An introduction to Prediction Markets[1]

Luca Iandoli
University of Naples Federico II, Italy

Abstract. Markets can be seen as an example of collective wisdom. In the following notes we provide some ground for this perspective. Furthermore we present Prediction Markets as the current technology that can be used to leverage the "wisdom of the crowds" and their applications for decision making.

Can groups be more intelligent than single individuals?

Under certain assumptions, groups can be smarter than their average member. Since the definition of what "smarter" means depends on the task to be carried out or the problem to be solved, in this paper we will focus our attention on one specific task: the prediction of future events. The accuracy of collective predictions is surprisingly high. Sometimes large random groups of people can be even more accurate than small teams of experts. Usually the following assumptions have to be satisfied for collective predictions to be accurate (Surowiecki, 2004):

1) people have diverse predictive models;
2) they are independent (they are not allowed to influence each other);
3) prediction process is decentralized.

Example of wise crowds - class experiment: guess my weight and age.

Why groups are accurate in collective prediction? Three possible models for information aggregation:

a) the competent minority ("who wants to be a millionaire" example)
b) aggregating pieces: the solution is the combination of partial answers
c) canceling out noise: random errors compensate each other

To work properly all the above models assume that: i) at least some people have (pieces of) correct information; ii) if they are not informed they will seek for information; iii) people are willing to disclose truthful information. Implicitly all the models assume some degree of diversity among people in the group. Implicitly all models assume some degree of accuracy for a crowd to be wise. An apparent paradox is why diversity leads to convergence to the best solution.

Example: designing pedestrian paths in Harvard courtyard, ants' systems, …

The diversity prediction theorem elaborated by Scott Page shows that collective wisdom is always a combination of individual accuracy and diversity. In other words we need diverse predictive model but this diversity is leveraged only if people in the crowd are smart enough. The theorem states that:

[1] These notes are a summary of different sources and have been created for class use only. Do not use or quote for other purposes without the author's permission.

Collective Error = Average Individual error – prediction diversity

If AIE is not small we need to counterbalance with some diversity. If people are all perfectly accurate, they do not need to be diverse (diversity = 0, but Individual error = 0). Since PD>0

Collective Error < Average Individual Error

Another interpretation of the above results can be provided by using a metaphor: the search of solution in a complex decision spaces. Diversity allows for the increase in the number of alternative dimensions of the problem (different people will look to the same problem from different complementary angles). Ability allows for: i) proper identification of relevant dimensions; ii) exploration skills (i.e. people are able to collect & filter relevant information and can recognize effective solutions when they see one).

When crowds are not smart?

When people are allowed to communicate freely we have a dilemma: deliberation can bring to the improvement of models through open criticism as well as to undesirable results like exasperation of conflicts or herding behavior. Common group pitfalls are: polarization, information cascades, hidden profiles, … Interestingly, polarized conflict, imitation and hidden profiles all imply a reduction of diversity.
Diversity may not be beneficial at all if it is about ends and not means. Diversity about ends means that people differ with respect to what is valuable, fair, and desirable. People with different values may have irreconcilable visions of give problems. Usually this kind of diversity brings to conflicts and polarization (in the debate people become even more extreme than they were before). Pro-Abort VS Con-abort is an example of such value-based debates.
Diversity about means implies that different people agree about the ends but use different problem solving strategies to achieve the same objectives. Diversity about means is the kind of diversity that has to be nurtured. Diversity about means also subsumes that people use a shared language and perhaps have some commonalities in their previous experience with a problem. We could find desirable to know what marketing people think of a product and have them cooperate with production guys, while it does not make much sense in most of the cases to have a biologist and an expert of ancient Greece literature to solve a problem in the area of biochemistry.

Definition of Prediction markets

Prediction markets are tools that leverage the wisdom of the crowds. They come through virtual trading software platforms and are used in distributed networks of decision makers, sometimes even in the form of public markets on the Internet.
PMs are a new—and emerging—form of financial market, often known as a prediction market, but also going by the name “information market” or “event futures.” Analytically, these are markets where participants trade in contracts whose payoff depends on unknown future events.

How they work
Substantially, like real markets. In a prediction market, payoffs are tied to the outcomes of future events. Participants trade contracts associated to the occurrence of a given event (e.g. Will Obama be re-elected in 2012? Who will win the Champions league? Will Italian GDP increase over 1% in 2010?). The market exchange of contracts determines their price: in general, the higher the price of

a contract, the higher the confidence of the market in the future occurrence of that event. Participants trade with real or play money; in any case they bet on the outcome they think is more likely (this does not exclude speculation attempts if some traders have hidden information or are just experienced traders whose objective is to make money through speculation).
In general 3 types of contract are available (see following table): winner-take-all, index and spread.

Table 1
Contract Types: Estimating Uncertain Quantities or Probabilities

Contract	*Example*	*Details*	*Reveals market expectation of . . .*
Winner-take-all	Event y: Al Gore wins the popular vote.	Contract costs $\$p$. Pays \$1 if and only if event y occurs. Bid according to value of $\$p$.	Probability that event y occurs, $p(y)$.
Index	Contract pays \$1 for every percentage point of the popular vote won by Al Gore.	Contract pays $\$y$.	Mean value of outcome y: $E[y]$.
Spread	Contract pays even money if Gore wins more than y^*% of the popular vote.	Contract costs \$1. Pays \$2 if $y > y^*$. Pays \$0 otherwise. Bid according to the value of y^*.	Median value of y.

In most prediction markets, the mechanism that matches buyers to sellers is a continuous double auction, with buyers submitting bids and sellers submitting asking prices, and with the mechanism executing a trade whenever the two sides of the market reach a mutually agreeable price.

Why they work

Here is a number of reasons:

Efficient market hypothesis. Much of the enthusiasm for prediction markets derives from the **efficient markets hypothesis**. In a truly efficient prediction market, the market price will be the best predictor of the event, and no combination of available polls or other information can be used to improve on the market-generated forecasts. This statement does not require that all individuals in a market be rational, as long as the marginal trade in the market is motivated by rational traders. Of course, it is unlikely that prediction markets are literally efficient, but a number of successes in these markets, both with regard to public events like presidential elections and within firms, have generated substantial interest.

Garbage-in/Garbage-out: the role of information. Another theoretical justification comes from studies on group decision making. The well known Condorcet theorem states that the probability a group of n individuals arrives at a correct decision increases with n if single voters are correct with a probability $p > ½$ (if instead $p < ½$ the probability of a wrong decision approaches 1 when n increase). If $p > ½$, the higher p the lower the number of people we need for a correct decisions Condorcet theorem and efficient market hypothesis have a lot in common: both are based on rationality and independence of participants. Both require large numbers to work properly (the market needs liquidity through a number of traders since only a minority of participant usually drive the market). Both require availability of information ($p > ½$ means that voters are more likely to be informed on what is the right outcome than not). However the market adds on the top of Condorcet theory some incentives: when people bet money, since they are rational and risk averse, they will have an incentive to collect good information. The incentive will cause p to grow. In other words,

the power of prediction markets derives from the fact that they provide incentives for truthful revelation, they provide incentives for research and information discovery, and the market provides an algorithm for aggregating opinions. As such, these markets are unlikely to perform well when there is little useful intelligence to aggregate or when public information is selective, inaccurate or misleading.

Uncertainty and diversity. Trade also requires some disagreement about likely outcomes. Disagreement is unlikely among fully rational traders with common priors. It is more likely when traders are overconfident in the quality of their private information or their ability to process public information or when they have priors that are sufficiently different to allow them to agree to disagree. Unlike polls, markets do not give the same weight to each person, but it's like they assign more weights to predictions made by people who are willing to bet more on an outcome. Those are usually the people that have more information and competence, while badly informed forecasters will be ruled out by the fear to loose their money. In this sense the market improves accuracy since it favors a sort of natural selection of the best models.
But market incentives may also encourage diversity when the solution is not so obvious. The nature of market reward is such that traders will gain a lot when they bet on unexpected outcomes: if the correct answer is obvious and everybody knows that, all bets will go in the same direction and the payoff will be virtually zero. If instead we have uncertainty and some buyers are more confident than others in a surprise and eventually they are right, they will get a very high payoff.
Unfortunately this is also the kind of reasoning that may trigger speculative behaviors (buy low to sell high), but speculation is less likely to occur in prediction markets and in any case it tends to have a short duration.
So the market provides both incentives for accuracy and for diversity.
While polls assign the same weight to everyone, and so they do not filter out bad information, expert-based methods do the opposite: they concentrate weight only on a few people (just one in the extreme case). An expert may be accurate but, by definition, it s not diverse from himself.
Following Page's theorem we miss all the benefits coming form diversity.
PMs are somewhere in the middle in the continuum between polls and the expert.

Applications

Examples:

- Public markets: Forecasting of political events (e.g. elections, decisions), economic statistics, movies successes, sports competitions
- Private companies markets: Increasing use of PM tools in companies for distributed forecasting related to internal or external events that are of interest for the organization: will project P end within the deadline? Will sales of product X be higher than y? Which product should we launch (opinions market)?

From a number of studies there is evidence that the market has both yielded very accurate predictions and also outperformed large- scale polling organizations (see in the political domain, Berg, Forsythe, Nelson and Reitz (2001)). Accuracy in many cases has proven to be higher than those produced by polls. See fig. 1 for an example: the graph also shows how the accuracy of the market prediction improves as information is revealed and absorbed as the election draws closer.

Figure 1

Information Revelation Through Time

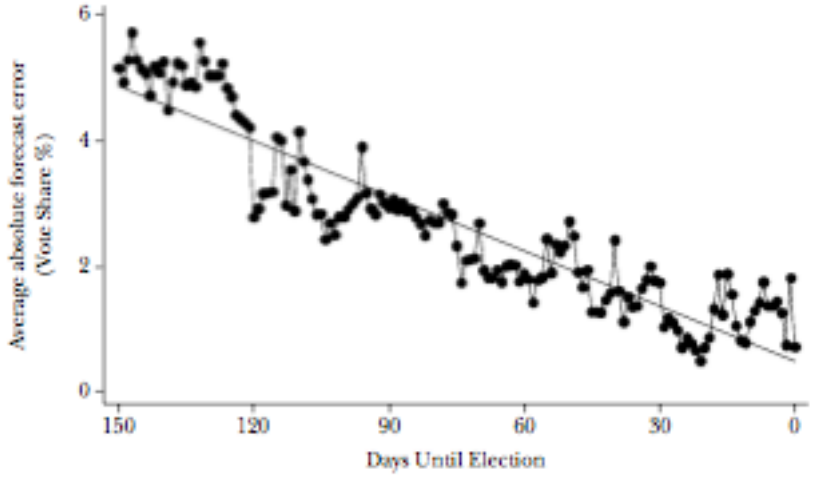

(Source Wolfers and Zitzewitz, 2004)

Prediction markets perform well also when compared with expert estimates (see the example in fig. 2 in which experts and market predictions are compared toward the likelihood of a war in Iraq to remove Saddam from power). This is a good aexample of how to link the time series of expectations elicited in prediction markets with time series of other variables (inferences). For instance, in Leigh, Wolfers and Zitzewitz (2003), we interpreted movements in the Saddam Security as an index for the risk of war and interpreted the comovement with the oil price shown in Figure 2 as a causal relationship, concluding that war led to a $10 per barrel increase in oil prices.

Figure 2

The Saddam Security

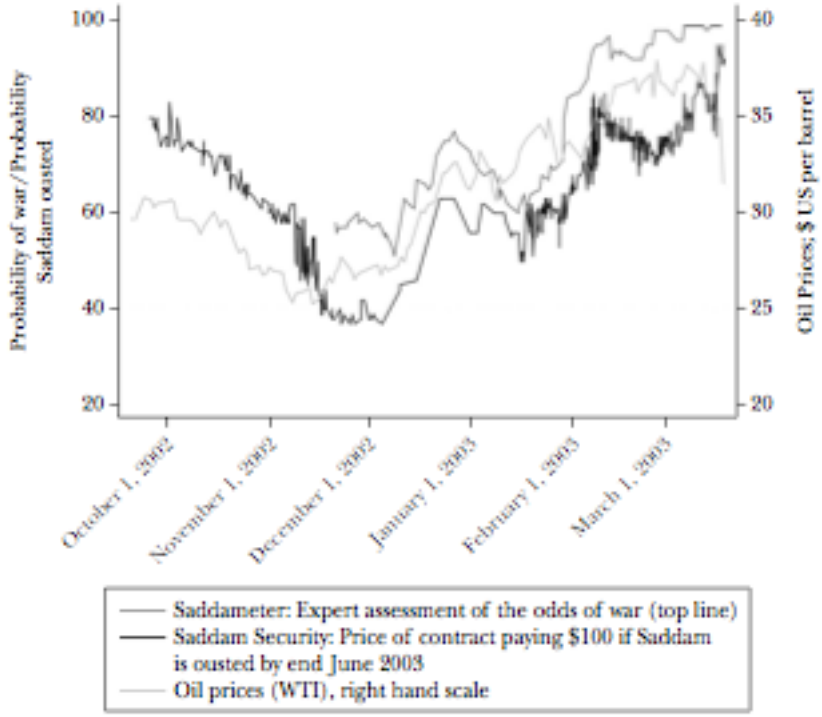

Sources: Trade-by-trade Saddam Security data provided by Tradesports.com; Saddameter from Will Saletan's daily column in Slate.com.

Evidence of performances and low cost of such tools have prompted many companies to adopt PM for business forecasting. A large number of company's employees can be involved in collective forecasts. Usually participants trade with virtual money that is eventually converted in prizes for best forecasters. Forecasts based on PMs in companies however may suffer from less accuracy (live over optmism, see Wolfers and Zitzewitz study on Google). One explanation is that in a closed group there is a higher chance to violate the assumption about the independence of traders (in other words traders may influence each other or just be less diverse than in an open market).

Futures market in finance are based on the same idea: a future contract is a standardized contract to buy or sell a specific commodity of standardized quality at a certain date in the future and at a market-determined price. Futures are not direct securities like companies stocks, but a type of derivative contract (i.e. title whose value depends on that of another assets, also called the underlying assets – e.g. oil). Futures can also been applied to non-commodities like currencies or other financial goods. A futures contract gives the holder the obligation to exercise the contract at the established date (e.g. sell the contract). Another kind of contract called "option" instead gives the right to exercise but not the obligation). Futures contracts may be sold or bought before the expiration date on the market.

Do PMs always work well?

Prediction markets do seem to display some of the deviations from perfect rationality that appear in other financial markets. There is substantial evidence from psychology and economics suggesting that people tend to overvalue small probabilities and undervalue near certainties.
Another behavioral bias reflects the tendency of market participants to trade according to their desires, rather than objective probability assessments. Strumpf (2004) provides evidence that certain New York gamblers are more likely to bet the Yankees.
A further possible limitation of prediction market pricing arises if speculative bubbles drive prices away from likely outcomes.

In addition to the possibility of irrational behavior the main limitation of a market is that it does not give any insight on what is the knowledge traders use behind their market decisions. In other words, a prediction market only aggregates information to tell us which is the probability of occurrence of an event in the future, but it will not say why the event is going to happen. While markets capture knowledge into process they do not capture the reasoning behind market behaviors. Sometime this can be obvious, sometime not so trivial. The market may process information effectively but the way this happens is a black box.

Suggested readings

Scott Page (2007), The difference: how the power of diversity creates better groups, firms, schools an societies. Princeton University Press.
James Surowiecki (2004). The Wisdom of Crowds. New York: Doubleday Press.
Justin Wolfers and Eric Zitzewitz (2004). Prediction Markets. Journal of Economic Perspectives—Volume 18, Number 2, Pages 107–126.