SILICON SIMULACRA

Post-humans of the Machine Worlds

Len Ellis

Chapter 6

Thinking Together

com/2007/07/17/social-network-aggregators/ reviews profile centralizers.

33 See Erica Naone, "Searching for Humans," *Technology Review,* August 20, 2007, www.technologyreview.com/communications/19270/; Paula J. Hane, "People Search Tools Populate the Web," *Information Today,* September 1, 2007, http://newsbreaks.infotoday.com/nbReader.asp?ArticleId=37403; Michael Arrington, "Spock's New People Engine," *TechCrunch,* April 11, 2007, www.techcrunch.com/2007/04/11/exclusive-screenshots-spocks-new-people-engine/.

34 See Michael A. Turner and Robin Varghese, *Making Sense of the Privacy Debate: A Comparative Analysis of Leading Consumer Privacy Studies* (New York: The Direct Marketing Association, 2001). The absence of control is likely one factor in consumers' "visceral reactions" to privacy that precede any balancing of benefits and risks, an observation reported by Fred H. Cate and Richard J. Varn, *The Public Record: Information Privacy and Access* (Des Moines, Iowa: Coalition for Sensible Public Records, 1999).

35 Those who argue along this line include Kenneth Gergen, *The Saturated Self: Dilemmas of Identity in Contemporary Life* (New York: Basic Books, 1991); Donna Harraway, *Simians, Cyborgs, Women: The Reinvention of Nature* (New York: Routledge, 1991), 144—181; Robert Jay Lifton, *The Protean Self: Human: Resilience in the Age of Fragmentation* (New York: Basic Books, 1993); Emily Martin, *Flexible Bodies: Tracking Immunity in American Culture from the Days of Polio to the Age of AIDS* (Boston: Beacon Press, 1994); Sherry Turkle, *Life on the Screen: Identity in the Age of the Internet*_(New York: Simon & Schuster, 1995), 180 and 256; Poster (2001), 37 and 75 and Hayles, 26—32.

CHAPTER 6

THINKING TOGETHER

Our ready-made individuality, our identity is no
more than an accidental cohesion in the flux of time.

—D. H. Lawrence, "The Crown" (1915)

The network knows what the nodes don't.

—Kim Rachmeier, Internet advisor and investor (2008)

The mission of all computing is to augment human intelligence, and the Internet holds out a particular potential in this regard. It enables many users to connect with many other users and is unlike other communications media. The telephone network connects one person with one other person; it's a *one-to-one* medium. The Internet does that, too, with e-mail and telephony. TV, radio, magazines and newspapers are *one-to-many* media; they enable one "person," the broadcaster or publisher, to connect with many persons—viewers, listeners and readers. Web sites and blogs also enable the one to connect with the many. Alone among media the Internet enables *many-to-many* communication and holds out the prospect that we can learn how to think better together.

This has already brought power to the people. Even before the Web came along, the Internet was famous for enabling people to find each other and create conversational communities on bulletin board systems, in chat rooms, in discussion forums, even via the e-mail lists known as LISTSERVs.[1]

Today, with the advent of Web 2.0 applications, which enable users to generate and share content, forming groups in the grassroots has become easy, quick and cheap. One result is a proliferation of groups that have narrow social value, low commercial value or both. Economically impossible under the old cost structure of communications, such groups flourish now that communications power is in the hands of the people.[2]

A grander understanding of this empowerment sees the surface of our planet wrapped in a cloak of trunk lines, radio waves and electronic circuits, humming with the communication of humankind, and proposes that something positive will emerge. The humanistic version anticipates a future in which meaning will emerge from an evolutionary ebb and flow of discourse within parameters and a discourse-based process of always-provisional outputs will augment not merely our intelligence but more profoundly our humanity.[3] The post-humanist version foresees the emergence of a super-organism—a networked intelligence, sensate, self-governing and capable of learning—displaying what *Wired* editor Kevin Kelly calls "intelligence without reason."[4] Both humanist and post-humanist visions center on the concept of emergence, and to understand the promise of many-to-many connectivity requires that term be defined.

Emergence occurs when the interactions of simple entities, acting locally and autonomously, turn into complex entities that act globally and holistically. In the natural world, the flocking of birds and the schooling of fish are examples of how many simple entities become one complex entity. The flock and the school emerge from the behaviors of birds and fish, but the complex entities have capabilities and functions of their own, greater than and different from what individual birds and fish have. Bees and ants are even more impressive. They build hives and colonies that have permanence but they don't have master plans. Those complex entities emerge from the interactions of bees and ants acting locally and autonomously; the emergent hives and colonies function

and respond globally as wholes in ways the bees and ants cannot. In short, emergence occurs when a whole emerges from but is different from the sum of the parts.

The World Wide Web also displays emergent properties. There's no master plan here either, but the Web does take on a shape that emerges from the links its users insert to connect one node to another. According to Internet impresario Tim O'Reilly:

> Hyperlinking is the foundation of the web. As users add new content, and new sites, it is bound in to the structure of the web by other users discovering the content and linking to it. Much as synapses form in the brain, with associations becoming stronger through repetition or intensity, the web of connections grows organically as an output of the collective activity of all web users.[5]

This ongoing activity is self-reinforcing. Some nodes get more inbound links than others do. Those better-linked nodes are more findable, attracting more visits, from which those nodes will get even more inbound links. The search engines reflect these decisions. They count the number of inbound links as a measure of a node's "authority" and rank nodes with more links higher in their lists of results. And they reinforce these decisions: the higher ranked nodes get even more visits, even more inbound links and even higher ranks in the search results. Some people disparage this bottom-up "authority" structure and suggest it's better defined as "popularity according to the Web's most prolific linkers." Judgments aside, a functional and responsive global order emerges within the Web from the independent linking of its autonomous users.[6]

In other words, cyberspace seems a particularly hospitable context for attempts to generate wisdom from crowds. In a network where each autonomous entity can connect to every other and where many-to-many connections can be enabled quickly, easily and at near-zero cost, it's only logical to ask: can we configure modes of interaction among these nodes so that the outputs of their interactions are greater than and different from the sum of their inputs? The three applications examined here take on a modest version of this quest. Specifically, crowdsourcing, recommender systems and prediction markets enable the many to produce and share content. As always, the specific mechanics of each application determines whether our thinking together produces intelligence, groupthink or something else. Before addressing those specifics, however, certain common features can and should be delineated.

In general, collective intelligence can be defined as a process in which many "agents" work, each in their own time and at their own pace, on parts of a larger project without a plan or supervision and produce knowledge. The common attributes are:

Purpose	The output of the collective endeavor is specific and understood.
Divisibility	The output is divisible such that agents can produce inputs toward its completion.
Parallelism	Agents work autonomously and concurrently.
Equality	Agents start off with the same privileges and responsibilities.
Hierarchy	In some configurations, agents can earn more privileges and responsibilities.
Transparency	The machinery of decision making and its use are visible.
Diversity	Difference in the agents' skills, expertise or perspectives is required if the output requires balance, innovation or reliability.
Systemic Effect	The output is attributable to the system that connects the agents at a higher level than the agents themselves and is not achievable by individual agents even with infinite resources.

The last attribute is the difference between aggregating individual intelligences and catalyzing collective intelligence. The following analysis suggests where this does and does not occur.

Crowdsourcing

Crowdsourcing refers to the use of Web 2.0 applications to enable a large group of people outside an organization to contribute to the completion of a work task. It's not a new production model. The *Oxford English Dictionary* was compiled that way back in the 19th century. What's new are the scale and low cost of online connectivity, and three types of crowdsourcing can already be distinguished: worksourcing, ideasourcing and expertsourcing.

Worksourcing involves tapping a crowd of people to complete repetitive tasks that require human intelligence; otherwise some machine would be performing the tasks. Two volunteer projects, for example, leverage the fact that humans are better than machines at pattern recognition. Four research universities set up Galaxy Zoo, a Web site where volunteers are classifying over a million galaxies as elliptical or spiral and by the directions of their spin. Similarly, Stardust@Home, hosted by The Planetary Society, invites volunteers to find interstellar dust particles from among some 700,000 3-D images of the Stardust spacecraft's collector plate.[7] Not to be outdone, the social sciences are trying to access what's called local or situational knowledge—information specific to a time and place. The Library of Congress recently posted thousands of its photographs to photo-hosting site Flickr and asked the public to answer: *What is this? Who is this? When was it taken?* The National Archives launched a similar project with online versions of its documents.

In the for-profit sector, Amazon.com has a worksourcing initiative called the Mechanical Turk, honoring the

18th century chess-playing automaton. It pays a few pennies for trivial tasks that only human can do, like identifying performers on music CDs and choosing appropriate categories for certain products. Not all worksourcing tasks are so trivial, however, and the potential consequences of worksourcing in the marketplace are not entirely benign.

Some worksourcing sites operate reverse auctions in which freelancers compete for projects by underbidding each other. Such bid-for-work markets have been set up in such categories as graphic design and copywriting, sales and marketing, financial and managerial services and computer programming. A few sites offer "on spec" assignments, which means doing the work without any guarantee of payment just to win the assignment and, if selected, get paid. The "winner" gets a nominal fee, all the others get zip and the company gets the work at a fraction of what a service firm might charge. The logical endpoint has already been reached. IStockphoto.com quickly became the third-largest seller in the $2 billion stock photography business by selling images at rock-bottom prices; it can do so because more than 90 percent of its photographers are amateurs who sign away their royalty rights. At the end of the day, piecework, whether by volunteers on science projects or low-bidding online freelancers, is neither collective nor intelligent.

A second type of crowdsourcing is ideasourcing, organized at an "ideagora"—a Web site where companies post problems, users propose solutions and the one who solves the problem wins a prize. Despite the fancy labels, the structure is an inducement prize contest and also quite old. The most famous of such contests is the Longitude Prize, announced by the British government in 1714; the most famous winner of such a contest is aviator Charles Lindberg, who claimed the Orteig Prize in 1927 for the first solo flight across the Atlantic. The structure still works. The Ansara X Prize, $10 million for the first nongovernmental organization to launch a reusable manned spacecraft into space twice within two weeks, was awarded in 2004. In 2009, a team of mathematicians won a $1 million prize from DVD rental company Netflix for improving its recommender system by 10 percent. In another recent contest, Canadian mining company Goldcorp made available online some 400 megabytes of geological data about its Ontario property and offered $575,000 to anyone whose data analysis could pinpoint where gold could be found. A small consultancy in Australia won. So did Goldcorp: it collected from all contest entries 110 targets, of which over 80 percent proved productive, yielding 8 million ounces of gold worth more than $3 billion.

What's new about ideagoras is not the inducement prize contest but the cost efficiency of offering multiple smaller challenges and, more important, the ability to engage diverse perspectives on the same problem. At InnoCentive.com, an ideagora spun off from drug company Eli Lilly, *seekers* such as Boeing, Dow, DuPont and Novartis post specific technical problems, or *challenges.* The site's *solvers* are more than 100,000 scientists from 175 countries, in such fields as chemistry, biochemistry, biology and material sciences. Their solutions have included a compound for skin tanning, a method for preventing snack chip breakage, a mini-extruder design for brick making and many others. InnoCentive collects a fee from the seeker for posting its challenge and, if solved, a finder's fee, roughly 40 percent of the cash prize, usually between $10,000 and $25,000.

InnoCentive's high solve rate is due largely to the fact that the solvers come from a very broad range of disciplines, far more diverse than the disciplines inside the seeker company. For example, Colgate-Palmolive's internal R&D team, unable to find a more efficient way to fill toothpaste tubes, posted that challenge. An electrical engineer won; he proposed putting a positive charge on fluoride powder and then grounding the tube, collecting $25,000 for a few hours work. As Harvard Business School professor Karim R. Lahkani, who studied InnoCentive, explained, "The further the problem was from the solver's expertise, the more likely they were to solve it," often by applying specialized knowledge or instruments developed for other purposes.[8] This leverages a well-known aspect of human intelligence: applying the perspective of one discipline to a problem in another always sparks previously unasked questions and, sometimes, good answers.[9] The inducement prize contest, whether offering one challenge or many, online or offline, does not generate collective intelligence, however. On the contrary, it's designed to attract the uniquely qualified mind to step forward out from the crowd.

A third type of crowdsourcing, expertsourcing, aims to harvest subject matter expertise. Commercial sites like Yahoo Answers, Ask Metafilter, AllExperts (née Expert Central) and Windows Live QnA (beta) enable users to ask and answer questions but are small successes at best. The shining example is the nonprofit Wikipedia, the online encyclopedia created by an all-volunteer army of contributors. They build it using a collaborative authoring application called a wiki. This software enables multiple users to make edits to the same text, to store and access all versions with all changes and to restore a previous version. On Wikipedia anyone can access these capabilities, that is, anyone can initiate and write an article and can add to, edit, rewrite entirely or restore articles without supervision or approval.

Wikipedia's growth has been explosive. From 31 articles in January 2001, it amassed more than 17,000 within the first 12 months and today has more than 3 million articles in English and 14 million articles in more than 260 languages. It's free and popular; in 2009 it was the 5th most visited site on the Web. As for quality, it's spotty but still impressive, given how the content is sourced. A comparison with Encyclopedia Britannica conducted by *Nature* magazine found 162 factual errors, omissions or misleading statements in Wikipedia and 123 in the Britannica. The magazine concluded that Wikipedia came close in terms of accuracy, to which the Britannica folks retorted that 30 percent more accurate was not insignificant. In theory, Wikipedians will fix these flaws in the normal course of using the site; indeed, they fixed all the

errors *Nature* found within just a few days. It's also dispiriting but not surprising that the entry for TV's cartoon series *The Flintstones* is twice as long as that for Thomas Jefferson. Time will tell how the site's content evolves, but an examination of its social production—the how of its user-generated content—suggests that Wikipedia is doing pretty much everything right.

The purpose is specific. It's not task- or problem-specific like worksourcing or ideasourcing; rather, there's a normative consensus about the output. All contributors are on the same page about what an encyclopedia entry *should* contain: one or more definitions of the topic, some historical account of origins and development, some assessment of best-known impacts in one or more domains and, if there's controversy, summaries of the different positions and their chief advocates.

If an article looks anything like that, others can improve it. A contributor can correct spelling, fix punctuation, add a citation, rewrite a sentence, flesh out a paragraph or write an entire article. The output is divisible and the inputs are scalable.

All contributors start off equal, but as in all volunteer endeavors, some do more than others. Indeed, the project's success has come to rely on a self-selecting cadre of highly active users who contribute more articles of their own, who edit more articles of others and who adopt and tend to a set of pages, eliminating vandalism and deciding on corrections.

Diversity can be discerned as an article evolves through successive iterations, but it's manifest in the pages behind each article. That's where contributors discuss and sometimes argue the merits of proposed changes to content. What's visible on the article's main page is the content on which they agree, at least at the moment.

When diversity turns into disputes, Wikipedia has a formal and complex system for collective adjudication of these matters that's transparent to all. If that doesn't work, the site's administrators can step in with a decision. They have even shut down public editing of articles on a handful of topics, such as the Iraq War and abortion, because they are so often vandalized by zealots.

Whether Wikipedia has a systemic effect is arguable. Wiki software enables multiple users to write and revise the same content, but largely because it enables each user to contribute independently of the others. The software doesn't connect contributors one to another in any novel way, and successive versioning, while yielding improvement over time, is not greater than or different from the sum of its inputs. Rather, multi-authored versioned content *is* the sum of its inputs, and in theory individual agents with infinite resources could produce equivalent output.[10]

Rather, Wikipedia's innovation is a set of attitudes about knowledge, specifically, that it should be always provisional and universally open. There is no last word. No one has final say. Of course, the creation of knowledge has always been an ongoing social process, and knowledge always a social output. Whether Wikipedia's decision to hold knowledge permanently open to everyone changes anything is yet to be determined.

Ratings, Rankings and Recommendations

Rating, ranking and filtering systems are all over the Web, and all have the same goal: to aid in selecting items from a set of items by distinguishing some items from others. Among these systems two types generate collective intelligence—a new type of higher-order content that is different from the sum of the inputs, that is useful and reliable and that can be produced in no other way. The particulars of these accomplishments are most easily understood in the context of the different benefits that each system delivers.

Rating systems invite users to score the items of a set, whether the system is offline, like Zagat's restaurant guides or, online, such as those offered at Amazon, Netflix and Trip Advisor, As aids to decision making, they have two weaknesses. First, the raters rate each item independently of,

not relative to, others. So, two similar items, such as a French cookbook by Julia Child and another by Jacques Pépin, could get the same score. Indeed, two dissimilar items, such as a French cookbook and a diet cookbook, could also get the same score. In short, items scored in isolation don't necessarily distinguish one item from another.

Second, the distributions of item scores can also be hard to make useful for decision making. The illustration shows four distributions of ratings, using a five-point scale, with 1 being the lowest score and 5 being the highest. The pattern of each is easy to discern, but only two are easy to put to use. Distribution 1 shows that a lot of raters like the item and would encourage a user to select it. Distribution 2 shows that a lot of raters dislike the item and would discourage a user from selecting it. Distribution 3 is also easy to interpret—raters are polarized about the item; they either love it or hate it—but that isn't especially helpful to a user in deciding whether or not to select it. Similarly, Distribution 4 looks like the normal bell curve, which doesn't help much in deciding about the item, either. In short, rating systems, whether they display scores or distributions of scores, consider each item in isolation and offer little help in making a decision to select this item rather than that item.

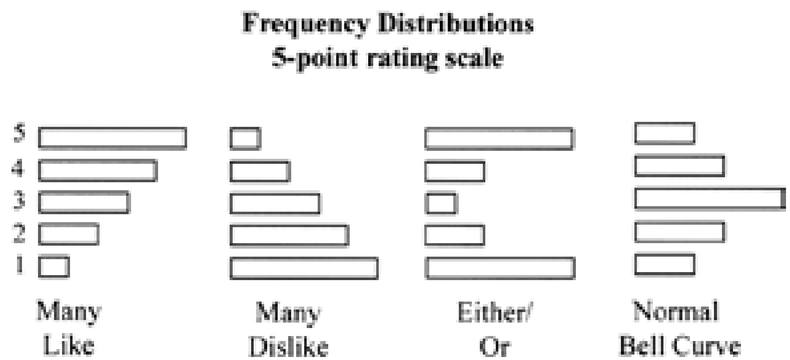

Measuring items relative to each other is the job of ranking systems, of which the biggest and best known by far is

Google. All search engines do this job. They enable a user to enter a word or phrase into a dialogue box, submit that to the engine and get back a list of Web pages related to that word or phrase. The challenge for the engines is not finding Web pages that are relevant to the word or phrase but figuring out an order in which to present those pages. That means ranking each relevant Web page relative to every other relevant Web page on the strengths of its correlations with the word or phrase that's being searched.

Google's founding innovation was a way to rank Web pages relative to each other that was unique to and meaningful on the Web. The PageRank algorithm, named for Google cofounder Larry Page, counts a link from one Web page to another as a "preferential attachment," as an implicit endorsement from the source page and a unit of authority for the target page. Web pages with more inbound links have more authority and get ranked higher than other pages in the list of results the engine presents to the user. This rank order is self-reinforcing. In the offline world, the rich get richer. In the online world, the higher ranked Web pages get more visits and from that more links, so the well-linked pages get even more well-linked.

Of course, search engines do more than count inbound links to determine the order of the Web pages they present. They apply text-mining to the content of each Web page to identify its keywords, scan and integrate the metatags (more keywords) attached to each Web page and factor in which links get clicked for different search terms, among many other techniques. The goal is to measure the relevance of each Web page to the search term and thereby rank the relevant Web pages into an order.

Although relative ranking is more helpful than isolated ratings, neither system can tell the user which item is better for the user. The results of both are nonspecific, applying to everyone in general and no one in particular, and there's validity in the criticism of both ratings and rankings as popularity contests. Tech Memorandum, for example, is a Web

site that aggregates the headlines of technology-related news stories. Like the search engines, it determines the order in which to present these headlines by tracking what links people click, what links they insert and so on. It then processes those choices to determine which headlines to show in what order on its home page. The home page changes every five minutes as the site processes new data. The result of this link-based popularity contest is predictable: the home page tends to publish what its readers already know.

While ratings and rankings systems present nonspecific results intended for everyone, recommender systems present results that suit the individual user. They proactively put in front of the user those items from the set that have not yet been seen by the user but in which the user is likely to be interested.[11] While search engines help *all* users *find* items they're looking for, recommender systems help *individual* users *discover* items they don't yet know about, would not likely have found on their own and will probably like. In short, recommender systems generate "relevant discovery."

There are two types of recommender systems. A content-based system recommends items to the user that are similar to items the user preferred or selected in the past. It has no collective dimension and is ignored here. A recommender system based on collaborative filtering puts items in front of the user that other users with similar tastes preferred or selected in the past. This approach is collective; it harvests inputs from the many—some, but not all, others—to predict the utility of an item to an individual user.

Amazon.com's "Customers Who Bought This Item Also Bought"function is the best-known collaborative filter, but many sites offering cultural products, such as Netflix (movies) and LastFM (music), use this approach because the number and variety of items and the users' preferences regarding them are vast.[12] Anyone who's used any of these sites more than a few times knows that relevant discovery is a genuine informational benefit. As a system that meets all the criteria for collective intelligence, it deserves a close look.

Collaborative filtering can be active or passive, depending on how user inputs are collected. Active collaborative filtering solicits and collects *expressed* inputs. It asks a user to *rate* an item on a sliding scale, *rank* a set of items from most favorite to least favorite or *list* those items he likes. The big advantage is that the collected data express actual cognitive activity; in short, the score is an actual assessment. The big downsides are that the data is limited to scores on items and that one can collect only so much expressed data before respondent fatigue sets in.[13]

Passive collaborative filtering solves both of those problems by gathering *behavioral* data. Instead of collecting a score for an item, these systems collect data on a broad range of item-specific user behaviors—what items a user views, saves, downloads, prints, shares, tags, comments on, links to and buys. Second, passive systems collect data from almost everyone. Typically, their default setting for data collection is opt-in; unless users actively opt out, their behaviors are recorded and put in the pool.

Most users do not opt out. Click-throughs are the currency of recommender systems. By allowing the site to capture and compare preference-revealing clicks, each user improves his own filter, and everyone else's, and gets increasingly better recommendations in the ordinary course of using the site. Moreover, there's a network effect: the more people who use the system, the better it gets. As the breadth and depth of the data increase, the predicted utility of a specific item to a specific user comes closer to the actual utility of that item to that user. In plain English, more recommendations (predicted items) become sales. That's why Netflix offered $1 million as an inducement to anyone who could improve its recommender system by 10 percent. Accurate recommendations drive sales because they deliver value—unknown but relevant items—to the user.

Of course, passive collaborative filtering is not perfect. Because filtering requires there be a flow to filter, two pump-priming problems can occur but typically don't last long. The "first rater" problem refers to a new *item,* which has not yet received any ratings from any users, cannot be correlated with rated items and is therefore never recommended. The "cold start" problem refers to a new *user* who has not yet rated any items, cannot be correlated with other users who have rated items and therefore never receives personalized recommendations. There's also a challenge in interpreting behaviors, such as items viewed, downloaded, shared and so forth as evaluations. Most sites are careful and sophisticated in making such imputations. Some enable users to rate the items the system recommends, thereby adding expressed judgments (active filtering) to behavioral data (passive filtering). Amazon, for example, invites users to fix each recommendation by selecting "I own it" or "Not interested" as well as to fix the prior behavior that prompted the recommendation by selecting "This was a gift" or "Don't Use for Recommendations." Amazon then removes these items from the user's filter, improving the accuracy of its recommendations.

Among the rating, ranking and recommender systems, collaborative filtering meets the criteria of collective intelligence. Its machinery processes inputs from the many–some but not all others—into discoveries for the one, and relevant discovery—items previously unknown but likely to interest the user—is substantively new knowledge that could not be created in other ways. It just needs pointing out that, as with other Web 2.0 applications, getting value from collaborative filtering requires a willingness to be visible to the machine. Our transparency provides the fuel from which the machine generates value.

Prediction Markets

If Big Brother had a little brother, he'd be retired U.S. Navy Vice Admiral John M. Poindexter. In the wake of the 9/11 terrorist attacks, Poindexter was appointed director of the Information Awareness Office (IAO), the data-mining unit of the Pentagon's R&D unit, the Defense Advanced

Research Projects Agency (DARPA). His first proposal, the Total Information Awareness plan, would have had the federal government scouring the country's databases for every consumer's credit card purchases, car rentals, library books, airline ticket purchases, hotel stays, telephone bills and the like, ultimately in real time, looking for patterns that might indicate terrorist planning or activity. The idea got nixed on privacy grounds.

His next idea was an online market in terrorism futures, and it cost him his job. Named the Policy Analysis Market (PAM), this e-market would have enabled traders to bet and win money by correctly predicting future terrorist-related events in the Middle East, such as assassinations, coups d'etat, bombings and missile attacks. Traders would express their predictions about such events by buying and selling futures contracts, just like commodities traders buy and sell contracts on pork bellies, coffee and other futures. Those who thought a politician would be assassinated or a government toppled by a certain date would buy the contract on that future and, if the event occurred before that date, would make money. Those who thought the murder or coup would not occur would sell the contract and, if the event did not occur, would make money. The Economist Intelligence Unit and Net Exchange, a company that makes software for trading financial derivatives, helped design the nuts and bolts.

Politicians elbowed their way to the nearest bully pulpit and editorial writers sharpened their pencils to denounce PAM as "grotesque," "ridiculous," "morally wrong," "bizarre," "ghoulish" and a "macabre bazaar of death and destruction." Some media, including both the liberal *New York Times* (July 31, 2003) and the conservative *Wall Street Journal* (August 1, 2003), defended PAM, but the outrage was too great. The market never launched and Poindexter resigned soon thereafter.[14]

This public relations disaster was unfortunate because prediction markets have an excellent track record. The oldest is the Iowa Electronic Markets. Organized by the

University of Iowa under special clearance from the Commodity Futures Trading Commission and operating continuously since 1988, these public markets have consistently proven more accurate than all traditional polling methods in not only predicting but also approximating the outcomes of U.S. and foreign elections, primaries and other political events. The Hollywood Stock Exchange, a prediction marketplace for movie fans, does equally well at both. It not only outperforms the studios' traditional methods for forecasting opening weekend box office receipts but also comes uncannily close to the actual receipts. This market does just as well in predicting the Oscar awards, not just the winners but the 40 individuals the industry selects as nominees across eight different categories. In short, these markets can not only predict but also approximate outcomes that are close to observed realities.

The U.S. business community is actively exploring the potential uses of this predictive machinery. According to various news reports, conference papers and prediction market software vendors, prediction markets are or have been used at Abbott Labs, Arcelor Mittal, Best Buy, Corning, Edmunds, Eli Lilly, General Electric, Google, Hewlett-Packard, Intel, Masterfoods, Microsoft, Motorola, Qualcomm, Pfizer, Renault, Siemens and Starwood Hotels. Their interest and activity is driven by two factors: one general, the other specific.

The general driver is a shift in the practice of knowledge management. Traditionally, knowledge managers focused on archiving, protecting and accessing their companies' patents, trademarks, business plans, research, customer lists and other intellectual property. But there's another type of business knowledge—the know-how in employees' heads that emerges from, exists within and evolves among interacting employees trying to get the day's work done—and knowledge management today tries to capture this *situated* cognition. Askme.com, for example, uses a version of expertsourcing to enable companies to locate knowledge within their organizations, but most companies rely on "social network analysis." This practice involves surveilling and mapping day-to-day contacts among employees to identify who talks with whom, when and about what; who has expertise, who has influence and who has connections and who depends on whom to solve which problems in order to identify and locate existing situated cognition.[15] Prediction markets for employees also tap situated cognition, but the purpose is to generate new and higher order intelligence.

That's the specific driver behind businesses experimenting with prediction markets; they deliver reliable intelligence that businesses need. Internal markets for sales forecasting have been widely implemented with great success. Anonymous employees playing for nominal prizes have proved consistently more accurate than surveys of both salespersons and customers for predicting sales. In absolute terms, the prediction markets' results come close to marketplace's actual results. Sales forecasting involves a simple outcome, predicting more or less of a single continuous variable with historical data available to traders, and some companies are exploring whether similar predictive accuracy can be achieved for other types of outcomes.[16]

The results have not been disclosed but here's what's been tried. A top-10 pharmaceutical company has deployed markets to predict the dispositions of regulatory issues, the relative performance of same-category drugs in treating illnesses, changes in drug prices and healthcare benefits, the relative probabilities of different R&D projects coming in on time and on budget and the number of New Molecular Entities that will be approved by the Food and Drug Administration. Since 2008, Google has set up internal markets that offered contracts on some 275 outcomes on sales and other company performance variables, as well as on such external events as mergers and acquisitions. General Electric has used prediction markets to generate new business ideas; Starwood Hotels, to select marketing campaigns.[17] Start-up company

Inkling Markets enables a company's employees to generate probabilities on business risks, while MIT's Sloan School of Business has set up markets for trading product attributes in such categories as automobiles and computers and service features in such categories as ski resorts and hotels. These markets stabilize quickly, often in less than 20 minutes, revealing the relative value of competing sets of product attributes or service features.[18] Using older methods, that task would take a lot of money and many months to complete.

Prediction markets will not, of course, displace business' traditional methods, but any organization that can frame its informational needs as predictions can set up markets where the trading patterns of experts, frontline personnel or customers will generate reliable probabilities on outcomes. Such markets are especially useful in situations where information or knowledge is distributed among many, is hard to gather or is difficult to verbalize, where new information flows in continuously and updates are required and where there's little relevant or reliable historical data about the outcome. Such conditions are widespread across the business landscape and since the Internet makes setting up such markets easy, quick and inexpensive, these markets are likely to become even more widespread. They transform the incomplete information and diverse judgments of individuals about the probabilities of future events into collective assessments that accurately correlate with actuality.

How the public might benefit from such prediction machinery is not clear, and other than the university-based Iowa Electronic Markets, there are only three prediction markets in which ordinary folks can participate. As mentioned earlier, the Hollywood Stock Exchange enables movie fans to trade stocks and bonds in movie stars and movies. The owner, financial services firm Cantor Fitzgerald, uses the data in servicing its entertainment industry clients. The MIT magazine *Technology Review* hosts markets in which readers predict such tech industry milestones as the price per kilowatt/hour of solar energy and the percentage of U.S. households using the Internet for telephone service. The purpose is to demonstrate to advertisers that its readers are leading edge. Consensus Point, a company that designs and hosts prediction markets for private clients, operates the Foresight Exchange. This public site offers markets on such outcomes as the year-end price of gasoline, a major earthquake occurring on the U.S. West Coast and the end of celibacy for Roman Catholic priests; marketing is the purpose here as well, specifically to showcase the company's software. Various public sites enable users to make predictions on a wide variety of events in popular culture, such as sports competitions, awards shows and celebrity shenanigans, hoping to attract and sell eyeballs to advertisers. But they are not markets. Participants do not interact. They predict independently and the sites tally their predictions like a survey—this percentage predicts this outcome, that percentage predicts the other.

In contrast, markets are based on the interaction among participants, an interaction that turns divergent individual judgments about an outcome into a single assessment on which all agree, a price. The outcome they are predicting is always "fuzzy," meaning it has multiple causal variables that are only partially known, such as a company's future profits. The individuals have always incomplete, sometimes biased, always rational but nonetheless subjective judgments about the fuzzy outcome. Trading is their conversation. To trade, each must self-quantify his judgment about that outcome relative to the trading price, a combination of all others' similarly self-quantified judgments. Through trading, the diversity of individual judgments is unified into a higher order assessment upon which all agree: the trading price is *their* consensus from one moment to the next. In general, humans excel at solving ill-defined problems—those that have complex goals, multiple solutions, or a changing nature—through the application of knowledge and experience. Markets transform what each of us can do into something many of us can do even better together.

Futures markets in particular are organized so that the trading price reflects the current consensus about whether a specific outcome will occur on a settlement date. For example, sales of Product A will reach X units by the end of the month.[19] Traders buy and sell contracts on the outcome. The contract's price always starts at 50. Its final price is always either 100 if the outcome occurs or 0 if it does not. In between it varies. If traders think the outcome is increasingly likely to occur, the price goes up; if they think the outcome is decreasingly likely, the price goes down. At any moment the price expresses traders' collective consensus on the outcome's probability. If the contract "sales of Product A will reach X units" is trading at 38, the crowd at that time thinks there's a 38 percent chance that outcome will occur.[20]

Prediction markets can go disastrously wrong in myriad ways and have to be set up just right to generate reliable probabilities. One element that's essential is diversity among market participants.[21] Diversity is the raw material; the market creates value by synthesizing diversity into a higher order intelligence on which all agree: the probability of the outcome according to them all. That's why DARPA's market in terrorism futures intended to invite not only government intelligence analysts but also academics, business executives, journalists, economists and other experts. Expertise ensured that the information of each would be relevant and, because it was narrowly specialized, diverse by definition. For the same reason, when companies set up internal markets for sales forecasting, they typically restrict participation to salespeople but make sure to include old hands and new hires who are working different types of trade zones, geographic territories and target markets. Anonymity is usually added as well to escape organizational politics, peer influences and personal careerism that can impinge upon being honest with oneself in assessing outcomes and in openly expressing assessments that may be different from others.

What's new about prediction markets is not the mechanism. For most of the 20th century, commodities exchanges have offered contracts on futures.[22] Rather, the innovation is a small but transformative change: the prediction markets invert the relationship between money and information. In financial markets, the purpose of trading is to make money; the trading price is an informational byproduct. In prediction markets, making money is not the purpose. On the contrary, the amount of money one can win for correct predictions is always limited to nominal amounts, just enough to ensure the quality of the trading price. That is, having a stake in being right motivates traders to ignore what they may want to occur (preferences) and focus on what they think actually will occur (probabilities).[23] In prediction markets, money is a means; the end is accurate probabilities from the crowd.

Because the Internet enables many-to-many connectivity, its potential for creating collective intelligence is among its most important promises. Although many applications claim to do so, they don't hold up to scrutiny. Using the crowd as a pool of individual intelligences has nothing collective about it, no matter if it's volunteers picking out space dust, freelancers underbidding each other for project work or scientists solving technical problems for prizes. To be sure, these applications have benefits, but they don't create collective intelligence. The Web's rating and ranking systems count up individual-level data, the expressed preferences of scores and the preferential attachments of links, but like their offline analogs, these aggregations don't emerge from any systemic effect in the interaction among raters or linkers. Finally, among the filter-based recommender systems, the content-based method is based solely on the prior behavior of the individual and doesn't qualify.

Three applications do qualify. The first, expertsourcing, is borderline, however. Successive editing of the same text does yield improvement, but each iteration is the work of an individual. The accretive result is better than the sum of its inputs but not by much, and it's not different in kind from its inputs. Collaborative filters and predictive markets qualify because they connect their agents in ways that generate knowledge at a level higher than their agents' aggregated inputs. The former connects us analytically through algorithms that transform the behaviors of many into recommendations for one. The latter connect us as traders, and our trading transforms our many divergent judgments into unified consensual assessments.

Even claims that relevant discovery and reliable probabilities are collective intelligence must be carefully expressed. Their mechanisms are not new. The algorithms of collaborative filtering are advanced versions of the predictive modeling practiced by data-based marketing, operating in near real time on a lot more and a lot better data. As for futures markets, they first became famous in 1636 when Tulip mania swept over Holland, England and France. Using the Internet makes implementation faster, better and cheaper but not different.

Rather, what's new is a set of attitudes about knowledge that are applied to implementation. The innovation at Wikipedia is not sequential editing or the wiki software but the decisions that all articles will be held open and open to all. The always-provisional text and universal access articulate without words the views that there is no last word and no final authority on any topic. The innovation in collaborative filters is not the mathematics but the decision to treat all items as equal, ignoring the totalizing hierarchal worldviews traditionally used to organize knowledge into a single comprehensive system. The innovation at prediction markets is not the mechanism but the decision to invert the relationship between money and information, subordinating the former to the pursuit of the latter. In other words, the ability to generate wisdom from these crowds depends on seeing wisdom as a bottom-up social product rather than as top-down content from authorities.

Of course, there are dangers on this path; many have already been identified. We rely on search engines to tell us what's important, even though bloggers are the Web's most prolific and timely linkers and have a disproportionate role in shaping search engine results. We resort to a free, all-volunteer online encyclopedia for quick research, even though anyone who actually knows a topic also knows that these articles just pass muster. We publish just about anything online because the near-zero cost has displaced the substantive question Why publish ? with the rhetorical question Why not publish? We cut and paste our way through the Web's vast resources, as if everything were miscellaneous, even though we know that nuggets of information, taken out of their contexts, have lost most of their meaning and that our power browsing, a mile wide and an inch deep, is not at all the same as analysis and inference. Some even say we are losing our capacity for the concentrated attention that reading an actual book requires as well as our taste for the interior pleasures of self-reflection, contemplation and imagination. All these dangers are quite real.

Similarly, all knowledge always was and will be socially produced; that's not new. But we have entered an age in which tools are readily available that enable us to engineer that social production, and we certainly are figuring out how to process our thoughts with machines to yield relevant discoveries, reliable probabilities and perhaps other cognitive outputs in the future. Whether our engineering of collective intelligence is put to better uses that selecting CDs or predicting sales, time will tell. As long as the barriers are low, experimentation is likely to broaden. What's already certain is that, going forward, more of what each of us knows will be anchored directly and explicitly in what all of us know, presented with the authority of the crowd and a freshness date.

Notes

1 See Howard Rheingold, *The Virtual Community: Homesteading on the Electronic Frontier* (Reading: Addison-Wesley, 1993).

2 See Clay Skirky, *Here Comes Everybody: The Power of Organizing without Organizations* (New York: Penguin, 2008). Shirky rightly emphasizes that mobile devices that can transmit to and from the Web are central in enabling groups to coordinate and collaborate as well as share content.

3 Pierre Levy, *Collective Intelligence: Mankind's Emerging World in Cyberspace (*New York: Plenum, 1997).

4 Kevin Kelly, "Evidence of a SuperOrganism" The Technium, October 24, 2008, www.kk.org/thetechnium/archives/ 2008/10/evidence_of_a_g.php, which reprises themes from his *Out of Control: The New Biology of Machines, Social Systems and the Economic World* (Reading: Addison-Wesley, 1994).

5 Tim O'Reilly, "What is Web 2.0?" O'Reilly, September 30, 2005, http://oreilly.com/web2/archive/what-is-web-20.html.

6 The topography of networks using one-way links and fixed nodes is explained in this book's introduction.

7 Some projects look like worksourcing but aren't. The University of Washington, for example, created an online game called *Fold It* in which players compete in folding proteins, but the goal is to gain insights into the human brain's abilities at pattern recognition and apply the insights to improving protein-generating software. In short, *Fold It* is research. Worksourcing must also be distinguished from distributed computing projects in which idle computers are accessed to process data. The best known is the search for extraterrestrial intelligence pursued by SETI@home, run by the University of California, Berkeley. Volunteers download the SETI@home software; when their computers are idle, the software automatically downloads a "work unit" of radio telescope data, performs a "signal analysis" and returns the analyzed findings. Over 5 million computer users in more than 200 countries have contributed over 19 billion hours of processing time. Such distributed computing projects harvest processing power; humans and human intelligence are not involved.

8 Cornelia Dean, "If You Have a Problem, Ask Everyone," *The New York Times*, (July 22, 2008, www.nytimes.com/2008/07/22/science/22inno.html.

9 Scott Page, *The Difference: How the Power of Diversity Creates Better Groups, Firms, Schools and Society* (Princeton: Princeton University Press, 2007).

10 The positive assessments of collaborative content creation derive in part from the contrast with traditional methods. See Yochai Benkler, *The Wealth of Networks: How Social Production Transforms Markets and Freedom* (New Haven: Yale University Press, 2007) for legal and economic perspectives. Charles Leadbeater, *We-Think: Mass Innovation, Not Mass Production: The Power of Mass Creativity* (London: Profile Books, 2008) offers an assessment from the bottom up.

11 Information filters such as spam blockers remove items flowing to the user; recommender systems add items flowing toward the user.

12 Cross-system collaborative filtering, where a user's filters from multiple recommender systems are combined, has been implemented but has not yet gained public acceptance.

13 Like all research based on expressed responses, those who contribute responses, either voluntarily or on request, are not necessarily representative of the larger population.

14 Robin Hanson, "The Informed Press Favored the Policy Analysis Market" (August 8, 2005), http://hanson.gmu.edu/PAMpress.pdf.

15 A brief expert summary of social network analysis is "Social Network Analysis: A Brief Introduction," www.orgnet.com/sna.html. John Seely Brown and Paul Duguid, *The Social Life of Information* (Cambridge, MA: Harvard Business School Publishing Corporation, 2000) and Marleen Huysman and Volker Wulf, *Social Capital and Information Technology* (Cambridge, MA: The MIT Press, 2004) are notable analyses. The intersection of prediction markets and social network analysis is explored by Bo Cowgill, Justin Wolfers, Eric Zitzewitz, "Using Prediction Markets to Track Information Flows: Evidence from Google," http://www.bocowgill.com/GooglePredictionMarketPaper.pdf.

16 Martin Spann and Bernd Skiera, "Internet-Based Virtual Stock Markets for Business Forecasting," *Management Science* 49 (October 2003), 1,310—1,226 offers a technical assessment of the accuracy of prediction markets in business contexts.

17 See Renee Dye, "The Promise of Prediction Markets," *McKinsey Quarterly*, no. 2 (2008) for Google's uses, Michael Totty, "How to Decide. Create a Market," *Wall Street Journal*, June 19, 2006 for applications at General Electric and Hewlett-Packard and Phred Dvorak, "Best Buy Taps Prediction Market," *Wall Street Journal*, September 16, 2008. The uses by the unnamed top-10 pharmaceutical company is based on an unpublished presentation.

18 Barnaby Feder, "To Learn What People Like, Trade 'Idea Stocks'," *The New York Times*, February 10, 2002. For a detailed description of MIT's Virtual Customer project and its aspiration to revolutionize product development, see "Virtual Consumer," MIT Sloan Management, http://mitsloan.mit.edu/vc/. On the traditional statistical procedures for this analytic task, see John Jullens and Gregor Harter, "Tracking the Elusive Consumer," *Strategy + Business Resilience Report*, November 11, 2008, www.strategy-business.com/media/file/resilience-11-11-08.pdf.

19 This binary outcome format has proven effective in generating predictions that closely approximate actual results and is the most widely used format. See David M. Pennock et al., "The Real Power of Artificial Markets," *Science*, February 2001, 987—988. It should also be noted that the settlement date can be implicit, such as "Candidate A will win the election."

20 Charles F. Manski, "Interpreting the Predictions of Prediction Markets," August 2005, www.aeaweb.org/annual_mtg_papers/2006/0106_1015_0703.pdf and Justin Wolfers and Eric Zitzewitz, "Interpreting Prediction Market Prices as Probabilities," January 8, 2007, http://bpp.wharton.upenn.edu/jwolfers/Papers/InterpretingPredictionMarketPrices.pdf provide mathematical assessments of using trading prices as a stand-in for probabilities.

21 Jennifer Watkins, "Prediction Markets as an Aggregation Mechanism for Collective Intelligence," (speech, 2007 UCLA Lake Arrowhead Human Complex Systems Conference, April 25–29, 2007) http://repositories.cdlib.org/hcs/WorkingPapers2/JHW2007 is a particularly clear exposition of these requirements. James Surowiecki, *The Wisdom of Crowds* (New York: Anchor, 2005) emphasizes how hard and rare it is to get everything right.

22 To be accurate, the commodity exchanges enable companies to hedge against price changes in raw materials as well as enable traders to speculate.

23 The use of money as an incentive is not, however, necessary for prediction markets to generate accurate results; see Emile Servan-Schreiber et al., "Prediction Markets: Does Money Matter," *Electronic Markets* 14 (September 2004).

CONCLUSION

What is appearance for me now? Certainly not the opposite of some essence: what could I say about any essence except to name the attributes of its appearance?

—Friedrich Nietzsche, *The Gay Science* (1882)

In 1993 the *New Yorker* magazine ran a cartoon by Peter Steiner in which one dog, sitting in front of a home computer, says to another dog sitting beside him, "On the Internet nobody knows you're a dog." The caption alluded to the chat rooms and electronic bulletin boards of the text-only pre-Web Internet where users, logged on under their screen names, discussed topics of all sorts. Among the obvious benefits, then and now, ordinary folks in socially stigmatized situations, such as living with diabetes or filing for bankruptcy, can get or give something helpful without disclosing their real world identities. Even more broadly, ordinary folks in ordinary situations, such as shopping for a car or finding new dinner recipes, can get what they want without dealing with salespeople or being obliged to reciprocate.[1] On the downside, the use of screen names enables sexual predators, con artists and other deceivers to pretend to be someone they are not.

Internet users are largely lackadaisical about our online anonymity. Losing it would be a disaster, but few pay it attention. Large majorities tell pollsters that they are "concerned" or "very concerned" about privacy online. In actuality, however, vast numbers of us give up our birthdays, zip codes and other personally identifying data to any Web site that asks in exchange for ringtones, screen savers, horoscopes and other trinkets. Similarly, very few use the privacy-protecting tools that have long been available. Leaking out personal data has become just a routine feature of the contemporary consumer condition.