

**To appear: Proc. Eighth Intl. Conf. on Risk and Gambling, London, July 1990**

# Could Gambling Save Science

## Encouraging an Honest Consensus

by [Robin Hanson](#)

Visiting Researcher, The Foresight Institute  
P.O. Box 61058, Palo Alto, CA 94306 USA  
[hanson@charon.arc.nasa.gov](mailto:hanson@charon.arc.nasa.gov) 415-651-7483

The pace of scientific progress may be hindered by the tendency of our academic institutions to reward being popular, rather than being right. A market-based alternative, where scientists more formally "stake their reputation", is presented here. It offers clear incentives to be careful and honest while contributing to a visible, self-consistent consensus on controversial (or routine) scientific questions. In addition, it allows funders to choose questions to be researched without choosing people or methods. The bulk of this paper is spent examining potential problems with the proposed approach. After this examination, the idea still seems plausible and worth further study.

### Introduction

After reviewing the discrepancy between what we want from academic institutions and what we get from current institutions, a market-based alternative called "idea futures" is suggested. It is described through both a set of specific scenarios and a set of detailed procedures. Over thirty possible problems and objections are examined in detail. Finally, a development strategy is outlined and the possible advantages are summarized.

### The Problem

**THE SCIENTIFIC REVOLUTION** Four centuries ago, some Europeans complained that the existing academic institutions were biased against them. Insiders, it was said, were "inflated by letters" and shunned anyone who dared "speculate on anything out of the common way" [De]. Outsiders -- astrologers, chemists, and people like Bacon and Galileo -- argued that they and their theories should be judged only by how well they agreed with observations, and not by how they agreed with the authorities of the day [Gal]. This was the age of utopias [Wh], as these rebels debated possible academic reforms and imagined whole new social institutions, for both academia in particular and society in general.

Within a century or so, the intellectual descendants of these outsiders became the new insiders

in a process now called the "Scientific Revolution". They introduced a new respect for observations along with new social institutions, such as the Royal Society of England, inspired by those utopian ideals. Since then science has made impressive progress. Most controversial issues of four centuries ago seem long settled by now, and continued research may well settle most of the today's controversies. Academia can claim some credit for this, and academic institutions have continued to evolve in response to perceived problems, formalizing publication in journals, credit in citations, and evaluation in anonymous peer review.

**PROBLEMS WITH ACADEMIA** Yet little has really changed. Academia is still largely a medieval guild, with a few powerful elites, many slave-like apprentices, and members who hold a monopoly on the research patronage of princes and the teaching of their sons. Outsiders still complain about bias, saying their evidence is ignored, and many observers [Gh,Re,Syk,Tu] have noted some long-standing problems with the research component of academia. (Teaching is not considered here.)

Peer review is just another popularity contest, inducing familiar political games; savvy players criticize outsiders, praise insiders, follow the fashions insiders indicate, and avoid subjects between or outside the familiar subjects. It can take surprisingly long for outright lying by insiders to be exposed [Re]. There are too few incentives to correct for cognitive [Kah] and social [My] biases, such as wishful thinking, overconfidence, anchoring [He], and preferring people with a background similar to your own.

Publication quantity is often the major measure of success. This encourages redundant publication of "smallest publishable units" by many co-authors. The need to have one's research appear original gives too little incentive to see if it has already been done elsewhere, as is often the case, and neglects efforts to integrate previous research.

Perhaps the core problem is that academics are rewarded mainly for telling a good story, rather than for being right. (By "right" I include not only being literally correct, but also being on the right track, or enabling work on the right track.) Publications, grants, and tenure are based what other insiders think today, independent of whether one's ideas and results are proved correct or valuable later. Even for researchers with a good track record, grant proposals must usually describe in some detail exactly what will be discovered and how; true exploratory work is done on the sly. This emphasis on story-telling rewards the eloquent, who know how to persuade by ignoring evidence that goes against their view, and by other standard tricks [Ci].

Admittedly, someone who has published an unusual idea that has proven right is thought of more highly, all else being equal. But all else is usually not equal. Outsiders find it hard to get an unusual idea published, and being able to say "I told you so" is of little help to academics who have failed to gain tenure. The powerful often get credit for the successes of those under them [Re]. Only in the most experimental fields, where feedback is direct and frequent, can we expect people who are disliked -- but usually right -- to be rewarded through informal reputations.

Perhaps our biggest problem is the distortion evident when a science question becomes relevant

for public policy, as in the recent debates over "Star Wars" or the greenhouse effect. The popular media tend to focus on those scientists prone to hyperbole. Any honest consensus of relevant experts is usually lost from public view, as advocates on each side accuse the other of bias and self-interest. Public policy can suffer dramatically as a result, a consequence that becomes more serious as the pace of technological change quickens.

On the whole, current academic institutions seem less than ideal [Ki], with incentives that reward being popular, fashionable, and eloquent, instead of being right.

**INCENTIVES MATTER** Are these complaints just sour grapes? Those who do well by an existing system tend to believe problems are minor. But even if the best ideas eventually win, we should worry if the people who advocate those ideas don't win. The social organization of any human effort can have a tremendous effect on its efficiency. For example, the heated debate over national health care is mostly about which way to fund and organize health care provides the best incentives to promote the general health. And different past cultures with different ways of organizing have had very different rates of scientific progress; compare Europe with China over the last five centuries. Our rate of progress may be less than 2% of what it could be [Be].

Are we wasting precious resources? Imagine what would happen if we used academic peer review to decide what products to manufacture. Proposals for new products would be reviewed anonymously by powerful people who produce similar products. These reviewers would pass judgement without taking any personal risk, and those judged favorably would win regardless of how useful their product turned out to be.

I much prefer our current business system, with all of its problems, where investors must take a personal risk when they endorse a product. Institutions like the stock market are comparatively egalitarian and flexible, allowing most anyone to participate in the ongoing debate about the profit potential of any public business or the relative potential of various industries, management styles, etc. Why can't we have academic research institutions like these?

**ACADEMIC REFORMS** Most efforts to improve academic institutions focus on incremental reform. Should reviewers be anonymous? Should submissions be anonymous? Occasionally someone proposes a more radical reform within the current framework, such as abolishing tenure or government funding [Fe], or scrapping the whole thing in favor of some existing alternative like patents. And once in a while a whole new social institution is proposed.

For example, science courts [Kan] (also called "scientific adversary procedures") were invented to blunt hyperbole on science controversies by using court-like proceedings to encourage cross-examination and to document areas of agreement. Hypertext publishing [Dr,Han88] imagines an advanced electronic publishing media where any critic could directly link a criticism to any published item, and where readers could combine the direct evaluations of previous readers they respect to decide what is worth reading. A recent suggestion [Ts] imagines governments paying private labs for each citation of one of their employee's papers, allowing a decentralized market to set research priorities. And prizes are often suggested as a way to fund specific verifiable

achievements like sequencing the human genome.

In this paper, I propose a new academic institution, tentatively called "idea futures". Also market-based, it is intended to counter many existing problems. It is utopian in the sense of describing a coherent vision of how things might be rather different, but hopefully practical in the sense of considering what could go wrong and how to start small.

**WHAT WE WANT** Before considering specific mechanisms, let us reflect a moment on what we want from academic incentives. We want to encourage honesty and fair play; the game should be open to anyone to prove him/herself. Funders of research, either private foundations or governments, presumably want research to be directed toward the academic subjects and questions of interest to those funders. (Funders also include the researchers themselves, to the extent that reduced salaries are understood to be in exchange for some research autonomy.) On controversial questions, we want a clear measure of the current opinion of relevant experts, a measure which political advocates could not easily distort. And those who contribute to such a measure should have clear incentives to be careful and honest.

Presumably we want as much progress as possible per effort invested, at least in situations where the following notion of "progress" makes sense. Consider a well-posed question, such as "Is the Earth basically spherical?", with a handful of possible answers (such as "No, its flat"). Experience indicates that, with enough study and evidence, one of the answers will eventually stand out as best to most anyone who considers the question carefully. At least this seems to happen for most questions that have been traditionally labeled "scientific"; questions about the morality of abortion or the nature of God may not fare as well. Where there is such a limiting "right" answer, "progress" can mean the rate at which general scientific opinion converges to that answer. This definition of progress hopefully avoids debates about whether more knowledge is good, or whether there is really an ultimate truth.

Translating these goals to an individual level, we want our institutions to reward academics for pushing scientific opinion toward the "right" answer, presumably by somehow increasing their reputation, influence, or resources. Let us imagine an academic who, after some reflection or observation, comes to a tentative conclusion which he/she would like others to consider. If most everyone already agrees with this conclusion, even without seeing the new supporting evidence or analysis, the academic should receive little credit for just making an "obvious" claim. However, credit should be possible if the claim is surprising, i.e., if people who have not yet seen the evidence are not yet willing to agree. If, upon reviewing the evidence, most everyone now agrees with the surprising claim, then the academic should certainly receive some credit. And, in fact, peer review can handle this case. But what if there is not uniform agreement? It still seems that the academic should be rewarded, if the claim is eventually born out. And others who supported the claim in the face of disagreement should also gain credit [Led], since they helped push the general opinion in the right direction.

Why shouldn't savvy academics now win credit by supporting as many claims as possible, or by multiplying controversies? Clearly they should risk losing credit when they are wrong, so that

credit is in some ways conserved. The ratio of possible loss to gain should depend on how unusual one's position is. Siding with the majority and being right should gain one less than siding with a minority and being right. The total amount gained or lost should depend on how much of their reputation each academic has chosen to stake on this issue, as well as on how interesting the issue is to the ultimate research funders.

In summary, part of what we want from academic incentives is a fair game for staking our reputation, so that on questions of interest to funders, we converge as fast as possible to the "right" answer.

## THE PROPOSAL

Surprising as it may seem, such a social institution exists. It is relatively simple, cheap, decentralized, and egalitarian. It could create a consensus on disputed science questions that would be clear, expert, honest, and self-consistent across a wide range of issues. This consensus should respond quickly to new information, and predict at least as well as any other co-existing consensus mechanism. It is well-grounded in our best theories of decision and incentives.

And it is ancient. We need only revive and embellish a suggestion made back during the utopian scientific revolution. Chemical physicians, excluded by the standard physicians from teaching in the British schools, repeatedly offered challenges like the following (circa 1651):

Oh ye Schooles. ... Let us take out of the hospitals, out of the Camps, or from elsewhere, 200, or 500 poor People, that have Fevers, Pleurisies, etc. Let us divide them into halfes, let us cast lots, that one halfe of them may fall to my share, and the other to yours; ... we shall see how many Funerals both of us shall have: But let the reward of the contention or wager, be 300 Florens, deposited on both sides: Here your business is decided. [De]

They proposed to bet on the question, apparently believing bets to be a useful augmentation of the existing academic incentives! Bets are a long-established and robust reputation mechanism, widely seen as a cure for excessive verbal wrangling; you "put your money where your mouth is". In science and elsewhere, phrases like "you bet" are standard ways to express confidence. Offers to make token bets are particularly compelling, and scientists of equal stature often make and publicize such bets, with recent bets on resource depletion, computer chess, black holes [Hal], solar neutrinos, and cold fusion [Gar,Lew,WSJ].

Consider the example of Piers Corbyn, a London astrophysicist who has been unable to get academic metrologists interested in his unusual theory of long-term weather cycles [NS]. Since June 1988 he has been making bets to gain publicity, betting against the bookmaker William Hill, who uses odds posted by the British Metrological Service. Over the last six months alone, he has won 80% of his 25 bets a month, gaining an over 90% average rate of return per bet. (There is a one in 200 to  $10^{20}$  chance, depending on what independence you assume, of this happening randomly.) Yet the Service still refuses to take Piers seriously, or make even token bets against

him. Which doesn't seem quite fair; hasn't Pier earned the right to be considered? William Hill has taken on the bets for the publicity, but is tired of losing, and has adjusted their odds accordingly. Why shouldn't these be the odds used for official British agricultural policy, instead of the Service's predictions?

If the primary way that academics are now rewarded for being right, rather than popular, is an informal process for staking their reputation, which has various biases because of its informality, and if we want a better reputation game, why not literally make bets and formalize the process?

Imagine a betting pool or market on most disputed science questions, with the going odds available to the popular media, and treated socially as the current academic consensus. Imagine that academics are expected to "put up or shut up" and accompany claims with at least token bets, and that statistics are collected on how well people do. Imagine that funding agencies subsidize pools on questions of interest to them, and that research labs pay for much of their research with winnings from previous pools. And imagine that anyone could play, either to take a stand on an important issue, or to insure against technological risk.

This would be an "idea futures" market, which I offer as an alternative to existing academic social institutions. Somewhat like a corn futures market, where one can bet on the future price of corn, here one bets on the future settlement of a present scientific controversy. This is admittedly an unusual suggestion. But consider what might happen.

## Scenarios

**CONTINENTAL DRIFT** In 1915 German meteorologist Alfred Wegener published his theory of continental drift, which he had collected extensive evidence in support of. But contemporaries considered his theory to be "impossible", and Wegener died an intellectual outcast in 1930 [Mar]. Yet in the 1960's his theory began to be taken seriously, and is now the established view. Wegener eventually gained fame, but overall academia seems to discourage activity like his. Some of Wegener's peers, for example, probably found his thesis plausible, but decided that to say so publicly would be a poor career move.

With idea futures, Wegener could have opened a market for people to bet on his theory, perhaps to be judged by some official body of geologists in a century. He could have then offered to bet a token amount at, say, 1-4 odds, in effect saying there was at least at 20% chance his claim would be vindicated. His opponents would have had to either accept this estimate, and its implications about the importance of Wegener's research, or bet enough to drive the market odds down to something a little closer to "impossible". They could not suppress Wegener merely by silence or ridicule.

As Wegener increased his stake, buying more bets to move the price back up, his opponents would hopefully think just a little more carefully before betting even more to move the price back down. Others might find it in their interest to support Wegener; anyone who thought the consensus odds were wrong would expect to make money by betting, and would thereby move

the consensus toward what they believe. Everyone would have a clear incentive to be careful and honest!

The market would encourage more research related to continental drift, as one could make money by being the first to trade on new relevant information. Eventually the evidence would more clearly tip in Wegener's favor, and the price of his bets would rise. Wegener, or his children, could then sell those bets and reap some rewards. While those rewards would not make up for years of neglect, at least he would get something.

As the controversy became settled, and opinions converged, people would gradually sell and leave the market. Few people, if any, need be left for the final judging, which could usually be avoided (using mechanisms to be described below).

**COLD FUSION** A more recent controversy began in March 1989, when Pons and Fleishman announced "fusion in a jar" at a dramatic press conference. In the months that followed, media aftershocks of confirmation attempts were tracked by thousands of scientists and others, who argued with each other about the chances of cold fusion being real. Proposals to bet came up often, even in the public debates. Critics, uncomfortable with airing scientific disputes in public, complained that Pons and Fleishman broke the rules by going to the popular media instead of through normal peer review channels, unfairly gaining extra attention and funding. Supporters countered that popular media spread information quickly to other scientists; cold fusion, if right, was too important to wait for normal channels.

In the journal *Science*, Robert Pool speculated that a market in cold fusion might have gone something like Figure 1 [Poo]. If there really had been a betting market, then there would have been a market price that journalists like Pool could publish as news. A table of going prices might appear on the science page in the newspaper, much like the stock page in the business section, conveying current scientific opinion better than the current "balanced" interviews with extremists on all sides. It's been suggested [Ze] that the added information in betting market prices might have helped resolve the debate more quickly

#### Figure 1 A Hypothetical Market in Cold Fusion

There needn't be a conflict between going through slow proper channels and getting the word out, if a fast market were a proper channel. The effect of staged media events might be reduced as it might not be news if the price didn't change; advocates would have to convince, not the average listener, but those people willing to make bets. Remaining biases, such as the overconfidence evident in figure 1, would be reduced by technical and other trading specialists.

Cold fusion businesses would have been less risky to start. As it was, a new fusion business had to bet both that cold fusion was real, and that they were the best group to develop and market it in that case. With idea futures they could, by both starting a business and betting against cold fusion (essentially taking out insurance), really only be betting on their ability to develop cold fusion if it were real.

Insights from a great many people whose opinions on the cold fusion controversy were ignored, such as inarticulate people and those without Ph.Ds, could have been integrated in a decentralized manner. Popular play would end up subsidizing professional efforts on questions of popular interest, offering more "direct democracy" in setting research priorities.

**NEUTRINO MASS** Betting markets could also function in the absence of overt controversy, as in the following (hypothetical) story.

Once upon a time the Great Science Foundation decided it would be a "good thing" to know the mass of the electron neutrino. Instead of trying to figure out who would be a good person to work on this, or what a good research strategy would be, they decided to just subsidize betting markets on the neutrino mass. They spent millions.

Soon the market odds were about 5% that the mass was above 0.1 eV, and Gung Ho Labs became intrigued by the profits to be made. They estimated that for about \$300K spent on two researchers over 3 years, they could make a high confidence measurement of whether the mass was above 0.1 eV. So they went ahead with the project, and later got their result, which they kept very secret. While the market now estimated the chance of a mass over 0.1 eV at 4%, their experiment said the chance was at most 0.1%.

So they quietly bought bets against a high mass, moving the price down to 2.5% in the process. They then revealed their results to the world, and tried their best to convince people that their experiment was solid. After a few months they mostly succeeded, and when the price had dropped to 0.7% they began to sell their bets they had made. They made \$400K off of the information they had created, which more than covered their expenses to get that information.

Of course if Gung Ho Labs had failed to convince the world of their results, they would have faced the difficult choice of quitting at a loss, or holding out for the long-term. No doubt a careful internal review would be conducted before making such a decision.

Gung Ho would be free to use peer review, tenure, and fixed salaries internally, if they are effective ways to organize workers. The two researchers need not risk their life savings to be paid for their efforts. But the discipline of the external market should keep these internal institutions from degenerating into mere popularity contests.

**KILLER PEANUT BUTTER** Once upon another time, Munchem Biolabs found compelling evidence that peanut butter was more deadly than most pesticides, a conclusion that Lunch Industries Exclusive (LIE) wanted desperately to suppress. LIE's usual procedure was to fund a bunch of competing studies to come to opposite conclusions, which usually kept the waters muddy enough that legislators and customers would ignore it all. But this time they had to deal with an idea futures market on the question, and the public was beginning to take the odds in such markets seriously.

Munchem had moved the market odds of deadly peanut butter up rather high. LIE now had two choices; either they could use overwhelming cash to move the odds back down, or use



competing studies, advertising, etc. to persuade others to bet on their side.

If they bet alone, they would know they were throwing their money away with no obvious limit on future spending. Not only might Munchem find allies, but LIE employees who knew they were bluffing might be tempted to pick up a little free money with some anonymous bets. If word of Lunch's bluff got out, as insider information often does, investors would flock in and wipe out the effect of LIE's bets.

If LIE tried to throw away other people's money through a persuasion campaign, they would face a market dominated, as most liquid markets are, by battle-hardened speculators. These investors, not easily persuaded by clever jingles, would quickly hook up with research insiders, who generally know which labs tend to find whatever results their funders want.

So in the end, Lunch Industries accepted the market odds, and began research on non-toxic peanut butter.

## Procedures

Rather than just present an abstract utopian vision of market-based academic incentives, this paper aims to consider in some detail what problems might arise and possible approaches for dealing with them. The following is a core set of procedures tentatively selected to best deal with known problems, a core that will be expanded upon later in the paper. No doubt, experience with real idea futures markets will show many of these suggestions to have been naive. I offer them primarily to make plausible the idea that betting markets could be applied to a much wider range of scientific questions than is presently considered feasible. (This section is somewhat dense, and may be profitably skimmed on a first reading.)

**ASSETS** Imagine that John bets Mary \$5, at even odds, that it will rain next Monday. Since they don't entirely trust each other, John and Mary put the bet in writing and each give \$5 to Frank, a trusted third party. John has essentially paid \$5 for an I.O.U. that says "Worth \$10 If Rain Monday", since if he wins he gets \$5 from Mary and his own \$5 back. Mary's I.O.U. says "Worth \$10 If Not Rain Monday". On Tuesday one of them can cash in their I.O.U. for \$10 from Frank.

This standard betting scenario can be improved by breaking it into different transactions; first create the I.O.U.s and then sell them. Replace Frank with a stable financial institution, let's call it a "bank", which will sell a pair of "\$10 if rain", "\$10 if not rain" coupons to anyone for a price of \$10. The bank takes no risk, since exactly one of the coupons will be worth \$10 in the end. And since the bank holds the \$10 in the meantime, it can afford to offer interest on the \$10, and perhaps pay a local meteorologist to be an impartial judge. Now Mary can first buy a coupon pair from the bank for \$10 and then offer to sell her "\$10 if rain" coupon to John or anyone for \$5, retaining the "\$10 if not rain" for herself.

A central clearinghouse for such offers, which matched compatible offers and insured that people made good on their offers, would always hold a best current offer to sell and to buy. If the

transaction costs of processing an offer through the clearinghouse were small, as current technology allows, then the "spread" between these offers could be quite small, leaving a going "market price". A going price of \$3.20 for "\$10 if rain Monday" would represent a temporary consensus of a 32% chance of rain Monday.

In general, these markets trade assets of the form "X if A" (called "contingent assets"), where X is some pre-existing "base" asset and A is one of a set of mutually exclusive claims that some judging organization agrees to eventually choose from. The base X can be any stock, bond, currency, commodity, or even another compatible contingent asset. The set of claims is a "question", and each claim is one possible answer to the question. To trade on a question, we need an agreement between several parties - an author, a judge, and one or more banks, registries, clearinghouses, and randomness checkers.

An author carefully words a set of claims, and a judging organization agrees to, if necessary, offer a verdict in favor of one of these claims at some, perhaps indirectly specified, date. Registries hold records of public, i.e. not anonymous, trades made at clearinghouses. (Clearinghouses may be required to hold additional private records of all trades, to be subpoenaed by criminal investigators if necessary.)

Consider a question with possible answers {A,B,...}. Any bank authorized in the agreement on that question can "split" any allowed base X (usually anything) into the assets {"X if A", "X if B", ...}, or "join" those assets back into X. In the example above, \$10 was split into "\$10 if rain" and "\$10 if not rain". The bank is trusted to report the net effect of these transactions to a central agent, who keeps track of the net "market capital" that has been split along this question.

On the specified date, and after a short warning period, the judges are given an agreed-upon judging-fee in order to study the question and render their verdict. Verdicts assign a percentage of validity to each of the possible question answers. If the verdict is 98% in favor of A, then banks are authorized to let people exchange their "X if A" assets for 98% of X.

The judging-fee is obtained from the banks, who devalue the current assets contingent on that question by some percentage, a percentage which can be no more than a pre-specified max-judging-percentage. This devaluation creates an incentive for traders to "settle out of court" and sell before the judging date.

What if there is too little capital in the market to support the required judging fee? John and Mary's market only has \$10 in it, and with a 10% max-judging-fee, only \$1 is available for judging, short of the \$5 the meteorologist judge requires. In this case we can hold an "audit lottery" [Pol]. The current market capital, \$10, is gambled with whomever offers the best price, among those approved by the randomness checker. If the gamble is won, every asset contingent on this question increases in value, resulting in enough market capital for judging to proceed, in this case \$50. If the gamble is lost, all such assets become worthless and judging is not needed. Investors can insure against the added risk audit lotteries impose by putting money into a pot to be gambled in the same lottery, but on the other side.

Judges can be given more flexibility to deal better with uncertainties regarding when a question will be judgeable and how much that will cost. The max-judging-percentage can be spent in discrete units, each specified by a percentage-unit and a fee-unit. After spending each percentage-unit, the judges could choose to postpone judging to a later date and/or raise the next fee-unit. If necessary, an audit lottery would be held before each new unit.

If desired, judges can also be given a direct financial incentive to be careful and honest. "Appeals" markets can be created on the same question, but judged by an independent group much later and/or with a much higher judging-fee. For a limited period after a verdict is announced, an amount, up to a fixed fraction of the original judging-fee, would be spent trying to move the price in the appeals market toward the verdict specified. Judges would end up with some contingent assets saying their verdict would be upheld in the appeals market, assets they could sell immediately, at a loss, if they so chose.

Idea futures markets need no central management. Anyone could author a claim on any subject of interest to them, contract with different judging groups to judge that claim on different dates, and allow different banks to deal in each question. And anyone should be able to open a clearinghouse to sell any asset. All of these groups could compete openly for the attention and respect of investors.

**INVESTORS** Investors could be as diverse as they are in current markets, each focusing on some specialty while avoiding risk from other areas. For example, if the market odds are "incoherent", i.e., deviate from the standard axioms of probability, a trader who corrects that deviation can make better than the average rate of return without significant risk. Therefore coherence specialists should keep the market consensus roughly consistent over a wide range of subjects. Similarly, technical traders would keep the pattern of price changes close to the ideal random walk [Ma]. The market odds should also quickly reflect information contained in any co-existing consensus measures, such as opinion polls or reports of elite committees, as traders could make easy money if alternative measures were reliably better predictors than the market.

A contingent asset, like "X if F", that is split again creates conjunctive contingent assets like "X if F and A". Conjuncts which combine a great many claims should be popular, since they offer investors the greatest expected return. Conjunctive assets also allow one to bet the conditional probability of A given F and remain insensitive to the verdict on F. In this way diverse traders, each of whom has only local knowledge, could manage a large network of dependencies such as the currently popular "Bayes net" models [Pe].

**SOCIAL ATTITUDES** Some new social attitudes toward these new markets are important elements of the envisioned approach. As with current financial markets, the market odds should be treated as the current social consensus on a question by popular media and policy makers. While one may of course disagree with this consensus in conversation, it is not impolite for others to inquire whether one who so disagrees has made investments commensurate with their wealth and the fuss they are making. People who do so invest should receive the same sort of social credit now granted to "do-gooder" advocates who devote personal resources to changing current

opinion on some important issue. Like Pheleas Fogg, the hero of Verne's *Around the World in Eighty Days*, "a man who rather laid wagers for honor's sake than for the stake proposed" [Ve], these investors should not be treated as mere risk-loving gamblers.

Social credit should also go to philanthropists who choose to subsidize a market on some important question. By funding an automatic inventory-based [St] market-maker, which always offers to buy or sell at prices determined solely by its current inventory, one gives away money only to those who move the market price in the direction of its final verdict.

Reputation scores can be computed from each person's public trades, recorded at registries. A trade is considered "public" if the trader committed at trading time to a date at which the trade would be publicly revealed, and that date has passed. One simple reputation score would be the ratio of the current market value of assets held to their value when purchased, corrected for a few distortions. People with high reputation scores should be respected for having been right against the crowd, and such scores might even compete with G.P.A.s or number of papers published as an evaluation measure.

## Objections

The main difference between "blue sky" fantasies and serious but radical suggestions is in how well they handle the details. If you are like most readers, you will by now have thought of one or more problems with or objections to idea futures. If so, you are encouraged to scan this section and go directly to the issues of concern to you. (Most of these issues have been raised by at least three independent commentators in previous discussions.)

**ISN'T GAMBLING ILLEGAL?** Yes, betting markets on science questions appear to be only legal in Great Britain, where they are highly regulated. Even Nevada, which allows sports betting, prohibits general betting to avoid scandals that might "taint" the gambling industry. Which is a shame because most of the arguments against betting, discussed below, do not apply well to science betting. We allow scattered markets that give us rather good consensus estimates on horse races and football teams, yet not on important science and technology questions! In the long term perhaps we can persuade legislators to allow science bets because of their extra benefits and reduced problems. Science betting certainly seems easier to justify than the currently popular regressive taxation through state lotteries.

**ISN'T BETTING A USELESS ZERO-SUM GAME?** A standard argument for making betting illegal is to keep people from wasting their energies in unproductive activities. The only obvious value in betting on dice throws is entertainment, but laws to prohibit this usually also prohibit much more. Life insurance and commodity futures markets [Ro] were both prohibited by anti-gambling laws until advocates managed to obtain exemptions.

Being monetarily zero sum does not make betting useless. Betting markets allow traders to reduce risk, and create informative prices. In liquid markets most of the trading, liquidity, and price rationalization comes from speculators, for whom the market is basically a betting game.

Buying any particular stock in the stock market, for example, is basically a bet in a zero-sum game when compared to investing in the standard "market" combination of all assets in the same tax and risk category. (While, if the prices are irrational, such bets may help the economy as a whole, this "externality" also benefits people not betting on that question.)

In fact, a standard way to analyze financial portfolios is to break them into contingent assets, each of which has value in only one possible world [ShW]. A "complete" market, where one can bet on anything, is best, allowing investors to minimize risk and maximize expected return [La].

Science bets would not only allow corporations to more easily insure against technological risk, but they would create prices embodying the sort of valuable information that governments now fund research to obtain. When the betting stakes are invested in stocks, the money is hopefully being put into productive use by those companies. Therefore, ignoring transaction costs and judging fees, the average rate of return of contingent assets split from stocks would be the same as the return on those stocks.

**DOES ANYBODY EVER BET THIS WAY?** Liquid markets in contingent assets are a somewhat different betting mechanism from the usual bookies or pari-mutuels. But they are not untried. Such markets are widely used to teach MBA students about how markets work [Fo], and are usually done on elections. Some financial traders have used them to bet on basketball tournaments [Pow]. And I have developed a board game where players use such a market to bet on a murder mystery as it unfolds. Most players learn the mechanism very quickly, and have lots of fun!

**WHAT ABOUT COMPULSIVE GAMBLING?** About 2% of the population seems unable to resist the temptation to risk more than they can afford to lose [APA] in casinos, racetracks, and high risk financial markets. Lost in the thrill of "action" and the hope that all of their financial worries will soon be over, they often regret their excess later, and resort to desperate measures, like theft, to pay debts.

Compulsive gambling is encouraged by advertising and easy access to games with a quick and possibly large payoff. British law reduces this problem by requiring casino players to apply 48 hours in advance, by allowing them to sign up on lists of people to be excluded from all casinos, and by forbidding youth and on-site alcohol, entertainment, and credit [Ke]. Margin limits in financial markets serve some similar functions.

Governments may impose similar rules to discourage compulsive gambling in idea futures, though it is important that any advertising restrictions not prevent the wide dissemination of current market prices. More importantly, unless options (or investments on margin) are offered, science questions are generally too long term to be a problem, offering no more "action" than long-term stock investments. Traders who regret their purchase a few days later can sell and get most of their money back. And, given that many other options markets exist, it is not clear that allowing science options would increase opportunities for compulsive risky investing.

**IS THERE ENOUGH INTEREST IN SCIENCE QUESTIONS?** A recent science fiction [Br] novel imagined wide-spread betting on science and technology questions, supplanting horse racing in popularity. And it is possible that having a direct, if small, influence and personal stake in science would heighten the public's interest. At present, though, fewer people probably follow science than football.

We don't need to interest everyone, however, just enough to pay for the modest overheads involved. Few people have interest and opinions about the future price of corn, yet corn futures markets thrive. A great many people are now involved in scientific research, many more follow scientific journals, and even more follow science in the popular media. Many of these people have strong opinions on various science controversies and feel they have insufficient opportunity to express them. Idea futures would thrive if it tapped only a small fraction of current interest and effort.

Having a fraction of science funding channeled through betting markets would certainly accomplish this. So might basic attitude changes toward seeing markets as a legitimate place to "take a stand" on important issues, trading scores as indicators of who is right more often, and the market price as a valid consensus measure. Idea futures does not need large sums of money to be successful; even when there is only \$100 bet on a question, the market still offers the social benefit of a visible consensus and incentives for honesty.

**WILL THESE MARKETS BE TOO THIN?** In a market with low "liquidity", there are so few traders that you have to wait a while to find someone willing to trade with you. Automated market-makers [Hak], always ready to trade at prices determined by their current inventory, can increase liquidity and maintain a small "spread" between their buy and sell price offers. And they can be very cheap if the basic transaction costs are low, which they could be if thousands of markets shared the same computerized market place.

But the market might remain "thin" in the sense that prices could change quickly against a trader in response to each small amount traded, so they would have to wait to get a "reasonable" price. A lack of expected market thickness can be a self-fulfilling prophecy, since traders prefer thick markets [Ec]. This is a standard explanation for the limited number of futures and options markets currently available. On the other hand, thin markets are known for being good places to find overlooked bargains, and are less prone to speculative bubbles (a single rational person can squash one).

A thin idea futures market may actually seem better to some people, as the cost to change the current market consensus would be less. But a thicker market seems preferable over all. Funding channeled through market-makers would of course thicken the markets, as would consistency arbitrage and conditional offers that connect questions. Two people making a bet is a very thin market, but it happens all the time.

**DOESN'T BETTING ONLY WORK FOR CLEAR CUT QUESTIONS LIKE HORSE RACES?**

Most organized betting focuses on questions which, like sporting events, will become very clearly

resolved in a fixed time. This minimizes disputed verdicts and judging costs, and so it makes sense for risk and entertainment seeking bettors to focus on such subjects. But this does not imply that, given a specific subject area, betting markets are not a reasonable alternative to other consensus, reputation, and incentive mechanisms. Any incentive mechanism must pick some arbiter of quality, and subjects that are difficult for bets are also difficult for other approaches. For example, peer review, which uses averages of anonymous expert reviews as a quality measure, is widely believed to work better in the "hard sciences" than elsewhere.

Most scientific controversies seem to eventually get resolved enough to settle a bet. This resolvability is in fact central to popular notions of what defines science. Scientific claims are often defined as claims of "fact" which future evidence could possibly disprove [Pop], or at least alter our degree of confidence in. And science is widely believed to be "progressive", so that as evidence accumulates and relevant studies continue, opinions gradually converge. Beautiful theories killed by ugly facts are left behind.

Actually most people believe that opinions on most questions of fact usually converge with time, evidence, and sincere study. We hope that history will prove us right. We debate and discuss, essentially saying "I'll bet if we talked it out, you'd see I'm right". We take the advice of experts, indicating that we think we would come to believe what the experts believe, if only we were to study what the experts have studied.

Even if we aren't sure whether opinions will converge, we think there is a good chance they would converge if only a knowledgeable and detached enough group would spend enough effort to study and debate the question. And if that group is diverse and independent enough, we believe we would probably agree with them. If so, we should accept their verdict to settle a bet.

**HOW OFTEN DO BELIEFS REALLY CONVERGE?** Just because people believe their opinions converge, doesn't mean that they do. After all, there are strong social reasons to want to believe in convergence. Even if most questions that are settled today were once controversial, this doesn't mean that most old controversies are now settled. Perhaps yesterday's questions referred to concepts that are not even considered to make sense today. Historical studies, examining random scientific questions and claims of several centuries ago, should be done to shed light on these doubts.

But there are reasons to be optimistic. Standard decision theory, though it does not adequately account for the computational costs of deducing the implications of theories and evidence, is instructive and indicates that rational agents should come to agree [Se]. Consider an ideal decision theory agent who has a degree of belief in some particular claim A and continues to observe new evidence. Asymptotically, either all new evidence will be irrelevant and have no bearing on A, or the agent will become certain about whether A is true or false. Now imagine that the claim A specifies a detailed possible world, i.e. says that the real world is one particular world out of the many possible worlds. If two ideal agents start out with wildly different beliefs, but neither of them is completely certain about A, and if they both observe the same not asymptotically- irrelevant evidence, then they will asymptotically come to agree about A.

Studies indicate that people also have strong tendencies to conform and agree when exposed to each others opinions [Li] and arguments [My]. In fact, the rate at which they come to agree often seems faster than can be rationally justified by decision theory. Randomly selected legal juries usually come to a unanimous verdict on complex legal questions.

**WHAT IF BELIEFS NEVER CONVERGE?** Even if beliefs usually converged, idea futures might be unworkable if it dealt badly enough with situations where beliefs don't converge. One approach is to have mutually exclusive claim sets include a "this question too vague to judge" claim which the judges could choose if it seemed clear that no amount of study or time would ever allow a choice between the rest. Most people could then bet on the question conditional on it being resolved. This solution fails, however, if sincere beliefs never converge and yet it never becomes clear whether or not beliefs will converge. A deadline by which a question must be resolved could deal with this, but has other disadvantages.

If investors can reasonably estimate the chances that a question will be unresolvable in this manner, then the problem is manageable. High-risk questions will only be traded if there is enough disagreement [Ja] or subsidies to justify it, and for low-risk questions the problem can be ignored. And, it seems, resolvability can be estimated. Questions about religion and morals are more difficult, as are certain long-standing riddles like the nature of consciousness. On the other hand, a question about a physical property of a substance, like a bond angle of some new molecule, seems quite resolvable. As a rule, one should prefer questions closer to direct observations. And general claims for which relevant evidence will always be available should do better than claims like what someone had for breakfast ten years ago.

**WHAT DO CONVERGENT BELIEFS HAVE TO DO WITH TRUTH?** The philosopher Peirce claimed that "The opinion which is fated to be ultimately agreed to by all who investigate, is what we mean by the truth" [Th]. However, the question of whether the convergent opinion we might all come to with unlimited evidence, study, and debate is the way the world "really" is, is beyond the scope of the paper. Even if it isn't "truth", we are all interested in it, and it's hard to think of a better truth-estimate on which to base academic incentives.

**WHAT ABOUT BADLY WORDED CLAIMS?** Even if an issue becomes settled, a poorly worded claim on that issue may be unresolvable. To avoid this, we need techniques for avoiding ambiguity and incentives for players to use them.

Wording a claim so it is both relevant to some important issue and minimally ambiguous is a skill that is routinely learned in many professions. Lawyers and philosophers obtain clarity through standardized words and language, and scientists are adept at finding connections between abstract theories and specific observations. Claims should avoid slippery concepts and phrasing which allows many interpretations. Verbose annotations can also help by discussing motivations, examples, intended word meanings, judging criteria, etc.

If copyright laws are interpreted as applying to claim wordings, then claim authors may be able to charge an extra royalty fee for each join. Claim authors would then compete with each other for



royalties from investors, who would prefer authors with reputations for writing clear and interesting claims. Added incentives come if authors bet against their claim being judged too vague.

To avoid excessive costs in forming a claim, a question could hold a "clarification lottery". After a certain time, or when the market capital reached a certain amount, judges could be funded in the usual manner to replace a hastily worded claim with a more considered one.

Even when one cannot really word a good claim to bet on directly, markets offer other ways to bet on a subject. For example, if one believed that when physicists disagree with chemists, the chemists are usually right, one could invest in a "basket" or mutual fund which bets on the side of chemists in as many controversies as possible.

**CAN'T WRONG IDEAS STILL BE USEFUL?** Absolutely. If you think an idea is probably wrong, but is probably more like the right answer than anything else around, then bet on that. If you just think that work on the idea is likely to inspire something interesting, then bet on that. These questions will be harder to judge though.

**WHAT IF THE FINE PRINT DIFFERS FROM THE SUMMARY?** Verbose claims would probably be described by short summary sentences or phrases in price lists, offers, etc. As with contracts and political ballot initiatives, there are problems when a deceptive title differs from the fine print. In extreme cases people might sue for misrepresentation, but usually we can only encourage the buyer to beware.

**WHAT ABOUT SUCKER BETS?** If a stranger offers to bet you on an oddball subject, there is a good chance they are trying to trick you with a deceptive claim. Even if it looks like you couldn't lose, you are well- advised to decline; the fact that they are making an offer gives you information.

In markets on pre-existing controversies where many traders have already examined the claims, this is less likely, though still possible. In general, traders should look claims over carefully and not bet unless they honestly think they know better than than the other traders.

**DON'T SCIENCE QUESTIONS RESOLVE TOO SLOWLY?** The fundamental questions that get people interested in science, such as whether the universe is infinite, can take decades or even centuries to resolve. But this does not prevent markets in such questions. Most any newspaper will show that people regularly buy bonds scheduled to mature in forty years. Fifty year-olds who buy such bonds are not counting on living to be ninety; they know they can sell the bonds in the market at any time.

At present, you usually can't get a Ph.D. on whether the universe is infinite; you focus instead on a smaller question that is hopefully relevant for the bigger ones. Idea futures investors will similarly prefer shorter-term questions. A question that takes ten years to resolve (say starting at 50/50 and ending more than 90% certain 90% of the time) should have the same sort of daily price fluctuations (around 1.5%) as stocks do, and so support a similar mix of short-term speculators, and long-term fundamentals-oriented investors.

But for longer-term questions, investing in fundamentals is less attractive. Less information comes out per unit time in a long-term market, so there is less money to be made for a given market thickness. And if you must hold out for decades until other investors come to their senses, the extra rate of return above the market average that you get for your information may be very small, and so you may prefer to quit now if you have better opportunities elsewhere. To make things worse, this creates an opportunity for strategic behavior. Someone might move the price in some direction and try to hold it there in the hope that other traders will not be willing to hold out as long and therefore quit at a loss.

Finally, you may not trust the underlying financial institutions to remain stable over a century or more. Few people would probably bet that "Nuclear war will destroy most of civilization", even though many people would like to for insurance reasons. Even if the banks don't go bankrupt, uncertainties about the relative long-term value of different base assets the betting stakes could be invested in may completely swamp any added return from winning a bet. This problem might be minimized if the "market asset" [ShW], a maximally diversified world mutual fund, became the standard base asset.

Even with all these problems, there will probably be rather thick and well subsidized markets on a few very basic science questions, as funding agencies and amateurs seeking to influence important issues would focus on them. Such questions could be connected, through a network of conditional offers, to related shorter-term questions which research could more directly resolve, allowing researchers of simpler questions to obtain some of the subsidies on the basic questions.

In financial markets, the conventional wisdom is that longer-term price movements are less rational, as there is less incentive to correct irrational deviations. But there is still some incentive, and so idea futures may still offer an improvement over the existing situation.

**WHY SHOULD I TRUST THE JUDGES?** Even when sincere opinions would converge, investors may worry about judges being biased by bribes or various shared interests and associations. Fortunately, investors get to pick the assets they buy, and therefore the judges they bet on. So they can prefer long-lived judging organizations with reputations for fairness and avoiding scandals, and which use various available means to discourage foul play.

Incentives for traders to settle out of court and avoid judging altogether certainly help avoid judging foul play. So do clear-cut claims and judging criteria that leave little room for judging discretion. If we wait so long that the right verdict becomes "obvious" it would also be hard for judges to cheat. Also more trustworthy are juries of people who have never had a stake in the question, randomly selected from a large population, deliberating openly and offering to consider any relevant evidence.

The question of whether some proposed evidence is relevant for some deliberation could have its own betting market. If so, then juries could offer to consider any evidence for which the market odds of relevance were above some threshold.

Incentives to detect foul play could come from both the ability to sue cheating judges, and possibly from large bonds which judges might post payable to anyone who uncovers such corruption. Also, any persistent difference in the market odds on the same claim with different judges would constitute consensus about judging bias, flagging those judges for closer scrutiny. Judge rating agencies might form. Finally, "appeals" markets can give judges a direct incentive to be careful and honest, since judges must then bet that their verdict will be upheld on appeal.

**WON'T JUDGING COST TOO MUCH?** Through audit lotteries, one can keep the percentage taken by judges below any given threshold, and still afford to pay for very detailed judging, even going so far as to choose many jurors from widely different cultures and train them in one or more specialties before having them adjudicate some specific issue! This approach is mainly limited by risk aversion, which limits the attractiveness of large wins. Most people will not want to bet so much on any one question that the amount they might win would be much more than their total wealth. A one in a billion chance of winning a billion dollars is not worth as much to most people as a one in a thousand chance of winning a thousand dollars. If the amount one would need to bet to avoid this effect is too small, it is not worth the bother and people will bet nothing on the question.

**WON'T WEALTHY PEOPLE HAVE TOO MUCH INFLUENCE?** Markets are not opinion polls where the rich get more votes; to use market influence one must risk losing it. As in existing financial markets, rich investors who are not specialists in some particular area will prefer to get investment advice from someone who is a specialist, or avoid investing in that area entirely. This is similar to the way that powerful people defer to academic specialists now. Rich people who carelessly throw their weight around will lose their riches.

Even so, the wealthier social classes will have more influence, as they do now in most areas of life, including academia. If this is a problem which you are willing to invoke the force of government to solve (I am reluctant to do so), then the natural solution is general wealth redistribution. This is economically much more efficient than trying to crudely keep the rich out of any particular walk of life.

If you worry that markets would create large inequalities in academia, don't. Influence in academia, as measured for example by number of papers published [Pr], is far more concentrated than in most walks of life. It seems unlikely that markets would make things worse, and could well make things much better, as people would not need degrees or the blessing of the academic elites to play as equals.

**WON'T THE MARKET BE DOMINATED BY FOOLS?** Again, markets are not opinion polls. Anyone can invest in any open market, but they only choose to invest where they think they have special insight or insurance needs. Even if they are mistaken about their special insights into, say, the gold market, they are fairly quickly taught otherwise. Most people who play commodity markets, for example, lose their stake and quit within a year. Such markets are dominated by the minority who have managed to play and not go broke. If you believe otherwise, and know of some market where the prices are obviously wrong, I suggest that you to "put your money where

your mouth is" and take some of that free money you believe is sitting there. It's easy to bad-mouth the stupid public before you have tried to beat them.

**WON'T ADVERTISING MANIPULATE OPINION?** Advertising, in the sense of campaigns to persuade through evidence and arguments, exists now in academia and would certainly persist. Advertising, in the sense of clever jingles and sex appeal to grab the subconscious of the impulse buyer, should not be a problem, except perhaps if idea futures assets became a popular fashion accessory. People do not try to affect the price of corn futures with clever jingles; it would be like trying to sell cars by offering free balloons to Consumer Reports technicians. The savvy investors who dominate markets are smarter than that.

**AREN'T MARKETS FULL OF CHEATS AND THIEVES?** Yes, but this does not usually distort the incentives or the consensus price much. Most cheating is not "manipulating the price", which is rather hard to do in a liquid market, but conflicts of interest where someone who supposedly represents someone else uses information gained to act in their own interest.

Insider trading is mentioned below. But brokers and investment advisors are the worst case. In markets you win whenever you can get others to do what you just did, or when you predict what they will do and do it first. Brokers and investment advisors often tell you to buy whatever they would like to sell, and charge you large commissions for the "advice". Brokers often trade for themselves just before they execute trades for you; stop orders and margin calls are especially lucrative.

To avoid being cheated, be careful who you trust. Avoid brokers who trade for themselves, and advisors who do not take the same risk they advise for you.

As bets, idea futures markets cannot be cornered or monopolized. No matter how many bets have been made, other people are always free to bet more.

**WHAT ABOUT INSIDER TRADING?** When an employee of a company makes money by trading on inside information they have about that company, or by telling someone else so they can trade, that employee is considered to be going against the interest of the other stockholders who own the company. Employment contracts and laws can forbid this conflict of interest, though price movements just before major announcements show that a substantial amount of such trading happens anyway.

Fortunately nature has no insiders or employees. The only similar problem in idea futures is when a research lab is trying to keep a result temporarily secret before trading on it, and an employee sneaks out and trades first. This can be dealt with exactly as if it were stock insider trading, through private trading records accessible to criminal investigators.

What About "Moral Hazard"? One of the advantages of a market is that it offers incentives to anyone to come and contribute their knowledge about the world. A disadvantage is that, since changing the world can give one special knowledge about it, people may have an incentive to cause harm. If we allow anyone to bet on your lifespan, then someone may decide to kill you just

to win a bet. And this murder may be much harder to solve than most since, with anonymous trading, most anyone might be a potential suspect. (Though criminal investigators should be able to learn who really made what "anonymous" trades.) For this reason, there are usually restrictions on who may buy how much life insurance on you.

Moral hazard should be less of a problem for basic questions about nature that people cannot change, though it could conceivably be a problem for short-term trading and options that bet on when information will come out. We wouldn't want someone to blow up the latest accelerator to prevent information from coming out, or to kill some patients to slant a medical study.

Yet we shouldn't prevent open markets if the chance of foul play seems small. Anyone is allowed to trade stock, even though there is a possibility that someone will sell short the stock of the makers of Tylenol, and then poison some packages to depress Tylenol sales. Only for the rare claim where the risk of harm seemed particularly high might one justify a prior restraint limiting who could have how much stake on the different sides of a question.

**WHAT ABOUT INCENTIVES TO START FALSE RUMORS?** A "rumor" is just information, perhaps false, passed informally through a social network. Maliciously false rumors occur whenever people both have an interest in what other socially connected people think about a question, and when there is inadequate feedback for learning what rumors were false, so that people can discount unreliable sources.

In current academia, there is often enough feedback to discourage false rumors about what results are about to be published. Word of mouth which discredits a person, however, can keep him or her out without others ever really finding out if the rumor was right.

Markets both encourage and discourage false rumors. Markets give more people an interest in fooling other people, but also improve the feedback about what rumors were right. And the market price offers an alternative to informal information channels. Again, don't believe everything you hear; trust advisors with a good track record who take the same risk they advise you to.

**WHAT ABOUT INCENTIVES TO KEEP INFORMATION SECRET?** If you acquired a piece of information where it was clear which side of what questions the information favored, then your best strategy would be to buy on those sides, reveal and publicize the information (perhaps after selling it to other traders), wait for the price to rise, and then sell at a profit. If, however, the implications of the information are not clear, you might be tempted to sit tight and wait for further revelations, even though you risk other people stumbling on to your insight in the meantime. It is similar with incentives to publish. Unless you can connect your insight to currently popular issues, and package enough of them together to make a paper, you cannot get published and so you keep the idea to yourself.

One approach might be to formulate a question more closely related to your information, and then try to convince some funding agency that your question is interesting, even if its implications are

not clear. Or you could subsidize your question, in the hope that this would encourage others to figure out its implications and create conditional offers connecting it to other questions. Either approach might induce enough market thickness to make your information pay off.

**WON'T AN APPARENT CONSENSUS CREATE A CROWD MENTALITY?** People might think they agreed more than they actually did, defer to a consensus that had little thought behind it, and so create the social analogues of anchoring and overconfidence [Kah]. Would creativity be suppressed?

Markets with less thought behind them should give themselves away by being thinner. If not, and some of us catch wind of this trend, we could make money by correcting for it. And, for what it's worth, the market odds at horse races actually tend to be underconfident, being biased toward long- shots. Markets encourage people to be contrarian in the sense that by making any trade one is saying the consensus is probably wrong.

**WILL THE NEW INCENTIVES SLOW OR STOP CONVERGENCE?** This is the opposite of the above problem. People with a stake on a certain side will become mentally biased toward that side, resisting the rational implications of mounting evidence. This is of course not a new phenomenon in academia, and so it's hard to see why the problem would be worse. Except for issues closely connected to basic "ideologies" about which most everyone has an opinion, we can expect to find impartial jurors to bring people back to reality.

**WON'T DIFFERENT CLAIM WORDINGS, JUDGES, AND BASE ASSETS CONFUSE THE CONSENSUS?** Unless the performance of a base asset correlates with a claim, the price should be independent of base, and arbitrageurs can easily enforce this. If the prices on the same claim judged by different judges were persistently different, this would constitute consensus about judging bias, a situation that judges would want to avoid. If different claim wordings on an issue have very different prices, this represents consensus that there are really several different issues to be distinguished. For each distinguishable issue, traders seeking liquidity will probably congregate around one or a handful of base asset/wording/judge combinations, preventing things from getting too confusing.

**WON'T THE CONSENSUS REFLECT RISK PREFERENCES AS WELL AS BELIEFS?** Yes, the amount one should bet depends on one's beliefs, attitude toward risk, and the stake one already has in a question [Kad]. Risk-avoiders bet less than risk-takers, and bet less on the side that they already have a stake in. Price distortions from this should be minor, unless beliefs correlate significantly with risk attitudes and non-betting stakes, and if the stakes held approach each person's total wealth. One exception is that few people would bet for "Technology will soon make us all too rich to care about money", even if they believed it.

It might seem that questions with extremely lop-sided odds would also be a problem. Too few people might bet that "energy is conserved" (EC) if they very confidently expecting to win very little. But by splitting EC assets along other questions, people could jointly support EC, debate other questions, and get a higher average return.

Some people have worried that opinionated yet extremely risk averse people, unwilling to bet on anything, would be unfairly labeled "insincere" debaters. But it is hard for me to imagine that they could not afford to risk even \$10 a year so that we could develop a reputation score for them. If it is the risk of a low reputation score that scares them, perhaps they should not act so opinionated.

**WON'T BETTING CHALLENGES DISCOURAGE CREATIVITY?** If people were expected to bet on every idea that comes out of their mouth, they would be more reluctant to think up wild ideas, most of which are going to be bad. Hopefully we can maintain a distinction between saying "Here is an interesting idea to think about" and "This is the way it is, why won't you agree?", only expecting people to put up or shut up in the second case.

**WHAT'S THE POINT OF A "CONSENSUS" THAT PEOPLE DISAGREE WITH?** Regardless of the name used, people often want to pool their differing individual estimates on some issue into a composite estimate. This is most clearly needed in the "public choice" problem, where citizen estimates must be combined into government policy. But there is also a more general need for social institutions where experts combine their estimates on some subject into composite estimates that non-experts can use to making individual choices. Several such institutions may compete for attention, but the need remains.

Most work on consensus measures [Ge,Gr,Syn] focuses on various explicit functions to combine individual beliefs, and some simple variations of these [Man] are now used as academic consensus mechanisms. Compared to these, betting markets not only offer superior incentives [Ei] for people to bother to make their beliefs explicit and honest, but betting markets have the following unique claim to the word "consensus".

It is in the personal interest of an ideal decision theory agent to make all external actions as if they agreed with the market consensus [Kad], without any coercion. Agents should buy contingent assets up to the point where their marginal rates of substitution are the same, i.e. where they all agree on the relative value of getting one more dollar for sure vs. even more dollars in some contingency. An external observer, who can offer agents trades or choices but cannot tell how much each agent has already bet, cannot tell that the agents internally disagree.

Insurance-based proposals [Fa] are similar in spirit to the betting markets proposed here, as is the following proposal for dealing with the public choice problem [Mu]. If a government threatens to make a change, sells insurance on the change either way, and then makes the choice that is cheapest for them, they produce the most efficient "parteo optimal" result.

**ISN'T IT BETTER FOR PEOPLE TO ARGUE OUT THEIR OWN DISPUTES?** Yes, which is why we want incentives, such as with audit lotteries, to settle out of court and avoid judging. Idea futures is only intended to discourage insincere debaters.

Another way to avoid judging is to hold "argue lotteries" which are like audit lotteries except that judges are not invoked. The idea is to focus attention on a smaller number of markets where more is at stake. This should induce more discussion and examination of such questions,

perhaps resulting in more related questions being formed to reflect the argument structures proposed. Hopefully, opinions would naturally converge, and people would leave the market. Judges are really only there to discourage self-deception and strategic bargaining, so that the market odds eventually reflect the "obvious".

**WON'T THIS HAVE THE SAME PROBLEMS AS PATENTS?** No [Hir]. With patents we must decide who owns an idea, and so a centralized legal system must make a great many subtle decisions with insufficient evidence and expertise. We must examine history to decide who contributed how much to the idea. We must define some sharp legal boundaries that determine what it is to use the idea. With present patent law, we must also decide if an idea is true, if it is "original", if it is "obvious", if it is a "process", if it was revealed properly, etc. Bets are much more flexible; we need only decide if an idea is right, and we can each choose who is to judge that question. Government intervention and international agreement is not needed.

**WOULDN'T ANONYMOUS TRADING SCREW UP REPUTATION STATISTICS?** Perhaps people could make private trades to move prices out of line, and then make public trades on the other side to bring them back, so that those trades do better than average. This is somewhat like giving someone a wad of money by dropping it in the park and having them wander by an hour later to pick it up. If the park is crowded enough, someone else will have found it by then. In the market, anyone else could make money by stopping the price from moving out of line. The problem is more serious, however, if everyone accepts that only one trader has any information about a question, and so no one else wants to bet there. If possible, such markets should be excluded from reputation scores.

**IF THIS IS SO GREAT, WHY HASN'T IT HAPPENED ALREADY?** If it was in people's interest, wouldn't there be such markets by now? Well, if we always assumed this we might never make anything new, but it's an important question to ask. The fact that science bets have been legal only in Britain, and then only in the last three decades is only part of an explanation.

English bookmakers perceive little demand for science bets, and so take them mainly to induce popular articles mentioning the going odds on unusual subjects [ShG]. This publicity brings in new clients, who may then switch to the "real" betting on sports. Because of this, bookies prefer small bets on subjects "in good taste" that anyone can understand, like UFOs, Yetis, and Moon landings. They avoid subjects that seem too esoteric for the general public, like the recent "cold fusion" claims, and subjects that won't very clearly resolve themselves, as a judging industry has not yet evolved.

Bookmakers traditionally prefer to set prices and stick to them, rather than setting up markets, letting prices fluctuate, and playing market-maker. Because of this, they are usually unwilling to offer bets on claims where they do not know how to estimate the odds, and few bookies have advanced science educations. As a result, they mainly take safe bets, siding with the scientific establishment against "crazy" outside theories, which doesn't help the image problem betting has in many quarters.



English bookmakers do not seem to have seriously tried to sell imagine-conscious academics on science bets, through arguments like those in this paper. Nor, to my knowledge, has the possibility for betting markets as a funding mechanism been pointed out. Questions of interest to academics are now avoided and no visible influenceable consensus is formed; one cannot even subscribe to a publication listing the going prices on science questions. It should be possible to improve on this.

## Strategy

It's a lot easier to sketch a grand utopian vision than it is to figure out how to get there from here. An ideal development strategy would show how to grow incrementally, with each self-supporting step leading naturally to the next one. Most utopian visions fail because they, instead, require too many things to change all at once.

One advantage of idea futures is that, if not legally prohibited or socially shunned, it can co-exist with existing academic institutions and incrementally attract investors, patrons, and controversies. Papers would still be published and elite committees would still convene. Professors would gradually make more side bets, and begin to challenge each other to bets. Journalists would gradually use the market odds in news stories more, and funding agencies would gradually try larger amounts of subsidies. Idea futures could rise or fall on its own merits, as people studied how well its predictions compared to other consensus measures, and how the rate of progress in a field depends on the fraction of funding channeled through the markets.

Unfortunately, there also seem to be some obstacles to overcome before gradual growth is possible. Economies of scale in forming reputable judging organization or building secure computerized marketplaces may mean that certain levels of participation may be required before idea futures can "take off". But the major hurdle seems to be attitudes toward the very idea, attitudes reflected in the world-wide legal prohibitions. There are several possible strategies here.

One approach is more discussions, like those in this paper, of the need for alternative academic institutions, and of betting markets as a particular alternative. Perhaps idea futures must be disassociated from ordinary betting, as insurance and stock bets have been, though the metaphor of bets is very useful in explaining how it works.

Also helpful is further research on markets in conditional assets, such as recent attempts to show them superior to opinion polls at predicting elections [Fo].. Laboratory experiments [Sm] comparing betting markets to some mockup of existing peer review institutions would be very interesting, though not of course decisive.

A different approach, which I am also pursuing, would be to create an electronic mail-based reputation game, where people play for "bragging rights" instead of money. This would avoid legal problems and the discomfort academics have in dealing explicitly with money, and would allow many people from around the world to participate in a less-threatening partial test of

markets as an academic consensus mechanism. However, avoiding money makes the incentives suspect, and precludes many of the advantages, like insurance, that idea futures offer. If enough people played, the scores would mean something to observers, and so people would have an incentive to play and play well. But building a game up to this status would be hard, probably requiring some "big name" players to attract others.

If the basic idea became plausible enough to enough British (because that's where its legal) intellectuals in some field, it could be seriously tried. The initial field would preferably be one where bets are easier to settle, like number theory, though such subjects tend to be ones where existing institutions also work better, and so have less of a perceived need for change. A socially important question with minimal opportunities for conflict of interest would also be nice. Idea futures will have "made it" when it becomes known as a good place to find out the latest thinking on certain issues, reliably predicting what will later become consensus in other social contexts.

## **Advantages**

If its potential problems can be overcome, and a development path charted and followed, idea futures offers many advantages, most of which have already been mentioned.

There would be a clear incentive to be careful, honest, and expert, when making public statements. People could be rewarded for being right, rather than just for being liked by academic insiders. Those who invest wisely would accumulate capital and gain influence, which they could reinvest in discretionary research or in influencing future consensus.

Funding agencies would only need to pick important questions, not who would be good to research them. Diverse approaches could be tried to research a question, without arbitrary penalties for crossing disciplinary boundaries, ignoring fashion and insiders, integrating pre-existing knowledge, and using insights too small or inarticulate to make a publishable unit.

Anyone, not just Ph.D.s, could contribute directly to the world's corpus of knowledge. Easily published science odds and amateur betting might increase popular interest in science, subsidize professional efforts on questions of popular interest, and perhaps even increase the general savings rate. Clear market odds would ease science reporting.

A visible scientific consensus would be available to guide public policy, a consensus which would be self-consistent across a wide range of issues and harder for media campaigns to distort. Compared to competing consensus mechanisms, idea futures should be relatively simple, cheap, decentralized, egalitarian, responsive to new information, and at least as informative. This consensus should correct for many current biases, such as overconfidence.

The mere threat of betting challenges could improve incentives in discussions and debates. If the market consensus carried social weight, it could serve as a coordination point for thousand of independent conversations. A rejected visionary would have a new way to get publicity for his ideas, and a reward for being right against the establishment. True cranks would subsidize

leveler heads, and as debates became settled, they would leave a trail of agreed-upon statements which could be used to counter bogus claims made by those ignorant of solid expert consensus.

Businesses could make insurance hedges against technological risk, as in the cold fusion case. While such insurance may be legal now, the introduction of speculators would increase market thickness to a point where it might be practical.

Reputation scores offer an new way to evaluate people's ability to separate the wheat from the chaff in ideas and arguments, and these scores should depend less on whether one has curried favor from the right people.

Idea futures is well-grounded in our best theories of decision and incentives. Once legal and accepted, idea futures could grow incrementally, and perhaps dramatically increase our rate of scientific progress per funding spent.

## **Conclusion**

Markets in contingent assets, more commonly known as "bets", offer a needed alternative to existing academic institutions. Betting markets cannot solve all current problems, or replace all current institutions. But if this paper has been successful, the potential of such markets should be clear, and most of the obvious problems with such markets should have been addressed in enough detail that we can say the idea still seems plausible on a closer examination. If so, more serious intellectual discussion is justified, and perhaps some small-scale experiments. We could do much worse than having intellectual institutions as open, flexible, diverse, and egalitarian as the stock market, with incentives as well-grounded and with estimates on important issues as unbiased and predictive.

## **Acknowledgements**

These ideas germinated in the fertile ground of discussions with friends interested in similar problems, most of whom are associated in one way or another with the company Xanadu. Eric Drexler, Mark Miller, and Phil Salin have been particularly influential. And my wife Peggy Jackson has influenced me in more ways than I know.

Perhaps a hundred people, more than I can list here, have provided useful comments and criticisms on all aspects of the idea. Useful comments on this particular paper came additionally from Amara Graps, Martin Haerberli, and Chris Hibbert.

## **References**

[APA]

(1987) Diagnostic and Statistical Manual of Mental Disorders Third Edition, Revised, American Psychiatric Association, Washington D.C.

- [Be]  
Bernal, J. (1989) "After Twenty-Five Years", *Science and Public Policy*, 16:3, June, pp.143-151.
- [Br]  
Brunner, J. (1975) *The Shockwave Rider*, Harper & Row, NY.
- [Ci]  
Cialdini, R. (1988) *Influence, Science and Practice*, Scott, Foresman and Co., Boston.
- [De]  
Debus, A. (1970) *Science and Education in the Seventeenth Century*, MacDonal, London.
- [Dr]  
Drexler, K.E. (1986) *Engines of Creation*, Doubleday, New York.
- [Do]  
Dowie, J. (1982) "Gambling in Education and Education in Gambling", *Proc. Fifth Intl. Conf. on Gambling and Risk Taking*, ed. William Eadington, pp.23-39.
- [Ec]  
Economides, N., Siow, A. (1988) "The Division of Markets is Limited by the Extent of Liquidity (Spatial Competition with Externalities)", *American Economic Review*, March, pp.108-121.
- [Ei]  
Eisenberg, E., Gale, D. (1959) "Consensus of Subjective Probabilities. The Pari-mutuel Method" *Ann. Math. Statist.* 30. pp.165-168.
- [Fa]  
Fairley, W., Meyer, M., Chernick, P. (1984) "Insurance Market Assessment of Technological Risk", *Proc. Soc. for Risk Analysis Internl. Workshop on Uncertainty in Risk Assessment, Risk Management, and Decision Making*, Sept. 30, pp.89-102.
- [Fe]  
Feyerabend, P., (1975) "How to Defend Society Against Science", *Radical Philosophy*, 2, Summer, pp.4-8.
- [Fo]  
Forsythe, R., Nelson, F., Neumann, G., Wright, J. (1990) "The Explanation and Prediction of Presidential Elections: A Market Alternative to Polls" *Economics Working Paper 90-11*, April 12. Univ. of Iowa, Iowa City.
- [Gal]  
Galileo, *Dialogue Concerning the Two Chief World Systems*, transl. by Stillman Drake, 2nd ed., Berkeley, Univ. of California Press, 1967.
- [Gar]  
Garwin, R. (1989) "Consensus on Cold fusion still elusive" *Nature*, 338, April 20.
- [Ge]  
Genest, C., Zidek, J. (1986) "Combining Probability Distributions: A Critique and Annotated Bibliography", *Statistical Science* 1:1, pp.114- 148.
- [Gh]  
Ghiselin, Michael T. (1989) *Intellectual Compromise: The Bottom Line*, Paragon House

Publ., NY.

- [Gr]  
Grofman, B., Owen, G. eds. (1986) Information Pooling and Group Decision Making: Proc of the Second Univ. Cal. Irvine Conf. on Political Economy, JAI Press, Inc. London.
- [Hak]  
Hakansson, N., Beja, A., Kale, J. (1985) "On the Feasibility of Automated Market Making by a Programmed Specialist", J. Finance, March.
- [Hal]  
Hall, S. (1989) "Professor Thorne's Time Machine" California, October, pp.68-77.
- [Han88]  
Hanson, R. (1988) "Toward Hypertext Publishing, Issues and Choices in Database Design", ACM SIGIR Forum, 22:1,2 Winter 1988.
- [Han90]  
Hanson, R. (1990) "Market-Based Foresight - A Proposal", Foresight Update, 10, publ. of Foresight Institute, P.O. Box 61058, Palo Alto, CA 94306.
- [He]  
Henrion, M., Fischhoff, B. (1986) "Assessing Uncertainty in Physical Constants", Am. J. Phys. 54:9, September, pp.791-798.
- [Hir]  
Hirshleifer, J. (1971) "The Private and Social Value of Information and the Reward to Inventive Activity", American Economics Review, 61:4, Sept., pp. 561-74.
- [Ja]  
Jaffray, J. (1989) "Coherent Bets Under Partially Resolving Uncertainty and Belief Functions", Theory and Decision, 26, 99-105.
- [Kad]  
Kadane, J., Winkler, R. (1988) "Separating Probability Elicitation from Utilities" J. American Stat. Assoc., June, 83:402, Theory and Methods, pp. 357-363.
- [Kah]  
Kahneman, D., Tversky, A., eds., (1982) Judgment under uncertainty: Heuristics and biases, Cambridge Univ. Press, NY.
- [Kan]  
Kantrowitz, A. (1977) "The Science Court Experiment: Criticisms and Responses", Bulletin of the Atomic Scientists, April, pp.44-50.
- [Ke]  
Kelley, J. (1988) "Compulsive Gambling in Britain", J. Gambling Behavior, 4:4 Winter, p.291
- [Ki]  
Kitcher P. (1990) "The Division of Cognitive Labor", J. Philosophy, 87:1, January, pp.5-22.
- [La]  
Laffont, J.J. (1989) The Economics of Uncertainty and Information, MIT Press.
- [Led]  
Lederberg, J. (1989) "Does Scientific Progress Come from Projects or People?" Current

Contents, 48, November 27, pp.4-12.

[Lew]

Lewis, H.W., (1989) "Fusion isn't just around the corner" San Jose Mercury News, April 7, Opinion page.

[Li]

Linstone, L., Turoff, M., ed. (1975) The Delphi Method, Addison- Wesley, London.

[Mal]

Malkiel, B. (1989) "Is the Stock Market Efficient?" Science, 243, March 10, pp.1313-1318.

[Man]

Mann, C. (1990) "Meta-Analysis in the Breech" Science 249, August 3, pp.476-480.

[Mar]

Marshall, E. (1990) "Science Beyond the Pale", Science 249, July 6, pp.14-16.

[May]

Mayer, M. (1988) Markets, Norton & Co., NY.

[Mu]

Mueller, D. (1979) Public Choice, Cambridge Univ. Press, Cambridge.

[My]

Myers, D. (1983) Social Psychology, 2nd ed., McGraw-Hill.

[NS]

(1990) "Feedback Column", New Scientist, July 14. See also 2/10, 7/28.

[Pe]

Pearl, J. (1988) Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference, Morgan Kaufmann Publ., Inc., San Mateo, CA.

[Pol]

Polinsky, M, (1983) An Introduction to Law and Economics, LittleBrown and Co., Boston.

[Poo]

Pool, R. (1989) "Bulls Outpace Bears - For Now!" Science, 244, April 28, p.421.

[Pop]

Popper, K. (1985) Popper Selections, ed D. Miller, Princeton Univ. Press.

[Pow]

Power, W. (1989) "Hottest Commodity In Wall Street Pits? Georgetown Hoyas", Wall Street Journal, March 24, p.A1.

[Pr]

Price, D.J. de Solla (1963) Little Science, Big Science, and Beyond, Columbia Univ. Press, NY.

[Re]

Redner, H. (1987) "Pathologies of Science", Social Epistemology, 1:3, pp.215-247.

[Ro]

Rose, I.N. (1986) Gambling and the Law, Gambling Times Incorporated, Hollywood.

[Se]

Seidenfeld, T. (1990) "Two Perspectives on Consensus for (Bayesian) Inference and Decisions" Knowledge Representation and Defeasible Reasoning, H. Kyburg, et. al. eds.

<http://hanson.gmu.edu/gamble.html>

pp267-286.

[ShW]

Sharpe, W. (1985) *Investments*, 3rd Ed., Prentice Hall, NJ.

[ShG]

Sharpe, G. (1990) phone conversations, William Hill Org. Ltd. 19 VAentine Pl. London SE1 8QH, July.

[Sm]

Smith, V. (1986) "Experimental Methods in the Political Economy of Exchange", *Science*, 234, October 10, pp.167-173.

[St]

Stoll, H. (1989) "Inferring the Components of the Bid-Ask Spread: Theory and Empirical Tests", *J. Finance*, 44:1, March. pp.115-134.

[Syk]

Sykes, C. (1988) *ProfScam*, St. Martin's Press, NY.

[Syn]

(1985) *Synthese*, "Consensus" issue, 62:1, Jan

[Th]

Thayer, H. (1980) "Peirce on Truth", *Two Centuries of Philosophy in America*, ed. P. Caws. Towman & Littlefield, Totowa NJ.

[Ts]

Tsao, J. (1989) "Consumer preferences and funding priorities in scientific research", *Science and Public Policy*, 16:5, October, pp.294- 298.

[Tu]

Tullock, G. (1966) *The Organization of Inquiry*, Duke Univ. Press, London.

[Ve]

Verne, J. (1872) *Around the World in Eighty Days*

[Wh]

White, F. (1946) *Famous Utopias of the Renaissance*, Hendricks House Inc, Putney Vermont.

[WSJ]

(1989) "Fusion Fuss Is Turning Scientists Into Gamblers" *Wall Street Journal*, April 18.

[Ze]

Zeckhauser, R., Viscusi, W. (1990) "Risk Within Reason", *Science*, 248, May 4, pp.559-564

BackLinks (from AltaVista)