Social Network Effects on Information Aggregation

by

Chilongo D. Mulanda

Submitted to the Department of Electrical Engineering and Computer Science
in partial fulfillment of the requirements for the degree of

Master of Engineering in Computer Science and Engineering

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

September 2006

Author ....
Department of Electrical Engineering and Computer Science
July 28, 2006

Certified by ....
Alex (Sandy) Pentland
Toshiba Professor of Media Arts and Sciences, MIT
Thesis Supervisor

Accepted by ....
Arthur C. Smith
Chairman, Department Committee on Graduate Students

Social Network Effects on Information Aggregation

by

Chilongo D. Mulanda

Submitted to the Department of Electrical Engineering and Computer Science
on July 28, 2006, in partial fulfillment of the
requirements for the degree of
Master of Engineering in Computer Science and Engineering

Abstract

In this thesis, we investigated how sociometric information can be used to improve different methods of aggregating dispersed information. We specifically compared four different approaches of information aggregation: vanilla opinion poll, opinion polls where sociometric data is inferred from the population's own perception of social connectivity, opinion polls where sociometric data is obtained independent of the populations beliefs and data aggregation using market mechanisms. On comparing the entropy of the error of between the prediction of each of these different methods with the truth, preliminary results suggest that sociometric data does indeed improve the enterprise of information aggregation. The results also raise interesting questions about the relevance and application of different kinds of sociometric data as well as the somewhat surprising efficiency of information market mechanisms.

Thesis Supervisor: Alex (Sandy) Pentland
Title: Toshiba Professor of Media Arts and Sciences, MIT

Acknowledgments

I would like to foremost thank my advisor, Professor Alex (Sandy) Pentland. I was fortunate enough to get the opportunity to work in his group - The Human Dynamics Group - at the MIT Media Lab for more than a year, during which I have learned so much from him and from the other members of the group. He has been most instrumental in guiding me through this thesis, always ready to answer my questions and patiently providing insight on how to push forward whenever I felt stuck.

I am also grateful to the other members of the Human Dynamics Group - Juan-Carlos Barahona, Jonathan Gips, Anmol Madan, Will Stoltzmann, Katherine Yiu, Wen Dong and Daniel Olguin - who made my time at the group quite enjoyable and fruitful. Special thanks to Jonathan Gips and Katherine Yiu, whose help was invaluable in setting up the experiments and analyzing the data without which this thesis would not be complete. Between the two of them I was spared hours of hardware and matlab debugging, for which I am grateful.

Finally, I would like to thank my family and friends for their perennial support and encouragement. I dedicate this thesis to my parents, James P. L. Chilongo and Elizabeth N. Chilongo - your unconditional support and dedication to see me through has been unwavering, and for that I am eternally grateful.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

While many of the phenomena we observe is predictable, the accuracy and reliability of the prediction relies a great deal on the amount of information we have regarding that particular occurrence. However, it is seldom the case that the information regarding the phenomenon of interest is attainable from one source (we usually have to refer to multiple texts on a specific subject, talk to many different experts, etc in order to have confidence in the conclusions we draw on a subject).

When the information is dispersed across many people, there are several ways to go about collecting that information accurately. The best established of these is random sampling, where a representative subset of the universe is interviewed and the responses aggregated by straightforward counting to determine the population's belief on a particular occurrence. However, such sampling and counting fails to account for the fact that many times people give responses that do not reflect the amount of information available to them regarding the occurrence, either because they have some incentive to give misleading answers, or they have no incentive to put some effort towards accurately analyzing the data at their disposal.

This question of incentives has led to the emergence of information markets as a way of aggregating dispersed knowledge. The main idea behind information markets is to have people buy and sell stocks/securities reflecting their perceived beliefs in a particular outcome. Since the possibility of making and losing money is very real in such markets, all players have an incentive to use their information truthfully (albeit selfishly), with the end result that the final state of the market (prices and trade volumes) tend to reflect the underlying phenomenon.

However, consider a situation where we desire to predict an answer to the question "Will South Africa emerge victorious in the 2010 FIFA World Cup in South Africa?". Furthermore, suppose we have three people A, B and C, who are part of our sample population. By traditional sampling and counting, we might get the responses "yes", "yes" and "no" from A, B and C respectively, thereby concluding that it is 66% likely for South African to win. However, suppose persons A and B are the very best of friends - they spend lots of time together, work together, watch football games together at the bar every week, etc - and person C practically has no contact with A or B. Then it is plausible that A's response might just reflect B's beliefs or vice versa, which would suggest that we may want to weight A and B's total response slightly less, so that the final likelihoods might be, for instance a 50% chance for South Africa to win and a 50% to lose. This example serves to illustrate that the social information may play a greater role than revealed by the sample-and-count or the information market mechanisms of aggregating information. For this reason, it is worth explicating the role played by the social network in the accuracy of the final result, and comparing it to the results of the other methods. Towards this end, we set up experiments based on [?] and extended to include sociometric information collected from the users using wearable devices.

1.1 Contributions

The main contribution of this thesis is in attempting to explicate how to factor social information into the enterprise of information aggregation, along with suggestions of how that might be improved as social network analysis tools and models gain more sophistication.

1.2 Thesis Organization

The content of this thesis is grouped into six chapters. In broad terms, we provide background on related work in the areas of information aggregation and social network analysis, after which we describe the aggregation algorithms we employed, the tools we developed and used for the experiments, the experimental setup and finally an analysis of the result.

Chapter 1 Introduction: We motivate the objective of the thesis, state it's contributions and outline the progression of the rest of the thesis.

Chapter 2 Background: We give an overview of the related fields of information aggregation and social network analysis as is pertinent to this thesis.

Chapter 3 Methodology: We give an account of the algorithms and performance metrics we use in our analysis.

Chapter 4 Experimental Setup: We describe the setup of the experiment, along with an overview of all the software and hardware tools we developed and used.

Chapter 5 Results: This chapter discusses the results of our experiments, and how those relate with the expected results. We also discuss any possible biases we may have inadvertently introduced that might explain discrepancies between the expected and the actual results, and how those biases can be corrected in future.

Chapter 6 Applications: We discuss several situations where we expect the methodology we employed to be either superior or inferior to other methods and why we expect the respective result.

Chapter 7 Conclusion: The thesis is concluded with a restatement of the contributions of the thesis, the ramifications of the results of the thesis and the direction that future research in this area might progress.

Chapter 2

Background

2.1 Regular Sample-And-Count

This is the most widely used method of aggregating dispersed information because it is well understood and, once a representative sample of the population has been identified, is easy to implement. The only significant difficulty in using this method (aside from the cost of resources that may be required if the sample size is very large) is the determination of a sampling frame that is representative of the entire population under investigation. The question of how representative a frame is varies depending on the purposes of the study, and once such a frame has been established and a sampling size determined then one of many sampling methods can be: simple random sampling, stratified sampling, cluster sampling etc [?].

Note that we use the term "regular" in "regular sample-and-count" to distinguish this method from the method we employ in this thesis, which also uses a representative population (hence a sample of the population). The main difference is that in the regular sample-and-count, each data point usually contributes equally in the analysis; for instance, when trying to predict an election, the results are aggregated such that each response from the opinion poll is weighted equally (once). The method we use, in contrast, tries to factor in sociometric information by weighting the responses according to sociometric data collected from the population.

Regular sample-and-count suffers a number of shortcomings. First, the quality of the responses might not be as good as the investigator hopes; many people will give inaccurate responses due to a number of reasons:

1. They have no information whatsoever regarding some particular event, but nevertheless feel inclined to give some sort of answer (either because they do not want to appear ignorant, or feel obligated in some way).

2. They have all the information they need to answer some question accurately, but the information requires some thought and analysis which they have no incentive to perform.

3. They only have partial information, and have no incentive to try and make an educated guess based on that partial information.

This first problem boils down to a question of incentives, and as we will see below, information markets are designed specifically to counter this problem.

Second, a simple sample-and-count usually leaves out of the analysis sociometric information that might be of relevance to the enterprise. As pointed out earlier, two identical responses from individuals who interact frequently are probably not as independent as two responses from individuals who almost never interact. This might not matter in some instances, where the people's beliefs (regardless of how they are achieved) represent the ground truth (for instance, an opinion poll for an election); in other instances, the phenomena we're trying to deduce or predict (e.g determining the weight of a cow by asking random people [**?**]) may not be affected by the beliefs of the sample population and therefore correcting for the effects of their social network may improve results significantly.

One way of incorporating sociometric data might be to use it when selecting the sample population. This may work in some cases, but in others it may not be easy to obtain a sample population that meets all social networking criteria relevant to the study; furthermore, there may be unavoidable social contact during the duration of the study. Therefore, it may be prudent to measure sociometric information before and through the duration of the study and bringing it to bear during the analysis stage of the study.

2.2 Information Markets

In the recent past, the idea of setting up information markets (also know as prediction markets) as a means of aggregating dispersed information for predictive or analytic purposes has become popular, and preliminary results show that such information markets can do very well [?, ?]. In brief, the idea of information markets is to create securities whose cash value is anchored to the event(s) being predicted. Like any other securities market, people who buy when prices are low and sell when prices are high make money while those who buy when prices are high and sell when prices are low lose money. As the market proceeds, the prices of the securities is believed to approach the true distribution of beliefs.

As mentioned earlier, one of the main benefits of information markets is that they provide an incentive for the participants to report their beliefs accurately, (which in this case means buying stocks that they believe will perform well and selling those that they believe will not), since their payoff is directly linked to the prices of the stocks they trade. This helps counter one of the problems of the sample-and-count method where the participant have no incentive to accurately report their beliefs.

In addition, the infrastructure supporting the world wide web provides an excellent platform to run these information markets with numerous participants, making it much more economically feasible to execute such markets than it previously was. In fact, there are a number of such prediction markets operating online today (with both real money and play money), for instance the Hollywood Stock Exchange for predicting which movies will do well in theatres (http://www.hsx.com), The Iowa Electronic Markets for predicting economic and political events (http://www.biz.uiowa.edi/iem/), Tradesports for predicting sporting events (http://www.tradesports.com), etc.

However, information markets suffer their own shortcomings:

1. **Liquidity:** Liquidity is the degree to which an asset or security can be bought or sold in the market without affecting the assets price, and is characterized by high volumes of trade. If the market is not sufficiently liquid, then any one trade may adversely affect the market, which means that the prices will fluctuate wildly and

therefore not approach the true distribution of beliefs. This was of particular concern to us in this thesis, since the market we instituted was quite thin; nevertheless, the results of the market were surprisingly close to the true distribution.

2. **Manipulation:** A rogue trader could try and manipulate the market to his/her advantage, with total disregard of what they truly belief. This possibility of manipulation is linked to the issue of liquidity - the more traders (and trade volumes) there are, the more liquid the market is, and therefore the more difficult it is for a single person (or small group of people) to manipulate the market. Consider the following entry in *Wikipedia*:

"...In the Tradesports 2004 presidential markets there was an apparent manipulation effort (an anonymous trader sold short so many Bush 2004 presidential futures contracts that the price was driven to zero, implying a zero percent chance that Bush would win. The only rational purpose of such a trade would be an attempt to manipulate the market in a strategy called a "bear raid". The manipulation effort failed, however, as the price of the contract rebounded rapidly to its previous level.)..."

3. **Theoretical Challenges:** There have also been questions about the ability of such information markets to really perform well. First, Dr. Charles F. Manski [?] attempts to show that under a wide range of assumptions, predictions of these markets do not closely correspond to the actual probability beliefs of the market participants unless the probability is near either 0 or 1. However, Steven Gjerstad [?] has shown that prediction market prices are very close to the mean beliefs of the market participants if the distribution of beliefs is smooth. Justin Wolfers and Eric Zitzewits [?] have obtained similar results.

Second, there has been widespread belief that markets operated using "play money" cannot generate credible predictions, possibly because play money supposedly does not provide "real" incentive. However, currently accumulated data suggests otherwise; Pennock et al [?] analyzed data from the Hollywood Stock Exchange and the Foresight Exchange (both play money markets) and concluded that the market prices predicted outcomes and/or outcome frequencies in the real world.

Third, markets tend to inadvertently measure things not intended to be factored into the final determination made by the market, such as the risk tendencies of the participants. Risk neutral participants tend to undervalue the securities while risk loving individuals tend to over-value them, and these biases show up in the final market prices.

4. **Social Information:** Finally, these markets do not explicitly deal with social network information. There is no doubt, however, that social connections do play a role in influencing the beliefs of the players. We believe that in the course of the market, there is some social interaction between the participants which the market converts to a price that can be easily interpreted by other players (for instance, if from previous history player 1 seems to make lots of money, then if player 1 offers to buy lots of stock A and some other player(s) get to know of this, they may themselves be tempted to buy more of stock A). One of the goals of this thesis is to use a method that explicates the social connections between participants and maybe shed some light on what goes on behind the scenes in these prediction markets with regard of interpretation of social signals.

2.3 Alternative Methods

Other aggregation schemes have been suggested that try to leverage the wisdom of the crowds without the pitfalls of information markets. For instance, [**?**] suggests that directly asking a group of participants to estimate probabilities may lead to better results. As for incentives, the participants can be compensated based on how close they are to the truth, which provides some incentive to be right. In fact, these "suggestions" form the basis of the experimental method developed by Chen et. al. [**?**] that this thesis borrows from.

Chapter 3

New Approaches

3.1 Chen, Fine and Huberman Approach [1]

Kay-Yut Chen, Leslie R. Fine and Bernardo A. Huberman of HP Laboratories developed a novel approach of aggregating information that uses a small number of individuals participating in an imperfect information market [**?**]. What follows is an extended quote from pages 47 - 50 of the Chen et. al. paper describing the methods and analysis done in [**?**] (with different section and equation numbering), since these methods and analysis are essential to this thesis.

3.1.1 Introduction

Information markets generally involve the trading of state-contingent securities. If these markets are large enough and properly designed, they can be more accurate than other techniques for extracting diffuse information, such as surveys and opinions polls. There are problems however, with information markets, as they tend to suffer from information traps (Camerer and Weigelt, 1991; Noth, et al., 1999), illiquidity (Sunder, 1992), manipulation (Forsythe and Lundholm, 1990; Noth and Weber, 1998), and lack of equilibrium (Anderson and Holt, 1997; Scharfstein and Stein, 1990).1

[1] this section is quoted verbatim from pages 47-50 of the paper by Chen, Fine and Huberman [**?**] since the methods and analysis described therein are essential to this thesis

These problems are exacerbated when the groups involved are small and not very experienced at playing in these markets. Even when possible, proper market design is very expensive, fragile, and context specific.

In spite of these obstacles, it is worth noting that certain participants in information markets can have either superior knowledge of the information being sought, or are better processors of the knowledge harnessed by the information market itself. By keeping track of the profits and final holdings of the members, one can determine which participants have these talents, along with their risk attitudes.

In this paper, *[the original paper by Chen, Fine and Huberman]* we propose a method of harnessing the distributed knowledge of a group of individuals by using a two-stage mechanism. In the first stage, an information market is run among members of the group in order to extract risk attitudes from the participants, as well as their ability at predicting a given outcome. This information is used to construct a nonlinear aggregation function that allows for collective predictions of uncertain events. In the second stage, individuals are simply asked to provide forecasts about an uncertain event, and they are rewarded according the accuracy of their forecasts. These individual forecasts are aggregated using the nonlinear function and used to predict the outcome. As we show empirically, this nonlinear aggregation mechanism vastly outperforms both the imperfect market and the best of the participants.

However, these results are achieved in a very particular environment, that of no public information. Public information is bound to introduce strong correlations in the knowledge possessed by members of the group, correlations that are not explicitly taken into account by the above-described aggregation algorithm. So, we propose a set of suitable modifications that would allow the detection of the amount of public information present in a group so as to subtract it. Assuming that subjects can differentiate between the public and private information they hold, that the private aspect of their information is truly private (held only by one individual), and that the public information is truly public (held by at least two individuals), we create a coordination variant of the mechanism which allows for the identification of public information within a group and its subtraction when aggregating individual predictions about uncertain outcomes. Experiments in the laboratory show that this aggregation mechanism outperforms the market, the best player in the group, and the initially proposed aggregation mechanism.

3.1.2 Aggregation Mechanism Design

We consider first an environment in which a set of N people have purely private information about a future event. If all players had the same amount of information about the event and were perfectly risk-neutral, then it would be easy to compute the true posterior probabilities using Bayes's rule. If individuals receive independent information conditioned on the true outcome, their prior beliefs are uniform (no other information is available other than the event sequence), and they each report the true posterior probabilities given their information, then the probability of an outcome s, conditioned on all of their observed information I , is given by:

$$P(s|I) = \frac{p_{s_1} p_{s_2} \ldots p_{s_N}}{\sum_{\forall_s} p_{s_1} p_{s_2} \ldots p_{s_N}} \tag{3.1}$$

where p_{s_i} is the probability that individual $i(i = 1, ..., N)$ assigns to outcome s. This result allows us simply to take the individual predictions, multiply them together, and normalize them in order to get an aggregate probability distribution. However, this will only work under the extremely restrictive constraints enumerated above. The first of these issues we will consider is how to design a mechanism that elicits truthful reporting from individuals. ... [T]he following mechanism will induce risk neutral utility maximizing individuals to report their prior probabilities truthfully.

We ask each player to report a vector of perceived state-probabilities, $q_1, q_2, ...q_N$ with the constraint that the vector sums to one. Then the true state x is revealed and each player paid $c_1 + c_2 \times log(q_x)$, where c_1 and c_2 are positive numbers.

While this very simple method might seem to aggregate dispersed information well, it suffers from the fact that, due to their risk attitude, most individuals do not necessarily report their true posterior probabilities conditioned on their information. In most realistic situations, a risk averse person will report a probability distribution that is flatter than her true beliefs as she tends to spread her bets among all possible outcomes. In the extreme case of risk aversion, an individual will report a uniform probability distribution regardless of her information. In this case, no predictive information is revealed by her report. Conversely, a risk-loving individual will tend to report a probability distribution that is more sharply peaked around a particular prediction, and in the extreme case of risk loving behavior a subject's optimal response will be to put all his weight on the most probable state according to his observations. In this case, his report will contain some, but not all the information contained in his observations.

In order to account for both the diverse levels of risk aversion and information strengths, we add a stage to the mechanism. Before individuals are asked to report their beliefs, they participate in an information market designed to elicit their risk attitudes and other relevant behavioral information. This information market is driven by the same information structure in the reporting game. We use information markets to capture the behavioral information that is needed to derive the correct aggregation function. Note that, although the participant pool is too small for the market to act perfectly efficiently, it is a powerful enough mechanism to help us illicit the needed information.

The nonlinear aggregation function that we constructed is of the form:

$$P(s|I) = \frac{p_{s_1}^{\beta_1} p_{s_2}^{\beta_2} ... p_{s_N}^{\beta_N}}{\sum_{\forall_s} p_{s_1}^{\beta_1} p_{s_2}^{\beta_2} ... p_{s_N}^{\beta_N}} \tag{3.2}$$

where β_i is the exponent assigned to individual i. The role of β_i is to help recover the true posterior probabilities from individual i_s report. The value of β for a risk neutral individual is one, as he should report the true probabilities coming out of his information. For a risk averse individual, β_i is greater than one so as to compensate for the flat distribution that he reports. The reverse, namely β_i smaller than one, applies to risk loving individuals. In terms of both the market performance and the individual holdings and risk behavior, a simple functional form for β_i is given by

$$\beta_i = r(\frac{V_i}{\sigma_i})c \tag{3.3}$$

where r is a parameter that captures the risk attitude of the whole market and is reflected in the market prices of the assets, V_i is the utility of individual i , and σ_i is the variance of his holdings over time. We use c as a normalization factor so that if $r = 1$, $\sum \beta_i$ equals the number of individuals. Thus the problem lies in the actual determination of both the risk attitudes of the market as a whole and of the individual players.

To do so, notice that if the market is perfectly efficient then the sum of the prices of the securities should be exactly equal to the payoff of the winning security. However, in the thin markets characterized here, this efficiency condition was rarely met. Moreover, although prices that do not sum to the winning payoff indicate an arbitrage opportunity, it was rarely possible to realize this opportunity with a portfolio purchase (once again, due to the thinness of the market). However, we can use these facts to our advantage. If the sum of the prices is below the winning payoff, then we can infer that the market is risk-averse, while if the price is above this payoff then the market exhibits risk-loving behavior. Thus, the ratio of the winning payoff to the sum of the prices provides a proxy for the risk attitude of the market as a whole.

The ratio of value to risk, (V_i/β_i), captures individual risk attitudes and predictive power. An individual's value V_i is given by the market prices multiplied by his holdings, summed over all the securities. As in portfolio theory (Markowitz, 1959), his amount of risk can be measured by the variance of his values using normalized market prices as probabilities of the possible outcomes.

3.1.3 Experimental Design for Private Information Experiments

In order to test this mechanism we conducted a number of experiments at Hewlett-Packard Laboratories, in Palo Alto, California. The subjects were undergraduate and graduate students at Stanford University and knew the experimental parameters discussed below, as they were part of the instructions and training for the sessions. The five sessions were run with eight to thirteen subjects in each.

We implemented the two-stage mechanism in a laboratory setting. Possible outcomes were referred to as "states" in the experiments. There were 10 possible states, A through J, in all the experiments. Each had an Arrow-Debreu state security associated with it. The information available to the subjects consisted of observed sets of random draws from an urn with replacement. After privately drawing the state for the ensuing period, we filled the urn with one ball for each state, plus an additional two balls for the just-drawn true state security. Thus it is slightly more likely to observe a ball for the true state than others.

We allowed subjects to observe different number of draws from the urn in order to control the amount of information given to the subjects. Three types of information structures were used to ensure that the results obtained were robust. In the first treatment, each subject received three draws from the urn, with replacement. In the second treatment, half of the subjects received five draws with replacement, and the other half received one. In a third treatment, half of the subjects received a random number of draws (averaging three, and also set such that the total number of draws in the community was $3N$) and the other half received three, again with replacement.

The information market we constructed consists of an artificial call market in which the securities are traded. The states were equally likely and randomly drawn. If a state occurred, the associated state security paid off at a value of 1,000 francs. Hence, the expected value of any given security, a priori, was 100 francs. Subjects were provided with some securities and francs at the beginning of each period.

Each period consists of six rounds, lasting 90 seconds each. At the end of each round, our system gathers the bids and asks and determines market price and volume. The transactions are then completed and another call round began. At the end of six trading rounds the period is over, the true state security is revealed, and subjects are paid according to the holdings of that security. This procedure is then repeated in the next period, with no correlation between the states drawn in each period.

In the second-stage, every subject played under the same information structure as in the first stage, although the draws and the true states were independent from those in the first. Each period they received their draws from the urn and 100 tickets. They were asked to distribute these tickets across the 10 states with the constraint that all 100 tickets must be spent each period and that at least one ticket is spent on each state. Since the fraction of tickets spent determines p_{s_i}, this implies that p_{s_i} is never zero. The subjects were given a chart that told them how many francs they would earn upon the realization of the true state as a function of the number of tickets spent on the true state security. The payoff is a linear function of the log of the percentage of tickets placed in the winning state.

3.2 Augmenting the Chen, Fine and Huberman approach with Social Information

From the discussion above, we see that the Chen, Fine and Huberman two-stage approach aims to tackle two main weaknesses previously discussed: correcting for the risk tendencies that are usually embedded in information markets, and accounting for social connectivity. The method discounts the effects of social connectivity by requesting users to provide vectors (q_i) denoting their belief of what information is in the public domain, as opposed to privately held information (p_i). Asking the participants to provide these q vectors assumes that the participants are in general adept at distinguishing what is public from what is private, and interpreting how the information in the public domain might affect other participants' bets in the experiment (in their experiment design, Chen, Fine and Huberman ask participants to match their $q_i s$ with another participants' $q_j s$ in such a way that will profit them most).

In many cases, however, it is reasonable to expect that if we ask the participants to provide such vectors q_i, the vectors may not be representative of the amount of public information in the market, especially when it is not clear to the participants what is public and what is truly private (the Chen, Fine and Huberman experiments used public and private draws from an urn so that there was no doubt what was truly private and what was public). We believe that we might be able to do a good job of estimating the amount of public information indirectly by analyzing data showing the social connectivity of the participants - who talks to whom, how often they talk, how often one person's opinions influence the other in general and in particular subjects, etc.

This sociometric data may be available from a number of sources such as previous studies and interviews with the participants. However, with the increasing prevalence of portable and wearable computer devices, we are also be able to get representative sociometric data by using these devices to directly measure both the quality and quantity of contact between individuals and using that as a proxy for social connectivity [?], which in turn should aid in the enterprise of determining how much of the information is public and how much is private.

Chapter 4

Experimental Setup

We set up our experiment with a number of goals in mind. First, we sought to determine the applicability of the analysis of the Chen et. al. approach to real world situations. To this end, we ran two parallel experiments; the first was a replica of the Chen et. al. experiment where subjects were asked to draw from urns - both publicly and privately - and subsequently asked to give probability vectors that were then analyzed (see chapter 5) to determine the suitability of the approach as an information aggregation mechanism. Our "real world"equivalent of the Chen et. al. experiment was a three-stage experiment designed to have a similar information structure to the Chen et. al. experiment, except with less control over the information revealed to the participants as we would expect in real-world situations. The three stages of the experiment were a "treasure hunt stage", a "betting stage" and a "trading stage".

4.1 Treasure Hunt

The participants (six per run) were allowed 20 minutes to roam a specified area of the MIT Media Lab, with instruction to look for as many clues as possible. Each such clue contained a photo of someone holding an electronic gadgets of some sort (PDAs, bluetooth headsets, led-infested gizmos, etc). An example of such a clue is shown in *Figure 4-1*

Information (similar to the urn drawings of the Chen et. al. setup) was dissem-

Figure 4-1: A sample clue

inated to the subjects by way of these clues. The participants were instructed to pay attention to the contents of each such clue (without altering/moving it), with the understanding that the each clue contained information that would be useful in the betting and trading stages of the experiment to follow; however, they were not told exactly what to look for in the clues. In this way, we attempted to approximate real-world situations where people may be asked to provide information from their recollection, except they did not initially have specific instructions what information to gather. (See the section on applications, §**??** for an extended example of such a scenario).

The six participants in each run were divided into two groups/teams of three each. The purpose of this division was to create distinguishable groups with more intra-group contact than inter-group contact, for the purposes of measuring how such contact influenced the eventual results. To ensure the distinction we desired, we added another artificial information boundary - group members were allowed to tell each other where some of the clues were hidden without revealing the contents of the clues themselves; however, group members from different groups were not allowed to communicate. In addition, each subject was informed that they would be ranked according to the combined performance (entropy) of both the group and the individual in the betting and the market stage; this ensured that there was some incentive to share information between group members while discouraging team members from sharing everything, since we desired that the participants should eventually have

	Team One			Team Two		
clue number	user1	user2	user3	user4	user5	user6
1						
2						

Table 4.1: A sample of the information sheets provided for recording observations

different information, some of which would be private and some public.

To aid the subjects in recollection, we provided them with blank sheets of paper on which to record whatever they thought was necessary. Furthermore, since we as the experimenters needed to know exactly what clues each subject saw, we gave them a table to fill, an excerpt of which is shown in *Table 4.1*.

Under the column corresponding to him/her, each subject checked all the clues that they saw, leaving everything else empty. We filled in the rest of the table as described in §4.2

4.1.1 Measuring Social Contact

As mentioned previously in §3.2, we also desired to estimate public information independently of the responses of the subject. To do this, we needed to observe and measure 'social contact' between the subjects, and from these measurements try and construct a picture of what information was public, and test if we get similar or better results than relying only on the subjects' reports. To accomplish this goal, each subject was fitted with an *UbER Badge* [**?**].

The *UbER Badge* (*Figure 4-2*) is a wearable electronic device that is capable of both radio frequency and infra red communication as well as displaying simple graphics and scrolling text on an LED display. The UbER Badge can also be augmented with fixed infrared beacons (*Figure 4-3*) that constantly broadcast an IR signal, so that whenever a person stands close to one such beacon, an event is registered on the badge and transmitted over radio frequency to a base station that logs all such events.

Using the UbER Badge, we were able to detect every time two subjects were facing each other (presumably in communication of some sort) via the IR bookmarking

Figure 4-2: One subject wearing the UbER Badge

Figure 4-3: The IR Beacon

capabilities of the badges. In addition, an IR beacon was fixed next to each clue, so that we could also reliably tell which clues the subject saw independently of what the subjects write on the information sheet in *Figure 4.1*.

4.2 Betting Stage

Once the subjects were done with the treasure hunt stage (§4.1), we collected the information sheets and filled in the remaining columns; under each subjects' name, we checked against clue numbers to indicate that a particular subject saw the corresponding clue, so that eventually, all the information sheets for one run of the treasure hunt had identical information. Using that information, we expected that the subjects should be able to distinguish truly private information from public (where public is defined as information shared between two or more individuals) in a relatively straightforward manner: a clue is truly private if it was seen by only one individual, otherwise it is public.

At this point, we then informed the subjects that the *people* in the clues each appeared a specified number of times that formed a specific distribution, with one of the people appearing more times than anyone else and therefore being the 'true' state. The subjects had been introduced to all these people before, so we did not expect any problems in differentiating one from the other, but to avoid any such ambiguity, each of the five people appearing in the clues was assigned one of five colors - blue, green, white, red and yellow - and the subjects were each given photos of the people to aid in recollection.

The subjects were then asked to supply the same information as in the Chen et. al. experiment: vectors p of the posterior probabilities of the frequency of appearance of the people in the clues given their information, and vectors q of bets placed by someone else in the room regarding the perceived state probabilities. To create an incentive to ensure truthful reporting, the payoff was as in [?] - a linear function of the log percentage of bets placed in the true/winning state (for the first set of bets p) added to a scaled log function of the players q bets matched with another player in the game:

$$P = c_1 + c_2 \times \log(p_{ix}) + f(\overrightarrow{q_i}, \overrightarrow{q_j}) \times (c_3 + c_4 \times \log(q_{xj})) \tag{4.1}$$

where c_1, c_2, c_3 and c_4 are positive constants, j is chosen such that $f(\overrightarrow{q_i}, \overrightarrow{q_j}) \geq f(\overrightarrow{q_i}, \overrightarrow{q_k})$ for all k, and the function $f(.)$ is given by:

$$f(\overrightarrow{x}, \overrightarrow{y}) = \left(\frac{1 - [\sum_s |x_s - y_s|]}{2}\right)^2 \cdot \overrightarrow{y} \tag{4.2}$$

4.3 Trading Stage

Finally, the subjects participated in a market similar to that in [?] to determine the market prediction of the distribution of the different people (henceforth *states*) in the urn, whereby in our case, the virtual urn was supplied by the treasure hunt stage. Instead of running a physical call market (where the subjects write their bids/asks which are physically collected by a trader, who determines the market price and completes the transactions), we implemented an online call market with exactly those same capabilities, i.e. an online call market which allowed subjects to sell and buy securities based on their perceived distribution of the true distribution of securities. *Figure 4-4* shows a screenshot of the online market.

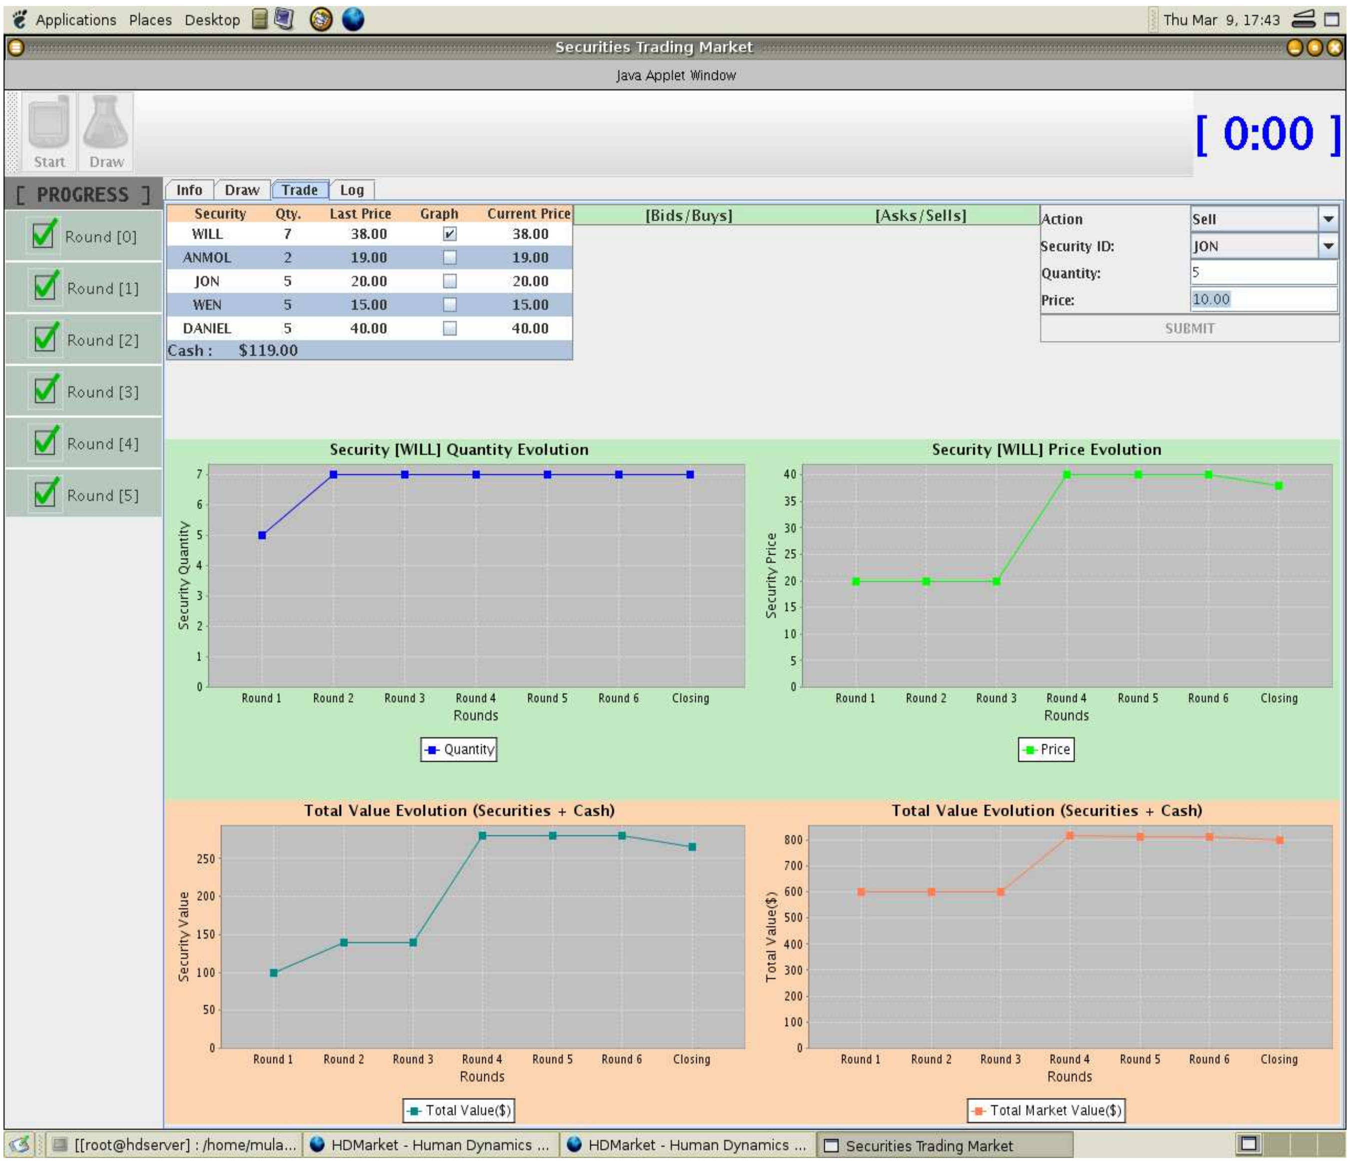

Figure 4-4: A Screenshot of the Online Securities Trading Application

Chapter 5

Analysis and Results

5.1 Analysis

We compared the performance of the different aggregation schemes against the true distribution of the securities in the urn using a Kullback-Leibler entropy measure as in [?], which gives the energy in the error between two distributions. Specifically, if p is the true distribution and q is the prediction using any one of the aggregation methods, then the entropy between the two distributions is given by:

$$KL(p, q) = \sum_{\forall s} p_s \log \left(\frac{p_s}{q_s} \right) \tag{5.1}$$

$KL(p, q)$ is zero only if p and q are identical, otherwise $KL(p, q) \geq 0$. In our analysis, we compared the following results against the truth:

1. *The prediction made by the market.*

2. *The prediction due only to the private bets placed by the subjects, the probability of each outcome s conditioned on all observations I is estimated by:*

$$P(s|I) = \frac{p_{s_1} p_{s_2} \ldots p_{s_N}}{\sum_{\forall_s} p_{s_1} p_{s_2} \ldots p_{s_N}} \tag{5.2}$$

3. *The prediction due only to the private bets, modified by the risk information (gathered from the market stage as explained in §3.1.2) and aggregated by:*

$$P(s|I) = \frac{p_{s_1}^{\beta_1} p_{s_2}^{\beta_2} ... p_{s_N}^{\beta_N}}{\sum_{\forall_s} p_{s_1}^{\beta_1} p_{s_2}^{\beta_2} ... p_{s_N}^{\beta_N}} \tag{5.3}$$

where β_i is the risk-dependent exponent assigned to individual i. The role of β_i is to help recover the true posterior probabilities from individual i_s report. The value of β for a risk neutral individual is one, as he should report the true probabilities coming out of his information. For a risk averse individual, β_i is greater than one so as to compensate for the flat distribution that he reports. The reverse, namely β_i smaller than one, applies to risk loving individuals. In terms of both the market performance and the individual holdings and risk behavior, a simple functional form for β_i is given by

$$\beta_i = r(\frac{V_i}{\sigma_i})c \tag{5.4}$$

where r is a parameter that captures the risk attitude of the whole market and is reflected in the market prices of the assets, V_i is the utility of individual i , σ_i is the variance of his holdings over time and c is a normalization factor so that if $r = 1$, $\sum \beta_i$ equals the number of individuals.

4. *The prediction due to the private bets and the public bets, where each person's public vector was modified by their private vector and the result augmented by risk information in calculating the state probabilities as follows:*

$$P(s|I) = \frac{\left(\frac{p_{s_1}}{q_{s_1}}\right)^{\beta_1} \left(\frac{p_{s_2}}{q_{s_2}}\right)^{\beta_2} ... \left(\frac{p_{s_N}}{q_{s_N}}\right)^{\beta_N}}{\sum_{\forall_s} \left(\frac{p_{s_1}}{q_{s_1}}\right)^{\beta_1} \left(\frac{p_{s_2}}{q_{s_2}}\right)^{\beta_2} ... \left(\frac{p_{s_N}}{q_{s_N}}\right)^{\beta_N}} \tag{5.5}$$

5. *The prediction due to the private bets and the public bets, where the public bets were aggregated together into a single public q_s vector according to:*

$$q_s = \frac{\sum_{i=1}^{N} \beta_i q_{si}}{\sum_{i=1}^{N} \beta_i} \tag{5.6}$$

And the resulting q_s used to estimate the outcome probabilities conditioned on observations as follows:

$$P(s|I) = \frac{\left(\frac{p_{s_1}}{q_s}\right)^{\beta_1}\left(\frac{p_{s_2}}{q_s}\right)^{\beta_2} \dots \left(\frac{p_{s_N}}{q_s}\right)^{\beta_N}}{\sum_{\forall_s}\left(\frac{p_{s_1}}{q_s}\right)^{\beta_1}\left(\frac{p_{s_2}}{q_s}\right)^{\beta_2} \dots \left(\frac{p_{s_N}}{q_s}\right)^{\beta_N}} \tag{5.7}$$

5.2 Results

The graphs that follow use the following key for x-axis labels:

1. *P - Prediction with the private bet vector (p) only*
2. *P/R - Prediction with the private vector augmented by risk information*
3. *P/R/M - Private vector prediction information corrected using each persons public vector of bets, and augmented with risk information*
4. *P/R/G - Private vector prediction information corrected using a weighted sum of the public vector and augmented with risk information*
5. *M - The market prediction*

The entropy of the truth is zero by definition and therefore sits on the x-axis.

5.2.1 The experimental setup by Chen et. al. and our treasure hunt setup are good approximations of each other

As can be seen in *Figure 5-1* (the corresponding data is in *Table* **??** and *Table* **??** in *Appendix A*) the mean of the entropy of all the trials of our implementation of the Chen et. al. setup and the treasure hunt setup show similar trends as different social information is added to the aggregation function. This suggests that the Chen et. al. setup is in fact a good approximation for some real world scenarios.

In our treasure hunt experiment, we took some extra precaution (splitting each run of the experiment into two teams, and restricting communication between the

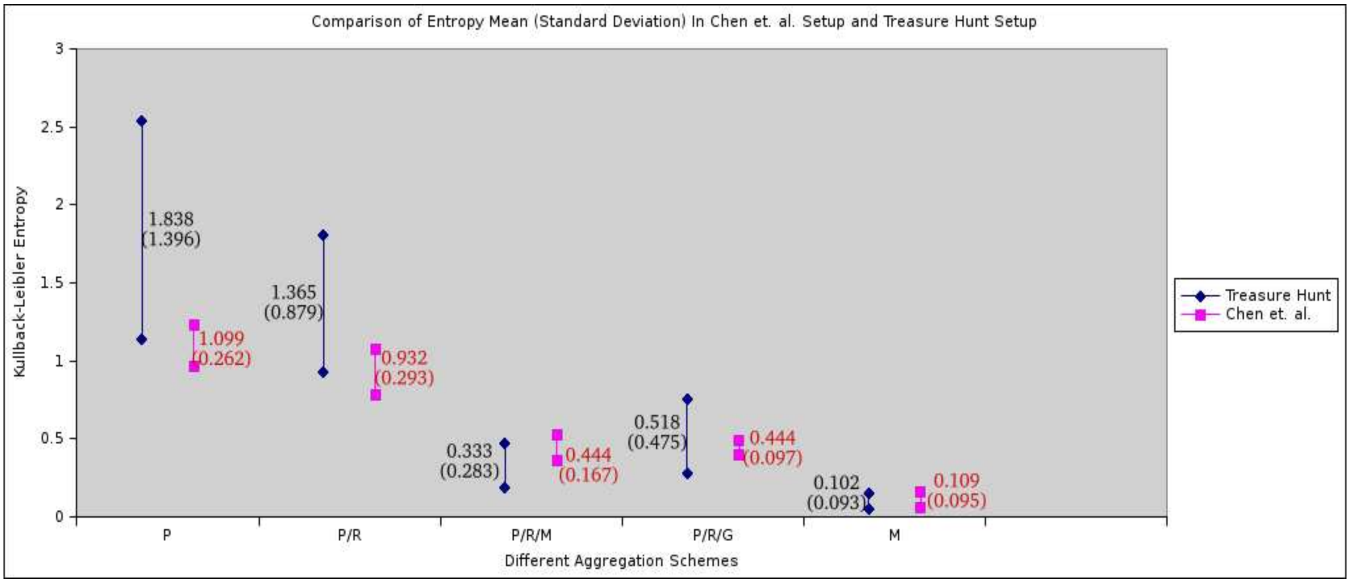

Figure 5-1: Comparison of the mean (standard deviation) of entropy between the Chen et. al. setup and our treasure hunt setup

teams - see §4.1) to ensure that the general information structure resembled that in the original Chen et. al. experiment. Nevertheless, we had less control over the information that was revealed to the subjects or how it was shared between the subjects. We expect that there are many real-world situations with similar conditions - people have imperfect information regarding an event, with different amounts of private and public information available to each person. This result suggests that in such situations, the analysis developed by Chen et. al. can be employed with reasonable results.

5.2.2 Adding more social information improves the result of aggregation

An additional result evident from *Figure 5-1* is that the adding more information of a social nature (the risk preferences of the subjects, the information that is public between the subjects) improves the accuracy of the aggregation mechanism. Using only the private bets of the perceived probability distribution, the treasure hunt prediction has an entropy of 1.838 (with a standard deviation of 1.396); this improves to a mean of 0.388 (0.192 standard deviation) when we correct each subject's private bets with their individual bets and augment that with information on their risk tendencies. This leads us to hypothesize that the market is also somehow incorporating social information in it's aggregation, something we discuss further in §5.2.3

To further test the hypothesis that including more social information results in more accurate predictions, we decided to take advantage of the team structures we enforced. In theory, the treasure hunt subjects should be able to tell more accurately what information was public within their team, because they worked together throughout the hunt, albeit with individualistic goals. Therefore we expect that performing the analysis along team divisions should improve the results:

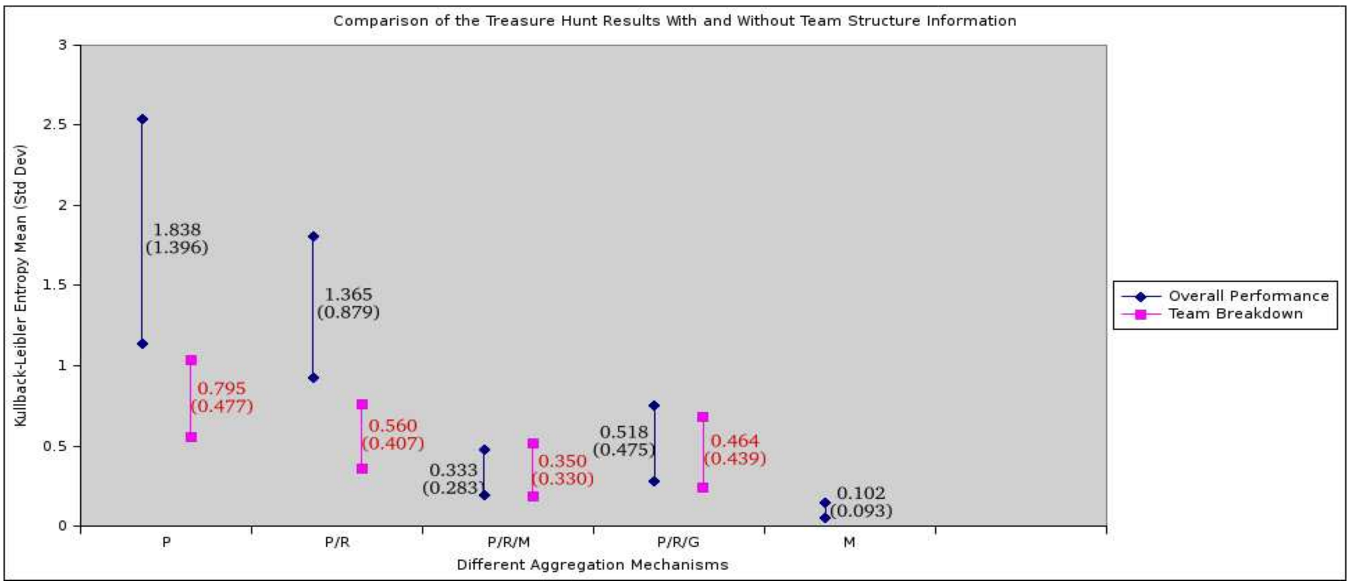

Figure 5-2: Comparison of the mean (standard deviation) of entropy with and without accounting for the team structure of the treasure hunt data

As seen in *Figure 5-2* that in fact turns out to be true; in all cases, the mean entropy when we perform the analysis along teams is lower than when we do not take the team structure into account (the corresponding team data is in *Table* **??**).

As explained earlier in §4.1.1, one of our objectives was to determine if the social information can be measured independently of the reported probabilities and then factored into the analysis. Even though we knew the team breakdown because we enforced it in the beginning, it turns out that from the data collected using the *UbER Badge* and the infra-red beacons, we could actually deduce the team structure:

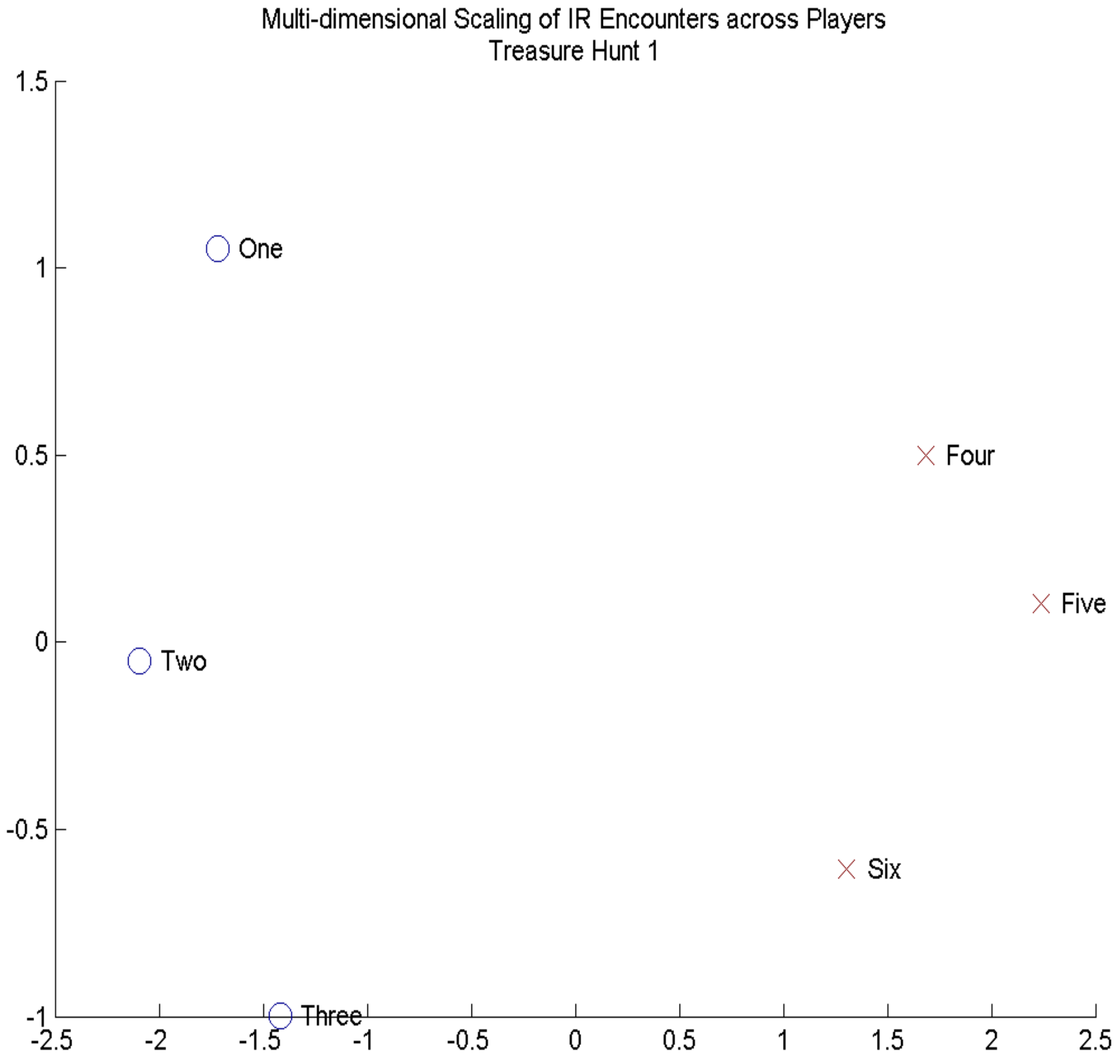

Figure 5-3: Multi-dimensional scaled distance of encounters between the players from the first treasure hunt trial

Figure 5-3 shows a plot of the multi-dimensional scaled distance between the subjects for one instance of the treasure hunt experiment based on the data of infra red logging of encounters between the players. From this data, we can infer which teams particular players belonged to without prior knowledge of the team structure; the inferred teams could then be used in the analysis to yield results identical to *Figure 5-2*.

This result is useful because it suggests that if we can collect reasonable sociometric data about the population - say using various wearable systems (such as mobile phones), mining e-mail communications, etc - the data may be used to augment the Chen et. al. information aggregation mechanism to yield more accurate results.

5.2.3 The market intrinsically performs some aggregation of the social information

Finally, we analyzed the evolution of the entropy of the market aggregation over successive call rounds (as explained in §4.3, the market is organized as a call market, where users place bids and asks in each round; the bids and asks are collected, a pseudo-demand curve is constructed and a market price for each security is determined at which sales and purchases occur). The result of this analysis is shown in *Figure 5-4*.

This shows that over successive rounds, the market prices generally approach the true distribution of probabilities; over the course of the rounds, the market is aggregating information from the subjects into a more and more accurate result. We believe that part of this is the aggregation of the social information that is explicitly accounted for in the other aggregation mechanisms (through the public information) but is not explicated in the market operation. We hypothesize that over time, as the subjects buy and sell securities representing the different states, they are constantly expressing their beliefs as well as adjusting them by observing and reacting to the stock prices, so that by communicating the prices to the subjects and affecting their behavior, the market is implicitly incorporating the social information present in the other methods.

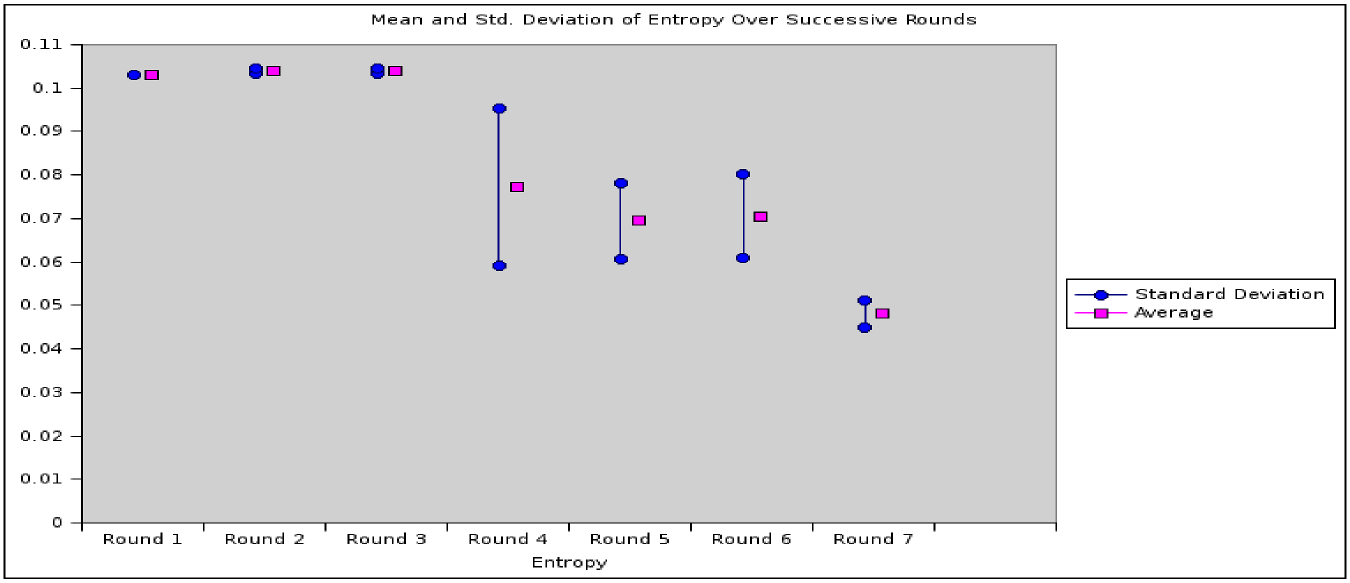

Figure 5-4: Evolution of the entropy (as determined by market prices) over successive rounds of the call market

5.2.4 Correlation Map Of The Individual Bets (as modified by the different mechanisms) compared to the market

Key:

1. p_x : *The prediction using private bets only*
2. p_r : *The prediction using private bets plus risk information*
3. p_m : *The prediction using private bets plus risk information plus my public bets*
4. p_g :*The prediction using private bets plus risk information plus group public bets*
5. p_t : *The market prediction*

Explanation to Prof. Pentland; should probably remove this section later. Please let me know if this makes sense *Figure 5-5* shows the correlation of the bets as determined by the different mechanisms with the market. I calculated this by using the probabilities for each player's security bets *modified* by the additional information in successive aggregation mechanisms. So for instance, if a players private bets were $\overrightarrow{P} = p_1, p_2, p_3, p_4, p_5$ and his/her public bets are $\overrightarrow{Q} = q_1, q_2, q_3, q_4, q_5$, then for that player:

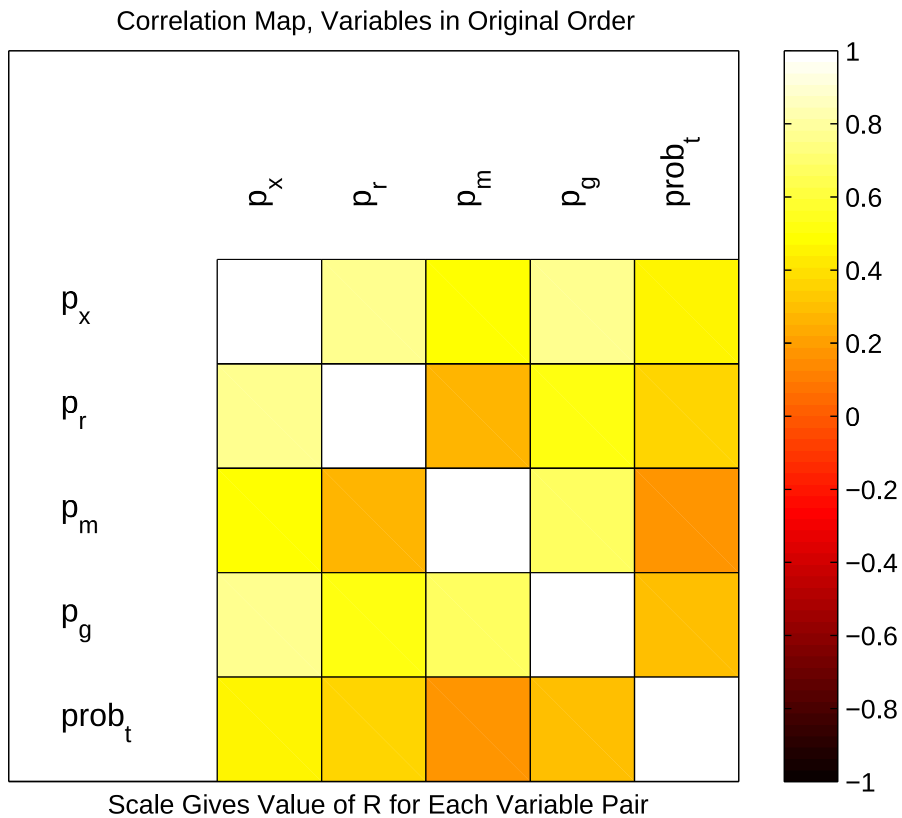

Figure 5-5: The correlation matrix of the individual bets, the urn truth and the market for the 1^{st} treasure hunt session

$p_x = P$ (no modifications whatsoever, as in *Equation 5.2*)

$p_r = P^{\beta_i}$ (each bet is exponentiated with the subjects risk tendencies, similar to *Equation 5.3*

$p_m = (P./Q)^{\beta_i}$ (each bet is first scaled by their public probability for the same bet, and the result exponentiated by risk as in equation *Equation 5.5*

$p_q = (P./\mathbf{Q_s})^{\beta_i}$ (each bet is first scaled by the aggregate public vector value for that security/bet and the result exponentiated by the risk, as in equation *Equation 5.7*. The aggregate Q_s is determined as in *Equation 5.6*.

This is what I had initially understood you suggested I do for the correlation. However, on looking at the resultant map (*Figure 5-5*) there was no evident trend. On thinking about this, I realized that this correlation is comparing the individual bets (modified but NOT aggregated) with the market (which is an aggregate), and I thought that didn't sound right. I thought it might be better if we compared the predicted *aggregate* distribution of each of those methods with that of the market, because then we are comparing two aggregate measures. Which is what I do in the next section.

5.2.5 Correlation map of the aggregate prediction of the different mechanisms compared to the market

	$prob_x$	$prob_r$	$prob_m$	$prob_g$	$prob_t$
$prob_x$	1.0000	0.9740	0.3724	0.4173	0.5990
$prob_r$	0.9740	1.0000	0.5411	0.5846	0.7033
$prob_m$	0.3724	0.5411	1.0000	0.9953	0.7663
$prob_g$	0.4173	0.5846	0.9953	1.0000	0.7886
$prob_u$	0.5990	0.7033	0.7663	0.7886	1.0000

Table 5.1: The correlation matrix of the different aggregation mechanisms for the all treasure hunt sessions

Key:

1. *$prob_x$: The prediction using private bets only*

2. *$prob_r$: The prediction using private bets plus risk information*

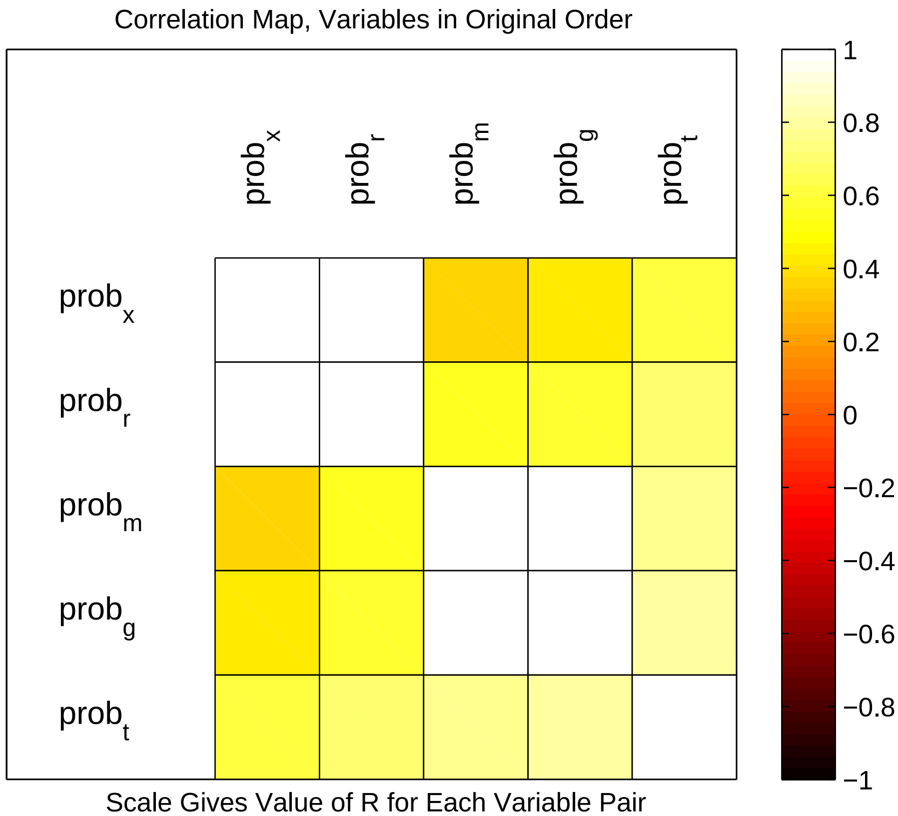

Figure 5-6: The correlation matrix of the different aggregation mechanisms for the all treasure hunt sessions

3. *$prob_m$: The prediction using private bets plus risk information plus my public bets*

4. *$prob_g$: The prediction using private bets plus risk information plus group public bets*

5. *$prob_t$: The market prediction*

The difference between this section and the previous one is that for this section, I calcuated the correlation using the probability distribution that was the result of the aggregate of the individual bets (as opposed to each users bets modified by risk or public information).

This particular correlation map actually shows a trend like the one we expected - the more social information we add, the higher the correlation with the market prediction. Please let me know if this is indeed correct so I can re-write this section accordingly.

Chapter 6

Applications

Having analyzed the performance of the aggregation mechanism proposed by Chen et. al. applied to a "real world" situation and augmented with sociometric data, it is important to point out instances where such an approach might be ideal for aggregating information.

The results show that the aggregation mechanism performs quite well in situations of imperfect private and public information. To illustrate an example where this approach might be ideal, consider soldiers heading out to the battlefield for reconnaissance purposes as they try to acclimatize to new surroundings. The soldiers will typically move out in different units, with each unit surveying a different part of the area of interest. There is relatively frequent contact and communication between members of the same unit, while the contact between soldiers in different units is generally infrequent during the duration of the reconnaissance. In addition, we expect that although the different units are surveying different areas, there is a significant possibility of overlap in the areas surveyed by different units, i.e. one unit might cover some ground previously covered by another unit, or two units might even end up at the same place at the same time (although this second scenario is easier to avoid by co-ordination of the different teams).

This reconnaissance scenario closely mirrors the treasure hunt experimental setup - different teams set out to gather information and, depending on where they go and who they interact with, the information is dispersed among the people, with some of it being public while and the rest private. Therefore, it is reasonable to expect that aggregating the reports of the soldiers using the methods discussed should give a good approximation of the true distribution of states, where those states could be anything such as the number of enemy tanks spotted, the likelihood of an area being a major enemy hide-out, etc. In fact, remembering that the subjects in the treasure hunt were not informed of the object of interest in the clues until the betting stage, we see that the method would do reasonably well in aggregating information that the soldiers were not paying particular attention to but which proves relevant later on.

There are many other situations that resemble the reconnaissance and the treasure hunt enough that the methods discussed should be directly applicable:

1. *Marine divers exploring underwater habitats, collecting information on the animal species inhabiting a specific locations*

2. *Predicting the weather in an area - different people make different observations that help them determine what the weather might be in future - bird migration patterns, previous weather history etc. The people beliefs can be aggregated to determine the likelihood of particular weather patterns in future. It is worth noting that there are already online markets that trade in securities to predict the weather (for instance http://www.theweathermarket.com), and the results show that the prediction of the aggregation mechanisms used rival those made by such markets.*

3. *etc.*

Chapter 7

Conclusion

In conclusion, this thesis set out to investigate how sociometric information can be used to improve the aggregation of dispersed information. We ran two parallel experiment - the treasure hunt and the experiment described in [**?**] - with similar information structures and from the reported perceived probabilities of different states of an event, constructed and compared the predictions of different aggregation mechanisms, explicating the role of social information in some cases (the public bets and team structures) or ignoring such information in other cases. Our analysis led us to a number of conclusions:

1. *The experimental setup described in [**?**] is a good approximation of many real-world situations, therefore the methods developed in [**?**] should yield reasonably accurate results when aggregating information in these scenarios.*

2. *Incorporating more sociometric data into the aggregation mechanisms improves the results significantly.*

3. *The data suggests that through the prices, the market communicates some social information to the users, and by reacting to this information they affect the prices in some way; therefore, the market also performs an aggregation that includes sociometric data - albeit implicitly - in making its prediction.*

4. **there's one more result related to the correlation between the market**

prediction and that of other schemes that I'd need to discuss before stating here

7.1 Future Work

Based on the results of this work, we expect that social networks play a significant role affecting information flow and aggregation. Therefore, it is prudent to investigate further the role of social networks especially in the dissemination of information. The events analyzed in this thesis were simple enough (draws from an urn, observations of clues) and straight-foward enough that it was reasonable to expect that if two people were exposed to the same information, it would affect their beliefs in the same way (in our case, their belief of the distribution of states). However, there are many instances where the effect of social network is not clear cut. Consider trying to predict the probability of adoption of a set of products. If we restrict ourselves to one individual, we expect that that individual's likelihood of adopting the product depends to some extent on whether his/her friends have adopted the said product (social pressure). A simple aggregation mechanism may suggest that the more of the person's friends have the product, the more likely the person is to adopt the product (linear model). However, it fails to account for some relevant attributes of the person or of their social network that may result in highly non-linear behavior. For instance, an individual may not adopt some product unless 20 of his/her classmates have adopted it, but that number may go down to 1 if one of those classmates is also their best friend. Many such examples exist, illustrating that different people are affected by their social networks in different ways, depending on the attributes of the person, the network and possible the prevailing environment. This suggests that while we adopted a somewhat simplistic view of how social networks may affect information flow and aggregation, the reality is more complicated. While it may be impossible to characterize a person's social network in its entirety, future work might seek to come up with slightly more sophisticated models of how social networks impact information flow, which should increase our ability to control for these effects when aggregating information.

From our data, it is evident that the market performed rather well, despite being rather thin (6 subjects per market run). We have several hypotheses for why this is so: (1) each market session consisted of 7 call rounds lasting 120 seconds each, therefore despite the thinness of the market, there were enough rounds for the probabilities to settle closer to the truth; (2) the market is also implicitly including social information in its aggregation as described in §5.2.3; (3) by starting off the market prices at a uniform distribution, we may have already introduces a bias since the entropy of the uniform distribution compared to the truth was not very high. Future work might look further into this issue of market performance and possibly try to explicate the role of sociometric information, the duration (length) of the market, etc.

Appendix A

Tables

	Session 1	Session 2	Session 3	Mean	Variance
private	3.437	0.868	1.207	1.838	1.396
private/risk	2.299	0.556	1.239	1.365	0.879
private/risk/my public	0.516	0.006	0.476	0.333	0.2831
private/risk/group public	0.951	0.010	0.593	0.518	0.475
market	0.2102	0.0525	0.044	0.102	0.0937

Table A.1: Data for the entropy of the prediction due to the treasure hunt experiment

	Session 4	Session 5	Session 6	Mean	Variance
private	1.312	1.180	0.807	1.099	0.262
private/risk	0.820	1.265	0.712	0.932	0.293
private/risk/my public	0.252	0.524	0.556	0.444	0.167
private/risk/group public	0.458	0.340	0.533	0.444	0.097
market	0.198	0.120	0.008	0.109	0.095

Table A.2: Data for the entropy of the prediction of the Chen et. al. setup

	Team 1A	Team 1B	Team 2A	Team 2B	Team 3A	Team 3B	Mean	Variance
private	1.135	1.487	0.230	0.328	0.738	0.853	0.795	0.477
private/risk	1.025	0.677	0.055	0.241	1.007	0.353	0.560	0.407
private/risk/my public	0.491	0.100	0.090	0.020	0.840	0.558	0.350	0.329
private/risk/group public	0.920	0.088	0.096	0.021	0.915	0.745	.464	0.438

Table A.3: Entropy for treasure hunt divided according to team structure

Bibliography

[1] Kay-Yut Chen, Leslie R. Fine, and Bernarde A. Huberman. Predicting the future. *Information Systems Frontiers*, 5(1):47–61, January 2003.

[2] Nathan Norfleet Eagle. *Machine Perception and Learning of Complex Social Systems*. PhD dissertation, Massachusetts Institude of Technology, MIT Media Lab, May 2005.

[3] Steven Gjerstad. Risk aversion, beliefs and prediction market equilibrium. Working Paper, January 2005.

[4] Albert Goodman. Introduction to data collection and analysis. http://www.deakin.edu.au/ agoodman/sci101/index.php, July 1996. Online Book.

[5] Mathew Laibowitz and Joseph A. Paradiso. The uber badge - a versatile platform at the juncture between wearable and social computing. Responsive Environments Group, MIT Media Laboratory, Cambridge, MA, USA, 2003.

[6] Charlse F. Manski. Interpreting the predictions of prediction markets. *Economic Letters*, 91(3):425–429, June 2006.

[7] David M. Pennock, Steve Lawrence, C. Lee Giles, and Finn Arup Nielsen. The real power of artificial markets. *Science*, 291(2):987–988, February 2001.

[8] James Surowiecki. *The Wisdom Of Crowds*, chapter 1.2. DOUBLEDAY, 2003.

[9] Justin Wolfers and Eric Zitzewitz. Interpreting prediction market prices as probabilities. Working Paper, April 2006.