HARVARD LAW SCHOOL | ENVIRONMENTAL & ENERGY LAW PROGRAM

Transcript of CleanLaw Episode 18: Bill Hogan Talks with Ari Peskoe about Recent PJM Pricing Proposal, April 22, 2019

To return to our website [click here](#).

Robin Just:	Welcome to this podcast from the environment and energy law program at Harvard Law School. Today, Ari Peskoe interviews Bill Hogan, Raymond Plank Professor of Global Energy Policy and Research Director of the Harvard Electricity Policy Group at the Harvard Kennedy School. They discussed electricity market design and PJM's proposal to reform price formation in its reserve markets. We hope you enjoy the podcast.
Ari Peskoe:	This is Ari Peskoe, Director of the Electricity Law Initiative, and I'm pleased to be joined by Bill Hogan, who is going to help us understand PJM's recent 848-page filing at FERC, detailing its reserve market pricing reform proposal. Bill, thank you for doing this.
Bill Hogan:	Thank you, Ari.
Ari:	So, before we get into PJM's filing, which includes a paper by Bill and his colleague, Susan Pope, I want to discuss some of the principles that inform electricity market design. The goal here is to unite physics with economics. We have power flows that behave according to physical laws, and we have generators and consumers that we assume act rationally in the face of economic incentives. I've heard you say, Bill, numerous times that market design has to get the prices right. So, let's start with that. What does it mean to get the prices right?
Bill:	Well, the basic framework is an economic efficiency framework maximizing the benefits minus the cost. Within that framework, the simple-minded case, which is the single product is the place where you get the efficient outcome. It's where the supply and demand intersect, and you get a price, which is the marginal cost of supply equals the marginal benefit of demand. That basic conceptual idea extends to multiple products, multiple locations, multiple time period, all the other complications that arise in the electricity system. So, on the first instance, I would say getting the prices right is trying to get prices to reflect the marginal cost of economic dispatch.
Ari:	So, we start with this concept of setting the price of the marginal cost of the marginal generator. How do we then incorporate? How does market design then incorporate the physical realities of the power system? The fact that we have

hundreds of generators interconnected across a transmission network that has its own physical parameters, how is that then reflected in the market design?

Bill: Well, we're very fortunate that over the years, before we tried to introduce open access in competitive markets, the people running the utility system had developed this general framework of economic dispatch. The basic idea is that you have the cost of different kinds of generators, they're located at different places, they're connected by a transmission grid, and you have tools, which can describe what are those interactions, and these were readily available.

Bill: Then you can choose for any level of demand. You can choose the dispatch of all of the power plants that incorporates the effects of transmission constraints and ramping constraints and all the other factors that have an influence on the ability of the system to provide power to meet the demand.

Bill: In the process of doing that efficiently by minimizing the cost associated with doing that, one of the byproducts that comes out of that is this estimate of what we refer to as the marginal cost over the locational price, which reflects at a relocation the marginal cost of meeting an increment of demand of that location if we redispatch the entire system. This is conceptually not a trivial calculation, but conceptually, it's not a big leap from just a simple supply-demand story I mentioned before.

Ari: So, when you're able to mathematically represent this physical system, how do we then connect that to getting consumers and producers to do what we want them to do? How do we connect that physical system to the prices that people would pay?

Bill: Well, there are many levels of detail packed into that question, but the first thing, and the most important thing is to look at the short-term. So, we have short-term and scheduling, and then longer term investment decisions, but the short-term is literally realtime, right now, here with the lights staying on.

Bill: We want the prices to reflect the real economic opportunities. Under some reasonable simplifying assumptions, if you solve for the economic dispatch and get these locational prices, and then you use those prices in the settlements, so generators get paid these prices and loads gets charged these prices, you will have the situation where everyone is better off, individually, given those prices, everyone is better off following the economic dispatch. So, they don't have incentives to deviate and do other things that are possibly beneficial to them but harmful to the aggregate system.

Bill: That's the characteristic of prices, which are sometimes referred to as supporting the solution. So, they actually are consistent with the economic dispatch solution. So, that's a realtime, and that's the first thing.

Ari: So, that's the realtime market. That's the critical point that the price is in the right incentives to all the actors, to do what they're supposed to do. In these markets and at least in this country, we have a realtime market, and we also have the day ahead market. So, how does that come in to play?

Bill: So, the day ahead market is essentially a scheduling and hedging market. I mean, no power is delivered day ahead. It's all delivered at realtime. The basic principle it's followed is to run a day ahead market that's consistent with the realtime. So, the basic structure is very similar to the economic dispatch. We have generators make offers for how much it's going to cost for them to produce, and loads make estimates of their total demand.

Bill: We solve that day ahead schedule. That gives energy schedules and prices. There's a settlement the day ahead. Then you could think of those as, for the most part, financial contracts, which are then going to be settled relative to the realtime. So, we get to the realtime. If you generate more or less than you were scheduled in the day ahead, you buy yourself, sell or buy, actually the reverse, from the generator at the realtime prices.

Bill: So, in the unusual case where the realtime is exactly the same quantities as the day ahead schedule for someone, then they'll pay anything or get paid anything different than the realtime. If they deviate, then they get charged at the realtime prices.

Ari: Do we have good data to suggest that what people say they're going to do day ahead actually reflects what they really do in the realtime market the next day?

Bill: Well, we hope not because the realtime market is, of course, uncertain day ahead. So, we want them to adapt to the actual conditions that exist in realtime as opposed to stick to where they are with their day ahead schedule. So, what you really would like to have is people adjusting depending on the weather and all the other things that happen between day ahead and realtime.

Bill: A more relevant test is how do the prices turn out to be on average. Do the prices tend to converge, so that on average the day ahead prices are approximately equal to the expected value of the realtime prices. The evidence there is that it's not perfect because there's risk aversion and other things, but largely speaking, the answer is yes, they do converge, and the better the market design, the more likelihood to converge. That's a sign that it's healthy, that markets are consistent, healthy, and working well.

Ari: You characterized the day ahead. You said you could think of it as a financial market. Is it also fair to say that the one purpose of the day ahead market is just to allow generators and other market participants to prepare just as a matter of these are business practice and making sure that they're ready to perform in the realtime market and can make those commitments?

Bill: Yeah. So, there actually are some decisions that actually have to be made the day ahead or earlier. For some plants unit, commitment decisions have to be made to make sure that they're going to be online. Fuel procurement may be part of the story as is being discussed in New England.

Bill: So, yes, certainly gives them signals and information, and they should be taking actions to prepare, and in some case, making decisions which are then they're going to have to live with because it's too late in the realtime to change those decisions.

Ari: So, I think this is all background to help us get to what's really an issue in the PJM filing, which is the reserve market. So, tell us a little about what is that market, what purpose is it serving.

Bill: Well, the discussion we've been having so far at a high level of generality is really focused on the energy production of electricity, and the energy market, and energy prices, and all that kind of thing. Because of the uncertainty that I talked about, there are in addition to the actual production of power from generators, over short horizons, you need to keep a certain amount of capacity on reserve, not producing energy, but that could produce energy relatively quickly in a few minutes later or an hour later or something depending on what happens.

Bill: So, we're doing an economic dispatch. It's realtime. It's 2:00 in the afternoon, and we're scheduling for the next hour, let's say. Then at 2:10, something happens. If you have a demand goes up, the sun goes over and solar stops generating, a transmission line falls down, all kinds of things can happen. Then you have to respond very quickly in order to be able to keep supply and demand in constant balance.

Bill: So, all electrical systems, in addition to the generators that are producing energy, have generators and other demand response reserves that can respond with varying lag times, but often some of them are synchronized and instantaneously available that have to be maintained in order to meet these forecasts, essentially, that are going to happen over a short period of time.

Ari: So, it's really where we as consumers are willing to pay for this reserve product because we value the reliability of the system that it's always going to work, and its reserve market is there to help make sure that that's a reality.

Bill: Right. It's necessary in order to keep the lights on. If you didn't have reserves, then as soon as something happens, a generator tripped off and suddenly wasn't available unexpected and you didn't have anything to quickly respond to it, well, that would then create a cascading failure and propagate through the system. So, we can't allow that to happen.

Ari: So, what we're talking about in this proposal is specifically how PJM, which is the market expands across 13 states in mid-Atlantic part of the country, is how they procure these reserve products. In their recent filing, they are asking FERC to find that the current mechanisms they use for buying these reserve products is unjust and unreasonable. That's the standard in federal law that FERC has to follow.

Ari: So, as I understand it and please correct me here when I get this wrong, right now, there's at least two different products right now in this reserve market. There's what PJM calls tier one and tier two reserves. The tier one products are from generators who are actually already part of the energy market but have room to spare. They can generate more energy if PJM asks them to and they can do it I think within 10 minutes. That's the tier one product.

Ari: Then there's this tier two product, which I think are generators that are just waiting in reserve but aren't actually part of the energy market right now. Do I have that roughly correct so far?

Bill: I think if I were to describe tier one, it's pretty much the way you described it. Tier two, the principal difference is that these generators actually would find it economic to participate in the generation, to generate electricity, but we need to hold them because of locational requirements or other special conditions. We need to hold them off because we need them for reserves.

Bill: So, they're losing money. They have an opportunity cost because we're not allowing them to produce electricity and generate, and we're not paying them to produce electricity and generate. So, there's an opportunity cost.

Bill: So, in theory, tier one resources don't face the same opportunity cost, although that's arguable on their part, but by definition, tier two resources do face this opportunity cost. These are short-term synchronized reserves. There's also some longer term reserves that are import.

Ari: So, just on this point about what tier two reserve is. They are getting paid for this reserve product there, right? They would just be paid more if they were allowed to sell into the energy market.

Bill: Well, the tier two reserves are paid based on an opportunity cost story. It's an estimate of how much they're losing and then trying to compensate them, but the tier one resources are not paid. The tier two resources have an obligation to respond. The tier one resources don't have an obligation to response. The performance of the tier two reserves is much better than the performance of the tier one reserves.

Bill: So, for all of those reasons, we have a breakdown in the theory of market efficiency. In principle, just as with energy, people who are providing reserves would receive the same marginal price. They would receive the same marginal price. It was a system-wide requirement. There would be a single price for reserves and everybody would get paid, not some get paid and some don't get paid. Everyone would have the same obligations for the reserves that they're providing, not some have one set of obligations, and the performance metrics would be more closely aligned and more connected to the physics rather than the poor incentives.

Bill: So, as I understand the numbers, essentially only about half of the synchronized reserves are actually getting compensated in this process. Suppose we had energy suppliers and we said, "Well, you're a renewable energy supplier. You're a wind resource. Your marginal costs are zero. So, we're not going to pay you for your energy. We're just going to take it." That would be the way we treat tier one resources. You can see by analogy this would be problematic.

Ari: So, there's multiple inefficiencies right now is what you're saying. Resources are not being treated similarly. In FERC speak, this is unjust and unreasonable. From your description, you might say this is also unduly discriminatory in a sense.

Bill: It's discriminatory, it's inefficient. It produces low-performance results. I mean, it produces all the bad things that you would think if you had an inefficient market.

Ari: From the consumer perspective, lights are staying on in PJM despite these inefficiencies. Do we think right now that consumers are paying too much for this reserve service or too little? Is there a way to look at it from that perspective?

Bill: There's a couple of things that are packed into that question. So, there's a distinction between ... I would make a very important distinction between what consumers are paying and the cost of society, and those are not the same thing.

Bill: So, you could have a situation where the costs are higher because of the way we're running the system, but the prices are lower. Both of those things can happen at the same time. So, imagine we said, "Well, what we'll do is we'll subsidize some very expensive resources, and then we won't charge customers for the prices for those resources at the margin, and we will then take the increased cost of those resources and socialize them and spread them across everybody so it's hidden," and then you would end up and you say, "Gee, prices are low, and consumers are paying less."

Bill: Well, the total costs of the system are actually higher than you need to do because you're acting essentially like, in this case, a monopsony or a monopoly behavior on a generator side. That's inconsistent with efficient markets. What we want is efficient markets not minimum payments by consumers. We want minimum cost for society. The prices and the incentives that people see at the margin could be all different.

Bill: So, the problem right now is that the costs are higher than they need to be because the system operators then have to compensate for the fact that the tier one resources are not performing. So, they then order other generators out of merit and they're essentially equivalent to hiring more tier two resources than they need in order to make sure they meet the reliability standards. Some of that is done through out of market payments, which the federal regulators have long recognized as problematic. You have to do it sometimes, but it's a symptom of a fundamental problem in the market.

Bill: Those out of market payments are charged to customers through various kinds of uplifts and socialized. So, it's not so much that the lights go out more often, it's just that the total cost of society of keeping the lights on is increased.

Ari: Right, because the bottom line for PJM is that it's going to do what it needs to do to keep the lights on, and it can, I guess as what you're saying, is it can either try to do that in an economically efficient way or it can do what it's doing now and have to take various actions that aren't quite reflected in prices and ultimately costs more to society.

Ari: So, we have the tier one. The tier two, we've mentioned those. You mentioned briefly that there's this other thing, the non-synchronized reserves, and by synchronized you mean, in a sense, they're slower to be able to respond to PJM's dispatch signals. Are those non-synchronized reserve, are they part of the problem as well or is that just a separate issue that we're not touching on?

Bill: Well, the tier one, tier two story is the primary focus, and then the treatment of other kinds of reserves and pricing of that is it's a similar set of problems, but it's quite as acute as for the synchronized reserves.

Ari: Okay. So, that's the case. I mean, you've talked about the case then about why the current system is unjust and unreasonable. That's the first thing that FERC has to conclude. Do you think there's going to be any pushback on that point or do you feel like there's broad recognition in the market that whatever they're doing now is just, as you've described, it's inefficient, it's not working, this should be a slam dunk case in your mind or do you think there'll be any pushback on this?

Bill: Those are not mutually exclusively options. So, it's long been recognized, I mean, since certainly more than 10 years ago, maybe more than 15 years ago that there was a problem in the way we were pricing reserves and pricing scarcity in the realtime market, in particular, and then by connection to the day ahead market and internal inconsistencies.

Bill: This is not news, but it's been deferred and deferred that the problems have been deferred while we were dealing with other problems. The policy agenda was crowded, and they couldn't do everything at the same time. So, a lot of attention, for example, has been devoted to improving capacity markets, other longterm-

Ari: It's a separate podcast. We'll get to that some other time.

Bill: Right. So, this is not news that there's a problem. The problem is becoming more acute as time goes on, and it becomes especially more a matter of concern as we get more and more intermittent resources coming on to the system because that creates the short run volatility that I talked about, where you're looking ahead and five minutes or 10 minutes or an hour, and then a lot more things can happen, and you have to respond to those. So, the reserves situation becomes more and more important.

Bill: It would have been my preference to have addressed, to make the PJM filing 10 years ago, but it's now. They've done it now, and it's certainly not too early.

Ari: So, before we get in to what the proposed solution is, there's one more thing I want to try to connect to this, which is something that was on FERC's agenda, and FERC did issue an order about it, which is called Order 825, which had to do with scarcity pricing or shortage pricing. My understanding of that was that it required PJM and the other market operators to make sure that the dispatch interval was matching the settlement interval.

Ari: So, I think PJM, for example, I think the other markets do this as well, is they may dispatch resources every five minutes because as you've mentioned, things change very quickly and they have to be able to react quickly to what's happening on the system, and they also wanted to make sure the prices reflected the value that these resources were providing. So, is there any connection between that shortage pricing issue and the reserve pricing issue we're talking about here or am I barking up the wrong tree here?

Bill: Well, they're not the same issue, but, obviously, if you're changing the dispatch every five minutes, then you're changing the prices every five minutes, certainly for energy. If you have a settlement system, which is based on an hourly average, then this is going to create a lot of funny incentives, particularly around the end of the hour or something like that, where people might have a chance, the prices, they think they're going to be high or low on average and they start having an incentive to deviate even though the actual price shouldn't be the same as the average price at this time.

Bill: So, that creates lots of problems. They're not insurmountable, but the obvious solution is to make the settlement system consistent with the pricing system.

Ari: That's what FERC did in 825.

Bill: Right. That logic applies to energy, and it applies to reserves. So, they're connected in that regard. So, underpinning all of this and it's part of what PJM has done and is planning to continue is what's called co-optimization. So, when they're dispatching energy and dispatching reserves, they do those simultaneously deciding on who's generating on whose reserves, and they're trying to get the overall benefit maximizing solution for each one of those things. Then they get a set of prices that are consistent for reserves and for energy.

Ari: That co-optimization, that's part of the filing and discussing this now.

Bill: Well, it's already been part of PJM's practice and its existing procedures. So, what they've said in their filing is, "We're going to continue to do the co-optimization," but they're just changing the description of the reserve pricing part of the story.

Ari: So, what's the key features then of the proposal? How are they changing the reserve market? They're getting rid of tier one and tier two, and what's going to replace it?

Bill: They'll have synchronized reserves, non-synchronized reserves. There are different categories that are staged. The different categories in realtime I'm talking about are based on the lead time for response. So, some generators are going to have to be ready right away. Then some others can be ready in somewhere up to 10 minutes and then others in 30 minutes. Then they're looking at the uncertainties over those periods of time, and they're also, which they've done things like this all along. I mean, the notion that there are different time periods and different types of response is not new. It's just trying to be a little bit more formal about this, and to connect it to the uncertainties.

Bill: Then the pricing mechanism is what's referred to as a cascade model. So, synchronized reserves can meet the synchronized reserve requirements. They can also meet the 10-minute reserve requirements. They can also meet the 30-minute reserve requirements as long as they haven't been used, I mean, as long as they're still there. So, there's a cascade of facts that the price of the synchronized reserves, which are the most valuable, will never be lower than the price of some of the other two and so forth. So, you don't get this problem they had in California a few years back, where they got price reversals because they didn't recognize this cascading effect. So, that's all part of the design.

Bill: Then there are two critical issues here that are addressed in the filing. One is what's referred to in PJM as the penalty factor. So, for each one of these categories, there's a minimum reserve requirement. Then the question is if you're now going to go below that minimum reserve requirement, what's the penalty factor? Because they're going to take action in order to get you back up, so that you have the minimum reserve requirement, and then the question is how much you're paying for those actions.

Bill: The penalty factor, they've selected based on the kinds of actions they're taking, other pricing rules that come from their tariff is the new factor they're proposing is $2,000 per megawatt hour.

Ari: That's a penalty that a generator would pay if it's not performing?

Bill: No. This is not for generators not performing. It's called a penalty factor for historical reasons, but you can think of it as the value of the operating reserves when you get into how much would you be willing to pay to get an increment of operating reserves? You'd be willing to pay whatever you're doing to avoid the penalty, in some sense, the system as a whole. The penalty is loosely is the emergency actions that they have to take or other kinds of things.

Bill: So, think of it as how much you'd be willing to pay for an increment of reserves when you get down to a very low level. So, that's one thing. There's an explanation of both the history and how it compares with other emergency actions and other prices that can set energy prices and all of those details.

Bill: Then the second question, which is problematic is that the existing penalty factor structure has a minimum level of operating reserves and a low penalty factor. So, they're raising it from $850 to $2,000 for the penalty factor. They also have a cliff at the minimum reserves. So, if you're one megawatt below the minimum reserves under the current system, the charge is $850. That's allocated to the price of reserves. That's what you pay for reserves. If you're one megawatt above the $850, then the price falls dramatically because you hit this cliff.

Bill: This is another fundamental problem and thinking about that because if you think about the question, "How much would I be willing to pay for an increment of reserves if I go above the minimum reserves?" Well, the answer isn't going to be zero. There must be some value to those reserves. It's relevant for a variety of reasons why that you should be trying to represent that.

Bill: The way I think about it and the way people have thought about this for a very long time and this is not a new idea is that, "Well, over the next period that we're talking about, 10 minutes for 10-minute reserves, 30 minutes, an hour for the 30-minute reserves, and these different uncertainties." There's some chance that things are going to happen. If things happen, we'll get into the penalty range, again, during that period of time.

Bill: Whereas if I could buy an increment of reserves now, then I would reduce the chance that I would actually get into that penalty situation over the horizon that we're planning for. So, the marginal value, how much would I be willing to pay now for reserves, should be the penalty factor times the probability that we're going to have to pay the penalty factor. That's where you get those loss of load probability story above the minimum reserve level, and multiplied, times the penalty factor.

Bill: What this does is it produces this funny shape, but it's because of the way the systems operate. If you get below the minimum reserve levels, they will act now. If you're above the minimum reserve levels, you may have to act later, but in answer to the question that you asked the same question in both situations, which is, "How much would I be willing to pay now for an increment of operating reserves?" When you're below the level, you'd be willing to pay the penalty factor. When you're above the minimum level, it should be based on the probability that you'll get in this situation later, and that's when you get this shape.

Bill: That produces something that looks like a demand curve that falls off gradually, and it has a number of advantages because one of the just practical problems is that if you don't have something like this and you're requiring them to decide you're on a shortage condition or you're not. You get a lot of this question that gets imposed on the system operators about, "Are we in shortage conditions or not?" Then there's a lot of pressure for them to say no. So, that tends to depress the price even further.

Bill: With this version of the story in the loss of load probability, it's there all the time. Sometimes it's very small, but that doesn't mean it goes away. So, you don't have to worry about this declaring of shortage question. It fits very naturally with the co-optimization with energy.

Ari: So, this curve that you're talking about, this is called the ORDC.

Bill: Operating Reserve Demand Curve.

Ari: So, what this is telling PJM is how much of this product to procure. Is that a fair summary?

Bill: Not quite, but it's close. So, think about it as co-optimized. So, let's suppose I have fixed amount of generating capacity, but I have flexible energy demand, which we don't, but-

Ari: It's hypothetical.

Bill: Hypothetical. Then you would say, "I can lower the price and then demand will go up, and then I'll have to use more of the capacity of produced energy and less for operative reserve." If the operating reserves are worth more, I actually want more operating reserves and less energy. So, I want to get this balance where the marginal value of operating reserves and energy were just indifferent in terms of using the generators in those two cases.

Bill: So, the co-optimization seeks that balance. It provides that balance between energy and operating reserves, and calculates prices, so that the price of energy is the price of the variable cost of energy. So, it's costing $40 per megawatt hour to produce the energy, plus the scarcity price, which is $20 per megawatt hour for operating reserves at the margin. So, the total price of energy will be $50. The price of operating reserves will be $20 in that situation, and a generator who was sitting there, the marginal generator who's looking at this is indifferent between producing energy and producing operating reserves, which is the condition we want to get.

Ari: So, that's the critical feature of this.

Bill: Right. That's a critical feature, co-optimization and then the prices are consistent between generation and operating reserves. If you think about that, and this is what would made it appealing in the first instance when we first started talking about this is that it also solves the problem of setting scarcity prices for energy because now, as operating reserves are reducing, their value is going up with the margin. This is being added to the variable cost of energy, so that the total price of energy is going up to reflect the scarcity.

Bill: So, you're getting the right signals for load. You're getting the right signals for other generators, for generators at the margin, and solving this problem of how do you get scarcity pricing in the realtime.

Ari: So, how do you calculate this ORDC? What goes into figuring it out? It seems like that can make a big difference on how the system works.

Bill: Well, there are basically three components. One is the minimum reserve requirement. So, there's a level, 1,500 megawatts of synchronized reserves, for example, just to make up a number, and that comes from NERC standards and contingency requirements and all of that complicated electrical engineering story that has been evolved over decades. I'm picking that as given.

Ari: Right. That may or may not be based on any sound economic principles, but it's there, and so you take it as it is.

Bill: For the purpose of what we're doing now, we're taking that as given. They've been doing it for decades, and they will continue to do it in the immediate future.

Bill: So, that's the threshold level, the minimum reserve. Then the penalty factor that applies when reserves fall below that minimum or would apply when that falls below, and that's the $2,000. So, there's an explanation of the history of that and how it compares with other emergency actions and so on. This is partly a policy decision. I mean, this is not a precise science thing here because there's actually 50 different things they do as emergency actions, and you're trying to get the average, and it gets something, it's representative, and keep it relatively simple. So, those are two out of the three.

Bill: Then the third one is this if I have this many reserves now, what's the probability in the next 10 minutes or the next hour? We're going to run out of reserves because of unforeseen events, and then have to invoke emergency actions and then pay the penalty factor. So, that's the loss of load probability that we're going to calculate.

Ari: So, is that just based on historical record in PJM of how the system actually works?

Bill: Yes. The proposal is ... I don't remember off the top of my head all of the details, but they're looking back over three years, and then as other people have done, they divided it up. So, weekends are different than weekdays and different hours during the day because of different peak loads and all these kinds of things. So, it's broken up into various intervals, historically. Then we look at the change in the load, the change in the solar, the change in the wind, the change in conventional-

Ari: So, this has to evolve. As the system evolves, it's going to evolve with it.

Bill: It's going to evolve with it, right. It's an empirical story. They have records on all of this information. The PJM filing goes through that story about what they're using and I think it's perfectly reasonable. It's been done by others. Texas, for example, has a similar approach to this.

Bill: So, the alternative approach would be a forward-looking modeling story, simulation of what we think is going to happen based on all kinds of analysis. This turns out to be quite hard. I don't think it's any reason to believe it's better than just taking the empirical and just going over the different time periods, have different expected effects, and different probabilities, and adjusting that empirically going forward. I that's a much more practical and transparent to what's going on, and that's what they're doing.

Ari: So, one problem this should solve is the generator incentives, that the generator should be indifferent as to whether they're in the energy or the reserve market. Are there any other improvements to the system that they should make? Any other key problems this might solve?

Bill: There is one that is underappreciated, but I think it's actually quite important. So, the operating reserve demand curve idea was motivated in part by the problem of inadequate scarcity pricing. If we had load bidding and you could get enough load to bid for five-minute intervals and they would adjust and so forth, a lot of the problems would be solved that way, but we don't have that very much. So, it makes the dealing with the operating reserves story more important.

Bill: One of the things that it does it is solves the conundrum that we had before, which is high prices during tight situations could be reflect, well, tight situations or high prices and tight situations could reflect market power and economic withholding.

Bill: If you leave it to generators to put in high offers as the means for getting high prices when you have scarcity, what you're relying on is essentially economic withholding and market power. It's always been a problem, which is how do we tell the difference between prices are high because they should be because we're in a scarce situation versus prices are high because people are withholding.

Bill: The operating reserve demand curve provides an answer to that question. So, as prices can go high because of scarcity of operating reserves, but it's completely consistent with generator offers being low. So, if you think generators have market power and you say, "We have an offer cap, and the offer cap for you is $1,000," or something, whatever the number is that you choose, well, then you would say, "If we didn't have operating reserve demand curve, then the prices can't go above $1,000 because you have an offer cap of $1,000."

Bill: With the operating reserve demand curve we say, "No. The offer cap is fine." Now, the price of reserves goes to $2,000 or $3,000, and you get this cascading effect, and that's added to the energy price. So, the price that generators get paid is the $1,000 variable cost plus the $3,000 scarcity price. So, they get paid $4,000, but they're not withholding. So, it's not market power. It's scarcity pricing.

Bill: That's a very good thing. We like high prices. Economists like high prices when the system is very tight because it gives everybody the right signals. We don't like high prices when it's caused by generators withholding supply in order to jack up the price. So, this is a cleaned and principled way to recognize the difference. I think that will be a big help going forward.

Ari: There's different types of generators on the system, and one way you might think about this is whether generators are flexible and able to respond very quickly to PJM system dispatch instructions versus an inflexible generator, the extreme can basically only go on or off and doesn't really have an ability to respond. Does this change in how the reserve market is going to be structured? Does this affect the two extremes differently, the flexible versus the inflexible resources?

Bill: The flexibility questions are addressed in part by the changes in the operating reserve pricing model here, but not completely. So, for example, we have different types of reserves that can respond at different times. If you can respond instantly because you're synchronized, then you're going to get paid more because they're higher valued reserve than if you don't. So, that helps in this regard.

Bill: Second is the amount of capacity that you will get credit for in the operating reserve calculation is a function of the speed with which you can ramp. So, generators can increase their output, but they can't typically increase their output instantaneously, but it comes in gradually. So, over the 10 minutes, you can get a certain amount, and then that's the amount that we say in your capacity.

Bill: If you increase that ramp rate by making investments in the generator, then, obviously, you get credit for higher levels of operating reserves, and then you can sell more and so on. So, both of those things help provide. The more flexible resources can both sell more operating reserves and sell higher quality operating reserves and get higher prices associated with it.

Bill: Then there's a second issue, which is not addressed in this filing, which is the dynamic dispatch question, which is prices going over time. So, if we thought about it as multiple dispatch periods, and we say the prices are going to be starting and they're going to go up, and go up, and go up, and go up over this timeframe, and we may want to hold off some of the regenerations so that's available later when the prices are higher. That's a separate problem that deals with the flexibility question as well, and that's not addressed by this. Although what's going on with the operating reserve demand curve will not be inconsistent with that. It's just that that's a separate issue.

Ari: Then there's also intermittent resources. Typically, wind and solar, obviously, is what we're talking about here, and those resources have typically tried to put as much energy into the market as they possibly can whenever they have it available, and that's motivated by energy prices and also the fact that they have certain policy incentives to produce as much as possible. If you could imagine a world in the future where we have more of these resources and where actually, these resources, of course, they can always curtail. They can always produce less, but if they had the right incentives, they might choose actually not to produce as much as possible. They might intentionally curtail so they had flexibility if that was something the market valued. Does this proposal at all connect to that issue for incentives or intermittent resources to be able to be flexible if that's something the system really needs?

Bill: There's no reason in principle why renewable generating resources can't be treated like any other kind of generating resource as long as it meets the same requirements. So, if the sun is shining, and they're not producing electricity, but they are providing reserves, we could pay them for their reserves that they could flip a switch and suddenly start producing electricity if we needed them in the next five minutes.

Bill: There would be an issue to be worried about, which is if you wait five minutes and then the sun isn't shining, they can't provide what you thought they could provide. So, there's an issue there. What the probability is that they'll be able to provide. That's not a new issue. That comes up under the same discussions and capacity markets on how much capacity you're going to get credit for.

Bill: It will be a much less serious problem in the very short run because our forecasting is pretty good, the very short one about wind and solar. The problem is not forecasting for the next half hour or hour. The problem is forecasting for tomorrow.

Ari: Last thing is I know you don't spend a lot of time on capacity markets, but does this proposal interact at all with capacity markets? One thing that we've seen in PJM in particular is that capacity market revenues are an increasing percentage of the overall market. Some people like that, some people don't like that, but does this proposal counteract that trend at all?

Bill: Well, certainly, it will interact with the capacity markets because the capacity markets are net of payments or net, eventually, of the expected energy revenues that are going to be earned, and this is going to affect energy revenues. So, in principle, the directions would be that more of the total cost will be embedded in their realtime energy and reserve market, and the last will be in the capacity market going forward, but it's certainly not a one-to-one story. It has to do with forecasts, and when you forecast, and how fast these things adjust, but they definitely will interact with each other.

Ari: So, comments on this are going to be due I believe in mid May. Is there anything else you think people should know about this that we haven't already talked about?

Bill: I think just in summary, this is a reformed proposal from PJM, but the idea of trying to compensate for the value of reserves by looking at how much it costs when you don't have them and the probability that you're going to get into that situation is an old idea that's been used in many, many different contexts. The only thing that's new about it in here or in the Texas situation or in MISO and other places where they have different versions of this is the focus on the realtime for doing that as opposed to some long period, long forecast ahead. That is a critical innovation, and a very good idea, and it deals with the real incentive problems that we actually have as opposed to our ...

Bill: I mean, the problem associated with the capacity markets, which is to provide incentives for resource adequacy is that you're looking very far ahead. It's impossible to design the capacity markets that produces exactly the right incentives every five minutes, but it's actually important to have the right incentives every five minutes. What we're trying to do here is get the incentives much closer to the real operating conditions, which might be on days when load in the aggregate is low, but it's all relative, right? So, if load is low and generation low, your in a scarce situation. So, that's what you want to deal with.

Ari: So, PJM's full filing is in FERC docket EL19-58. You can find Bill's paper that he wrote with his colleague, Susan Pope, in that docket, and it's also available on whogan.com, which is a wealth of information on all things electricity markets. You could spend the next many months of your life just reading all the papers on that.

Ari: Bill, thank you so much for explaining this to us.

Bill: Thank you.

To return to our website click here.