IDEA FUTURES
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 415-651-7483

INTRODUCTION

Are you fascinated by some basic questions about science, technology, and our future? Questions like: Is cryonics technically feasible? When will nano-assemblers be feasible and how quickly will resulting changes come? Does a larger population help or hinder the world environment and economy? Will uploading be possible, and if so when? When can I live in space? Where will I be able to live free from tyranny? When will A.I.s be bucking for my job? Is there intelligent life beyond earth? If you are like most *Extropy* readers, such question matter to you.

Now how do we, as a society, go about answering such questions? People who have an appropriate background, and who are interested enough in a particular question, can research that subject in depth themselves, and come to a considered opinion. And people who happen to know, respect, and trust such a person can simply take those opinions as their own, avoiding all the hard work. But what is everyone else to do, people whose actions often implicitly depend on such questions?

In practice, people usually defer to larger social institutions on most questions, institutions which combine and evaluate contributions from many specialists, and which offer apparent institutional consensus estimates on many different questions. These consensuses may be uncertain and temporary, and individuals may prefer to combine the results of several institutions, but basic need for such estimates remains.

For example, popular media chose what they consider to be ranges of reasonable and noteworthy opinions on noteworthy issues for presentation. Peer-reviewed academic journals and societies offer more detailed, though less accessible, consensuses about which opinions are reasonable and in favor. Government agencies often try to form and act on such consensuses, as the U.S. E.P.A. does for pesticide health risks. And there are many other such consensus institutions, such as opinion polls.

How well do these institutions work? How many of us are confident that, when a technical controversy arises, a widely visible consensus will quickly emerge representing society's honest best estimate on the issue, reflecting the relevant insights of the best relevant experts? Or that those with foresight will eventually be rewarded for advocating positions which later become accepted?

People who have little contact with an existing social institution, or who have a position of power within it, may feel things are basically okay. But those "in the trenches" typically voice more skepticism. Your opinion on the trustworthiness of newspapers probably changed for the worse the last time you read about an event in which you were personally involved. Since

subjects like cryonics and nanotechnology have often been unfairly treated by most current institutions, I expect at least a few *Extropy* readers to be dissatisfied with such institutions.

Skeptics about current institutions are not typically focused on methods, often the center of philosophical discussions, but on incentives. Skeptics see too many rewards for bias, and too few for honesty and care, and so distrust official statements. People often promote beliefs which serve their self-interest, and try to appear more confident, original, and knowledgeable than they are. People don't correct for standard human biases [Kah], such as wishful thinking, overconfidence, and belief fixation. They massage evidence, suppress criticism, and just plain lie. Fashion, eloquence, and politics often dominate expert consensus. Rewards often go to those whose ideas are popular now, rather than those who are later proved correct [Tul]. Paid advocates distort the consensus we perceive by using raw media exposure, bribes, and by exploiting human biases [Cia]. An honest consensus of relevant experts is often lost from public view.

Many existing social institutions, such as investigative reporting, due process, public debates, and peer review, claim to address these problems. But there is room for improvement. To improve the way ideas evolve, many people try to reform existing social institutions, and a few try to invent new ones [Li, Kan,Han88]. In this paper I suggest a new social institution, called "idea futures", which can create a visible expert consensus with clear incentives for honest contribution.

CONCEPT

Idea Futures is intended to aid the evolution of a wide range of ideas, from public policy to the nature of the universe, and in particular should be able to help us predict and understand our future. The basic concept is to combine two phenomena, convergence and markets, and so make "a futures market in ideas". Let me explain.

Disagreement is rarely as fundamental as it seems. In the long run, beliefs often converge. For example, in science the steady accumulation of evidence eventually settles most debates. We take the advice of experts, indicating that we think we would come to believe what they believe, if only we were to study what they have studied. Randomly selected juries usually reach a unanimous verdict, even more often than seems rational. Theory [Se] and experiment [Li] indicate that the people's beliefs should and do converge. In sum, we generally trust in a convergence of human judgment.

Figure 1 Some Controversial Claims

- By 2030, the greenhouse effect and other causes will have raised the average world sea levels by 1 meter.
- Cold fusion of deuterium in palladium can produce over 10 watts/cc. net power at STP.
- By the time we have surveyed our galaxy in the infrared to the 25th magnitude, there will not be any evidence of another technological civilization in our galaxy.
- By the time world GNP is four times the 1990 level, 1000 people will have physically lived in space over 90% of the time for the previous seven years.
- By the year 2000, over 20,000 people in the U.S. will commute to work in Vertical Take-Off and Landing (VTOL) aircraft they park in their garage and drive down the street.
- The rest mass of the electron neutrino is greater than .01eV in ordinary space.
- If labor saving device X were widely used in industry Y, industry employment five years later would be less.
- Death-bed confessions or other evidence will eventually show that person X was murdered.

If people wait long enough for evidence to surface and then take enough effort to study and debate a specific enough claim, they often come to agree. When the people are reasonable, knowledgeable, and detached enough, and when they avoid subjects like religion, they usually agree. When such a group is diverse and independent enough, we believe we would probably also agree.

Markets are a way to create a consensus about the value of an ownable item, i.e., the "price". Futures markets are a way to create an immediate consensus about future consensus. For example, a market in corn creates a price in corn, so that most buyers pay about that price. A futures market in corn creates a futures price, which is an immediate estimate of what the actual price of corn will be in, say, nine months. Traders have clear incentives to make honest contributions to the consensus; you "put your money where your mouth is". A trader who believes the future price will be higher than the market indicates buys, and in so doing raises the consensus price. Those who are right make money from those who are wrong.

Of course markets have limitations. Ideally, items to value should be of wide interest, exclusively ownable, cheaply transferable, and have many identical copies. How can we apply this to ideas? By creating coupons whose value depends on whether an idea is validated. For example, a coupon which says "Exchangeable for $1 should a person land on Mars by 2020" is a direct tie between an idea, people on Mars, and money, a well-known unit of incentive. Such coupons can be thought of metaphorically as futures, and more literally as bets, a metaphor often used to describe both investments and science.

If convergence creates a future consensus in ideas, and if futures markets can create an immediate and honest consensus about a future consensus, then futures markets might be able to create an immediate and honest consensus in ideas. If the market price for a "$1 if person on Mars" coupon were 23¢, then that would typically represent a consensus that there was about a 23% chance of this happening. Anyone could express their opinion on the subject by trading coupons, or could just read the "market odds" to see the best present estimate. This market consensus would compare favorably to other methods of forming perceived consensus, such as by advertising, opinion polls, or elite committees. An idea futures consensus could be simultaneously open, egalitarian, universal, expert, honest, self-consistent, and cheap.

A mature idea futures market could offer coupons on many claims about the future of technology and society. The consensus prices would give describe a consistent set of probabilities for various possible future events, and conditional probabilities for some events given others. Investors there should be as diverse as investors are elsewhere, with a mix of short-term and long-term focuses, large investment houses and daring do-it-yourself individuals, each contributing their specialized knowledge about an issue or the connection between two issues to the total consensus.

Like cryonics, idea futures is another way to take advantage now of the fact that the future should be rich with power and knowledge. We create good incentives now by letting the future settle our bets.

To make the whole idea more vivid let us consider an simple (fictional) scenario.

SCENARIO

Pat Thgisni was not a model student. A knack for making experiments work is probably what got him into graduate school -- it certainly wasn't his grades. Worse, he

was unkempt and had a disturbing habit of bending people's ears with one harebrained idea after another. Definitely not one of the rising stars of the University of Toledo Physics Department.

In his second year, 1992, Pat hit upon his best idea yet, "superscattering". If a neutrino could scatter off all the nuclei in a crystal at once, the interaction could be a billion billion times more powerful, perhaps allowing neutrino telescopes [Web]. Pat showed his calculations to Prof. Ezra Puccuts, a local and renowned neutrino expert, though rusty on scattering. Prof. Puccuts explained to Pat that a similar idea had occurred to him, but he had found it conflicted with an accepted formula. Such a negative conclusion wasn't worth a publication.

Pat persisted, however, bringing out his pages of calculations. After ten minutes of going through the first page, and finding three glaring, though irrelevant, math errors, Prof. Puccuts lost his patience. "I do not have the time to correct your math for you", he declared, and shut the door. Over the next few months Pat redid his calculations several times, but Prof. Puccuts was not interested and other professors referred Pat back to him. Pat submitted his work for publication anyway, and then waited; he did not have the $100K he figured it would cost to do an experiment.

That Christmas in L.A., Pat told the story to his family. His brother Al, a sports fanatic, suggested that Pat dare them to make a bet. Before Pat could object, Al described how idea futures were revolutionizing the oil industry, and were a new way for the little guy to contribute to the world of ideas. After a few more drinks, Pat saw the light.

Pat wrote up a precise statement of his claim, and then stopped by the idea futures mart in Las Vegas{footnote: Actually, science bets are illegal in Nevada} on his way back to Toledo. He paid $100 to have a reputable judging group decide if it was precise enough for them to judge in 2013 (which it was), $20 to the Bank of Vegas so they would issue coupons on it, and another $20 to have a computer market set up. Finally, he funded an automated market maker with $200 in seed capital, and set the initial market odds at 30%. Back at the university, Pat set his computer up to track the market, and then spread the word, causing an epidemic of giggles. One of Prof. Puccuts smirking students agreed to put up $20 against him, and a half dozen other students joined in, mostly at $2 each.

It worked like this. While Pat bought coupons which said "$1 if superscattering" from the market maker for around 30¢, the other students bought "$1 if no superscattering" from the market maker for around 70¢. Whoever was right in the end would make money on the deal, receiving the $1 plus 7% interest per year, compounded. Every time Pat bought a coupon the market odds went up, and every time the other students bought the odds went down. The market maker got these coupons from the bank, who would sell the pair of coupons "$1 if superscattering" and "$1 if no superscattering" to anyone for $1 a pair. The bank made money on transaction charges, and risked nothing because exactly one of the pair will be worth $1 (plus interest) in the end. The bank also bought such pairs back for $1, allowing people to sell coupons back to the market maker. Many students took advantage of this feature in the next few weeks, as some professors made it clear they were not amused.

In 1995, Pat earned an early Ph.D. Like most students, he could not find an academic position and went to work in industry. A year later he finally published his superscattering article in a vanity journal.

Over the years, Pat had tried to follow the literature to see if anyone else had the same idea, but without success. Meanwhile bets slowly trickled in, with the odds hovering around 15%. In 1997 the market told Pat of another bettor, in Peru, who made a number of publicly declared purchases, or "public bets", of superscattering coupons. The Peruvian had also published on the idea, but in an obscure Peruvian journal. Someone else created a market on whether there would be a compact neutrino telescope, which became popular with amateur astronomy clubs. Certain traders even specialized in keeping these two markets in rough correspondence. By 1998, the total value of all coupons out on superscattering, the "market capital", reached $8K.

That year an aide to Nevada Senator Sue Toshgib, member of the Senate Committee on High Tech, noticed the high odds for neutrino telescopes, and traced it to Pat's claim. Sue saw an opportunity to push Nevada's fledgling idea futures industry, and made an issue of the fact that the markets had apparently discovered a number of potential new technologies. For example, she said, if there was a 15% chance of superscattering, why wasn't this possibility being pursued in the labs?

Wheels turned, but Prof. Puccuts' technical explanations did not satisfy Senator Toshgib. Federal funding agencies wanted to avoid a confrontation, but also wanted to protect their turf from every senator's whim. So they prodded the administration of the University of Toledo, to quietly make a few bets with university overhead funds. This infusion of capital overwhelmed what Pat and a few others could scrape up. The odds fell to 2%, and the issue was dropped.

But the $100K now in the game raised the interest of a few experienced speculators looking for an angle. They hired a few grad students to try the superscattering calculations, and the odds crept up to 6% over the next year. About the time the university realized there might be no limit to the capital required to keep the odds down, Prof. Puccuts published a paper showing why superscattering was impossible. The hired grad students were intimidated, and the odds fell to 1%. Six months later a student from across the hall over heard Prof. Puccuts mention that an equation in his paper was an ad-hoc approximation. So he re-tried the calculation, and got a rather different result. He told his old professor, Prof. Yikkul, and they jointly wrote a paper. As rumors

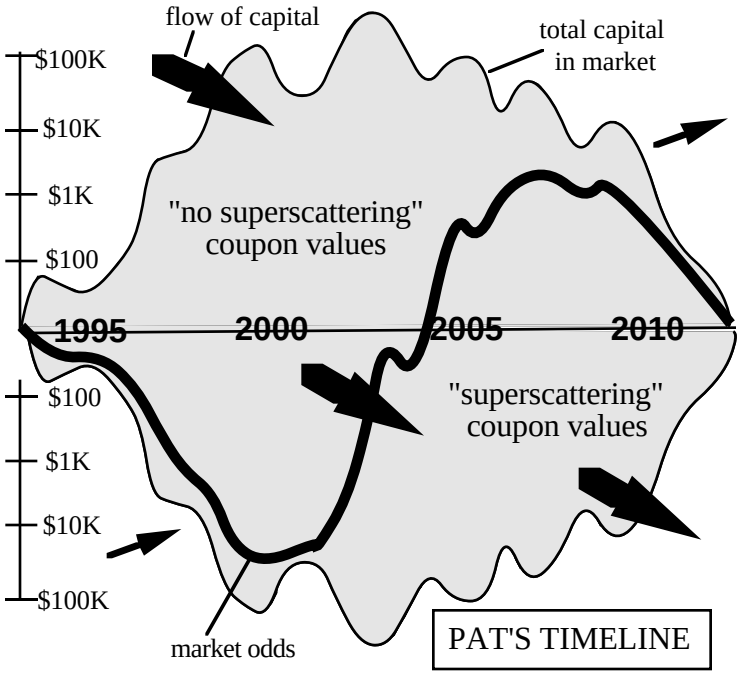

PAT'S TIMELINE

spread, the market odds shot to 20%.

In 2001 the first experiment was started. The odds fluctuated under the influence of some false rumors, and some experimenter's friends made extra income by trading before the experimental results were revealed. In 2005 the market odds were at 70%, and by 2008 the issue seemed pretty much settled, with the market odds at 98%. Pat, who had doggedly stuck with superscattering, sold his coupons to reap a total profit of $700K on his $20K investment (which was all his spouse would let him risk). This profit came from selling "$1 if superscattering" coupons for 98¢, after buying them for as little as 1¢, and from the 7% interest the Bank of Vegas had agreed to pay on the money it held. Just before the coupons were to be judged in 2013, the last traders sold their coupons to avoid paying a judging fee. The market was closed down and the judges were never needed.

Pat was still not offered an academic position, as Prof. Yikkul became the celebrated discoverer of superscattering. So Pat started a new market, to be judged by a detailed historical study in thirty years, on who was the first discoverer of superscattering. Certain universities vowed to let professors defend their own ideas. Prof. Puccuts, who never bet any of his own money, still has tenure.

SCOPE

Mechanisms like idea futures have been used for a long time to create consensus about corn prices, stock dividends, life expectancy, marine accident risks, horse races, and football games. So clearly it can work for some topics. But the vision offered here is to make much wider use of such mechanisms. Some areas of science and technology seem similar enough to horse races to suggest betting will work there also, but what about everything else?

There are a number of parameters which indicate when a claim will be more difficult to handle, including the time and expense required to resolve a claim, the probability it will never be resolved, the strength of emotions on the issue, the lopsidedness of the odds, and the scarcity of interested traders. The procedures described below are intended to allow idea futures to handle as many claims as possible. I hope to make the case for wide applicability plausible enough to inspire interest and experiments.

PROCEDURES

In idea futures markets people would exchange coupons like those attached to the beginning of this paper (Anti-gambling laws require that these be void, unfortunately.) Each coupon is issued by a bank, and specifies a judging organization who will decide the issue "beyond a reasonable doubt". Coupons have the basic form "$X if A", where A is a claim and $X is a "conditional value". A claim includes a sentence, such as those in Figure 1, and any clarifications on word meanings. The claim part of a coupon also specifies who will judge the claim, a judging date, and any declarations about the decision criteria or process to be used. There can be many coupons on the same claim, each to be judged by different judges on different judging dates.

Coupons also specify a total judging fee and a maximum percentage judging fee. The judging fee is obtained by reducing the face value of each coupon in the market on the judging date by whatever percentage is necessary given the total market capital. If this would violate the maximum percentage judging fee, then the banks must try to create enough market capital by gambling the

existing capital in an "audit lottery". If the coupons win the lottery, enough capital is created to support judging, and coupon face values are increased. If not, coupons are worthless and judging is not needed.

The judging fee creates an incentive to "settle out of court" by selling before the judging date, as happened in Pat's superscattering market. Audit lotteries preserve incentives for honest evaluation even when an issue is of only limited interest and very expensive to judge [Pol], such as whether your daughter would make a good doctor, if only someone would pay her way through medical school. Pat specified that an audit lottery be used, if necessary, to fund the historical study on who discovered superscattering.

Judge's verdicts should be "beyond a reasonable doubt" and are expressed as percentages to be paid off to each side. Judges have additional incentives to be careful if they agree to spend some fraction of their judging fee to keep the market price of an "appeals" coupons near that of their verdict. Appeals coupons are on the same claim, but judged much later by independent judges with a much larger budget.

If there is not yet enough evidence or funding to decide a question, judges may be allowed to postpone judging to a new date with a new, perhaps larger judging fee. They could in the process offer some clarification of the question, and would use up some portion of the max percentage judging fee to pay for their trouble. Markets can also be set up so that if no decision can be agreed on, judges can declare "all bets are off".

The conditional value part of a coupon, the $X, specifies a standard investment instrument, such as a stock, bond, or mutual fund, and gives specifics like amount, date of purchase, interest rate, etc. There should be a liquid market in such instruments, so that it is always clear what the equivalent cash value is. In the superscattering example, bonds issued by the Bank of Vegas paying 7% interest were used. There can be coupons with different investment instruments for each claim and judge combination described above. By building on standard investments, an idea futures investor can expect a better rate of return than he could get with any standard investment alone.

Banks are long-lived financial institutions trusted to properly report judging fees. A bank's main function is to split and join

Figure 2. Idea Futures Home Version 1) Choose a claim like "I will win this hand of stud poker". 2) Get a pencil, and some chips. Let red chips be "$10 if claim", blue chips be "$10 if not claim", and something else be money. 3) Give each person $100. 4) At any time players may buy or sell pairs of red and blue chips for $10. 5) Place the pencil across the bar below between 5 and 6 on the CLAIM side. This means the market maker will sell one red chip for $6 or buy it for $5. 6) Whoever yells first, such as "buy red for 6", can trade one chip at the offered price. The pencil then immediately moves in that direction, such as to 6-7. Repeat till market settles. 7) Reveal new information, like the next card, and repeat step 6. 8) If the claim was right in the end, reds are worth $10, blue $0. If wrong, blue is $10, red $0. Have fun!

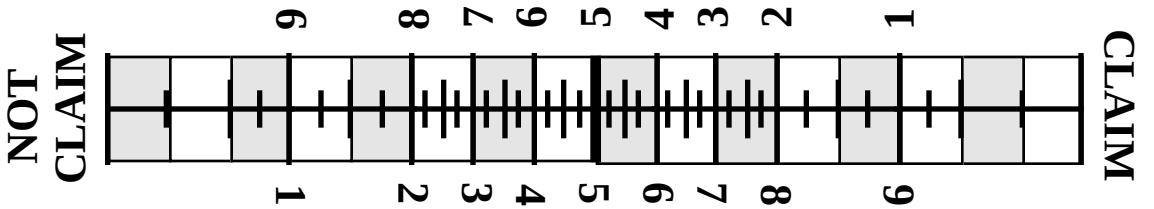

coupons. For example, the claim "True" can be split into "A" and "not A". Imagine giving a bank one share of General Motors (GM) stock. The bank would see this as a "1 GM share stock if True" coupon and exchange it for the coupon pair "1 GM share if A" and "1 GM share if not A". The bank would hold on to the pair and if A wins, give the 1 GM share to any holder of a "1 GM share if A" coupon. While the bank held the 1 GM share betting stakes, the wealth it represents would, we hope, be put to productive use by GM. An "A" coupon can be further split into "A and B" and "A and not B". Using certain combinations of such coupons, one can bet on the conditional probability of "B given A" and be insensitive to the verdict on A.

Each type of coupon must have at least one public market for trading coupons. Preferably, such markets will be continuous bid/ask markets allowing anyone to post or take offers via computer. A single computer could implement thousands of low-activity markets.

To increase liquidity and reduce price fluctuations and spreads, anyone can fund automated market makers [Bl], computer programs always available for trading. A simple market maker algorithm exists which can function indefinitely and not be cheated by clever combinations of traders (see Math Appendix). The degree of price smoothing it provides depends on the amount of sacrificial cash it starts with. This provides a way to subsidize a market, as does offering to pay part of the judging fee. Pat funded such a market maker to promote trading in his market.

If the odds on different claims are inconsistent, i.e., do not adhere to the standard axioms of probability, then arbitrageurs can make money by buying or selling "Dutch books" [De]. This profit comes at no risk if the final verdicts can be trusted to be consistent. Therefore arbitrage activity should keep the total social consensus roughly self-consistent.

Other market innovations, such as options, baskets, and hedges, allow investors to specialize in details they think they know about and ignore other issues. Options allow bets on price volatility, independent of the way the price moves. Baskets allow one to ignore differences; one can buy a basket of all types of coupons on a certain claim, and ignore differences in judges, investment instruments, etc. Hedges allow bets on price differences, such as when investors kept the odds on superscattering and compact neutrino telescopes in rough correspondence. For example, one could correct for the human bias of overconfidence by betting that on average the odds are not as extreme as the market odds.

While Pat had to risk a substantial portion of his wealth on one question, a more typical scenario would include larger private research labs whose salaried employees direct investment in many questions.

Idea futures markets could be integrated with one or more publishing media or "registries". People could make "public bets", where they buy a coupon for a claim, write a statement of support, and commit to having registries reveal both of these at a pre-specified date. Track records could be compiled from such information and used as reputation scores. People with high scores could become investment advisors, making a public bet with each piece of advice. A advanced publishing medium [Han88] would allow anyone to post evidence and arguments and link them to the disputed claim.

ADVANTAGES

Idea futures offers many possible advantages. A visible consensus would im-

mediately form on a wide range of hotly debated issues. This consensus would be relatively universal, expert, honest, self-consistent, and cheap. Such a consistent consensus might allow society as a whole to approach the level of rational consistency that is now only expected of individuals.

The market consensus could carry social weight, serving as a coordination point for thousands of independent conversations. In each discussion, the market odds on an issue could be assumed as the default unless specific arguments were presented to the contrary. Dissenters could be given the time-honored challenge to "put up or shut up". In the same tradition, those willing to put themselves on the line would be given due respect and attention. I have observed that the challenge of a bet makes people noticeably more cautious about what exactly they are claiming.

As debates become settled, they would leave a trail of agreed upon statements. These could be used to counter bogus statements, often made by those ignorant of solid expert consensus. Visionaries like Pat would have a new way to try to convince others of a revolutionary claim; they could throw all available capital into bets. If this were enough to change the market odds, they could point to these odds in arguments. If not, they would at least expect to make a healthy profit, and gain social credit for being serious. True cranks would end up subsidizing leveler heads.

The social weight of consensus could help damp many presently distorting biases. It would be harder for popular media to create consensus by sheer repetition of a claim; they would have to convince those willing to bet. A sincere public relations campaign could make a public bet, but an insincere one would know they were throwing money away. And an insincere attempt to throw enough money away to change the market odds runs the risk of the word getting out and the market ending right back where it started. Finally, hedge bettors can correct for standard biases in individual judgment.

Individuals would have clear monetary incentives to be honest and careful in contributing to the market consensus. If the odds you believe are different enough from the market odds, you believe you will on average make money, even more than with a standard investment like a stock index fund. And compared to stocks, idea future bets are precise and modular. In stock bets one must usually bet on a combination of ideas, such as the company's product, marketing strategy, production techniques, etc. In idea futures you can bet on exactly the issue you think you know something about.

It might be fun! Imagine a page in the newspaper like the stock page, showing this week's odds on controversial issues. Imagine coming home from an enlightening discussion to change your investments. Imagine reading something you disagreed with, and stopping for a minute to make a bet against it. The knowledge you created while reading would be directly useful to society and yourself, instead of thrown away as is usual now.

Non-scientists could have a direct, even if small, influence and personal stake in science to heighten their interest, like the amateur astronomy clubs in the story of Pat. Amateur trading would induce scientific research by traders seeking an edge, subsidize professionals who can better predict, and might even fund research by judges. Savings would be encouraged and research would be directed more at issues of general interest. Capital and hence intellectual effort would flow to markets where there is broad interest, strong disagreement, and relevant data obtainable for a modest effort or a short wait.

Idea futures markets create information, combining what individuals know. A market with more capital will probably have better information, as people will see there is more to win by figuring out the answer. By subsidizing a market you can pay to create information, though you won't get exclusive access. This might be a better way for government to fund scientific research, instead of the usual grant-giving approach [Han90]. In fact, governments might use odds from subsidized markets as factual input for government decisions. We could all have our say about whether projected usage would justify a new mass transit system, or whether a death was suspicious enough to justify an autopsy. Schools might even admit students based on the market odds of candidates getting a high G.P.A. if they attend.

Idea futures can also provide insurance. A risky business venture based on a new technology might bet against that technology to reduce total risk. Idea futures can be a foundation for reputation systems, providing another way to encourage experts to give honest advice, and allow other experts to disagree. Idea futures offers all these benefits without requiring any coercion or taxation. Unlike patents, it requires no international enforcement or litigation about the origins of an idea.

CRITICISMS

By now you probably have in mind at least one objection to idea futures, and will not be entirely comfortable with it until this objection is addressed. Longer papers on this subject [Han90] consist largely of detailed responses to such objections. Space limitations preclude such detail here, so figure 3 just gives a list of some issues addressed in those papers. The editor of *Extropy* willing, I could respond to the specific concerns of *Extropy* readers in future issues.

RELATED WORK

In Bayesian decision theory, an agent's degree of belief in A is often defined to be the price they would be willing to pay for a "$1 if A" coupon [DeF]. Idea futures just applies this definition to a society as a whole to find our consensus degree of belief. In the presence of a market, agents appear to agree [Kad].

As was mentioned before, markets similar to idea futures in commodities, finance, insurance, and sports betting. Science and technology bets are frequently made between individuals [Hal,Ti,WSJ], as they have been through history [Deb]. The idea of betting on a wide range of legislative and technological issues is raised in a recent science fiction novel [Bru,Ve], and scattered proposals [Fa,Ho,Lea,Ze] have been made to formalize bets in science.

Business schools widely use such markets to teach M.B.A. students about markets [Fo]. In economic theory, the coupons I have been describing are called "contingent assets", and are often used as a foundation in analyzing financial investments [ShW] and the effect of uncertainty [La]. Ideally, there would be a "complete market", with assets contingent on every possible state of the world. In reality markets are not complete, and various sorts of "market failure" are traced to this fact.

Incompleteness is usually [Hir] explained as due to judging difficulties, finite transaction costs, and market thinness. In fact, these authors are often unaware that such markets are almost universally prohibited by anti-gambling laws, as joint-stock companies, life insurance, and commodity futures [Ros] were prohibited before special interests managed to obtain exemptions.

Though unevenly enforced, such laws prohibit public science bets between strangers in all of the U.S. and in most of the world.

Only Great Britain, to my knowledge, allows such bets, and then only for the last three decades. 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. One cannot even subscribe to a publication listing the going prices on science questions. It should be possible to improve on this.

AN APPEAL

Idea futures is mainly just a curiosity to most people, even those convinced of its feasibility and desirability, . It would require substantial effort to implement, and in some sense is an trivially obvious idea, given the appropriate theoretical back-

Figure 3 A Few Concerns about Idea Futures

Isn't Gambling Illegal? Isn't Betting a Useless Zero-Sum Game? Does Anyone Ever Bet This Way? What About Compulsive Gambling? Is There Enough Interest in Science Questions? Will These Markets Be Too Thin? Doesn't Betting Only Work For Clear Cut Questions Like Horse Races? How Often Do Beliefs Really Converge? What If Beliefs Never Converge? What Do Convergent Beliefs Have To Do With Truth? What About Badly Worded Claims? Can't Wrong Ideas Still Be Useful? What If The Fine Print Differs From The Summary? What About Sucker Bets? Don't Science Questions Resolve Too Slowly? Why Should I Trust The Judges? Won't Judging Cost Too Much? Won't Wealthy People Have Too Much Influence? Won't The Market Be Dominated By Fools? Won't Advertising Manipulate Opinion? Aren't Markets Full Of Cheats And Thieves? What About Insider Trading? What About "Moral Hazard"? What About Incentives To Start False Rumors? What About Incentives To Keep Information Secret? Won't An Apparent Consensus Create A Crowd Mentality? Will The New Incentives Slow Or Stop Convergence? Won't Different Claim Wordings, Judges, And Base Assets Confuse The Consensus? Won't The Consensus Reflect Risk Preferences As Well As Beliefs? Won't Betting Challenges Discourage Creativity? What's The Point Of A "Consensus" That People Disagree With? Isn't It Better For People To Argue Out Their Own Disputes? Won't This Have The Same Problems As Patents? Wouldn't Anonymous Trading Screw Up Reputation Statistics? If This Is So Great, Why Hasn't It Happened Already? Won't greed sully the pure pursuit of ideas? Does a few dollars of compensation in the end really help a rejected visionary? Doesn't this presume there is some absolute truth? Won't convergence be culturally relative? Isn't consistency unhealthy? Doesn't organized crime take over anything having to do with gambling? What about libel and national security? What about "Nuclear war will destroy 90% of the world by 2020"? Won't different claim wordings, judges, and investment instruments fragment the market? Why should verdicts be consistent with one another? Won't judges be reluctant to contradict the market? What if the probabilities get very small? Why not do without judges?

ground. I think the only people who might actually be willing to work to make it happen are people who are particularly unhappy with current methods of forming and communicating scientific consensus, and how those methods have treated issues dear to them. People perhaps like *Extropy* readers, sympathetic to markets and subjects, like cryonics and uploading, which current consensus institutions deal poorly with. I fear it will require more effort than I alone can muster to make it real. It may well be that if you don't do it, no one will; what do you say?

There are many options for pursuing idea futures. I have worked to gain the attention of "science policy" academics [Han92], and idea futures will soon be a known, if oddball, suggested alternative mechanism for science funding. I have mostly developed a board game and to a lesser extent an email reputation game. Legal research is probably the most important task, but it is on hold for a lack of funds.

CONCLUSION

I have argued that futures markets in ideas could help the evolution of ideas by creating a visible consensus of relevant experts, and better incentives for honesty and care when making contributions. Idea futures might offer these and many other benefits cheaply and without coercion. Though some problems remain, it seems worth further study.

I leave the reader with this challenge: Can you think of a question where 1) you think the answer will eventually become clear, or would with enough study, and 2) you think you disagree with some generally perceived majority opinion? If so, imagine creating a market in that claim and then making a few trades.

MATH APPENDIX

VARIABLES:
P(A) = Market probability of A
J(A) = Judge's verdict probability of A
$X = An investment with a current market value of X dollars.
C = Total value of distributed coupons on a claim
F = Total budget available for judging
f = Maximum percentage judging fee
IDENTITIES: $X = "$X if True"
"($X if A) if B" = "$X if (A and B)"
EXCHANGES: (These remain valid if change $X to "$X if A", or multiply all $ by a constant)
Split/Join: $1 <-> "$1 if B", "$1 if not B"
Trading on A: $P(A) <-> "$1 if A"
On A given B: $P(A given B) <-> "$P if not B", "$1 if B and A"
Cash in with Judges: "$1 if A" -> $J(A) "$1 if not A" -> $(1 - J(A))
DUTCH BOOK EXAMPLE: If P(A) + P(not A) < 1, then can buy "$1 if A" and "$1 if not A" for less than $1, sell the pair to the bank for $1, and make a profit.
JUDGING FEES:
1) If f*C >= F, pay $F to judges, reduce coupon values $X -> $X*(1-(F/C))
2) If f*C = 0, no judging happens
3) If 0 < f*C < F, Take C and play a lottery: With probability C*f/F, increase value of coupons $X -> $X*F/(C*f) and do 1) Otherwise $X -> $0 and do 2)
MARKET MAKER ALGORITHM: (See Figure 2)
Choose a function M(i) from integers to [0,1] such that M(i) > M(i + 1), M(0) = 1/2.
Choose a transaction quantity Q.
Market starts at j = 0.
Offer "$Q if A" -> $(Q*M(j)) and if taken j -> j+1
Offer "$Q if A" <- $(Q*M(j+1)) and if taken j -> j-1
If M(i) = 1/(1 + exp(i/k)), total loss <~ $Q*k/2.

REFERENCES

[Bl] Black, F. (1971) "Towards a fully automated exchange", *Financial Analyst Journal*, July and November.

[Bru] Brunner, J. (1975) *The Shockwave Rider*, Harper & Row, NY.

[Cia] Cialdini, R. (1988) *Influence, Science and Practice*, Scott, Foresman and Co., Boston.

[DeF] De Finetti, Bruno, "Probability: Beware of Falsifications", 1976, in Studies in Subjective Probability, ed. Kyburg, H., Smokler, H., NY, Robert Krieger Publ. Co., 1980, pp 192-224.

[Deb] Debus, A. (1970) *Science and Education in the Seventeenth Century*, MacDonald, London.

[Dr] Drexler, K.E. (1986) *Engines of Creation*, Doubleday, New York.

[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.

[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.

[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) "Could Gambling Save Science? Encouraging an Honest Consensus" *Proc. Eighth Intl. Conf. on Risk and Gambling*, July, London.

[Han92] Hanson, R. (1992) forthcoming in *Social Epistemology*.

[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.

[Ho] Hofstee, W. (1984) "Methodological Decision Rules As Research Policies: A Betting Reconstruction of Empirical Research", *Acta Psychologica*, 56, pp93-109.

[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.

[La] Laffont, J.J. (1989) *The Economics of Uncertainty and Information*, MIT Press.

[Lea] Leamer, E. (1986) "Bid-Ask Spreads For Subjective Probabilities", *Bayesian Inference and Decision Techniques*, ed. P. Goel, A. Zellner, Elsevier Sci. Publ., pp.217-232

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

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

[Ros] 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. pp267-286.

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

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

[Ti] Tierney, J. (1991) "A Bet On Planet Earth", *Reader's Digest*, March, pp.61-64.

[Tul] Tullock, G. (1966) The Organization of Inquiry, Duke Univ. Press, London.

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

[Web] Weber, J., "Apparent observation of abnormally large coherent scattering cross sections using keV and MeV range antineutrinos, and solar neutrinos", Physical Review D, (38, 1) July 1, 1988.

[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.