

Betting on a better world

How might prediction markets benefit international relations?

Presented at the annual meeting of the International Studies Association
New York, February 18, 2009

Emile Servan-Schreiber
NewsFutures, Inc.

At the end of 2007, a favorite magazine of the global elite, *The Economist*, famously declared that "the most heeded futurists these days are not individuals, but prediction markets, where the informed guesswork of many is consolidated into hard probability."¹ Thus was the "wisdom of crowds"² officially ushered into the rarefied field of respectable futurology where exclusive think tanks and charismatic pundits formerly ruled unchallenged.

Prediction markets leverage the wisdom of crowds to predict the future. I realize that this short sentence may set off several alarm bells at once. In today's context of sudden, unforeseen, market-based economic crisis, the terms "market", "crowd" and "predict", each already suspect individually, might inspire more than skepticism when mixed together.

However, in the two decades since the first modern prediction market began operating at the University of Iowa's Tippie College of Business,³ this form of collective intelligence has produced an impressive track record of accurate forecasting in fields as diverse as sports,⁴ movies,⁵ politics,⁶ business,⁷ and medicine.⁸ So the question is, might this technology be of value in the field of international relations as well?

¹ Cottrell, 2007

² Surowiecki, 2004

³ Foresythe et al, 1992

⁴ Servan-Schreiber et al, 2004

⁵ Pennock et al, 2001

⁶ Berg et al, in press

⁷ Chen & Plott, 2002; Hopman, 2007; Cowgill et al, 2008

⁸ Polgreen et al, 2007

Being no expert in International Studies, I won't presume to have a firm opinion, leaving it instead to the professionals that you are to evaluate for yourselves if prediction markets have anything to contribute to your field, and how they may best be applied.

But to get the discussion going, I will illustrate the promise and the difficulty to the exercise with several examples, including the failure to find weapons of mass destruction in Iraq, and the failure to launch the infamous DARPA-funded "Policy Analysis Market". Then I'll conclude with an actionable proposal.

However, for the benefit of those not yet intimately familiar with prediction markets, it is useful to first describe what prediction markets are, how they work, why they work, and how accurate their proponents really claim them to be.

PREDICTION MARKETS 101

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.

Buy low, sell high

The most popular prediction markets propose "binary contracts" that pay \$1 if an event happens or \$0 if it does not. An event might be, for instance, whether Vladimir Putin reclaims the Russian Presidency in 2009, or whether Iran and the US will hold a summit meeting before the end of the year.

Traders 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 100 contracts at \$0.33 a piece, while another offers to sell 25 contracts at \$0.37. 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).

Importantly, 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 conversely, selling high first (on credit) and buying later at a lower price.

Be right, before others

As *long* as a trader disagrees that the current market price captures the contract's fair value, she has incentives to buy it (if undervalued) or sell it (if overvalued). 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

mispricing 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

Importantly, each trade will tend to move the market price closer to the trader's own estimate of the contract's fair value. 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 publicly 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 estimates, and different degrees of confidence in their estimates, the trading price eventually settles down where people "agree to disagree". That equilibrium price holds until some new traders join the market, or some new information makes existing traders change their assessments, which leads to a new equilibrium. The more disturbing the new information, the more dramatic the change in price.

Contract Price = Event Probability

The most remarkable property of a binary prediction market is that at any point in time the price of a contract between \$0 and \$1 can be readily interpreted as the probability of the event being traded. So if the contract is priced at 75 cents, it means the event has a .75 probability of happening. If the contract is priced at 30 cents, it means 30% chance that the event occurs.

That interpretation is not just for convenience, nor is it wishful thinking. It is true calibration with reality. Numerous studies of prediction markets in various fields have shown that there is a high correlation between contract prices and the number of events at each price level that do occur in the real world. For instance, about 75% of the events that are priced at 75 cents really do occur, while only about 30% of the events priced around 30 cents really do occur.

This price-probability calibration captures both the power and the limits of market-based predictions. On the one hand, the ability to accurately estimate the probability of an event is very impressive. On the other hand, the non-intuitive nature of probabilities can lead to confusion, surprise and disappointment when, for instance, a highly probable event fails to occur, or when an unlikely event takes place. The natural inclination of the human brain is to wrongly interpret 90% likely as "sure thing", and 10% likely as "not a chance". We will come back to this when discussing the failure to find WMDs in Iraq.

How accurate ?

People are often curious about how accurate prediction markets are. They'll ask: "what's the accuracy rate?" Unfortunately, this is impossible to quantify in general, for it of course depends on the intrinsic difficulty of each prediction. The best we can say is that prediction markets are more accurate than certain common alternatives, especially individual experts.

Specifically, the data⁹ show that:

- (i) Over the long run and many predictions, markets outperform most individuals.
- (ii) The more traders there are in a market, the more accurate it is (although there are diminishing returns).
- (iii) The more traders there are in a market, the fewer individuals are able to outperform its accuracy.

Three sources of power

Prediction markets derive their predictive powers from three sources:

The first source is cognitive diversity. In a nutshell, when dealing with complex issues involving many variables or moving parts, no one can claim to have a complete model or theory from which to make fail-safe predictions. More likely everyone has a partial understanding of the situation, further clouded by his own biases. But when all these partial, biased models are put together, a wonderful thing happens: knowledge accumulates, gaps get filled, while the various biases cancel each other. The group's collective model is better and more complete than any individual model.¹⁰ Prediction markets attract diversity like powerful magnets, because anyone with a model that disagrees with the current consensus has a profit motive to participate in the market.

The second source of market accuracy is that the process encourages the voicing of informed contrarian opinions rather than calculated conformity or respectful consensus. Indeed, the only possibility for profit lies in disagreeing with the consensus, publicly. This makes sure that all informed points of views are included and aggregated. Importantly, the possibility of financial loss also discourages and penalizes the voicing of non-informed opinions.

The third source of market power is that framing predictions as wagers makes people "think different", literally: Brain imaging studies show that when it contemplates a gamble, the brain becomes more risk averse¹¹ and tunes out the emotional signals that

⁹ Servan-Schreiber et al, 2003; Pennock, 2007; Reeves & Pennock, 2007

¹⁰ Page, 2007

¹¹ Tom et al, 2007

might interfere with cognitive performance.¹² In short, our thinking becomes more objective and our judgements less clouded by our passions and preferences.

BENEFITS FOR IR FORECASTING

It's relatively easy to imagine how prediction markets might, in theory, help enhance forecasting in the field of international relations.

Likelihood estimates

The first benefit might simply be quantification of otherwise vague conjectures. For example, last november the Russian business daily *Vedomosti* floated the theory that Russian Prime Minister Vladimir Putin *may* be planning a return to the presidency, perhaps as early as 2009 through a mix of constitutional changes and the stepping down of the current puppet President Medvedev. But even as respected western media such as the *London Times* and the *Foreign Policy* blog picked up the story and spread the rumor further on the internet¹³, the actual likelihood of Putin's machiavellian maneuver was never assessed. So how should one treat this second-hand report based on a single anonymous source in an otherwise respected business daily? As an amusing thought experiment in Russian constitutional law? As a red herring, a trial ballon, or a real possibility?

While experts were at a loss to provide a quantitative or even a qualitative estimate how serious the return of President Putin might be, the NewsFutures prediction market¹⁴ - a popular prediction market among news aficionados in the US and France - simply had to list a contract for an answer to appear almost instantly: Since mid-december 2008, the price of the "Putin will reclaim the Russian Presidency in 2009" contract has been trading between 30 and 40 cents, implying a 30% to 40% likelihood for this outcome.¹⁵ Something, then, to take seriously.

Skeptics may howl, and this prediction may of course be taken with a grain of salt in the absence of any qualification of the traders' foreign policy chops. But that's not the point. The point is how quick and easy it is to get a precise consensus prediction about anything through a prediction market, even when experts are at a loss to voice an actionable opinion. (Another example: will there be an Iran-U.S. summit this year? Market says 50% chance.) Obviously, the more confident we are that knowledgeable people are trading, the more confident we should be in the market's collective estimates.

¹² Pochon et al, 2002

¹³ <http://blog.foreignpolicy.com/node/10232>

¹⁴ <http://us.newsutures.com>

¹⁵ <http://us.newsutures.com/market/market.html?symbol=PUTINBAK>

Dynamic updates

An appealing property of prediction markets - markets in general - is that prices react in real time to news developments. If the market is “efficient”, the latest trading price always incorporates all the relevant and latest information available to the traders. When the price moves, it means new information has come up, and the importance of this new information may be measured by the size of the price change. For instance, when Hillary Clinton failed to crush Obama on the Democratic primary’s Super Tuesday, the price of the “Obama will win the Democratic Nomination” contract on NewsFutures shot up from 35 cents to 75 cents within a week. On the other hand, when, a month later Clinton won both Ohio and Texas, Obama’s stock fell just 5 points, from 90 cents to 85 cents.

Reality check

As traditional media outlets sink like Titanic in the unwelcoming waters of the blogosphere, taking with them the last specimen of the tough-minded fact-checking reporters breed, we become more and more vulnerable to false beliefs amplified by the “group think” of misinformed talking heads. However, to participants in a prediction market, the false beliefs of others simply mean more profit opportunities... which is why prediction markets will tend to hug reality tighter than mainstream news media.

For instance, recall how, in the early days of the invasion of Iraq, everyone was certain that the weapons of mass destruction which were used to justify the attack would be found in Saddam’s arsenal. Previously, the CIA chief himself, presumably the best informed man on the planet, called the case for war a “slam dunk”, and General Colin Powell put all his gravitas on the line while dangling little vials of dangerous chemicals before the eyes of the world at the UN. Who could resist such an onslaught of righteous confidence?

A few weeks into the invasion, the Intrade¹⁶ prediction market started listing a contract on whether the WMDs would be found in Iraq in the next few months. At its peak, in late April, the price of this contract implied an 80% likelihood for this outcome. No weapons were found, and to most observers, in hindsight, the 80% prediction just seems flat wrong. Skeptics put it forward as a fine example of the uselessness of prediction markets. But this criticism is misguided on three counts:

- (i) A single probability prediction cannot, in principle, be right or wrong. About one in five of all “80%” predictions will in fact not occur, and that should not come as a surprise.
- (ii) Even if one wants to consider the prediction “wrong”, it at least had the merit to capture the consensus opinion based on the best *available* information.
- (iii) The market’s “80% likely” estimate was actually much closer to the truth than “slam dunk”, although perhaps a bit less actionable... which in this case may have been a virtue!

¹⁶ www.intrade.com

The point is that someone brainwashed by the media and the Government into thinking that WMDs would be found in Iraq would consider the 80 cents trading price on Intrade and have to wonder: why is it not trading closer to \$1? Why is it that some people are willing to sell me this contract at such a *low* price? What do they know that I don't? By forcing us to examine the basis and solidity of our beliefs, the prediction market makes us smarter and more in touch with reality.

Scaling

Other methods may be designed or may already be in use for aggregating the predictions of several experts. Polls or variants of the Delphi method may do a reasonable job of it. But prediction markets have the advantage of scale: the ability to involve hundreds or thousands of knowledgeable participants rather painlessly, and asynchronously. In the context of international affairs, where the relevant information and points of view might be spread widely around the planet, this property of scale might prove essential.

PRACTICAL CHALLENGES

Despite the potential benefits listed above, there remain practical challenges to the establishment of large scale prediction markets on international affairs.

Legal hurdles

In the western world, betting is tightly regulated. In the United States, federal law bans it outright. In Europe, outside of the United Kingdom and Ireland, betting on anything else than sports and horses is essentially forbidden as well. These legal hurdles strongly inhibit the creation of real-money markets in international affairs on both sides of the Atlantic.

The closest thing there is to such a market now is Intrade. Safely hosted in Ireland, from which it targets primarily U.S. customers, its development is nonetheless hampered by U.S. laws requiring banks and credit card companies not to process payments to offshore betting websites. Nor is Intrade allowed to promote its services with advertising in the U.S. and the bigger European countries.

Other companies such as NewsFutures have fielded play-money markets in order to operate legally in all jurisdictions. Although the accuracy of play-money markets may rival that of real-money markets when the issues traded have wide appeal¹⁷ - sports and presidential elections are prime examples - they are often unable to provide enough incentives to attract expert opinions on narrower topics.

¹⁷ Servan-Schreiber et al, 2004

Taste

In 2003, the U.S. Congress aborted the launch of a government-sponsored prediction market that was intended to deal specifically in international affairs. The so-called “Policy Analysis Market”, sponsored by the Defense Research Projects Agency (DARPA), was to propose bets primarily on various economic indicators and “political instability” in several Middle-Eastern countries.¹⁸ Bets would also cover more bloody topics such as the number of U.S. military casualties and terrorist-driven casualties in the western world, which, in some sense, is the bottom line.

Despite the fresh failure of the U.S. intelligence agencies to prevent the wild-goose chase for non-existent WMDs in Iraq, and the very obvious need for alternative ways to gather and share intelligence, U.S. lawmakers could not overcome their distaste for the concept of overt government-backed gambling on deaths and “instability”.¹⁹

The ethical issues raised by the PAM project are rather thorny and it is beyond the scope of this article to sort them out. But to this day, the sorry fate of this creative initiative demonstrates the practical difficulty of applying prediction markets to non-trivial issues.

A PROPOSAL

Between the real-money markets such as Intrade and PAM, which raise tough ethical and legal issues, and play-money markets such as NewsFutures, which lack the utility and cachet of their real-money counterparts, there may exist a “third way” to set up a useful yet ethical prediction market on international affairs.

Such a market may operate under the authority of a respected non-profit organization such as the International Studies Association. Only member of the ISA would be allowed to participate. Traders would be required to finance their trades with their own money, perhaps up to a limited amount to avoid excessive personal risk, but the profits could only be used to finance research grants or other relevant budgets such as conferences, books and work-related travel.

Such a market would enable participants to take advantage of their superior foresight in international affairs to raise money for their own research from the pockets of more hubristic colleagues. If that proved a popular proposition, it would allow the umbrella organization to draw widely on the brain power of its members to make quantitative, synthetic, consensus predictions about international affairs, perhaps thus contributing to more enlightened policy making and better world governance.

¹⁸ <http://hanson.gmu.edu/policyanalysismarket.html>

¹⁹ Surowiecki, 2003

REFERENCES

- Berg, J.E., Nelson, F., & Rietz, T.A. (in press). Prediction market accuracy in the long run. *International Journal of Forecasting*, in press.
- 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.
- 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>
- Foresythe, R., Nelson, F.D., Neumann, G.R., & Wright, J. (1992). Anatomy of an Experimental Political Stock Market, *American Economic Rev.*, 82(5)
- Hopman, J.W. (2007). Using forecasting markets to manage demand risk. *Intel Technology Journal*, 11(2)
- 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., Nielsen, F.A., & Giles, C.L. (2001) Extracting collective probabilistic forecasts from web games. *Proceedings of the 7th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*
- 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)
- Surowiecki, J. (2003). Damn the Slam PAM Plan! *Slate*. <http://slate.com/id/2086427/>
- Surowiecki, J. (2004). *The Wisdom of Crowds*.
- 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)