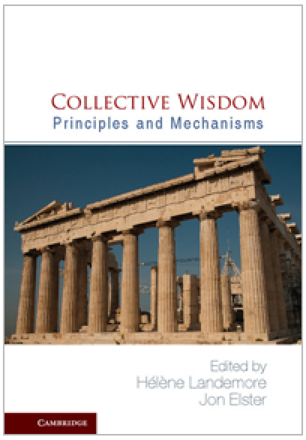

Publishing Reference
Servan-Schreiber, E. (2012). Prediction Markets: Trading Uncertainty for Collective Wisdom. In *Collective Wisdom: Principles and Mechanisms*, edited by Landemore, H. & Elster, J., Cambridge University Press

Prediction Markets
Trading Uncertainty for Collective Wisdom

By Emile Servan-Schreiber[1]

Collective intelligence in insect societies had long been a source of wonder for the educated mind, but in the first decade of this new century, something remarkable happened: the wide cultural realization that collective intelligence can be efficiently and successfully leveraged in human societies as well.

Perhaps because it knocks the human brain off its pedestal as the most intelligent thing in the universe, because, in some ways, it devalues the *individual* intellect, the notion that, as the Japanese proverb says, "none of us is smarter than all of us" is disturbing to many. Indeed, the "discovery" of human collective intelligence may yet launch a revolution of thought as profound as those initiated by Copernicus and Darwin, although, appropriately, no single individual's name may be attached to it.

It is the World Wide Web that has enabled human collective intelligence to burst into the public consciousness. Google, Wikipedia, and online prediction markets are the three poster children of what James Surowiecki (2004) termed "The Wisdom of Crowds".

For years now, billions of humans have experienced the power of Google's PageRank search algorithm, which relies primarily on the rich network of links from one website to another that results from countless individual webmasters' assessments of which pieces of information are related and worth linking to. Wikipedia, the collaborative encyclopaedia, has confounded sceptics with its breadth and depth, so much so that its entries are often found at the top of Google searches.

---

[1] Managing Director of Lumenogic, LLC (formerly NewsFutures, Inc), a consulting firm specialized in collective intelligence solutions to business problems.

But whereas Google and Wikipedia are content to aggregate existing knowledge, prediction markets take up the challenge of generating reliable knowledge about what is fundamentally unknowable: the future. At this ancestral human pursuit they have proven consistently more adept than the existing alternatives, such as individual experts and opinion polls. This *relative* success has captured the public's imagination and given them a slight aura of "magic" (Pethokoukis, 2004; Stuart, 2005). It is the purpose of this paper to investigate the sound economic, mathematical and neurological foundations of this particular form of collective wisdom.

The rest of this paper describes various prediction market designs, documents their record of successes, and identifies the ultimate drivers of their performance.

PREDICTION MARKET DESIGN

A prediction market is, at its core, a betting venue. Where it differs from a classic bookmaker's operation is that it removes the middle man and allows people to bet against each other through a trading system borrowed from the financial stock markets.

A classic design: binary contracts

One of the most popular market designs proposes "binary contracts" that pay $1 if an event happens or $0 if it does not. People can offer buy or sell this contract in various quantities at their preferred prices, in a process known as a "continuous double auction". For instance, one trader may offer to buy 15 contracts at $0.60 a piece, while another offers to sell 30 contracts at $0.65. When a buyer and a seller agree on the quantity and the price, the transaction happens. The buyer can then later turn around and resell her contracts at a higher price if she finds a willing buyer, or she can hold on to them until the outcome is decided and the market operator buys back all outstanding contracts at the expiry price of $1 (event happened) or $0 (event did not happen).

While the market is open for trading, speculative profits can be made by anticipating shifts in the collective opinion, buying low and then selling high, or selling high (on credit, also called "shorting") and buying low later.

Rational trading behavior

How would a rational trader behave in such a market? At any point in time, she could calculate the expected value of the contract as follows:

*Expected value of contract = Probability that the event happens * expiry price if it does*

So, for instance, if she gave the event a 65% chance of happening,

*Expected value of contract = .65 * $1 = 65 cents*

Accordingly, she would be willing to buy contracts that are offered at lower prices, and to sell contracts which receive higher bids.

How many contracts she would be willing to buy or sell at the market price would then depend on at least on five things: (i) how much her estimate of the contract's expected value differs from the market price, (ii) how much confidence she has in her estimate, (iii) how many contracts are actually available to buy or sell at the market price, (iv) how much cash (to buy) or how many contracts (to sell) she has at her disposal, and (v) the opportunity cost of trading this particular outcome's contracts rather than contracts for other outcomes that might be listed concurrently. Different combinations of these factors may yield exactly the same trading behavior, making it difficult for an observer to tell them apart. In practice it is very challenging to infer with any precision what a trader's belief actually is from just observing his trading behavior.

For instance, if a trader buys contracts at 65 cents, all we may infer is that she expects those contracts to be worth more, but we can't tell how much more. All other things being equal, she may be *very confident* that they will be worth *a little bit more*, or she may be *somewhat confident* that they will be worth *a lot more*. In reality, the situation is even more complex because the trade may just be speculative, that is, driven by expectations of how the market will move in the near term rather than by anticipation of the ultimate expiry price.

To profit: be right, before others

As *long* as a trader disagrees with the market price, she has incentives to trade. But she also has incentives to trade as *soon* as she disagrees with the market price, because the more she waits, the more others are likely to find out what she knows and to take advantage of the market themselves before she does, thereby erasing the profit opportunity. In a prediction market, as in any financial market, it is not enough to *be right* to make a profit; you must also be right *before others*. In this way, the prediction market provides incentives for both timely and truthful revelation of trader opinions.

Market price as consensus

As noted before, as long as a trader disagrees with the market price, she has incentives to trade. But, importantly, each of those trades will tend to move the market price closer to her estimate. In this way, just by trying to profit from a "wrong" market price, the trader shares her opinion with the other traders. In effect, it is impossible to seek to profit from the market without broadcasting information to other traders. That is what enables markets to aggregate the knowledge that is privately held by each trader.

With many traders coming to the market with different probability assessments, and different degrees of confidence in their estimates, the trading price eventually settles down where people "agree to disagree": the equilibrium price at which no one is willing or able to buy or sell. That consensus price holds until some new traders join the market, or some new information makes existing traders change their probability assessments, which leads to a new equilibrium. The more disturbing the new information, the more dramatic the change in price.

An example

Figure 1 displays the actual trading history of the "Obama to win the Democratic nomination" contract on NewsFutures, a popular "play-money" prediction market where people trade for prizes and bragging rights. The price of this contract experienced some dramatic changes as the primary race developed and as Obama's various successes and setbacks impacted traders' perception of his prospects.

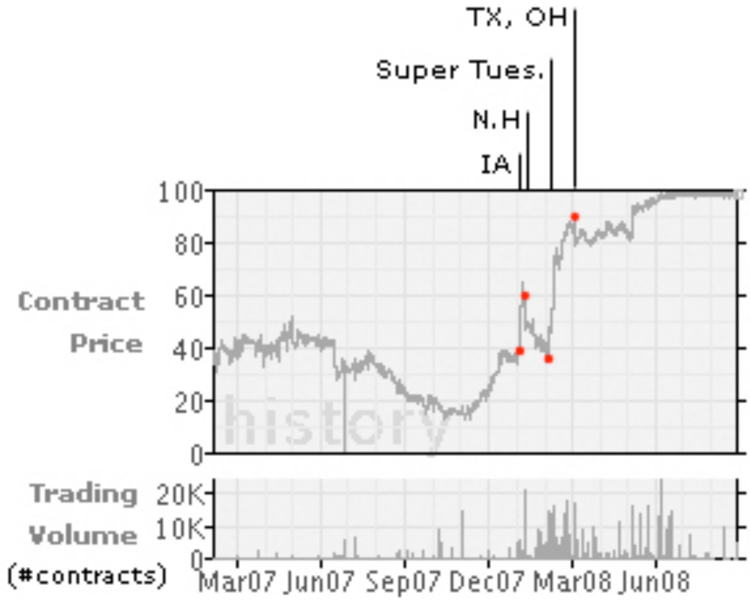

Figure 1: Trading history of the "Obama to win the Democratic nomination" contract on NewsFutures, from early February 2007 to the Democratic Convention at the end of August 2008. The contract would be worth 100 points if Obama became the nominee or 0 otherwise. Key milestones in the campaign were the Iowa primary (IA) which Obama unexpectedly won, then New Hampshire (NH) which he unexpectedly lost just a week later, then Super Tuesday, when Clinton failed to crush him, then Ohio and Texas, which Clinton both won, but too late. Red points indicate trading prices just *before* each milestone.

Binary contract price = event probability

What is the proper way to extract a prediction from the price of a binary contract? Well, if at any point in time the trading price captures the collective estimation of the contract's expected value, we can just directly derive from it the crowd's estimation of the event's probability. For instance, when the contract is trading at 65 cents (or points), it must mean that the market currently "believes" that the event has a 65% chance of happening. Empirical evidence that such predictions correlate well with observed event-frequencies in the real world will be presented further down.

When looking at a price chart such as Figure 1 that exhibits dramatic price changes, one might wonder if the market is really predicting anything. Since Obama did finally win the nomination, shouldn't his contracts have traded very high, at least higher than 50%, or perhaps just higher than all its rivals, *throughout* the campaign? Wasn't the market wrong when it priced Obama in the teens in the fall of 2007? Well, no. At any point in time, the market cannot be faulted for not taking into account *future* developments. In other words, only in a fully deterministic Universe could you fault the market for not divining right away the ultimate outcome. Rather, at each point in time there were possible futures in which Obama would win, and others in which he would lose, and all of these futures were just as valid, if not as probable. So it may well be that the market was right that, in November 2007, Obama had less than 1 in 5 chances of winning, and then was right again a few weeks later when, after his victory in Iowa, it gave him 2 in 3 chances.

Other designs: *index* contracts and *winner-take-all* markets

Other prediction market designs are meant to extract other types of predictions. The two other common designs are the "index" contract and the "winner-take-all" market.

The expiry price of an index contract, instead of being a binary 0-or-1 variable, is a continuous variable that depends on the value of an outcome. For instance, the Iowa Electronic Markets[2] (aka IEM), a political prediction market run by the University of Iowa's Tippie College of Business, features "vote-share" contracts: each candidate's contract pays 1 cent per percentage of the vote that he obtains. So if 56% vote for this candidate, the contract expires at 56 cents (Foresythe et al, 1992). In this design, the trading price captures the mean consensus estimate for a specific variable. Besides vote share, it has been used in various contexts to forecast a company's quarterly sales, the number of bug reports for a software product, or the completion date a large industrial project. Any continuous quantity can be forecasted in this manner.

In some situations, however, it may be insufficient to forecast just the mean estimate for a variable. One may instead want to extract the full probability distribution for the outcome. The winner-take-all design enables that through a combination of several binary contracts. For instance, to predict the sales of a widget, the full range of possible sales outcomes is divided into adjacent intervals, and the market proposes a binary contract for each interval that pays $1 if the interval includes actual sales or $0 otherwise. Since each contract is traded according to the probability that it captures actual sales, the set of prices reveals a probability distribution over the continuum of possible sales outcomes (Chen & Plott, 2002).

Whichever design is chosen, the prediction market performs three tasks: provide incentives for research and knowledge discovery (the more informed you are the more you can profit), provide incentives for timely and truthful revelation (the sooner you act on your information, the more profit you can make), and provide an algorithm for aggregating opinions (into the trading price). Its accuracy depends on how well it performs these various tasks.

EVIDENCE OF MARKET PREDICTION ACCURACY

In its prospective special issue, "the world in 2008", the magazine *The Economist* anointed prediction markets as today's "most heeded futurists," to the chagrin of the certified pundit whose glory days now seem past (Cottrell, 2007). Theory alone, or mere

---

[2] http://www.biz.uiowa.edu/iem/

laboratory experiments, could not buy this level of public respect. It has been earned with a record of success in the field.

Prediction markets have been available on-line to the general public since the mid-1990's, in both real-money (gambling) and play-money (game) formats, and a few have developed large communities of regular traders. Researchers have closely studied the predictions implied by prices in these markets; using the various designs described above, and have found them to be remarkably accurate, whether they operate with real-money or play-money.

Price-Probability Calibration

The prices of binary contracts, which imply event probabilities, are found to be closely correlated with observed event frequencies. This is verified by pooling a number of events that have contracts trading around a particular price; say 30 cents, and observing that about 30% of these events actually do happen. Figure 2 illustrates this with data from Servan-Schreiber et al (2004) who looked at market predictions for the winner of 208 NFL football games. Two markets were actually investigated, one using real money (TradeSports[3]) and the other using play money (NewsFutures), but in both cases the correlation was very high: 0.96 for TradeSports, and 0.94 for NewsFutures.

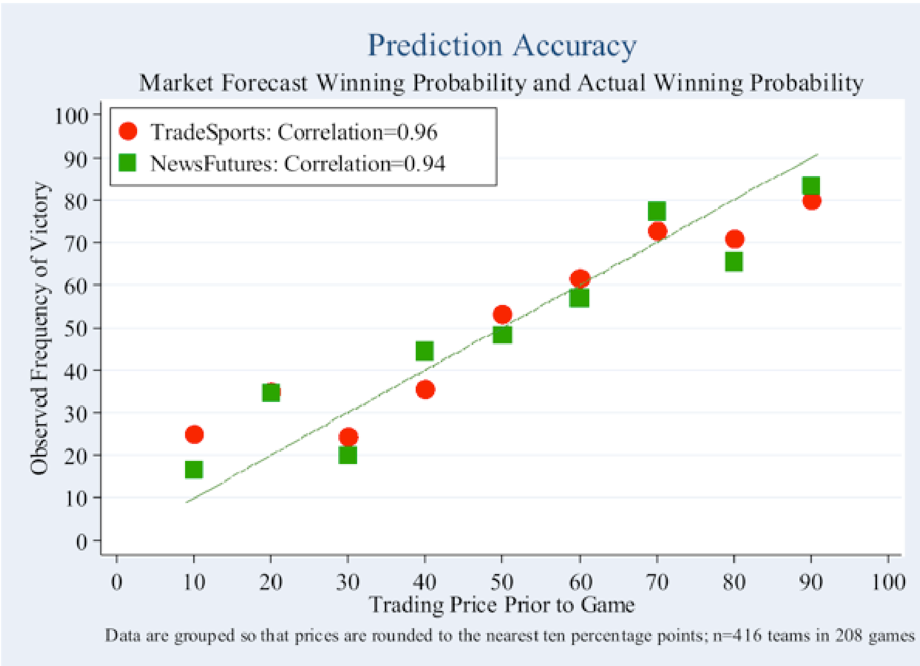

Figure 2: Calibration of market prices to observed event frequencies. Pre-game home-team prices for 208 games are rounded to the nearest ten percentage points, and the observed frequency of victory is plotted against these prices.

The calibration of contract price to observed frequency is a robust finding that has been reproduced in a variety of domains and prediction market venues. For instance, Pennock et al (2001a, b) documented it in Hollywood Stock Exchange[4] (aka HSX) predictions

[3] http://www.tradesports.com
[4] http://www.hsx.com

about Oscar, Emmy and Grammy awards, and also in Foresight Exchange[5] (aka FX) predictions about future developments in science and technology. An internal market operated by Google to pool the guesswork of its employees about various business issues also exhibits this calibration (Cowgill et al, 2008).

Note, however, that prediction markets are not immune to the classic "favorite-long shot bias", whereby high probability events are somewhat under-priced while low probability events are somewhat over-priced. This bias is apparent in the TradeSports and NewsFutures data in Figure 2, as well as in the HSX, FX, and Google data mentioned above. Erikson & Wlezien (2008) have also detected it in IEM election prices. Snowberg & Wolfers (2008) argue that this miscalibration is due to traders' actual misperceptions of the probabilities of very high and very low frequency events rather than to some form of "risk love" that might be encouraged by the wagering format. In any case, it is clear that the probabilities implied by very high or very low prices are to be taken with a grain of salt.

Index contracts, which directly forecast specific quantities, have also performed well in the real world. For instance, HSX "movie stocks" are indexed on the box office receipts of a movie over the first four weeks of its release. Looking at data from 280 movies, Elberse & Anand (2007) found a 0.94 correlation between forecasted and actual receipts. Evidently, despite the fact that, as everyone knows in Holywood, *nobody* can consistently identify blockbusters or flops in advance, HSX is able to do just that by aggregating the insights of many nobodies.

Political markets vs. polls

In the political domain, the IEM's "vote share" contracts have become famous for outperforming polls, both on election eve and in the long run (Berg et al, 2008). Market forecasts are also typically more stable over time than poll numbers. Although the advantage is small – a fraction of a percentage point, on average – these results are impressive because, contrary to polls, markets make no effort to recruit traders who are "representative" of the U.S. voting population as a whole: IEM traders, are typically young, male, Caucasian, well educated, etc.

Arguably, the further back in time from election day one looks at the data, the more awkward the comparison of market predictions with raw poll numbers becomes, because the polls are only meant to provide a snapshot of today's voting preferences rather a prediction of the ultimate outcome. However, in practice, this comparison is fair game because the latest poll numbers are so often reported by the media, and understood by the public, as being themselves predictive of the outcome.

A less naïve interpretation of polling data, Erikson & Wlezien (2008) argue, entirely dissipates the market's advantage, and even reverses it. Drawing poll *projections* from simple linear regressions on polling data gathered during all U.S. presidential elections since 1952, they show them to be superior forecasts of vote share than the IEM prices. Again, the advantage is small: a trader armed with their model would be able to profit from trading on the IEM by a modest 1.4% return on investment. A similar comparison

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

of IEM implied probabilities for winning an election versus educated projections drawn from historical polling data confirms the wisdom of trusting the poll-based projections rather than the markets.

What this critique demonstrates is that there is a wealth of information in the historical record of polling data that is not priced into the IEM, and that the popular perception that IEM prices are more predictive than polls misses part of the story: the IEM may beat the *raw* poll numbers, but it doesn't beat educated *projections* from the polling data. In fact, the advantage that the IEM has claimed over polls in the United-States has not been consistently reproduced elsewhere in the world (Bruggelambert, 1999; Servan-Schreiber, 2007), presumably because U.S. polling organizations report *raw* poll numbers whereas polling organizations in other countries publish educated *projections* instead.

However, if Erikson & Wlezien's critique helps dissipate some of the magic aura that has surrounded the political stock markets, it hardly reduces their practical merit. On a daily basis during a presidential campaign, the IEM and other markets of its kind still provide a better calibrated summary of the candidates' prospects than what is reported by polling organizations and the media.

Markets vs. individual experts

Outperforming individuals is what collective intelligence is all about. Servan-Schreiber et al (2004) repeatedly matched the probabilistic predictions of NewsFutures and TradeSports against those of 1,947 individual prognosticators of 208 NFL game winners. As the season developed, week by week, the predictions were scored by a proper scoring rule, which allowed the participants and the markets to be ranked according to their cumulative performance. In the first week, the markets were already performing better than 85% of the participants. By week 12, they were both in the top 1%, and they ended the 21-week season ranked 6th and 8th, outdone by only 0.26% of the participants.

Enterprise markets vs. official company forecasts

As the Google study mentioned earlier suggests, leading companies have started adopting prediction markets to pool the guesswork of their employees, and such markets have been consistently found to improve on the company's internal forecasts.[6] Popular market applications include forecasting sales of existing products, identifying promising new products, ranking projects, monitoring project deadlines, and measuring various business risks.

In a series of experiments that would later inspire many other companies to kick the tires of this technology, researchers at Hewlett-Packard enrolled a few dozen of the company's employees as prediction traders, and found that their forecasts of product sales outperformed the official ones 75% of the time (Chen et al, 2002).

---

[6] Besides Google, some of the better known companies that have implemented internal markets to date includes Hewlett-Packard, Eli Lilly, Pfizer, Siemens, Cisco, Best Buy, Motorola, Google, General Electric, Microsoft, and Electronic Arts.

Over many years of practice as a provider of collective intelligence solutions for businesses, Lumenogic has found that H-P's pioneering result is quite robust across companies and products. Within Eli Lilly, for instance, a dozen sales people were recruited in November 2004 to trade contracts indexed to sales of three of the company's drugs 2 weeks ahead, 6 weeks ahead, and 10 weeks ahead (so nine contracts in all). As Figure 3 illustrates, the market beat the official forecast 2 times out 3 in each time period, and this advantage was larger for forecasting targets that were farther into the future. The H-P and Eli Lilly experiments also show that, remarkably, prediction markets can work well even when the number of traders is quite small (compared to the public markets that were discussed earlier).

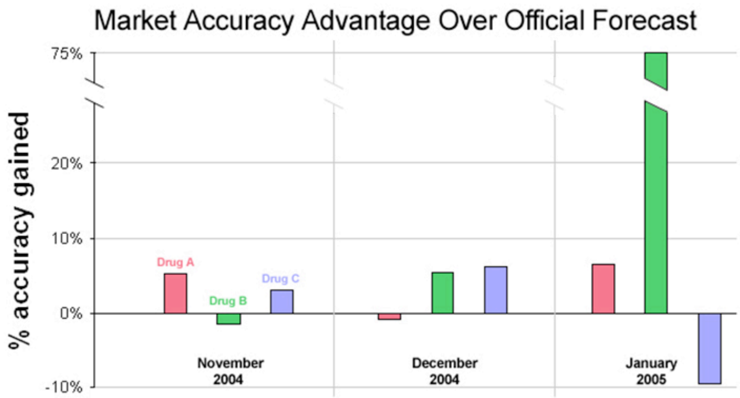

Figure 3. Sales forecasting at Eli Lilly for three drugs A (red), B (green), C (blue). All forecasts were made in early November 2004, for that same month, and for the two following months (December and January). The market beats the official forecast 2/3 of the time, and its advantage, combined over the three drugs, grows for targets farther into the future.

WHAT DRIVES MARKET PREDICTION ACCURACY?

Most of the research to date has focused on documenting accuracy in various contexts rather than trying to understand what produces it. In fact, trying to understand performance by looking at the trading data is a very hard problem because, despite being specifically designed to draw informed traders and to have them trade truthfully, real-world prediction markets like IEM, NewsFutures, TradeSports, FX and HSX are populated by human beings who behave not at all like the rational trader that was discussed earlier. That makes market performance all the more remarkable, but also more mysterious.

Irrational trading

One of the best documented biases in people's decision-making under risk (aka Prospect Theory, see Kahneman & Tversky, 1979), is called "loss aversion": people would rather avoid a risk than seek a reward of equivalent size. Typically, the subjective impact of losing money is perceived as *twice* that of winning an equivalent amount. In a prediction market context it means that traders will, for instance, avoid trades that present a 50/50 chance of winning $1 if their exposure is more than 50 cents. Or they will require a profit opportunity of at least $2 if they take a 50/50 chance of losing $1.

Another source of behavioral asymmetry is known as the “endowment effect” (Thaler et al, 1992): people often demand much more to give up an object they have received for free, than they would have been willing to pay to acquire it in the first place. Consistent with this, Oliven & Rietz (2004) have observed that traders in the IEM make fewer pricing mistakes on the selling side than they do on the buying side.

In a prediction market, even these well known biases are hard to isolate because they interact with risk/reward computations that are dynamic in time. In contrast to a standard betting venue where one risks something and then waits for the outcome to decide if he loses the entire investment or wins the entire reward, a prediction market allows the trader much more risk control: if the price starts going down, he can resell at a loss without losing all his investment. Similarly, he can take a partial profit on his investment by reselling at any time when the price is rising. For instance, a trader who buys a binary contract at 60 cents might consider that, for a chance to win 40 cents, she is really risking only 10 cents, rather than 60 cents, because, if worst came to worst, she could always liquidate her position at 50 cents instead of taking a complete loss. An observer, however, would not be able to deduce what risk/reward computation the trader conducted before buying the contract.

Another source of asymmetry in the market is that some traders usually command more, sometimes vastly more, monetary resources than others, and that gives them more power to move the price in their preferred direction. When the resources that a trader commands are unrelated to how informed or knowledgeable she is, knowledge aggregation should suffer. Servan-Schreiber et al (2004) cite this as a possible reason why some play-money prediction markets, which correlate trader wealth with past performance, can achieve a prediction accuracy comparable to that of real-money markets, where this correlation is unlikely to hold. Consistent with that, Oliven & Rietz (2004) report that higher income participants to the IEM tend to be more careless with their trades, presumably because for them the limited stakes[7] are comparatively lower. In general, when monetary resources differ, it is not reasonable to assume that a $10 bet represents the same perceived risk for everyone.

Market makers

Despite the daunting complexity real-world prediction market ecosystems, a few researchers have delved deep inside trading data to try to understand what drives performance. An early proposal from Foresythe et al (1992) is the “Marginal Trader Hypothesis”: in a sea of biased, emotional, error-prone traders, there exists a core group of traders whose behavior is more rational, and they are the ones who truly set the market prices. For instance, Oliven & Rietz (2004) observe that in the IEM there are essentially two types of traders: those who tend to propose prices to others (aka *market makers, aka marginal traders*), and those who tend to accept those offers (aka *price takers*). The roles are self-selected and may shift back and forth in the course of the market, but most people tend to stick to their preferred role. Looking closely at

[7] Eventhough the IEM is a real-money market, its stakes are limited: each trader can only invest up to $500.

thousands of orders from hundreds of traders, they found that market makers tend to behave much more rationally than price takers, making five times fewer trading errors (i.e., sub-optimal trades). The researchers conclude that because market makers have, by definition, a much more durable impact on market prices than price takers, their superior rationality helps the market as a whole overcome the wilder behavior of the other traders.

However, as Berg et al (2008) note, a rational trader who doesn't possess good information cannot take advantage of others' biases, so the "marginal trader hypothesis" isn't sufficient to explain market performance. To complete the explanation, they propose that markets are more accurate than publicly available predictions *when* those predictions possess a known bias. People who are aware of this bias are naturally attracted to a market where they can take advantage of their superior knowledge for profit. They self-select into the rational, marginal trader role and help the market settle on the appropriate price.

It's not about the market mechanism itself

On higher grounds, there is reason to doubt that any explanation which, like the marginal-trader hypothesis, is rooted in particular trading mechanics, can really get at the heart of a phenomenon that is reproducible across a diverse set of trading schemes. Indeed, while markets like IEM, TradeSports, NewsFutures and FX let traders negotiate prices directly with each other through a continuous double auction, HSX and a host of more recent venues rely on various automated market makers to set the price in response to traders' buy and sell orders.[8] Although the accuracy of these various trading schemes has yet to be formally compared, the wisdom of the trading crowds seems to emerge in every case.

Several non-market *betting* schemes have also been shown to perform just as well as prediction markets when matched head to head. For instance, Chen et al (2005) compared the prediction accuracy of NewsFutures and Tradesports to the simple averages of predictions elicited from 1966 individuals regarding the outcomes of 210 American football games. Those predictions were scored by the quadratic scoring rule, one of the so-called "proper" scoring rules designed elicit honest forecasts. Although the collective intelligence emerged forcefully in each case – outperforming all but a few individuals – there was no advantage for the markets over a simple arithmetic average of the individual predictions. Limited laboratory experiments have shown that when the number of people in a crowd is particularly small, say around a dozen, some simple betting schemes may even beat the market (Chen et al, 2003).

So, if prediction accuracy isn't manufactured by the market mechanism itself, where does it come from?

Attracting diversity

Scott Page's *Diversity Theorem* provides an elegant mathematical foundation for the wisdom of a crowd (Page, 2007):

---

[8] See for instance the Logarithmic Market Scoring Rule (Hanson, 2003) and the Dynamic Pari Mutuel algorithm (Mangold et al, 2005).

Collective Error = Average Individual Error – Diversity of Predictions

Or more precisely:

$$(P - O)^2 = average(p_i - O)^2 - average(p_i - P)^2$$

Where p_i are the individual predictions, P is the average of those predictions, aka the consensus, and O is the observed outcome.

In layman terms, when we average individual predictions, each of which contains bits of truth mixed with various misconceptions, the bits of truth add up to a larger truth whereas the misconceptions and biases cancel each other. (That is because bits of truth tend to be positively correlated whereas biases tend to be independent.) So the more diverse opinions are represented, the more complementary bits of truth can be combined, while the extra biases still get cancelled.

One implication is that you may obtain the same prediction accuracy from a crowd of amateurs that you can from a much smaller group of experts, as long as the individuals in the crowd hold sufficiently diverse opinions. In other words, the diversity in a crowd can make up for lapses in individual expertise. This trade-off is illustrated in Figure 4.

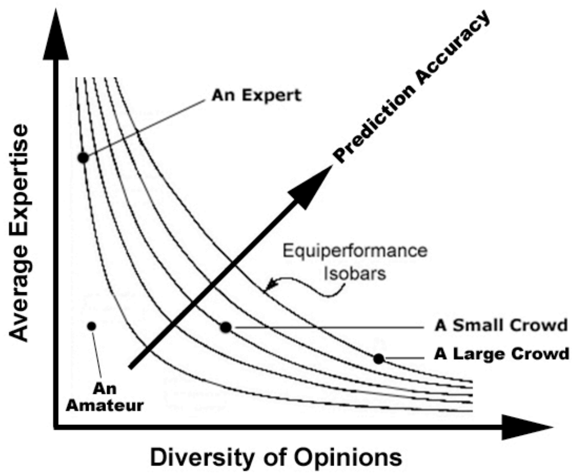

Figure 4: Page's Diversity Theorem applied to forecasting. The curves represent equiperformance isobars, so that every point on a particular isobar is a combination of average expertise and total diversity that yields the same prediction accuracy. In a crowd of one, diversity is null and an expert will out predict an amateur. However, a small group of amateurs endowed with diverse opinions may out predict an expert. Even as the average amount of expertise decreases when a crowd grows, it may more than make up for it with increased diversity.

The power of diversity is beautifully illustrated by a study of 2,231 football amateurs making predictions about 267 NFL games (Pennock, 2007, and Reeves & Pennock, 2007). Again, the predictions were scored by the quadratic scoring rule. They found that the number of individuals able to out predict a group decreased as a power function of the size of the group. For instance, while 64% of the participants scored above the average individual score, 16% scored better than the average group of five, and only 5% scored better than the average group of ten.

So one fundamental factor in the performance of any forecasting system is its ability to attract a crowd of people with diverse information and opinions. However, in the real world, it is often difficult or expensive to recruit diverse expertise. Budgets may severely constrain a manager's ability to solicit and aggregate the opinions of more than a handful of experts. Polling organizations need to keep up with fast changing communication trends, such as cell-phones, email, instant-messaging, to maintain their ability to recruit "representative" panels of consumers or voters.

In contrast, prediction markets excel at attracting diverse expertise because they incentivize disagreement and confrontation rather than conformity. If anyone has any reason to disagree with the market price, she has incentives to self-select into the pool of traders to contribute her contrary opinion. And, as noted earlier, the more she disagrees with the market, the more incentives she has to join it because of the larger profit opportunity.

Promoting objectivity

Perhaps the "explanation" that is given most often for the power of prediction markets is that, contrary to punditry and polls, they require participants to "put their money where their mouth is" (e.g., Hanson, 1999). As valid an insight as that may be, it needs some fleshing out before it may claim explanatory power.

First of all, it is necessary to qualify the use of the loaded word "money". As the comparative studies of play-money and real-money markets have shown, the presence of real currency in the market is not necessary to obtain the best performance: Play-money markets may perform as well as real-money ones (Servan-Schreiber et al, 2004; Gruca et al, 2008), as long as they provide *real incentives*, of which there are various sorts besides currency: for instance prizes, knowledge, professional advancement, or social recognition.

From a different perspective, brain imaging studies show that the mere prospect of risk and reward recruits into cognition specific brain modules. In particular, behavioral "loss aversion" seems to be driven by the activation or inhibition of specific neurons deep in the brain (striatum) as well as in the prefrontal cortex (Tom et al, 2007). Neuroscientists have also found evidence that, when under stress to perform a demanding task for monetary rewards, the brain inhibits signals from its emotional modules in order to maximize cognitive performance: the more important the rewards are, the more inhibition is observed (Pochon et al, 2002). In summary, when the brain contemplates a gamble, it literally *thinks differently* than when there's no perceived risk or reward: It becomes more risk averse and tunes out the emotional signals that might interfere with cognitive performance. "A true Englishman doesn't joke when he's talking about such serious a thing as a wager," remarks the hero of Jules Verne's *Around the World in 80 days*. Nor, apparently, do the rest of us.

Much of the advantage of prediction markets – indeed, of any betting institution – over pundits and polls may then simply come from framing the predictions as wagers rather than unaccountable opinions, thereby triggering more objective, less passionate thinking, and increasing the quality of the opinions that the market aggregates.

CONCLUSION

Prediction markets have captured the public's imagination with their ability to pool the guesswork of many to outperform venerable institutions such as individual experts and polls. Typical explanations for why they work so well put forward the efficiency of markets, or the insight that participants must "put their money where their mouth is."

Ultimately, though, it seems that the forecasting accuracy of prediction markets has little to do with any specific trading mechanism or even with trading in any way. Nor is hard currency necessarily involved.

The key driver of accuracy seems to be the betting proposition itself: on the one hand, a wager attracts contrarians, which enhances the *diversity* of opinions that are aggregated. On the other hand, the mere prospect of reward or loss promotes more objective, less passionate thinking, thereby enhancing the *quality* of the opinions that are aggregated.

We can speculate about the reasons why this conclusion is rarely voiced by advocates of this technology. In the United States, where most of the research on prediction markets has been done, and where most of the applications have been fielded, "bets and wagers" have much social stigma attached to them – they destroy families – and anti-gambling laws are strictly enforced. Markets, on the other hand, are respected institutions at the core of the national economy.[9] This encourages proponents to argue, or let believe, that more than betting is involved in a prediction "market," as when a score of high-profile economists, including several winners of the Nobel Prize, recently called for their legalization (Arrow et al, 2008). In Europe, where betting is more institutionalized, a more visible part of the social fabric, people simply use the phrase "betting exchange" to refer to the same thing.

Vocabulary aside, many practitioners have already recognized that the financial trading metaphor, which is so much part of the concept of a prediction market, is not always the best approach to harness the collective wisdom. On the one hand, it is customary nowadays to encounter forecasting systems that loosely call themselves "prediction markets" or "forecasting markets" for marketing purposes but which really are just sophisticated betting venues involving no trading. The chip maker Intel and the steel maker Arcelor Mittal, for instance, have relied on such systems to forecast product sales with apparent success (Llera, 2006; Hopman, 2007). On the other hand, even some of the most dedicated operators of prediction markets are themselves fielding applications that go to great lengths to hide the trading mechanics in order to reduce the user experience to a simple bet. For instance, the creators of the IEM have chosen this strategy when fielding their Influenza Prediction Market which is populated by doctors who are mostly unfamiliar with financial trading (Polgreen et al, 2007).

[9] Although the level of public respect for markets may have taken a hit as a result of Wall Street's subprime-fueled implosion.

A prediction market, then, is one of many betting-based methods for aggregating forecasts. In some contexts, for some purposes, it is an elegant solution, while in other situations it may be cumbersome. In any case, the ability of betting crowds to predict the future is a robust phenomenon that doesn't seem to be bottled up in any one particular contraption.

REFERENCES

Arrow, Forsythe, Gorham, Hahn, Hanson, Ledyard, Levmore, Litan, Milgrom, Nelson, Neumann, Ottaviani, Schelling, Shiller, Smith, Snowberg, Sunstein, Tetlock, Tetlock, Varian, Wolfers and Zitzewitz (2008). The promise of prediction markets. *Science*, 320, p.877

Berg, J.E., Nelson, F., & Rietz, T.A. (2008). Prediction market accuracy in the long run. *International Journal of Forecasting*, 24(2).

Bruggelambert, G. (1999). Information efficiency in political stock markets: An electronic market experiment. Working paper, University of Essen.

Chen, K.Y., & Plott, C.R. (2002*).* Information Aggregation Mechanisms: Concept, Design and Implementation for a Sales Forecasting Problem. Social Science Working Paper No.1131, California Institute of Technology, Pasadena.

Chen, K.Y., Fine, L.R., & Huberman, B.A. (2003). Predicting the future. *Information Systems Frontiers* 5(1)

Chen, Y., Chu, C.H., Mullen, T., & Pennock, D.M (2005). Information Markets vs. Opinion Pools: An Empirical Comparison. *Proc. of the 6th ACM Conference on Electronic Commerce (EC)*, Vancouver, BC Canada, June 2005.

Cottrell, R. (2007). The future of futurology. *The Economist – The Word in 2008*

Cowgill, B., Wolfers, J., Zitzewitz, E. (2008) Using prediction markets to track information flows: Evidence from Google.
http://www.bocowgill.com/GooglePredictionMarketPaper.pdf

Elberse, A, and Anand, B. (2007). The effectiveness of pre-release advertising for motion pictures: An empirical investigation using a simulated market. Information Economics and Policy, 19(3-4)

Erikson, R.S., & Wlezien, C. (2008). Are political markets really superior to polls as election predictors? *Public Opinion Quarterly*, 72(2)

Foresythe, R., Nelson, F.D., Neumann, G.R., & Wright, J. (1992). Anatomy of an Experimental Political Stock Market, *American Economic Rev.*, 82(5)

Gruca, T.S., Berg, J.E., & Cipriano, M. (2008). Incentive and accuracy issues in movie prediction markets. *Journal of Predicion Markets*, 2(1)

Hanson, R. (1999). Decision Markets. *IEEE Intelligent Systems*, May/June 1999.

Hanson, R. (2003). Combinatorial Information Market design. *Information Systems Frontiers* 5(1):107–119

Hopman, J.W. (2007). Using forecasting markets to manage demand risk. *Intel Technology Journal*, 11(2)

Jacobson, J., Dobbs-Marsh, J., Liberman, V., and Minson, J. (2011). Predicting Civil Jury Verdicts: How Ajorneys Use (and Misuse) a Second Opinion. *Journal of Empirical Legal Studies,* 8(S1)

Kahneman, D., & Tversky, A. (1979). Prospect Theory: An Analysis of Decision under Risk. *Econometrica, 47*(2)

Llera, J. (2006). Sparc: Arcelor's Prediction Market. *Presentation at the 1st European Prediction Market Summit – Vienna, November 2006*.

Mangold, B., Dooley, M., Flake, G.W., Hoffman, H., Kasturi, T., & Pennock, D.M. (2005). The Tech Buzz Game. *IEEE Computer*, 38(7)

Oliven, K. and Rietz, T. (2004). Suckers are born but markets are made: Individual rationality, arbitrage and market efficiency on an electronic futures market. *Management Science*, 50(3)

Page, S. (2007). *The Difference: How the Power of Diversity Creates Better Groups, Firms, Schools, and Societies*, Princeton University Press, Princeton, NJ

Pennock, D.M. (2007). The wisdom of the ProbabilitySports crowd. Oddhead Blog, http://blog.oddhead.com/2007/01/04/the-wisdom-of-the-probabilitysports-crowd/

Pennock, D.M., Lawrence, S., Giles, C.L., & Nielsen, F.A. (2001a) The real power of artificial markets. *Science*, 291(5506)

Pennock, D.M., Lawrence, S., Nielsen, F.A., & Giles, C.L. (2001b) Extracting collective probabilistic forecasts from web games. *Proceedings of the 7th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*

Pethokoukis, J.M. (2004) All seeing, all knowing. *US News & World Report*, Aug 22, 2004.

Pochon, J.B., Levy, R., Fosati, P., Lehericy, S., Poline, J.B., Pillon, B., Le Bihan, D., & Dubois, B. (2002) The neural system that bridges reward and cognition in humans : An fMRI study. *Proceedings of the National Academy of Sciences*, 99(8)

Polgreen, P.M., Nelson, F.D., Neumann, G.R. (2007). Use of prediction markets to forecast infectious disease activity. *Clinical Infectious Diseases*. 44(2)

Reeves, D., & Pennock D.M. (2007). How and when to listen to the crowd. *Overcoming Bias Blog*, http://www.overcomingbias.com/2007/02/how_and_when_to.html

Servan-Schreiber, E.J., Wolfers, J., Pennock, D.M., & Galebach, B. (2004) Prediction Markets: Does money matter? *Electronic Markets*, 14(3)

Servan-Schreiber, E.J. (2007). Results of the Dutch political stock market. NewsFutures blog,
http://newsfutures.wordpress.com/2007/02/22/results-of-the-dutch-political-stock-market/

Snowberg, E., and Wolfers, J. (2008). The Favorite-Longshot Bias - Examining Explanations of a Market Anomaly: Preferences or Perceptions? In Hausch, D.B. & Ziemba, W.T. (eds) *Handbook of Sports and Lottery Markets*, Elsevier, pp. 103-136.

Stuart, A.N. (2005). Market Magic. *CFO*, Nov 1, 2005.

Surowiecki, J. (2004). *The Wisdom of Crowds*. Doubleday

Thaler, R.H.., Kahneman, D. & Knetsch, J.L. (1992). The endowment effect, loss aversion and status quo bias. *R. Thaler The Winner's Curse.* Princeton University Press.

Tom, S.M., Fox, C.R., Trepel, C., Poldrak, R.A. (2007). The neural basis of loss aversion in decision-making under risk. *Science*, 315(5811)