Technical Proposal (MDA972-03-T-0002)
Scaleable Bayesian Updating of Combinatorial Information Markets

Submitted by

Interdisciplinary Center for Economic Science
George Mason University
MSN1B2
4400 University Drive Fairfax, VA 22030
(Contact: David Porter 703-993-4841 email: dporter4@gmu.edu)

Overview

We propose to support the FutureMAP program, which is investigating the ability of information markets to aggregate information from diverse sources on topics of interest to the Department of Defense. In particular, we propose to extend the market design developed by the Net Exchange team, a design that enables trading not only of the probabilities of various basic events, such as the economic, military, and political status of many nations over time, but also on combinations of these events, such as conditional probabilities. The Net Exchange design uses combinatorial information market makers, and coordinates these via arbitrage. Arbitrage coordination is simple and robust, but it does not change probabilities according to Bayes' rule, and so may hinder information aggregation. We plan to use Bayesian network technology to coordinate such market makers, so that updates follow Bayes' rule. We plan to then test whether this design addition improves system accuracy in laboratory experiments.

In order to complete this ambitious project we have assembled a world-renowned team of economists, market designers and operations research experts. The team will be led by Professors Vernon Smith (2002 Nobel Prize winner in Economics) and David Porter. These two individuals are recognized as world leaders in the design, development and testing of new markets. Professor Robin Hanson, a central figure in the design and analysis of information markets and the primary author of the Net Exchange design, will lead the analytical design effort of this project. Professor Kathryn Laskey is an expert in the computational probability and decision theory. Professor Laskey will provide the integration of Bayes Network algorithms into the market design. With the assistance of

Information Extraction and Transport (IET Inc.) scaleable algorithms for Markov Networks will be designed and implemented. In addition to these principals, the technical personnel associated the ICES Experimental Laboratory at George Mason University will assist in testing the new designs with laboratory experiments.

Scientific and Technical Elements of the Proposal: Scaleable Bayesian Updating of Combinatorial Information Markets

In a standard information market, such as found at the Iowa Electronic Markets, traders typically aggregate their information into a consensus probability distribution over a single variable with a handful of (e.g., less than a dozen) possible values. In a combinatorial information market, in contrast, a few traders can aggregate their information into a full joint probability distribution over all combinations of values of some set of variables (e.g., thousands or more possible states). This allows the same set of traders to estimate more probabilities, including joint and conditional probabilities. While combinatorial information markets are relatively new, they seem reasonably well understood in the small-scale case where all of the possible states are explicitly represented in a computer. The more general case, however, remains an open area of research.

Bayesian and Markov networks are technologies that allow computers to reason about probability distributions over sets of states that are too large for computers to explicitly represent. These *graphical statistical models* represent a joint probability distribution over a set of random variables as a graph representing dependencies and a set of parameters representing local probability relations. A Markov graph says that each state is independent of all other states, given its neighboring variables. Such independence assumptions are often plausible, and seem common in human reasoning. Special procedures have been developed to propagate belief updates in graphical models so that when one part of the network changes, the rest of the network also changes according to Bayes' rule. Off-the-shelf methods give exact answers in the case where the network is nearly a tree. The more general case requires approximations, however, and remains an open area of research.

The two ongoing FutureMAP efforts, by Neoteric Technologies and Net Exchange, could both benefit from a closer integration of Bayesian/Markov networks and combinatorial information markets. Neoteric Technologies plans to have selected small groups use standard information markets to estimate particular probability numbers within a Bayesian network. With integrated combinatorial information markets, the same small group could fill in more of the network.

Net Exchange team has pioneered both call market and market maker versions of combinatorial information markets, and has demonstrated the efficacy of all-possible-state versions in human laboratory experiments (with 256 states). To deal with the vast number of states in its FutureMAP field experiment ($>10^{60}$), Net Exchange plans to use overlapping all-possible-state combinatorial information market makers, and to coordinate these via arbitrage. Arbitrage coordination is simple, robust to manipulation, and guaranteed to converge, <u>but does not update according to Bayes' rule</u>.

Such overlapping markets form a Markov network. When updates follow Bayes' rule, each trader can focus on only one part of the network, changing it to reflect his information about that part. All the changes he sees there, even those due to changes in other parts of the network, will directly and transparently reflect information about his part of the network. If updates do not follow Bayes' rule, however, then noise and distortion will cloud the relation between the changes he sees and the local information those changes represent. Thus information aggregation seems better served by Bayes' rule updating.

We propose to test this hypothesis. That is, we propose to advance the state of the art of Markov networks by developing a way for overlapping combinatorial information market makers to be approximately updated according to Bayes' rule. (Since the Net Exchange field experiment network is nothing like a tree, off-the-shelf network updating methods are insufficient.) We also propose to conduct human laboratory experiments comparing information aggregation performance under this updating process to performance under the arbitrage coordination process.

We now discuss combinatorial information market makers, arbitrage coordination, updating via Markov networks, further computational issues, and laboratory experiments in a little more detail.

Combinatorial Information Market Makers

To estimate the chances of the event E = "Terror attack in US in 2004 kills over 100 people," one might convene a committee, or ask an agency to produce a report. Alternatively, one might create a simple information market, where people trade assets of the form "Pays $1 if E". If the going price were $0.30, that would suggest that those willing to "put their money where their mouth is" estimate a probability of about **p**(E) = 0.30.

Similarly, to estimate the chance of the event E conditional on the event F = "US invades Iraq in 2003," one could create a market where people trade assets of the form "Pays $1 if E & F" for assets of the form "Pays $1 if F." If the going price were one of the latter assets for every four of the former assets, that would suggest a conditional probability of **p**(E|F) = 0.25. Taken together with the previous estimate, this suggests that invading Iraq would reduce the chance of a US terror attack.

Simple information markets have many obvious advantages over committee or agency reports, including precision, frequency of updates, openness to challenge, and avoiding disputes about relative expertise. Empirical studies also suggest such markets have advantages in accuracy and robustness to manipulation. A limitation of such markets, however, is that market estimates will not even exist unless each market has at least a handful of traders. This limits the number of estimates one can hope to get out of any given set of participants. Yet once one starts to consider various combinations of scenarios, the number of desired estimates increases very rapidly.

Combinatorial information market makers are designed to overcome this limitation, while retaining the other advantages of information markets. Consider a set of variables, such as A = "Iraq's level of cooperation with UN inspectors in first quarter of 2003," each of which as a handful of values, such as "1 = no cooperation," "2 = minimal cooperation," and so on. Ideally, a combinatorial information market maker always has a number for any joint or conditional probability expressible in terms of these variables, such as **p**(A=1&B=2), **p**(A=1|C=0), or **p**(A=1&B=2|C=0&D=1).

A user can browse these probabilities looking for one he disagrees with, and then change it to a new number by making a matching trade (such as the trades above that match **p**(E) and **p**(E|F)). He can do this even if no other user ever trades on exactly the same combination. There is a financial cost to create such a market maker, but this cost goes as the number of variables allowed, not the number of combinations allowed.

A simple combinatorial information market maker (SCIMM) explicitly represents all possible combinations of the values of its variables, and maintains an exact consistent probability distribution over these combinations. The Net Exchange team has conducted human laboratory experiments showing effective information aggregation by six people in five minutes for the case of a SCIMM with 256 combinations. SCIMMs are simply not feasible, however, for sets of more than about 30 variables (which have at least 10^9 combinations).

Arbitrage Coordination

Current plans for the Net Exchange field experiment call for over 200 variables (5 parameters for each of 8 nations over 4 quarters). This implies at least 10^{60} combinations, which is far more than computers can explicitly represent. To deal with this, Net Exchange plans on creating hundreds or more SCIMMs, each of which deals in some set of roughly 10-15 variables, and then only allowing a trade when a SCIMM can be found that deals in all of the variables mentioned in that trade. The SCIMM variable sets chosen will be those that seem most likely to appear in trades, such variables that are close in space or time. This will allow users to trade on a great many, though not all possible, combinations.

Net Exchange plans to coordinate these simple SCIMMs via arbitrage. When two SCIMMs have variables in common, then there are some trades that they can both handle. If they offer different probabilities on those common trades, then one can extract cash from them by having them trade with each other until they agree. Repeated arbitrage must eventually converge to a situation where all SCIMMs agree, and there is a bound on the profits users can make by finding inconsistencies before the system does.

Unfortunately, using arbitrage to coordinate a set of SCIMMs does not result in the same dynamic behavior as if there were one big SCIMM for all the variables. If there were one big SCIMM, then when a user changed **p**(B=2), the value of **p**(A=1|B=2) would not change, and when a user changed **p**(A=1|B=2), the value of **p**(B=2) would not change. There would be a change in the probability corresponding to the risk the user took in making the trade, and the only other changes made would be according to Bayes' rule. That is, it would be as if the user had introduced new evidence on the variables he changed, and rational statistics were used to update all other probabilities based on this evidence.

Arbitrage coordination does not generally update according to Bayes' rule. For example, a user who changes either **p**(B=2) or **p**(A=1|B=2) will typically change both of these numbers. Because of this, changes to a particular probability will often be due not to information that users have about that probability, but instead to unintended and uninformative side effects of changes made elsewhere. This noise and distortion would seem to interfere with the essential process of information aggregation, wherein users learn from and respond to the information embodied in the trades of others.

Updating Markov Networks

When SCIMMs agree with each other on the probabilities of shared variables, they in effect describe a full joint probability distribution over all combinations of all their variables. And it turns out that a Markov network describes this probability distribution. If one takes each variable in a SCIMM and links to all of the other variables in that SCIMM, it turns out that in the full joint distribution each variable is independent of all other variables, given the variables that it is linked to.

This is exactly the defining assumption of a Markov network. Markov networks are a popular way to represent probability distributions. They are popular because they allow a small number of parameters to describe a distribution over a very large space of possibilities, and because the independence assumptions they make often seem plausible, and seem common in human reasoning.

Ideally the SCIMMs would all agree with each other before a user made a trade. During the trade, the user would specify a new value for some probability. After the trade, all the SCIMMs would agree with each other and this new probability. Furthermore, all of the changes would be according to Bayes' rule. It turns out that a great deal of research over the last few decades, much of it funded by DARPA, has gone into this computational task, which is called inference in Markov networks.

There is now a standard method (join tree propagation), embodied in many available software packages, for doing exact Markov network inference in the case where the Markov network forms a tree. (This method is also standard for Bayesian networks, which are closely related to Markov networks.) This method can also be used in networks that are sparsely connected, by merging variables into small clusters arranged in a tree structure.

Unfortunately, the network corresponding to the set of overlapping SCIMMs envisioned for the Net Exchange field experiment is nothing like a tree, and could not be plausibly made into a tree with reasonable-sized variable clusters. When variables evolve in time, even when variables in an individual time slice can be arranged in a tree of small clusters, correlations over time destroy this tree structure. Thus non-tree networks will probably be typical in desired applications of combinatorial information markets.

There is no standard method for doing inference in non-tree Markov networks, and one can show that exact computation is intractable in the general case. Therefore, one must settle for approximations. There has been extensive research, however, much of it funded by DARPA, resulting in many promising approximation methods. It thus seems possible for researchers experienced in this area to try to identify an approximation appropriate to updating sets of overlapping SCIMMs.

Non-Standard Computational Issues

It is not enough to simply find an approximation for updating the probabilities in a non-tree Markov network of overlapping SCIMMs, however. User assets must also be updated, and special attention must be paid to avoiding becoming a money-pump.

In addition to managing a probability distribution, a combinatorial information market maker must also manage the assets held by each trader that are available to trade in these markets. It must verify that users have assets to pay for any proposed trade, it must change their assets in response to a trade, and it should be able to cheaply calculate both the market value of the user's portfolio and how far long or short the user is regarding any given trade. These asset calculations may benefit from close integration with probability calculations.

If an approximation method does not result in an exactly self-consistent probability distribution, then in principle users who detect such inconsistencies could make free money via trades that takes advantage of such inconsistencies. While small profits like this going to users is tolerable, it is important to think carefully about how to ensure that large profits wouldn't be involved.

Finally, it would be nice to think about whether to and how to let users change the network structure.

Experiments

This proposal has been motivated by the hypothesis that updating the probabilities of a combinatorial information market maker via an approximation to Bayes' rule would produce superior information aggregation than via arbitrage. This hypothesis seems plausible but is hardly obvious, and should be checked in controlled experiments.

In a controlled laboratory experiment, we would privately determine the values of a certain set of variables, and then we would give random human subjects some private clues about those values. (Usually, everyone would be told how the values and clues are determined.) Subjects would be given an initial stake of cash, would be allowed to trade that stake with a combinatorial information market maker dealing in these variables, and would take home their net trading gains. We would measure performance by looking at the distance between the market's probability distribution and the actual distribution of values. We would repeat this whole process as many times as possible, in order to make as precise as possible a performance estimate.

The above procedure should be done both for the arbitrage process, and for some version of an approximation to Bayes' rule updating. While making this one change, all other features of the experiments should be held as constant as possible. Ideally, subjects would use the same user interface, have the same time to trade, and be given the same clues about the same actual values. While the same subjects could not be used, since they will have learned about the variable values, subjects should be drawn from the same pool of available subjects.

If time and resources permit, we might make this comparison in several different environments, varying in the number and complexity of the variables to estimate, or in the number of subjects and the time they are given. Subjects might be given other activities that distract them from trading, or be given reasons to want to bias the market estimates.

Contribution and Relevance to DARPA Mission

DARPA's FutureMAP (Futures Markets Applied to Prediction) program is investigating the use of market-based techniques for avoiding surprise and predicting future events. This program has two main efforts. Neoteric Technologies plans to have selected small groups use standard information markets to estimate particular probability numbers within a Bayesian network. Net Exchange plans to use combinatorial markets to allow aggregation on conditional probabilities and other complex relations in international military, political, and economic activity. We propose to complement both these efforts by better integrating combinatorial information markets and Bayesian networks.

Such integration can both allow markets to fill in more parameters of a Bayesian network, and it can allow quicker and more accurate updating across markets that allow traders to express conditional probabilities and other complex relations. In particular, since the Net Exchange market design uses arbitrage to coordinate prices, a procedure that does not follow Bayes' rule, there is the question of whether we can do better by updating according to Bayes' rule.

To answer this question and contribute to DARPA's Mission of FutureMAP we plan to develop robust and fast Bayesian Network algorithms to be incorporated into the Net Exchange market maker system. These algorithms will be tested on large-scale problems to determine their accuracy and speed of update to make the markets continuous. In addition to such simulations, we plan to address the behavioral question of whether such a mechanism can better help humans to aggregate information in complex environments containing many related variables. The experiments that we will design will provide DARPA with information about whether a more sophisticated algorithm can obtain faster and more accurate conditional probability updates and movements in market prices.

Capabilities and Related Experience

The qualifications and achievements of the team are unparalleled. A list of each of the principals is provided below along with a brief description of their abilities as related to this project.

Dr. Vernon Smith is the 2002 Nobel Prize winner in Economics and Professor of Law and Economics at George Mason University. Dr Smith earned his bachelor's degree in electrical engineering at the California Institute of Technology in 1949, a Masters in Economics from the University of Kansas in 1952 and a Ph.D. from Harvard University in 1955. In 1962, while a Visiting Professor at Stanford University, published his first experimental findings in the Journal of Political Economy. The article, "*An Experimental Study of Market Behavior*," is today considered the landmark paper on experimental economics. Dr. Smith has published and co-published numerous seminal works exploring, and defining, experimental economics as well as other economic disciplines. His "*The Principle of Unanimity and Voluntary Consent in Social Choice*" published in the Journal of Political Economy in 1977 initiated the systematic study of institutional design for public choice decisions. The 1982 "*Microeconomic Systems as an Experimental Science*" in the American Economic Review marked the still adhered-to methodology for experimental economics. His 1982 "*A Combinatorial Auction Mechanism for Airport Time Slot Allocation*" in the Bell Journal of Economics provided a real-world application of experimental economics on economic systems design. The 1988 "*Bubbles, Crashes and Endogenous Expectations in Experimental Spot Asset Markets*" published in Econometrica examined stock market bubbles and rational expectations. The 1994 "*Preferences, Property Rights and Anonymity in Bargaining Games*" in Games and Economic Behavior started the systematic study of personal exchange. Dr. Smith has formed the Interdisciplinary Center for Experimental Science (ICES) at George Mason University, which Dr. Smith now directs. At ICES Dr. Smith and colleagues continue to conduct economic experiments and solidify the application of developed knowledge. Current research is focused on the design and testing of markets for electric power, water and spectrum licenses. Dr. Smith and his colleagues have also worked with the Australian and New Zealand governments on privatization issues, developed market designs for the Arizona stock exchange, and designed an electronic market for water in California.

Dr. David Porter is a Professor of Arts and Sciences and IFREE Gilder Fellow at the Interdisciplinary Center for Economic Science at George Mason University. Dr. Porter was the lead designer of the Cassini Resource Exchange that allowed scientists on the mission to Saturn to barter resources over the Internet. The project won a Ford Foundation prize for innovations in American Government. Dr. Porter has also been involved in developing electronic exchange systems for over twenty years. He was involved in the first logistics auction for trucking services. The auction used combined value methods and resulted in substantial savings for Sears Logistic Services. The Southern California Emission markets RECLAIM, was design and developed by Dr. Porter while he was a faculty member at the California Institute of Technology. This exchange system was the first ever two-sided combinatorial trading market. In addition to designing and implementing exchanges systems, Dr Porter is also widely published in the area of experimental economics testbedding. His tesbedding activities included the testing of combinatorial designs for the allocation of personal communications spectrum for the Federal Communications Commission. He has also developed testbed activities for the scheduling of resources that include Deep Space Network antennae, software licenses, and science observations on spacecrafts. Dr. Porter received his M.S. in Mathematics and Ph.D. in Economics from the University of Arizona.

Dr. Robin Hanson is an Assistant Professor of Economics at George Mason University. Previously, Dr. Hanson was a Robert Wood Johnson Foundation Health Policy Scholar at the University of California at Berkeley. Dr. Hanson received his Ph.D. in Social Science from California Institute of Technology, and his M.S. in Physics and M.A. in Conceptual Foundations of Science from the University of Chicago. Between Caltech and Chicago (from 1984 to1993), Dr. Hanson was a computer researcher in artificial intelligence and

Bayesian statistics, first at Lockheed (now Lockheed-Martin), and then at NASA Ames Research Center. During that period he was also a consultant to Xanadu, Inc. a famous precursor to the World Wide Web.

Dr. Hanson has published in *CATO Journal, Communications of the ACM, Economics Letters, Econometrica, Extropy, IEEE Intelligent Systems, Information Systems Frontiers, International Joint Conference on Artificial Intelligence, Journal of Evolution and Technology, Journal of Public Economics, Social Epistemology, Social Philosophy and Policy, Theory and Decision,* and *Workshop on Physics and Computation.*

Dr. Hanson has spoken and written advocating wider use of speculative markets to aggregate information for 14 years. His writings influenced the construction of several web-based betting markets, one of which won Dr. Hanson a Prix Ars Electronica Golden Nica electronic art prize in 1995. As part of the Net Exchange FutureMAP team, Dr. Hanson chose the domain of the field experiment (international instability), designed the market mechanism Net Exchange will use (combinatorial information market makers), designed the arbitrage process for coordinating these market makers, and designed the environments they used to compare market mechanisms in lab experiments.

Dr. Kathryn Laskey is an Associate Professor of Systems Engineering and Operations Research at George Mason University. Dr. Laskey is a recognized expert in extraction, representation, instantiation, and maintenance of complex knowledge schema. She has published extensively in the areas of inference, knowledge representation, learning, and the integration of symbolic and sub-symbolic knowledge. She has developed an application-neutral, object-oriented framework for rapidly developing computational representations of domain knowledge. She has worked with colleagues to apply this framework to a number of problems involving reasoning about military entities and their behavior. Dr. Laskey also specializes in fusing all-source intelligence message product with generic network fragments to produce situation updates for operational force commanders. She assisted in the development of an antiterrorist risk management system and a system for detecting large-scale coordinated cyber-attacks. She was co-program chair in 1999 and general chair in 2000 for the Conference on Uncertainty in Artificial Intelligence, has been an associate editor of *IEEE Transactions on Systems, Man, and Cybernetics*, and has served on several panels of the Committee on National Statistics of the National Academy of Sciences. She teaches courses in Bayesian reasoning, decision support, and systems engineering at George Mason University and has supervised doctoral dissertations in knowledge representation, inference and decision support, and Bayesian learning.

Information Extraction and Transport (IET) is a world leader in development and application of probabilistic inferencing technology. IET has over 12 years experience in the development and application of decision algorithms – all invented within the last 15 years. IET is in transition to become the industry leader in providing decision algorithm and probabilistic inference software and consulting services. IET currently provides software, training and consulting on the use of these technologies in real-world systems. IET technology is realized through the use of *Bayesian Belief Networks (BBNs)*. BBN models have seen wide application in expert system, decision support, and policy analysis problems. IET has applied BBN technology on a number of DARPA programs, including Warbreaker, Dynamic Database, and Information Assurance. They provide a parsimonious and computationally efficient way to represent probability models over large numbers of interdependent hypotheses. The scientists from IET who will be assisting on this project will be *Drs. Bikash Sabata and Masami Takikawa.* Dr. Sabata has proven experience in leading teams in design and implementation of advanced algorithms and analysis engines for pattern recognition, machine learning, modeling, distributed multimedia systems, video processing, game theory, and automated resource allocation. He has led teams that built: 1) pattern recognition and machine learning algorithms to model and control the web infrastructure to assure quality of service of web-based applications, 2) algorithms for protecting web sites from denial of service attacks, and 3) adaptive mechanisms for quality of service in distributed applications (e.g., content transcoding, resource allocation, admission control and protocols for reliable multicast). Dr. Sabata received a B.Tech from the Indian Institute of Technology, Bombay in 1986, an M.S. in Electrical and Computer Engineering in 1987 and a Ph.D. in Electrical and Computer Engineering in 1993 both from The University of Texas at Austin. Dr. Takikawa has sixteen years of research and industrial experience in software design and development. He has a strong background in probabilistic inference, artificial intelligence, and programming language design and implementation. In his work at IET, Dr. Takikawa has been the primary designer and developer of the Java Symbolic Probabilistic Inference (JSPI) package, one of the most efficient and feature-rich BN inference engines in the world, JSPIScript, an object-oriented scripting language equipped with various constructs for BN modeling and inference, and Java Probabilistic Frame (JPF), a frame-based probabilistic knowledge representation language. Dr. Takikawa has also developed unique capabilities to scale the dynamic instantiation and processing of very large data sets applied to a large number of probabilistic models associated with disparate and dynamically changing geographical-temporal regions.

In addition to the principals listed above, there is access to technical personnel, graduate students and laboratory facilities at the ICES Laboratory. The laboratory consists of 40 high-end workstations for subjects. The monitor room facilities include 4 dedicated servers that have a minimum of 2 processors with 1.5GB of memory and 2 terabytes of storage capacity. The laboratory is maintained and run by technical personnel who maintain programs and networks.