Modeling Prediction Markets with Dynamic Bayesian Networks

Ethan Z. Shen
Department of Computer Science, Stanford Univesity
Cole R. Winstanley
Department of Symbolic Systems, Stanford University
{ezshen, colew}@stanford.edu

Abstract

Prediction markets have shown a remarkable ability to predict outcomes. Here, we propose a Dynamic Bayesian Network model to extract information and infer prediction market prices by modeling interactions between agents. We validate our methods using poll and price data from the 2012 presidential election, and show that this model is more effective at predicting price trends than previous published methods. We finally explore the robustness of the model to variations in agent information and noise. Given the recent recent resurgence of prediction markets, our work builds upon the current literature on prediction market analysis, which has implications for large-scale, self-incentivized outcome prediction.

1. Background

Prediction markets are platforms on which users can trade or price securities based on the outcome of an event. For example, one might buy a security that pays one dollar if Trump is the Republican Party nominee in 2020, and zero dollars otherwise. The key features here is that the outcome of the event is currently uncertain, but there will come a time when the outcome of the event is known to everyone. Presumably, there also needs to be some potential indicators of what the outcome of the event will be, i.e. the outcome cannot be totally random.

Prediction markets are useful as faithful extractors of information from the general public [1]. In general, markets aggregate information in many cases better than experts can, presumably due to greater overall access to information that all of the market participants have, even when compared to experts in a given domain. For an application of this fact, see Figure 1. As an example, orange juice futures are a better predictor of the weather west of Orlando, Florida than actual weather forecasts issued by the National Weather Service [2], a government agency that receives over a billion dollars in annual budget [3]. Additionally, prediction markets are usually much more accurate than poll data or statistical aggregations of polling data [4].

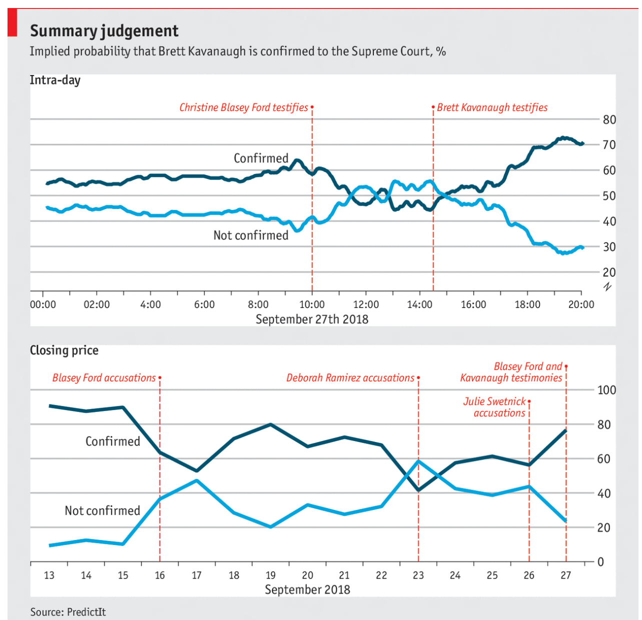

The Economist

Figure 1. *The Economist* uses prediction market data to show changing attitudes surrounding the 2018 congressional testimony of Christine Blasey Ford and Brett Kavanaugh [5].

The power of markets lies in the alignment of incentives in the participants. People trading on a security based on the outcome of an event presumably have some belief about the outcome; their ability to make a profit in the market depends on the accuracy of this belief. Therefore, markets select for the participation of agents who are well-informed about the topic of the market. This accuracy leads a virtuous cycle that only selects more and more for the most qualified agents and traders. However, it may also lead to bias among the participants if the pool becomes too small [6].

Due to regulatory pressures, there are few public prediction markets in use today. This is in large part due to resistance to betting on elections and other events in the United States [4]. Many companies use internal prediction markets [7], but the associated data is proprietary and so not useful for research purposes. The two public prediction markets that have meaningful volume are the Iowa Electronic Markets [8] and PredictIt [9]. Both are research projects at non-profit universities that have received conditional no-action letters from the United States Commodity Futures Trading Commission [10] [11], and

The mathematical backing of modern prediction markets relies on one of two concepts: continuous double auctions or market scoring rules [12]. Market scoring rules are more of a theoretical idea, since both the Iowa Electronic Markets and PredictIt use a continuous double auction for trading. Continuous double auctions operate using the same mechanism as the New York Stock Exchange. Market participants submit buy or sell orders at a certain price for a given security. If an order can be fulfilled, meaning that there is a seller willing to sell at the same or a lower price than a buyer is willing to buy at, the order is executed immediately. Otherwise, orders remain pending and are executed in the order they came in. The result is that there is always a market clearing price, which is the price at which all higher buy orders and lower sell orders are executed. This price (for a security that pays $1) can be interpreted as a consensus probability of whether the event is going to occur or not.

2. Task definition

Prediction markets have proven to be extremely useful research tools, particularly in the previous few years. The strong form of the efficient markets hypothesis [13] would have that the prices in a prediction market accurately reflect the information available to market participants at any given time. However, there are some potential modes of error that prediction markets face, particularly with agents who form coalitions to manipulate the market, agents who trade based on maliciously inaccurate information, and other problems. These potential issues are only exacerbated by the fact that current-day prediction markets have a capped number of participants and total investment, amplifying any abnormalities [14].

2.1. Evaluation Metrics

To evaluate our model, we first define an evaluation metric to compare our inference probability distribution with the actual distribution of the prediction market at each time point. We seek to quantify the information our prediction provides with respect to entropy

$$H = -\sum_{i}^{n} p(x_i) \log p(x_i) [15] \tag{1}$$

Finding the expectation of entropic difference between our prediction and the actual prices lands us at KL divergence:

$$D_{KL}(p||q) = \sum_{i} p(x_i)(\log p(x_i) - \log q(x_i)) \tag{2}$$

We define our loss as the total entropic difference, i.e. the sum of the KL divergence scores over all the time points:

$$\text{Loss}_{KL}(P||Q) = \sum_{t} \sum_{i} P_t(i) \log \left(\frac{P_t(i)}{Q_t(i)} \right) \tag{3}$$

Beyond the quantitative KL loss that we've defined above, we can also see whether our model can roughly approximate specific turning points in prediction market prices in response to significant events, such as scandals, successful rallies or announcements.

3. Infrastructure

We scrape existing polling data from public websites, and build Bayesian Networks using `pomegranate`, a Python package which supports building and inference on discrete Bayesian Networks.

4. Literature Review

In this section, we briefly recount the background of prediction markets.

In 1906, there was a weight-judging competition where eight hundred competitors bought numbered cards for 6 pence to inscribe their estimate of the weight of a chosen ox. The vox populi, or voice of the people, was astonishingly accurate - the average was within 0.08%. This story became immortalized as an early example of what we now call the wisdom of the crowds. The ox story offers a glimpse of of the concept that the aggregation of opinions can be surprisingly good predictors of outcomes, even where individuals are not considered experts [16].

Recent advances in modern machine learning to predict outcomes with high accuracy may suggest the irrelevance of prediction markets. However, in applications where the relevant feature space is too large (political climate), or where data is too costly to aggregate (political polling), prediction markets may present an opportunity for researchers to automatically gain information at a low cost and without domain expertise. [17] Recent interest in prediction markets has manifested in different ways. Current relevant mechanisms include Thompson Reuters lab which uses prediction markets to incentivize entities to reveal unique data in niche markets [17].

The relevance of prediction markets calls for a more nuanced model. To our knowledge, the current state-of-the-art model proposed by Lee et al attempts to capture the interactions between investors, which we will discuss in detail below [18].

5. Model

In this section we will present our final model and discuss its advantages over current models of prediction markets. We will also discuss areas that could serve as targets for improvement.

5.1. Approach

Modeling continuous double auction prediction markets with a Bayesian network is a natural idea because the prices can be directly interpreted as probabilities of the relevant event being realized. This has also been the implicit approach of previous models in the literature since they use Bayesian updating of the agents over time in their models. However, interpreted as Bayesian Networks, these models are overly simplistic when compared to any reasonable conception of the actual flow of information through a market in the real world. This becomes clear when one tries to use them to predict the behavior of prediction markets over a long period of time: they consistently fail to incorporate new information at at a reasonable level, either weighing it too high or too low.

For our model, we will focus on a prediction market that pays off if Barack Obama wins the 2012 election. We will assume that this result is based on the political climate of the time of the election. For us, this represents the hidden variable that is driving the market prices and outcomes. More specifically, we assume that at any given point there is an actual distribution over vote share Δ for the scenario of the election being held at that time. This distribution can change over time, and it drives the signals that market participants use to trade on. For example, the vote share impacts polling (in a noisy way), and it can also impact other market signals, such as the opinions of experts.

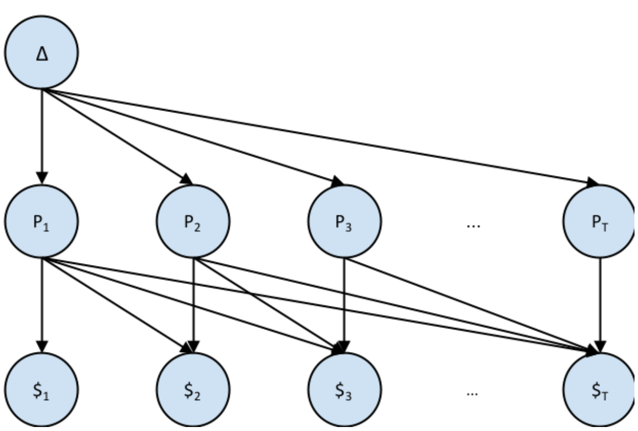

Figure 2. The Bayesian updating model presented in [18].

Starting with the basic Bayesian updating model from [18] that is shown in Figure 2, we targeted the specific shortcomings of that model. In particular, we noted the following:

1. The Bayesian updating model does not account for potential changes in the political climate over time.
2. The Bayesian updating model does not attempt to account for the impact of agents on the market prices, rather going directly from probability to price.
3. The Bayesian updating model does not consider how past prices are themselves a signal that affects the future market prices [19].

For our oracle, we used the actual prediction market prices over the time period. Our evaluation metric is performed in comparison to this oracle.

5.2. Modeling Changes of Political Climate over Time

To reason about how the political climate changes over time, we first decided to reason about how signals of the political climate are affected by the political climate itself. The resulting idea was to have a hidden Markov model for how the political climate changes over time, with emissions visible at each timestep. (In our tests, we use polling data as a proxy for the political climate; see the Data section for details.) We take the political climate at each time step, labeled Δ_t, to be a continuous distribution over possible vote shares:

$$\Delta_t \sim \mathcal{N}(\delta_t, 1/h) \tag{4}$$

where δ_t is the mean and $1/h$ is the precision. We don't explicitly model the change in the distribution from Δ_t to Δ_{t+1} beyond saying that

$$\Delta_t \approx \Delta_{t+1} \tag{5}$$

and noting the relationship between the underlying political climate and the polling results:

$$P_t = \mathbf{Pr}[\Delta + e_t > 0] \text{ where } e_t \sim \mathcal{N}(0, 1/\sqrt{N_t}). \tag{6}$$

This assumes that the polling data is a sample from the actual distribution Δ_t, so the sampling error is e_t, and the sample size of the poll is N_t.

Instead of explicitly modeling the distribution changes between time steps, we directly model the relationships between the polling emissions, which are at the level of real-world data. This is reasonable because we are interested in the relationship between the real-world signals and the prices in the prediction markets. There is already significant amount of work on the relationship between polling data and the underlying distribution of the vote share, which

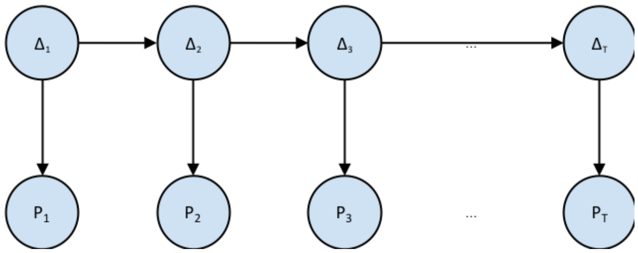

Figure 3. A Bayesian network model for how emissions from the political climate (here, polling data) changes over time.

we are not interested in. Additionally, having the continuous distributions Δ_t in the model adds significant computational complexity, which we can eliminate by only using the binary nodes P_t. We model those by saying that

$$\begin{aligned}\mathbf{Pr}[P_t = A | P_{t-1}] = (1-w)\,\mathbf{Pr}[P_{t-1} = A] \\ + w\,\mathbf{Pr}[P_{t-1} \neq A]\end{aligned} \quad (7)$$

This models the political signals with a random walk between the two results and parameter w, which we can tune in our model.

5.3. Modeling Agents' Effects on Prices

In real-world prediction markets, participants react to information in the domain of the market by affecting the price of the security in the relevant way. However, not all participants have the same information or even interpret the same information in the same way. Individuals may have some personal bias or receive the signal plus some noise. Additionally, some agents may only take market signals from other agents or from past market prices.

Let there be n agents participating in the market. Keeping the number of agents in the market fixed over time may seem unreasonable at first, but is actually accurate in the case of current-day prediction markets. We assume that agents receive some noisy signal of the political climate emission P_t, which we denote Z_{ti}. We use the same scheme as with the P_t to P_{t-1} signal, with some noise parameter γ:

$$\begin{aligned}\mathbf{Pr}[Z_{ti} = A | P_t] = (1-\gamma)\,\mathbf{Pr}[P_t = A] \\ + \gamma\,\mathbf{Pr}[P_t \neq A]\end{aligned} \quad (8)$$

Agents also observe the previous price $\$_{t-1}$, and incorporate it into their predictions. We model the overall prediction belonging to an agent as a linear combination (with parameter e) between the previous price and the new belief based on the data:

$$\begin{aligned}\mathbf{Pr}[X_{ti} = A | Z_{ti}, P_{t-1}] = \mathbf{Pr}[P_{t-1} = X_{ti} = A] \\ + (1-e)\,\mathbf{Pr}[P_{t-1} = A \neq Z_{ti}] \\ + (e)\,\mathbf{Pr}[Z_{ti} = A \neq P_{t-1}]\end{aligned} \quad (9)$$

We assume that all of the market participants have the same total purchasing power in this particular market, so that the price takes into account all of the agents' beliefs equally. The result is that we need to set the price as the mean probability distribution over the agents' beliefs, and so to achieve this we use the majority operator:

$$\mathbf{Pr}[P_t = A] = \text{majority}(\{X_{ti} = A\}_{i \leq n}) \quad (10)$$

The complete Bayesian Network model is shown in Figure 4.

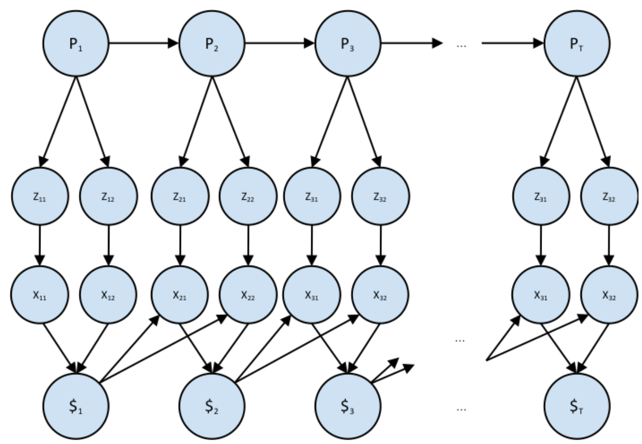

Figure 4. Our complete dynamic Bayesian Network model. In this figure, $n = 2$ and there are 4 timesteps shown. Each X node is an agent, each Z node is their corresponding noisy signal at that timestep, and each $\$_t$ is the predicted price at time t.

5.4. Using the Model

This model can be used for several tasks relating to prediction markets. The uses fall into two main categories: learning on the Bayesian network, and inference using the resulting model. In general, learning is used when political climate data is available, and inference is used to predict missing or future data.

5.4.1 Learning

While our model is mostly hand-designed to match the real-world dynamics of prediction markets, there are several elements which are best learned or determined by data for the given market. In particular, the parameters w and e are best determined by matching a sampling of the data from the actual market (or a similar market, if no such data is available) so that we have the most accurate model. As noted in [20], this is particularly crucial for how the agents are interpreting information because it is probable that different prediction markets have different such parameters. For example, an obscure House of Representatives race probably has many fewer polls than a presidential general election, and so agents may need to update their beliefs differently at each time a poll is released. The flexible parameters in our model allow us to model such differences with the same Bayesian network.

5.4.2 Inference

Inference is the crucial task that our model performs. Like any Bayesian Network, our model is highly effective in determining the probability distributions of unknown nodes given data at other nodes. Two useful inferences that we can make using our model are

$$\mathbf{Pr}[\$_i = A | P_1, P_2, P_3, \ldots P_k] \text{ where } k < i, \quad (11)$$

predicting the price of the market a certain number of periods after the latest polling data, and

$$\mathbf{Pr}[P_i = A | \$_1, \ldots, \$_k] \text{ where } k < i, \quad (12)$$

predicting the political climate at a certain time given the prediction market prices.

Since all of our market prices and polls are interepreted as probability distributions, we are performing inference given a *distribution* over the known nodes, rather than the actual resolution of the node as is more common. This does not change the algorithms for inference over Bayesian Networks except for trivial changes to take distributions as input.

To perform the inference, we used a Python library called `pomegranate`. This library uses Gibbs sampling to perform inference on a given node given distributions on the other nodes. Since it produces concrete results which can be measured against the real world, we evaluate our model mainly based on its ability to perform inference accurately. The results from these evaluations are presented in following sections.

6. Data

Polling data were scraped from `RealClearPolitics.com` [21] for the period between January 2011 to November 2012. The data was sparse for the early timeframes and for dates without polls, forward filling was used to upsample and match the price data. The frequency of the polls is shown in Figure 5. We have preprocessed the data to only include the vote shares of Barack Obama and Mitt Romney, so that third-party candidates do not affect the probabilty estimates.

It is evident that the polling data is very noisy, which is the nature of polling data in general. Different pollsters have different biases and different ways of normalizing their results to demographics and population [22]. Since the polls are obviously a noisy signal of the political climate in the first place, it can be expected that our model will reproduce this noise somewhat, and indeed this is true as shown in the Model Validation section.

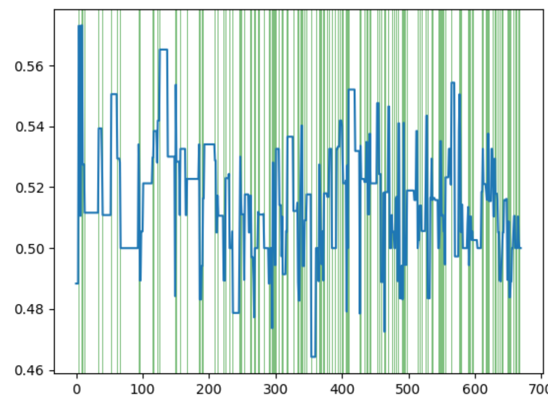

Figure 5. Poll data sparsity. Green vertical lines represent days when poll data were available.

We make the simplifying assumption that all of the polls have the same sample size, and we make this assumption reasonable by only considering polls with sample size of 800 or more as useful in our model. We assume that we can interpret the polls as independent samples from the overall voter population so that we can infer a probability of one candidate having more votes directly from the polling data.

The closing prices of each day in the prediction market were obtained from `PredictIt.com` [9] for the period between January 2011 to November 2012. These are shown, in Figure 5, along with a visualization of the sparsity of the data. We interpret these directly as estimated probabilities of each candidate winning.

7. Experiments

In this section we first present quantitative and qualitative methods of model evaluation. We show that we can validate our Bayesian Network model by inferring prediction market data from general political sentiment (polls), and show both qualitative and quantitative effectiveness based on the evaluation methods above. We also perform hyperparameter tuning using these metrics. Finally, we use our tuned model to make predictions without any prior knowledge of poll data.

7.1. Model Validation

Our model performs moderately well when compared to our two baselines: the Bayesian Updating model in [18] and the raw use of polls as probability estimates as in Figure 7

The results of the quantitative evaluation of each of these models is shown in Table 1. Our model fares better than the two baselines, and in Figure 6 it is easy to see that our model follows the price curve well, albeit with significant noise induced from the noise in the polling data. As expected, the Bayesian updating model from [18] does not account for the changes in political climate over time, and that model is essentially flat at about 51.5% for most of the timesteps, and the direct polling model varies wildly with the noise in the polling data.

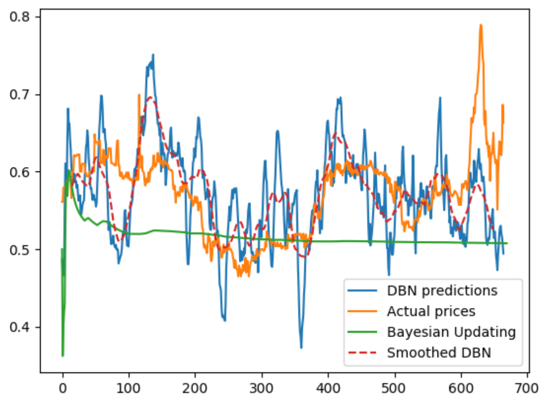

Figure 6. The predictions of our model versus predictions of the Bayesian updating baseline. It is clear that our model handles the passage of time much more gracefully than the single-state Bayesian updating model.

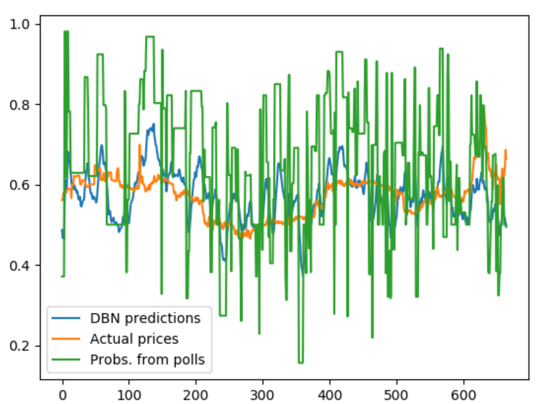

Figure 7. The predictions of our model compared to the baseline that uses only polling data for each time step. Clearly, the result from our model has much less noise.

Table 1. Total KL divergence loss for our model versus each of the baselines.

Polls-only Prediction	54.2625744602
Bayesian Updating	8.24015852064
Dynamic Bayesian Network	6.0164734637

7.2. Parameter Tuning

From tuning the parameters w and e on this particular prediction market's data, it was easy to settle on an optimal value for both, since the curves for model performance relative to each of these values has an obvious optimum in both cases. Here, we are using the KL divergence loss defined previously as a proxy to the maximum-likelihood estimate for these parameters.

The particular values were $w = 0.08$ (see Figure 9) and $e = 0.58$ (see Figure 8). These indicate that there is in fact a fair amount of bias and noise among the agents in the real-world prediction market, and additionally suggests that agents weigh the past prices about equally with the current updated information in the market. These parameters gave us our most accurate model, with which we were able to evaluate market reactions to various adverse conditions. These experiments are detailed in the next subsection.

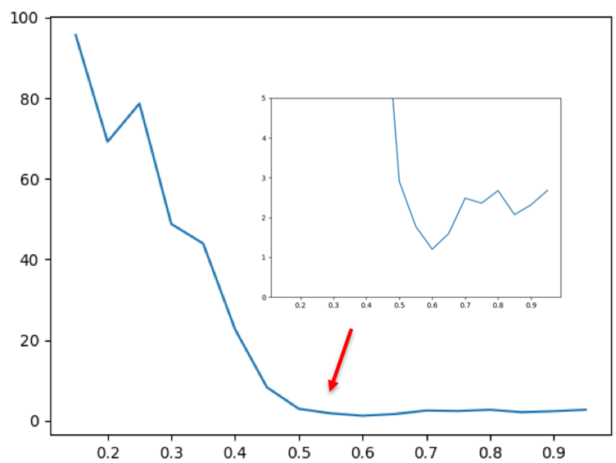

Figure 8. Values for the parameter e and how they affect the KL-divergence loss of the model. The loss is quite asymmetric on each side of the optimum, but the optimum is still clear.

7.3. Prediction

We also tested our model on a prediction task. Given the prediction market prices over a certain length of time, we used the dynamic Bayesian network to infer the probability that Obama would win the election n days out from the last known price of the market. With some noise, the model predicts what we would expect: if Obama has somewhat high odds on the prediction market a number of days in a row, this is a very good indicator that the underlying political climate parameter will have Obama winning in the next few timesteps. However, after a longer time, there is more opportunity for that parameter to move a bit in either direction, resulting in more uncertainty. This is reflected in the model, as seen in Figure 10. In real life, this corresponds to the fact that the further out the election, the more time

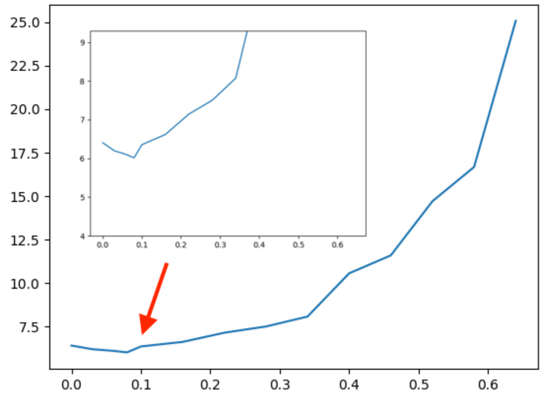

Figure 9. Values for the parameter w and how they affect the KL-divergence loss of the model.

there is for major scandals and other intervening factors to corrupt the indicators that are driving the predictions at the current time.

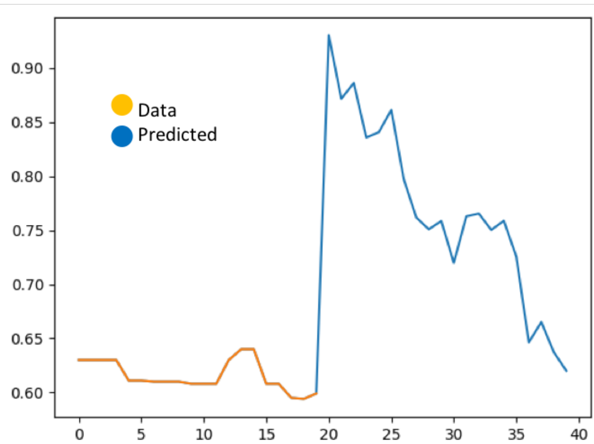

Figure 10. Predicted probabilities of an Obama win n days after the last known price of the prediction market.

7.4. Qualitative Performance

Qualitatively, the most prominent strength of our model was that it preserved the signal of turns in the prediction market very easily based on very noisy polling data. For example, in Figure 6, we see the model matching a spike around day 100, and even anticipating a fall in Obama's odds around days 440-500. This means that the model is effectively processing information contained in the noisy polling data, modeling how the market participants are using that information, and then incorporating that into a prediction.

The model shows some weakness closer to the election, with a much less exaggerated expression of the political climate when compared to the actual prices. This is a result from the fact that our model does not take into account the number of days until the election. In order to use the model to take that into effect, one would have to use the prediction task outlined in Section 7.3.

8. Error analysis

In this section we seek to test the robustness of our model by finding ways to "break" its inference ability. In particular, we first focus on agent information: Do we need to model how informed the agents are? Is there a threshold at which the model becomes uninformative if the number of uninformed agents becomes too high? We next test model robustness to agent noise by raising the noise parameter for each of the agents. Can the model still have inference power if each agent receives a very noisy representation of the world?

8.1. Modeling Agent Information

One of the assumptions that our model makes is that all n agents receive some noisy signal of the data Z_{ti}. We can examine the effects of this property by defining an uninformed agent which models its prediction with:

$$\mathbf{Pr}[X_{ti} = A | P_{t-1}] = \mathbf{Pr}[P_{t-1} = X_{ti} = A] \tag{13}$$

and informed agents whose prediction is described in (Equation 9). We implement this new model, and vary the number of uninformed agents while keeping the total number of agents constant at 5. The results are below:

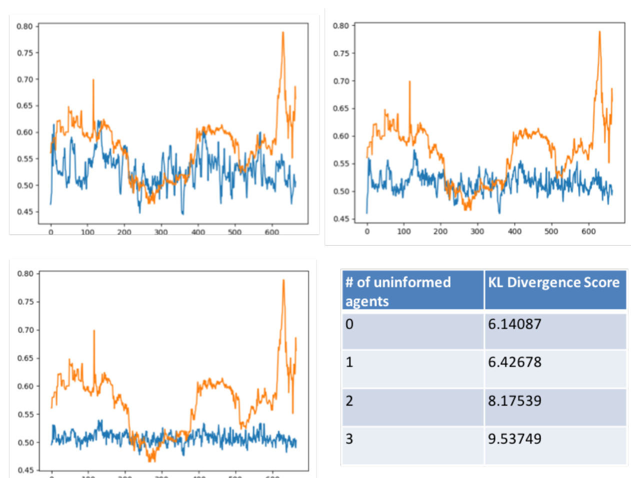

# of uninformed agents	KL Divergence Score
0	6.14087
1	6.42678
2	8.17539
3	9.53749

Figure 11. The predictions of our model with 1 (top left), 2 (top right), and 3 (bottom left) uninformed agents out of 5. The KL scores are displayed (bottom right).

As the number of uninformed agents increases, the KL score increases quite rapidly, and our model predictions become quite uninformative of the market price. When 3 out of 5 agents are uninformed, the uninformed agents' predictions will outvote the predictions of the informed agents in the majority operator function (Equation 10), and the predictions hover around 50%, meaning the model is not incorporating any information from the world.

This experiment reveals the sensitivity of our model to agent informed-ness, which has implications for real-world implementation. For example, in the real world there often are "expert" agents who have privileged information in niche spaces, as well as traders who are more reactionary and have less information about the world. Careful parameter tuning and weighting of agent votes may lead to a much more powerful model, which is a further area of exploration.

8.2. Modeling Agent Noise

Another potential area of error is how the amount of noisy information that each agent receives, represented by N_{ti} affects the output prediction. We can easily perform error analysis experiments by varying the e in Equation 9. In Figure 12 we increase e from 0.08 to 0.4 and see that though the predictions become noisier, the model is still able to capture the general trends of the data, and therefore can still be effective given simple smoothing techniques. This suggests that our model is robust to agent noise, and refers to the de-noising properties of prediction markets in aggregating public opinion.

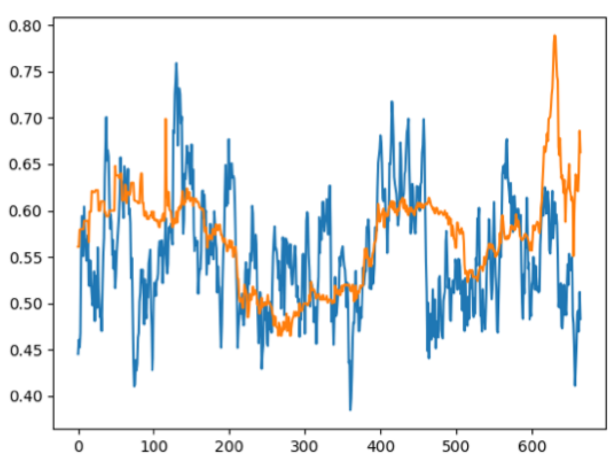

Figure 12. The predictions of our model when agents receive high-noise signals.

9. Conclusion and Future Work

In this paper, we first propose a Dynamic Bayesian Network to model a prediction market, and validate this model with data from the 2012 presidential polls and market prices. We show a significant increase in effectiveness measured in KL divergence over the baseline model, improving in both trend identification and noise smoothing ability. We show that this model has predictive power, even without polling data, suggesting that the latent information about political climate can be inferred from market prices. We finally analyze model robustness to potentially erroneous assumptions about how informed agents are, and the noisy information that agents are getting, and show that too many uninformed agents may "break" the market's predictive ability.

As mentioned in sections above, future work may include a deeper exploration of hyperparameters involved in agent knowledge: How do we distinguish, and extract information from "expert" agents with niche, informative priors? Another important step to validate our model would be to test its generalizability to other prediction markets: Can we use this same model to predict senate races, or sports games?

Overall, we identify prediction markets as a high-potential economic mechanism to extract information about the world. This mechanism allows us to go against the canonical, data-focused machine learning paradigm of data extraction, featurization, and learning. By utilizing markets to incentivize agent participation, we can learn about the world by beginning with the aggregated "wisdom of the crowd", and infer the information that each agent contributes autonomously. The autonomous nature of the market allows rapid and distributed applications to real-world problems with large feature sets which require high domain expertise to extract relevant data - by using prediction markets, we leave the data crunching to the people.

10. Contributions and Acknowledgements

We would like to thank Andrea Shulman for providing guidance.

References

[1] R. J. Lundholm, "What affects the efficiency of a market? some answers from the laboratory," *The Accounting Review*, vol. 66, no. 3, pp. 486–515, 1991. [Online]. Available: http://www.jstor.org/stable/247806

[2] R. Roll, "Orange juice and weather," *The American Economic Review*, vol. 74, no. 5, pp. 861–880, 1984. [Online]. Available: http://www.jstor.org/stable/549

[3] N. Oceanic and A. Administration, "Fy2018 noaa budget summary," 2018. [Online]. Available: https://www.corporateservices.noaa.gov/nbo/fy18_bluebook/FY18-BlueBook.pdf

[4] A. Ozimek, "The regulation and value of prediction markets," 03 2014, working paper at the Mercatus Center, George Mason University. [Online]. Available: https://www.mercatus.org/system/files/Ozimek_PredictionMarkets_v1.pdf

[5] T. Economist, "Daily chart: Brett kavanaugh may have fared better with senators than voters," 09 2018. [Online]. Available:

https://www.economist.com/graphic-detail/2018/09/28/brett-kavanaugh-may-have-fared-better-with-senators-than-voters

[6] T. Gruca and J. Berg, "Public information bias and prediction market accuracy," *Journal of Prediction Markets*, vol. 1, pp. 219–231, 01 2007.

[7] C. F. Horn and B. S. Ivens, *Corporate Prediction Markets for Innovation Management*. Cham: Springer International Publishing, 2015, pp. 11–23. [Online]. Available: https://doi.org/10.1007/978-3-319-14523-5_2

[8] U. o. I. Henry B. Tippie College of Business, "Iowa electronic markets." [Online]. Available: https://iemweb.biz.uiowa.edu/

[9] PredictIt, "Predictit." [Online]. Available: https://www.predictit.org/

[10] U. S. C. F. T. C. Andrea M. Corcoran, "No-action letter 93-66: To the university of iowa regarding the iowa electronic markets," 06 1993. [Online]. Available: https://www.cftc.gov/sites/default/files/idc/groups/public/@lrlettergeneral/documents/letter/93-66.pdf

[11] U. S. C. F. T. C. Vincent McGonagle, "No-action letter 14-130: Re: Victoria university of wellington's request for no-action letter regarding the operation of a small-scale, not-for-profit market for the trading of event contracts for educational purposes," 10 2014. [Online]. Available: https://www.cftc.gov/sites/default/files/idc/groups/public/@lrlettergeneral/documents/letter/14-130.pdf

[12] T. Roughgarden, "Cs269i: Incentives in computer science lecture #18: Prediction markets," 11 2016. [Online]. Available: http://theory.stanford.edu/~tim/f16/l/l18.pdf

[13] E. F. Fama, "Efficient capital markets: A review of theory and empirical work," *The Journal of Finance*, vol. 25, no. 2, pp. 383–417, 1970. [Online]. Available: http://www.jstor.org/stable/2325486

[14] X. F. Zhang, "Information uncertainty and stock returns," *The Journal of Finance*, vol. 61, no. 1, pp. 105–137, 2006. [Online]. Available: https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1540-6261.2006.00831.x

[15] D. Tse, "Lecture notes: Ee 376a information theory," Winter 2017.

[16] F. Galton, "Vox populi," *Nature*, vol. 75, no. 1949, pp. 450–451, 03 1907. [Online]. Available: https://doi.org/10.1038/075450a0

[17] W. W. A. P. David Shrier, Dhaval Adjodah, "Prediction markets," 2016. [Online]. Available: https://www.getsmarter.com/career-advice/wp-content/uploads/2017/07/mit_prediction_markets_report.pdf

[18] E. M. David S. Lee, "Bayesian learning and the pricing of new information: Evidence from prediction markets," *American Economic Review: Papers & Proceedings*, vol. 99, no. 2, p. 330336, 2009. [Online]. Available: http://www.aeaweb.org/articles.php?doi=10.1257/aer.99.2.330

[19] K. Daniel, D. Hirshleifer, and A. Subrahmanyam, "Investor psychology and security market under- and overreactions," *The Journal of Finance*, vol. 53, no. 6, pp. 1839–1885. [Online]. Available: https://onlinelibrary.wiley.com/doi/abs/10.1111/0022-1082.00077

[20] D. Heckerman, D. Geiger, and D. M. Chickering, "Learning bayesian networks: The combination of knowledge and statistical data," *Machine Learning*, vol. 20, no. 3, pp. 197–243, Sep 1995. [Online]. Available: https://doi.org/10.1023/A:1022623210503

[21] RealClearPolitics, "General election: Romney vs. obama polling data," 11 2012. [Online]. Available: https://www.realclearpolitics.com/epolls/2012/president/us/general_election_romney_vs_obama-1171.html#polls

[22] R. Jowell, B. Hedges, P. Lynn, G. Farrant, and A. Heath, "Review: The 1992 british election: The failure of the polls," *The Public Opinion Quarterly*, vol. 57, no. 2, pp. 238–263, 1993. [Online]. Available: http://www.jstor.org/stable/2749470