# Zocalo: An Open-Source Platform for Deploying Prediction Markets

Chris Hibbert

CommerceNet Labs Technical Report 05-02
February 2005

# Table of Contents

# Executive Summary

We propose developing an open-source toolkit for creating markets, called *Zocalo*, in order to catalyze broader adoption of markets in academia, industry, and throughout society. We are primarily interested in *prediction* markets, which allow traders to buy and sell securities that pay out based on the outcome of some future event, but Zocalo would also be useful for creating markets in other goods.

Markets are the most effective mechanism for assigning resources to their most valuable uses without relying on a single agent with complete information. [Hayek] They are mostly known as a human-scale institution for transferring control of physical goods ranging from a single sandwich to a tanker of crude oil. However, economists have shown over the last several decades that another function of markets — making prices visible, or "*Price Discovery*" — is valuable in its own right as mechanism for information aggregation.

By constructing financial securities that pay the holder an amount that depends on the outcome of particular events, markets in such artificial commodities generate forecasts that have been shown to improve on other predictions whenever buyers and sellers have access to information that bears on the likelihood of the events. More accurate predictions are valuable whenever we have to make choices in the face of uncertainty.

There is an active body of research into the structure and effects of these markets. Applications include forecasting (project schedules, level of sales, election outcomes) and making decisions (which pharmaceuticals to take to trial, which flu vaccine to manufacture this season, what mix of products to build). Trials of prediction markets are beginning to demonstrate their usefulness, but proponents are handicapped by not having access to tools that make them straightforward to use.

Zocalo will fill that gap.

To succeed, it must be easy to deploy in a variety of environments, and must support a variety of market institutions that can be configured through a simple interface. If we achieve these goals, it will enable a wider community to investigate the effects of innovative market institutions. Our initial users will be experimental economists, because that is a forgiving environment, but one that will push us to build the market infrastructure on a principled foundation. Later applications will expand to include uses from trials in industrial settings to public markets for prizes or real money.

CommerceNet is ideally positioned to catalyze creation of a community around these ideas that will, we believe, accelerate their adoption.

# 1. Catalyzing Adoption of Prediction Markets

> A market designed from the outset for information gathering and forecasting is called an *information market*. Information markets can be used to elicit a collective estimate of the expected value or probability of a random variable, reflecting information dispersed across an entire population of traders. The market prediction is not usually an average or median of individual opinions, but is a complex summarization reflecting the game-theoretic interplay of traders as they obtain and leverage information, and as they react to the actions of others obtaining and leveraging their own information, etc. In the best case scenario, the market price reflects a forecast that is a perfect Bayesian integration of all the information spread across all of the traders, properly accounting even for redundancy.
>
> — Announcement for the DIMACS Workshop on Markets as Predictive Devices [DIMACS]

This paper proposes that CommerceNet fund an initiative to build Zocalo,[1] an open-source toolkit for markets. We are primarily interested in prediction markets, but the same tools should be useful for creating markets in other interchangeable goods. To be successful, the toolkit must be easy to deploy in a variety of environments, and must support a variety of market institutions that can be configured through a simple interface.

Our goal is to enable broader adoption of prediction markets in academia and industry in the near term. In the long-term, we hope this will lead to more use throughout society to support public-policy decisions.

The existing toolkits we have found [Walras, Gambit, JASA, jMarkets, *etc*.] support experiments but aren't intended for deployment as actual markets where valuable goods are exchanged. Without consummating exchanges, they merely report the results of abstract transactions. Each one emphasizes a particular format of market and depends on its own built-in interface.

Zocalo will be designed to support markets for goods with actual value, and will thus emphasize robustness beyond the minimal level expected of experimental software. It will support a wide range of goods and a variety of market formats. Standard software engineering practice, which is also not typically emphasized in software developed by domain specialists, can ensure that the user interface is separate from the underlying model.

Zocalo will support real markets that transfer ownership and control of goods (see [WebMart] for an example), so more attention will be focused on implementing the goods to be transacted — currency, coupons of various types, securities and derivatives — and ensuring that the ownership of the goods changes hands as the markets direct. This should pay off because new goods will have interfaces that would be compatible with those built previously, so the design of the basic institutions won't have to change as new combinations are added.

---

[1] *Zocalo* is the term used in Mexico for the open public market in a town's square. It plays the same role as the *agora* did in ancient Greece.

Our first versions will be suitable for experimental situations in economics labs. Later on we expect to be able to support internal prediction markets for business partners, and eventually, we expect the software to evolve to be suitable for real-money public policy prediction markets. As an open source toolkit, other developers will be able to make use of whichever portions of the code base they find to be useful.

CommerceNet is in an ideal situation for sponsoring this development. CommerceNet Foundation is a non-profit, and it is using some of its funds to start a new research lab focused on decentralizing ecommerce. This gives us the opportunity to invest in developing software systems that will enhance commerce in new ways. From this position, we hope to be able to recruit collaborators who are interested in operating markets.

# 2. Understanding Prediction Markets

Markets have long been considered to be the best way to allocate resources among competing uses in the absence of an arbiter with complete information [Hayek]. Prices ensure that the different agents competing for access have a common standard for comparison across different choices. The market clearing process ensures that each resource is assigned to its best use.

Different market designs satisfy different purposes [Reynolds]. Continuous double auction markets provide goods on demand to buyers who are willing to pay the going rate, while call markets improve prices for buyers and sellers when time is not the most crucial factor. When buyers' needs for different goods are interdependent, combination markets may be necessary. For example, someone may want to reserve airline tickets and hotel rooms at the lower advance-purchase, no-changes rates, but only if both are available at the same time.

Prediction markets are a relatively new form of market in which the goods being traded are securities whose values will be determined by the outcomes of future events. The securities are structured so that trading between buyers and sellers causes the price to reflect the probability of the underlying event. When a trader sees a market price (probability) that is less than her expected probability for the event, she will see a profit opportunity in buying more, thus driving the price up. The new price reflects a higher probability to others monitoring the market.

## 2.1 Prediction Markets are an Active Research Area

Since prediction markets were first proposed in 1988, they have attracted a fair amount of attention ([Wolfers & Zitzewitz A] provides a survey of recent work). Several variations on the basic idea have been described that ought to make it possible to find answers to more kinds of questions, or apply the institution in a wider set of circumstances. Robin Hanson [Hanson2003] suggested ways to support conditional bets, which would elicit estimates about the relative chances that alternative approaches would have of achieving certain outcomes. [Pennock] suggested combining pari-mutuel style betting with standard double-auction prediction markets to increase the liquidity of the markets at the cost of having prices that don't directly reflect probabilities. Several automated market makers have been proposed that would offer different ways to improve the predictive ability of markets with small numbers of traders ("thin markets"). In markets for general commodities, similar approaches can make thin markets more liquid and make it easier for traders to find counterparties to trade with.

So far, the public markets have been limited to questions of very broad interest like national elections and the outcome of professional sporting events. The new institutions that have been proposed should make it possible to offer a wider variety of questions and provide greater detail in the answers. We suspect that a large part of the reason these improved institutions are not being used is that independent development by each provider has meant that each developer has to spend most of their time solving the first order problems, and they never find the time to add more institutional variants. Our hope is that shared access to open source will provide incentive and opportunity to get beyond the first stage of system development.

## 2.2 Public Prediction Markets are Difficult to Deploy

One of the obstacles to the introduction of prediction markets in the U.S. has been that the gambling laws in most jurisdictions are written broadly enough that they could be construed to prohibit use of these markets with real money ([Bell, Hahn & Tetlock 2004b]). Few people have been willing to risk the possible penalties without clearer indications that these markets won't attract prosecution. The fact that the laws vary from one state to the next makes the legal research prohibitively expensive. A recent decision by the World Trade Organization states that the American ban is an unfair trade practice. The U.S. is fighting the ruling, and may not actually have sufficient incentive to force it to comply.

In this environment there have been a variety of alternatives introduced in order to explore the ideas without facing the unknown legal risks. Some of these have been "play money" markets [HSX, FX] or limited the amount participants could risk [IEM], others have awarded prizes but not required participants to put up money to join [Time, NewsFutures], and one awards the winnings to the winner's preferred charity rather than to the winner [LongBets]. There is continuing discussion of whether play money produces results that are as reliable as real money markets [Schrieber].

# 3. Prediction Markets Aid Decision-Making

[Berg & Rietz p80] cite prior research supporting six reasons why prediction markets are a good choice for Decision Support applications:

1) They "give continuously updated dynamic forecasts."
2) They "aggregate information across traders."
3) These "markets give unbiased, relatively accurate forecases well in advance of outcomes."
4) They "can outperform existing alternatives."
5) "market dynamics overcome biases that individual traders may have, effectively eliminating [the biases] from forecasts."
6) They "can be designed to forecast a variety of issues and provide a variety of types of information."

These attributes fit well with zLab's focus on decentralizing e-commerce. Prediction markets can enable "agile networks of firms [to] collaborate dynamically", as proposed in the introduction to zLab's mission statement. [zLabs] They can also be used as critical tools within companies to allow the business to react more spontaneously to changing conditions before information has a chance to percolate to the traditional decision makers. In addition, they will constitute an important component of some of the institutions (decentralized advertising markets and peer-to-peer auctions) that zLabs is pursuing.

## 3.1 Markets Aid Decentralized Communication

When an organization has internal expertise that isn't being reflected in the information that reaches the decision makers, this new tool can draw together the knowledge that is spread around the organization, and circumvent the incentives that sometimes cause official channels to distort messages. Some managers pad their estimates so they aren't squeezed when project requirements change, while others shave the numbers they provide to make proposals more attractive, even if less reliable. Executives often have to guess which distortions have been applied and don't know whether multiple layers of intermediaries have compounded the problem or cancelled it out.

Pharmaceutical companies, for example, have many people with considerable expertise on all the different factors relating to successful introduction of a new drug. They spend much of their time presenting much of it to other people in the organization. But every time the information is repeated, it has to be summarized or rearranged to fit the constraints of the presenter and the audience, diluting or filtering the information content.

The result is that even when everyone means well, executives can't rely on the information that reaches them. Several experiments have demonstrated that prediction markets can remedy this problem. The examples cover corporate sales performance, efficacy for drugs that are still in trial, and likelihood of software development projects meeting their delivery schedule.

## 3.2 Corporations are Adopting Markets

Markets can be used to find the most efficient ways to allocate resources or work among different divisions, even when the resources are exotic or no one is sure to begin with what a reasonable approach would be.

> There have been several publicized uses of formal markets inside businesses [Time, Chen & Plott, Proebsting]. One, at British Petroleum [Malone], allowed divisions of the company to trade internal pollution credits (which each allowed a certain amount of $CO_2$ emissions) in a market that encouraged decentralized coordination on the most cost effective ways to reduce BP's overall output of $CO_2$.

In 1998, BP's CEO had made a commitment to reduce the company's emissions by 10 percent from 1990 levels by 2010. The divisions were each assigned a certain number of permits, and they were allowed to buy and sell permits among the divisions. Each year, the divisions were constrained to not emit more $CO_2$ than their respective permits allowed. Each division had a mandate to reduce their emissions, so they each explored various ways of reaching their targets.

As they developed new ways to make the changes, they could publish to the other divisions what it would cost to reduce emissions by a certain amount, and how much they could reduce beyond what their permits required. The different divisions could then compare results and shift the permits to the divisions with the highest cost to make changes, while shifting the reductions to the divisions with the lowest costs of change. The divisions could then also share notes on what strategies were most cost effective, and adopt each other's processes when they made sense. By 2001, 9 years ahead of the CEO's commitment, the company had met its original goals.

## 3.3 Corporations are Adopting Prediction Markets

In other cases, companies intended to improve on information they were already generating and relying on for planning purposes. All of these reports indicate that the information produced gave better answers than the mechanisms the companies had previously relied on:

- Hewlett-Packard (HP) ran markets predicting the level of sales expected for some of its product lines. The markets have consistently produced more reliable predictions than the official sales forecasts that management used to rely on. [Plott & Chen]

- Eli Lilly used prediction markets to evaluate candidate drugs. The market correctly picked out the three most successful drugs. The resulting price levels gave better indications of the traders strength of convictions than other mechanisms had. [Time]

- Siemens ran internal markets on whether software development projects would meet their schedules. The market gave early indications that development would not be finished by the time that management was projecting. [Ortner]

- Intel used markets to make some decisions about which manufacturing plants to assign particular production runs to. Microsoft was reported to be considering setting up markets to get more information on the level of interest among developers for new products. [Time]

The largest area of current use for prediction market software is among experimental economists. There is also a growing demand for these and other markets inside businesses.

## 3.4 Potential Long-Term Impact on Society

If we have developed the toolkits sufficiently, they will make it possible to run very small scale markets that nonetheless produce reliable information. This will enable promoters of these markets to take advantage of an effect identified by Wired Magazine [Anderson]. They described the competitive advantage of on-line electronic retailers as the ability to make enormous catalogues of goods available and organize them so that every customer can find a niche market customized to his tastes. The possible uses of these markets will expand enormously if we make it cost effective and valuable to build many markets covering large numbers of questions, some of which may be important only to a small number of people.

In a broader social context, Robert Hahn and Paul Tetlock [Hahn & Tetlock 2004a] have described in detail how prediction markets could be used to improve the effectiveness of aid to developing nations. Their proposal is that prediction markets should be used to estimate the comparative benefits of alternative approaches to resolving a particular problem. In parallel, the funder would auction off the right to be paid a specified performance bonus for implementing the proposed program. The two parallel markets give separate, unbiased estimates of the expected benefit and specify the amount the funding agency would have to pay. This allows them to make trade-offs between different approaches without having to assess the biases of experts who are likely to have an interest in which outcome is chosen.

Similar approaches could be used to produce cost-effective solutions to many other public policy challenges. [Leigh & Wolfers] In addition, the establishment of stable, legal prediction markets could provide access to ongoing, reliable estimates of the relative importance of widely discussed issues whose severity is currently debated by competing experts without resolution. DARPA proposed setting up a market to investigate the effectiveness of these ideas in foreign policy settings. [PAM]

# 4. Our Design Approach

*"Nothing seems as clear as a design you haven't written down."* — E. Dean Tribble

Our goal is to establish Zocalo as the obvious tool to start from when considering the deployment of a new market, whatever the purpose. As serious software developers with a background in providing reliable systems and services, we expect to be able to offer stability, adaptability, and continuity. These should be valuable both to experimental economists looking for a quick turnaround on a simple solution and to entrepreneurs who want to build a reliable market to allow participants to exchange items of real value.

Since we are claiming that design differences will make Zocalo more robust and reliable than other tools that practitioners might use, it behooves us to discuss the design approach that will lead to these benefits. The most important aspect of the design is that the market institutions will have polymorphic interfaces [Page-Jones] for dealing with the commodities being traded. This will ensure that new forms of goods and new institutions can be created and introduced into the exchange without major rewrites. This should enable experimenters to make minor (or major) changes to their experimental design without revamping the software in a way that renders previous experiments inoperable or requiring a complete rewrite for each new use of the toolkit.

There will be a clear separation between implementation of the market institutions and the user interfaces (UIs). The first UIs will be developed for deployment over the web. If there is any reason to develop versions that run as applications only the UI code will be affected. It is envisioned that versions will be developed that support both web interfaces and direct TCP/IP connections as well as interfaces for email. None of these possibilities will effect the initial implementations except in enforcing the clean separation between the market model and support for user interfaces. These clean interfaces will also allow software agents to interact with the markets. There is a significant amount of market research that relies on agents to provide liquidity or to investigate specific trading strategies.

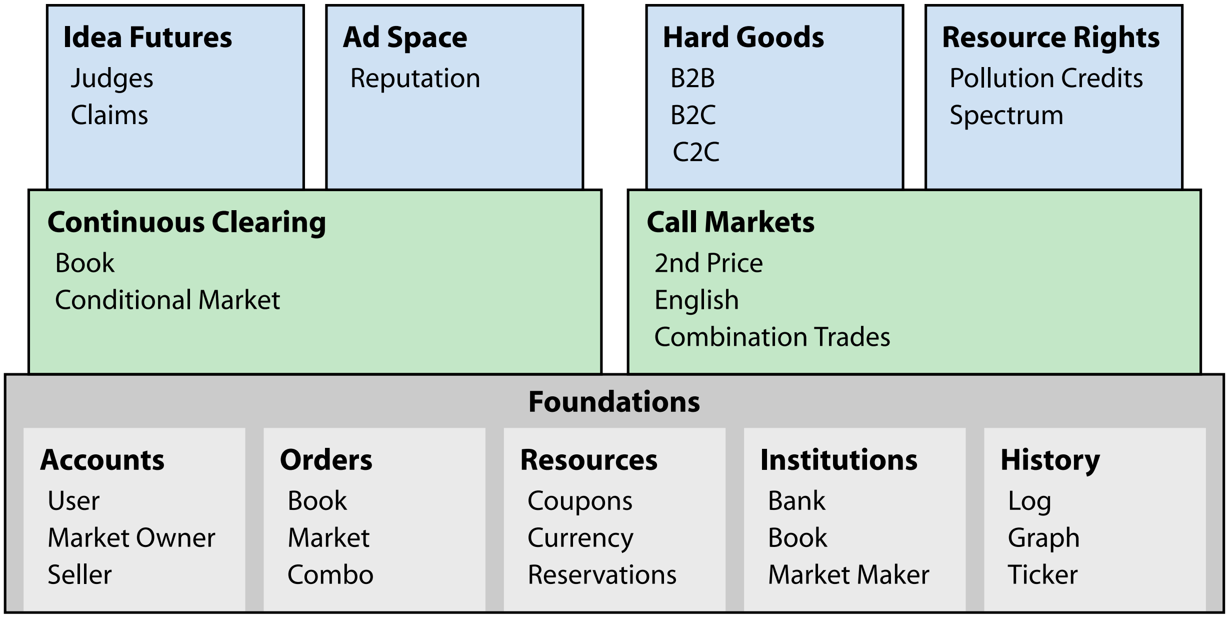


*Figure A: Zocalo Toolkit Architecture*

The foundation of the system is a layer of classes that support abstract trading in a market. There are objects representing the Resources being traded, their Owners, the Orders used to specify the desired trade, and the institutions (OrderBook, generic Auction, etc.) which process orders in order to cause trades to happen. Built on top of those foundation classes are more specific implementations (call market, continuous double auction) that use the same interfaces. These institutions will then need to have particular Resources specified in order to deploy a particular prediction market, or a market in rights to some abstract resource like pollution credits or network bandwidth.

The approach is based on previous work by the author, including prediction market prototypes and auction systems [WebMart, SMLI-TR, Miller] that supported dynamic markets in network bandwidth, charged for use of a high-end color printer, and metered access to a network jukebox. We have previously built two prototypes for prediction markets: one that integrates a combinatorial market with an order book, and another that integrates an order book with a market maker using a logarithmic scoring rule [Hibbert].

## 4.1 Our Development Process

We expect to start on this project by building a simple working version of a prediction market, sufficient to support modest experiments. We are discussing requirements with experimenters at George Mason University who expect to run an experiment in March, and are preparing to begin development. It will require simple versions of some of the classes in the foundation layer of figure A, and a few specific solutions situated along the left edge of the higher layers. We also hope to take advantage of the opportunity to develop software customized for particular experiments by inviting developers to work with us at CommerceNet where we will be able to work closely with them to learn from their expertise and show them how to take advantage of the software's features.

Once we have built the software foundations, refactoring [Fowler] to support different requirements should be straightforward. Object-Oriented Design using polymorphic objects [Page-Jones] readily lends itself to the Extreme Programming (XP) style of rapid prototyping and rapidly reconfiguring existing code to suit additional uses. When combined with XP's emphasis on test-driven development [Beck], it is common to be able continue to support all working features as new abilities are added incrementally.

The initial requirements for the software are to support:

- currency
- binary prediction market claims
- multiple accounts
- multiple markets
- a continuous double auction market based on an order book
- logging events to allow activity analysis

It will need user interfaces to support experimenters creating claims and accounts and allocating funds; as well as participants viewing prices, entering bids, and accepting offers.

Particular experiments will, of course, impose yet other requirements, such as limited access to information, existence of market makers or introduction of arbitrage agents to the system. Other directions the software will be able to adapt to quickly include:

- multiple accounts
- multi-way claims (extending beyond binary claims)
- more auction formats
- integrated order book
- market scoring rules
- conditional claims

Once the software has been demonstrated to support experiments, and the extensibility mechanisms have been documented, we expect to investigate approaches to decentralization of the system to increase redundancy and scaling in order to be able to support larger markets.

One goal is that the software support decentralized markets sponsored by different parties. We also expect to support markets in other commodities. One example of interest within CommerceNet's zLabs is building decentralized advertising networks ("zClassifieds") intended to allow publishers to regain control of the ads they display on their sites.

## 5. Collaborating Across Communities

We expect Zocalo to be valuable to several groups; some will benefit from the early uses, while others will see the gains as the project develops a significant collection of relevant software, and others demonstrate the value of the techniques. We expect economists, particularly those investigating the properties of different forms of prediction markets, to be among the early groups to find uses for our software. The existence of the software, as a demonstration of the feasibility, and as a foundation for deployment, should lower the cost of using markets within firms. The toolkit should find uses in a wider range of environments, including in public policy discussions that will have a larger effect on society, but may take longer to emerge. We hope to work with organizations that would like to explore these uses in order to ensure the tools are suitable in those environments.

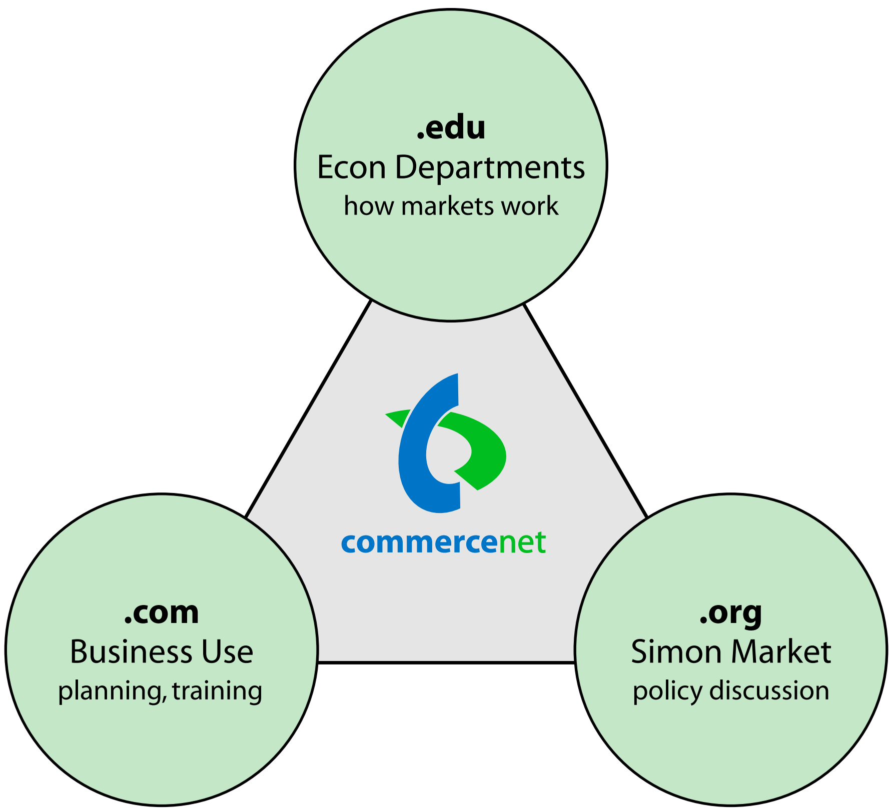


*Figure B: Our Target Audiences*

We have started discussions with experimental economists to find out whether they would be receptive to using Zocalo in their experiments, and how they currently develop software for their experiments. Some respected researchers have engaged in discussions with us of their requirements and current approach. Robin Hanson and David Porter's research group at George Mason University is planning an experiment for this spring, and has welcomed our offer to provide software. At the DIMACS Workshop on Prediction Markets [DIMACS], we talked with several other groups that were looking for alternative sources of software. Some are not currently able to run pilot projects because companies that support this kind of project are charging substantial fees. Experimenters that we contacted were often unsatisfied with the generic experimental software they are currently using and expressed serious interest in alternatives that would be more readily adapted to investigation of prediction markets.

From the few groups we have contacted, it appears likely that there are many departments who don't have access to a reusable toolkit for their experiments. These groups would be likely to welcome support from an organization that specializes in professional software development, and could provide extensible foundations that would reduce the costs and time required to develop tools for their experiments.

Over the longer term, the existence of this software as a focal point for research into the properties of prediction markets, and for development of advanced market mechanisms should drive deployment of prediction markets into more environments. The open source Apache web server platform, for example, makes it possible for many organizations and individuals to deploy sophisticated web sites at a reasonable cost. The large number of uses and users has attracted other developers to create commercial alternatives, which has not kept the usage of Apache from continuing to grow.

If studies of these markets continue to demonstrate that these markets contribute unique value, this will feed back into further interest in these techniques. In the medium term, this would lead to increased use of markets inside businesses and non-profit organizations to improve decision making and optimize the allocation of effort and resources. The internal markets pioneered by BP, HP, and Eli Lilly are indicative of some of the possibilities.

## 5.1 The Case for Sharing Infrastructure

Experimental economists have a continuing need to test new or familiar market institutions while varying the rules about participants' goals or what information they have access to. The NetLab Workshop sponsored by the NSF in 1997 [NetLab] called for development of shared software for experiments in economics and other social scientists to make it easier to test new hypotheses, but little seems to have come of that. The one example that we've found is that U. C. Berkeley has set up an Experimental Social Science lab [Xlab], that is open to students and professors across departments on campus. It is intended to make it easy for social scientists there to set up computer-based laboratory experiments. They are providing access to experimental economics software provided by HP Labs [Xlab Software, MUMS].

Another source of deployable software is the growing number of companies selling prediction market services[2]. So far, all of these companies are keeping their source code proprietary. Occasionally, an experiment can be run using public prediction markets as they are deployed for other purposes [Wolfers & Zitzewitz B], but the more normal case is that some modifica-

[2] Examples include: Common Knowledge Markets, Foresight Exchange, NewsFutures, IncentiveMarkets, and EconOne [CommerceNetWiki].

tion is necessary in order to test a specific hypothesis. The commercial companies are seldom willing to make changes within the budget of academic economists.

These companies are competing to provide software to support internal markets in the business context. We hope and expect that the fact that our code will be open source and designed for customizability, and that we are looking for test cases to prove those properties will make us a more viable choice.

# 6. Conclusion

Markets have been shown to be the best way to allocate resources in a large variety of situations. Indeed, arguments favoring alternative mechanisms often bear the burden of proof. Prediction Markets are not yet as established, but evidence that they are valuable tools is accumulating rapidly. We hope to catalyze the creation of a community around an easily deployable software system, both to improve research in this area, and to enable broader adoption in the situations in which their value has already been established.

CommerceNet is uniquely positioned to support this project. Other groups that are interested in using prediction markets are focused on exploring the effects of variations in market design or selling services to customers, and don't have an incentive to provide a general framework. As a result, every potential user of these tools must start by acquiring expertise in developing software. Since their interests are elsewhere, the software is not built in a way that lends itself to re-use, and the barriers to adoption remain high. CommerceNet could create a foundational toolkit that would be valuable to many people, and enable vastly increased use of this tool.

We have presented our approach to building open source prediction market software. We are looking for collaborators who would make use of the software for experiments in the economics of markets or the functioning of prediction markets. The contribution we hope to make to the field lies in expanding the use and usability of these markets so they will be deployed in more places. We hope to catalyze growth of a community interested in promoting the use of this tool

# 7. Bibliography


[Anderson] Chris Anderson, "The Long Tail", *Wired Magazine*, October 2004
http://www.wired.com/wired/archive/12.10/tail.html

[Beck] Kent Beck, *Test Driven Development: By Example*, Addison-Wesley, 2002

[Beck & Andres] Kent Beck and Cynthia Andres, *Extreme Programming Explained: Embrace Change*, Addison-Wesley, 2004

[Bell] Tom Bell, "Gambling for the Good, Trading for the Future: The Legality of Markets in Science Claims", 5 Chapman Law Review 159, 2002
http://www.tomwbell.com/writings/Gambling4Good.pdf

[Berg & Rietz] Berg, Joyce and Thomas Rietz, "Prediction Markets as Decision Support Systems", *Information Systems Frontiers*, 5(1), 2003
http://www.biz.uiowa.edu/iem/archive/decisionsupport.pdf

[Berg et. al.] Joyce Berg, Forrest Nelson, Robert Forsythe, and Thomas Rietz, 2003. March, "Results from a Dozen Years of Election Futures Markets Research"
http://www.biz.uiowa.edu/faculty/trietz/papers/iemresults.pdf

[Chen & Plott] Chen, Kay-Yut and Charles Plott, "Information Aggregation Mechanisms: Concept, Design and Field Implementation for a Sales Forecasting Problem", Social Science Working Paper No. 1131, March, 2002
http://www.hss.caltech.edu/SSPapers/wp1131.pdf

[CommerceNetWiki] http://labs.commerce.net/wiki/index.php?title=PredictionMarketLinks

[DIMACS] "Workshop on Markets as Predictive Devices (Information Markets)", Rutgers University, February 2-4, 2005
http://dimacs.rutgers.edu/Workshops/Markets/announcement.html

[Economist] *The Economist*, "Guessing Games", Nov 18th 2004
http://www.economist.com/printedition/displayStory.cfm?Story_ID=3400241

[Fowler] Martin Fowler, *Refactoring: Improving the Design of Existing Code*, Addison-Wesley, 1999

[FX] Foresight Exchange
http://www.ideosphere.com/fx/

[Gambit] McKelvey, Richard D., McLennan, Andrew M., and Turocy, Theodore L., "Gambit: Software Tools for Game Theory", Version 0.97.0.7
http://econweb.tamu.edu/gambit

[GoldmanSachs] Goldman Sachs Economic Derivatives Market
http://www.gs.com/econderivs

[Hahn & Tetlock 2004a] Robert W. Hahn, Paul C. Tetlock, "Harnessing the Power of Information: A New Approach to Economic Development", AEI-Brookings Joint Center for Regulatory Studies, Working Paper, September, 2004
http://www.aei-brookings.org/publications/abstract.php?pid=846

[Hahn & Tetlock 2004b] Robert W. Hahn and Paul C. Tetlock, "A New Approach for Regulating Information Markets", *Regulatory Analysis*, December, 2004
http://www.aei-brookings.org/publications/abstract.php?pid=881

[Hanson 2003] Hanson, Robin, "Combinatorial Information Market Design", *Information Systems Frontiers*, 5(1), 105-119, January 2003
http://hanson.gmu.edu/combobet.pdf

[Hanson1999] Hanson, Robin (1999), "Decision Markets", *IEEE Intelligent Systems*, 14(3), 16-19
http://hanson.gmu.edu/decisionmarkets.pdf

[Hayek] Hayek, F. A., "The Use of Knowledge in Society", *The American Economic Review*, Vol. 35 p. 519-530, September 1945 http://www.econlib.org/library/Essays/hykKnw1.html

[HedgeStreet] http://www.hedgestreet.com

[Hibbert] Prototype code
http://mydruthers.com/IF-code/index.html

[HSX] Hollywood Stock Exchange
http://www.hsx.com/

[IEEE] IEEE Spectrum Online, "Toward a Bettor Future", *IEEE Spectrum Online*, 8/8/2003
http://www.spectrum.ieee.org/WEBONLY/wonews/aug03/future.html

[IEM] Iowa Electronic Markets
http://www.biz.uiowa.edu/iem/

[JASA] http://www.csc.liv.ac.uk/~sphelps/jasa/

[jMarkets] http://jmarkets.ssel.caltech.edu/downloads.html

[Kittlitz] K. Kittlitz, "Experiences with the Foresight Exchange", *Extropy Online*, 1999
http://www.ideosphere.com/fx/papers/extropyonline/fx.html

[Leigh & Wolfers] "Policy Improves by Putting Rhetoric on Trial", *Sydney Morning Herald*, March 5, 2003
http://bpp.wharton.upenn.edu/jwolfers/press.shtml#Policytrials

[LongBets] http://longbets.org

[Malone] Thomas W. Malone, *The Future of Work*, Harvard Business School Press, 2004, pp. 91-92

[Miller] Mark Miller, David Krieger, Norman Hardy, Chris Hibbert, and E. Dean Tribble, "An Automated Auction in ATM Network Bandwidth", in *Market-Based Control: A Paradigm for Distributed Resource Allocation*, Scott Clearwater (Ed) 1996

[MUMS] "Game Theory and Experimental Economics", HP Labs web site,
http://www.hpl.hp.com/research/ssrc/services/architecture/economics/

[NetLab] National Science Foundation, "NetLab Workshop Report", 1997
http://www.uiowa.edu/~grpproc/netlab.htm

[NewsFutures] http://www.newsfutures.com

[Ortner] Ortner, Gerhard, "Forecasting Markets: An Industrial Application", mimeo, Technical University of Vienna, 1998
http://ebweb.tuwien.ac.at/apsm/fmaia2.pdf

[Page-Jones] Meilir Page-Jones, *What Every Programmer Should Know about Object-Oriented Design*, Dorset House, 1995

[PAM] "Policy Analysis Market Archive"
http://hanson.gmu.edu/policyanalysismarket.html

[Pennock] David M. Pennock, "A Dynamic Pari-Mutuel Market for Hedging, Wagering, and Information Aggregation", *Proceedings of the 5th ACM Conference on Electronic Commerce*, Pages: 170 – 179, 2004
http://dpennock.com/papers/pennock-ec-2004-dynamic-parimutuel.pdf

[Proebsting] Todd Proebsting, "Tee Time with Admiral Poindexter", Presentation at Workshop on Markets as Predictive Devices (Information Markets), Rutgers, February 2-4, 2005
http://dimacs.rutgers.edu/Workshops/Markets/program.html

[Reynolds] Kate Reynolds, "Going, Going, Gone! A Survey of Auction Types", Agorics, Inc., 1996
http://www.agorics.com/Library/auctions.html

[Schrieber] Emile-Servan Schrieber, "Does money matter?", Presentation at Workshop on Markets as Predictive Devices (Information Markets), Rutgers, February 2-4, 2005
http://dimacs.rutgers.edu/Workshops/Markets/program.html

[SMLI-TR] SunLabs Tech Report TR-95-50, "Fiscal 1995 Project Portfolio Report" http://www.sunlabs.com/techrep/1995/annualreport95/ist.html

[Stiegler] Marc Stiegler, *Earthweb*, Baen Books, 1999

[Time] *Time Magazine*, "The End Of Management?", Jul. 06, 2004
http://www.time.com/time/insidebiz/printout/0,8816,1101040712-660965-1,00.html

[TradeSports] http://tradesports.com/

[Walras] JQ Cheng and MP Wellman, "The WALRAS algorithm: A convergent distributed implementation of general equilibrium outcomes", *Computational Economics*, 12:1-24, 1998
http://kapis.www.wkap.nl/article.pdf?137532

[WebMart] http://www.sun.com/960201/cover/webmart.html

[Wolfers & Zitzewitz A] Wolfers, Justin and Eric Zitzewitz, "Prediction Markets", *Journal of Economic Perspectives* 18(2) Spring 2004 Pages 107-126
http://www.aei-brookings.org/admin/authorpdfs/page.php?id=1027

[Wolfers & Zitzewitz B] Wolfers, Justin and Eric Zitzewitz, "Experimental Political Betting Markets and the 2004 Election", AEI-Brookings Joint Center for Regulatory Studies, Related Publication, Oct. 2004
http://www.aei-brookings.org/admin/authorpdfs/page.php?id=1027

[Xlab] http://xlab.berkeley.edu

[Xlab Software] http://xlab.berkeley.edu/protocol1.shtml#software

[zLabs] CommerceNet, "zLabs Mission Statement", 2004
http://labs.commerce.net/wiki/index.php/About