
Prediction Markets: The Collective Knowledge of Market Participants

Justin Wolfers

Associate Professor of Business and Public Policy

Wharton School, University of Pennsylvania

Philadelphia

Prediction markets provide an information-aggregation technology applicable to a variety of topics, including political and financial risk. Because of the human idiosyncrasies identified by behavioral finance, prediction markets can fail, but historical data show them to be as accurate as traditional polling methodologies and far less expensive to establish and maintain.

In this discussion, I will make three substantive claims about prediction markets. First, if markets really are efficient, as the efficient market hypothesis asserts, then the prices that come out of any market contain valuable information. To make the case for this claim, I will use data from sports betting markets. Second, prediction markets can be used to track political risk, which can be a key factor driving investment performance. Third, prediction markets can fail, so I will conclude my discussion by describing why prediction markets work and what causes them to fail.

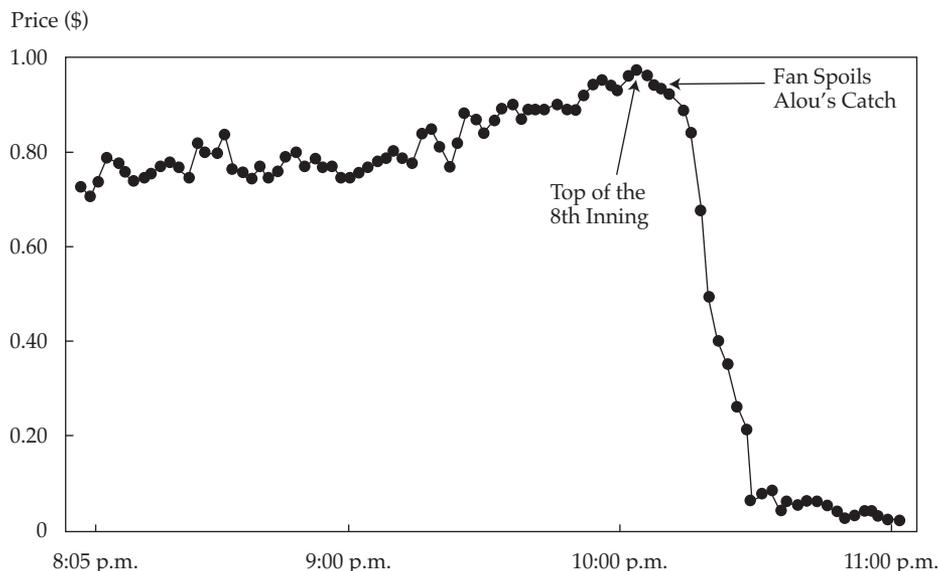
Prediction Markets and the Efficient Market Hypothesis

The efficient market hypothesis asserts that market prices fully reflect available information. But the degree to which information is available determines the form in which the hypothesis reveals itself. In its strong form, prices reflect all information; in its semi-strong form, prices reflect all publicly available information; in its weak form, prices reflect only past data. In any one of these forms, however, prices can reflect, summarize, and aggregate a huge amount of information, but the unfortunate implication of the hypothesis is that excess returns are unpredictable. Thus, prices can be said to follow a random walk.

This presentation comes from the 2008 CFA Institute Annual Conference held in Vancouver, British Columbia, Canada, on 11–14 May 2008.

Sports Prediction Markets. Interestingly, sports betting markets manifest many of the characteristics expected of an efficient market, and they now allow participants to trade stock during the game on the likelihood of a particular team winning. For example, **Figure 1** shows real-time betting during Game 6 of the 2003 National League Championship Series (NLCS). This prediction market from Intrade pays \$1 if and only if the Chicago Cubs win this particular game. At the beginning of the game, it looks as if the Cubs have about a 75 percent chance of winning. In the top of the eighth inning, the Cubs go ahead of the Florida Marlins 3–0, and the chance of the Cubs winning the game and going to the World Series reaches 95 percent. Unfortunately, in the bottom of that same inning, one of the great collapses in sporting history occurred. A fan reaches out from the stands and prevents a catch by Cubs outfielder Moises Alou. The Marlins proceed to score eight runs in that inning, and the market responds immediately to this new information. In a matter of minutes, the market price declines from 95 cents to about 5 cents.

The question more important than whether prediction markets respond quickly to news, though, is whether these market prices are accurate. To answer that question, I have charted data from all 21,885 major league baseball games played from 1991 through 2000, and the data show that the predictions are quite close to the actual outcomes. For instance, if the market says a team has a 33 percent chance of winning, it actually wins 31 percent of the

Figure 1. Real-Time Betting on the Chicago Cubs: 2003 NLCS, Game 6

Note: Time is given in central (Chicago) time.

Source: Based on data from Intrade.

time. When the market says a team has a 67 percent chance of winning, it wins 65 percent of the time, and when the market says a team has a 90 percent chance of winning, it wins 88 percent of the time. The difference is how Las Vegas makes money. But the real point is that if a person wants accurate information on the likely outcomes of baseball games, the betting odds do an excellent job of providing that information.

The same point holds true for other sports. For example, based on market expectations and game outcomes for the National Football League from 1984 to 2000 (i.e., for 3,791 games), a regression analysis shows a slope of almost 1.0 (0.997 to be exact). Thus, when the market says that a team is likely to win and gives a margin of a prespecified number of points (the "point spread"), that team, on average, wins by the point spread.

Do not, however, think of these findings as mere conversation topics. Rather, think of prediction markets as a technology that can be used to aggregate opinions about anything of importance.

Socially Useful Prediction Markets. Consider the firm called the Hollywood Stock Exchange. It runs prediction markets on the opening weekend box office take of various movies. Essentially, it allows people to trade, or bet, on how well each movie will perform. Data from that market can show the aggregated forecast and the actual open-

ing take, which are, again, quite accurate. Therefore, an analyst covering the movie industry may find these forecasts quite useful, which means that such an exchange moves us toward the realm of socially useful prediction markets.

For instance, Deutsche Bank and Goldman Sachs briefly started a new market (now folded, unfortunately) that they called the "economic derivatives market." It allowed investors to trade on the future outcomes of particular economic indicators, such as retail sales and business confidence. The economic derivatives market provided forecasts for nonfarm payrolls, initial unemployment claims, retail sales, and business confidence, and based on data that Refet Gürkaynak and I gathered, these forecasts were more accurate than an average or consensus forecast of 35 economists.

My conclusion, therefore, is that rather than averaging the forecasts of 35 economists, those same economists should be put in a room and told to bet or trade against each other. Taking the market-aggregated price from this trading would provide a more efficient way of aggregating forecasts. The task, then, is to harness this technology to more challenging issues. One such issue is politics. In fact, prediction markets have been used quite extensively to predict political risk, which is a key factor driving investment performance.

Prediction Markets and Political Risk

In 2003, a recall election was held to remove California Governor Gray Davis from office and replace him with one of several contenders, one of whom was Arnold Schwarzenegger. Prediction markets allowed people to trade on the question of whether Schwarzenegger would win the governorship. A contract paying \$100 if Schwarzenegger won could be traded on two exchanges—Intrade, which is based in Ireland, or the World Sports Exchange, which is based in the Caribbean. Attempts to recall other governors had occurred in California, but none had been successful. Nevertheless, I set my computer to look for arbitrage opportunities between Intrade and the World Sports Exchange. Every four hours, my computer checked these two websites and searched for potential profits. By the end of the exercise, I had made no money, but I had learned an important research lesson. When I followed the chart of the two exchange prices, the bid price in one market was never above the asking price in the other, which means that the two markets are closely linked. The price pattern that we are so familiar with in other financial markets appears to hold true for political prediction markets.

Accuracy of Political Prediction Markets.

These markets have a tremendous historical record. In 1988, faculty from the University of Iowa created a trading market now called the Iowa Electronic Markets. This market has followed hundreds of elections in the United States and elsewhere, so an extensive record exists to chart its predictive accuracy. Except for a few misses, the predicted outcome and the actual outcome of most elections converge along the same line.

Consider in particular the U.S. presidential elections. When I imagine trying to predict the winner of the next election, I typically think of two methods. The first is to create a traditional polling operation. I would hire hundreds of staff members, set up massive phone banks, call thousands of Americans around the country, and then have top-notch statisticians analyze the resulting data to come up with a forecast. The second method is to set up a website and let people log on and trade to their heart's content. I would then use the resulting prices to establish a prediction. The first method is quite accurate. It tends to predict U.S. presidential elections by roughly ± 2.0 percentage points, but it is also expensive to operate. The second method—that is, a political prediction market—is far less expensive, and it forecasts U.S. presidential elections to within ± 1.4 percentage points. Prediction markets,

therefore, represent a valuable gain in accuracy and a significant reduction in costs.

■ *Results from recent cycles.* In the 2004 U.S. election cycle, investors could trade not only which candidate would become president but also who would win the Electoral College votes of each state. Based on data from Intrade, the prediction market price accurately predicted every state that went for George W. Bush and every state that went against him. For the 2004 U.S. Senate races, the Intrade price was correct for every state but Alaska. In the 2006 election cycle, most knowledgeable observers and many Republicans did not expect the Democrats to win enough seats to gain the majority in the Senate. But every single Senate race went exactly the way the election-eve prediction market said it would.

■ *History of political prediction markets.* Such prediction markets seem like a marvelous new technology for forecasting election results, but as it turns out, betting on politics is as old as politics itself. Two enterprising economic historians—Rhode and Strumpf—have gathered data on the amount of betting that has occurred on U.S. presidential elections from 1884 through 1928. Opinion polling did not exist in any systematic or nationwide manner at the time, but betting on elections occurred regularly on the curb where the American Stock Exchange now stands. In fact, large sums of money were traded. For the 1884–1928 period, an average of \$37 million (in 2002 dollars) was bet per election, which turns out to be a \$2.28 bet per vote and the equivalent of 54 percent of total campaign spending. By using archival data, one can follow the performance of political prediction markets for well over 100 years.

In fact, in a comparison of the election-eve prediction markets for 1880 through 2004 with the actual results, once again, the prediction markets accurately indicated—with two exceptions—whether a Republican would win the election. In 2000, the prediction markets picked George Bush to win, which he did, but he did not win the popular vote as expected. The other exception occurred in the 1948 election and is, I think, the most famous exception in U.S. political history—the election in which Harry Truman unexpectedly defeated Thomas Dewey. Those two results aside, however, the prediction markets have been decidedly accurate.

Utility of Prediction Markets. Based on the historical evidence, the prediction market seems to be a wonderful tool, if for no other reason than it makes all of us political forecasters. The efficient market hypothesis says that all available information is embodied in the market price. Prediction markets, such as Intrade or the Iowa Electronic Markets, offer access to that very market price. Therefore, all I have

to do is look up one price and I become a political forecaster. I do not need to read the *Wall Street Journal*, the *Washington Post*, or the *New York Times*. In fact, I have begun acting on this very idea. I am now writing political commentary on the current election cycle for the *Wall Street Journal*, and the only data I am allowed to use are prediction market data. If the political prediction market actually aggregates all available information, that is the only information source any of us should be watching. We can ignore all the other commentary.

Figure 2 illustrates the sort of narrative I have been presenting so far. It shows the fluctuating price of a security during the 2008 Democratic nomination race that pays \$100 if Barack Obama becomes the nominee. Until his surprising victory in the Iowa caucuses, the price on the security was only \$20, but it quickly moved upward, to around \$70 at the time of the New Hampshire primary. At that point, however, Hillary Clinton's dramatic comeback caused Obama's stock to fall but not to pre-Iowa levels. The Michigan primary turned out to be no news for anyone, but Super Tuesday proved to be good news for Obama, and he got more good news with the Potomac primary. After Obama lost Ohio and Texas, Obama's stock declined again slightly.

Even after Obama lost the Pennsylvania primary, his stock value remained high. According to the *New York Times*, Clinton's victory in Pennsylvania was clearly important; however, seen through the lens of political prediction markets, she won Pennsylvania by the margin expected in a state with Pennsylvania's characteristics. As a result, her

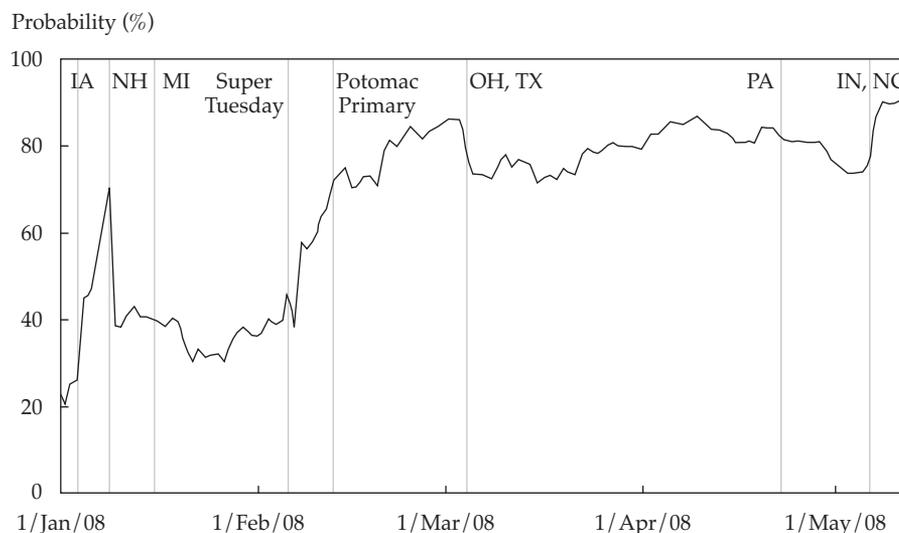
chances of winning the nomination did not change. The primaries in May led to a split decision—Indiana for Clinton, North Carolina for Obama—so both could claim a victory. But at this point, the political prediction markets decided that splitting victories was better news for Obama than for Clinton, so his stock price climbed dramatically.

The key point is that prediction markets provide an excellent way to track political risk. But the ability to track political risk is not an end in itself; its ultimate purpose, in this case at least, is its utility as a valuation factor.

Political Risks and Financial Markets. As a social scientist, if I wanted to find out whether Democrats or Republicans are good for stocks, I would like to do the same thing a physical scientist would do: Run an experiment in which I randomly assign the presidency half the time to Republicans and half the time to Democrats and then measure how stocks performed under each party. Fortunately for me, such an experiment actually occurred on Election Day 2004.

By following the changing probability of who—George Bush or John Kerry—would become president and comparing that probability with the simultaneous rise and fall of the S&P 500 Index, I was able to perceive the likely effect of either presidency on the stock market. For example, on 2 November 2004, from noon (when the probability of Bush retaining the presidency was 60 percent) to 6:00 p.m.—when exit polls put Bush's fortunes at

Figure 2. Probability of Obama Winning the 2008 Democratic Nomination, 1 January 2008 to 12 May 2008



their lowest point (i.e., a 30 percent probability that he would retain the presidency)—the S&P 500 fell 0.7 percent. Thus, if reducing the likelihood of a Bush presidency by 30 percentage points causes stocks to decline 0.7 percent, I can infer that removing George Bush entirely would cause stocks to decline 2.3 percent (−0.7 percent/−30 percentage points). As the actual vote counting began, however, the probability of Bush retaining the presidency rose 65 percentage points in just a few hours. Everyone had already voted, and the S&P 500 rose 1.3 percent, thus suggesting that stocks would be worth perhaps 2 percent more (1.3 percent/65 percentage points) under Bush than under Kerry. Republicans, I am sure, will be reassured to learn that their candidate is good for the stock market, and Democrats will be reassured to see that the difference between the two is actually rather small.

Equally interesting is the issue of which companies are expected to do well under which candidates. According to a paper by Knight (2006), a George Bush stock index includes such businesses as big oil, big tobacco, and big pharmaceuticals. By contrast, a John Kerry index includes environmentally friendly manufacturing companies and generic pharmaceuticals. From March to October 2004, as each nominee's fortunes rose and fell, so did their indices. As Bush's fortunes rose, so did his index, and when his fortunes fell, his stock index also fell.

Prediction Markets and Broad Political Issues. Prediction markets can be used to measure other forms of political risk. For example, before the United States went to war in Iraq, Intrade had a security that was worth \$100 if Saddam Hussein was no longer the leader of Iraq by June 2003. Because the United States appeared less likely from September through November 2002 to go to war, the Saddam Hussein security dropped from a value of about \$80 to about \$40. Then, as the likelihood of war rose for the next several months, so did the Saddam Hussein security, reaching a value of about \$75 in February 2003, a couple of months before the war actually began.

In contrast, the S&P 500 moved largely in the opposite direction of the Saddam Hussein security—a clearly negative correlation. One could interpret this movement as a warning from the markets, in contrast to the views expressed by most political and economic commentators, that going to war in Iraq would be bad for the U.S. economy. The data suggest that the difference between peace and war was worth as many as 15 percentage points for U.S. stocks. Based on these and similar results, it seems clear that political risk is a key factor in investment valuations.

The challenge is to find creative ways to deploy this technology to better forecast our future. By using prediction markets, qualitative issues that are hard to value can now be quantified and aggregated in a useful way.

Why Prediction Markets Work and What Causes Them to Fail

When it comes to large social issues—such as war or global warming—media commentators may err by taking popular rather than accurate positions. After all, they are not taking positions with their own assets. But that is exactly what people in the prediction markets, like those in the financial markets, do every day. So, prediction markets are quite effective at removing incentives to be popular instead of accurate. Similarly, prediction markets, like financial markets, encourage participants to seek out better information and act on that information, thus increasing their chances of outperforming other participants.

If no useful information is available to be aggregated, however, prediction markets are not especially helpful. After all, they are nothing more than an efficient means of aggregating available information. For example, when the prediction markets were asked whether weapons of mass destruction would be discovered in Iraq by May 2003, June 2003, July 2003, or September 2003, they were no more accurate than the rest of us.

Behavioral Finance and Long-Shot Bias.

Besides the question of available information, we must also keep in mind recent research in behavioral finance. Such research demonstrates that people do not always act in their apparent best interests. Therefore, human behavior can prevent prediction markets from being as efficient as possible. Consider one of the best-documented mistakes that individuals make in markets. The difference between a small probability and a tiny probability can be difficult to perceive, even in situations like horse racing, where reasonably reliable odds are known in advance. Despite this knowledge, individuals tend to bet too much on an extremely low-probability event relative to what they are prepared to wager on an outcome with better odds. Anyone who follows the markets should be aware of this reality.

Behavioral Finance and Human Preconceptions. Another problem in markets is that we all bring our preconceived notions to the markets. To help people understand this problem, I sometimes conduct a little thought experiment. When I ask people to imagine a great CEO, most

people will choose a tall, white male wearing a pinstripe suit. Thus, female CEOs do not fit the typical representation of a great CEO, and investors may be likely to undervalue companies directed by a female CEO. To investigate this hypothesis, I gathered all the earnings announcements in the United States during the last decade and organized them by the size of the earnings surprise that occurred. Because women represent about 1 percent of all CEOs, they also represent about 1 percent of all earnings announcements. But when I look at the largest positive earnings surprises, women are dramatically overrepresented, and when I look at the largest negative earnings surprises, women are dramatically underrepresented. This finding implies to me that market participants systematically underestimate the value of companies directed by women. Not only that, but further analysis of the data shows that female CEOs have a tremendous propensity to surprise male analysts.

Conclusion

A few things need to be kept in mind when dealing with prediction markets. First, if one assumes that the efficient market hypothesis is correct, market

prices often provide efficient forecasts, including forecasts of political risks. Making money in the markets may be difficult, but prices offer a useful summary statistic of available information. Second, prediction market prices can be useful inputs for decision making in finance, management, and public policy. Third, the success of prediction markets is context specific, and the idiosyncrasies of human behavior found in behavioral finance may provide hints of profit opportunities. Finally, prediction markets offer two particularly interesting uses to investment professionals: (1) providing market-aggregated forecasts of political risk as an input to the investment process and (2) providing a tool for information aggregation within investment firms themselves. Most companies have an information-aggregation mechanism that includes neither accountability nor a reward system. It is called a “staff meeting.” A prediction market within a company could be used to aggregate information from staff in a far more efficient manner.

Editor’s Note: Justin Wolfers provides consulting services to several prediction markets.

This article qualifies for 0.5 CE credits.

REFERENCES

- Gürkaynak, Refet, and Justin Wolfers. 2006. “Macroeconomic Derivatives: An Initial Analysis of Market-Based Macro Forecasts, Uncertainty, and Risk.” NBER Working Paper 11929 (January).
- Knight, Brian. 2006. “Are Policy Platforms Capitalized into Equity Prices? Evidence from the Bush/Gore 2000 Presidential Election.” *Journal of Public Economics*, vol. 90, no. 4–5 (May): 751–773.
- Rhode, Paul, and Kolemian Strumpf. 2004. “Historical Presidential Betting Markets.” *Journal of Economic Perspectives*, vol. 18, no. 2 (Spring):127–141.

Question and Answer Session

Justin Wolfers

Question: Are prediction markets subject to manipulation?

Wolfers: Many attempts have been made to manipulate prediction market prices. A good example occurred when Pat Buchanan was running for the Republican presidential nomination. He encouraged his supporters to buy Pat Buchanan stock. Their behavior, however, did not change the reality that Pat Buchanan had little chance of winning the nomination. Therefore, if someone says I am willing to buy stock in Pat Buchanan, someone in the market will sell. But the very nature of free entry and the supply response meant that Pat Buchanan's supporters could not affect the price for more than a couple of hours.

So, manipulation attempts sometimes occur, but I am aware of no attempt that has succeeded for a sustained period. There were manipulation attempts again in the 2008 presidential markets, but these too ultimately were unsuccessful, as prices reverted to fundamentals over the course of a couple of weeks.

Question: Can widespread use of prediction markets influence subsequent behavior in the sense of persistence momentum?

Wolfers: At this point in the political cycle, prediction markets are still a fringe phenomenon. The number of words devoted by the press to prediction markets compared with those devoted to polls is sufficiently small that we probably have little need to worry about persistence momentum.

Question: Is there a difference in the behavior of prediction markets that use real money and those that use play money?

Wolfers: Because I am an economist, I would lose my union card if I did not say that real money always makes a difference. One of the important things about prediction markets is that they ask participants to take positions with their own money.

Having said that, I once spoke with someone who had created a prediction market company that uses play money, and he bet me that his play money market could do as well as a real money market. So, we set up a competition in which his market predicted the winner of Saturday football games while I used Intrade, a real money market, to predict winners. By the end of the season, the two markets had amassed comparable prediction records. In that instance, therefore, money did not make a big difference.

I can offer two explanations for this result. First, prediction markets—whether using play money or real money—do something useful that polls do not do. Prediction markets ask the right question. They do not ask, as polls do, “Who would you like to win?” or “Who do you plan to vote for?” Prediction markets ask, “Who do you think will win?” Second, there is no way to get rich in play money prediction markets. Therefore, there are no big players influencing such markets, which is not the case in real money markets.

For example, Warren Buffett has a big influence on U.S. equity markets. That is probably a good thing, and it makes the market more efficient. But there may also be large pools of uninformed money in real money markets, which may make the real market less efficient. Without uninformed money in a play money

prediction market, a play money market tends to be more efficient and thus may approach the efficiency of a real money market.

Question: Why are some of the prediction markets used for economic analysis no longer available?

Wolfers: Which markets survive and which do not is a function of which markets are profitable to run for those who run the market. For example, Goldman Sachs and Deutsche Bank were charging what I considered unbelievably high fees, which proved to be a disincentive to trade. So, it was not a profitable business for them to stay in. Nevertheless, if such a market existed, many of us would find it valuable.

Many governments around the world are thinking about running such markets within their bureaucracies so that, as I mentioned earlier, their economists will have a more efficient aggregation tool than staff meetings. I foresee such markets being used within organizations more often than between organizations.

Question: How far in advance do prediction markets work?

Wolfers: Let me make two observations. First, the closer we come to any event, the easier it is to predict the outcome. Second, at any given horizon, prediction markets seem to do better than the alternative.

Question: What impact do tail events have on prediction markets?

Wolfers: Every time a tail event happens, I get calls asking whether the prediction markets are wrong. This happened after

the 2008 New Hampshire primary. The New Hampshire primary markets said there was only a 7 percent chance Clinton would win; Clinton won. When such things happen, we tend to forget a truism about 7 percent probabilities, which is that they will happen 7 times out of 100. If they happen less than that or more than that, then the markets are in trouble.

But also keep in mind that prediction markets are not well calibrated over very small probabilities. When markets are predicting between a 10 and 90 percent chance of an event occurring, I take that as a fairly good assessment of probability. Nevertheless, prediction markets seem to be fairly bad at telling

the difference between a 5 percent chance, a 2 percent chance, and a 1 percent chance.

Question: What are some of the better prediction markets?

Wolfers: Let me say up front that I have a conflict of interest here because I advise many of these companies. That said, Intrade has a great deal of good data and offers perhaps the best U.S. political coverage. Another excellent one is Betfair, although it has more British than U.S. political coverage. The Iowa Electronic Markets is predominantly political, so it has a narrower range than the others.

Question: How much volume is required for the prediction markets to be seen as accurate?

Wolfers: Remember how well play money prediction markets do with no money? Then you should remember, too, that high volume times zero price equals zero dollar volume. Yet, play money markets do quite well despite their zero dollar volume. In fact, my observation—and this is just an intuitive answer, not a scientific one—is that prediction markets do surprisingly well with only a small number of participants.

I have seen within corporations perhaps a dozen people trade in a market and do quite well. And remember, the right metric is the accuracy of the forecast. Is it more accurate than you would have had otherwise?