ASSESSING INFORMATION QUALITY USING PREDICTION MARKETS:

(Completed Paper – IQ Assessment)

Elizabeth Pierce
University of Arkansas at Little Rock, USA
mailto:expierce@ualr.edu

Londraies Thomas
University of Arkansas at Little Rock, USA
mailto:ldthomas1@ualr.edu

Abstract: Measuring the quality of one's information can be a major challenge. Previous studies have used a variety of techniques including data tracking, expert opinion, data profiling, surveys, customer complaints, and comparison of data values to their real-world counterparts to evaluate information quality. This study introduces a new technique to assess the quality of data. Prediction markets are speculative markets similar to stock exchanges except in prediction markets individuals place a value upon the outcome of some future event rather than on some company or commodity. In a prediction market, the market prices are interpreted as probability estimates of the outcomes' happening. This study will demonstrate how prediction markets can be applied to the assessment of information quality using newspapers corrections as an example.

Key Words: Information Quality Assessment, Prediction Markets, Newspaper Corrections

INTRODUCTION

Today's economy runs on data, pervading people's lives privately, professionally, and publicly. Data determines the decisions and actions that people make as either consumers or as producers of products and services. Prior to the advent of the information age, a consumer's trust in an organization was established through personal interactions. Nowadays many of these interactions are based on the exchange of data and consumers will distrust organizations with bad data. Unfortunately in most organizations the quality of data is low [7]. The consequences of relying on poor data quality can range from trivial annoyances to insidious catastrophes. For example, a customer could be billed twice for an invoice that has been paid or a patient's drug prescription could be fatally switched with another.

Situations like the examples mentioned above are why an understanding of data quality is crucial to the survival of an organization. Such understanding; however, is dependent on one's ability to evaluate the quality of data and this have proven challenging in practice. Redman [8] in his analysis of measurement systems for data accuracy notes that there are six basic types of measurement devices.

- *Data tracking* – This method is based on the sampling of data as they move through the information process chain, tracking changes at each stage of the process, and applying business rules to identify nonconformities.

- *Inspection by experts* – This method uses people familiar with the data to spot quality issues.
- *Comparison of data values to their real-world counterparts* – This method involves taking data and tracking down the real-world entities associated with these values to verify their accuracy.
- *Comparison of data values to their domains of allowed values/business rules* – This method involves data profiling, that is, the use of software or other tools to count data values that do not conform to allowed ranges or other established business rules.
- *Customer complaints* – This method relies on using customer feedback to track the level of complaints concerning data problems.
- *Customer surveys* – This method using customers' opinions to gauge the quality of the data.

All of these techniques have their advantages and disadvantages. Data tracking is a very useful aid in discovering the root causes of data error, but it requires resources and some expertise to successfully implement. Inspection by experts often provides reasonably accurate error counts; however, identifying experts can be hard, there is the issue of incentives, and combining opinions from multiple experts can be difficult. Comparison of data values to their real-world counterparts often provides a superior assessment of data quality, but this approach is generally time-consuming, complicated, and expensive. Comparison of data values to their domains of allowed ranges or compliance with business rules is usually straightforward in terms of providing a profile of the current state of data quality; however, it may be tricky to incorporate recent new information. Relying on customer complaints provides insight as to the end-user view of data quality, but this method does not provide feedback in a timely enough fashion to prevent data problems from reaching the consumer in the first place. Finally while customer surveys often provide good attitudinal measures regarding data quality, one must deal with sampling issues, whether or not survey participants have the appropriate incentives to be truthful, the timeliness of survey responses, and how best to weight the various responses one receives into an overall appraisal of data quality.

Although these six devices for gauging data quality work well for a variety of situations, one may argue that the more the merrier. Thus this study seeks to add to the research regarding data quality measurement by investigating the use of prediction markets as a seventh tool for estimating information quality. This investigation explains how prediction markets work and then demonstrates how prediction markets can be applied to the assessment of information quality using newspapers corrections as an example.

BACKGROUND

Prediction markets are speculative markets similar to financial exchanges except in prediction markets individuals place a value upon some outcome of a future event rather than on some company or commodity. A prediction market begins when someone turns an uncertain event of interest into a random variable. For example, suppose one wanted to know "Who will be the Democratic Presidential Nominee in 2008?" The next step in the creation of a prediction market is to designate two or more financial contracts (also known as shares) that represent the various outcomes associated with that random variable. In this example, one could present to the market three possible contracts: (A) Barack Obama as nominee, (B) Hilary Clinton as nominee, and (C) Someone else as nominee. The next step is to open the prediction market to participants interesting in trading on these "commodities". Based on the speculation activity of the traders, the prediction market averages the trading price of each contract. The average trading price of a contract is interpreted as the market's judgment as to the probability of that outcome occurring.

To illustrate imagine an individual who purchases a single contract "Barack Obama will be the

Democratic Presidential Nominee in 2008" for a price of $45 per share. The prediction market reads this as that individual believing this contract as having at least a 45% probability of becoming true. If this contract does become true, then the contract would be worth $100 per share because this outcome has now occurred with 100% certainty. Thus this trader has made a $55 profit from that purchased contract. If the contract does not become true then the contract price drops to $0 per share resulting in a $45 loss for that trader.

Prediction markets are not new. Although most commonly known as prediction markets, this technique goes by many names (information markets, virtual stock markets, decision markets, betting markets, and contingent claim markets) and has been in use even during the election of George Washington [9]. In fact between 1868 and 1940 prediction markets for betting on presidential and state elections commonly operated in the United States and were considered accurate prediction tools [9]. According to Wikipedia [6], a number of organizations currently utilize prediction markets; a partial list is included below.

Examples of Real Money Prediction Markets

- Iowa Electronic Markets (http://www.biz.uiowa.edu/iem/) used for prediction in small scale election markets.
- TradeSports (http://www.tradesports.com) used to trade in political futures, financial contracts, current events, sports, and entertainment.
- InTrade (http://www.intrade.com) used to trade in political, current, financial, weather and unique events.

Examples of Play Money Prediction Markets

- Hollywood Stock Exchange (http://www.hsx.com/) used to predict the success of movies, movie stars, winners of awards. Its data is used for market research purposes.
- NewsFutures (http://www.newsfutures.com) used to make predictions on political, finance, current events and sports market.
- Foresight Exchange (http://www.ideosphere.com/) used to predict political, finance, current events, science, and technology events.

Examples of Internal Prediction Markets

- Google uses an internal market to predict project-completion and project-launch dates.
- Microsoft is piloting prediction markets internally.
- Hewlett Packard uses prediction markets in several business units for sales forecasting.

The successful application of prediction markets for many types of events has led scholars to investigate the science behind how prediction markets work. Researchers such as Polgreen [5] and Wolfers and Zitewitz [12] have identified several reasons why prediction markets often produce useful predictions:

- Prediction market aggregate information from all participants, each of whom has different information about the issue in question.
- Prediction markets provide incentives either in the form of real money, gifts, or psychological rewards to encourage knowledgeable participants to reveal true information in their trades. In addition, prediction markets provide anonymity to its traders. Thus participants can signal market information privately that they might not be willing to do so publicly.
- Prediction markets provide real time feedback to participants through market prices which represent the beliefs of other traders. Participants are often motivated to collect more information in order to "beat the market."

It should be noted that prediction markets are not infallible. While Wolfers and Zitzewitz [13] conclude that prediction market prices normally do a good job of aggregating beliefs about a given outcome, they note that useful forecasts result only if the traders are well-informed. In addition, the efficacy of these forecasts may be undermined for market prices close to 0% or 100% certainty, when the distribution of beliefs is especially diverse, or when trading volumes are somehow constrained or motivated by an unusual degree of risk-acceptance [13].

Another factor that can influence the quality of the forecasts is the design of the prediction market itself. For example, the contracts associated with an uncertain event of interest can be phrased in several ways ranging from a winner-take-all approach (e.g. Clinton wins the popular vote), an index (contract pays $1 for every percentage point of the popular vote won by Clinton), or a spread (contract pays even money if Clinton wins more than y% of the popular vote). In addition, considerations like the type of reward (play vs. real money), the reward structure (e.g. payment based on the size of an individual's market portfolio, rewards for first, second, or third place), the algorithm used to determine market price, the length of time between distribution of awards, determining the minimal group size necessary to generate accurate predictions, and behavioral biases on the part of traders can all impact the ability of the prediction market to generate accurate forecasts.

RATIONAL AND PURPOSE

Prediction markets typically ask questions about some upcoming event that is tangible in nature such as "Will a category 4 hurricane hit Florida in 2007?" Questions like these work well in a prediction market for several reasons. First, the outcome will be known with certainty at some point in time. Second, people participating in this prediction market are likely to have some understanding of this topic so they can make a reasonable value assessment for this prediction. In addition people participating in the prediction market for this type of question are unlikely to have advanced knowledge of the outcome (i.e. insider trading) or be in a position to manipulate the final outcome. Finally the prediction should be something that is of interest to the group sponsoring the question as well as representing some prediction that cannot readily be ascertained by the sponsor using other forecasting methods.

In order to use prediction markets to assess data quality, questions would need to focus on some phenomenon resulting from the use of poor quality data such as making a poor decision based on misinformation. For instance suppose one wished to develop a prediction market to assess the quality of data used by a call center. One could ask "How many times will call center employees contact a client and find the telephone number on file is incorrect during the month of November"? The contracts associated with this event could express various ranges of outcomes such as "less than 10 times", "between 10 and 20 times", "between 20 and 30 times", and "more than 30 times". A market like this would work well for several reasons: (1) It is concrete and the final outcome can be verified, (2) It is tied to a data quality issue, and (3) the call center can ask this question in advance of the month in question so as to get a leading indicator of data quality. In addition the people participating in the prediction market could be restricted to those familiar with the workings of the call center since they would know best the nature of the data thus increasing the accuracy of the prediction.

Suppose at the close of this particular prediction market, the average value associated with the contract "more than 30 wrong client calls" is $75, that is, call center employees on average feel this contract as having a 75% chance of becoming true. This suggests that the quality of the call center telephone database may need to be improved and the organization can act accordingly. The call center can also make use of this type of prediction market as part of its on-going efforts for data quality improvement by running this market on a monthly basis. Similar to a financial market, the call center can give each employee participating in the prediction market an initial pool of "data quality dollars" to invest. Employees could buy and trade contract shares on a monthly basis where the share price represents the probability associated with that contract coming true and the number of shares purchased reflect the confidence in that probability assessment. Employees who predict correctly will be rewarded by earning additional "data quality dollars". For this example, if indeed more than 30 wrong telephone numbers are dialed as a result of poor customer data during the month in question then employees who bought this contract would receive $100 data quality dollars, resulting on average in a $25 profit per purchased share. As an extra incentive to employees who do the best job trading, the call center could give small prizes to those employees who have earned the most "data quality dollars" at the end of the year.

By tracking the value of data quality contracts traded in the prediction market, the call center can also determine which employees in the call center have the best insights into the quality of the data being used by the call center. In addition, trends in the prediction market values such as rising market values associated with poor data quality predictions can be an early indicator that the call center needs to make improvements in its data. As time goes on and the call center continues to build upon its prediction market, the call center can ascertain how effective the prediction market is at evaluating the true quality of its customer database.

It should be noted that measuring data quality using prediction markets involves several obstacles. Under any circumstances, judging data quality is difficult because data quality is multidimensional in nature. Whether using intuitive techniques, a system definition approach [10], or an empirical study [11], data quality can be defined in terms of a wide variety of characteristics such as its accuracy, timeliness, availability, or consistency. Second, even after one has identified the quality dimensions one wishes to estimate about their data, one still faces the challenge of framing a prediction market question that will address those data quality dimensions in a meaningful way. Third, the values generated by the prediction market are not an absolute indicator of data quality. Instead these values represent a quantified, subjective perception of data quality based on the needs and experiences of those familiar with the data.

In many respects prediction markets represent a back-door approach to information quality assessment. Rather than directly asking respondents to rate the various quality aspects of the data with which respondents work, prediction markets ask respondents to estimate the chance of some data quality related event occurring whereby the appraisal process is facilitated by a market environment that promotes information aggregation. Prediction markets are not a replacement for other information quality assessment techniques, but they do represent an additional resource to consider. For example a survey of data consumers may indicate that 90% are satisfied with the security of their data. However if an internal company prediction market indicates that the contract entitled "A major data loss will occur in the next year" is selling for 90 dollars a share then this is a signal for the company to reexamine its data protection policies. In order to further explore the issues associated with using prediction markets to evaluate data quality; this paper next presents a pilot study that attempts to measure the quality of an information source (e.g. the New York Times) using a prediction market based on the number of errors detected and corrected.

METHODS

One data quality question that is often asked is "How good is this information source?" A metric that one might use to answer this question is to study the number of errors detected. The more errors detected the less confidence one might have in the quality of the information provided by that source. One information source that regularly publishes its error corrections is newspapers. The website:

http://www.regrettheerror.com which reports on corrections, retractions, clarifications, and other accuracy trends in the media is a useful place for finding newspapers that publish their corrections.

Wikipedia [3] defines a correction in a newspaper as typically the posting of the notice of a typographical error or mistake that appeared in a past issue of a newspaper. Usually a correction notice appears in its own column. Newspapers generally have specific policies for readers to report factual errors. Often it involves the reader contacting an editor (either by phone, letter, email, or in-person visit), pointing out the mistake and providing the correct information. Sometimes, an editor or affected reporter will be asked to refer to a note or press release to determine how the mistake was made.

A correction differs from a clarification, which clears up a statement that — while factually correct — may result in a misunderstanding or an unfair assumption. Most corrections are the result of reporting errors, although sometimes the newspaper was provided incorrect information. Most newspaper errors are relatively minor and involve one of the following:

- Names — A name was misspelled, someone was misidentified (e.g., in a photograph), a professional title was incorrect, etc.
- Figures – Usually, the result of a typographical error, although it can adversely affect a story (e.g., "the lawsuit was for $8 million, not $8 billion").
- Time/date/place – Usually regarding an event (e.g., "the event will be on Friday in the auditorium, not Saturday in the atrium").
- Other information – Other corrections may involve prices, URL, telephone numbers, misquotes, sports scores, or even published lottery numbers.

It should be noted that not every mistake gets corrected. Only those that the newspaper considers significant enough to warrant correction or clarification get published.

Corrections at the New York Times

One newspaper that has a long history of publishing its corrections, retractions, and clarifications is the New York Times. According to Wikipedia [4], the New York Times is a daily newspaper published in New York City by Arthur Ochs Sulzberger Jr. and distributed internationally. It is owned by The New York Times Company, which publishes 15 other newspapers, including the International Herald Tribune and The Boston Globe. It is the largest metropolitan newspaper in the United States. Nicknamed the "Gray Lady" for its staid appearance and style, it is often regarded as a national newspaper of record, meaning that it is frequently relied upon as the official and authoritative reference for modern events. Founded in 1851, the newspaper has won 95 Pulitzer Prizes as of 2007, more than any other newspaper.

Each day the New York Times prints (both online and in its printed editions) a list of corrections, clarifications, and retractions that have recently been discovered. An examination of corrections and clarifications posted between April 1, 2007 and July 31, 2007 reveals that typically some 0 to 21 corrected or clarified items were listed each day with a mean of 12 items per day and a standard deviation of 4 items per day (See Figure 1). About 99% of the erroneous news items were corrected on average within 6 days of their original publication date with the standard deviation for the correction time also being about six days. It should be noted that about 1% of the corrections were not posted until 6 weeks or more after the source article was published (See Figure 2). In a few instances corrections were not discovered and published until months or even years after the fact. Most corrections were linked to only one article but sometimes an article could contain two, three, or more significant errors requiring correction.

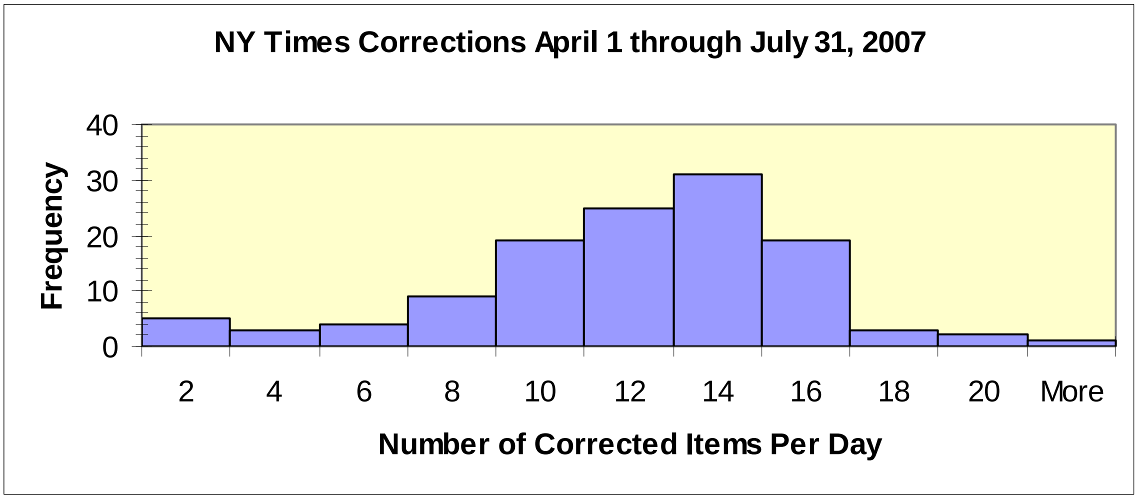

Figure 1: Frequency Distribution of Corrected Items Published Per Day

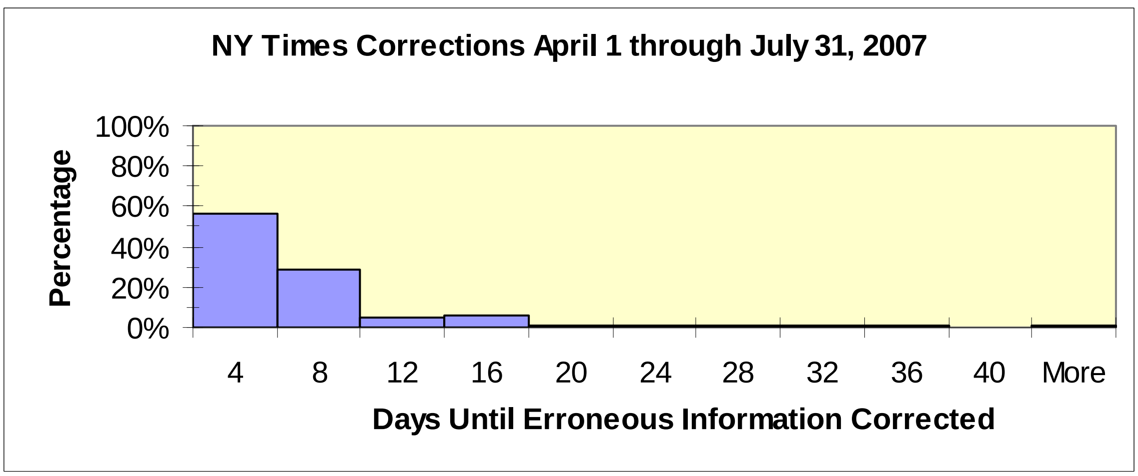

Figure 2: Frequency Distribution of Time Elapse between Date of Corrected Item and Source Article

Based on the analysis of recent corrections posted by the New York Times, this paper proposes the following prediction market questions.

Question 1: How many corrections will be generated concerning material published in the NY Times between July 23 and July 29, 2007? (Market opened on July 21, 2007 and closed on August 20)

- *Choice 1: More than 140 Corrections*
- *Choice 2: Between 106 to 140 Corrections*
- *Choice 3: Between 71 to 105 Corrections*
- *Choice 4: Between 35 to 70 Corrections*
- *Choice 5: Less than 35 Corrections*

Question 2: How many misspelling errors will be reported in the New York Times between July 30, 2007 and August 3rd, 2007? (Market opened on July 27 and closed on August 4.)

- *Choice 1: 0 to 10*
- *Choice 2: 11 to 20*
- *Choice 3: 21 or more*

Question 3: How many omission errors will the New York Times report on August 3rd, 2007? (Market opened on July 27 and closed on August 4.)

- *Choice 1: 0 to 7*
- *Choice 2; 8 to 15*
- *Choice 3: 16 to 22*
- *Choice 4: 23 or more*

All three of these questions deal with the trading public's perception of the quality of the error detection and correction process used by the New York Times. Question 1 is more focused on determining the public's perception of the quality of the information generated by the New York Times over a given period, while Questions 2 and 3 are more focused on the frequency of specific types of errors that are reported over a given time period..

Setting up a Prediction Market

To set up the prediction market, this study used the services of InklingMarkets.com (http://www.inklingmarkets.com). Inkling Markets is a web-based hosting service that allows registered individuals to trade on public prediction markets, to run trial markets for free, and to obtain prediction markets services. Inkling Markets encourages people to learn about prediction markets by making free registration quick and easy (participants are given 5,000 inkles as starter money), providing educational information online, and employing a simple email request system for persons interested in setting up their own prediction markets.

Once Inkling Markets has processed a request from an individual to organize a prediction market, it then designates that individual as the administrator of his or her own prediction marketplace. Inkling Markets allows administrators to configure their marketplaces by setting login information, personal profile information, alerts, homepage content, as well as selecting the marketplace management's options. The marketplace management settings allows administrators to control their marketplace appearance, its security and access (e.g. will this marketplace be open to the public or to only an invited list of participants), and communication links (e.g. Discussion Boards, Contact Us, Submit Bug, About Us)

Once the marketplace configuration is complete, administrators can begin listing questions to be traded upon in the virtual marketplace. Building a prediction market involves these five steps.

1. Design a Question: Inkling Markets allows administrators to create questions based on the following three types of predictions:
    a. The probability of a single event occurring (e.g. Which one of these products will sell the most in January?)
    b. The probability of multiple events occurring (e.g. Which of these products will debut by January of next year?)
    c. A specific dollar amount or number (e.g. What will be the price of this product when it debuts in January of next year?)

    For our questions, we have chosen the probability of a single event occurring.
2. Stock Information: Inkling Markets allows administrators to create the answer categories that participants can trade upon. For example, for question 1, there will be 5 stocks, one for each range of correction results.
3. Details: Administrators can post additional information about the subject to assist participants to trade

more knowledgeably.

4. Security: Inkling Markets allows administrators to restrict access to a particular question to an invited list of participants rather than the public at large. For this case, the market will be open to the public.
5. Publish: This is the final step that permits the prediction market to be open for trading.

Once participants begin trading predictions for a given question, Inkling Markets monitors the trades being made and posts them to a price chart so that both participants and the administrators can track the progress of the market. Similar to a stock exchange, the prediction market provides information about the volume of shares traded and the price of the shares over time. The share itself represents the outcome that a trader believes is most likely to occur for a given event. The share price is an aggregate measure indicating how likely the market as a whole believes an outcome will occur while the number of shares purchased by traders reflects their confidence in those prediction shares.

Tabulating the Results

To tabulate the various types of corrections, the authors examined the list of corrected items published in the New York Times' Corrections/For the Record Section on a daily basis during the months of July and August and counted the number of items corrected whose dates met the time period and correction category as specified by a given question.

There were several complications associated with this tabulation process.

- For this study we counted any correction published for any material produced by the New York Times. This included photographs along with their captions, articles, columns, listings, charts, special reports, etc.)
- For questions regarding the source of the correction, it often took several weeks before a reasonably accurate tally could be made. This is a result of the long time delay between the publication of an article and the publication of a correction associated with that article. Although most corrections and clarifications are detected, verified, and published within a few days of the published material, some corrections may take place months or even years later. As a result, the prediction market associated with Question 1 allowed approximately 4 weeks to elapse before closing this market.
- Sometimes the same article would generate multiple corrections. For example if an article had a name error, date error, and a location error listed among its corrections, then the authors counted this as three corrections generated by the same article.
- The number of corrections was also determined in part by the logical grouping of the data. For example, a telephone number that was published in error with several digits reversed was counted as one correction since most people would view the telephone number as a single entity. In addition, if the captions for two pictures were reversed then this will be counted as one error since both pictures are affected by the same error. On the other hand, if two people within the same article have their names misspelled then this would be counted as two errors since two different entities are involved requiring two different corrections to resolve the problem.
- Sometimes there would be a correction made to a correction. This was counted as an additional error since that too was information published in the New York Times. As a related complication, sometimes the same error was repeated multiple times before it was caught and corrected. Since the error was repeated on multiple, separate occasions, it was counted multiple times.

RESULTS

Throughout the open trading period for each question, the authors monitored the prediction market on a daily basis, responding to inquiries from participants as well as tabulating the corrections listed by the New York Times. Besides the administrator, 17 individuals registered to trade. Of the 17 registered traders, 14 were listed by the system as active traders.

Question 1

For question 1, the authors recorded 100 corrections to information published in the New York Times between July 23, 2007 and July 29, 2007. The initial share price for each prediction contract started at 20% (i.e. $20). The final share prices and trading volumes are listed below. For this question, the market did not predict the correct outcome. Much of the trading activity centered on shares associated with fewer corrections being found. In addition both share price and trading volume were much lower for the outcomes associated with higher numbers of corrections being found. Interestingly although the correct outcome (71 to 105 corrections) had the second highest share price, it did have the heaviest trading volume. Based on these results it seems that the market perceives the error rate of information published in the New York Times to be lower than it actually is.

Question 1: How many corrections will be generated concerning material published in the NY Times between July 23 and July 29, 2007? **(Market opened on July 21, 2007 and closed on August 20)**			
Prediction	**Last Trading Value (Date)**	**Final Value**	**Shares Traded**
More than 140 Corrections	14.4% (July 27, 2007)	0	90
Between 106 to 140 Corrections	12.3% (July 26, 2007)	0	10
Between 71 to 105 Corrections	24.8% (August 6, 2007)	100	335
Between 35 to 70 Corrections	32.9% (August 10, 2007)	0	296
Less than 35 Corrections	15.6% (August 6, 2007)	0	250

Question 2

For question 2, the authors recorded 9 incidences of misspellings that were reported between July 30, 2007 and August 3, 2007 by the New York Times. The initial share price for each prediction contract started at 33.33% (i.e. $33.33). The final share prices and volumes are listed below. For this question, the market did not correctly predict the outcome. In general traders believed that the final tally of reported misspellings would be much higher (21 or more) than was actually recorded (0 to 10). In fact the volume of shares traded indicates that traders did not choose the correct option at all. It is interesting to note that misspellings are a particularly common problem at the New York Times. By its own records, the New York Times has published corrections for 269 misspelled names of people between January 1, 2007 and August 12, 2007 [2]. Since the New York Times only publishes corrections for reported name misspellings, this may indicate a signal from the market that the actual problem of misspellings is perceived to be a much wider problem by the public.

Question 2: How many misspelling errors will be reported in the New York Times between July 30, 2007 and August 3rd, 2007? **(Market opened on July 27 and closed on August 4.)**			
Prediction	**Last Trading Value (Date)**	**Final Value**	**Shares Traded**
0 to 10	29.4% (Not Traded)	100%	0
11 to 20	34.6% (July 31, 2007)	0%	40
21 or more	36.0% (August 1, 2007)	0%	60

Question 3

For question 3, the authors recorded one incident of an omission error that was reported on August 3, 2007. The initial share price for each prediction contract started at 25% (i.e. $25). The final share prices and volumes are listed below. For this question, the market correctly predicted the outcome. The shares associated with the outcome, "0 to 7 omissions" had the highest share value (29.9%) as well as the highest volume. The second most commonly traded share was the next category, "8 to 15 omissions" with a share value of 29.3%. No one traded in shares associated with any higher outcomes. This result suggests that the market believes that the reporting of omission errors by the New York Times to be a relatively rare occurrence.

***Question 3: How many omission errors will the New York Times report on August 3rd, 2007?* (Market opened on July 27 and closed on August 4.)**			
Prediction	**Last Trading Value (Date)**	**Final Value**	**Shares Traded**
0 to 7	29.9% (August 1, 2007)	100%	95
8 to 15	29.3% (August 1, 2007)	0%	90
16 to 22	20.4% (Not Traded)	0%	0
23 or more	20.4% (Not Traded)	0%	0

DISCUSSION

This pilot study investigating the use of prediction markets for estimating information quality reveals several notable features.

- Because prediction markets are based on some future event, the market values associated with a given outcome can be interpreted as a gauge of what traders believe will happen. Because prediction markets are intended to run over time, prediction markets allow one to observe both the current level of belief in the data consumer population regarding some future data quality related event, as well as any changes in beliefs that may be occurring. Thus prediction markets serve as a dynamic, leading indicator of perceived data quality, allowing one to see where market values are trending over time. Other forms of data quality assessments such as surveys and opinion by experts tend to be static in nature, allowing one to capture only a snapshot of beliefs at one particular time.
- Prediction markets are easy to set up and operate thanks to the widespread availability of prediction market software both from vendors as well as open source providers.
- Prediction markets are an indirect way of subjectively evaluating data quality because the types of questions one asks are based on events driven by data quality concerns, rather than on a specific data quality dimension. Thus while in a survey approach, one might ask "Do you trust the data contained in the hospital pharmaceutical database?" a prediction market question might ask "Will at least one hospital patient receive a prescription with the wrong information prescription next month?" Because of its indirect nature, the prediction market value reflects a compilation of factors (some data quality related, some not) that influence both the magnitude and the direction of data consumers' beliefs.

LIMITATIONS

The major limitation of this study is its small pilot size both in terms of the number of participants and the number of questions posed over a relatively short period of time. More investigation is needed using a larger number of participants predicting a greater variety of data quality events to more fully explore the use of this technique as a data quality assessment tool. Despite the limited nature of this study, initial results appear promising and the authors are encouraged to explore further research on this topic.

CONCLUSION

Prediction markets represent an intriguing new tool in data quality assessment. While this technique is not applicable to every evaluation of data quality; in those circumstances where events related to data quality can be articulated, the prediction market may yield more insights than a conventional customer satisfaction survey on data quality. To the authors' knowledge, the application of prediction markets to the estimation of information quality is a novel idea. Additional research is needed to more fully explore how best to use prediction markets for appraising data quality. Future research questions to be investigated include the following:

- What question type and marketplace design will maximize the effectiveness of using the prediction market technique to assess information quality issues? More research is needed to guide practitioners in how best to formulate questions to better capture the various dimensions of data quality. Also of concern is whether prediction markets work well for only certain types of data quality dimensions and if so, which dimensions are most amenable to this approach?
- Does the prediction market technique provide a reasonably accurate judgment of people's perceptions regarding information quality? Further research is needed to validate if the market value determined by the prediction market is an accurate representation of people's true feelings regarding the data quality characteristic in question. Also at issue is how to best interpret and make use of prediction market results for data quality improvement when they can reflect factors other than data quality concerns.
- How efficient is the prediction market technique compared to other methods for evaluating data quality? Currently there exists a variety of ways to assess data quality including data tracking, expert opinion, data profiling, surveys, customer complaints, and comparison of data values to their real world counterparts to measure data quality. How does the amount of effort (i.e. materials, software, hardware, network, people, training, etc.) and results of these techniques compare to the efforts and results associated with prediction markets when it comes to appraising data quality?
- Can prediction markets be used for making data quality comparisons? Suppose one were to develop a series of questions based on judging the quality of several different information sources. Would the market values generated by the prediction markets provide an accurate reflective of the level of trust that consumers place in one information source versus another?

There are no doubt other questions that one might ask about using prediction markets to evaluate data quality. This study is meant to be a starting point and hopefully will encourage others to further investigate the use of prediction markets for information quality assessment.

REFERENCES

[1] Frequently asked questions. (2007). Retrieved June, 15, 2007, from http://inklingmarkets.com/homes/faq.

[2] Hoyt, C. (2007). So Many Names, So Many Corrections. Retrieved August 30, 2007 from http://www.nytimes.com/2007/08/12/opinion/12pubed.html?pagewanted=1&ei=5070&en=b5d6c9907c6849ec &ex=1188619200

[3] Newspaper Corrections (2007). Retrieved July 2, 2007 from http://en.wikipedia.org/wiki/Correction_ (newspaper).

[4] New York Times (2007). Retrieved July 4, 2007 from http://en.wikipedia.org/wiki/NY_Times

[5] Polgreen, P. (2007). Frequently Asked Questions about the Influenza Prediction Markets. Retrieved July 16, 2007 from http://fluprediction.uiowa.edu/fluhome/FAQ.html.

[6] Prediction market. (2007). Retrieved May 20, 2007, from http://en.wikipedia.org/wiki/Prediction_market

[7] Redman, T. (2001). Data quality: The Field Guide (1st ed.). United States of America: Digital press.

[8] Redman, T. (2005). Measuring Data Accuracy, Information Quality – Part of the Advances in MIS Monograph Series. United States of America: M. E. Sharpe, Inc., 21-36.

[9] Rhode, P., & Strumph, K. S. (2003), Historical Prediction Markets: Wagering on Presidential Elections. Retrieved August 27, 2007, from http://www.unc.edu/~cigar/papers/BettingPaper_10Nov2003_long2.pdf.

[10] Wand, Y. and Wang, R. Y. (1996). Anchoring Data Quality Dimensions in Ontological Foundations. Communications of the ACM, 39, 11 (November), 86-95.

[11] Wang, R. Y. and Strong, D. M. (1996). Beyond Accuracy: What Data Quality Means to Data Consumers. Journal of Management Information Systems, 12 (4-Spring), 5-33.

[12] Wolfers, J., & Zitzewitz, E. (2004). Prediction markets. The Journal of Economic Perspectives, 18(2), 107-126.

[13] Wolfers, J., & Zitzewitz, E. (January 8, 2007). Interpreting prediction market prices as probabilities. , 1(1), June 9, 2007.Retrieved June 9, 2007, from http://bpp.wharton.upen.edu/jwolfers/Papers/ InterpretingPredictionMarketPrices.pdf

BIBLIOGRAPHY

The following publications represent additional sources used for this research.

Berg, J., Forsythe, R., Nelson, F., & Rietz, T. (2000). Results from a dozen years of election futures markets research., June 9, 2007 . Retrieved June 9, 2007,

Hanson, R. (2003). Shall we vote on values, but bet on beliefs? Department of Economics, George Mason University.

Pipino, L., Lee, Y., & Wang, R. (2002). Data quality Assessment. Communications of the ACM, 45(4), 211- 218.

Schreiber, E., Wolfers, J., Pennock, D., & Galebach, B. (2004). Prediction markets: Does money matter? Electronic Markets, 14(3), June 11, 2007 . Retrieved June 11, 2007, from http://www.newsfutures.com/pdf/does_money_matter.pdf

Spann, M., & Skiera, B. (2003). Internet-based virtual stock markets for business forecasting. Informs, 49(10), June 7, 2007. Retrieved June 7, 2007, from www.hsxresearch.com/assets/VSM_Mgmt_Sci.pdf

Wang, R. (1998). A Product Perspective on Total Data Quality Management. Communications of the ACM, 41(2), 59-65.

ACKNOWLEDGEMENTS

The authors would like to thank Dr. Herbert Remidez, Assistant Professor of Management at the College of Business at the University of Arkansas at Little Rock for his advice and assistance on prediction markets.

In addition the authors would like to dedicate this paper to Ms. Yvonne Croston, Coordinator of the McNair Program, who died unexpectedly in July 2007 and whose dedication to the Ronald E. McNair Scholars Program helped to make this research collaboration possible.