<< Back to Article

WIRED MAGAZINE: 17.06

Secret of Googlenomics: Data-Fueled Recipe Brews Profitability

By Steven Levy 05.22.09

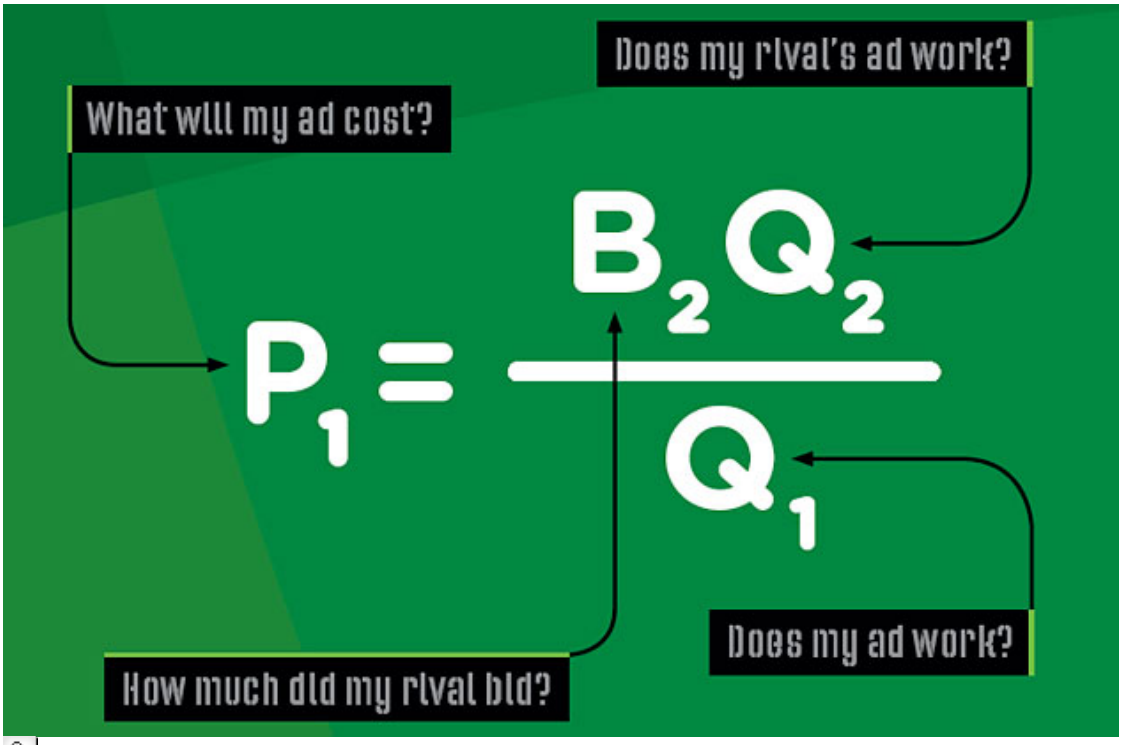

As the amount of data at Google's disposal grows, the opportunities to exploit it multiply.

THE NEW NEW ECONOMY


More Startups, Fewer Giants, Infinite Opportunity


Beyond Detroit: On the Road to Recovery, Let the Little Guys Drive


The New Socialism: Global Collectivist Society Is Coming Online

INFOGRAPHIC

Anatomy of an Auction

In the midst of financial apocalypse, the gadflies and gurus of the global marketplace are gathered at the San Francisco Hilton for the annual meeting of the American Economics Association. The mood is similar to a seismologist convention in the wake of the Big One. Yet surprisingly, one of the most popular sessions has nothing to do with toxic assets, derivatives, or unemployment curves.

"I'm going to talk about online auctions," says Hal Varian, the session's first speaker. Varian is a lanky 62-year-old professor at UC Berkeley's Haas School of Business and School of Information, but these days he's best known as Google's chief economist. This morning's crowd hasn't come for predictions about the credit market; they want to hear about Google's secret sauce.

Varian is an expert on what may be the most successful business idea in history: AdWords, Google's unique method for selling online advertising. AdWords analyzes every Google search to determine which advertisers get each of up to 11 "sponsored links" on every results page. It's the world's biggest, fastest auction, a never-ending, automated, self-service version of Tokyo's boisterous Tsukiji fish market, and it takes place, Varian says, "every time you search." He never mentions how much revenue advertising brings in. But Google is a public company, so anyone can find the number: It was $21 billion last year.

His talk quickly becomes technical. There's the difference between the Generalized Second Price auction model and the Vickrey-Clark-Groves alternative. Game theory takes a turn; so does the Nash Equilibrium. Terms involving the c-word—as in clicks—get tossed around like beach balls at a summer rock festival. Clickthrough rate. Cost per click. Supply curve of clicks. The audience is enthralled.

During the question-and-answer period, a man wearing a camel-colored corduroy blazer raises his hand. "Let me understand this," he begins, half skeptical, half unsure. "You say that an auction happens every time a search takes place? That would mean millions of times a day!"

Varian smiles. "Millions," he says, "is actually quite an understatement."

Why does Google even need a chief economist? The simplest reason is that the company is an economy unto itself. The ad auction, marinated in that special sauce, is a seething laboratory of fiduciary forensics, with customers ranging from giant multinationals to dorm-room entrepreneurs, all billed by the world's largest micropayment system.

Google depends on economic principles to hone what has become the search engine of choice for more than 60 percent of all Internet surfers, and the company uses auction theory to grease the skids of its own operations. All these calculations require an army of math geeks, algorithms of Ramanujanian complexity, and a sales force more comfortable with whiteboard markers than fairway irons.

Varian, an upbeat, avuncular presence at the Googleplex in Mountain View, California, serves as the Adam Smith of the new discipline of Googlenomics. His job is to provide a theoretical framework for Google's business practices while leading a team of quants to enforce bottom-line discipline, reining in the more propellerhead propensities of the company's dominant engineering culture.

Googlenomics actually comes in two flavors: macro and micro. The macroeconomic side involves some of the company's seemingly altruistic behavior, which often baffles observers. Why does Google give away products like its browser, its apps, and the Android operating system for mobile phones? Anything that increases Internet use ultimately enriches Google, Varian says. And since using the Web without using Google is like dining at In-N-Out without ordering a hamburger, more eyeballs on the Web lead inexorably to more ad sales for Google.

The microeconomics of Google is more complicated. Selling ads doesn't generate only profits; it also generates torrents of data about users' tastes and habits, data that Google then sifts and processes in order to predict future consumer behavior, find ways to improve its products, and sell more ads. This is the heart and soul of Googlenomics. It's a system of constant self-analysis: a data-fueled feedback loop that defines not only Google's future but the future of anyone who does business online.

When the American Economics Association meets next year, the financial crisis may still be topic A. But one of the keynote speakers has already been chosen: Googlenomist Hal Varian.

Ironically, economics was a distant focus in the first days of Google. After Larry Page and Sergey Brin founded the company in 1998, they channeled their energy into its free search product and left much of the business planning to a 22-year-old Stanford graduate named Salar Kamangar, Google's ninth employee. The early assumption was that although ads would be an important source of revenue, licensing search technology and selling servers would be just as lucrative. Page and Brin also believed that ads should be useful and welcome—not annoying intrusions. Kamangar and another early Googler, Eric Veach, set out to implement that ideal. Neither had a background in business or economics. Kamangar had been a biology major, and Veach's field of study was computer science.

Hal Varian, high priest of Googlenomics.
Photo: Joe Pugliese

Google's ads were always plain blocks of text relevant to the search query. But at first, there were two kinds. Ads at the top of the page were sold the old-fashioned way, by a crew of human beings headquartered largely in New York City. Salespeople wooed big customers over dinner, explaining what keywords meant and what the prices were. Advertisers were then billed by the number of user views, or impressions, regardless of whether anyone clicked on the ad. Down the right side were other ads that smaller businesses could buy directly online. The first of these, for live mail-order lobsters, was sold in 2000, just minutes after Google deployed a link reading SEE YOUR AD HERE.

But as the business grew, Kamangar and Veach decided to price the slots on the side of the page by means of an auction. Not an eBay-style auction that unfolds over days or minutes as bids are raised or abandoned, but a huge marketplace of virtual auctions in which sealed bids are submitted in advance and winners are determined algorithmically in fractions of a second. Google hoped that millions of small and medium companies would take part in the market, so it was essential that

the process be self-service. Advertisers bid on search terms, or keywords, but instead of bidding on the price per impression, they were bidding on a price they were willing to pay each time a user *clicked* on the ad. (The bid would be accompanied by a budget of how many clicks the advertiser was willing to pay for.) The new system was called AdWords Select, while the ads at the top of the page, with prices still set by humans, was renamed AdWords Premium.

One key innovation was that all the sidebar slots on the results page were sold off in a single auction. (Compare that to an early pioneer of auction-driven search ads, Overture, which held a separate auction for each slot.) The problem with an all-at-once auction, however, was that advertisers might be inclined to lowball their bids to avoid the sucker's trap of paying a huge amount more than the guy just below them on the page. So the Googlers decided that the winner of each auction would pay the amount (plus a penny) of the bid from the advertiser with the next-highest offer. (If Joe bids $10, Alice bids $9, and Sue bids $6, Joe gets the top slot and pays $9.01. Alice gets the next slot for $6.01, and so on.) Since competitors didn't have to worry about costly overbidding errors, the paradoxical result was that it encouraged higher bids.

"Eric Veach did the math independently," Kamangar says. "We found out along the way that second-price auctions had existed in other forms in the past and were used at one time in Treasury auctions." (Another crucial innovation had to do with ad quality, but more on that later.)

Google's homemade solution to its ad problem impressed even Paul Milgrom, the Stanford economist who is to auction theory what Letitia Baldridge is to etiquette. "I've begun to realize that Google somehow stumbled on a level of simplification in ad auctions that was not included before," he says. And applying a variation on second-price auctions wasn't just a theoretical advance. "Google immediately started getting higher prices for advertising than Overture was getting."

Google hired Varian in May 2002, a few months after implementing the auction- based version of AdWords. The offer came about when Google's then-new CEO, Eric Schmidt, ran into Varian at the Aspen Institute and they struck up a conversation about Internet issues. Schmidt was with Larry Page, who was pushing his own notions about how some of the big problems in business and science could be solved by using computation and analysis on an unprecedented scale. Varian remembers thinking, "Why did Eric bring his high-school nephew?"

Schmidt, whose father was an economist, invited Varian to spend a day or two a week at Google. On his first visit, Varian asked Schmidt what he should do. "Why don't you take a look at the ad auction?" Schmidt said.

Google had already developed the basics of AdWords, but there was still plenty of tweaking to do, and Varian was uniquely qualified to "take a look." As head of the information school at UC Berkeley and coauthor (with Carl Shapiro) of a popular book called *Information Rules: A Strategic Guide to the Network Economy*, he was already the go-to economist on ecommerce.

At the time, most online companies were still selling advertising the way it was done in the days of *Mad Men*. But Varian saw immediately that Google's ad business was less like buying traditional spots and more like computer dating. "The theory was Google as yenta—matchmaker," he says. He also realized there was another old idea underlying the new approach: A 1983 paper by Harvard economist Herman Leonard described using marketplace mechanisms to assign job candidates to slots in a corporation, or students to dorm rooms. It was called a two-sided matching market. "The mathematical structure of the Google auction," Varian says, "is the same as those two-sided matching markets."

Varian tried to understand the process better by applying game theory. "I think I was the first person to do that," he says. After just a few weeks at Google, he went back to Schmidt. "It's amazing!" Varian said. "You've managed to design an auction perfectly."

To Schmidt, who had been at Google barely a year, this was an incredible relief. "Remember, this was when the company had 200 employees and no cash," he says. "All of a sudden we realized we were in the auction business."

It wasn't long before the success of AdWords Select began to dwarf that of its sister system, the more traditional AdWords Premium. Inevitably, Veach and Kamangar argued that *all* the ad slots should be auctioned off. In search, Google had already used scale, power, and clever algorithms to change the way people accessed information. By turning over its sales process entirely to an auction-based system, the company could similarly upend the world of advertising, removing human guesswork from the equation.

The move was risky. Going ahead with the phaseout—nicknamed Premium Sunset—meant giving up campaigns that were selling for hundreds of thousands of dollars, for the unproven possibility that the auction process would generate even bigger sums. "We were going to erase a huge part of the company's revenue," says Tim Armstrong, then head of direct sales in the US. (This March, Armstrong left Google to become AOL's new chair and CEO.) "Ninety-nine percent of companies would have said, 'Hold on, don't make that change.' But we had Larry, Sergey, and Eric saying, 'Let's go for it.'"

News of the switch jacked up the Maalox consumption among Google's salespeople. Instead of selling to corporate giants, their job would now be to get them to *place bids in an auction?* "We thought it was a little half-cocked," says Jeff Levick, an early leader of the Google sales team. The young company wasn't getting rid of its sales force (though the system certainly helped Google run with far fewer salespeople than a traditional media company) but was asking them to get geekier, helping big customers shape online strategies as opposed to simply selling ad space.

Levick tells a story of visiting three big customers to inform them of the new system: "The guy in California almost threw us out of his office and told us to fuck ourselves. The guy in Chicago said, 'This is going to be the worst business move you

ever made.' But the guy in Massachusetts said, 'I trust you.'"

That client knew math, says Levick, whose secret weapon was the numbers. When the data was crunched—and Google worked hard to give clients the tools needed to run the numbers themselves—advertisers saw that the new system paid off for them, too.

AdWords was such a hit that Google went auction-crazy. The company used auctions to place ads on other Web sites (that program was dubbed AdSense). "But the really gutsy move," Varian says, "was using it in the IPO." In 2004, Google used a variation of a Dutch auction for its IPO; Brin and Page loved that the process leveled the playing field between small investors and powerful brokerage houses. And in 2008, the company couldn't resist participating in the FCC's auction to reallocate portions of the radio spectrum.

Google even uses auctions for internal operations, like allocating servers among its various business units. Since moving a product's storage and computation to a new data center is disruptive, engineers often put it off. "I suggested we run an auction similar to what the airlines do when they oversell a flight. They keep offering bigger vouchers until enough customers give up their seats," Varian says. "In our case, we offer more machines in exchange for moving to new servers. One group might do it for 50 new ones, another for 100, and another won't move unless we give them 300. So we give them to the lowest bidder—they get their extra capacity, and we get computation shifted to the new data center."

The transition to an all-auction sales model was a milestone for Google, ensuring that its entire revenue engine would run with the same computer-science fervor as its search operation. Now, when Google recruits alpha geeks, it is just as likely to have them focus on AdWords as on search or apps.

The across-the-board emphasis on engineering, mathematical formulas, and data-mining has made Google a new kind of company. But to fully understand why, you have to go back and look under AdWords' hood.

Most people think of the Google ad auction as a straightforward affair. In fact, there's a key component that few users know about and even sophisticated advertisers don't fully understand. The bids themselves are only a part of what ultimately determines the auction winners. The other major determinant is something called the quality score. This metric strives to ensure that the ads Google shows on its results page are true, high-caliber matches for what users are querying. If they aren't, the whole system suffers and Google makes less money.

Google determines quality scores by calculating multiple factors, including the relevance of the ad to the specific keyword or keywords, the quality of the landing page the ad is linked to, and, above all, the percentage of times users actually click on a given ad when it appears on a results page. (Other factors, Google won't even discuss.) There's also a penalty invoked when the ad quality is too low—in such cases, the company slaps a minimum bid on the advertiser. Google explains that this practice—reviled by many companies affected by it—protects users from being exposed to irrelevant or annoying ads that would sour people on sponsored links in general. Several lawsuits have been filed by would-be advertisers who claim that they are victims of an arbitrary process by a quasi monopoly.

You can argue about fairness, but arbitrary it ain't. To figure out the quality score, Google needs to estimate in advance how many users will click on an ad. That's very tricky, especially since we're talking about billions of auctions. But since the ad model depends on predicting clickthroughs as perfectly as possible, the company must quantify and analyze every twist and turn of the data. Susan Wojcicki, who oversees Google's advertising, refers to it as "the physics of clicks."

During Varian's second summer in Mountain View, when he was still coming in only a day or two a week, he asked a recently hired computer scientist from Stanford named Diane Tang to create the Google equivalent of the Consumer Price Index, called the Keyword Pricing Index. "Instead of a basket of goods like diapers and beer and doughnuts, we have keywords," says Tang, who is known internally as the Queen of Clicks.

The Keyword Pricing Index is a reality check. It alerts Google to any anomalous price bubbles, a sure sign that an auction isn't working properly. Categories are ranked by the cost per click that advertisers generally have to pay, weighted by distribution, and then separated into three bundles: high cap, mid cap, and low cap. "The high caps are very competitive keywords, like 'flowers' and 'hotels,'" Tang says. In the mid-cap realm you have keywords that may vary seasonally—the price to place ads alongside results for "snowboarding" skyrockets during the winter. Low caps like "Massachusetts buggy whips" are the stuff of long tails.

Tang's index is just one example of a much broader effort. As the amount of data at the company's disposal grows, the opportunities to exploit it multiply, which ends up further extending the range and scope of the Google economy. So it's utterly essential to calculate correctly the quality scores that prop up AdWords.

"The people working for me are generally econometricians—sort of a cross between statisticians and economists," says Varian, who moved to Google full-time in 2007 (he's on leave from Berkeley) and leads two teams, one of them focused on analysis.

"Google needs mathematical types that have a rich tool set for looking for signals in noise," says statistician Daryl Pregibon, who joined Google in 2003 after 23 years as a top scientist at Bell Labs and AT&T Labs. "The rough rule of thumb is one statistician for every 100 computer scientists."

Keywords and click rates are their bread and butter. "We are trying to understand the mechanisms behind the metrics," says Qing Wu, one of Varian's minions. His specialty is forecasting, so now he predicts patterns of queries based on the season, the climate, international holidays, even the time of day. "We have temperature data, weather data, and queries data, so we

can do correlation and statistical modeling," Wu says. The results all feed into Google's backend system, helping advertisers devise more-efficient campaigns.

To track and test their predictions, Wu and his colleagues use dozens of onscreen dashboards that continuously stream information, a sort of Bloomberg terminal for the Googlesphere. Wu checks obsessively to see whether reality is matching the forecasts: "With a dashboard, you can monitor the queries, the amount of money you make, how many advertisers you have, how many keywords they're bidding on, what the rate of return is for each advertiser."

Wu calls Google "the barometer of the world." Indeed, studying the clicks is like looking through a window with a panoramic view of everything. You can see the change of seasons—clicks gravitating toward skiing and heavy clothes in winter, bikinis and sunscreen in summer—and you can track who's up and down in pop culture. Most of us remember news events from television or newspapers; Googlers recall them as spikes in their graphs. "One of the big things a few years ago was the SARS epidemic," Tang says. Wu didn't even have to read the papers to know about the financial meltdown—he saw the jump in people Googling for *gold*. And since prediction and analysis are so crucial to AdWords, every bit of data, no matter how seemingly trivial, has potential value.

Since Google hired Varian, other companies, like Yahoo, have decided that they, too, must have a chief economist heading a division that scrutinizes auctions, dashboards, and econometric models to fine-tune their business plan. In 2007, Harvard economist Susan Athey was surprised to get a summons to Redmond to meet with Steve Ballmer. "That's a call you take," she says. Athey spent last year working in Microsoft's Cambridge, Massachusetts, office.

Can the rest of the world be far behind? Although Eric Schmidt doesn't think it will happen as quickly as some believe, he does think that Google-style auctions are applicable to all sorts of transactions. The solution to the glut in auto inventory? Put the entire supply of unsold cars up for bid. That'll clear out the lot. Housing, too: "People use auctions now in cases of distress, like auctioning a house when there are no buyers," Schmidt says. "But you can imagine a situation in which it was a normal and routine way of doing things."

Varian believes that a new era is dawning for what you might call the datarati—and it's all about harnessing supply and demand. "What's ubiquitous and cheap?" Varian asks. "Data." And what is scarce? The analytic ability to utilize that data. As a result, he believes that the kind of technical person who once would have wound up working for a hedge fund on Wall Street will now work at a firm whose business hinges on making smart, daring choices—decisions based on surprising results gleaned from algorithmic spelunking and executed with the confidence that comes from really doing the math.

It's a satisfying development for Varian, a guy whose career as an economist was inspired by a sci-fi novel he read in junior high. "In Isaac Asimov's first *Foundation Trilogy*, there was a character who basically constructed mathematical models of society, and I thought this was a really exciting idea. When I went to college, I looked around for that subject. It turned out to be economics." Varian is telling this story from his pied-è0-Plex, where he sometimes stays during the week to avoid driving the 40-some miles from Google headquarters to his home in the East Bay. It happens to be **the ranch-style house**, which Google now owns, where Brin and Page started the company.

There's a wild contrast between this sparsely furnished residence and what it has spawned—dozens of millionaire geeks, billions of auctions, and new ground rules for businesses in a data-driven society that is far weirder than the one Asimov envisioned nearly 60 years ago. What could be more baffling than a capitalist corporation that gives away its best services, doesn't set the prices for the ads that support it, and turns away customers because their ads don't measure up to its complex formulas? Varian, of course, knows that his employer's success is not the result of inspired craziness but of an early recognition that the Internet rewards fanatical focus on scale, speed, data analysis, and customer satisfaction. (A bit of auction theory doesn't hurt, either.) Today we have a name for those rules: Googlenomics. Learn them, or pay the price.

Senior writer Steven Levy (steven_levy@wired.com) *wrote about the Kryptos sculpture at CIA headquarters in issue 17.05.*