

By Marti A. Hearst
University of California, Berkeley
hearst@sims.berkeley.edu

Building intelligent systems one e-citizen at a time

For decades, the AI community attempted to build intelligent systems by writing software that mimics human thought processes. More recently, researchers have advocated solving complex problems by building networks of interacting autonomous agents. But some researchers are now proposing that the key to building intelligent systems is leveraging the power of individual human thought. The idea is that now that networked information systems are widely available, they can be used to link together the mental efforts of individuals in innovative ways.

In this installment of "Trends and Controversies," we discuss two variations on this idea. First, Robin Hanson describes his idea of decision markets, also known as idea futures, in which matters of public debate are resolved via a mechanism like a commodities market.

This idea rests on the assumption that the answers to pressing social questions (such as, "Does violence in video games cause violence in children?") are not generally agreed on because those in the best position to know the answer can't or won't share this information. The proposed solution is to create a market in which getting the answer right or wrong has financial consequences. The hypothesis is that a market like this will tend to draw out the voices of those who are in good positions to know the correct answer. Just as online auctions let buyers and sellers find acceptable prices for their merchandise, an idea market allows knowledge-seekers to find a well-motivated answer for their question. The market might also compel people to invest in research to find out the correct answer.

The reason this idea might work is that those who are not well-informed will quickly lose their investments and will be motivated to stop participating in the market, while those who have more accurate information stand to gain more from the market and will direct its outcome. Hanson cites research in which an idea-based market was better at predicting election results than polls and expert forecasts.

In the second essay, David Stork presents what he calls the Open Mind Initiative. This is a broader formulation of an idea that subsumes decision markets. Stork argues that the newly networked society lets us capture and pool the results of human mental cycles that would otherwise be lost if simply expended in individual pursuits. For example, rather than simply playing a video game, why not play one that has a side-effect of producing training data for a machine-learning algorithm? Rather than simply finding and bookmarking a useful Web page, why not add it to an online Web directory? Many piecewise variations of these ideas are currently being bantered about; Stork puts these within a framework that lets us think about them more generally. He also draws parallels and contrasts between Open Mind and the Open Source software development effort.

These essays are complementary to those that appeared in the January/February 1999 "Trends and Controversies" on the unforeseen social consequences of networked information systems. The earlier essays described how networked information systems are changing society; this issue addresses how society might help build intelligent systems.

—Marti Hearst

Decision markets

Robin D. Hanson, George Mason University

Engineers' love of technology often gets in the way of their being useful. Consider Post-it Notes or, better yet, plain paper notepads. These probably seemed like trivial ideas, but they turned out to be terribly useful. Why? Because the marvel that is the human brain has a horrible short-term memory, which means that dumb-as-dirt memory aids can make people substantially smarter.

The human part of any large intelligent system is by far the most intelligent part. As long as this remains true, the biggest system advancements will come from aids that fill big holes in human abilities, rather than from artifacts that stretch engineers' abilities.

I mention all this because I want you to consider a simple, not very technically challenging idea—one that might nevertheless fill a gaping hole in our collective intelligence, similar to the way notepads fill a

gaping hole in our individual memories. I am talking in general about speculative markets, and in particular about *decision markets*. Decision markets might allow us to more accurately estimate the consequences of important decisions, by helping us to better share relevant information.

Consider, for example, a clearly important policy question such as,

How would crime rates change if more citizens could legally carry hidden guns?

Many observers say hidden guns obviously increase crime, while many others strongly disagree (see John Lott's *More Guns, Less Crime*, Univ. of Chicago Press).

I suspect that the existence of such divergent opinions reflects the fact that we suffer from a serious failure to share information. The so-called Information Revolution has greatly improved our ability to find out what others have said. However, it has done much less to improve our ability to find out what other people know. We can now find a blizzard of words on a topic such as the interplay between guns and crime, but we know that most of those words are written by people with axes to grind. The real problem is not finding more words, but judging who really knows about the topic and whether these experts are saying what they know. We are in many ways bit-full, yet information-poor.

I suggest that speculative markets are a neglected way to help us find out what people know. Such markets pool the information that is known to diverse individuals into a common resource, and have many advantages over standard institutions for information aggregation, such as news media, peer review, trials, and opinion polls. Speculative markets are decentralized and relatively egalitarian, and can offer direct, concise, timely, and precise estimates in answer to questions we pose.

These estimates are self-consistent across a wide range of issues and respond quickly to new information. They also seem to be cheap and relatively accurate.

In particular, for questions such as hidden guns and crime, I suggest we consider decision markets, which are speculative markets focusing on particular decisions.

How decision markets work

Imagine that we created markets where people could bet on future crime rates, conditional on allowing or not allowing more hidden guns. That is, if the market prices predicted that murder rates would be 10% higher should a certain hidden-gun bill pass, then anyone who thought this estimate too high could easily identify a particular profitable trade. If you made this trade and the market estimate then fell to 5%, for example, you could undo this trade for a profit.

Imagine further that if market prices said that crime rates would be 10% higher given more hidden guns, most nonexperts would accept this estimate as our “best answer” or as a neutral “consensus.” In particular, a state legislature might accept this estimate when considering whether or not to pass this hidden-gun bill.

I call such a set of markets a decision market. In this situation, advocates for each side of an issue would be forced to influence speculators if they wanted to influence general opinion. Speculators, in turn, would have a clear incentive to be careful and honest in contributing what they know and in judging what advocates know. This is because speculators must “put their money where their mouth is.”

Six steps are required to create a decision market to help us better share information on a topic such as the effect of hidden guns on crime.

First, you must state your claim clearly. For example, you might focus on a particular bill B before your state legislature, which would allow more citizens to carry hidden guns. You might decide to focus on your state’s murder rate, using some standard government statistic M as your official measure of it. You should choose a lowest and highest relevant murder rate, and scale M so that $M = 0$ at the lowest rate and $M = 1$ at the highest rate. (You might choose, for example, the lowest rate to be zero murders and the highest rate to be the population size, which is the highest conceivable murder rate.)



Robin D. Hanson is a Robert Wood Johnson Foundation Scholar in Health Policy Research at the University of California at Berkeley. This fall he will be an assistant professor of economics at George Mason University. He spent nine years as a researcher in AI and Bayesian statistics, at Lockheed and then at NASA. His research interests include explaining various puzzling human behaviors in terms of information asymmetries. He has a BS in physics from the University of California at Irvine, an MS in physics and an MA in philosophy of science from the University of Chicago, and a PhD in social science from the California Institute of Technology. Contact him at 140 Warren Hall, School of Public Health, UC Berkeley, 94720-7360; hanson@econ.berkeley.edu; <http://hanson.berkeley.edu>.



David G. Stork is the chief scientist at Ricoh Silicon Valley and a consulting associate professor of electrical engineering at Stanford University. A graduate of MIT and the University of Maryland, he has written and edited several books, including *HAL's Legacy* and (with R.O. Duda and P.E. Hart) the forthcoming second edition of *Pattern Classification*. His central interests are in pattern recognition in machines and humans. Contact him at Ricoh Silicon Valley, 2882 Sand Hill Rd., #15, Menlo Park, CA 94025-7022; stork@rsv.ricoh.com.

Next, you must choose some particular trusted third party who will finally declare a murder rate M within $[0,1]$ and determine whether bill B passed. (This party might be a jury randomly drawn from some pool.) You must also either pick a date by which these judges are to decide, tie the judging to some other event like the release of murder statistics, or grant the judges discretion to choose this date for themselves.

Third, you must choose what asset A you will bet. If the bet is going to last any substantial time, asset A ought to give a reasonable rate of return, to induce speculators to invest in it. You might pick a safe government bond, or you might pick a broad-index stock mutual fund.

Also, given M, B, and this asset A, you authorize some financial institution to make exchanges between units of A and the following set of four assets:

1. M units of A if B passes,
2. M units of A if B does not pass,
3. $1 - M$ units of A if B passes,
4. $1 - M$ units of A if B does not pass.

Note that because $M + (1 - M) = 1$, and because B either will or won't pass, this financial institution takes no risk from these exchanges; each set of four assets will be worth exactly one unit of A in the end.

Fifth, you create markets in which people can trade various combinations of these assets with each other. In particular, if people trade asset 1 for the bundle of assets 1 and 3, the market price (asset ratio in trades) is an estimate of the murder rate conditional on the bill passing. Similarly, the price in trades of asset 2 for the bundle of 2 and 4 is an estimate of the murder rate

conditional on the bill not passing. Moreover, comparing these two estimates tells you whether, and by how much, speculators expect this bill to increase or decrease the murder rate.

Finally, you have to decide how much to subsidize this market. If interest in your topic is strong enough, simply creating these markets might induce people to trade in them. Failing that, sufficient interest might be induced if someone committed to make a policy choice based on the market estimate. A state legislature, for example, might commit to pass the bill or not depending on the market estimate of their effect on the murder rate.

You can also safely and directly subsidize a market to induce more participation. Doing this in effect creates an *information prize* offered to those who first make the market price better reflect relevant information. (In econo-speak, one way to do this is to create a market maker whose bid and ask prices are monotonic functions of its assets held.)

In addition to estimating the effect of hidden guns on crime, decision markets might give us estimates on

- Murder rates—with or without capital punishment?
- Average mortality rates—with or without national health insurance?
- Health-care spending—with expanded or curtailed use of health maintenance organizations?
- Employment change—raise minimum wage or rescind NAFTA?
- Global sea-level and temperature changes—impose or not impose a carbon dioxide tax?

- Military casualties—with a Republican or a Democratic president?
- Stock prices—with a Republican- or Democrat-controlled US Congress?
- World per-capita food consumption—raise or lower average tariffs?
- Student test scores—with or without school choice or voucher reform?
- Future national economic growth—raise or lower interest rates, or with or without an education subsidy?

Science fiction writers have posited even more ubiquitous betting markets (see John Brunner's *Shock-wave Rider*, Del Rey Books, and Marc Steigler's *Earthweb*). In general, decision markets can estimate the net effect of any policy choice of interest on any outcome of interest, as long as there is a decent chance that, after the fact, we can reasonably verify what outcome happened and what policy was chosen.


How well do markets work?

By its nature, a betting-market estimate is decentralized, direct, concise, timely, precise, self-consistent, and responds quickly to new information. It is also egalitarian, if everyone is allowed to participate. But how clear, cheap, and accurate are such market estimates?

On accuracy, decades of research on the efficiency of financial markets have found little price-relevant information that is not reflected in market prices. Any inefficiencies seem to be weak and to go away with publicity, because they represent a profit opportunity. If you think the current price is too low, you expect to profit by buying now and selling later, and buying now will raise the price, partially correcting the error you perceived.

Speculative markets have done well in direct tests against standard information-aggregation institutions. For example, orange juice futures prices have been shown to improve on government weather forecasts.¹ Also, markets where traders can bet on election results predict vote totals better than opinion polls.²

How do markets do so well? After all, aren't they made of the same fallible humans as other institutions? A study of



those election markets found that while most traders tended to suffer from cognitive biases such as expecting others to agree with them, the most active traders were not biased this way—and active traders set the prices. Speculative markets thus seem to induce the real experts to self-select and participate more. Lab experiments also indicate that speculative markets tend to aggregate information when traders are experienced with their roles and know the payoffs for other roles.³

Older economics writings sometimes give the impression that speculative markets will not function unless you have thousands of traders frequently trading millions of dollars worth of goods. Recent Web markets, however, show clearly that markets can be much smaller and slower than this. Furthermore, a subsidized market can function over any time period with only one trader. If an information prize is offered, and only one person is induced to learn enough to only once correct the initial market price to something else, the market has still served a valuable information role.

With subsidies, the key question is not whether we can create speculative markets, or whether such markets can induce people to learn and reveal information. The key questions are whether the information gained is worth the costs paid and whether a similar benefit could have come cheaper via some other institution.

What's the holdup?

If speculative markets are so great at information aggregation, why don't we al-

ready use them to form consensus on topics such as the correlation between hidden guns and crime?

Until the last few centuries, the cost of simply handling trades was enough to sharply limit the number of speculative markets that could be made widely available. Yet today *ebay.com* routinely sells \$10 items by having a handful of people bid a few times each over a period of a week. Moreover, play-money Web betting games have shown that just a handful of people, each making a few small trades over several years, can create reasonable estimates on a wide variety of questions. (See www.hsx.com, myhand.com, and especially www.ideosphere.com.)

More important, most speculative markets are now illegal. The short history of financial market regulation is that everything was once illegal, until limited exemptions were granted for specific purposes. Betting on cards was a foolish waste of money; only fools would invest in a business they did not closely monitor; and it was the height of folly to let people bet on the death of others. So casinos, stocks, and insurance were all banned.⁴

Gradually, exceptions were granted for what came to be seen as worthy purposes, such as teaching people about horses (horseracing) or raising state revenue (lotteries). Stocks were allowed for the purpose of capitalizing firms, and insurance was allowed to let individuals hedge risks. More recently, commodity futures and financial derivatives were allowed to let firms hedge more risks. All these areas are highly regulated, however, in part to prevent limited exemptions from devolving into general gambling.

Accepted functions of markets now include entertainment, capitalization, and hedging, but not *information aggregation*. Thus while it is widely recognized that markets created for other purposes accomplish IA, we're prevented from creating a market whose primary function is IA. So we cannot create a market whose legal price would inform raging policy debates, such as the interplay between crime rates and hidden guns.

Okay, betting markets are mainly illegal. But if economists have data suggesting that that speculative markets do well at IA, why aren't lots of economists pushing the idea

of better IA via more markets?

Actually, it's worse than you think; economists also have sophisticated theory that suggests that IA should not be that hard on factual topics like the effect of hidden guns on crime. Rational agents should not even be able to agree on which one of them thinks hidden guns cause more crime.⁵ (I won't say more, as the editor wisely advises against using more econo-speak.) Economic theory thus really does suggest we humans have a gaping hole in our social intelligence.

So why aren't economists pushing IA markets? One answer is that economists are just spread too thin. Economic theory suggests many policy improvements over the status quo, and there are few economists that anyone else will listen to. These few economists thus have to choose their battles carefully.

The relevant theory for IA is also recent, and most economists don't yet know about it. Worse, the IA functions of markets seem too complex to model in much generality. When systems become too complex to model in detail, engineers usually resort to building and testing theory-inspired prototypes. However, economic-theorist types who understand this area are reluctant to move that far away from theory. (Capitalization and hedging functions of markets are easier to model, and economists do use theory to design market prototypes for these functions.)

It thus seems to fall to a few economics-savvy and engineering-minded folks like me to think of using prototypes to explore the idea of using more speculative markets for IA.

A promising direction: internal corporate markets

On the types of topics to which they have been applied so far, IA markets have looked promising. There remain, however, many legitimate concerns. For example, does the existence of speculative markets discourage communication via other channels, and is this a net benefit or loss?

Tests of prototypes might help us answer such questions. But how can we test prototypes, if IA markets are generally illegal? Well, there is one plausible loophole (besides offshore gambling), which I have saved for those of you who are still reading this far: internal corporate markets. Corporations have great leeway in what they make

employee bonuses depend on, and a contingent bonus is pretty close to a bet. So several companies, including Hewlett-Packard and Siemens, have begun experimenting with real-money internal speculative markets for estimating things such as future sales.

Corporations also need to make decisions, and often have problems inducing relevant parties to reveal information about the consequences of those decisions. Furthermore, companies have a good rough-and-ready measure of "good for the company"—the stock price. Thus, you could create decision markets that estimate whether any particular decision, such as introducing a new product, is better or worse than some alternative for the stock price. Alternatively, you might predict the sales of some product contingent on some important product-design decision.

Just as notepads fill a gaping hole in our individual cognitive abilities, speculative markets might fill a gaping hole in our collective ability to share information. Economic theory suggests that IA should not be that hard, at least for factual policy questions like the effect of hidden guns on crime rates. Speculative markets seem to work well at such tasks. Let us thus develop prototypes to explore this potential, in the hopes of someday lifting current legal barriers to widespread use of more effective institutions for IA.

For more information on this topic, see <http://hanson.berkeley.edu/ideaofutures.html>.

References

1. R. Roll, "Orange Juice and Weather," *American Economic Rev.*, Vol. 74, No. 5, Dec. 1984, pp. 861–880.
2. R. Foresythe et al., "Anatomy of an Experimental Political Stock Market," *American Economic Rev.*, Vol. 82, No. 5, Dec. 1992, pp. 1142–1161.
3. R. Foresythe and R. Lundholm, "Information Aggregation in an Experimental Market," *Econometrica*, Vol. 58, No. 2, Mar. 1990, pp. 309–347.
4. I.N. Rose, *Gambling and the Law*, Gambling Times Inc., Hollywood, Calif., 1986.
5. R. Hanson, *Four Puzzles in Information & Politics*, PhD thesis, California Inst. of Technology, Pasadena, Calif., 1997.

The Open Mind Initiative

David G. Stork, Ricoh Silicon Valley

After decades of research in pattern recognition and components of intelligent systems, the AI community has shifted its focus from fundamental concepts and mathematical techniques to large-scale data acquisition and knowledge engineering. In contest after contest in academic and commercial realms, the best systems for optical character, speech and face recognition, and so on are the ones trained with the most data. Collecting very large, high-quality datasets is evidently vital to progress in these and several other areas. Such data is informal—known by everyone who can read, speak, or hear, or has a commonsense understanding of the world.

Consider, too, the qualified but increasingly compelling success of the Open Source methodology—which promotes software reliability and quality by supporting independent peer review and rapid evolution of freely distributed source code. Linux, SendMail, Apache, the Mozilla version of the Netscape Web browser, and other high-quality software testify to the viability of this collaborative approach to software engineering.

Together, these developments suggest a new approach to building components of intelligent systems—the Open Mind Initiative. This initiative relies on three types of participants:

- domain experts who contribute libraries of algorithms,
- tool developers who contribute and refine the enabling software, and
- lay e-citizens who contribute data via the Internet.

Users with an interest and expertise in a particular domain, such as speech, vision,

**Coming Next
Issue**

**Quantum Computing
and AI**




Figure 1. In an Open Mind project on OCR: handwritten characters are presented to e-citizens whose judgements (here, 4 vs. 9) are entered for instance by means of buttons. Such responses are returned to the Open Mind host and used to train the classifier.

language, or commonsense, serve as reviewers or moderators.

Consider the development of an optical character-recognition system through Open Mind (see Figure 1). A host machine presents pixel images of handwritten characters, transformed by distortions, warping, line thickening, and so forth, on browsers of e-citizens, perhaps in a game interface. E-citizens classify these images by means of a button response; the responses are aggregated, screened for significant outliers at the host machine, and used as training data to improve the classifier.

To explore the novel infrastructure required by Open Mind, Stanford graduate student Chuck Lam and I have developed a Java- and Web-based version of *Animals*, an elementary interactive children's program for classifying animals, dating from the late 1970s. The child (e-citizen) thinks of an animal. Using the child's responses to a series of questions (for example, two-legged or four-legged?), the program tries to determine this animal's identity. If the program guesses wrong, the child must enter a question that distinguishes her animal from the program's guess. After a number of children have played this guessing game, *Animals* has learned a simple tree-based classifier for animals. While not a deep or particularly useful core program, *Animals* provides an excellent platform for solving important problems in Open Mind.

It provides an interface and protocol design for

- efficiently extracting the maximum information from e-citizens,
- detecting and eliminating significant errors and statistical "outliers,"
- repelling hostile attacks, and
- automatically listing contributors according to the amount of information they contributed.

The Open Mind Initiative differs from the Free Software Foundations and the Open Source organization in a number of ways. First, while Open Source draws its support almost entirely from a hacker culture (for example, roughly 10^5 programmers have contributed to *Linux*), Open Mind relies heavily on an e-citizen and business culture (10^9 nonprogrammers on the Web). While most of the work in Open Source is directly on the final released source code, most of the effort in Open Mind focuses on the tools, infrastructure, and data gathering. An expert arbitrates final decisions in Open Source; in Open Mind, much information is accepted or rejected automatically by the infrastructure software. Table 1 summarizes some of these differences.

While domain experts and infrastructure developers are likely to have the same motivations as contributors to Open Source, the

motivations of e-citizens deserve special consideration. E-citizens seek benefit from the resulting Open Mind software, including software that would be very difficult to develop in other ways, such as commonsense knowledge. E-citizens would enjoy game interfaces and seek the public recognition of their contributions. There could be financial incentives such as lotteries, discounts, or frequent-flyer awards provided by corporations seeking new customers.

Others have discussed components of Open Mind and its use in different contexts—online interactive data acquisition and voting, collaboration, and machine-learning and pattern-recognition algorithms—yet the integration proposed in Open Mind seems not to have been discussed. A number of existing projects would fit under an Open Mind umbrella and would profit from the initiative. One example is *Newhoo*, where nonspecialist e-citizens propose keyword and index information about Web pages. Their contributions are reviewed by volunteer referee and editors (currently 10,000) and made available to all. We can imagine Open Mind projects in numerous pattern-recognition domains or knowledge engineering to improve navigating news groups, Web sites, or FAQs. Later, these systems can be integrated, for instance, to incorporate natural language or commonsense constraints in speech recognition, OCR, or web searching.


Physics has had its atom smashers, microbiology its Human Genome Project, and aeronautics and astronautics its space missions. Now is the time for computer science and cognitive science to have their big science—one that harvests informal knowledge from a large number of e-citizens for building useful software for next-generation systems. Given the conjunction of several forces—the need for natural human-machine interfaces and improved Web searching, the existence of good learning algorithms and Web infrastructure, and the demonstrated success of the Open Source methodology—the time is right for the Open Mind Initiative.

Reference

1. D.G. Stork, "Character and Document Recognition in the Open Mind Initiative," *Proc. Int'l Conf. Document Analysis and Recognition (ICDAR '99)*, to be published by IEEE Computer Society Press, Los Alamitos, Calif., 1999.

Table 1. Comparisons between Open Source and Open Mind.

OPEN SOURCE	OPEN MIND
No e-citizens	E-citizens crucial
Expert knowledge	Informal knowledge
Machine learning irrelevant	Machine learning essential
Web useful but optional	Web essential
Most work is directly on the end-user software	Most work is on infrastructure, not on the end-user software
Hacker culture ($<10^5$)	E-citizen/business culture ($<10^9$)
Separate functions contributed (device drivers in Linux)	Single functional goal (for example, recognition rate in OCR)



By Marti A. Hearst
University of California, Berkeley
hearst@sims.berkeley.edu

Building intelligent systems one e-citizen at a time

For decades, the AI community attempted to build intelligent systems by writing software that mimics human thought processes. More recently, researchers have advocated solving complex problems by building networks of interacting autonomous agents. But some researchers are now proposing that the key to building intelligent systems is leveraging the power of individual human thought. The idea is that now that networked information systems are widely available, they can be used to link together the mental efforts of individuals in innovative ways.

In this installment of "Trends and Controversies," we discuss two variations on this idea. First, Robin Hanson describes his idea of decision markets, also known as idea futures, in which matters of public debate are resolved via a mechanism like a commodities market.

This idea rests on the assumption that the answers to pressing social questions (such as, "Does violence in video games cause violence in children?") are not generally agreed on because those in the best position to know the answer can't or won't share this information. The proposed solution is to create a market in which getting the answer right or wrong has financial consequences. The hypothesis is that a market like this will tend to draw out the voices of those who are in good positions to know the correct answer. Just as online auctions let buyers and sellers find acceptable prices for their merchandise, an idea market allows knowledge-seekers to find a well-motivated answer for their question. The market might also compel people to invest in research to find out the correct answer.

The reason this idea might work is that those who are not well-informed will quickly lose their investments and will be motivated to stop participating in the market, while those who have more accurate information stand to gain more from the market and will direct its outcome. Hanson cites research in which an idea-based market was better at predicting election results than polls and expert forecasts.

In the second essay, David Stork presents what he calls the Open Mind Initiative. This is a broader formulation of an idea that subsumes decision markets. Stork argues that the newly networked society lets us capture and pool the results of human mental cycles that would otherwise be lost if simply expended in individual pursuits. For example, rather than simply playing a video game, why not play one that has a side-effect of producing training data for a machine-learning algorithm? Rather than simply finding and bookmarking a useful Web page, why not add it to an online Web directory? Many piecemeal variations of these ideas are currently being bantered about; Stork puts these within a framework that lets us think about them more generally. He also draws parallels and contrasts between Open Mind and the Open Source software development effort.

These essays are complementary to those that appeared in the January/February 1999 "Trends and Controversies" on the unforeseen social consequences of networked information systems. The earlier essays described how networked information systems are changing society; this issue addresses how society might help build intelligent systems.

—Marti Hearst

Decision markets

Robin D. Hanson, George Mason University

Engineers' love of technology often gets in the way of their being useful. Consider Post-it Notes or, better yet, plain paper notepads. These probably seemed like trivial ideas, but they turned out to be terribly useful. Why? Because the marvel that is the human brain has a horrible short-term memory, which means that dumb-as-dirt memory aids can make people substantially smarter.

The human part of any large intelligent system is by far the most intelligent part. As long as this remains true, the biggest system advancements will come from aids that fill big holes in human abilities, rather than from artifacts that stretch engineers' abilities.

I mention all this because I want you to consider a simple, not very technically challenging idea—one that might nevertheless fill a gaping hole in our collective intelligence, similar to the way notepads fill a

gaping hole in our individual memories. I am talking in general about speculative markets, and in particular about *decision markets*. Decision markets might allow us to more accurately estimate the consequences of important decisions, by helping us to better share relevant information.

Consider, for example, a clearly important policy question such as,

How would crime rates change if more citizens could legally carry hidden guns?

Many observers say hidden guns obviously increase crime, while many others strongly disagree (see John Lott's *More Guns, Less Crime*, Univ. of Chicago Press).

I suspect that the existence of such divergent opinions reflects the fact that we suffer from a serious failure to share information. The so-called Information Revolution has greatly improved our ability to find out what others have said. However, it has done much less to improve our ability to find out what other people know. We can now find a blizzard of words on a topic such as the interplay between guns and crime, but we know that most of those words are written by people with axes to grind. The real problem is not finding more words, but judging who really knows about the topic and whether these experts are saying what they know. We are in many ways bit-full, yet information-poor.

I suggest that speculative markets are a neglected way to help us find out what people know. Such markets pool the information that is known to diverse individuals into a common resource, and have many advantages over standard institutions for information aggregation, such as news media, peer review, trials, and opinion polls. Speculative markets are decentralized and relatively egalitarian, and can offer direct, concise, timely, and precise estimates in answer to questions we pose.

These estimates are self-consistent across a wide range of issues and respond quickly to new information. They also seem to be cheap and relatively accurate.

In particular, for questions such as hidden guns and crime, I suggest we consider decision markets, which are speculative markets focusing on particular decisions.

How decision markets work

Imagine that we created markets where people could bet on future crime rates, conditional on allowing or not allowing more hidden guns. That is, if the market prices predicted that murder rates would be 10% higher should a certain hidden-gun bill pass, then anyone who thought this estimate too high could easily identify a particular profitable trade. If you made this trade and the market estimate then fell to 5%, for example, you could undo this trade for a profit.

Imagine further that if market prices said that crime rates would be 10% higher given more hidden guns, most nonexperts would accept this estimate as our “best answer” or as a neutral “consensus.” In particular, a state legislature might accept this estimate when considering whether or not to pass this hidden-gun bill.

I call such a set of markets a decision market. In this situation, advocates for each side of an issue would be forced to influence speculators if they wanted to influence general opinion. Speculators, in turn, would have a clear incentive to be careful and honest in contributing what they know and in judging what advocates know. This is because speculators must “put their money where their mouth is.”

Six steps are required to create a decision market to help us better share information on a topic such as the effect of hidden guns on crime.

First, you must state your claim clearly. For example, you might focus on a particular bill B before your state legislature, which would allow more citizens to carry hidden guns. You might decide to focus on your state’s murder rate, using some standard government statistic M as your official measure of it. You should choose a lowest and highest relevant murder rate, and scale M so that $M = 0$ at the lowest rate and $M = 1$ at the highest rate. (You might choose, for example, the lowest rate to be zero murders and the highest rate to be the population size, which is the highest conceivable murder rate.)



Robin D. Hanson is a Robert Wood Johnson Foundation Scholar in Health Policy Research at the University of California at Berkeley. This fall he will be an assistant professor of economics at George Mason University. He spent nine years as a researcher in AI and Bayesian statistics, at Lockheed and then at NASA. His research interests include explaining various puzzling human behaviors in terms of information asymmetries. He has a BS in physics from the University of California at Irvine, an MS in physics and an MA in philosophy of science from the University of Chicago, and a PhD in social science from the California Institute of Technology. Contact him at 140 Warren Hall, School of Public Health, UC Berkeley, 94720-7360; hanson@econ.berkeley.edu; <http://hanson.berkeley.edu>.



David G. Stork is the chief scientist at Ricoh Silicon Valley and a consulting associate professor of electrical engineering at Stanford University. A graduate of MIT and the University of Maryland, he has written and edited several books, including *HAL's Legacy* and (with R.O. Duda and P.E. Hart) the forthcoming second edition of *Pattern Classification*. His central interests are in pattern recognition in machines and humans. Contact him at Ricoh Silicon Valley, 2882 Sand Hill Rd., #15, Menlo Park, CA 94025-7022; stork@rsv.ricoh.com.

Next, you must choose some particular trusted third party who will finally declare a murder rate M within $[0,1]$ and determine whether bill B passed. (This party might be a jury randomly drawn from some pool.) You must also either pick a date by which these judges are to decide, tie the judging to some other event like the release of murder statistics, or grant the judges discretion to choose this date for themselves.

Third, you must choose what asset A you will bet. If the bet is going to last any substantial time, asset A ought to give a reasonable rate of return, to induce speculators to invest in it. You might pick a safe government bond, or you might pick a broad-index stock mutual fund.

Also, given M, B, and this asset A, you authorize some financial institution to make exchanges between units of A and the following set of four assets:

1. M units of A if B passes,
2. M units of A if B does not pass,
3. $1 - M$ units of A if B passes,
4. $1 - M$ units of A if B does not pass.

Note that because $M + (1 - M) = 1$, and because B either will or won't pass, this financial institution takes no risk from these exchanges; each set of four assets will be worth exactly one unit of A in the end.

Fifth, you create markets in which people can trade various combinations of these assets with each other. In particular, if people trade asset 1 for the bundle of assets 1 and 3, the market price (asset ratio in trades) is an estimate of the murder rate conditional on the bill passing. Similarly, the price in trades of asset 2 for the bundle of 2 and 4 is an estimate of the murder rate

conditional on the bill not passing. Moreover, comparing these two estimates tells you whether, and by how much, speculators expect this bill to increase or decrease the murder rate.

Finally, you have to decide how much to subsidize this market. If interest in your topic is strong enough, simply creating these markets might induce people to trade in them. Failing that, sufficient interest might be induced if someone committed to make a policy choice based on the market estimate. A state legislature, for example, might commit to pass the bill or not depending on the market estimate of their effect on the murder rate.

You can also safely and directly subsidize a market to induce more participation. Doing this in effect creates an *information prize* offered to those who first make the market price better reflect relevant information. (In econo-speak, one way to do this is to create a market maker whose bid and ask prices are monotonic functions of its assets held.)

In addition to estimating the effect of hidden guns on crime, decision markets might give us estimates on

- Murder rates—with or without capital punishment?
- Average mortality rates—with or without national health insurance?
- Health-care spending—with expanded or curtailed use of health maintenance organizations?
- Employment change—raise minimum wage or rescind NAFTA?
- Global sea-level and temperature changes—impose or not impose a carbon dioxide tax?

- Military casualties—with a Republican or a Democratic president?
- Stock prices—with a Republican- or Democrat-controlled US Congress?
- World per-capita food consumption—raise or lower average tariffs?
- Student test scores—with or without school choice or voucher reform?
- Future national economic growth—raise or lower interest rates, or with or without an education subsidy?

Science fiction writers have posited even more ubiquitous betting markets (see John Brunner's *Shock-wave Rider*, Del Rey Books, and Marc Steigler's *Earthweb*). In general, decision markets can estimate the net effect of any policy choice of interest on any outcome of interest, as long as there is a decent chance that, after the fact, we can reasonably verify what outcome happened and what policy was chosen.


How well do markets work?

By its nature, a betting-market estimate is decentralized, direct, concise, timely, precise, self-consistent, and responds quickly to new information. It is also egalitarian, if everyone is allowed to participate. But how clear, cheap, and accurate are such market estimates?

On accuracy, decades of research on the efficiency of financial markets have found little price-relevant information that is not reflected in market prices. Any inefficiencies seem to be weak and to go away with publicity, because they represent a profit opportunity. If you think the current price is too low, you expect to profit by buying now and selling later, and buying now will raise the price, partially correcting the error you perceived.

Speculative markets have done well in direct tests against standard information-aggregation institutions. For example, orange juice futures prices have been shown to improve on government weather forecasts.¹ Also, markets where traders can bet on election results predict vote totals better than opinion polls.²

How do markets do so well? After all, aren't they made of the same fallible humans as other institutions? A study of



those election markets found that while most traders tended to suffer from cognitive biases such as expecting others to agree with them, the most active traders were not biased this way—and active traders set the prices. Speculative markets thus seem to induce the real experts to self-select and participate more. Lab experiments also indicate that speculative markets tend to aggregate information when traders are experienced with their roles and know the payoffs for other roles.³

Older economics writings sometimes give the impression that speculative markets will not function unless you have thousands of traders frequently trading millions of dollars worth of goods. Recent Web markets, however, show clearly that markets can be much smaller and slower than this. Furthermore, a subsidized market can function over any time period with only one trader. If an information prize is offered, and only one person is induced to learn enough to only once correct the initial market price to something else, the market has still served a valuable information role.

With subsidies, the key question is not whether we can create speculative markets, or whether such markets can induce people to learn and reveal information. The key questions are whether the information gained is worth the costs paid and whether a similar benefit could have come cheaper via some other institution.

What's the holdup?

If speculative markets are so great at information aggregation, why don't we al-

ready use them to form consensus on topics such as the correlation between hidden guns and crime?

Until the last few centuries, the cost of simply handling trades was enough to sharply limit the number of speculative markets that could be made widely available. Yet today *ebay.com* routinely sells \$10 items by having a handful of people bid a few times each over a period of a week. Moreover, play-money Web betting games have shown that just a handful of people, each making a few small trades over several years, can create reasonable estimates on a wide variety of questions. (See www.hsx.com, myhand.com, and especially www.ideosphere.com.)

More important, most speculative markets are now illegal. The short history of financial market regulation is that everything was once illegal, until limited exemptions were granted for specific purposes. Betting on cards was a foolish waste of money; only fools would invest in a business they did not closely monitor; and it was the height of folly to let people bet on the death of others. So casinos, stocks, and insurance were all banned.⁴

Gradually, exceptions were granted for what came to be seen as worthy purposes, such as teaching people about horses (horseracing) or raising state revenue (lotteries). Stocks were allowed for the purpose of capitalizing firms, and insurance was allowed to let individuals hedge risks. More recently, commodity futures and financial derivatives were allowed to let firms hedge more risks. All these areas are highly regulated, however, in part to prevent limited exemptions from devolving into general gambling.

Accepted functions of markets now include entertainment, capitalization, and hedging, but not *information aggregation*. Thus while it is widely recognized that markets created for other purposes accomplish IA, we're prevented from creating a market whose primary function is IA. So we cannot create a market whose legal price would inform raging policy debates, such as the interplay between crime rates and hidden guns.

Okay, betting markets are mainly illegal. But if economists have data suggesting that that speculative markets do well at IA, why aren't lots of economists pushing the idea

of better IA via more markets?

Actually, it's worse than you think; economists also have sophisticated theory that suggests that IA should not be that hard on factual topics like the effect of hidden guns on crime. Rational agents should not even be able to agree on which one of them thinks hidden guns cause more crime.⁵ (I won't say more, as the editor wisely advises against using more econo-speak.) Economic theory thus really does suggest we humans have a gaping hole in our social intelligence.

So why aren't economists pushing IA markets? One answer is that economists are just spread too thin. Economic theory suggests many policy improvements over the status quo, and there are few economists that anyone else will listen to. These few economists thus have to choose their battles carefully.

The relevant theory for IA is also recent, and most economists don't yet know about it. Worse, the IA functions of markets seem too complex to model in much generality. When systems become too complex to model in detail, engineers usually resort to building and testing theory-inspired prototypes. However, economic-theorist types who understand this area are reluctant to move that far away from theory. (Capitalization and hedging functions of markets are easier to model, and economists do use theory to design market prototypes for these functions.)

It thus seems to fall to a few economics-savvy and engineering-minded folks like me to think of using prototypes to explore the idea of using more speculative markets for IA.

A promising direction: internal corporate markets

On the types of topics to which they have been applied so far, IA markets have looked promising. There remain, however, many legitimate concerns. For example, does the existence of speculative markets discourage communication via other channels, and is this a net benefit or loss?

Tests of prototypes might help us answer such questions. But how can we test prototypes, if IA markets are generally illegal? Well, there is one plausible loophole (besides offshore gambling), which I have saved for those of you who are still reading this far: internal corporate markets. Corporations have great leeway in what they make

employee bonuses depend on, and a contingent bonus is pretty close to a bet. So several companies, including Hewlett-Packard and Siemens, have begun experimenting with real-money internal speculative markets for estimating things such as future sales.

Corporations also need to make decisions, and often have problems inducing relevant parties to reveal information about the consequences of those decisions. Furthermore, companies have a good rough-and-ready measure of "good for the company"—the stock price. Thus, you could create decision markets that estimate whether any particular decision, such as introducing a new product, is better or worse than some alternative for the stock price. Alternatively, you might predict the sales of some product contingent on some important product-design decision.

Just as notepads fill a gaping hole in our individual cognitive abilities, speculative markets might fill a gaping hole in our collective ability to share information. Economic theory suggests that IA should not be that hard, at least for factual policy questions like the effect of hidden guns on crime rates. Speculative markets seem to work well at such tasks. Let us thus develop prototypes to explore this potential, in the hopes of someday lifting current legal barriers to widespread use of more effective institutions for IA.

For more information on this topic, see <http://hanson.berkeley.edu/ideaofutures.html>.

References

1. R. Roll, "Orange Juice and Weather," *American Economic Rev.*, Vol. 74, No. 5, Dec. 1984, pp. 861–880.
2. R. Foresythe et al., "Anatomy of an Experimental Political Stock Market," *American Economic Rev.*, Vol. 82, No. 5, Dec. 1992, pp. 1142–1161.
3. R. Foresythe and R. Lundholm, "Information Aggregation in an Experimental Market," *Econometrica*, Vol. 58, No. 2, Mar. 1990, pp. 309–347.
4. I.N. Rose, *Gambling and the Law*, Gambling Times Inc., Hollywood, Calif., 1986.
5. R. Hanson, *Four Puzzles in Information & Politics*, PhD thesis, California Inst. of Technology, Pasadena, Calif., 1997.

The Open Mind Initiative

David G. Stork, Ricoh Silicon Valley

After decades of research in pattern recognition and components of intelligent systems, the AI community has shifted its focus from fundamental concepts and mathematical techniques to large-scale data acquisition and knowledge engineering. In contest after contest in academic and commercial realms, the best systems for optical character, speech and face recognition, and so on are the ones trained with the most data. Collecting very large, high-quality datasets is evidently vital to progress in these and several other areas. Such data is informal—known by everyone who can read, speak, or hear, or has a commonsense understanding of the world.

Consider, too, the qualified but increasingly compelling success of the Open Source methodology—which promotes software reliability and quality by supporting independent peer review and rapid evolution of freely distributed source code. Linux, SendMail, Apache, the Mozilla version of the Netscape Web browser, and other high-quality software testify to the viability of this collaborative approach to software engineering.

Together, these developments suggest a new approach to building components of intelligent systems—the Open Mind Initiative. This initiative relies on three types of participants:

- domain experts who contribute libraries of algorithms,
- tool developers who contribute and refine the enabling software, and
- lay e-citizens who contribute data via the Internet.

Users with an interest and expertise in a particular domain, such as speech, vision,

**Coming Next
Issue**

**Quantum Computing
and AI**




Figure 1. In an Open Mind project on OCR: handwritten characters are presented to e-citizens whose judgements (here, 4 vs. 9) are entered for instance by means of buttons. Such responses are returned to the Open Mind host and used to train the classifier.

language, or commonsense, serve as reviewers or moderators.

Consider the development of an optical character-recognition system through Open Mind (see Figure 1). A host machine presents pixel images of handwritten characters, transformed by distortions, warping, line thickening, and so forth, on browsers of e-citizens, perhaps in a game interface. E-citizens classify these images by means of a button response; the responses are aggregated, screened for significant outliers at the host machine, and used as training data to improve the classifier.

To explore the novel infrastructure required by Open Mind, Stanford graduate student Chuck Lam and I have developed a Java- and Web-based version of *Animals*, an elementary interactive children's program for classifying animals, dating from the late 1970s. The child (e-citizen) thinks of an animal. Using the child's responses to a series of questions (for example, two-legged or four-legged?), the program tries to determine this animal's identity. If the program guesses wrong, the child must enter a question that distinguishes her animal from the program's guess. After a number of children have played this guessing game, *Animals* has learned a simple tree-based classifier for animals. While not a deep or particularly useful core program, *Animals* provides an excellent platform for solving important problems in Open Mind.

It provides an interface and protocol design for

- efficiently extracting the maximum information from e-citizens,
- detecting and eliminating significant errors and statistical "outliers,"
- repelling hostile attacks, and
- automatically listing contributors according to the amount of information they contributed.

The Open Mind Initiative differs from the Free Software Foundations and the Open Source organization in a number of ways. First, while Open Source draws its support almost entirely from a hacker culture (for example, roughly 10^5 programmers have contributed to *Linux*), Open Mind relies heavily on an e-citizen and business culture (10^9 nonprogrammers on the Web). While most of the work in Open Source is directly on the final released source code, most of the effort in Open Mind focuses on the tools, infrastructure, and data gathering. An expert arbitrates final decisions in Open Source; in Open Mind, much information is accepted or rejected automatically by the infrastructure software. Table 1 summarizes some of these differences.

While domain experts and infrastructure developers are likely to have the same motivations as contributors to Open Source, the

motivations of e-citizens deserve special consideration. E-citizens seek benefit from the resulting Open Mind software, including software that would be very difficult to develop in other ways, such as commonsense knowledge. E-citizens would enjoy game interfaces and seek the public recognition of their contributions. There could be financial incentives such as lotteries, discounts, or frequent-flyer awards provided by corporations seeking new customers.

Others have discussed components of Open Mind and its use in different contexts—online interactive data acquisition and voting, collaboration, and machine-learning and pattern-recognition algorithms—yet the integration proposed in Open Mind seems not to have been discussed. A number of existing projects would fit under an Open Mind umbrella and would profit from the initiative. One example is *Newhoo*, where nonspecialist e-citizens propose keyword and index information about Web pages. Their contributions are reviewed by volunteer referee and editors (currently 10,000) and made available to all. We can imagine Open Mind projects in numerous pattern-recognition domains or knowledge engineering to improve navigating news groups, Web sites, or FAQs. Later, these systems can be integrated, for instance, to incorporate natural language or commonsense constraints in speech recognition, OCR, or web searching.

Physics has had its atom smashers, microbiology its Human Genome Project, and aeronautics and astronautics its space missions. Now is the time for computer science and cognitive science to have their big science—one that harvests informal knowledge from a large number of e-citizens for building useful software for next-generation systems. Given the conjunction of several forces—the need for natural human-machine interfaces and improved Web searching, the existence of good learning algorithms and Web infrastructure, and the demonstrated success of the Open Source methodology—the time is right for the Open Mind Initiative.

Reference

1. D.G. Stork, "Character and Document Recognition in the Open Mind Initiative," *Proc. Int'l Conf. Document Analysis and Recognition (ICDAR '99)*, to be published by IEEE Computer Society Press, Los Alamitos, Calif., 1999.

Table 1. Comparisons between Open Source and Open Mind.

OPEN SOURCE	OPEN MIND
No e-citizens	E-citizens crucial
Expert knowledge	Informal knowledge
Machine learning irrelevant	Machine learning essential
Web useful but optional	Web essential
Most work is directly on the end-user software	Most work is on infrastructure, not on the end-user software
Hacker culture ($<10^5$)	E-citizen/business culture ($<10^9$)
Separate functions contributed (device drivers in Linux)	Single functional goal (for example, recognition rate in OCR)