being no exception) are determined on a trading period (i.e. hourly for PJM) or daily basis, and summed over the billing period. From a calculation perspective, these amounts are compatible with a billing period as short as one day.

Payments and charges with a period greater than the billing period could be addressed in a number of possible ways, including:

- Changing calculations to be compatible with the accelerated billing cycle.
- Calculating approximate amounts within the billing period, with later true-up[32].
- Settling them on a separate billing cycle to the main market charges[33].

Pros and cons of weekly versus other billing periods are discussed in Section 7.3.

4.3.2.2 Shorter Payment Lag

Shorter payment lag means fewer liabilities can be incurred before a default is discovered.

***Recommendation**: Reduce the lag between the end of the billing period and payment to five (5) calendar days or less.*

The main limiting factor in reducing payment lag is the timeliness of the information required to perform settlement calculations. This is of less concern for markets which settle on traded quantities such as the DAM, but is more problematic for markets which settle on actual quantities, such as the RTM. Even for the latter, however, payment lag could be reduced appreciably, from current timeframes to around 3-5 days, without significant process implications[34].

To reduce payment lag below this, to around 1-2 days, would require more difficult process changes, such as:

- Acceleration of submission timeframes for data such as metering, which could involve industry-wide process changes.

[31] A detailed list of payment and charges can be found in: *Customer Guide to PJM Billing*, located at http://www.pjm.com/markets/settlements/downloads/custgd.pdf.
[32] Charges of this sort are typically non-market in nature, and generally have lower volatility than market-based charges.
[33] Many futures markets, for example, operate on a daily billing period, but bill their own fees on a monthly schedule.
[34] Though system impacts could be appreciable, depending on current system capabilities.

- Performing an initial settlement on less accurate (but still acceptable) data, with re-settlement a few days afterwards, as more accurate data becomes available[35].

Also important to remember in the discussion of payment lag is that, no matter how quickly the settlement calculations can be performed, payment itself cannot take place on weekends and bank holidays, as the wire transfer systems (e.g. Fedwire) are closed.

Another important, but somewhat less tangible, obstacle to accelerated settlement is potential resistance from participants whose 'float' benefit would be reduced. Participants who owe money to the market presently benefit from an average float of 36 days[36]. In a regime of weekly billing, paid five calendar days in arrears, this average float would be reduced to 8.5 days. Some participants will be content to accept this change as a trade-off for reduced credit risk (and reduced collateral, for those required to post collateral), while others will not. There is little that can be done to negate this impact, other than a cross-subsidy from net sellers to net buyers (perhaps with a sunset clause). This is not recommended.

Some participants may also be averse to changing their settlement processes to accommodate more frequent settlement. An argument frequently advanced is: "I bill my retail customers monthly, so wholesale settlement needs to be monthly too." From a process point-of-view, however, this argument is fallacious. Wholesale and retail billing periods do not need to be aligned – in fact, in many markets around the world, and some in the US, they are not. e.g. in Australia the wholesale billing period is weekly, whereas the typical domestic customer is billed quarterly. Additionally, even where both billing periods are monthly, retail customers do not all pay on one day of the month, just before wholesale settlement; they pay every day throughout the month. Management of cashflow, therefore, is a process inherent to market participation, that must be addressed regardless of billing period. Ultimately, if a participant only wishes to pay on a monthly basis for the sake of process convenience, this can be accommodated in a credit risk-friendly manner, through pre-payment.

4.4 Trading/Submission Limits

The objective of applying credit risk limits to trading and other submissions is to prevent a participant entering into a trade, or otherwise making a change in their position, that would cause the potential exposure of their net position to exceed their credit cover.

Such limits, however, are only useful where the participant's activity is controllable (i.e. the market operator can prevent them from incurring further liability) – such as in the Day-Ahead Market, FTR auctions and eSchedule submissions, where settlement (and hence credit exposure) is based on transacted or submitted quantity. In markets that settle on actual quantity – such as the Real-Time Market – trading limits only prevent a participant from taking part in the price-setting process, and do not prevent them from continuing to consume physical electricity, and thus amass liability.

***Recommendation**: Apply credit risk limits to all day-ahead market trading, FTR trading, eSchedule submissions and any other 'controllable' market activity.*

[35] This is a trade-off of accuracy for timeliness. While there is a risk of material changes in resettlement, having some money sooner, even if not all of it, reduces exposure vis-à-vis waiting a few extra days for any payment.
[36] Payment is an average 21 calendar days in arrears from the end of the billing period. Averaging the start and end of a billing period of ~30 days gives 15 days. The sum is 36 days.

4.4.1 Limits as a Statistical Measure

The methodology for calculating various risk limits is discussed in Sections 6.1.3, 7.1.1 and 7.2.1. However, regardless of the specific limit or calculation methodology, almost all limits involve some form of statistical estimation, and therefore do not provide a 100% guarantee that they will not be exceeded. For example, a limit on eSchedule submission might be based on submitted quantity, multiplied by price at that location/time to x% statistical confidence, based on historical prices for that hour-of-day/ location. If prices exceed this range, credit cover may be insufficient.

4.4.2 Use of Bid Limits vs. Trading Limits

Ideally, credit risk limits when evaluating a potential trade should be based on the potential exposure of the participant if the trade is done; a function of the participant's current position plus the additional quantity to be traded (+ve or -ve). There are serious difficulties, however, incorporating such limits within the optimization problem for multilateral auction markets.[37] As a result, a number of market operators use bid limits as a proxy for trade limits in certain markets. For example, PJM applies bid limits in its FTR auctions, and to certain virtual bidders in its day-ahead markets.

Bid limits, however, are a blunt instrument. They ensure that posted margin is sufficient to cover the largest position a given participant could acquire, based on bids submitted. On the other hand, if a participant submits a large number of bids it could be required to post significant collateral, even though only a fraction of its bids might be expected to clear. This presents a dichotomy, as the posting of large numbers of bids and offers, across multiple pricing points, is an activity that the market generally seeks to encourage, as it promotes liquidity and price discovery.

> ***Recommendation**: Ideally, use trading limits rather than bid limits. Where this is not practicable, institute a process to assess required collateral soon after an auction, and release unused collateral (or provide a streamlined process by which participants can request its release).*

4.5 Credit Cover/Collateral (Part 2)

The nature of credit risk differs appreciably between the cash and forward markets, and as a consequence, so does the way it is assessed and managed.

4.5.1 Forward Markets

In electricity forward markets, the credit risk for much of a contract's life is ostensibly a 'forward risk' or 'replacement cost risk', based on the cost of replacing the open contracts of a participant in default. This risk is covered by 'initial margin'. For contracts that are in delivery there is also a 'settlement risk', that a participant will not perform against the settlement obligations associated with the 'delivery' of its open position at expiry[38]. This risk is covered by 'delivery margin'.

[37] This is related to the fact that the trades executed form part of both the input to the problem – the credit constraint (where it contributes to potential exposure) – and the output of the problem.

[38] In the context of a centrally-operated electricity market, this 'delivery' is the nominal delivery of a financial instrument against the operator's cash markets – with settlement of delivery taking place in these cash markets, and against their indices – rather than actual physical delivery of electricity, which is the exclusive province of the real-time market.

4.5.1.1 Initial Margin

Initial margin requirements are generally calculated based upon the potential exposure of the *x*% worst possible move between when the position was last marked-to-market and when it could be liquidated in the event of default. If there is no variation margining, this move will be from when the position was initiated until when it goes to delivery, and if no liquidation, delivery margin should be assessed up-front, instead of initial margin.

***Recommendation**: Assuming the use of both variation margining and liquidation, set initial margin requirements to cover the x% worst move over two auction periods.*

It helps to illustrate this with an example. Assume Party A has an FTR position which was successfully marked-to-market after the latest auction. After the next auction, Party A fails to pay variation margin it is assessed. It is declared to be in default, and its position is then liquidated in the following auction (an inherent assumption is that Party A's position is not so large that it cannot be effectively liquidated in one auction). If auctions are held monthly, this represents collateral to cover a two-month move, though shorter auction periods result in a reduced impost (a topic discussed further in Section 6.2.1).

Some variations on this theme are also possible. For example, if two auctions were held in each round, with variation margin taken after each, anyone who didn't pay in the first round could be liquidated in the second round, a day or two later. This still runs the risk, however, that someone who survived the first round gets into trouble in the second round (though the time between mark-to-market and liquidation is still less than it otherwise would have been).

Issue: Determining Volatilities

A key issue to be resolved is how to determine an *x*% worst move. In futures markets this is typically based on implied or historical volatilities. For the FTR market, however, implied volatilities are not available[39], and there are concerns regarding whether there is a sufficient diversity of trading history for historical volatilities to be sufficient. A more reliable method for deriving volatilities would be to analyze the results of a range of simulated scenarios. This raises the challenge, however, of how many and which scenarios to pick, and how computationally intensive such a process would be.

Issue: Path or Portfolio?

Another key question is, what is this an *x*% worst move of? It is possible to determine the potential exposure associated with each FTR, assess initial margin to cover this exposure, and add everything up. Such an approach, however, ignores any risk correlation between FTRs, and would tend to over-collateralize a diverse portfolio.

The alternative, and more theoretically desirable, approach is to assess the potential exposure of each participant on a portfolio basis. This, however, could present practical difficulties depending upon the approach taken to calculating volatility. e.g. a requirement to maintain and apply massive correlation tables, or requiring the use of a simulation-based approach.

Both of these issues are addressed in greater detail in Section 6.2.1.

[39] Even though 'FTR Options' exist, these are not analogous to the call options used in futures markets for deriving implied volatility.

4.5.1.2 Delivery Margin

Delivery margin is applied to a forward position as it approaches 'delivery', to protect against a failure of the party to perform. For the 'long' side of a position (i.e. nominal buyers) , this represents the settlement risk associated with non-payment, which is generally based on the full value of the contract at the time it ceases trading in the forward market and moves to delivery (i.e. the price it was last marked-to-market at, also known as the 'final settlement price').

For FTRs, this value is the FTR auction price – as marked-to-market through succeeding auctions, if variation margining is in use. However, what the FTR 'delivers' is a financial flow in the day-ahead market – the FTR Payout – which serves as an offset to the purchase price. This FTR Payout is not known until the DAM transacts, however, and in considering the margin offset to allow, it is prudent, as with other margin calculations, to consider the *x*% worst-case scenario. Depending on the typical distribution of values for a path, this could be negative (as shown in Figure 2 below), and therefore would serve to increase delivery margin requirements.

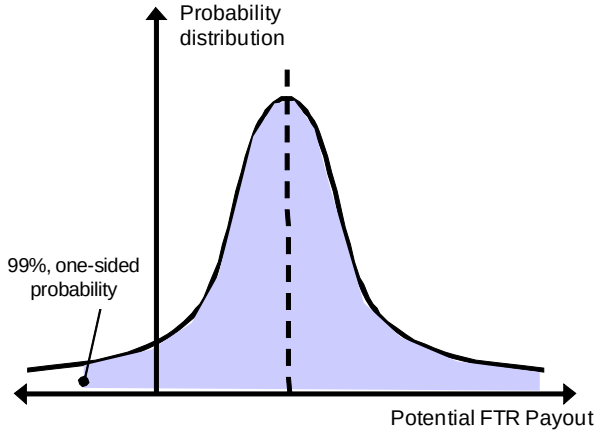

Figure 2 – Example Distribution of a Potential FTR Payout

As with initial margin, it must be decided whether to evaluate the offset from Potential FTR Payouts on a path-by-path or portfolio basis (with the same arguments applicable).

> ***Recommendation****: Determine delivery margin requirements for FTRs based on the auction value of the FTR portfolio, as marked-to-market, minus the x% worst-case payout of the FTR portfolio (which could be negative).*

One of the key questions concerning delivery margin is when to require it to be posted. Standard practice in many futures markets is to require delivery margin to be posted in the last couple of trading days prior to delivery (e.g. some on the penultimate day, and the residual immediately following the final day). The analogous situation for FTRs would be to require some delivery margin before the final auction, and the residual immediately afterwards.

This, however, begs the question – what happens if the participant fails to post the required delivery margin? Once the final auction has taken place, there is no longer a trading venue to liquidate the participant's position[40]. In order to avoid being under-collateralized heading into the delivery period, it is necessary to collect some margin prior to the final auction – 'pre-auction delivery margin'.

[40] In futures markets the EFP (Exchange of Futures for Physical) mechanism would be used to liquidate the position, but there is no equivalent mechanism available in the FTR market.

Pre-auction delivery margin must be sufficient that it, along with initial margin held, equals or exceeds likely delivery margin requirements. If the participant fails to post the required pre-auction delivery margin, its position would be liquidated. However, this presents a liquidity risk, particularly if the position is large. A variation on this approach, that mitigates some of the liquidity risk, is to require some or all delivery margin to be posted earlier. This could range from requiring some margin to be posted before the second-last auction, to requiring all delivery margin to be posted up-front (effectively the process in place at PJM today[41]).

***Recommendation**: Require 50% of estimated delivery margin requirements to be posted prior to the penultimate auction before delivery, with the remaining 50% posted prior to the final auction before delivery, and true-up after the final auction based on the final settlement price. Initial margin held against these positions should offset delivery margin requirements.*

Issue: Why 'Deliver' FTRs? Why Not Treat as a Cash-Settled Swap?

FTRs are financial instruments, that ultimately settle against prices in the day-ahead market. Why then can't they just be treated like normal cash-settled swaps[42], rather than going through a nominal 'delivery' process?

The answer lies in consideration of who the buyer and seller are. In a normal future, every contract has a buyer and seller, with long and short positions matching. For FTRs, participants are for the most part buyers (excluding FTR re-sales) – even in the case of counter-flow FTRs, which are just bought at a negative price. The nominal seller is the day-ahead market itself[43]. For any given long positions that are owed money, this is funded by short positions – in this case, the congestion revenues in the DAM (which, being for a future date, are not available yet)[44].

The FTR market is not in financial balance by itself, requiring funds from the DAM to achieve this balance. The simplest way to effect this is to 'deliver' open FTR positions against the DAM, where congestion revenues and FTR payouts can be settled together.

The corollary to the posting of delivery margin is the subject of when and how participants get it back. It is possible to recalculate delivery margin requirements throughout the delivery period. When settlement calculations are performed for a given trading day, the FTR Payout (and potentially the FTR Auction Payment[45]) for that day will form part of those calculations. As a result, these amounts become part of the Actual Exposures in the cash market (see Section 4.5.2.1 below), and no longer part of the potential exposures against which the delivery margin is provisioning. Delivery margin can be adjusted accordingly.

[41] It should be noted that if variation margining is used, and delivery margin is posted in full before the final auction, the level of delivery margin held will need to be adjusted after each auction (potentially up or down) to take account of the revaluation of positions. Some markets, such as New York, are presently doing this.
[42] A 'swap' is the exchange of one payment for another – usually a fixed payment (the strike price of the forward contract) for a floating payment (the underlying index, generally from a cash market). They are also referred to as two-way contracts for differences (CfDs).
[43] It is possible to think of ARR holders as the seller for FTRs sold in the annual auction, but this still does not account for additional FTRs made available in monthly auctions.
[44] The converse is true for long positions that owe money, with funds flowing to the congestion revenues.
[45] This presumes that the settlement for FTR purchase price (as marked-to-market) is split across the days it is in delivery. If the FTR is paid for in full at the beginning of the delivery period, this process changes slightly, but the daily recalculation of delivery margin remains valid.

***Recommendation**: Adjust delivery margin requirements during the delivery period, to take account of those days that are 'delivered', and for which the FTR Payout (and Auction Payment) form part of the Actual Exposures in the Cash Market.*

4.5.2 Cash Markets

In electricity cash markets the credit risk is ostensibly a settlement or receivables risk, associated with the failure of a defaulting participant to pay its settlement obligations. At any point in time, this risk is equal to actual exposures incurred, plus potential exposures that could be incurred before a default is detected and the participant prevented from incurring any further liabilities. It consists of the following components:

Actual Exposures

1. ***Outstanding Bills***: Settlement charges/payments that have been billed, but not yet paid. These amounts are known at the time of credit assessment.
2. ***Calculated Amounts***: Settlement charges/payments that have been calculated, but not yet billed. These amounts are known at the time of credit assessment.

Future Potential Exposures

3. ***Incurred Amounts***: Liabilities that have been incurred in the market, but have not yet been calculated (due to unavailability of information, process issues or some other reason). These amounts are unknown at the time of credit assessment.
4. ***Liabilities for Next Assessment Period***: Future liabilities that might be incurred in the time until credit is next assessed or payment is next due (whichever is sooner). These amounts are unknown at the time of credit assessment.
5. ***Liabilities for Time-to-Post/Time-to-Remedy***: Future liabilities that might be incurred in the 'time to post' a participant has if a call is made for more collateral, or the 'time to remedy' if it fails to make a payment on the due date. If these durations are not the same, the risk is the greater of the two. These amounts are unknown at the time of credit assessment.
6. ***Liabilities for Time-to-Transfer***: Future liabilities that might be incurred in the time a Load-Serving Entity (LSE) is allowed to continue incurring liabilities after being declared in default. This is a result of the time it takes to transfer the LSE's customers to another retailer, which itself is a consequence of the public policy 'obligation to serve'.[46] These amounts are unknown at the time of credit assessment.

Resettlement Exposures

7. ***Resettlement Amounts***: Future liabilities due to re-settlement of the market(s). These amounts are unknown at the time of credit assessment, and for some time thereafter.

Actual and Future Potential Exposures are depicted in Figure 3 below, for a market with monthly billing. In this diagram the current credit assessment is being performed mid-month. Actual Exposures will continue to increase each day until payment is made.

[46] This does not address issues associated with the default of a Provider of Last Resort (PoLR), for which most markets have no established procedure for termination, and could therefore result in virtually unlimited liabilities.

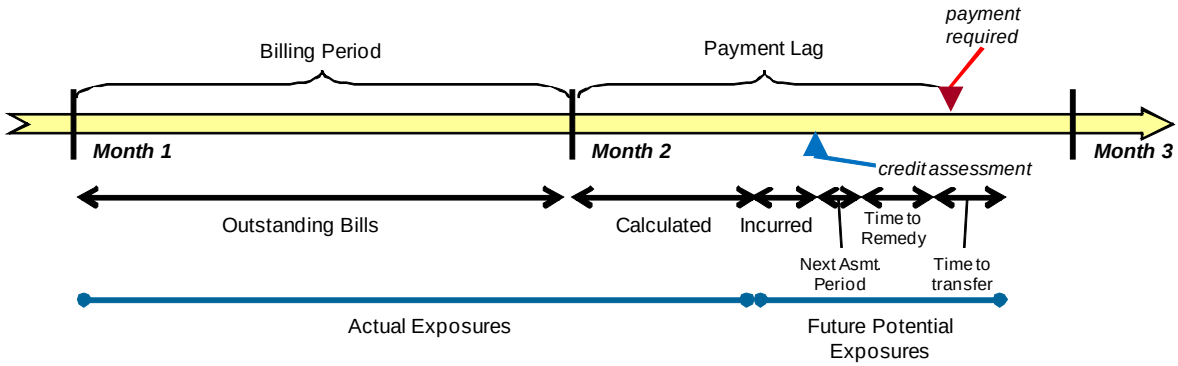

Figure 3 – Actual and Future Potential Exposures

4.5.2.1 Determining, and Reducing, Actual Exposures

Actual Exposures are determined directly from the settlement calculations performed by the settlement system, and bill payments made by participants. These exposures will fluctuate throughout the billing cycle, as liabilities increase with each new day's activity, and then decrease as a bill is paid.

The key to reducing Actual Exposures – and by extension, margin requirements – is to reduce the total number of days' liability outstanding at any point in time. As discussed in Section 4.3.2, this is achieved through shortening the billing period, and the payment lag between liabilities being incurred and being settled.

4.5.2.2 Determining, and Reducing, Future Potential Exposures

The Future Potential Exposure for a participant is the potential exposure which it might incur before its ability to amass further liability is curtailed (including any liabilities it might already have incurred but which remain undetermined). It is an attempt to answer the question "how much damage could this participant do before we could stop it?" to a degree of statistical likelihood.

As these exposures span a series of consecutive days, they can be assessed as a single contiguous risk, rather than a set of separate and independent days.

> ***Recommendation**: Calculate Future Potential Exposure as the worst contiguous n-day exposure that the participant might incur over the future exposure period (n), to x% likelihood.*

For example, in PJM at present:

- The lag between liabilities being incurred and being calculated – which includes liabilities already incurred (Incurred Amounts), plus those that might be incurred before the next assessment (Liabilities for Next Assessment Period) – can, at certain times of the billing cycle, be as long as eleven (11) days; potentially higher if there is a bank holiday.
- Time-to-remedy is three (3) days, for both a credit and a payment default.
- Time-to-transfer customers away from a defaulted participant is often one (1), but up to three (3) days.

In this case, the Future Potential Exposure can be calculated as the *x*% worst contiguous 17-day exposure that the participant might incur[47].

[47] For entities that do not serve load this could potentially be reduced to a 14-day exposure.

The key to reducing Future Potential Exposures is to reduce the time period these exposures span:

- ***Reduce time lag in calculating incurred liabilities***: Liabilities incurred in the market should be calculated as soon as possible after the event. While these results might not be perfect, and hence subject to later revision, imperfect knowledge of actual exposures is better than total ignorance.[48] Many electricity cash markets conduct a preliminary settlement calculation shortly after the event, often specifically to support credit assessment.

  Recommendation*: Perform an indicative settlement calculation the morning after the trading day, using approximated data as required.*

- ***Reduce time between credit assessments***: Reduce the time between credit assessment reduces the uncertainty concerning what has taken place in the intervening period.

  Recommendation*: Re-evaluate participant exposures, and associated credit cover requirements every business day, and possibly every calendar day*[49]*.*

Compared to current practice, this may require additional staffing and automation of processes, but based on the experience of other markets should be readily achievable.

- ***Reduce time-to-remedy***: There is little value in having processes in place for rapid assessment of credit exposure, and then providing days to remedy a non-payment.

  Recommendation*: Any failure to pay, or to post required credit cover, should result in an immediate 'margin call'. If the margin call is not satisfied within 24 hours, default proceedings should commence.*

A time-to-remedy of one (1) day is practical, and already used in some electricity markets. Experience in the futures markets – where the typical time allowed for responding to a margin call is only 1-2 hours – indicates that such a goal should be readily achievable.

- ***Reduce time-to-transfer***: Future Potential Exposure is driven by the maximum time allowed to move load off to a new retailer. In PJM at present this maximum is three (3) days, though in recent cases PJM has generally been able to transfer load to the Provider of Last Resort (PoLR) in around one (1) day.

  Recommendation*: Require the transfer of load responsibility to a PoLR to be completed within a maximum of two (2) business days.*

All of these improvements together could conceivably reduce Future Potential Exposure to five (5) business days, equaling 5-7 calendar days in most cases[50].

This discussion begs the question of how to determine Future Potential Exposure. Typical practice would be to use history as an indicator. Here, however, there is an issue…

[48] Collateral requirements are also reduced as, by design, the actual exposure will be lower than potential exposure in almost all cases (i.e. to *x*% statistical confidence).
[49] Assessment on a non-business day is less useful, as Fedwire is not operational and funds cannot be moved. However such an assessment can provide advance indication of a deteriorating participant position, allowing action to be prepared, and taken as soon as Fedwire re-opens (e.g. first thing Monday morning).
[50] Bank holidays would add to this. If margin is being managed dynamically, however, additional margin to cover the bank holiday could be requested shortly before the event, and returned afterward.

Issue: Past Performance as an Indicator of Future Performance

History is a perfectly valid indicator of future performance for some exposures – particularly those where the exposure is directly related to the underlying physical production or consumption of energy (or ancillary services), which follows certain statistical (and seasonal) trends.

However, some exposures are freely controllable by the participant, who can choose whether to incur them or not. For example, in PJM a participant may choose to participate in the Day-Ahead Market on one day, and not to participate the following day – or lodge an eSchedule for one week, and not the next. While for some participants these exposures may follow an historical trend, there is nothing compelling them to act the same way in the future, and therefore, historical performance is not valid as an indicator of future performance – at least for risk management purposes. So, what can be done to manage these exposures?

The solution is to use trading limits[51]. Although participants have discretion to decide whether and how much they wish to trade, it is up to the market operator to determine if they are allowed to trade. Trading limits ensure that participants are only be able to incur *controllable exposures* to the extent that they have posted residual margin to cover them. i.e. margin in excess of that needed to cover actual and *historically predictable exposures*[52].

Application to Day-Ahead and Real-Time Markets

The most important application of this approach is in considering exposures resulting from the day-ahead and real-time markets. Under the settlement approach in use by PJM (and other electricity cash markets in the US), settlement amounts for the day-ahead market (DAMS) and real-time market (RTMS) are calculated as:

DAMS = DAQ.DAP (summed across location and time)

RTMS = (RTQ-DAQ).RTP (summed across location and time)

…where:

DAQ ≡ day-ahead quantity, which is not historically predictable

DAP ≡ day-ahead price, which is historically predictable

RTQ ≡ real-time quantity, which is historically predictable in most circumstances (see discussion later in this section)

RTP ≡ real-time price, which is historically predictable

As defined, neither DAMS nor RTMS amounts are predictable based on historical performance. Additionally, as RTQ is not controllable through trading limits (see discussion in Section 4.4), RTMS is both uncontrollable and unpredictable.

From a credit point of view, however, it is useful to look at the aggregate liability across the two markets, and then reorganize terms:

$$DAQ.DAP + (RTQ-DAQ).RTP = \underbrace{RTQ.RTP}_{RTMS'} + \underbrace{DAQ.(DAP-RTP)}_{DAMS'}$$

(cont…)

[51] Or bid limits as a proxy.

[52] It should be remembered that 'historically predictable' does not mean that the market operator has prescient foresight; just that the likely range of outcomes is determinable to a statistical degree of confidence.

RTMS' is equivalent to the underlying physical market, and is historically predictable. DAMS' is equivalent to a day-ahead-hourly financial hedge around real-time (which, by no accident, equals settlement outcomes for virtual bidding), and is controllable through trading limits in the DAM.

Major Shifts in Consumption/Production

The final wrinkle to consider is what happens when there are material shifts in consumption (e.g. acquisition of major new loads through retail competition), or production (e.g. a new plant) for a given participant. Such changes can invalidate historical data for quantities that would normally be historically predictable (e.g. RTQ).

There are two basic solutions to this problem. In circumstances where responsibility for an existing facility transfers between participants (e.g. transfer of an existing load), it may be possible to also adjust each participant's historical data accordingly. Where this is not possible, including for all new facilities, it will be necessary to have participants submit details of their expected and potential quantities[53]. These estimates would then be used, in conjunction with projected prices to *x*% confidence, to determine potential exposure.

4.5.2.3 Determining, and Reducing, Resettlement Exposures

Resettlement is principally driven by changes in actual quantities, such as metering data. As such, resettlement exposures should be able to be estimated to a level of statistical confidence, based on historical trends. Where historical resettlement information for a given participant is not available, market-wide historical trends for the ratio between resettlement and initial settlement might also be usable – though less specific to any individual participant's circumstances (and therefore, probably less accurate).

The overall resettlement exposure for a participant, at any point in time, will be dependent upon its resettlement exposure across all billing periods which are still open to be resettled. The most effective way to reduce total resettlement exposure, therefore, is to reduce the timeframe in which resettlement can occur. i.e. if all resettlement must be complete within two months, exposures would be lower than if resettlement could occur for six months.

It is important to note that all participants may incur adverse exposures related to resettlement, even those otherwise in the habit of always receiving money from the market. For example, if a generator sells only into the real-time market it will always be a net payee in the initial settlement, but at a later date may incur resettlement liabilities, and therefore has a resettlement exposure that must be covered.

***Recommendation**: Calculate resettlement exposures based on historical trends for the participant, where available, or for the market otherwise.*

***Investigate**: the possibility of shortening the resettlement window for some or all charges and payments.*

[53] It is of course possible that participants might misstate these quantities in order to reduce their collateral requirements. It is impossible to guard against all such eventualities, and it should be remembered that such actions could constitute a deliberate fraud (and if such a fraud left PJM under-collateralized, might give it the ability to pierce the corporate veil and seek recompense from the individuals responsible as well as their company).

4.5.2.4 Credit Cover Requirements

At a minimum, a participant must post sufficient credit cover to cover those exposures which are known, or can be statistically estimated based on historic performance:

> ***Recommendation**: Set the Minimum Credit Cover Requirement for a participant equal to Actual Exposures + Resettlement Exposures + historically-predictable Future Potential Exposures + any controllable Future Potential Exposures already accepted (e.g. future-dated component of submitted eSchedules).*

However, posting the minimum credit cover would not allow for any further trading activity. In order for a participant to incur additional controllable exposures – e.g. through trading in the day-ahead market – it would need to post additional 'uncommitted' credit cover.

> ***Recommendation**: Set the Working Credit Cover Requirement equal to 110% of the Minimum Credit Cover Requirement. This is the level to which credit cover must be restored in the event of a margin call[54], and ensures that uncommitted credit cover is available for additional controllable exposures.*
>
> ***Recommendation**: Validate any new controllable exposures, using bid/trading limits, to ensure they do not exceed the participant's uncommitted credit cover.*

Additionally, actual exposures will vary on a daily basis throughout the billing cycle, increasing as settlement amounts are calculated, and decreasing as bills are paid. If only the minimum credit cover is posted at a point in time, this could result in the need to dynamically adjust collateral postings on a daily basis[55]. Some participants will be comfortable with such a process (e.g. those familiar with futures markets, where daily adjustment of collateral is standard practice), while it could be a source of consternation for others.

For the latter, this problem could be avoided, in all but extreme circumstances, by electing to post additional credit cover, such that their postings equal their Maximum Predicted Exposure – determined by replacing Actual Exposures with the *x*% worst exposure the participant might incur over a contiguous period of *d* days, where *d* is the sum of calendar days in the billing period and payment lag (e.g. for weekly billing, paid five calendar days in arrears, *d* equals 12).

> ***Recommendation**: Provide each participant with an estimate of their Maximum Predicted Exposure, as part of the daily credit process.*

If the participant elects to post credit cover equal to their Maximum Predicted Exposure they can have a high degree of confidence that they are sufficiently collateralized throughout the billing period, though ultimately they would not be obliged to maintain more than the Minimum Credit Cover Requirement.

In the event that a participant ceases to participate actively in the market, most of its credit cover can be returned, as the settlement and billing process runs its course. It will still be necessary, however, to hold some credit cover until all potential exposures from resettlement have been run-off. Some markets choose to set a participant's collateral floor – the minimum collateral the participant must post in order to participate – so that it equals or exceeds likely resettlement exposures.

[54] This is analogous to 'initial margin' and 'maintenance margin' for futures. When a position is established, initial margin must be posted. Thereafter, as exposures change, a margin call will not be made as long as posted collateral exceeds maintenance margin requirements. If and when this happens, collateral must be brought back to initial margin levels.

[55] Strictly stated it is credit cover that would need to adjust on a daily basis, but as unsecured credit allowances would tend to remain static for several months at a time, it is collateral that would be the balancing item,

***Recommendation**: For a terminating participant, continue to hold credit cover for resettlement exposures, until the resettlement window for their trading activity has closed.*

4.5.3 Credit Assessment Across Markets

A participant does not go bankrupt in the cash markets but remain solvent in the forward markets, or fail in the real-time market but not the day-ahead market. It goes broke all at once, in all the markets in which it participates. Therefore, the credit risk associated with a participant is by definition based upon its aggregate exposure across all these markets. The collateral it posts, as security for this risk, should therefore also span all markets.

For the sake of calculation simplicity it can be convenient to separately assess the exposures associated with each market, or each product within a given market. This is a theoretically valid approach when exposures are independent of each other, and can be a workable approach, that helps to simplify calculations, when correlation is low. As correlation increases, however, such simplifications become increasingly invalid, and the risk associated with the portfolio as a whole, including any risk offsets that correlation provides, should be assessed where practicable. Otherwise, credit cover requirements may be under-estimated, increasing market risk, or more likely over-estimated, requiring excess collateral to be posted by participants (which has a cost impact, and can itself lead to pressure not to fully collateralize, thus increasing risk).

4.5.4 Credit Assessment Across Related Entities

It is possible for a parent organization to have child entities with exposures capable of offsetting each other, the classic example being an organization that owns a generation subsidiary and a load-serving subsidiary. Assuming an appropriate legal framework is in place – e.g. that the parent is responsible for the child entities, and that it is subject to the market rules – it should be permissible to offset these exposures and assess credit cover requirements at the parent level. However, for the portfolio assessment of credit at the parent level to be of greatest utility, the parent would need to be wholly responsible for the child entities – and therefore be able to net settlement monies in and out – not just the provider of a limited parent guarantee.

Portfolio assessment of credit risk across related entities obviously has practical limits. For example, if one child has 400MW of load, and the other is a 300 MW generator, it should be straight-forward to assess the portfolio risk. On the other hand, if both entities hold different sets of FTRs, the assessment of this risk as a single, virtual FTR portfolio is likely to be impractical.

4.5.5 Degree of Confidence

This report makes frequent reference to an '*x*% adverse move' as a basis for determining potential exposures, and hence credit cover requirements. '*x*' is the degree-of-confidence to which potential exposures are determined, based on a one-sided probability distribution – or stated differently, it is the statistical likelihood that actual exposures (once known) will not exceed estimated potential exposures. Much of this report avoids stating a value for *x*, other than by way of example, as it is a topic that can be philosophically fraught, and does not have a single correct answer (though obviously should be large).

"What we anticipate seldom occurs; what we least expected generally happens"
Benjamin Disraeli

Values for *x* in electricity cash markets typically range from 95% to 99%. Futures market clearing houses tend to be more conservative, with NYMEX typically using a value of 99%, and LCH 99.7%. It is useful to think about these statistics in terms of the likelihood that they will be exceeded: 95% representing 1-in-20 (about once a month); 99% representing 1-in-100 (a few times a year); and 99.7% representing 3-in-1000 (about once a year). The capital adequacy standards which the US Federal Reserve Bank requires banks to maintain are typically based on a 99% confidence level[56]. Ultimately, the value chosen by a market will be reflective of the appetite for risk of the market operator (in order to protect its reputation and integrity), its members (to protect themselves) and its regulators (to protect the public interest).

4.5.6 Determination of Unsecured Credit Allowance

Section 3.2 recommends that "All positions should be 'fully collateralized', to cover potential exposures to a high degree of statistical confidence. Unsecured credit should not be extended to any participant." Nevertheless, this is a contentious issue, in which a number of parties will have strong opinions driven by factors other than the reduction of credit risk.

Should PJM make a policy decision to continue offering unsecured credit, a methodology which determines unsecured credit allowances based on counter-party credit rating and tangible net worth seems as justifiable as any other. However, one important consideration that has not been necessary to address within PJM to-date, due to the diversity of its membership base, is that of concentration limits. As stated by Fitch Ratings, "appropriate levels of exposure should not be based solely on the financial strength of the counterparty but also the size, credit quality and diversification of the membership base."[57] With mergers always in the offing, it is certainly conceivable that at some date in the future a single participant to accumulate a position which, in the absence of full collateralization, could threaten the viability of the market as a whole if a default were to occur.

The only way to address this is to impose additional constraints on credit, based not on the participant's own wherewithal but on the ability of the market as a whole to absorb loss. ISO New England and NY ISO, for example, both set concentration limits on unsecured credit for any single participant at 20% of market activity. Even if such constraints were not presently binding, it is better to put them in place ahead of the event, so that any large parties choosing to merge would know what to expect, rather than being surprised by a future change they could view as arbitrary. Of course, none of this discussion negates the clear finding of this report that, from a credit risk management perspective, full collateralization is both superior and preferable.

[56] For liquid risks. Standards can be higher for illiquid risks.
[57] Fitch Ratings, *Independent Grid Operators, A Credit View of RTOs and ISOs*, May 9, 2005.

5 DEFAULT MANAGEMENT STRUCTURES AND PROCESSES

Defaults fall into two basic categories:

- Payment default – the failure to pay a bill
- Credit cover default – the failure to post required credit cover

A key requirement of any clearing house – or organization performing these functions – is to have appropriate processes and structures in place to manage such events, so that the impact to non-defaulting participants is minimized, and the financial integrity of the market is preserved.

This section examines opportunities for improving the management of defaults within PJM's markets – broken into three principal topics: (1) the triggering of a default; (2) 'standard' default management, and; (3) extraordinary default management, including multi-tier guarantee structures.

5.1 Triggering a Default

As discussed in Section 4.5.2, participant exposures and credit cover levels should be re-evaluated daily. Clear credit limits must be established for every participant, based on credit cover posted.[58] If a participant's aggregate exposure (and hence its Minimum Credit Cover Requirement) exceeds its credit limit, a 'margin call' must be made, requiring that additional credit cover be posted, in order to restore levels to equal or exceed the Working Credit Cover Requirement.

> ***Recommendation****: Failure to respond to a margin call within one (1) day should result in a participant being put into 'credit cover default'.*

This should immediately activate the default management processes described below. i.e. there should be no further delay to 'give them a chance to fix the problem', as each additional day equals additional liabilities, and the margin call is the participant's chance to fix the problem.

It is often also useful to establish warning levels. When these levels are hit, a warning notice to the participant, and potentially to others, is generated that notifies them that their exposures are moving into riskier territory. Ultimately, though, these warning levels should be strictly advisory – any 'warning' that is more than advisory is, in effect, a limit. Participants should be given the opportunity to address warnings they receive in ways other than posting collateral, such as cessation of trading (where practical).

> ***Recommendation****: Provide formal warnings to a participant when its aggregate exposure reaches 80% of credit limit, and again when it reaches 90%.*

Payment defaults are fairly cut-and-dry. Participants have clear advance notice of when a payment is due, and failure to pay by the due date is a serious failure.

> ***Recommendation****: Failure to remedy a failure to pay within one (1) day should result in a participant being declared to be in payment default.*

In this event, default management processes should be immediately activated.

[58] These limits might equal the posted credit cover, or could be a percentage (<100%) thereof, in order to provide some 'headroom' – though the gap between Working Credit Cover Requirement and Minimum Credit Cover Requirement (see <<Section x.x>>) already provides some of this headroom.

5.2 Default Management – Standard

Firstly, what classifies a default as 'standard' versus 'extraordinary'? For the purposes of this analysis, standard default management processes apply to all defaults, regardless of severity, whereas extraordinary processes are those that apply in circumstances where the defaulting participant's liabilities, post liquidation, exceed the credit cover it has posted.

The Bank for International Settlements (BIS) and the International Organization of Securities Commissions (IOSCO), in their *Recommendations for Central Counterparties*, advise that:

> "A CCP's default procedures should be clearly stated, and they should ensure that the CCP can take timely action to contain losses and liquidity pressures and to continue meeting its obligations. Key aspects of the default procedures should be publicly available."[59]

In the event of a default, steps must be taken against the defaulting participant to limit the incurrence of further liability, liquidate their forward positions and seize any security posted – while also minimizing the impact to non-defaulting participants.

5.2.1 Limit Incurrence of Further Liability

In most markets the terms 'limit the incurrence of further liability' and 'prevent further trading' would be synonymous. One of the features that makes electricity cash markets unusual, however, is that because of the 'obligation to serve' retail load, certain participants must be allowed to continue incurring obligations for a period of time, even though they have been declared to be in default.

Immediately a default is declared, the participant should be suspended from trading in all markets – cash and forward. This will prevent all participation in markets where settlement is based on scheduled quantity[60], such as the day-ahead market. It will also prevent the participant from taking part in the price-setting process for the real-time market, and other markets which settle against actual consumption.

> ***Recommendation**: In the event of default, PJM should terminate all eSchedules, and other transfers of settlement obligation, which serve to increase the participant's exposure (but not those which decrease it), from the current day forward.*

Even though a participant is suspended from trading, and therefore taking part in price-setting, it can continue to incur liabilities as a price-taker for the load it serves – in the real-time market and any other market where obligations are based on actual quantities – until its load responsibility is transferred off to other retailers, generally referred to as Providers of Last Resort (PLRs). In PJM this can presently take up to three (3) days.

Issue: Default of a Provider of Last Resort

If the PLR serves as backstop to the default of a normal Load-Serving Entity (LSE), who serves as the backstop to the PLR? The answer, in many electricity cash markets at present, is: no-one. This involves an implicit – and very dangerous – assumption that a PLR cannot default. However, the organizations which serve as PLRs are generally large Investor-Owned Utilities (IOUs), and as shown in California, such organizations can default.

(cont…)

[59] Committee on Payment and Settlement Systems, *Recommendations for Central Counterparties*, Bank for International Settlements, November 2004.

[60] Though the participant might still receive settlement obligations in the DAM, from the 'payout' of negatively-priced FTRs for the period currently 'in delivery'.

The current standing instruction in the event of a PLR default it to take the issue to FERC. However, there is no indication of timeframe for FERC to consider the matter, nor of what it might do as a result. Meanwhile, the PLR can continue to amass liabilities related to serving its retail customers, constituting what is effectively an open-ended and unquantifiable credit risk[61]. There is little that any clearing house can do to protect against such a risk.

The best solution to this problem is to establish clear processes and timeframes for the transfer of load in the event of default by a PLR, consistent with the mechanisms which apply to other LSEs. Another, less appealing, option could be the establishment of some form of mutual assurance scheme amongst PLRs (perhaps coupled with some insurance), though ultimately this would be a substitute for an inability to tackle the root cause of the problem.

5.2.2 Liquidation of Forward Positions

Liquidation of forward positions has two objectives – to realize any remaining value in the positions which can be applied to the default[62], and more importantly, to limit further losses from those positions. The process for, and constraints on, liquidation of positions are discussed at length in Sections 4.2.1 and 6.1.1.1.

5.2.3 Seizure of Security

The purpose of requiring collateral to be posted is so that, in the event of a participant default, it can be seized for the purpose of satisfying any outstanding financial obligations. In order to adequately fulfill this function, in a manner that does not interfere with other participants' cashflow, it is essential that collateral is both secure and highly liquid – cash, or instruments rapidly redeemable for cash, such as treasury bills or an appropriate Letter of Credit (LC). The latter could do with exploring further.

Many variations of LC are possible, only a subset of which are useful as collateral in a volatile commodity market. To be acceptable, an LC should be:

- Payable on demand: The clearing house should not need to prove that certain conditions have been fulfilled in order to call the LC – it is payable when called.
- Payable the same-day: The issuer of the LC must perform the same day the LC is called (or if called late afternoon, within a specified number of business hours).[63]

Procedures must be in place for promptly calling on LCs, as the need arises.

***Recommendation**: PJM should continue to employ standards which all LCs must comply with, and stipulate them in a pro-forma LC. These standards should include that the LC is payable on demand, and payable the same day.*

***Recommendation**: PJM should continue to maintain minimum financial standards (including required credit ratings) for LC providers. A list of accredited LC providers should be maintained and regularly reviewed, with participants able to petition to have additional providers added.*

[61] Which is, in actuality, a secondary effect of an underlying regulatory risk.
[62] Though the value of the positions is most likely diminishing, and a contributory reason for the participant being in default in the first place.
[63] Most futures market clearing houses, such as NYMEX, require performance within 1-2 hours, so such LCs are certainly available.

Issue: Insurance as a Substitute for Collateral

In a number of electricity markets, over the years, insurance has been suggested as an alternative mechanism for the management of credit risk – and not infrequently as the panacea to all credit woes. Very few markets, however, have ended up incorporating insurance into their credit structure, and none as a solution for the entire problem. Why is this?

The starting point for many markets has been to seek insurance that provides a total guarantee on all market activity – in other words, covers the default of any participant, at any time, for the participant's entire exposure. This is a difficult risk for an insurance company to assess, involving a multitude of parties of different financial strength – some rated and some not, and exposures varying from period to period. Not surprisingly, the quotations were prohibitively expensive, and often sought to impose unacceptable restrictions on the market.

Insurance companies have been more enthusiastic about offering insurance limited to highly-rated parties. All other participants in the market – these ironically being the riskiest participants, over whom credit protection is most desired – must be dealt with some other way (i.e. collateralization). There is also the issue of the 'deductible' on the policy, which must still be recovered from participants in the event of default. This poses a dichotomy, as a low deductible can encourage the market operator to become complacent in its management of credit risk – something the insurance company would not wish, whereas a high deductible could result in sizable socialized charges, one of the very problems the insurance is meant to mitigate.

There is also the question of who should pay. If some participants are required to post full collateral, is it right that they also pay for this insurance? The most common approach is for this insurance to be paid for by the insured parties[64]. However, this begs the question: if the insured entity is paying for insurance on itself, what is the functional difference between this and posting a letter of credit? The answer is – not much, except LCs are generally simpler in form, easier to claim, quicker to perform and cheaper for participants to obtain. Finally, in the event of a default it may prove difficult to renew insurance in later years, at least without significant rate increases, whereas there should be far less difficulty for non-defaulting parties (i.e. those remaining) to continue to obtain LCs.

None of this discussion should imply that there is no valid role for insurance – in fact, as discussed in Section 5.3, insurance can fill a useful place in the market's trade guarantee structure. Its application, however, is limited not universal, and it does not serve as an effective substitute for collateral.

5.2.4 Minimizing Cashflow Impact to Non-Defaulting Participants

A key aim of any default process is to ensure that payments to/from non-defaulting participants continue to perform as expected. An issue common to many electricity markets, though, is that even in situations which are promptly resolved, such as:

- a payment default, which is resolved through seizure of security;
- a non-payment (or short-payment), which is resolved by payment before a default is declared;

… non-offending participants are still short-paid, or must pay extra. Even though these participants are 'made whole', it can pose short-term cashflow difficulties for some. While they are eventually 'made whole', this can pose short-term cashflow difficulties, as well as undermining confidence in the market.

[64] This, for example, is the approach adopted by ISO New England.

***Recommendation**: PJM should maintain a line-of-credit, or similar facility, to provide bridging finance for near-term shortfalls (to a reasonable limit).*

Similar facilities have been used effectively by other markets – including PJM, which has used similar arrangements on an improvised basis.

Such a mechanism should serve solely as a cashflow management tool, to avoid unnecessary socialization. Any shortfalls remaining after this period should be dealt with using the structures and processes for extraordinary default management.

5.3 Default Management – Extraordinary

A general aim of any central clearing mechanism is to ensure that a defaulting participant's liabilities (post liquidation) can be satisfied by the security it has posted. It is not reasonable (and probably not feasible), however, to collateralize to cover all conceivable scenarios. As a result, collateral is determined on a probabilistic basis, to a high degree of confidence – with 99% being a typical measure. This means there is a 1% chance of being under-collateralized. If a participant going into default happens to be one of these under-collateralized parties, then the default cannot be resolved by the defaulting participant's posted security alone, and thus the default is 'extraordinary'.

It is essential to have mechanisms in place that, in such circumstances, preserve the financial integrity of the market as a whole. Historically, for most electricity spot markets this has meant the socialization of any residual default amounts. However, these amounts are unpredictable and unhedgeable, and the use of the socialized default mechanism serves to undermine the confidence of participants, and potential participants, in the market.

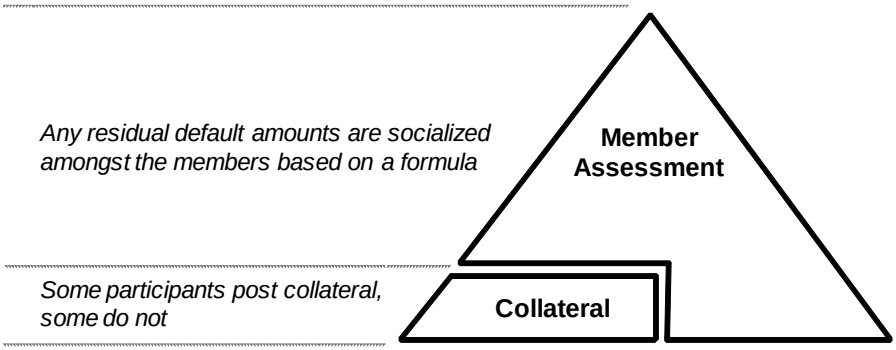

Figure 4 - Typical Electricity Cash Market Guarantee Structure

The BIS and IOSCO recommend that:

> "A CCP should maintain sufficient financial resources to withstand, at a minimum, a default by the participant to which it has the largest exposure in extreme but plausible market conditions."[59]

In other words, the central counter-party should maintain additional financial resources to cover the difference between what is collateralized (based on potential exposure to *x*% confidence) and the most extreme event that could plausibly happen. e.g. default of the largest participant in the market during a period of high-priced activity.

To understand the types of financial resources which might be drawn upon it is instructive to look at futures market clearing houses, which tend to utilize an escalating sequence of trade guarantees. Elements included in these guarantee structures can include:

- ***Guarantee Fund***: Many clearing houses utilize a guarantee fund[65], which all members contribute to (as a condition of membership), and serves as a monetary reserve to cover default, until depleted.
- ***Clearing House Profits***: The profits of the clearing house, where applicable, are used to cover default, usually to some pre-defined limit
- ***Insurance***: An insurance policy is taken out to cover amounts not covered by other financial resources. This usually comes later in the trade guarantee structure. Monies paid out from other resources, such as the guarantee fund, serve as the de-facto deductible (constituting the clearing house's 'skin in the game'), thus helping to keep premiums reasonable.
- ***Clearing House Capital***: The capital reserves of the clearing house are used to satisfy the default. Because this compromises the ongoing ability of the clearing house to function, this is typically a last resort in the markets in which it is used.
- ***Member Assessment***: Members are assessed a socialized share of residual default amounts. This (along with the use of clearing house capital) is usually a last resort.

Figure 5 and Figure 6 illustrate how these elements are used in the guarantee structures for NYMEX and LCH, the two largest energy clearing houses in the world[66].

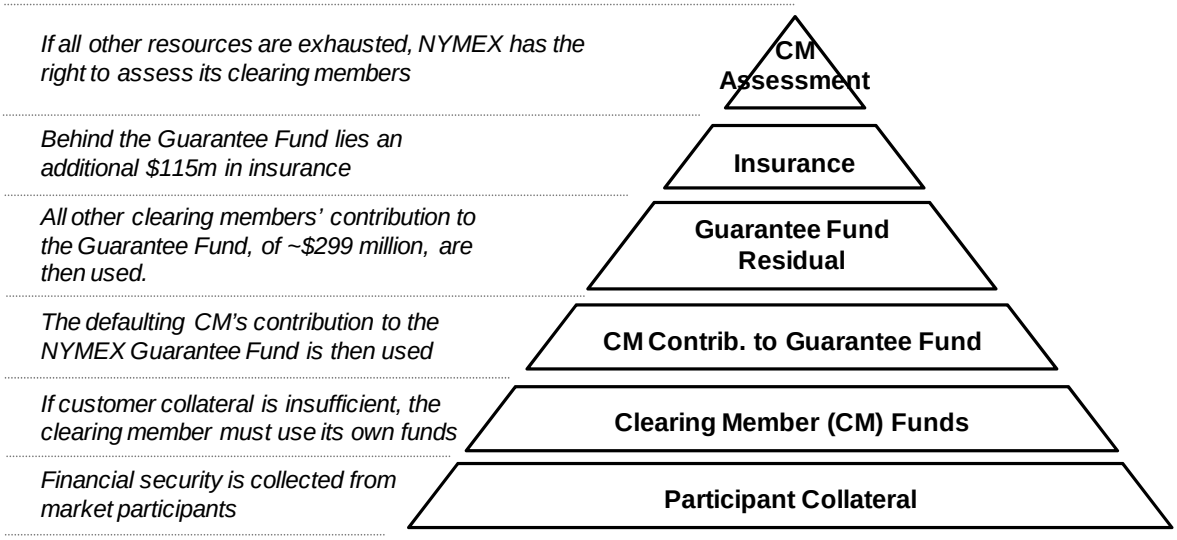

Figure 5 - NYMEX Guarantee Structure[67]

[65] Also known as a default fund, reserve fund, etc.
[66] The financial resources shown are used to support all trading, not just the energy product suite.
[67] As of December 31, 2007, NYMEX's Guarantee Fund contained ~ $299.3 million. Default insurance of $115 million was held, callable for defaults in excess of $250 million. Source: NYMEX 10-K Annual Report, filed March 3, 2008.

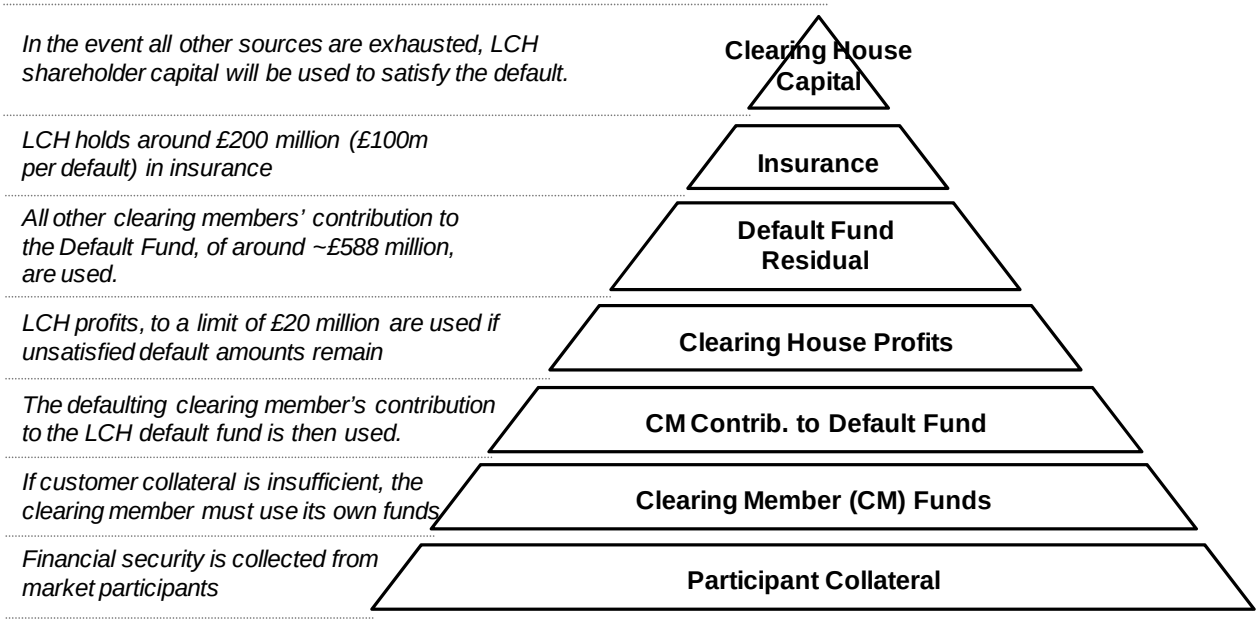

Figure 6 - LCH Guarantee Structure[68]

In a market such as PJM, which runs as a not-for-profit mutual, the concepts of clearing house profit and clearing house capital are not readily applicable. The use of a guarantee fund and residual default insurance, though, are readily translatable into the PJM context. Some electricity markets also have residual funds available from charges such as non-compliance penalties, late fees and the like, some of which might also be usable for such a purpose.

***Recommendation**: PJM should explore the creation of an escalating guarantee structure, including residual funds held by PJM, a guarantee fund contributed to by the members, and default insurance. Socialization of default to members should be the last step in this structure.*

[68] All numbers in this diagram are sourced from the undated LCH.Clearnet publication: "*LCH.Clearnet Limited's Default Protections*", referenced on the LCH website on April 23, 2008.

6 PROCESS AND CALCULATION: FORWARD MARKETS

This section discusses some of the key options, and their associated pros and cons, related to:

- Changes that could be made to PJM's business processes to reduce credit risk. This addresses not just explicit credit risk management processes, but also process options in other areas which have a credit risk impact.
- The methodology for performing credit risk calculations.

The focus of this section is on PJM's forward markets.

6.1 FTR Trading

6.1.1 Increasing Auction Frequency

Under a regime where FTRs are auctioned monthly – such as that in place at PJM presently – the value of FTR positions can move significantly between consecutive auctions (including in the time between the incidence of default and when a position can be liquidated). This seriously limits the usefulness of variation margining as a tool for realizing 'incremental price adjustments', and requires high levels of initial margin, to cover the '*x*% worst-case adverse move'.

These issues can be addressed, at least in part, through more frequent FTR auctions. However, there are two principal limiting considerations to such an approach:

6.1.1.1 Liquidity

Periodic auctions serve to concentrate liquidity to a certain point-in-time. If auctions are conducted too frequently, it may serve to split this liquidity, making it more difficult to achieve good liquidity in any one auction – a particularly important consideration for the liquidation of defaulted positions.

Options for addressing liquidity are limited. Over the medium-term, as a market matures, liquidity will tend to progressively develop – provided the product itself is something that is of commercial interest to participants to trade. Sometimes this development can be accelerated by reducing inertia, such as ensuring that the tools for large-scale trading are convenient to use, collateral is managed efficiently, etc.. 'Market-maker' schemes – which offer incentives or fee breaks for providing liquidity – are also used in some forward markets.

***Recommendation**: Review PJM's FTR bidding systems to ensure they provide effective support for entry and cancellation of large numbers of bids. Review the system used to download auction results to ensure it provides efficient straight-through-processing.*

6.1.1.2 Auction Execution Performance

Execution of FTR auctions is a computationally intensive process , with serious performance limitations on how quickly the auction can be run, which could become limiting if auctions are run more frequently than they are at present.

Options for increasing the frequency of auctions generally involve separation of auctions into two classes:

- 'Full' Auction: which would be the same as those run today, and take place on a similar timeframe (e.g. monthly).
- 'Reduced' Auction: which would utilize various functional compromises to limit solution time. These auctions would take place more frequently.

Some of the functional compromises for Reduced Auctions might include:

- Trading of FTR Obligations only; no FTR Options.
- Using a reduced set of nodes. e.g. hub/zone-to-hub/zone.
- Trade a limited set of time periods. e.g. only the next few months.

However, these compromises would reintroduce, for those products/paths not included in the reduced auction, some of the limitations that more frequent auctions sought to address, and potential create some new complications:

- The ability to liquidate open positions would be limited to only those paths/products that traded in the reduced auction – with all others not able to be liquidated until the full auction.
- It would not be possible to directly perform mark-to-market, though for the first two alternatives it would likely be possible to extrapolate prices from those paths/products that did trade. This, however, starts to edge away from true mark-to-market towards a 'mark-to-model' approach.
- As a consequence, the '*x*% worst adverse move' between mark-to-market and liquidation would have a different basis for different products. This would make it extremely problematic to determine initial margin on a portfolio basis. This might not be such a problem if only FTR options were excluded from the reduced auction, but would likely be highly problematic if certain paths were excluded (e.g. hub-to-hub).

The two obvious candidates for accelerating auction frequency are to move to weekly or daily auctions. These are contrast below with monthly auctions.

Table 1 - Comparison of Options for FTR Auction Frequency

Monthly	**Weekly**	**Daily**
Can run full auction each time.	Full auction might be possible. Elimination of only FTR options is very likely to work.	Would require major reductions in problem size, such as no FTR options and hub-to-hub.
Variation between marks could be quite substantial.	Reduces variation likely to occur between marks.	Minimizes variation likely to occur between marks.
Maximum initial margin requirements.	Reduced initial margin requirements.	Minimum initial margin requirements.
Liquidity same as today.	Not likely to face liquidity challenges.	Likely to face liquidity challenges.
No changes to existing bidding and trading processes.	Modest changes to bidding and trading processes.	Substantial changes to bidding and trading processes

***Investigate**: whether a full auction can be run on a weekly basis.*

***Investigate**: whether a reduced auction that only eliminates FTR options is feasible to run on a daily basis, and confirm it is viable for a weekly basis.*

Although daily auctions are theoretically more desirable, there are likely to be some significant practical constraints on making them work, both in terms of systems and liquidity. A weekly auction is likely to be more practical in the near-term, and does not preclude a move to daily in the longer-term (by which time solution performance improvements can also be expected, which may require fewer functionality compromises to be made).

> ***Recommendation**: Pending the investigation into performance of reduced auction models, move to the weekly auctioning of FTRs. Re-examine the feasibility of daily auctions within two years of this move.*

6.1.2 Reducing Contract Tenor

The 'tenor' of a forward contract is the period of time which the contract spans. PJM presently offers FTR contracts with monthly and annual tenors, with the annual contracts 'decomposing' into individual months which are separately tradable. Currently if a default occurs, even assuming the market operator has liquidation rights, there is no way to liquidate the residual component of the month that is 'in delivery' – meaning that, although other exposures in the cash market may have been terminated, the FTR payout will continue until the month is complete. This is not an issue if the FTR payout is positive, but is problematic if the payout is negative.

One way to limit the amount of this exposure is to decompose the contract into smaller blocks, that could continue to trade, and therefore be liquidated, into the delivery month. As weeks do not divide neatly into months, the most logical solution would be for the delivery month to decompose into a strip of daily contracts a few days before the commencement of the delivery month. These contracts could be separately traded, allowing participants to fine tune any residual exposure due to short-term variations. To be effective this would require FTR auctions to take place every (business) day, though conceivably these auctions could be limited to only FTR obligations, and only daily contracts for the current month. Whether this is feasible from a systems performance point-of-view is yet to be determined.

> ***Investigate**: the feasibility of providing a separate auction for the trading of daily FTRs in the delivery month, including the systems performance issues that would be involved.*

Reducing the contract tenor would effectively mean that most of the current month is not really 'in delivery', but is still available for forward trading. If this short-term forward trading of FTRs became sufficiently liquid – and this is a big 'if' – it could ultimately mean that delivery margin requirements could be reduced to cover only the next few days. This, however, might be some time off.

6.1.3 FTR Bidding/Trading Limits

As discussed in Section 4.4, trading limits result in more efficient usage of credit cover, and are preferable to the use of bid limits, but may face serious implementation difficulties.

> ***Recommendation**: PJM should discuss the feasibility of incorporating credit constraints within its FTR auction algorithm with the supplier of this software (AREVA).*

Assuming that this is not practically achievable, at least in the near-term, it is important to ensure that bid limits make as efficient a use of credit cover as possible. It is understood that PJM is already in the process of making changes to its bid limit methodology, to assess the exposure from submitted FTR bids for a given path based on the maximum exposure of any quantity of FTRs that might be awarded for that path (given that all accepted bids for a given FTR path will trade at the same clearing price), rather than the current practice, where exposure is calculated as the total area under the bid price-quantity curve, resulting in over-collateralization.

As also discussed in Section 4.4.2, exposures should be re-assessed immediately after completion of an FTR auction, based on the participant's aggregate portfolio, including any new awards, with any collateral held to cover potential bid exposures returned.

6.2 Credit Cover

6.2.1 Initial Margin for FTRs

Section 4.5.1.1 discussed two key challenges in calculating credit exposure, and hence initial margin requirements, for the FTR position held by a participant – more reliable determination of volatility, and determination of exposure on a portfolio basis.

There are two general approaches to addressing this problem:

- ***Historical***: Determine the exposure of each path, separately for peak and off-peak, based on an analysis of historical auction prices. Maintain a (large) matrix of the risk correlation between paths, and apply these correlations as offsets to the individual path exposures of the participant's positions, in order to determine a portfolio exposure.
- ***Simulation***: Run the FTR auction (or a simplified version thereof), for a range of different scenarios, based upon perturbation of various inputs such as which transmission lines are in service, load, etc.. Determine the value of each participant's FTR portfolio under each of these scenarios, and the 'portfolio value differential' between this and its current value (as determined in the auction that just occurred). Based on a statistical analysis of these portfolio value differentials, calculate potential exposure based on the '*x*% worst move'.

The characteristics and implications of these two options are contrasted in the table below.

Table 2 – Comparison of Options for Calculating Initial Margin

Historical	Simulation
Trading history for FTR auctions not presently deep enough to determine worst-case move with reasonable confidence.	Likely to result in a more accurate picture of potential portfolio movement, provided scenarios are chosen appropriately
Volatilities and correlations can be determined statistically from historical data, to the extent it is available.	Some effort, and judgement, will be required to determine how many and which scenarios to study (see Issue discussion below).
Would need an alternative mechanism to handle paths (and correlations) not included in history, due to network augmentation or re-configuration.	Uses current auction model configuration, so no problems if paths don't have history.
Margin offsets are inexact, resulting in coarse portfolio margining.	As the entire auction is simulated each time, this is inherently a portfolio-based approach.
Need to maintain and apply very large correlation matrices (both peak and off-peak).	Simulating large numbers of scenarios is likely to be highly computationally intensive, unless workable simplifications are available.

***Recommendation**: Simulation should be used for calculating the potential exposure (and hence initial margin) of FTR positions, provided performance issues can be addressed satisfactorily.*

Issue: Determining Scenarios to Simulate

An immediate concern when considering what scenarios to simulate is that this could be something of a black art, that nevertheless has real consequences for the credit cover that most be posted by participants. However, there is a solid basis of information upon which scenarios can be determined.

The major source of perturbation from expected outcomes, as represented by the results of the FTR auction, is likely to be unplanned transmission outages (planned transmission outages will be included in the transmission topology incorporated in the auction, and planned generation outages should be factored into participants bids). There are well-developed historical performance indices for the incidence and duration of transmission equipment outages – the System Average Frequency Index (SAFI) and System Average Duration Index (SADI) – which are compiled by NERC, and incorporate a diversity of information related to types of equipment, region and other factors.

While sound engineering judgement will always be required in their application, these indices provide an empirical basis for the selection of a range of appropriate scenarios for simulation.

6.2.2 Delivery Margin for FTRs

Section 4.5.1.2 recommends that delivery margin requirements for FTRs be calculated as "the auction value of the FTR portfolio, as marked-to-market, minus the x% worst-case payout of the FTR portfolio (which could be negative)." The auction value of the portfolio can be calculated explicitly after each auction. The offset based on the 'x% worst-case' payout of the FTR portfolio (referred to as 'Worst Projected FTR Payout' in this discussion), however, will need to be estimated. But how?

FTR Payouts are cashflows in the day-ahead market, based upon the price differential between two nodes. One basis for calculating Worst Projected FTR Payout, therefore, is to look at historical days in the DAM. Two particular options based on an analysis of historical DAM activity are described below, though others would no doubt be possible:

6.2.2.1 DAM Path/Price Approach

This approach calculates the Worst Projected FTR Payout based on a statistical analysis of historical prices per path:

- Determine the *x*% worst-case *n*-day payout for each path (where *n* is the number of contiguous days in the FTR), with separate totals for peak and off-peak.
- Sum these FTR payouts for the portfolio held by each participant.
- Apply credits based on the risk correlations between paths (this assumes a satisfactory risk correlation matrix for this purpose can be derived).

This approach requires the creation of a large database of price and risk correlation information, though much of this data could be gathered and/or calculated well before the time of the credit analysis, and would be re-usable in the analysis for each participant.

The usage of the worst-case for each path is likely to result in very conservative offset levels. While credits based on risk correlation may adjust for this in part, they are likely to be inexact, especially given that a number of simplifications and limiting assumptions will be required to keep the correlation matrix to a reasonable size.

By looking at *n*-day payout, this approach does take proper consideration of diversity between days (see further discussion below). During the delivery period it would be possible to release delivery margin. This would require the re-calculation of the Worst Project FTR Payout each day, based on the *x*% worst case (*n*-*d*)-day payout (replacing the original *n*-day payout), where *d* is the number of days into the delivery period – or a suitable approximation in lieu (which, given its use for risk management purposes, would need to be conservative)

6.2.2.2 DAM Participant/Portfolio Approach

This approach calculates the Worst Projected FTR Payout based on the application of a participant's FTR portfolio against a series of historic days, and performing a statistical analysis on the results:

- Select a series of historical days in the day-ahead market.
- For the FTR portfolio held by each participant, determine the FTR payout that would have occurred on each of these days.
- Take the *x*% worst case of payouts.
- Multiply by *n*, the number of days in the FTR.

As it is based upon the price differential between nodes in actual historic days, this approach takes proper consideration of the portfolio relationship between different FTR paths. However, it is based on a limited set of historical days, which may limit its usefulness as an indicator of future performance.

Additionally, this approach does not consider diversity between days. The Worst Projected FTR Payout is calculated based on the *x*% worst payout for a single day, and multiplying this by the *n* day period of the FTR. However, the worst event won't happen every day over an *n*-day contiguous period, and *n* times the worst single-day payout will produce a smaller offset than the worst *n*-day payout. The end result is that allowed offsets, and therefore delivery margin taken, could be more conservative than they need to be. One advantage, though, is that during the delivery period it is valid to return 1/n of the delivery margin as each day's FTR Payout is 'delivered', as every day stands alone, rather than re-calculating for the remaining *n*-*d* days.

An alternative approach, that addresses some of the issues above, is available, but only in the case where a simulation-based approach is taken to determining initial margin (see Section 6.2.1).

6.2.2.3 Simulated Scenario Approach

Under the simulation approach to determining initial margin, the FTR auction calculation is re-run for a number of different scenarios. The nominal purchase price for an FTR in each scenario equals the expected FTR Payout over the *n*-day period of the FTR, based on the initial conditions defined for the scenario. The Worst Projected FTR Payout for each participant could be determined by:

- Determining the payout for the participant's portfolio in each scenario.
- Calculating the *x*% worst payout across these scenarios.

This would take account of both the risk correlation between FTRs in the portfolio, and diversity across days.

The scenario simulations required for this approach will already be available from the determination of initial margin, and therefore would not produce a requirement for additional processing. However, when returning delivery margin there are no scenario calculations for *x*% worst case (*n*-*d*)-day payout (where *d* is the number of days into the delivery), as this is not required for initial margin purposes. It would be necessary, therefore, to use an approximation to determine delivery margin release, or to perform additional simulations.

***Recommendation**: The Simulated Scenario approach should be used for determining the Worst Projected FTR Payout component of delivery margin requirements - predicated upon a simulation approach also being followed for initial margin calculation.*

***Investigate**: If a simulation approach was not followed, further investigation of alternative approaches is required – examining both their accuracy and processing complexity.*

6.3 Settlement/Variation Margin

6.3.1 Variation Margin for FTRs

Assuming variation margining is used, it is possible to determine variation margin requirements for each participant as soon as auction results are finalized. Once these amounts are calculated it makes no sense, from a credit risk point of view, to then wait until the cash market bill is issued.

***Recommendation**: The settlement and billing of variation margin should be de-coupled from the cash market billing cycle, with bills issued as soon as variation amounts are calculated, and payable within one (1) business day, or before the next auction, whichever is sooner.*

The option premium component of the purchase price for FTR Options (discussed in Section 4.3.1.4) would form part of these amounts. The final settlement of FTRs 'in delivery' would not be affected, as these settle against prices in the day-ahead market, and therefore should be included on the same bill as day-ahead market amounts.

6.3.2 Payment Imbalance Account

As discussed in Section 4.3.1.3, an imbalance account would be required to hold any settlement imbalances resulting from the variation margining process (prior to delivery into the DAM). This could conceivably by the same account presently used for managed the under/over-funding of FTRs delivered into the DAM.

The size of the float, or overdraft facility, required on this account will be dependent on how far 'in the red' the account could go. This in turn will be dependent on the extent to which participants hold short positions – in particular, whether holders of ARRs are deemed to hold short positions, and are variation margined – versus PJM (or more precisely, the pool of congestion revenues in the DAM)[69] is deemed to hold them. If ARR hold the short side of FTRs, the size of potential imbalances will be less. It would also stop potential aberrations such as a participant holding a balanced ARR and FTR position, yet paying/receiving variation amounts on only side of the position.

***Recommendation**: ARR holders should be deemed to hold the short-side of FTRs, and be margined accordingly.*

[69] This is an important distinction, as it should be reinforced that PJM takes no position in the market. By virtue of the way the FTR auction mechanism works, the congestion revenues in the DAM fund FTR auction payouts, and this manifests itself in the form of auction-to-auction cashflow imbalances if variation margining is used. At no point does PJM have a beneficial interest in the price of FTRs.

7 PROCESS AND CALCULATIONS: CASH MARKETS

This section discusses some of the key options, and their associated pros and cons, related to:

- Changes that could be made to PJM's business processes to reduce credit risk. This addresses not just explicit credit risk management processes, but also process options in other areas which have a credit risk impact.
- The methodology for performing credit risk calculations.

The focus of this section is on PJM's cash markets.

7.1 Day-Ahead Market

Bid/trading limits should apply to all markets where past performance is not necessarily an indicator of future performance, and to all participants in those markets.

Ideally, credit limits would form part of the day-ahead market auction algorithm, allowing participants to enter bids and offers without limitation, but not scheduling them beyond their posted credit cover. If this is impractical to achieve, however, it is recommended (as discussed in Section 4.4.2) that bid limits be imposed initially, with collateral requirements re-assessed after the day-ahead auction is run, based on actual amounts scheduled.

7.1.1 Bid Limits

The credit risk associated with the day-ahead market is the risk that the day-ahead quantity is not 'delivered', and therefore needs to be offset in the real-time market (see "Issue: Past Performance as an Indicator of Future Performance" in Section 4.5.2.2). As such, it is a based upon the price differential between the day-ahead and real-time markets.

When new bids are posted, the basic credit check is that:

Current Exposures + New Exposures (from the bid) ≤ posted credit cover

…where:

Current Exposures = Actual Exposures + Resettlement Exposures + historically-predictable Future Potential Exposures + any controllable Future Potential Exposures already accepted.

New exposures can be determined by:

New Exposures = $\sum$(DAQ(submitted).Px(ΔP))

…where:

DAQ equals the submitted bid quantity for the participant at a node.

Px(ΔP) is the projected value of ΔP, the price deferential between the real-time and day-ahead markets at the node, to x% confidence (based upon a probability distribution of historical prices).

The exposure for the participant is summed across all its bids.

Px(ΔP) could be based on:

- A distribution over the day: This would result in one Px(ΔP) value for the node being multiplied by the quantity in all hours, which would give fairly high values for low demand times of day.
- A distribution by hour: This would result in an hourly Px(ΔP) value for the node being multiplied by the quantity in that hour, and each hour then summed. This would tend to be more accurate, but computationally intensive.

- A distribution by time-block (e.g. 4 or 6 hours): This would result in a value being determined for P*x*(ΔP) for the node for the time block, and multiplied by the quantity for that time block. This would account for intra-day price variations (at peaks, etc.), while being less computationally intensive.

***Investigate**: the options for calculating Px(ΔP), in order to determine the materiality of the differences between solutions, versus computation requirements.*

One outstanding question is: if a participant submits DAM bids for multiple days into the future, how much of it is validated on submission? There are two basic options:

Table 3 - Comparison of Options for DAM Bid/Offer Validation

Validate all Submitted Days	Validate Bids for Next Auction Only
All days that are submitted are validated.	Only bids/offers for next DAM auction are validated.
Credit cover must be sufficient to cover all days submitted.	Credit cover must be sufficient only for the next auction.
Bids submitted in advance are already validated and cannot be rejected as the DAM approaches (other than for reasons of default).	Bids submitted in advance could be rejected as they come up to the DAM and are validated.

Options in-between these two are possible, such as validating a week at a time. However, any option that does not require the validation of all days has the issue that bids submitted in advance could be rejected as the actual auction approaches. Ultimately there is no right answer to this question, as adequate credit protections can be put in place for each – the issue is one of participant preference.

***Investigate**: participant preferences for timing of DAM bid/offer validation.*

7.1.2 Post-Auction Exposure Recalculation

Once the day-ahead market is run, but before the real-time market has been run (and therefore real-time price is known), the DAM exposure can be recalculated, replacing bid quantities with scheduled quantities, but otherwise using the formulae above.:

New Exposures = ∑(DAQ(scheduled).P*x*(ΔP))

Once the real-time market is run, values are available (even if not final) for all prices and quantities, and these 'controllable Future Potential Exposures' should be replaced by Actual Exposures.

7.2 Reassignments of Settlement Obligation

PJM can reassign their settlement obligations in the day-ahead and real-time energy markets – through eSchedules – as well as in other markets such as regulation and synchronous reserve. For the sake of simplicity, these will all be referred to as 'eSchedules' in the following discussion. General recommendations regarding management of the credit risk associated with eSchedules are discussed in Section 4.2.3. The sub-sections below address process and calculation specifics.

7.2.1 Submission Limits

Section 4.2.3 recommends the use of submission limits upon eSchedules, to ensure that:

Current Exposures + New Exposures (from eSchedule) ≤ posted credit cover

…where:

Current Exposures = Actual Exposures + Resettlement Exposures + historically-predictable Future Potential Exposures + any controllable Future Potential Exposures already accepted.

New exposures can be determined by:

New Exposures = ∑(ESQ(submitted).P*x*(MP))

…where:

ESQ equals the eSchedule quantity for the participant at a node.

P*x*(MP) is the projected value of MP, the market price for the particular market the eSchedule relates to at that node, to x% confidence (based upon a probability distribution of historical prices).

The exposure for the participant is summed across all hours in the eSchedule, and all eSchedules submitted.

As with bid limits on the day-ahead market, it must be decided whether P*x*(MP) is based on a distribution over the day, by hour or by time-block, with the same decision criteria applying.

> ***Investigate****: the options for calculating Px(MP), in order to determine the materiality of the differences between solutions, versus computation requirements.*

Treatment would be slightly different for eSchedules submitted after the fact, with the actual market price (MP) at that location, in each hour being used instead of P*x*(MP). i.e.:

New Exposures (submitted after the fact) = ∑(ESQ(submitted).MP)

7.2.2 Cancellation of eSchedules

Participants also have the ability to cancel previously submitted eSchedules. This has the effect of increasing the exposure of one party, and decreasing the exposure of the other. This could potentially cause the party with increased exposure to exceed the credit cover it has posted. As such, it is necessary to validate not just the submission of eSchedules, but also their cancellation. The same formulae could apply as for submission limits, except ESQ would be the opposite sign to when it was submitted.

> ***Recommendation****: Perform credit validation on any transaction to cancel an eSchedule. If the validation causes either party's exposure to exceed their credit cover, reject the transaction.*

The same considerations would apply to any eSchedule cancellation resulting from default in PJM's markets. eSchedules that decrease a participant's exposure could be cancelled only to the extent that the participant has credit cover posted in excess of requirements (some of which might result from the cancellation of other eSchedules that serve to increase their exposures).

This does raise the possibility that participants might be obliged to make good on an eSchedule in the future, even though the underlying bilateral contract between the two parties has terminated due to default. e.g. if Generator A has a 100MW eSchedule with LSE B, and B defaults on the underlying deal. However, it must be remembered that, because this eSchedule served to reduce its exposure, LSE B has been extended credit – in this case, allowed to incur additional exposures – as a direct result. It is not appropriate to cancel the eSchedule and drop these exposures back onto the socialized risk pool, as they are rightly the consequence of a bilateral deal between A and B.[70]

If it is desirable to limit situations such as that illustrated above from occurring, some broad options include:

- Prohibit the entry of eSchedules that stretch beyond the next day. This would prevent participants receiving exposure offsets for eSchedules beyond that period. However, it would also limit the convenience which participants currently have to enter eSchedules for time periods well into the future.
- Deem that, no matter the period on the eSchedule submission, only the next day – plus any current and after-the-fact days included – are considered for credit purposes. These would be the only days validated, and as a result, the only days for which exposure offsets related to eSchedules would be applicable. Any days beyond this could be cancelled without credit impact. Each day, one more day from the eSchedule would be 'peeled off' and validated, with the eSchedule being rejected from that point forwards if credit validation failed.
- Don't allow credit exposure offsets based on eSchedules. i.e. an eSchedule could serve to increase a participant's exposure, but not to decrease it. This would eliminate any problems with cancelling the eSchedule, but would result in a number of other aberrant effects. e.g. if a participant has offsetting eSchedule positions. It would be virtually impossible to write exception rules around all such potential aberrations.

The first and second options are more readily achievable. The third option would have a more substantial market impact and could be quite complicated to implement (given all the special cases likely to be necessary), making that option inadvisable.

7.3 Settlement

7.3.1 Billing Period

In Section 4.3.2.1 it is recommended that the cash market billing period be reduced to a duration of one (1) week or less. The two most obvious choices are a billing period of one week or one day. These are contrast in the table below:

Table 4 – Comparison of Weekly vs. Daily Billing Period

Weekly	Daily
A bill is issued for each trading week.	A bill is issued for each trading day.
The majority of settlement amounts are determined at sufficient granularity.	The majority of settlement amounts are determined at sufficient granularity.
Would require changes to PJM and participant processes and systems.	Would require greater changes to PJM and participant processes and systems.
Does not necessarily require weekend settlement/ credit operations.	Would require weekend settlement/credit operations, or catch-up afterwards.

[70] There are a also number of gaming opportunities that would be possible if it was permissible to cancel eSchedules in such circumstances.

Weekly	Daily
Higher credit exposure, and hence credit cover requirements.	Lower credit exposure, and hence credit cover requirements.
Common practice in a number of electricity cash markets, including some in the US.	Not common in electricity cash markets; common practice in derivative markets.

Daily billing is a superior solution from a credit point of view, as well as being more conservative of collateral, and should represent an ultimate goal. However, a move to daily billing would involve substantial changes to current settlement practices, particularly for those participants not presently engaged in trading cleared derivatives. A move to weekly billing would achieve substantial credit benefits vis-à-vis monthly billing, without as profound a process change. As a result, this report does not go so far as to recommend the immediate adoption of daily billing, but certainly would not discourage it, and recommends that the billing period be reduced to at least weekly.

7.3.2 Payment Lag

Section 4.3.2.2 recommends that PJM reduce the lag between the end of the billing period and payment to five (5) calendar days or less. The table below contrasts two such options – T+1 and T+3 settlement – with current practice of payment ~20-22 calendar days in arrears.

Table 5 – Comparison of Payment Lags

T+1	T+3	T+16
Payment the business day after the trading day.	Payment three business (five calendar) days after the trading day.	Payment 16 business (~22 calendar) days after the trading day
Settlement based on actuals (e.g. metering) may require approximation, and near-term true-up.	Actuals should be sufficiently accurate that near-term true-up is not likely.	Actuals should be sufficiently accurate that near-term true-up is not likely.
Lowest exposure and credit cover requirements.	Lower exposure and credit cover requirements.	High exposure and credit cover requirements.
Best practice in the financial markets	Practice common in weekly-billed markets	Practice common in monthly-billed markets
Significant change to existing PJM and participant processes.	Some change to existing PJM and participant processes.	No change to existing PJM and participant processes.

T+1 settlement is the best of these options from a credit risk point-of-view. Even if settlement inputs requiring actuals (such as metering data) are not as accurate as they could be, it is better to have data that is 80% accurate (which can be compensated for) than no data. However, T+1 settlement would also involve significant process complexities, including:

- multiple resettlement runs performed soon after initial settlement;
- use of substitute data, in lieu of metering, in initial settlement runs – such as state estimator data, scheduled quantities, etc. – with replacement in later runs.

As a result, T+3 settlement may be more achievable in the short-term, with an objective to move towards T+1 settlement at a later juncture.

8 UTILIZATION OF AN EXTERNAL CLEARING HOUSE

One of the options available to PJM for the management of credit risk is to engage the services of an external clearing house, which would replace PJM as the entity responsible for all credit functions and a number of settlement functions, such as funds transfer[71]. Given the importance of these functions to the effective operation of PJM markets, any organization performing them must be experienced, reputable and robust. Therefore, for the purposes of this analysis, an 'external clearing house' shall be taken to be an organization accredited as a Designated Clearing Organization (DCO) by the US Commodity Futures Trading Commission (CFTC).

This section discusses the disciplines likely to be enforced by a DCO if it were to clear PJM's markets, key considerations in the selection of an external clearing house, and some of the more logical candidates for this role.

8.1 Clearing House Disciplines

Many of the key structures, policies and processes utilized by clearing houses have been discussed in Section 3 through 7, and as such, their mechanics (e.g. how variation margining works) are not re-explained in detail below. This section focuses on the typical or likely application of these disciplines by a DCO. Also discussed are some additional disciplines commonly utilized by DCO's, and not previously discussed, such as the use of credit intermediaries.

8.1.1 DCO as Central Counter-Party

DCOs act as the central counter-party to all trading in the markets they clear.

In order for PJM to appoint a DCO to clear its markets, it is likely that PJM itself would need to be the central counter-party to this trade, giving it the legal right to then novate these positions off to a DCO. PJM is investigating whether this would impose an obligation on PJM itself to seek registration with the CFTC as an exchange or DCO[72].

8.1.2 Credit Intermediaries

DCOs typically utilize financial intermediaries known as Clearing Members (CMs) – or in the US, Futures Commission Merchants (FCMs). Under this structure, Clearing Members are responsible to the clearing house for financial performance, with participants, in turn, responsible to their CM (see Figure 7).

This structure provides an intermediate layer of credit protection between the participant and the clearing house, reducing the risk of both. It is also deliberately structured to diversify the risk placed under its care, so that risk is not just transferred, it is reduced. The intermediated clearing model can also serve to promote competition for various services, and encourage flexibility in meeting individual participant needs.

[71] Given the complexity of PJM's settlement calculations, it is unlikely that the actual calculation of settlement amounts would be taken on by an external entity, as they would effectively need to replicate PJM's settlement system.

[72] The purpose of such registration would not be to extend the current scope of PJM's functions into futures markets, simply to protect existing practices.

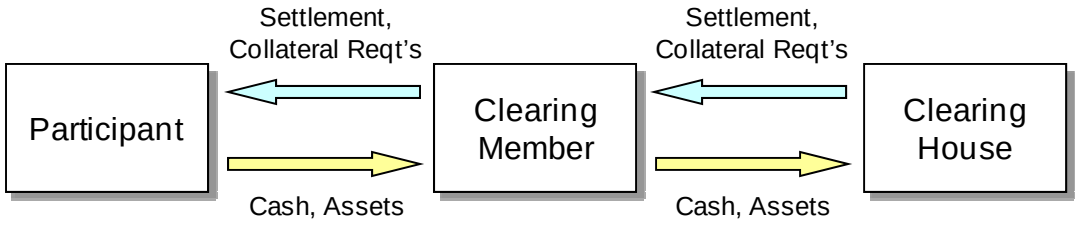

Figure 7 – Intermediated Clearing Structure

8.1.3 Credit Cover and Collateral

DCOs operate on the principle of 'full collateralization', requiring collateral to be posted to cover a participant's exposure to a given degree of statistical confidence. Unsecured credit is not accepted.

Collateral is generally levied to cover the *x*% worst exposure over an *r*-day risk period, with *x*% typically equaling 99% or even 99.7%. For liquid products that trade on a daily basis (such as many of the futures products cleared by DCOs), the risk period is typically one day, this being the time between when the position was last marked-to-market and when it could be liquidated in the event of participant inability to meet a margin call. For less liquid products, or those with longer billing periods, payments lags, gaps between trading or time to liquidate, *r* will be longer. e.g. many clearing houses assess initial margins for electricity futures based on a 'worst two-day move' because of liquidity concerns.

The length of this *r*-day period will be an extremely important consideration to any clearing house, as the longer this period, the greater the risk exposure. While collateral protects the clearing house to *x*% confidence, this is ultimately a statistical measure and the clearing house guarantee structure takes on the residual (100-*x*)% risk. Obviously a clearing house will seek to reduce the quantum of this risk, and will have a threshold for the amount of residual risk it is willing to take into its clearing structure, versus the revenue opportunity it sees from taking on the risk. PJM has sizable business to offer, with trade exceeding $20 billion per annum. It is important to remember, however, that many DCOs, such as NYMEX and LCH, clear notional trade worth trillions of dollars per year. While no doubt keen for PJM's business, prudent clearing houses will also be keen to ensure that their existing business is not put at undue risk. Though the particular requirements of a clearing house will no doubt be the subject of negotiation, the discussion below indicates some of the factors that are likely to be considered.

8.1.3.1 Cash Markets

As discussed in Section 4.5.2, the total exposure at any point in time is comprised of:

- **Actual Exposures**: These equal amounts billed but not paid, and calculated but not billed, and as such their quantum is driven by the billing period and payment lag. PJM's current billing cycle, of monthly billing paid 20+ days in arrears, will almost certainly be unacceptable to an external clearing house. Daily billing paid 2-3 calendar days in arrears would almost certainly be acceptable. A key topic for discussion with external clearing house candidates will be whether a middle-ground solution, such as weekly billing paid five calendar days in arrears, is acceptable to them from a risk perspective.
- **Future Potential Exposures**: These exposure are driven by the *n*-day contiguous period covering Incurred Amounts, Liabilities for Next Assessment Period, Liabilities for Time-to-Post/Time-to-Remedy, Liabilities for Time-to-Transfer (see discussion in Section 4.5.2.2). An external clearing houses will no doubt wish to quantify this risk, and keep *n* as low as possible.

An important consideration in this regard is the treatment of PoLR default, where there is no fallback entity to take on its retail load, and as such, the PoLR must be allowed to continue amassing liability for some indeterminate period (as the Time to Transfer is undefined). It is highly unlikely that any clearing house will agree to take on such an open-ended risk (and one that, at its origin, is really a regulatory risk). As a result, this risk would continue to be borne by PJM's members (perhaps beyond a threshold that the clearing house agreed to take on), a mechanism would be required for limiting the risk, or some alternate risk management mechanism (such as a mutual insurance scheme) would be required.

- **Resettlement Exposures**: The size of the aggregate resettlement exposure for a participant is driven by three factors – the accuracy of initial settlement, the frequency of resettlement and the cut-off beyond which resettlement can no longer occur. The accuracy of settlement is often a trade off with timeliness, making this a difficult topic to address as increased timeliness reduces Incurred Amounts and Future Potential Exposure. However, clearing houses will be keen to see frequent resettlement and a relatively short resettlement window, in order to limit the length that individual settlement amounts can run, and therefore the total that can accrue over the entire resettlement period.

An external clearing house would be likely to have significant dependencies on PJM for the provision of information to support the calculation of these exposures. Actual exposure information will be required from PJM's settlement process. Future Potential Exposures and Resettlement Exposures will either need to be determined by PJM, or require an extensive database of historical settlement information to be shared with the clearing house, to enable it to perform the statistical analysis to determine these amounts.

8.1.3.2 Forward Markets

Clearing of forward markets is the 'stock in trade' of a DCO. FTRs, however, are not 'stock standard' forward products, and an external clearing house is likely to have some concerns regarding the infrequency of FTR auctions, and how significantly the value of a participant's position could move between auctions. It is likely that an external clearing house would seek for FTR auctions to be carried out more frequently, with the exact period a subject for negotiation. An external clearing house may also have concerns regarding the perceived illiquid nature of FTRs, though it should be possible to address these (at least to some extent) through explanation of the auction mechanism, and the fact that its transfer function allows reconfiguration of paths (i.e. for a path to be sold it does not require a buyer for the exact same path).

Initial margin will typically be assessed based on the '*x*% worst potential move' between when the position was last marked-to-market and when it could be liquidated in the event of default (i.e. over two auction periods in this case). As discussed in Section 6.2.1, the preferred methodology for determining this move, and hence initial margin requirements, is through the use of simulated scenarios. An external clearing house would need to endorse this methodology and would almost certainly rely upon PJM to carry out these simulations.

It is fairly standard practice for clearing houses to assess delivery margin on long positions based on their full value, as marked-to-market, at the time the contract ceases trading in the forward market and moves to delivery. The more unusual aspect of FTR margining is the offset to this delivery margin based on the '*x*% worst-case FTR Payout' – a conservative projection of the cash market payout of the FTR. Given this is based on calculation to a conservative degree-of-confidence, it should not pose serious risk concerns for a clearing house, though it is likely the clearing house will need PJM to calculate these offsets, or provide it with an extensive database of historical information, depending on the method of calculation (see discussion in Section 6.2.2). External clearing houses are also likely to be conservative with respect to collateral return during the delivery period.

8.1.3.3 Portfolio Margining

Portfolio margining is standard practice for a typical DCO, and it should be possible for an external clearing house to provide margin offsets based upon the risk correlations between a participant's positions in the various markets operated by PJM.

The clearing house may also be able to provide offsets between PJM positions and positions in related markets, such as:

- positions in other RTO-operated cash and forward markets cleared by the DCO.
- electricity forward/futures contracts
- contracts for fuel sources, such as natural gas, coal and fuel oil
- other commodity contracts with a significant linkage to electricity, such as weather and emissions credits.

This, of course, is only possible to the extent the clearing house has meaningful standing in those markets. e.g. it is no use a clearing house promising offsets against natural gas futures positions if it does no meaningful clearing of natural gas.

Exposures will typically be calculated, and compared against posted collateral, on a daily basis, and in some cases intra-day. For the daily process, collateral requirements are typically calculated overnight, and any additional collateral required must be posted by a specified time in the morning (10a.m. being fairly typical). Failure to post must generally be remedied within one (1) hour to avoid default proceedings including seizure of collateral and liquidation of positions.

It is also fairly standard practice for clearing houses to maintain two different margin levels – initial margin and maintenance margin. Whenever a margin call is made, margin levels must be restored to be equal to (or greater than) initial margin. Margin levels can fall to the maintenance margin level, however, before a call for additional collateral is made. Depending on the type of participant, initial margin levels are typically set between 1.1 and 1.35 times maintenance margin, which is set based on exposure. The leeway provided by this differential can reduce the frequency with which participants are asked to adjust their collateral levels.

8.1.4 Forms of Collateral

Almost all clearing houses accept cash, treasury instruments and letters of credit (LCs) as collateral. Typically there is a limit on the percentage of assets that may be deposited in the form of LCs (50% is common). Instruments other than cash are generally also subject to a 'haircut' discount to their face value, to represent potential loss of value if the instrument needs to be liquidated. Some clearing houses also accept other assets as collateral, such as holdings in money market mutual funds.

8.1.5 Liquidation Rights

It is standard practice in most forward markets to liquidate the outstanding positions of participants in default. FCM's are required to have the authority to liquidate their customer's positions in the event of default, and the clearing house will similarly have authority to liquidate the positions of a defaulting FCM.

8.1.6 Settlement/Variation Margining

External clearing houses will typically settle on a T+1 (next business day) basis. i.e. they have a payment lag of one business day. The frequency with which the settlement process is carried out (i.e. the billing period) will depend on the particulars of the market, though for many of the markets cleared by DCOs this also takes place on a daily basis.

8.1.6.1 Forward Markets

The settlement of forward markets is typically carried out by clearing houses on an incremental basis, through a variation margining (mark-to-market) process. This is driven by the frequency with which the market transacts, which is a principal concern for FTR markets, as there is no use doing daily mark-to-market if there is no new settlement price to mark the position to. As discussed previously in the section on collateralization, an external clearing house may seek to have PJM increase the frequency of FTR auctions.

When compared to most forward markets, which operate on the principal of zero sum gain, a unique consideration for FTR auctions is that the net value of all outstanding contracts may vary from auction to auction. The clearing house will want all positions to balance (i.e. monies out = monies in), meaning that PJM will need to manage an imbalance account that holds this difference[73].

8.1.6.2 Cash Markets

As discussed in the previous section on Credit Cover and Collateral (Section 8.1.3), an external clearing house would almost certainly require a reduction in both the billing period and payment lag in PJM's markets.

With respect to billing period, the majority of payments and charges in PJM's markets are determined on a trading period or daily basis, and are therefore compatible with a billing period as short as one day. For those amounts with a period greater than the billing period, an agreed methodology for billing would need to be determined (some options are discussed in Section 4.3.2.1).

More problematic is the reduction of payment lag. Payment lag could probably be reduced to T+3 (five calendar days) without significant changes to settlement process. Reduction to T+1, however, may require initial settlement to be performed with approximate or alternate data (e.g. real-time market dispatch quantities in place of metering) with near-term true-up as soon as more accurate measurement data becomes available. This would require PJM to make significant changes to its settlement methodology. Payment lag will no doubt be an extremely important topic of discussion with any external clearing house candidate.

Given the complexity of cash market settlement calculations, and the existing capabilities of PJM in this area, it is highly likely that PJM would continue to perform these calculations, and send the results to the clearing house, which would then be responsible for the remainder of the settlement, billing and payment process.

[73] This account is effectively held on behalf of the pool of settlement residues in the day-ahead market, which fund the FTR Payouts, but which are not available until the day-ahead market actually transacts.

Clearing houses typically maintain an extensive network of relationships with money-market banks, to facilitate the efficient daily and intra-day movement of funds. Settlement amounts must be paid by the FCM by a specified time each morning (10a.m. being typical), though some FCMs will give their customers a little longer (the daily close of the Fedwire system being common)[74]. Failure to pay must generally be remedied within one (1) hour to avoid default proceedings. One advantage of this approach is a single, standardized settlement across all markets cleared by the clearing house.

8.1.7 Trading/Bid/Submission Limits

Clearing houses will generally require an exchange to impose some form of credit risk limit validation on all unpredictable or controllable exposures. i.e. any exposure that cannot be predicted from history, and is at the participant's discretion to incur or not. PJM already imposes such limits on FTR trading. Other areas where similar limits would need to be imposed are day-ahead trading (not just virtual bidding) and eSchedule submissions.

In some markets these limits will be calculated based on free collateral (i.e. collateral posted and not otherwise used to cover existing exposures), the aim being to prevent any new exposure being incurred that would cause the exposure of the participant's net position to exceed their posted margin. Some other markets, rather than calculate such limits directly, allow the participant's FCM to specify daily risk limits on trading activity. In either case, PJM would need to make changes to its systems to validate against these limits, and implement appropriate interfaces to the clearing house (to obtain details of collateral holdings and exposures) or the FCMs (to allow entry of daily limits).

Specifics regarding how some limits are imposed – such as whether credit validation can be applied in the auction process, or on bids prior to auction – would require further investigation, and be the subject of negotiation between PJM and the clearing house.

8.1.8 Participation

Clearing houses impose relatively strict requirements on their clearing members, related to capital adequacy, liquidity, etc., but do not impose any requirements directly on market participants who are not clearing members[75]. In order to manage its risk an FCM may require its customers to display a certain wherewithal to participate in the market, but this is a matter between the participant and its FCM.

In order to ensure that the clearing of PJM's markets receives favorable regulatory treatment under the Commodities Exchanges Act (CEA), it may be necessary for participants in some products (e.g. FTRs) to be qualified as Eligible Commercial Participants (ECPs), or to receive a no-action letter from the CFTC waiving this requirement, or stating that they consider this requirement satisfied.

8.1.9 Trade Guarantee Structure

External clearing houses will typically have an extensive trade guarantee structure to protect the integrity of the clearing house in the event of default. These often include a Guarantee/Default Fund, insurance, and other mechanisms. Examples of the NYMEX and LCH guarantee structures are provided in Section 5.3.

[74] FCMs cannot generally allow longer for customer payment, as this crosses the line from cashflow management to margin lending, which FCMs are not allowed to engage in (though their banking affiliates may).
[75] Many of the futures markets associated with these clearing houses require traders to undertake training and satisfy defined qualification criteria, but these are provisions of the exchange (as opposed to the clearing house) and would not automatically apply to PJM.

8.1.10 Financial and Market Surveillance

Clearing houses typically conduct extensive financial and market surveillance.

- *Financial Surveillance*: Monitors the financial integrity of clearing members and customers, in order to ensure the integrity of the Clearing House. Functions performed include periodic audits of clearing members, analysis of position limits, review of margin call submissions, etc.
- *Market Surveillance*: Normally conducted in conjunction with the exchange being cleared, market surveillance ensures that the markets are operating in an orderly fashion, and free from distortion or manipulation. Functions performed include ensuring the convergence of cash markets with forward markets, and monitoring of large trader positions.

8.2 Key Considerations in Selection of an External Clearing House

A number of factors should be considered by PJM in the assessment of potential candidates to clear its markets, should this be an option it wishes to pursue further.

Relevant Expertise and Experience

- Does the candidate have expertise and a proven track record in providing clearing services?
- Does it have experience providing clearing services for markets/exchanges other than its own?
- Does the candidate have expertise in clearing electricity and other energy products?
- Does the candidate have any experience dealing with PJM, or products that cover the PJM market? What is its experience dealing with PJM participants?

Financial Safeguards

- Does the clearing house have a strong set of credit risk management processes?
- Does it have a robust system of trade guarantees?
- Does it have a solid capital structure?
- Is the clearing house rated by a credit rating agency, and if so, what is its rating?

Credibility of the Clearing Solution and Provider

- Is the clearing solution proven? (Note: it can be assumed that entities accredited as DCOs by the CFTC have a solution sufficiently proven to satisfy the CFTC in this regard).
- Is the solution dependent on external factors, and if so, how reliable are they? Is the solution dependent on variables outside the clearing provider's control?
- How financially solid is the candidate? What is its financial history? If revenues from clearing PJM's markets are less than expected, what guarantees are there that the candidate can fulfill their contractual obligations?

Ability to Provide Margin Offsets

- Does the candidate offer collateral offsets of PJM cash and forward market positions against positions the participant might have in other markets? If so, in which markets?
- What standing does the candidate have in these markets to credibly claim it can provide such offsets?

Acceptance of all PJM Members

- Is the candidate able to provide clearing services which cover all PJM participants? If not, what participants are excluded, and on what basis? How would these participants need to be handled?
- Does the candidate treat all PJM participants on an equal basis?

Ability to Accommodate PJM Operational Practice

- What level of change does the candidate require to PJM's current processes?
- What level of system change is the candidate's solution likely to require?

8.3 High-Level Discussion of Potential Candidates

The CFTC presently lists 15 organizations as DCOs, of which 11 currently have a status of 'Registered' (the others having been declared 'Dormant' or having 'Vacated' their registration). These entities are[76]:

- *AE Clearinghouse, LLC*: A subsidiary of The Actuarials Exchange, LLC
- *Chicago Board of Trade (CBOT)*: CBOT contracts are now cleared by the CME Clearing House under a clearing services agreement.
- *The Clearing Corporation (TCC)*: Formerly the Board of Trade Clearing Corporation (BOTCC). Lost the clearing business of CBOT and presently clears a number of smaller markets.
- *CME Clearing House*: The clearing house of the Chicago Mercantile Exchange (CME). Also clears CBOT.
- *HedgeStreet, Inc.*: Clears exclusively for its affiliated contract market, Hedge Street, Inc.
- *ICE Clear US, Inc.*: Formerly the New York Clearing Corporation (NYCC), responsible for clearing trade on the New York Board of Trade (NYBOT), which was acquired by the Intercontinental Exchange in 2007. ICE intends to take back clearing of its contracts some time in 2008.
- *Kansas City Board of Trade Clearing Corporation*: The clearing house for the Kansas City Board of Trade.
- *LCH.Clearnet Ltd.*: LCH.Clearnet Ltd., based in London, clears a number of markets worldwide, and holds DCO status in the US. Amongst other markets, it presently clears products for ICE, in the US and UK, though is due to lose that business sometime in 2008.
- *MGE Clearing House*: The clearing house of the Minneapolis Grain Exchange (MGE).
- *NYMEX Clearing House*: The clearing house of the New York Mercantile Exchange (NYMEX). Also clears a few other markets, some of which NYMEX has an ownership interest in.

Of these markets, only three presently have, or are likely in the near-term to have, a material presence in the energy markets: NYMEX, LCH.Clearnet and ICE Clear US. These are discussed in a little more detail below:

[76] Source: http://services.cftc.gov/SIRT/SIRT.aspx?Topic=ClearingOrganizations. The CFTC also lists two organizations defined as Multilateral Clearing Organizations (MCOs) – NetThruPut and NOS. Both are, in the words of the CFTC, "supervised by a foreign financial regulator that the CFTC, or one of several other United States financial regulators, has determined satisfies appropriate standards." It is also possible that other existing clearing entities registered by the SEC, such as the Depository Trust & Clearing Corporation (DTCC), could seek accreditation by the CFTC.

8.3.1 NYMEX

NYMEX is the world's largest clearing house for energy products – principally clearing contracts traded on NYMEX, or traded off-exchange and brought to NYMEX for clearing. NYMEX's products span a wide range of energy and related commodities, including electricity, natural gas, crude oil, refine products, coal, uranium and emissions allowances. These products are global, though the majority are US-based.

NYMEX lists a number of forward contracts in electricity, including a few based on hubs within PJM. These contracts, however, have not tended to attract significant liquidity – certainly nothing on the scale of some of NYMEX's more successful natural gas contracts. Nevertheless, NYMEX's wide range of products is likely to provide good opportunities for portfolio margin benefits with PJM's markets.

Currently NYMEX is in the process of being acquired by the CME. It is unclear how much this merger process will distract NYMEX from opportunities such as clearing PJM's markets. It is also unclear how important energy clearing customers – such as PJM and its participants – will be to the combined entity, in which the current NYMEX will represent less than 20% of its business, and therefore how responsive it will be to their needs.

8.3.2 LCH.Clearnet

LCH.Clearnet is one of the world's largest clearing companies. In the UK, LCH.Clearnet Ltd. clears ICE Futures Europe (formerly the International Petroleum Exchange, or IPE), as well as providing a clearing services for OTC UK power. Its continental affiliate, LCH.Clearnet SA, clears the Powernext electricity market for France. In the US, LCH.Clearnet Ltd. currently clears US products for the Intercontinental Exchange (ICE).

LCH.Clearnet is presently the second largest clearing house for energy in the world, behind NYMEX. However, LCH has lost the clearing business of a couple of small energy exchanges in recent years – Endex and APX Power UK (formerly UKPX) – and this year is due to lose the business of its largest energy customer, ICE, which is seeking to move this trade to its own clearing house, ICE Clear Europe. While LCH.Clearnet will no doubt retain institutional knowledge concerning the clearing of energy products, this will significantly reduce its abilities to offer portfolio margining benefits.

Recently LCH.Clearnet announced an intent to clear electricity products for Nodal Exchange, based on forward trading around nodes in PJM and other US electricity markets[77]. Whether this will provide an opportunity for meaningful risk offsets against PJM's markets, however, is impossible to determine, as Nodal Exchange is not presently operational, and how successful it will be is unknown.

8.3.3 ICE Clear

ICE Clear US is the former New York Clearing Corporation (NYCC), the clearing house for the New York Board of Trade (NYBOT), which ICE acquired in 2007. Presently ICE Clear US clears the contracts of the former NYBOT (now ICE Futures US), which are primarily agricultural products such as coffee, sugar, cocoa and cotton. ICE also has a European clearing subsidiary, ICE Clear Europe, which is not registered as a DCO, but is registered by the UK Financial Services Authority (FSA). ICE has announced plans to move all of its energy contracts currently cleared by LCH to ICE Clear Europe by July 2008, including its US OTC energy contracts.

ICE Clear US does not presently clear any energy products[78], and has no published plans to do so. As such, its capability to clear PJM's markets is unclear, and it would not be able to provide any portfolio margining benefit.

[77] Joint press release from LCH.Clearnet and Nodal Exchange dated 28 November 2007. Refer to: http://www.lchclearnet.com/press_and_publications/press_releases/2007-11-28.asp.
[78] Unless its ethanol contract is counted as an energy product.

ICE Clear Europe is presently "in discussions with the U.S. Commodity and Futures Trading Commission in order to ensure that it is able to provide clearing services to ICE's OTC products." Not being a DCO, it is presently unclear whether ICE Clear Europe could clear PJM's markets. If it can do so, its significant standing in the energy markets could allow meaningful portfolio margining benefits – presuming it is successful in bringing this cleared business over from LCH.