In *Workshop on Economic Agents, Models, and Mechanisms at the 17th International Joint Conference on Artificial Intelligence (IJCAI-2001)*, August 2001.

NP Markets, or How To Get Everyone Else to Solve Your Intractable Problems

David M. Pennock
NEC Research Institute, 4 Independence Way, Princeton, NJ 08540
dpennock@research.nj.nec.com

Abstract

I discuss the prospects of opening securities markets in hard computational problems, including satisfiability, counting problems, and Bayesian inference problems. Such *NP markets* would offer direct monetary incentives for the development of better algorithms. Market prices would serve as collective approximate solutions, and bid-ask spreads may reflect problem difficulty. Some markets offer controlled settings for investigating the speed of information incorporation in markets, and exploring evidence of bounded rationality and imprecise subjective probabilities among market participants.

Going once . . . Going twice . . . Sold! Congratulations!
You are now the proud owner of . . .

. . . one unit of one instance of three-satisfiability problem #V300:C1000:ID854!

1 Introduction

In May 1997, IBM's Deep Blue became the first computer to defeat a reigning world chess champion. Along with the $700,000 winner's purse, the programming team won the $100,000 Fredkin prize, an amount set aside in 1980 to stimulate computer chess research. A $1.5 million prize still awaits the first computer champion of the Chinese board game Go.[1] Every year the Loebner prize,[2] and $2000, goes to the computer program that is judged most human-like; a $100,000 prize is earmarked for the first computer to fool a judge into believing that it is human, a gold standard for artificial intelligence first popularized by Turing [28]. In January 1997, RSA Data Security sponsored a $10,000 prize for the first person to decode a particular 56-bit DES-encrypted message.[3] A distributed collection of computers, coordinated across the Internet, accomplished the task, and collected the prize money, in June of the same year.[4]

By sponsoring contests, a funding agent can provide incentive for researchers to tackle the problems that it wants solved. However, significant rewards of this type are rare, and are certainly negligible compared to the number of challenging problems of interest. Few individuals or groups can afford to back a meaningful prize and advertise it sufficiently. Others may simply not desire such publicity, especially if the problem itself is proprietary.

[1] The Ing prize, sponsored by Acer Incorporated and the Ing Chang-Ki Wei-Chi (Go) Educational Foundation.
[2] `http://www.loebner.net/Prizef/loebner-prize.html`
[3] `http://www.rsa.com/des/`
[4] `http://www.frii.com/ rcv/deschall.htm`

Hanson [15, 14] proposes an *Idea Futures* market where participants can bet on future developments in science, technology, and other arenas of public interest. He argues that the reward structure of such a market encourages honest revelation of opinions among scientists, yielding prices that form accurate forecasts for use by funding agencies, public policy leaders, the media, and other interested parties. The concept is operational as a Web game (run with play money) called the Foresight Exchange.[5] In practice, however, it is often difficult to precisely define each security's payoff-triggering event in a way that foresees all possible eventualities. As a result, an unbiased human judge is generally required for every security. When a security is vaguely defined or the judge is not trusted, agents are usually wary of trading in it, and its price may have little or no informative value.

There *is* a class of problems that are both interesting and difficult to solve, and yet any proposed solution can be easily and precisely verified: namely, the class of NP-complete problems. In Section 3, I discuss the prospect of opening securities markets that pay off contingent on the discovery of solutions to particular instances of an NP-complete problems. Such *NP markets* would provide direct monetary incentives for developers to test and improve their algorithms, and allow funding agents to target rewards to the designers of the best algorithms for the most interesting problems. In Sections 4 and 5, I discuss markets in #P-complete problems, where prices serve as collective approximate bounds on the number of solutions, and bid-ask spreads may indicate problem difficulty. Markets in Bayesian inference problems may prove a natural testbed for controlled experiments to measure the speed of information incorporation in markets. Markets in SAT counting problems and Bayesian inference problems may yield evidence of bounded rationality or imprecise subjective probabilities among participating agents.

2 Markets as information aggregation devices

When markets attract broad participation, prices can encode the sum total of a large amount of disparate and distributed information. The prices reflect, in a very real sense, the collective opinion of a myriad of informed and well-motivated traders [20, 21]. Informative prices often translate directly into accurate forecasts of future events. For example, prices of financial options are good probability assessments of the future prices of the underlying assets [26]; prices in political stock markets, like the Iowa Electronic Market (IEM), can furnish better estimates of likely election outcomes than traditional polls [10, 11]; odds in horse races, determined solely by how much is bet on which horses, match very closely with the horses' actual frequencies of winning [27, 30]; and point-spread betting markets yield unbiased predictions of sporting event outcomes [12].

2.1 Securities markets and no arbitrage

Almost all economic theories of equilibrium assume, at a minimum, that equivalent portfolios are priced consistently with one another, such that arbitrage opportunities do not exist [2, 8, 17, 29].[6]

An example arises in the context of a *securities market*. In the parlance of economic theory, a *security* is defined as a lottery ticket that pays off \$1 if some uncertain event A occurs, and pays off nothing if A does not occur.[7] For example, the owner of a security "\$1 if and only if (iff) it rains tomorrow" will be paid \$1 if it rains tomorrow, and nothing otherwise. In general, we use $\langle A \rangle$ as shorthand for the security "\$1 iff A". Now imagine a market of two disjoint and exhaustive securities: "\$1 iff it rains tomorrow" and "\$1 iff it does not rain tomorrow". Owning both securities guarantees the holder a payoff of exactly \$1 regardless of whether it rains. Thus the total price to buy both securities should never dip below \$1—otherwise, the buyer can obtain a risk-free profit. Similarly, in the absence of arbitrage, the total price to sell both securities can never exceed \$1. More generally, the prices of a collection of such securities must conform to a legal probability distribution (modulo the

[5] `http://www.ideosphere.com/`

[6] Pareto efficiency, a common and mild assumption, implies no-arbitrage.

[7] Insurance contracts, futures, options, derivatives, and even stocks can be modeled as portfolios of such atomic securities.

bid-ask spread), otherwise arbitrage is possible. The proof follows from the same argument de Finetti used in his famous *no Dutch book* justification for the existence of individual subjective probabilities [7].

A *conditional security* $\langle A_1|A_2\rangle$ pays off *contingent* on A_1 and *conditional* on A_2. That is, if A_2 occurs, then it pays out exactly as $\langle A_1\rangle$; on the other hand, if $\bar{A}_2$ occurs, then the bet is called off and any price paid for the security is refunded [7]. In an efficient (arbitrage-free) market, prices of conditional securities must also adhere to the laws of probability. So, for example, given three securities $\langle A_1 \wedge A_2\rangle$, $\langle A_1|A_2\rangle$, and $\langle A_2\rangle$, there must be prices $p^{\langle A_1 \wedge A_2\rangle}$, $p^{\langle A_1|A_2\rangle}$, and $p^{\langle A_2\rangle}$ within the three corresponding bid-ask ranges such that $p^{\langle A_1 \wedge A_2\rangle} = p^{\langle A_1|A_2\rangle} * p^{\langle A_2\rangle}$.

2.2 Forecast accuracy: Rational expectations theory

According to the theory of *rational expectations* (RE), securities prices are not only coherent (i.e., form a valid probability distribution), but are also accurate forecasts, reflecting the sum total of all information available to all market participants [13, 16]. Even when some agents have exclusive access to inside information, prices equilibrate exactly as if everyone had access to all information. The procedural explanation is that prices reveal to the ignorant agents any initially private information; that is, agents learn by observing prices.

Several studies demonstrate that, in a laboratory setting, securities markets are often able to aggregate information correctly, as postulated by RE theory [9, 23, 24, 22]. Beyond the controlled setting of the laboratory, empiricists have analyzed the forecast accuracy of public markets. Perhaps the most direct tests involve sports betting markets. Several studies demonstrate that odds on horses correlate well with the actual frequencies of victory at the track [27, 30]. Other sports betting markets, like the National Basketball Association point spread market, provide very accurate forecasts of likely game outcomes [12].

The Iowa Electronic Market (IEM)[8] supports trading in securities tied to the outcome of political and financial events. Their 1988 market, open only to University of Iowa students and employees, offered securities that paid off proportionally to the percentage of votes received by various candidates in that year's US Presidential election. The final prices matched Bush's final percent margin of victory more closely than any of the six major polls [10]. Since opening to the public, subsequent US Presidential election markets have attracted wide participation and following. Other election markets have now opened in Canada[9] and Austria.[10]

3 Markets in satisfiability

Satisfiability (SAT)—the problem of determining whether a propositional logic sentence has a satisfying instantiation—is the canonical NP-complete problem, and has found use in numerous applications, ranging from circuit design to theorem proving. Due mostly to its central role in logic, SAT has received much attention in the artificial intelligence community throughout the years.

I propose opening an online market of securities that pay off contingent on the discovery of valid instantiations to SAT problems. For example, suppose that one unit of a security pays off $1 iff a solution to a particular SAT instance is found by midnight EST tonight. The problem is posted on the Web in a standard format; the submission and verification of solutions are entirely automatic. As soon as a valid solution is received, anyone owning the security earns $1 per unit bought; if no solution is received by the expiration time, security holders get nothing. Conversely, anyone that short sold the security gets $1 if a solution is *not* found and nothing if it is. Before a solution is received and before the expiration time, traders buy and (short) sell the security based on their expectation of the

[8] `http://www.biz.uiowa.edu/iem/`
[9] `http://esm.ubc.ca`
[10] `http://ebweb.tuwien.ac.at/apsm/`

eventual outcome. The current price can be thought of as a collective assessment of the probability that a solution will be found.

The owners of good SAT solvers would be the most direct beneficiaries of such a market. Agents from around the world could devote excess CPU time to solving the problem. If someone finds a solution, he or she would buy up large quantities of the security and then submit the solution; similarly, if someone proves the problem unsatisfiable, he or she would sell *en masse*. Note, however, that trading in a securities market is a zero-sum game: in order to buy a security for the chance to make 1-p dollars (or lose p dollars), someone else must be willing to sell the same security for the chance to make p dollars (or lose 1-p dollars). Whatever one agent loses, other agents gain the same.

Finding a solution when the ask price is less than \$1 (or proving insolubility when the bid price is greater than \$0) is like discovering an arbitrage opportunity. If all participants traded only based on such risk-free arbitrage opportunities, then no two agents would ever agree to trade (assuming everyone's programs were correct). Liquidity in the market depends on the participation of *speculators* and *subsidizers*. Speculators are agents that do not have a solution (or an unsatisfiability proof) in hand, but trade anyway based on their *expectation* of earning money given their assessment of the probability that a solution will be found. Note that owners of good incomplete SAT algorithms or heuristic SAT algorithms may make successful speculators. Subsidizers are agents that buy and sell simply to encourage trading. Subsidies are incentives for algorithm developers to join the market, improve their algorithms, and solve the given problems. Subsidies may come from governments, universities, or companies with an interest in solving the particular problems at hand, or in simply funding SAT algorithms research and development. For example, a semiconductor maker with a library of hard SAT problems whose solution would help improve their circuit designs may contribute both problems and subsidies to the market.[11] Government agencies can feel fairly confident that their subsidies will go to the best algorithm developers, rather than to the best research salespeople.

Note that the source for SAT instances must be trusted, since it is possible to generate instances with a known solution that are very difficult to solve *a priori*.

4 Satisfiability counting

Now consider the following SAT market variant: instead of paying off iff a solution is found, securities pay off iff at least one of k *random* instantiation of the variables renders the sentence true. Then the probability of a \$1 payoff is equal to $1-(1-s/n)^k$, where s is the number of of satisfying instantiations and n is the number of total possible variable instantiations.[12] If the problem is unsatisfiable, then a bid price greater than \$0 still constitutes an arbitrage opportunity; similarly, if the problem is a tautology, then an ask price less than \$1 is incoherent. However, if the problem is neither unsatisfiable nor a tautology, and the price is between \$0 and \$1, then there are no opportunities to earn a risk-free profit, only opportunities to earn an expected (non-arbitrage) profit.

Such a market could encourage the development of good SAT *counting* algorithms. If the ask price is less than $1-(1-s/n)^k$, then the security is worth buying (assuming risk neutrality); if the bid price is more than $1-(1-s/n)^k$, then the security is worth (short) selling. Even if an agent cannot compute s exactly, as long as it can compute bounds on s (i.e., prove that there are at least s_l or at most s_u solutions), it can submit safe bids and asks (in terms of expected profit). Approximation algorithms can also yield reasonable bids and asks. The purpose of using a value of k greater than one is to discourage simple stochastic simulation algorithms that can compute s/n to very high precision. I envision k to be large enough so that the time required to establish the truth or falsity of the claim is much greater than the time allotted to solving the problem (e.g., 48 hours to test k random instantiations vs. one hour for market participants to try to solve for s). Figure 1 shows the relationship between expected payoff and s for $k = 10^9$ and $n = 10^{10}$. The current bid-ask spread

[11] The company may remain anonymous and the problem descriptions randomized.

[12] If the instance is a near-tautology, it would make more sense to pay of iff at least one of k random instantiations renders the sentence false.

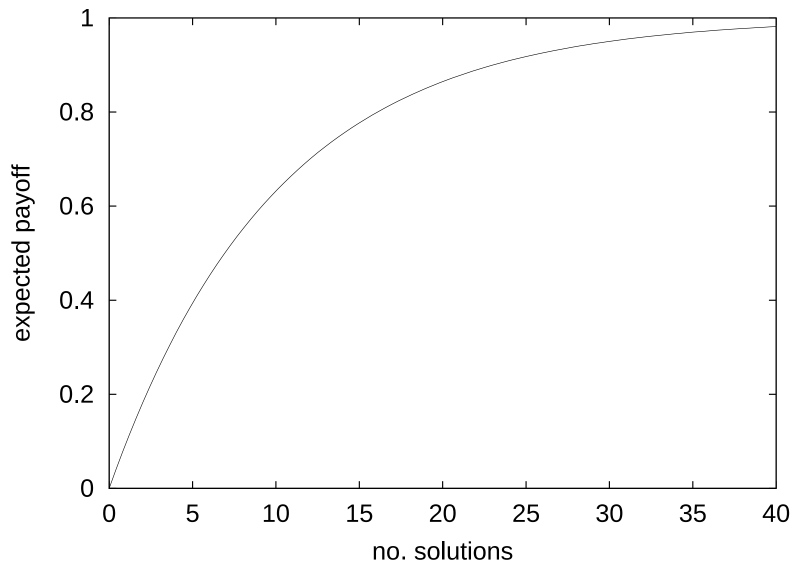

Figure 1: Expected payoff $1 - (1 - s/n)^k$ versus s when $k = 10^9$ and $n = 10^{10}$ in a SAT counting market.

can be considered the collective assessment of bounds on $1 - (1 - s/n)^k$, providing a distributed approximate solution to the given counting problem. Subsidies can again add liquidity to the market.

Note that agents do not need to consider the likelihood of someone else solving the problem, as they do in the SAT markets of Section 3. Agents need only concentrate on analyzing the problem itself. On the other hand, SAT counting is #P-complete, so there is no (known) way to verify solutions in polynomial time. Thus the markets yield only approximate and unverifiable solutions, unlike those in Section 3.

The source of problems and randomization must be trusted. One interesting possibility would be to tie the variable instantiations to events in the world. For example, the polarities of variables could depend on the outcomes of sporting events, perhaps incorporating point spreads to insure a roughly 50-50 split.

4.1 Measuring difficulty

As mentioned, the bid-ask spread encodes bounds on $1 - (1 - s/n)^k$. A larger bid-ask spread may signal a more difficult problem, since agents are unable to zero in on s. Then the market gives not only a bound on the number of solutions, but also an estimate of problem difficulty.

Prices in the basic SAT market of Section 3 also reflect problem difficulty, since if the problem is too hard to solve, the security will cash out at $0. However, it is nearly impossible to untangle the influence of problem difficulty from the potential that the problem is unsatisfiable. On the other hand, if a trusted source specifically generates problems with exactly one solution, all bidders know that there is a solution, and none of them know the solution itself, then prices below $1 only reflect uncertainty in whether the solution will be found.

4.2 Imprecise probabilities

Most axiomatic justifications for subjective probabilities insist that agents maintain *point* probabilities. So if an agent's probability that a coin turns up heads is 0.5, then it will be willing to buy at least some small amount (perhaps much less than one unit) of a security "$1 iff heads" for $0.499, and it will be willing to sell some small amount for $0.501. And if the same agent's probability that a stock price will go up is 0.5, it will do exactly the same, even though one might imagine that the agent is more confident in the probability assessment for the coin. Other theories distinguish between uncertainty and ignorance, so that an agent may be willing to bet more on the coin than on the stock. Many researchers study *imprecise probabilities*, where agents maintain probability intervals to encode "uncertainty ambiguity" instead of singular point estimates [6]. If a SAT counting market exhibited varying bid-ask spreads based on problem difficulty, that might constitute a form of empirical (if indirect) evidence for the existence of imprecise probabilities—and thus irrationality—among participating agents.

5 Bayesian inference

Bayesian inference is #P-complete and is intimately related to the SAT counting problem [5]. Bayesian inference markets might operate as follows. An instance of a Bayesian network is posted on the Web in a standard format. A security is then offered for a particular inference query; for example, $\langle A_1 | A_2 \wedge \bar{A}_3 \rangle$. Payoffs are determined based on either one or k random instantiation of all network variables, in accordance with the given conditional probability tables. According to the semantics of conditional securities, if the evidence variables in the query ($A_2 \wedge \bar{A}_3$ in the example) do not occur, then all buyers and sellers receive refunds precisely as if trading had never taken place. If the evidence variables do occur, then the security pays off $1 iff the non-evidence variable in the query (A_1 in the example) occurs.

Similar to SAT counting markets, if the bid (ask) price is greater than (less than) the exact probability of the query (perhaps modified according to k), then selling (buying) the security is worthwhile for risk-neutral rational agents. Then the bid-ask spread will tend to converge around the true probability, providing approximate bounds for the inference problem. Again, the size of the spread may correlate with problem difficulty.

5.1 Speed of evidence incorporation

An alternative way to handle evidence would be to incorporate it into the instantiation step: when variables in the network are instantiated, evidence variables are forced to their given state, and all bidders know this in advance. The query security is a standard security rather than a conditional one. Continuing with the above example, the query security is just $\langle A_1 \rangle$ and variables are instantiated stochastically under the constraint that A_2 is true and A_3 is false. This approach eliminates the possibility that all bets are canceled, though only allows one evidence set to be evaluated at a time.

It is also possible to announce evidence instantiations in the middle of trading, forcing (rational) agents to revise their expectations, and thus changing the price of the query security to reflect the new evidence. Economists are interested in how quickly evidence is incorporated into market prices (the *efficient markets hypothesis* asserts that it is essentially instantaneous). However in most markets, it is very difficult to control only for changes in evidence, since so many other factors are at play. Bayesian inference markets may provide an appropriate environment in which to examine such questions in a very controlled manner. In addition, the speed of evidence incorporation may depend on the difficulty of the inference problem, constituting evidence of bounded rationality among agents in the market.

6 More variations and other issues

6.1 Solver takes all

In the SAT markets of Section 3, all buyers receive a payoff regardless of who actually solves the problem. An alternative framework would be to pay only the first agent to submit a solution. Presumably only sellers who were unlucky enough to transact with the winning agent would lose money. It's not clear how this would affect prices or the interpretation of prices. Perhaps prices would fall, since sellers don't necessarily lose money when a solution is found, so would tend to sell more. Lower prices would make it a more attractive prospect for buyers; however, buyers don't necessarily win even when a solution is found.

Another possibility is simply to allow subsidizers to contribute to a "pot" for each problem, with the entire pot going to the first agent to solve the problem. This variation is simply a contest where payoffs depend on subsidizers' interests.

6.2 Optimization problems

Markets in optimization problems like the traveling salesman problem are also possible. For example, securities of the form "$1 iff a tour of length less than K is found" would operate just as SAT markets. Other possibilities include:

- only the first agent to submit a tour of length less than K is paid,
- all agents that submit tours of length less than K are paid,
- only the agent that submits the best solution is paid, and
- only agents that submit solutions improving on the current best submission are paid.

Again, the effect on price under each of these variations is unclear. Perhaps some types of markets are easier to subsidize than others. Or perhaps some provide greater incentives for algorithm developers, or are more attractive to speculators.

6.3 An agent playground

It is likely that people would interact with the market mostly via agent surrogates. Agent programs would download problems daily, attempt to solve the problems (perhaps during otherwise idle computer time), and transact in the markets accordingly. There is room for both simple trading strategies that exploit arbitrage opportunities only, and more sophisticated (and speculative) strategies that attempt to maximize expected utility across multiple markets.

6.4 Play-money markets

Throughout this paper we have assumed that payoffs are denominated in real money (e.g., US dollars), thus providing monetary incentives to agents. However, significant regulatory and legal hurdles would have to be overcome before any real-money NP markets could be established. A play-money version would be much easier to set up and operate in the short term. In this case, incentives would presumably derive from entertainment value, educational value, bragging rights, and/or other intangible sources. Our recent study [18, 19] suggests that intangible rewards can (to some extent) still drive information aggregation and forecast accuracy in markets. So play-money NP markets might afford some of the same benefits as real NP markets, at least until permission can be obtained for the latter.

7 Related ideas

Several researchers are investigating the possibility of solving hard problems by coordinating distributed computers across the Internet. For example, NASA developed a screen saver application that pulls SETI data from NASA's Web site, analyzes it locally, and sends back the results.[13]

Researchers have also proposed opening market in CPU cycles. Examples include the Popcorn Market [25], the Compute Power Market [4], and the Java Market [1]. NP markets would in a sense solicit CPU cycles to solve particular problems of interest (and encourage algorithm development), rather than sell CPU cycles for arbitrary private tasks.

Brewer [3] devises another interesting scheme to incent distributed agents to solve an NP-complete optimization problem—in this case, to compute a Pareto optimal allocation in a combinatorial auction. Traditionally, the auctioneer carries out the computation. Brewer suggests instead a procedure where bidders in the auction must settle for the current (possibly Pareto dominated) allocation unless someone demonstrates that a better solution, with smaller total consumer surplus, exists. Bidders whose utilities increase in the improved solution have incentive to report the solution if known. Bidders may also receive a percentage of the improvement when they report a better solution, with percentages increasing as the auction's clearing time approaches.

8 Conclusion

I discussed the promise (and some potential pitfalls) of opening markets in hard computation problems. I described markets in SAT problems, SAT counting problems, and Bayesian inference problems, though analogous markets could be opened in almost any NP-complete or #P-complete problem. The markets serve as a vehicle for subsidizing algorithms research, where rewards go directly to the best performers instead of the best solicitors. Prices in the markets are collective approximate solutions, and bid-ask spreads signal problem difficulty. Certain markets seem well suited for empirical studies of the speed of information aggregation, bounded rationality, and imprecise probabilities.

Acknowledgments

Thanks to Michael Wellman for ideas, insights, and encouragement.

References

[1] Yair Amir, Baruch Awerbuch, and Ryan S. Borgstrom. The Java Market: Transforming the Internet into a metacomputer. Technical Report CNDS-98-1, Johns Hopkins University, 1998.

[2] Kenneth J. Arrow. The role of securities in the optimal allocation of risk-bearing. *Review of Economic Studies*, 31(2):91–96, 1964.

[3] Paul J. Brewer. Decentralized computation procurement and computational robustness in a smart market. *Economic Theory*, 13:41–92, 1999.

[4] Rajkumar Buyya and Sudharshan Vazhkudai. Computer Power Market: Towards a market-oriented grid. In *First IEEE/ACM International Symposium on Cluster Computing and the Grid (CCGrid)*, 2001.

[5] G. Cooper. The computational complexity of probabilistic inference using Bayes belief networks. *Artificial Intelligence*, 42:393–405, 1990.

[13] `http://setiathome.ssl.berkeley.edu/`

[6] Gert de Cooman, editor. *First International Symposium on Imprecise Probabilities and Their Applications*. 1999.

[7] Bruno de Finetti. *Theory of Probability: A Critical Introductory Treatment*, volume 1. Wiley, New York, 1974.

[8] Jacques H. Dreze. Market allocation under uncertainty. In *Essays on Economic Decisions under Uncertainty*, pages 119–143. Cambridge University Press, 1987.

[9] Robert Forsythe and Russell Lundholm. Information aggregation in an experimental market. *Econometrica*, 58(2):309–347, 1990.

[10] Robert Forsythe, Forrest Nelson, George R. Neumann, and Jack Wright. Anatomy of an experimental political stock market. *American Economic Review*, 82(5):1142–1161, 1992.

[11] Robert Forsythe, Thomas A. Rietz, and Thomas W. Ross. Wishes, expectations, and actions: A survey on price formation in election stock markets. *Journal of Economic Behavior and Organization*, 39:83–110, 1999.

[12] John M. Gandar, William H. Dare, Craig R. Brown, and Richard A. Zuber. Informed traders and price variations in the betting market for professional basketball games. *Journal of Finance*, LIII(1):385–401, 1998.

[13] Sanford J. Grossman. An introduction to the theory of rational expectations under asymmetric information. *Review of Economic Studies*, 48(4):541–559, 1981.

[14] Robin Hanson. Decision markets. *IEEE Intelligent Systems*, 14(3):16–19, 1999.

[15] Robin D. Hanson. Could gambling save science? Encouraging an honest consensus. *Social Epistemology*, 9(1):3–33, 1995.

[16] Robert E. Lucas. Expectations and the neutrality of money. *Journal of Economic Theory*, 4(2):103–24, 1972.

[17] Andreu Mas-Colell, Michael D. Whinston, and Jerry R. Green. *Microeconomic Theory*. Oxford University Press, New York, 1995.

[18] David M. Pennock, Steve Lawrence, C. Lee Giles, and Finn Årup Nielsen. The power of play: Efficiency and forecast accuracy in web market games. Technical Report 2000-168, NEC Research Institute, 2000.

[19] David M. Pennock, Steve Lawrence, C. Lee Giles, and Finn Årup Nielsen. The real power of artificial markets. *Science*, 291:987–988, February 9 2001 (Letters).

[20] David M. Pennock and Michael P. Wellman. Representing aggregate belief through the competitive equilibrium of a securities market. In *Thirteenth Conference on Uncertainty in Artificial Intelligence*, pages 392–400, Providence, RI, USA, July 1997.

[21] David M. Pennock and Michael P. Wellman. Compact securities markets for Pareto optimal reallocation of risk. In *Sixteenth Conference on Uncertainty in Artificial Intelligence (UAI–2000)*, pages 481–488, July 2000.

[22] C. R. Plott, J. Wit, and W. C. Yang. Parimutuel betting markets as information aggregation devices: experimental results. Technical Report Social Science Working Paper 986, California Institute of Technology, April 1997.

[23] Charles R. Plott and Shyam Sunder. Efficiency of experimental security markets with insider information: An application of rational-expectations models. *Journal of Political Economy*, 90(4):663–98, 1982.

[24] Charles R. Plott and Shyam Sunder. Rational expectations and the aggregation of diverse information in laboratory security markets. *Econometrica*, 56(5):1085–1118, 1988.

[25] Ori Regev and Noam Nisan. The popcorn market—an online market for computational resources. In *First International Conference on Information and Computation Economies*, Charleston, SC, 1998.

[26] Bruce J. Sherrick, Philip Garcia, and Viswanath Tirupattur. Recovering probabilistic information from options markets: Tests of distributional assumptions. *Journal of Futures Markets*, 16(5):545–560, 1996.

[27] Richard H. Thaler and William T. Ziemba. Anomalies: Parimutuel betting markets: Racetracks and lotteries. *Journal of Economic Perspectives*, 2(2):161–174, 1988.

[28] Alan M. Turing. Computing machinery and intelligence. *Mind*, 59:433–560, 1950.

[29] Hal R. Varian. The arbitrage principle in financial economics. *J. Economic Perspectives*, 1(2):55–72, 1987.

[30] Martin Weitzman. Utility analysis and group behavior: An empirical study. *Journal of Political Economy*, 73(1):18–26, 1965.