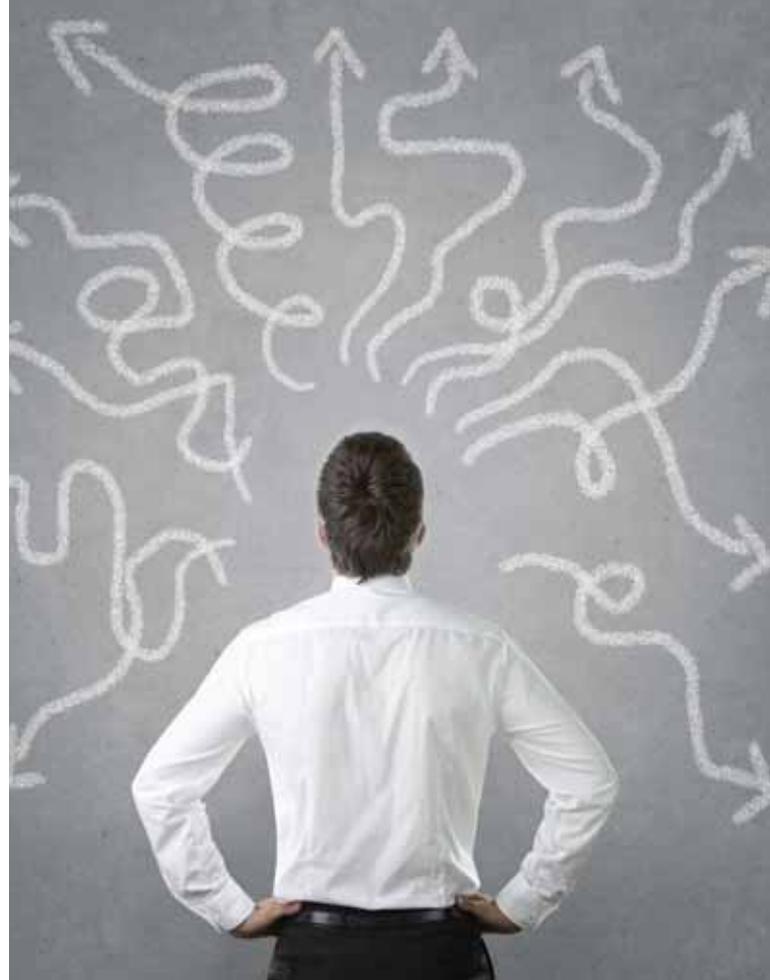# Outcomes, not intentions

The power of prediction markets

| By Julie Wittes Schlack

snapshot

What's the allure of prediction markets? Do they work? If so, how? Should researchers be interested?

If you've ever run a concept test, you've probably experienced the dull ache of having all of your concepts average a 3.5 on a five-point rating scale for appeal, uniqueness and purchase intent. You may even have suffered the acute pain of seeing a more highly-rated concept – one that clearly led the pack – fail miserably when the product actually went to market.

While the joy of collaborating with consumers lies in discovering their aspirations, frustrations and unmet needs, then in co-creating new concepts, campaigns and products with them – the suspense lies in seeing if the outcomes you hope for come to pass. And since marketers (and market researchers) hate suspense, a growing number are turning to prediction markets as a means of prioritizing, optimizing and ultimately predicting the success of products, messages, promotions and campaigns.

Beaten out fellow investors

Imagine that in 2010, you were presented with several potential investments, including a water-powered cellphone battery that didn't require a charger, a touch-screen-based tablet computer and a waterless washing machine. You might have decided to hedge your bets, putting equal amounts of money into all of these ideas. You might have decided that nobody in their right mind would turn down a waterless washing machine and put all of your chips on that product's success. Alternatively, you might have bet heavily against its rapid adoption and, early on, put much of your money on the touch-screen tablet, aka the iPad. If you had, you'd not only be a millionaire but you would have beaten out your fellow investors by investing heavily in the right predictions – that Xeros's waterless washing machine would be slow to release and win adoption in the U.S. and that the iPad would succeed instantly – and doing it early enough for your shares to gain in value.

The same principle applies to a prediction market, the purpose of which is to

anticipate the likely success of an idea, product or political candidate. Unlike traditional concept-test surveys, where you try to assemble a representative sample of likely consumers and ask about their personal preferences and intentions, a prediction market relies more on diversity in the participant pool than on representativeness. And rather than ask "What do you like" or "What are you likely to do," a prediction market asks "What is going to happen?"

For example, a brand might pose yes-or-no questions like "Will this product appeal to 16-to-18-year-old women?" or "Will Product A outsell Product B?" or "Will this offer motivate non-customers to sign up for a free trial?" Alternatively, they can pose multivariate questions such as "Which of these products will be most successful among 16-to-18-year-old women?" or "Which of these offers will motivate non-customers to sign up for a free trial?"

Each participant in the market is given play money or points to invest in their answers. They answer only the questions about which they have a strong opinion; can invest in the likely failure of an idea as well as in the likely success of one; and invest as few or as many points as they want based on the strength of their confidence in their own predictions. Furthermore, when they invest, they provide an explanation of why they're doing it. When the market closes, each prediction ends up with a numeric value representing the probability that it will come to pass. If someone invests early and the other investors in the market tend to agree with them (i.e., if the market consensus coalesces around their prediction) then their "net worth" will increase. On the other hand, if the market consensus runs counter to their prediction, the value of their "investment" will decline.

When real-world results are imminent (as is the case with political or sports markets), market "investors" can be rewarded based on the accuracy of their predictions. But in the context of new product development and market research, when it may be a year or more before a product is actually released to the real-world market, rewards can be based on participants' net worth when the market closes. For example, at Communispace, where members of our private online communities participate in proprietary prediction markets set up exclusively for a given client, we both motivate and reward participation by giving a prize to the top five traders – those with the greatest net worth when the market closes.

Employed internal prediction markets

Prediction markets originated in two places: the University of Iowa's markets have been used to predict the winners of elections since 1988; and a number of major companies, including Hewlett-Packard, Motorola, Intel, Best Buy, Microsoft, Google and Pfizer Animal Health, have employed internal prediction markets (where the traders are employees) to assess likely product shipment dates, predict sales figures and volumetrics and identify best-selling products.

Across the board, results have been as good as if not more accurate than other methods. For example, Hewlett-Packard's employee prediction markets produced sales forecasts that were generally more accurate than official company forecasts (Chen and Plott 2002) and Intel's markets for predicting product demand have been as much as 20 percent more accurate than official forecasts (Hopman 2007). When Iowa's Electronic Market predictions were compared to 964 polls over the course of five presidential elections, the market was more accurate than the polls 74 percent of the time overall and, more significantly, outperformed polling when the election was more than 100 days out (Berg, Nelson and Rietz 2008).

But what about in the consumer domain? Can consumers predict the likely behaviors of other consumers any more effectively than traditional quantitative tools can? According to a 2009 study:

Market research experts have conducted prediction markets in conjunction with conjoint analysis that yielded correlations between 0.715 and 0.885. In addition, features identified by the prediction markets as "Preferred by a Majority" appeared as features in new product launches later the same year, providing additional validation of prediction market accuracy. (Dahan, Soukhoroukova and Spann, 2009)

And a more recent study by BrainJuicer measuring the performance of prediction markets concluded "a .085 correlation versus monadic concept testing in 215 head-to-head experiments." (Kearon and Earls, ESOMAR, 2009)

Moreover, a significant body of academic and corporate research suggests that you need neither large nor necessarily representative sample sizes to arrive at accurate answers. Still, many researchers are wary of a method based on participants' knowledge and judgment rather than on their social and demographic qualifications. That's why with practically every prediction market that we've run, we or our clients have conducted at least one parallel study (Table 1). In every case, our results have been almost identical.

But the acid test is how well the market predicts real-world outcomes. Given the typical duration of a new-product development cycle, we don't yet have many results to report on. However, where we do, they have also been promising. Internally, when we ran an internal prediction market asking who would win the Grammy Awards (music), 68 of our demographically-diverse employees correctly predicted the winners in six out of eight categories, regardless of whether they listened to specific genres and artists.

More significant, though, was a recent experiment conducted by a media company client, who ran a prediction market with about 200 members of their consumer communities and roughly 50 employees and agency partners. They asked which of five 2012 products would outperform their 2011 counterparts and, roughly six weeks after asking, got the actual 2012 sales data. While we cannot share any specifics we can tell you that the overall prediction market was correct in three of their five predictions and for the two in which they weren't, there was no clear market consensus, meaning that the point spread

Table 1: Prediction Market-Based Studies and Their Parallel Alternates		
Client	**Prediction Market**	**Alternate Method**
Global food company	287 people not screened to be representative of the target market	Survey of over 3,700 people screened to be representative of target market
Global hospitality company	131 community members	Max-diff study with 3,192 respondents
Global financial services company	421 community members	Max-diff study with 611 respondents
Leading health and personal care company	165 community members	Community survey with 205 respondents
Leading insurance provider	82 people not screened to be representative of the target market	Community survey with 111 respondents

between the 2011 and 2012 versions was extremely small.

At first glance, being correct in three-out-of-five predictions may not seem especially exciting. But the three correct predictions pertained to new-product failures.

Our client's objective in this case wasn't to evaluate and test products so much as to answer two methodology questions: Did this method increase engagement for both community members and employees over a traditional survey? And did it highlight differences and similarities in how loyal consumers and employees evaluated the appeal of key new products? (The answer to both was a resounding yes.) But the weightier moral to this story is that had the prediction market been used earlier in the product life cycle (as opposed to when the product was already on store shelves), it could have been an invaluable tool for optimizing the most promising concepts while pulling the plug on the weakest ones. In either case, the risk-mitigation benefit could have been significant.

Get a truer picture

As with a stock market, the theory is that if you enable people to put their (play) money where their mouths are, to invest in their own predictions about the behavior of *others* (as opposed to what they *themselves* would be likely to buy or use) based on their confidence in the outcomes, you'll get a truer picture of what is likely to have

value in the real market. And it's true – gamification elements such as leaderboards and net-worth scores offer participants recognition, reward and status as well as an opportunity to compete. When people have a skin in the game, when they're invested in the outcome, they commit more thought and attention to the process.

Consensus Point, maker of the Huunu prediction market platform, distinguishes between prediction markets and surveys in this way:

	Surveys	Prediction Markets
Answers	Opinions	Judgment
Type of Knowledge	Explicit	Tacit
Cognitive Process	Disinterested	Invested

Asking about the behavior of others can often allow participants to bring greater knowledge and insight to bear on a question than simply asking about their own preferences or intentions. For example, as someone in my late fifties, I would never be surveyed by a manufacturer of baby products. But the fact is that I have young grandchildren and in my workplace am surrounded by new parents. Though not personally representative of the target market, I actually have knowledge of that market that I could bring to bear to a question if only I was given a chance.

Indeed, participant diversity can be a terrific asset. While a given

individual may have limited experience – only a piece of the puzzle – the collective knowledge and mutual influence that a prediction market can aggregate is what generates what James Surowiecki has famously called the wisdom of crowds. Or, to use a less trendy but no less true phrase, "The whole is often greater than the sum of its parts."

Can be valuable tools

Prediction markets comprising a diverse set of consumers can be valuable tools for companies spanning a wide range of industries and at every stage in the product or service life cycle and be used for: narrowing the new-product development funnel; concept testing; forecasting; pricing; message optimization; and promotion testing.

Of course there are some questions that employees are better equipped to answer than consumers, such as volumetric projections. And there are some categories – such as manufacturing or health sciences – where specialized knowledge that's simply not possessed by a broad consumer pool is essential.

But the biggest caution is to remember that picking winners and losers is only one small piece of the puzzle. In both incremental and breakthrough innovation, whether developing new products or promotions or tweaking existing ones, the magic lies in generating the insight that leads to a deeper understanding of consumers and where their unmet needs lie. Consumers want to be known as much as they want to be heard and will embrace any opportunity to meaningfully partner with brands.

So engage them in sharing their frustrations and dreams. Stimulate them to think creatively right along with you. Team up with them in ideating new solutions, even the ones that you'll eventually test with them. We predict that if you collaborate with your customers at every step of the way, you'll be glad you did. Ⓠ

Julie Wittes Schlack is senior vice president of innovation at Boston research firm Communispace. She can be reached at 617-316-4458 or at jwschlack@communispace.com.