Fiftieth Annual Allerton Conference
Allerton House, UIUC, Illinois, USA
October 1 - 5, 2012

Prediction Markets for Electricity Demand

Luciano I. de Castro   Peter Cramton

Abstract— **Load forecasting is a central task for operating, maintaining, and planning power systems. Because of this importance, many different methods are proposed to forecast load, but none of them is proved clearly superior. This paper proposes a prediction market to forecast electricity demand that has the advantage of allowing aggregation and competition among the many available methods. We describe how to implement a simple prediction market for continuous variables, using only contracts based on binary variables. We also discuss possible pitfalls in the implementation of such a market.**

Keywords: Forecast of electricity demand, power generation planning, implementation of prediction markets, forecast of future demand, mechanism design of capacity markets, prediction market for continuous variables, winner-takes-all contract, index contract.

I. Introduction

Load forecasting is a central process in the operation and planning of electric utilities. Indeed, an essential step of resource planning in electricity markets is assuring that there will be sufficient resources to meet future demand. While building capacity is costly and takes time, the economic consequences of an electricity shortage may be severe. If a system operator can accurately predict what the future demand will be, it can indicate the appropriate level of investments needed, without incurring in the risk of shortages. From this, it is clear that future electricity demand is of high interest, which explains why hundreds of papers have been written to propose different methods for load forecasting.[1]

These methods were compared by some authors (see references in section II), but their conclusions are not sufficiently compelling for pointing out a clearly superior one. Also, the very fact that many papers are still published on this topic shows that the problem is far from being settled. In fact, it is more likely that it is not possible to choose the best method since technological advances will continue to improve current approaches and specific conditions in different markets may favor one method over another. Given the value of load forecasting and the many different methods, what should a regulator do to obtain the best possible predictions?

This paper argues that a simple mechanism, based on prediction markets, is likely to produce better results on load forecasting. Prediction markets are platforms for trading contracts associated with well-specified events, whose main purpose is to reveal information about the likelihood (probability) of those events. Our proposal is to create a market where contracts associated with the demand of electricity in a given region and period of time are negotiated. For example, a contract could be specified to pay \$1 if the summer peak load in New England in 2013 is below 30,000 MW and zero otherwise. We argue in section III that if the price of such a contract is \$0.72, we can interpret this as indicating that there is a 72% chance that this event will happen. From a market with this kind of contracts we can obtain not only a load forecast, but also a good sense of how precise the forecast is.

To show that this proposal is sound and important, we begin in section II by discussing some specific characteristics of electricity demand that put in perspective some of the difficulties associated with load forecasting. In particular, we point out potential incentive problems with the current used methods. As we shall see, the analysis of the difficulties will lead us to our prediction markets solution. Although prediction markets have been extensively used in recent years to predict many uncertain events, we are not aware of any application to energy markets. Nevertheless, section III provides extensive evidence that prediction markets do work in many contexts.

In electricity markets, however, the implementation of prediction markets is not as simple as it may appear. Most prediction markets deal with binary events such as whether or not the Republican candidate wins the next American election, but electricity demand is a continuous random variable, and this may cause complications. Section IV-A reviews the literature and describes some market designs for the case of continuous variables. The paper then introduces one of its key contributions in section IV-B, which is a new and simple market design for continuous variables. We show that binary contracts are not only capable of dealing with continuous random variables, but also present a number of advantages in comparison to other designs. Among those advantages is one that is particularly valuable for electricity markets and has apparently not been emphasize before. Our prediction market allows us to obtain not only the expected value of the future demand, but also the variance and higher moments of this random variable. Since shortages are so costly, it is important to know more about the distribution of future demand than just its expected value.

Section V continues the discussion on implementation, by considering other practical issues, such as the definition of the contracts and the possibility of manipulation. While

We are grateful to Joisa Dutra, Carmen Liron-Espana, Marco Ottaviani, Mario Pereira, Charles Plott, Thomas Rietz, Thiago Veloso and Eric Zitzewitz for helpful comments.

Kellogg School of Management, Northwestern University. E-mail: decastro.luciano@gmail.com.

Department of Economics, University of Maryland. E-mail: pcramton@gmail.com

[1]Section II explains in more detail why load forecasting is useful and offers a brief account of the respective literature.

we recognize that more sophisticated mechanisms can be proposed, we argue in section VI that the implementation of these mechanisms may face some difficulties. Section VII is a brief conclusion, where we explain why this market solution is not subject to the problems usually associated with stock markets, such as speculation and the formation of bubbles.

II. Electricity demand forecasting

The planing and operation of electricity markets require forecasting many relevant variables: load (power requirements, in MW), energy consumption (in MWh), availability of generation and transmission systems, raining (in hydro systems), price of fuels, etc. The forecasting of these variables require different techniques and have diverse purposes. Although the ideas presented in this paper can in principle be adapted to the forecasting of these variables, we will focus on load or electricity demand forecasting.

The forecasting of electricity demand can be classified in short-term load forecasting (from a few minutes to one day), medium-term load forecasting (from one day to one year), and long-term load forecasting (for periods greater than one year).[2] Since different plants have different startup time and costs, short-term forecasting can be used to optimally allocate electricity generation among plants. It is also useful to evaluate net interchange and to make system security analysis. Medium-term forecasts are useful to maintenance planning, fuel scheduling and hydro reservoir management. In regulated environments, where electric utilities are required to sign contracts to cover their demand, these forecasts can also be used to adjust the utility's number of contracts, by bilateral trades. Long-term forecasts are required to assess the necessity of construction of new plants or new transmission systems, which is essential for preventing electricity shortages. Besides these reasons for doing load forecasting, some regulatory bodies, such as the North American Electric Reliability Corporation (NERC), require system operators to report demand forecasts (see [3]). The regulatory use can vary from just system security, to the determination of tariffs (when the regulated price is fixed by taking in account the total expected costs and the expected load and energy required during the year). For simplicity, our discussion will focus on long-term load forecasts (but see section V-A).

The importance of load forecasting is reflected in the number of papers published about the subject. The literature on this topic grows so fast that researchers have felt the need of producing periodic surveys, such as [4], [5], [6], [7], [8] and [2]. Also, [9] surveys only artificial intelligence (AI) methods for demand forecasting. [8] compares some of those methods.

[2] classify the methods for load forecasting into nine categories.[3] These categories reveal not only different approaches, but also specific assumptions and model choices. This is natural, since electricity markets can be very different across regions and what works well in one place may not work in other. It is useful to illustrate this general characteristic with some examples.

In a region with many electricity intensive plants, the load will be very correlated to the growth of those plants' industry. In another region where such plants do not exist, the performance of that industry is an irrelevant variable. If we consider a rich residential area in temperate climates, the temperature may be an important factor in predicting demand. However, in tropical climates, where temperatures are more stable along the year, where residential heating and cooling are not so common, the same variable is possibly not as informative as before. In most cases, electricity demand is correlated to GDP growth, but sometimes this dependence may not be simple and can vary across regions as [10] show. Also, some specific variables can be quite useful in particular regions, as the sales of refrigerators or other electrical equipment in developing countries, where many people are for the first time buying them.

All these examples suggest that the choice of a method can be much harder than just imitating what other regions or countries do. The underlying assumptions of each model and the specific conditions of each region do matter. This create the problem of how to take advantages of the constant technological advances in forecasting methods, without incurring in the high cost of testing every available model. But the problem is not only technical, as the discussion so far suggests. An important aspect of the load forecasting problem is exactly the design of the forecasting system, that is, who does the forecasts and what incentives they have.

Again, the organization of the forecasting systems varies across countries and regions.[4] The most common organizations can be broadly classified in private, public and mixed. In a private system, the forecast is prepared directly by an investor-owned utility. A forecasting system is public if the people involved in preparing the forecasts are technicians working with a public body. In a mixed system, there is a committee formed by participants working to electric utilities and other working to regulators. In all these organizations, we can have economic incentive problems.[5]

Most systems lack a clear incentive mechanism that compensates for the effort for getting more accurate forecasts. Even if the technicians responsible for the forecasts are well-intentioned and dedicated, they may have no incentives to experiment new methods or look for alternative models and sources of information. In this situation, they can systematically miss opportunities to improve their forecasts. Worse

[2] See [1] and [2].

[3] These categories are: (1) multiple regression, (2) exponential smoothing, (3) iterative reweighted least-squares, (4) adaptive load forecasting, (5) stochastic time series, (6) ARMAX models based on genetic algorithms, (7) fuzzy logic, (8) neural networks and (9) expert systems.

[4] Another aspect that varies and we will not discuss here is the use of the forecasts. This use can vary from an informative purpose, to a formally important number, to be used as a mandatory requirement for contracting new capacity. Also, in some places the construction of new capacity is a decision left exclusively to electric utilities, while in others it is a government decision, passing also by the intermediary case where a regulatory body enforces the realization of capacity auctions, where contracts for future provision of electricity are negotiated.

[5] Most papers published on load forecasting deal with technical issues and it is natural that they ignore economic incentives. We are unaware of papers pointing out to the incentive problems treated here.

than this, the forecasting system may be organized in such a way that there are incentives for producing biased forecasts.

In a private system, the companies responsible for the forecasts may benefit from a forecast that is biased in one direction or another. For instance, incumbent generators may be interested in keeping new generators out of the market and if a lower forecast is capable of reducing the likelihood of new generation being mandated or contracted, they have a clear incentive to reduce the forecasts.[6] On the other hand, if builders of new generators are involved in the forecast, they may be interested in maximizing the construction of new plants, in which case it is in their benefit to increase the forecasts. If no mechanism is introduced to correct for these incentives, it is likely that the forecasts based on a private system will be biased.[7]

In a public system, the technicians can have biased incentives as well. Indeed, if their forecasts are always above what should be, and therefore a shortage never happens because excessive capacity is built, then no questions are likely to be made. On the contrary, when there is shortage and the forecast was below the realized load, then the media can picture the forecast as a "mistake" and the body of technicians will have a hard time explaining it. Therefore, it is natural that the forecasts are biased in the direction of high values. If politicians are also involved in the decision to contract more capacity, as it is the case in many places, the influence that they may have in a public forecasting system can produce other kind of distortions. The politician can be connected to companies with interests in the electricity market, so that all the bad incentives described above for private systems can play a role also in public systems. On the other hand, if building new capacity requires public funds, and it does not have a good electoral payoff in the short-run, the politician may prefer to allocate funds to other more appealing projects. In this case, the politician would be interested in low forecasts, to justify its lack of investment. The incentives in a mixed forecasting system is also a mix of all incentives discussed for private and public systems.

As we can see, there are many problems with the current way of doing forecasts. From the society point of view, it will be desirable to have the forecasts of all technical models at the same time, making it possible to compare them. But this is not yet enough, since this creates the problem of aggregating those forecasts or choosing the best. It would be better to put those different forecasts to "compete" with each other in order to get an aggregated number. Also, it is desirable to have a system where the mentioned incentives can be confronted and hopefully canceled out. But what system could approximate this ideal? The question already suggests its answer: it would be important to create an environment where different forecast methods can be used, and the best ones are rewarded for their performance. This is a rough description of what a prediction market does. Our proposal is therefore a natural response to the problem of load forecasting.

Of course, the proposal needs clarification. This requires first to understand better how prediction markets work, which is the objective of section III. But this is still not enough, because load forecasting deals with continuous variables and the most successful prediction markets deal with binary variables. Section IV contains a central contribution of this paper, which is a proposal for a simple and effective continuous variable prediction market, suitable for electricity demand forecast. This proposal brings as additional advantage a better understanding of the entire cumulative distribution of the forecasted demand. To understand the importance of this aspect, consider figure 1 below.

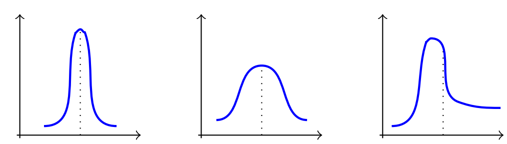

Figure 1: Graph of different p.d.f.'s of electricity demand, with the same expected demand X^*. The differences in shapes may be important for planning purposes.

Figure 1 depicts hypothetical distributions of the random variable "electricity demand." A forecast is usually understood as the expected value $X^* = E[X]$ of this random variable. As the different distributions in figure 1 illustrate, it is not enough to know the value of the expected demand, because this is the same in the three cases, but these cases are very different in terms of the risk associated with the realization of the expected value. This is particularly relevant for electricity demand, since the variance (and high order moments) can be used to access the risk of a shortage (a highly costly event). Also, it may give more reliable numbers to make a good trade-off between the risk of shortage and the costs of building new capacity. Although some forecasting methods are capable of providing information on these high order moments and some reports do include this information, it is not clear that the preoccupation for obtaining them are sufficiently emphasized.

III. A PRIMER ON PREDICTION MARKETS

The recent attention devoted to prediction markets both in academic circles and the media, including a best-selling book ([11]) may suggest that the use of prediction markets is a recent phenomenon. However, [12] reports that a market for US presidential elections has functioned on Wall Street for many years since the Civil War (1880-1944). In fact, prediction markets are essentially betting markets, and betting markets exist since primitive ages. [13, p. 78] observes that "betting on horse races is a custom dating back thousands of years" and adds the ironic comment that "one

[6] Of course there are limits for this manipulation and it is in general difficult to prove that the forecasts are manipulated in this way. For our discussion, however, it is sufficient to recognize the presence of the incentives for doing so.

[7] In Brazil, the regulators ask distribution companies to forecast their future demand, and require that they contract energy to meet such demand. There are penalties if these companies are over or under contracted when the time arrives. This is in the direction of giving incentives for correct forecasts.

may wonder whether the comparatively recent development of trading in corporate equity will prove to be as durable an institution." The fundamental characteristic of prediction markets that distinguish them from standard betting markets is the centrality of the information contained in the traded prices. Accordingly, the evidence on betting markets spans a long time period. In general, betting markets are accurate in predicting outcomes, except in the tails, where there is a longshot bias. See [14] and [15].

This increasing interest in prediction markets is reflected in the explosion of articles about such markets, as the surveys [16], [17] and [18] report.

In the following section, we briefly review the evidence on the good properties of prediction markets and, in the subsequent section we describe the debate about the theoretical foundations of these findings. Then, section III-C provides a (partial) justification for the notion that prices in prediction markets reflect probabilities, alongside the framework for our proposal.

A. Evidence that prediction markets perform well

The performance of prediction markets is better known in the outcomes of political processes, as the Iowa Electronic Market (IEM), functioning since 1988, illustrates. The IEM was first discussed by [19], but see also [20], [21] and [22]. In particular, the IEM has been shown to (a) predict well both shortly before an event ([21]) and through time ([22]); (b) forecast better than alternative means ([22]); and (c) be accurate not just on average, but on a case-by-case, contract-by-contract basis ([21]). In the working paper version of this article [23], we provide other sources of evidence that prediction markets perform well.

B. Theoretical foundations of prediction markets

As the discussion above shows, there is strong empirical and experimental evidence that predictions markets do work. But why does this happen? A few theoretical papers consider the explanation of this success. [24] shows in a simple model with investment constraints that prices *should not* be interpreted as probabilities. Trying to offer a theoretical justification, [25] find that the equilibrium price of an Arrow-Debreu security written on an event E is given by the mean or a certain quantile of the distribution of beliefs among traders about the probability of E. [26] assume that market participants have heterogenous beliefs and consider a rational expectations equilibrium. They show that prices under-react to the arrival of new information.

However, apparently we still do not have a good theory to explain *why* such markets perform so well. Although [20] forcefully argue in favor of prediction markets, they concede that theoretical explanations are limited and quote Vernon Smith, who wrote: "Things sometimes work better than we had a right to expect from our abstract interpretations of theory" ([27]).

C. Prices as probabilities

Prediction markets may be created with many different types of assets or contracts. See [17] for a description of the more common assets. Although we will also mention some of those below, our method will rely only on a simple set of assets: the "winner-takes-all" contract. A "winner-takes-all" contract is an asset that pays \$1 if a well-specified event has occurred up to some specific date and \$0 otherwise.[8]

Throughout the paper we will assume that the price of a "winner-takes-all" asset associated with an event reflects the probability of that event, given the aggregated information of market participants. It is useful to state this assumption formally. For this, we will need some notation.

Let E be a well specified event and let $\Pr(E)$ denote the best possible probability prediction for the occurrence E, that is, the probability that perfectly aggregates all the information possessed by market participants. Also, let p_E denote the market price of the winner-takes-all asset based in the event E. If the event is clear from the context, we will write p instead of p_E. The basic assumption that justifies the use of prediction markets is the following:

Assumption 1 (Information aggregation): For any well specified event E, $p_E = \Pr(E)$.

But why is assumption 1 reasonable? For a single individual, this assumption is reasonable almost by definition. Indeed, one of the most important contributors to the foundations of probability, *defined* probability as the price that the subject would consider equivalent to receiving a winner-takes-all contract as described above. See [29, Chapter 3] for an extensive discussion and justification. Things become more complicated when we want to aggregate the different beliefs of individuals in a single price. The theoretical debate mentioned in section III-B is exactly over how this aggregation occur and how it produces good estimates.

From Assumption 1, we will refer to p interchangeably as the market prediction price and as the probability that the event occurs. Assumption 1 is a basic working assumption for all applications of prediction markets. If it does not hold, the foundations of prediction markets are undermined. We stress, however, that this is an *assumption*: it is not true in all possible markets and a prediction market designer or operator must verify—examining accumulated data— that it remains at least approximately true.

IV. Prediction markets with continuous variables

While some uncertain outcomes are clearly binary in their nature, such as the event that the Republican candidate wins the American election, many others correspond to continuous random variables. For instance, the market share of a new product, the percentage of electoral vote by some candidate, or the electricity demand at a future time. At first sight, it is not obvious how to use markets with only winner-takes-all assets in order to obtain information about the probability distribution of these variables. In this section, we review how

[8] It is important that the contract clearly specifies a fixed date for the event, and the event is defined in detail. As an illustration of the potential problems, [28] reports a case of a prediction market on whether a software project would be delivered to the client on schedule, but the client changed the deadline.

previous papers have handle this question (section IV-A) and then propose our own solution (section IV-B).

A. Existing literature

[30] proposes market scoring rules, which combines ideas of scoring rules and the standard design of prediction markets.[9] At each time t, there is a *current* probability distribution p^t and a market participant reports some probability r^t that she thinks is the most correct one and this will lead to the probability for the next period, p^{t+1}. This is associated through some scoring function to a cost (or gain) related to the change of probabilities. The final payoff of the market participants will be the sum of gains and losses along the trades in the different periods. Many scoring rules can be adopted, but quadratic and logarithmic scoring rules are the more common. See [30] for more details.

It is possible to obtain the expectation of a continuous random variable using the index contracts, as [17] call them. Index contracts pay \$1 for each unit of the outcome that is realized. For example, if the units are percentage points of popular vote for the Democratic candidate, then the contract will pay \$44 if the Democratic candidate obtained 44% of the popular vote. The idea behind this kind of contracts is that the price of the contract should be equal to the expectation of X, which is denoted $E[X]$, that is, $p = E[X]$.

Of course this idea is in the same spirit of Assumption 1, since $E[X] = \Pr(\{X = 1\})$ for a binary variable. However, Assumption 1 *does not* imply that the price p of the index contract is equal to $E[X]$.[10] The reason is that this assumption refers only to events or binary variables. Of course we could require Assumption 1 to hold also for continuous variables but, as we stressed before, Assumption 1 is not free of controversy. Despite the justifications for that assumption, its conclusion is less tested for continuous (index) contracts than it is for binary (winner-takes-all) contracts. Perhaps there is no gap between the properties of winner-takes-all and index contracts prediction markets, but it is better to be conservative in this matter and work with the weakest conditions that can deliver the desired result.

Also, note that this approach only gives information about the *expectation* of X. As mentioned in the introduction, we often are interested in obtaining more information about the distribution of X. For addressing this problem, [17] propose the use of contracts that pay x^2 dollars if the outcome of X is x. In this way, the price of the contract will be $E[X^2]$ and from this and $E[X]$, one can obtain the standard deviation $\sigma_X = \sqrt{E[X^2] - (E[X])^2}$ of X. They conclude that "adding even more complicated index contracts can yield insight into higher-order moments of the distribution" (p. 110).

Yet another method to make evaluations of continuous variables with prediction markets is the one used by [31] in the evaluation of future sales of a new product developed by HP. They divided "the real line into about 10 or so (exact number depends on the event) intervals". If the final outcome fell in an interval, the corresponding security would pay \$1 per share at the end. Note that each asset is a winner-takes-all (binary) asset. Therefore, Assumption 1 applies. [31] report strong results in the HP experiment using this implementation. The described implementation was successful, but it has at least a potential limitation for other cases. If one is interested in more accuracy in the outcome, then the only way to achieve this accuracy is to increase the number of intervals. However, if there is a large number of intervals, then the problem of *thin markets* becomes important: there is not enough trade in each asset to make the information reliable. Another problem, which actually happened in [31]'s experiment, is that the sum of prices (that is, probabilities, from our assumption) may not equal 1. Of course this may be attributable to the lack of expertise or the lack of enough trade in the markets. However, the occurrence of this fact clearly undermines the argument for Assumption 1 and raises some doubts about the approach. As we will show in the next section, our solution avoids all these problems and yet remains simple.

B. Cumulative Distribution Function Prediction Market

In this section, we describe a simple method of implementing a prediction market for continuous variables using only winner-takes-all (binary) assets. In order to do so, let us take the values $x_1 < x_2 < \cdots < x_n$ in the set of possible values of the variable X. Now, for $k = 1, \ldots, n$, let E_k denote the event that the realization of X is not greater than x_k, that is, $E_k \equiv [X \leqslant x_k]$. The cumulative distribution function prediction market (CDF-PM) that we propose is simply a prediction market with n winner-takes-all contracts based in the events E_k, for $k = 1, \ldots n$. The following analysis justifies this name. Let p_k be the price of the asset k, that pays \$1 if E_k occurs and nothing otherwise. Then, by Assumption 1 we have: $p_k = \Pr\{E_k \text{ occurs }\} = \Pr[X \leqslant x_k] = F_X(x_k)$, where F_X denotes the cumulative distribution function (CDF) of X. This shows that the price p_k of the contract k is just the probability that the realization of X is below x_k, that is, p_k gives the value of the CDF for the chosen values $x_1, x_2, \ldots x_n$. Figure 2 below illustrates this.

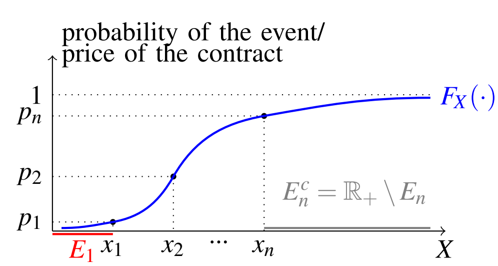

Figure 2: Graph of the CDF $F_X(\cdot)$, with the prices of the contracts.

Note that a simple arbitrage argument guarantees that the prices are increasing, that is, $p_k \leqslant p_{k+1}$ for all $k = 1, \ldots, n-$

[9] [30] is not restricted to continuous variables. In fact, his main concern is the case of binary variables. However, since he can treat any combination of a number of binary variables, the continuous case can be approximated.

[10] Assumption 1 does imply $p = E[X]$ under some assumptions on the preferences and the market.

1. Indeed, if $p_k > p_{k+1}$, a trader can guarantee making money by buying asset $k+1$ and selling asset k. Therefore the points in the CDF F_X produced by the market will be monotonic. Also, note that it is not necessary that $F_X(x_n) = 1$, since the event $E_n^c = [X > x_n]$ can have a positive probability.

It should be noted that when the contract ending date is closer, for instance, the end of the summer for contracts associated with the peak load in the summer, then there is more information about the realization x of the demand. As a result, the value of some contracts will go to zero (those contracts that specify $x_k < x$) and the value of others will go to to \$1 (those contracts with $x_k \geqslant x$). In the end, the market will have a degenerate CDF function, concentrated in one value, x.

Let us compare this procedure with the procedures previously described. First, while index contracts yield information only about the expectation of X, the above procedure will provide much more information. By adding additional points, we can get finer information about the distribution of X. Since this procedure gives a good approximation of the whole CDF, we can calculate all moments we need.

Note that we can create new contracts even after the market is initiated. In fact, this ability to create new contracts may even be desirable. Let us suppose that there is an important jump in the prices of the contracts k to $k+1$. This indicates that there is a large probability of X being between x_k and x_{k+1}. Since we may not know this before the experiment, this problem could not be anticipated. This large probability indicates that we may be interested in knowing the distribution between x_k and x_{k+1} with more detail. To obtain this information, we just need to choose a value v between x_k and x_{k+1} and create a new asset for the event $[X \leqslant v]$. The monotonicity property described above then implies that the price of the new asset will be between p_k and p_{k+1}. Note that nothing changes for the other contracts.

Note also that since prices of new contracts already come restricted to some intervals, the *thin markets* problem is reduced. Even if there is absolutely no trade in the asset k, we know that its price *is* between p_{k-1} and p_{k+1}. So, we can have a large number of contracts, with small trade in each of them, but the market still works sufficiently well—provided that there is sufficient trade overall.

These advantages also make clear why this method is better than the evaluation of frequencies used by [31]. Not only does the creation of new intervals become problematic with the frequency approach, but the *thin markets* problem may also be severe. As we discussed, our approach avoids these problems.

Other authors have previously discussed or directly used the kind of contracts proposed here. For instance, the site *intrade.com* has contracts similar to those considered here, that allow obtaining the correspondent CDF, and could be viewed as forming a CDF-PM. [32], [33] and [34], among others, use some contracts of the form proposed here. The point is not that our proposal was never used or considered, but that it presents advantages over the methods discussed in section IV-A. To the best of our knowledge, the received literature does not discuss those advantages.

V. Implementation of Prediction Markets for Electricity Demand

The implementation of a prediction market for future electricity demand, as we suggest in this paper, requires dealing with a number of issues that go beyond the accuracy of these markets. We discuss below the definition of the contracts, the possibility of manipulation of the market, some political aspects that may affect the implementation, the consideration of additional variables in the market and other issues.

A. Specification of contracts and market design

Using the method described in the previous section, we can design prediction markets for the future demand of electricity. It is necessary to begin by defining how the future demand will be measured and verified by establishing a clear, completely technical, non-manipulable procedure. The geographic area and the interval of time for which the future demand is considered are also important. Once these specifications are made, the market designer has to choose values $x_1 < \cdots < x_n$ that cover the likely values of the demand X. From this, the cumulative distribution function prediction market (CDF-PM) described in section IV-B can be implemented. It should be noted that the market designer has the freedom to create new contracts after the beginning of the actual implementation of the market, as we also discussed in section IV-B.

Of course, there are a number of details that have to be clarified. One of them is who can participate and what is the maximum amount traded by each market participant (if any). Many prediction markets have included limits in the participation, probably to avoid manipulation and excessive risk taking by some individuals. However, low limits may reduce the liquidity of the prediction market, undermining its function. The literature has not provided guidance so far in the proper way to evaluate this trade-off.

It should be noted that for short-term or medium-term load forecasting, the general organization of the market can be the same. The main difficult in these cases is to guarantee that the markets have enough participation for the desired aggregation of information to happen. In short-term forecasts, the requirements for participation should be higher because there is less time for the trades to occur. However, if the participation is high enough, the CDF prediction markets could also be used for short and medium-term forecasting.

B. Manipulation

The possibility of manipulation is a real and important concern. If some participants are sufficiently big and have special interests in the outcome of the market, then they may act to undermine the performance of the prediction market. This is an important problem that is still not completely understood. For example, [35] placed this problem in their list of five open questions regarding prediction markets, but report that known attempts to manipulate public prediction markets have largely failed. Further discussion about this issue can be found in the working paper version of this article [23].

C. Political aspects—necessity of a communication plan

In some countries, market-oriented institutions such as the one proposed here can face spontaneous opposition just for being "pro-capitalism." Even in pro-capitalist countries, such as the US, this kind of resistance can appear. A curious example was the proposal of (and public outcry against) a prediction market for terrorism, described in detail by [36].

Fortunately, the application suggested here is about an *economic* phenomena (electricity demand) and not a politically controversial topic as terrorism. This fact can help to avoid the main problems involved in the actual implementation of the market. However, in countries where market oriented proposals are negatively seen and can, therefore, be politically explored, some precautions should be taken and a "communication plan" seems desirable.[11]

D. Dependence on other variables

In some circumstances, it may be useful to obtain predictions conditional to the realization of relevant variables. For instance, it may be useful to know how the forecasts changes with variables such as temperatures, economic growth or even the price of electricity itself.[12] Indeed, although the demand for electricity is inelastic in the short-run, resource planning and capacity markets may operate five or more years in advance and, with this time interval, the demand can be more elastic. For example, plants that require large electricity inputs, such as aluminum smelters, may not be constructed depending on the future price of electricity. Thus, it may be desirable that the prediction market allows for the dependence of the demand with the price.

The objective of including additional variables can be achieved using conditional contracts, such as those proposed by [37] and [38]. We adapt their idea for future demand of electricity as follows.

Let $P_1, \dots, P_m$ be a set of events covering all relevant values of the additional variables. These sets can form a partition of all conceivable values, but this is not necessary. For example, P_j can be the event that the price of electricity will be between $\underline{y}_j$ and $\bar{y}_j$ or, alternatively, that the price of electricity is below (or above) y_j. In any case, let E_k be the event that the demand is below x_k, as described in section IV-B. Then, let us consider a market with winner-takes-all contracts based on the events P_j (whose price will be denoted p_j) and $E_k \cap P_j$ (whose price will be denoted p_{kj}), for all $j = 1, \dots, m$ and $k = 1, \dots, n$. Then, Assumption 1 implies that the conditional CDF $F(x_k|P_j) = \Pr(E_k|P_j) = \frac{\Pr(E_k \cap P_j)}{\Pr(P_j)}$ is given by p_{kj}/p_j. This gives the CDF of the demand of electricity conditional to electricity prices being on P_j, as we wished to obtain.[13]

E. Legality and other issues

One of the most important issues regarding prediction markets is its legality. Indeed, some states and the federal government in US have issued regulation against Internet gambling that might apply to prediction markets as well [39]. Giving the benefits that prediction markets may bring, a number of important researchers have proposed solutions to clarify the legality of prediction markets [40].

There are still other issues that need to be properly addressed in an actual implementation: duration of the contract and how to make its value persistent in time (maybe just inflation-free or with some small interest); how to ensure payments at the end of the contract; the custody mechanism; the amount of subsidy, if any;[14] limits for participation per individual or firm; and the mechanisms for monitoring activity (for detecting attempts of manipulation), without constraining truthful bets. Although the resolution of all these issues is important, the best implementation will probably vary from case to case.

VI. More sophisticated mechanism designs

Although we have considered prediction markets only to give information about future demand, of course it is possible to conceive more sophisticated mechanisms in which the demand may have a more active role.

In this way, the prediction of the future demand can enter directly the market for future capacity. For this, we must have double auctions, that is, auctions in which both demanders and suppliers have active participation.

There are, however, some practical problems in pursuing this idea. First, there may be resistance from distribution companies not used to having an active role in capacity markets. Second, it is not clear how susceptible to manipulation such a market would be. Third, political pressure and lack of confidence in the performance of an (unknown) market may make the actual implementation difficult.

VII. Concluding remarks

This paper proposes the use of prediction markets to inform the regulatory bodies that are responsible for resource adequacy. The proposed method is simple and easy to implement. Moreover, it provides much more information than just the expected demand, since it obtains higher order moments with any desired level of precision. We also described the limitations of the current approaches, both to load forecasting and to continuous random variable prediction markets.

Our solution is a market-based proposal. It is natural that market skeptics react negatively to such an idea. From this

[11] A more extended discussion of this topic can be found in the working paper version of this article [23].

[12] The economic concept of demand already takes in account its dependence with respect to price. Thus, the term *load* forecasting seems better than *electricity demand* forecasting, if the price is not considered. Despite this inaccuracy, we follow the standard practice and use both terms as synonymous in this paper.

[13] Note that this implementation also gives a forecast for the events P_j, which can be useful for market participants for independent reasons. For instance, if the variable is future electricity price, this information can be useful in the decision of investing on an electricity intensive industry.

[14] Experience has shown that prediction markets are inexpensive to create and maintain. Some authors advocate for the use of subsidies for promoting participation, but the necessary level of subsidies is probably small.

point of view, the recent problems in the financial markets would be a "proof" that prediction markets do not work.[15] While it is not convenient to engage in a far reaching debate over financial markets, we would like to point out that the contracts that are traded in the proposed cumulative distribution function prediction market (CDF-PM) are not subject to many of the problems that assets in the financial sector may have. Assets subject to "bubbles" are those in which arbitrarily high expectations about future price levels are sustained because prices can increase without bound for indefinite periods of time. The prediction market contracts do not have this characteristic, because they always have a maximum price ($1) and have a definitive date for clearing. In a sense, irrational prices cannot be sustained by beliefs about future prices. Thus, most of the possible problems with financial markets seem less relevant in the proposed prediction market.

The strong evidence on the performance of prediction markets presented in section III suggests that the CDF-PM can perform better than the current forecasting methods, even without considering the incentives problems that such methods may have. Thus, it is likely that CDF-PM can improve load forecasting, with significative gains for electricity markets.

REFERENCES

[1] D. Srinivasan and M. Lee, "Survey of hybrid fuzzy neural approaches to electric load forecasting," vol. 5, oct 1995, pp. 4004 –4008.

[2] H. K. Alfares and M. Nazeeruddin, "Electric load forecasting: literature survey and classification of methods," *International Journal of Systems Science*, vol. 33, no. 1, pp. 23–34, 2002.

[3] NERC. (2009) Standard mod-017-0.1 aggregated actual and forecast demands and net energy for load.

[4] P. Matthewman and H. Nicholson, "Techniques for load prediction in the electricity-supply industry," *Electrical Engineers, Proceedings of the Institution of*, vol. 115, no. 10, pp. 1451 –1457, october 1968.

[5] M. A. Abu-El-Magd and N. K. Sinha, "Short-term load demand modeling and forecasting: A review," *Systems, Man and Cybernetics, IEEE Transactions on*, vol. 12, no. 3, pp. 370 –382, may 1982.

[6] D. W. Bunn and E. D. Farmer, *Comparative models for electrical load forecasting*. Chichester [West Sussex]; New York: Wiley., 1985.

[7] G. Gross and F. Galiana, "Short-term load forecasting," *Proceedings of the IEEE*, vol. 75, no. 12, pp. 1558 – 1573, dec. 1987.

[8] I. Moghram and S. Rahman, "Analysis and evaluation of five short-term load forecasting techniques," *Power Systems, IEEE Transactions on*, vol. 4, no. 4, pp. 1484 –1491, nov 1989.

[9] K. Metaxiotis, A. Kagiannas, D. Askounis, and J. Psarras, "Artificial intelligence in short term electric load forecasting: a state-of-the-art survey for the researcher," *Energy Conversion and Management*, vol. 44, no. 9, pp. 1525 – 1534, 2003.

[10] S.-H. Yoo and J.-S. Lee, "Electricity consumption and economic growth: A cross-country analysis," *Energy Policy*, vol. 38, no. 1, pp. 622 – 625, 2010.

[11] J. Surowiecki, *The wisdom of crowds: Why the many are smarter than the few and how collective wisdom shapes business, economies, societies and nations*. Doubleday, 2004.

[12] P. W. Rhode and K. S. Strumpf, "Manipulating political stock markets: A field experiment and a century of observational data," *University of North Carolina*, 2006.

[15] Another possible doubt is whether the current markets for electricity cannot already work as prediction market. The problem is that standard markets or future markets for electricity have contracts based on prices, not directly in demand. Future prices clearly depend on the future demand, but also in future supply. Thus, a single price can correspond to different loads, which implies that future electricity prices are not sufficient for load forecasting.

[13] S. Figlewski, "Subjective information and market efficiency in a betting market," *The Journal of Political Economy*, vol. 87, no. 1, pp. 75–88, 1979.

[14] R. H. Thaler and W. T. Ziemba, "Anomalies: Parimutuel betting markets: Racetracks and lotteries," *The Journal of Economic Perspectives*, vol. 2, no. 2, pp. 161–174, 1988.

[15] R. D. Sauer, "The economics of wagering markets," *Journal of Economic Literature*, vol. 36, no. 4, pp. 2021–2064, 1998.

[16] G. Tziralis and I. Tatsiopoulos, "Prediction markets: An extended literature review," *Journal of Prediction Markets*, vol. 1, no. 1, 2006.

[17] J. Wolfers and E. Zitzewitz, "Prediction markets," *The Journal of Economic Perspectives*, vol. 18, no. 2, pp. 107–126, 2004.

[18] D. Pennock and R. Sami, "Computational aspects of prediction markets," *Algorithmic Game Theory*, pp. 651–674, 2007.

[19] R. Forsythe, F. Nelson, G. R. Neumann, and J. Wright, "Anatomy of an experimental political stock market," *The American Economic Review*, vol. 82, no. 5, pp. 1142–1161, 1992. [Online]. Available: http://www.jstor.org/stable/2117471

[20] J. Berg and T. Rietz, "The Iowa Electronic Markets: Stylized Facts and Open Issues," *Information Markets: A New Way of Making Decisions, edited by Robert W. Hahn and Paul C. Tetlock. AEI-Brookings Joint Center for Regulatory Studies*, pp. 142–169, 2006.

[21] J. Berg, R. Forsythe, F. Nelson, and T. Rietz, *Results from a dozen years of election futures markets research. In: Plott, C. R. and Smitt, V. L. Handbook of Experimental Economics Results, vol. 1,*. Elsevier, 2008, ch. 80, pp. 742–751.

[22] J. Berg, F. Nelson, and T. Rietz, "Prediction market accuracy in the long run," *International Journal of Forecasting*, vol. 24, no. 2, pp. 285–300, 2008.

[23] L. I. de Castro and P. Cramton, "Prediction markets to forecast electricity demand," University of Illinois at Urbana-Champaign, Tech. Rep., 2009.

[24] C. Manski, "Interpreting the predictions of prediction markets," *Economics Letters*, vol. 91, no. 3, pp. 425–429, 2006.

[25] J. Wolfers and E. Zitzewitz, "Interpreting prediction market prices as probabilities," *IZA Discussion Paper No. 2092*, 2006.

[26] M. Ottaviani and P. N. Sorensen, "Aggregation of information and beliefs in prediction markets: Lessons for asset pricing," Kellogg School of Management, Northwestern University, Tech. Rep., 2008.

[27] V. Smith, "Experimental economics: behavioral lessons for microeconomic theory and policy," *Frontiers of Research in Economic Theory: The Nancy L. Schwartz Memorial Lectures, 1983-1997*, p. 104, 1998.

[28] G. Ortner, "Forecasting Markets—An Industrial Application," *Universität Wien*, 1998.

[29] B. de Finetti, *Theory of Probability, Vol. 1-2*. Wiley, New York, 1974.

[30] R. Hanson, "Combinatorial information market design," *Information Systems Frontiers*, vol. 5, no. 1, pp. 107–119, 2003.

[31] K. Chen and C. Plott, "Information aggregation mechanisms: Concept, design and implementation for a sales forecasting problem," in *Lee Center Workshop*, 2002.

[32] M. K. Chen, J. E. Ingersoll, Jr., and E. H. Kaplan, "Modeling a presidential prediction market." *Management Science*, vol. 54, no. 8, pp. 1381 – 1394, 2008.

[33] J. Berg, G. Neumann, and T. Rietz, "Searching for Google's Value: Using Prediction Markets to Forecast Market Capitalization Prior to an Initial Public Offering," *Management Science*, vol. 55, no. 3, pp. 348–361, 2009.

[34] J. Berg, J. Geweke, and T. Rietz, "Memoirs of an Indifferent Trader: Estimating Forecast Distributions from Prediction Markets," University of Iowa, Tech. Rep., 2009.

[35] J. Wolfers and E. Zitzewitz, "Five open questions about prediction markets," *SSRN*, 2006.

[36] R. D. Hanson, "Designing real terrorism futures." *Public Choice*, vol. 128, no. 1/2, pp. 257 – 274, 2006.

[37] R. Hanson, "Decision markets," *IEEE Inteliggent Systems and*, vol. 14, no. 3, pp. 16–19, May-June 1999.

[38] J. Berg and T. Rietz, "Prediction markets as decision support systems," *Information Systems Frontiers*, vol. 5, no. 1, pp. 79–93, 2003.

[39] R. Hahn and P. Tetlock, "A new approach for regulating information markets," *Journal of Regulatory Economics*, vol. 29, no. 3, pp. 265–281, 2006.

[40] K. Arrow *et al.*, "The promise of prediction markets," *Science*, vol. 320, pp. 877–878, 2008.