John Favaro

The Crowd: Wisdom or Madness?

John Favaro

Viareggio, 10 April 2010

Introduction

A few years ago a brother gave me a book on musical temperament. Reading the book inspired the topic of my lecture in 2003. This year, instead of a book leading me to the topic of my lecture, it's the other way around: the topic of my lecture led me to a book.

With the rise of the Internet, Wikipedia, and social networking, the idea of talking about "collective intelligence" had been percolating in my mind for quite a while. Finally I decided to talk about it this year, and that led me to read a book called *The Wisdom of Crowds* by James Surowiecki – so now you know where part of the title of this lecture came from. It turned out that the title of that book was partly a homage to another book, called *The Madness of Crowds*. And there you have it, the title of this year's lecture –a homage to a homage. In any case, it does capture the essence of a question I want to talk about today: are we collectively smart or stupid?

I. The Wisdom

The *Wisdom of Crowds* starts out with a story about a country fair in England where a crowd correctly guesses the weight of a cow. But before I get to that story, I want to dwell a moment on the story that really had me spellbound: the story of this guy's life. I'm not talking about the author of the book. I'm talking about the man who reported on the cow-weight-guessing at the fair. His name was Francis Galton.

Francis Galton and Vox Populi

Peter Bernstein has written, "Francis Galton was one of those men of a Victorian age who roamed the earth as if he owned it." The Victorian Age – what a time that was! If you were lucky enough to be born into the right circumstances and had your wits about you and a solid dose of intelligence, the world really was your oyster. Francis Galton certainly fit the bill. He was a half-cousin of Charles Darwin, no slouch himself. Darwin's trip to the Galapagos Islands inspired Galton to travel to Africa, where he had his own share of adventures. One of his hobbies was admiring beautiful women and running up check marks on cards grading each woman he saw (he determined that the most beautiful women in Britain were in London, the ugliest in Aberdeen). While in Africa, he encountered the Hottentot tribe. The curvaceous figures of the women astonished him so much that he was dying to take the measurements of one of them. But he was understandably worried that a tribesman would grant him his wish (of dying, that is) if he caught him wrapping his arms around one of the women, measuring

John Favaro
Via Gamerra, 17
56123 Pisa - Italy
Tel +39 050 55 60 74
Fax +39 050 6143 1145
john@favaro.net
www.favaro.net/john

tape in hand. So he surveyed her – literally. He casually paced off several meters from where she was standing, and then used his surveying instruments to triangulate his measurements from a safe distance, thereby managing to obtain her dimensions without ever touching her (thereby saving himself a lot of trouble).

In other words, he was a genuine English eccentric. I have a friend and colleague who loves to collect stories about English eccentrics (another fan of English eccentrics is the author Bill Bryson). One thing they seem to have in common is an extraordinary intellect. Galton could read any book by the time he was four. He didn't just dabble in several scientific disciplines, he produced important results in them. In mathematics he invented the concept of "correlation," a keystone of modern statistics, and the concept of "reversion to the mean" (ditto). You may not immediately recognize the term, but "reversion to the mean" is actually ingrained in our everyday lives. Every time you say "this rain has to stop *sometime*" you're invoking reversion to the mean. Every time you say, "what goes up must come down," you're invoking reversion to the mean.

Galton liked to study late into the night, and to keep himself awake he invented a machine that would douse him with water if he dozed off. He also invented a device that would allow him to read underwater, and almost drowned in his bathtub one day. He was also the proud inventor of fingerprinting. It was he who got the idea that all fingerprints were unique and did not change throughout a person's life, and therefore would be useful for forensic purposes. On the somewhat darker side, he was also the inventor of eugenics, the study of heredity in human intelligence that gave Adolf Hitler some really nasty ideas several decades later. (To be fair, it doesn't seem that he shared Hitler's point of view, but things are admittedly somewhat fuzzy in the historical record.)

I said earlier that the reason Galton got into Surowiecki's book was that story about guessing the weight of a cow at a country fair. As you can see from the episode with the Hottentot women, Galton loved to measure things. It turned out by a stroke of luck that the various guesses made by the people had been written down on cards and saved. So suddenly Galton had a mass of data on his hands – irresistible for one of the founding fathers of modern statistics. (In statistics they have a saying: "Without the data, yo' chatta' don't matta'"). He took the guesses (just under 800 of them) and calculated the average. Then he noted in an article he wrote in *Nature* magazine in 1907:

> According to the democratic principle of "one vote one value," the middlemost estimate expresses the *vox populi* ...

(The title of the article was "Vox Populi," too.) In other words, as Surowiecki put it, "That number represented, you could say, the collective wisdom of the … crowd. If the crowd were a single person, that was how much it would have guessed the ox weighed."

The crowd had guessed 1,197 pounds. The right answer turned out to be 1,198 pounds. With typical English understatement, Galton summed it up with:

> This result is, I think, more creditable to the trustworthiness of a democratic judgment than might have been expected.

Certainly more than *he* might have expected. As the inventor of eugenics, whose goal was essentially to weed out inferior intellects from the population, he must have been shocked that those inferior intellects got it spot on.

Finding the USS Scorpion

Surowiecki recounts another impressive story regarding a lost submarine. On 22 May 1968 the *USS Scorpion* (SSN-589), a Skipjack-class submarine of the United States Navy, disappeared in the Atlantic, for reasons that remain unclear today. Now, the Atlantic Ocean is a big place; and when ships sink, often the simplest thing to do is just leave them there rather than go looking for them. After all, the *Titanic* sank in 1912 but they didn't go looking for it until 1985. Same with the German Battleship *Bismarck*: it went down in 1941, and they got around to looking for it in 1989. But the Scorpion was different: it was a nuclear submarine – and you can't just leave a nuclear submarine lying around. So they had to find it, period.

But as I said, the Atlantic Ocean is a big place. The Navy knew the last position of the submarine before radio contact was lost, but not much else – especially since they didn't know *why* it went down. Now, normally you would expect that the Navy would gather up the best experts they could find and pick the most plausible solution offered by one of them. Dr. John Craven, who was in charge of the search operation, decided on a different tack. For starters, he decided to use something called Bayesian Search Theory. This was not a coincidence: two years earlier, another nuclear device (this time a hydrogen bomb) had been lost when a B-52 plane had crashed in Palomares, Spain, and Bayesian Search Theory had been developed to help find it. (For my part, I find it rather unsettling that the military loses so many nuclear devices.)

Then, instead having specialists work *together*, Dr. Craven assembled the broadest multi-disciplinary team he could find – people who knew something about all kinds of different fields – and had them, *independently* of each other, guess what happened to the submarine. He even used bottles of whiskey as prizes in the guessing game. Using the Bayesian Search Theory, he collected all these guesses together and pieced them into a composite picture that was, "roughly speaking, the group's collective estimate of where the submarine was." As Surowiecki stresses, the estimated location didn't correspond to any of the individual guesses. None of these people individually had enough wisdom to guess right. But putting them all together: they found the submarine less than 200 meters away from where their collective estimate said it would be.

Collective Intelligence

How the heck did they do it? That is what has everybody so excited. Somehow, some way, it appears that a crowd of diverse, independently acting people possesses something called *collective intelligence* – an intelligence that is even more intelligent than the most intelligent individual in that crowd; and there are attempts being made everywhere to harness that collective intelligence.

A prestigious example is the Center for Collective Intelligence at the Massachusetts Institute of Technology. Here they have launched a number of projects to try to exploit collective intelligence in some very surprising applications. For example, one project has the title *The Climate Collaboratorium: Harnessing Collective Intelligence to*

Address Climate Change Issues. Another is called *Collective Intelligence In Healthcare*, and " … focuses on harnessing the collective intelligence of medical professionals, researchers, and others to provide better healthcare for individual patients." Yet another is entitled *Collective Prediction*, whose purpose is " … to make accurate predictions about future events such as product sales, political events, and outcomes of medical treatments." Let's talk about that last one now.

Prediction Markets

The lesson from Galton's cow-weight-guessing story is that a crowd can make a better guess than an individual. An obvious way to harness this phenomenon is to put the crowd to work in an area that, by its very nature, always involves guessing: predicting the future. It turns out that this is a huge business. There are so-called *prediction markets* today for just about anything. They are also called *event derivatives*. (You have just heard the word *derivative* for the first, but not the last time, in this talk.)

Some are familiar, such as horse-racing handicapping and election polls. The most famous operation is the Iowa Electronic Markets run by the University of Iowa. Others are less familiar, but no less successful. The Hollywood Stock Exchange is a virtual, Web-based game where the players effectively make predictions about things like box office success of films and this year's Oscar winners. They always seem to do better than the "experts". In 2007, they managed to predict 32 of the 39 major-category Oscar nominees and 7 out of 8 top-category winners. Corporations are very interested in harnessing the power of prediction markets to find out whether their products will sell.

But the oddest prediction market of all involved the military. On September 11, 2001, over and above the obvious human tragedy, the Pentagon also got a black eye for not seeing it coming, in spite of all its intelligence-gathering mechanisms. So they decided to try a more innovative approach and proposed the creation of what they called the Policy Analysis Market. It would be a prediction market like others, in the sense that it would be freely open to the general public; anybody could participate. But the topics didn't involve Hollywood actors or horses. Rather, they involved questions like "Will the leader of Syria be assassinated?" "When will the next terrorist attack in Baghdad occur?" Officially, the Policy Analysis Market was called "a market in the future of the Middle East," and any kind of possible political development in the Middle East was fair game for a bet.

One thing they didn't expect was the enormous backlash from all sides. Senator Ron Wyden stated in a press conference that "the idea of a federal betting parlor on atrocities and terrorism is ridiculous and it's grotesque." Others objected that betting on things like assassinations could actually create "assassination markets" that would encourage some to try and make the predictions come true. The Pentagon finally gave up on the idea – much to the chagrin of Surowiecki himself, a big supporter.

The interesting thing is that policy analysis markets as such didn't disappear as a consequence, they just migrated into the private sphere. Today, commercial prediction markets like Intrade allow bets on events like the capture of Osama Bin Laden or whether Iran will be bombed – very much the kind of thing envisioned for the original Policy Analysis Market.

Democracy

It's all very nice that the aggregating mechanisms of collective intelligence can be harnessed for horse racing, acting award ceremonies, and the like. But why not apply them to the noblest purpose of them all? Let's go back to that sentence of Francis Galton where he talks about the "…the democratic principle of 'one vote one value' …" Because in the end, this is what democracy is all about, isn't it? The people's choice, *vox populi*. And what could be more important than making **democracy** work as effectively as possible, so that the best collective choice is made by the people? And that brings us to the subject of the aggregating mechanism for political choices of the crowd: **voting**.

Galton reminds us that at the heart of democracy is the simple principle "One Person, One Vote." But that simple principle obscures a question with a surprisingly complicated answer: "What is the best way to implement the voting process?"

Many of you will have voted in the regional elections here in Tuscany two weeks ago [28 March 2010]. As usual in Italy, there were many parties on the ballot, and I'm sure that at least some of you didn't vote for the person you really wanted, because of a classic voting dilemma that might be called the "outsider" dilemma. It goes something like this: "I know that my candidate doesn't really have a chance, so if I vote for my candidate, I'm throwing away my vote – or worse, I'm playing into the hand of the candidate I really don't want to win – so I'm going to vote for this other candidate instead, who isn't the one I really want, but is better than the one I really don't want." The classic outsider presidential candidate in the United States for many years was Ralph Nader, the consumer rights crusader, and many people who would have liked to vote for him did not, for exactly the reasons I outlined above.

This problem of "not wanting to waste your vote" manifests itself in many ways. Last month [March 2010] people were pretty shocked when the historic U.S. health reform bill passed in the House *without one single Republican vote*. Aside from the obvious partisan issues, there was a feeling that the voting system itself is broken, so there's a lot of interest right now in finding better ways to vote. A few weeks ago [24 March 2010] *New York Times* columnist Thomas Friedman suggested introducing the Alternative Vote – which is actually known under a variety of names: Instant Runoff Voting, the preferential ballot, and ranked choice voting. The basic idea is that you can list more than one candidate, in order of preference. It has been in use in Australia, for example, for many years (although several readers reminded Friedman that another essential characteristic of voting in Australia is that it is compulsory).

Alternative Voting is only one of several different voting systems – exhaustive ballot, contingent votes, two round systems, and others. The point is that it's not just the collective intelligence of the crowd that matters, it's *how* you aggregate their choices in the best way to produce the best result. And as I said, where could that be more important than in our political systems?

Ants

Prediction markets and elections provide an example of harnessing the collective intelligence of human beings. But the classic examples of collective intelligence that

everybody is more familiar with are found in the animal world. Insects like ants and termites build these amazing structures by just following simple rules that they instinctively know. A well-known phenomenon is called the "Ant Spiral of Death". This happens when ants get lost for some reason. A lost ant obeys a simple rule like "follow the ant in front of you." But when they're all lost, they sometimes just end up forming a circle by following that rule, and keep walking and walking, sometimes for days, until they literally drop dead from exhaustion.

But normally their innate rules work very well, and allow them to create structures of jaw-dropping complexity and sophistication. There are several videos on YouTube that show excavations of large ant colonies by researchers. In one such video the narrator remarks,

> [After weeks of excavations], at last they begin to see the structure of the city-state. There are subterranean highways connecting the main chambers, and off the main routes are side roads. The paths branch and lead to many fungus gardens and rubbish pits. The tunnels are designed to ensure good ventilation and provide the shortest transport routes. **Everything looks like it has been designed by a single architect, a single mind, but of course that isn't true. This colossal and complex city was created by the collective will of the ant colony – the super-organism.** The structure covers 50 square meters and goes 8 meters into the earth. In its construction, the colony moved 40 tons of soil. Billions of antloads of soil were brought to the surface. Each load weighed four times as much as the worker, and in human terms, was carried a kilometer to the surface. It is the equivalent of building the Great Wall of China. It is truly a wonder of the world.

Emergence

The giant ant colony is an example of what is called *emergent behavior*. Through the individual behavior of numerous "agents" (e.g. ants), a kind of collective behavior emerges that can only be associated with the whole – the super-organism, as the narrator above called it. Here, too, the military has poked its head through the window to see if there is anything it can make use of. Sure enough, there is active research going on in the military on what is known as *swarming*.

A swarm of insects is an object of envy to military aviation planners. We've all seen how they move around in formations that compose and recompose themselves without any seeming instructions from anywhere. How does their "command and control" system work? The Air Force would like to know, and it would like to harness it for its own use. They have this idea of creating swarms of Unmanned Aerial Vehicles (UAVs) that they send out into the field against the enemy. Like insect swarms, there would be lots of them (thousands); they would be self-sufficient; their behavior would not be entirely predictable; and they would be resistant to individual losses.

If you're nervous about the idea of a bunch of deadly air vehicles swarming around, with nobody sure about just what they're going to do next, you're in good company. The military's biggest problem is convincing people this won't get out of control.

Indeed, this is the flip side of emergent behavior: it's not always a good thing. Sometimes you don't *want* emergent behavior. Sometimes, when you combine things, you just want them to do what they know how to do and nothing else. In my work in safety-critical systems (e.g. automobiles, trains, space systems), emergent behavior is a huge problem. Engineers are always worried that, even after putting a bunch of well-

tested components together in, say, a car, they'll interact in such a way that some totally unexpected (and ruinous) behavior emerges out of the whole. You want your car to do what you tell it to do. You don't want it to acquire a mind of its own.

Mitchel Resnick of MIT has created a programming system called StarLogo for studying environments with lots of simple agents (ants, fish, termites, etc.) obeying simple rules. StarLogo is based on the original Logo system, which was developed by Seymour Papert of MIT for teaching children to program. It can be freely downloaded from the Internet and has been successfully used in high school educational programs. Based upon his work studying the "massively parallel microworlds" of ants and the like, Mitchel Resnick has written a book entitled *Turtles, Termites, and Traffic Jams* – which brings me to the next subject.

Traffic

Three things have always puzzled me; after all these years, decades, sometimes even centuries of research:

- **We still can't predict the weather**. Actually, we do finally have a plausible explanation from Chaos theory regarding why this is the case (and why it will always be the case). You can find a discussion of that in my previous lecture on Chaos theory.
- **We still don't know why we sleep**. Incredibly, scientists still don't really know why we need sleep. We sleep for a third of our lives and we still don't know why! To a great extent, scientists have given up for now on the *why*, and are instead concentrating on finding ways to make sure that we get *enough* of it.
- **We still don't know what causes traffic jams.**

I would have thought that, by now, people would have figured out what causes traffic jams and done something about it. It turns out it's not that simple, and incredibly, there are enormous academic disputes about just where the problem lies. There are a number of perfectly respectable scientists who even say that traffic jams have no cause at all – they just appear "spontaneously".

If ants can (mostly) manage their traffic smoothly, why can't humans? A little reflection yields an intuitive idea of the reason: the ants are all following some kind of common set of rules that yields a high degree of coordination. But human drivers all have minds of their own. Certainly they all have to follow the basic rules of traffic, but aside from that, they pretty much do what they want to do. So one idea is to have drivers behave more like ants – that is, get them to coordinate better.

The so-called *automated highway system* (also called a Smart Road) is an example of this. They actually built one of these about 20 years ago, in Southern California. The road had sensors built into it that basically took control away from the driver and gave it to the cars. The result was that the cars all lined up beautifully and traffic flow was smooth as silk.

The problem is that drivers don't like to give up their autonomy. They don't want the road – smart or not – telling them what to do. They want to be in the driver's seat, so to

speak. So people are looking for ways to provide coordination without taking control away from the driver. I have actually come into contact with projects in which they have done this. The HIDENETS project, sponsored by the European Commission, studied the concept of car *platooning*. This consists of cars connected with each other with wireless communication so that each knows what the other is doing and they can try to coordinate, kind of like a flock of birds. Time will tell how well that idea works out – in the meantime we'll have to get used to sitting in those traffic jams.

We've spent a lot of time so far talking about the different ways in which people are harnessing the wisdom of crowds, whether in finding lost things, or predicting the future, or getting traffic untangled. I think it's time to take a look at the other side now.

II. The Madness

We had that great British eccentric Francis Galton to introduce us to the wisdom of crowds. And now we are fortunate to have his Scottish contemporary, Charles Mackay, to introduce us to the madness of crowds. Was Mackay also an eccentric? Well, he certainly shared with Galton the characteristic of being brilliant in a number of different fields. In fact, in his life he was most famous for his songwriting (like the popular *Cheer, Boys, Cheer*). He also wrote a dictionary of Lowland Scotch.

But he is known today for a book entitled *Extraordinary Popular Delusions and the Madness of Crowds*, published when he was only 27 years old. The book is still in print– in no less than six different editions. (And by the way, it is in the public domain and you can download it free on the Internet; I did just that, and put it on my Kindle to read.) The book gets off to a roaring start like this:

> Men, it has been well said, think in herds; it will be seen that they go mad in herds, while they only recover their senses slowly, and one by one.

So why is the book so famous? It actually covered a broad swath of topics that ranged from witch hunts to alchemists to the crusades. But the part of the book that earned Mackay immortality concerned *economic bubbles*. It's probably not surprising to most of you that there is so much interest in financial craziness – after all, it's been in the news a lot lately with the worldwide economic crisis. What may be surprising to you (and it certainly was to me) is how long this craziness has been going on. Even without the advantages of modern technology, transport, and communication, people have been managing to create enormous financial messes for centuries.

One bubble described by Mackay involved the South Sea Company, which was founded in 1711. The company had a monopoly on trading in Spain's South American colonies and speculation in the company's stock eventually led to a huge crash. Then there was the Mississippi Company bubble around the same time. This one also had to do with speculating on the riches of New World colonies of European countries – but this time the European country was France and the New World colony was Louisiana. In any case, the result was the same: a huge crash and a lot of ruined lives.

But the most famous bubble described by Mackay occurred nearly a century *earlier* than the other two, and this time a European country was able to make a perfectly good financial mess all on its own, without any help from New World colonies.

The Dutch Tulip Mania

The economic bubble and crash that has the dubious honor of being the first on record occurred in 17th Century Holland. Some things are so closely associated with a country that we forget they weren't always there. It's hard to imagine Italy without the tomato (in the region of Puglia where my wife is from, they call it *oro rosso* – red gold), but of course it has only been there since the discovery of the New World. Likewise, think "Holland" and the first thing that comes to mind is the tulip. But the tulip was only introduced into Holland in the mid-16th century from Turkey. Tulips were soon a big hit in Holland and everybody wanted them. Prices got higher and higher. In order to satisfy demand, an exchange was created in which people could buy "tulip futures" – contracts on future deliveries of tulips. These first-of-a-kind financial "instruments" were then further refined as *options* – the opportunity but not the obligation to buy tulips. Options are a form of – and here's that word again – *derivatives*.

By early 1637, a tulip contract was selling for more than ten times the annual income of a skilled craftsman. But as we all know now from Francis Galton and the principle of reversion to the mean, "what goes up must come down." In February 1637 the market crashed, and a legend was born: the Dutch Tulip Mania.

To this day, financial bubbles – for example, the dotcom bubble in the late 1990s that burst in the year 2000 – are commonly referred to as Tulip Mania.

To be fair, modern-day research into the Dutch Tulip Mania has led to other opinions about just how maniacal it all was. For one thing, some think that the reporting wasn't always unbiased; for example, it is thought that some reporters like Mackay might have had a religious and moral agenda, and overhyped their stories to warn people about the evils of financial speculation (we could probably use some of that today). Furthermore, there are those such as Earl Thompson of the University of California at Los Angeles who argue that not only was the Dutch Tulip Mania not maniacal, but on the contrary could be explained in perfectly rational terms, as an excellent illustration of what is known as the *efficient market hypothesis*. So what's that? It has a lot to do with the (surprisingly) fascinating story of modern finance, which I'd like to talk about now.

Efficient Markets

Economics was rocked to the core in the twentieth century, starting in its very first year. In 1900, a young French mathematician named Louis Bachelier published his doctoral thesis, called *The Theory of Speculation*. I have to admit that this time the person we are dealing with is neither British nor eccentric (unless having a wine merchant for a father and a poet for a mother counts), but young Bachelier certainly did have extraordinary talent. His thesis involved something called Brownian Motion (named after the Scottish botanist Robert Brown), which describes the motion of particles in a fluid. The interesting thing about this is that in doing so, he anticipated the Nobel Prize winning work of none other than Albert Einstein by five years – which is a pretty cool thing to be able to say about your work. Bachelier used Brownian motion to study the

movements of stock markets; Einstein merely used it to prove the existence of atoms. In addition, Bachelier's thesis involved the study of stock options, which we now know from the Dutch Tulip Mania are a form of *derivatives* (that word again). He invented a way of describing the value of options that still bears his name: Bachelier diagrams.

With his sophisticated use of Brownian motion and the like in his work, Bachelier could arguably be called the father of "mathematical finance," which eventually came to dominate most of the twentieth century, as we'll see. But the part of his work I want to talk about now is a statement he made as a conclusion of his research.

Bachelier claimed that the movements of stock markets are random.

When you read that claim, your first thought is, "Maybe Bachelier really was an eccentric; no, maybe he was downright crazy." Because it doesn't make any sense. It's obvious to anybody that stock markets don't just wander around randomly like a bunch of drunks; they move in response to real, significant events, like company reports, mergers, acquisitions, all those things. They move in patterns. Only a fool would make such an absurd claim as Bachelier's. And so he was ignored by pretty much everyone.

The years passed. A Great War devastated Europe (Bachelier himself was drafted into the French Army as a private). We lived through the greatest financial madness of all time and the subsequent stock market crash of '29, followed by the Great Depression. Then the Second World War, which gave us, among much else, the electronic computer (invented for military purposes, of course).

Then, in 1953, more than half a century after Bachelier's original work on the stock market (Bachelier was dead by that time), a British fellow named Kendall, armed with piles of stock market data (without which, his chatta' didn't matta') and all the modern tools of computing, decided to go about finding and categorizing all the patterns in that data. Who knows, maybe he even secretly harbored the hope of making some easy money – there is an active branch of finance called *technical analysis*, which tries to identify patterns in the stock market and exploit them for financial gain, and if Kendall were able to identify and categorize all the patterns, he could be a step ahead of the pack. Whether or not he had that in mind (unlikely, actually), Kendall took his piles of data and "crunched the numbers," and then tallied up the grand total of patterns he had identified.

Zero.

Not one single pattern. Nothing. After a lot of head-scratching, people began to remember Bachelier's work. But this time they couldn't ignore it like before – they had hard data that confirmed it. Stock market movements really *are* random.

But *knowing* the fact didn't help people understand *why* it was true. Economists had to wait another decade before the first plausible explanation arrived, which is still with us today. It was provided by Professor Eugene Fama of the University of Chicago, and has become known as the Efficient Market Hypothesis. Here's the basic idea. Think of the stock market as a pond full of voracious, greedy piranha (come to think of it, the analogy is probably pretty good from several points of view). When any new information arrives – like a company reporting great sales, or acquiring another

company – that sort of thing, the piranha jump on the information, analyzing it, looking at it from every perspective. They "strip it to the bones," leaving no fact undigested, and then they act on it. If it's good news, they buy; if it's bad news, they sell. And here's the important part: they not only act on the immediate information, but on all of its implications for the future: after all, if you think a company's stock price is going to go up in a month, why wait a month to buy it? You buy it *now*. The result is that the market always reflects all the information that is currently available. Its next reaction will be to the next *new* information, and of course, by definition, nobody knows when that will arrive – in other words, the arrival of new information is random and thus so are the market's movements.

So now you know why market movements are random. But the characteristic of efficient markets that is relevant to this lecture is that this "piranha pond" consists of investors of all types, from all backgrounds, all shapes and sizes, all levels of intelligence, all interpreting financial information from their own point of view and drawing their own conclusions. Like a kind of large, sprawling financial **democracy**, each investor, through the act of buying or selling a stock, is effectively casting a **vote** that reflects his opinion of that stock's value. Sound familiar? Yes indeed, we are once again back to the wisdom of the crowd, and a corollary of the efficient market hypothesis is that the crowd of piranha, in its collective wisdom, gives us the best possible estimate of the value of a stock – better than even the smartest individual can do. Or, in layman's terms: "you can't beat the market."

And that's why people are always so angry at professional money managers. Only the smallest percentage of them ever manages to do better than you would by just spreading your money around in the market. (The popular mathematician and author J.A. Paolos calls that situation "slightly scandalous" and it's hard to disagree with him.)

Behavioral Finance and the Human Factor

The efficient market hypothesis dominated finance for the second half of the twentieth century, and reinforced the impression of stock markets as rational, well-functioning investment mechanisms – the finest possible illustration of the Wisdom of the Crowd. But a nagging question remained: if the stock market is so rational and wise, what about all those things Charles Mackay was writing about? The Dutch Tulip Mania? The Mississippi Company? For that matter, what about the Great Crash of 1929?

In the last decade of the twentieth century, some people began to question whether the wisdom of the stock market hadn't maybe been a bit exaggerated, and its madness a bit downplayed. One of those people was Robert Shiller, an economics professor from my own *alma mater*, Yale University. Like Francis Galton, Shiller also preferred to work with real data (instead of just chatta') and basically said, "Okay, if stock market prices are supposed to be the wisest estimate of the real value of companies, let's go verify whether that has turned out to be true in the past." After all, data was lying around from over a hundred years of stock market activity and company financial reports. All he had to do was compare them with each other to see whether the crowd had guessed right – in retrospect, it's surprising that nobody had thought of it before. Sure enough: Shiller found that, time and again, those "wise" investors were way off the right estimate of a company's real value. His explanation? Simple: human foibles – the madness of the crowd.

Shiller became one of the leading proponents of a new discipline known as *behavioral finance*, which explicitly acknowledges the influence in financial decision-making of overconfidence, overreaction, and all those other quirks of humanity we know and love. His comment: "When we started doing behavioral finance we were total outcasts … nobody appreciated us. I had tenure so I could do it …"

In the year 2000, right at the very height of the dotcom boom, Shiller published a book called *Irrational Exuberance*, warning about the current madness of dotcom investors. (This madness was going on in Italy, too, by the way – one company was stampeded by crazed investors desperately trying to buy its shares because it had a product with "Net" in its name. It turned out that the product, WC-Net, was a detergent for cleaning toilets). When the market crashed shortly thereafter, a lot of believers in behavioral finance were born.

In 2005, Shiller published a revised version of *Irrational Exuberance*. This time he included a discussion of the growing housing bubble. He pointed out, for example, that median home prices in the United States had grown to be as much as nine times greater than median income (remember that discussion about Dutch tulips costing ten times the annual income of a skilled worker?). And as we now know, he was right – the madness of the crowd had struck again and we are still now enduring its consequences. This time the culprit wasn't the Internet. It was something far more sinister – that word again: *derivatives*. The legendary American investor Warren Buffett has called derivatives "financial weapons of mass destruction." You may have also seen the news last month [March 2010] that four banks were being put on trial for their handling of derivative transactions with municipalities in Italy.

As a matter of fact, though, Shiller isn't against derivatives *per se*. He shares the view of James Morgan, a columnist from the *Financial Times*, who said, "A derivative is like a razor. You can use it to shave yourself … Or you can use it to commit suicide." Shiller is in favor of derivatives when they are used responsibly to diminish risk. But derivatives can also be used for very wild and aggressive speculation, and seem to bring out the worst madness in the crowd. Witness the mess the world is in right now.

The current financial disaster has thrown the field of economics into a serious identity crisis. The century that began with brilliant economists like Louis Bachelier laying the foundation for a rational, crystalline mathematical model of financial behavior ended with brilliant economists like Robert Shiller tearing down that very foundation. In an article last week [26 March 2010] in the *New York Times*, David Brooks discussed the causes and nature of all this soul-searching among economists. The principal cause is that so few economists saw the crash coming.

> "Where were the intellectual agenda-setters when this crisis was building?" asked Barry Eichengreen of the University of California, Berkeley, in The National Interest. "Why did they fail to see the train wreck coming?"

I mean, it's pretty embarrassing when the worst economic collapse since the Great Depression happens and nobody (except one or two like Shiller) foresees it. It kind of makes you reflect. Right now, in the tug-of-war between the Mathematical Economists and the Behavioral Economists, the behavioralists seem to have the upper hand. As Brooks concludes,

> Economics achieved coherence as a science by amputating most of human nature. Now economists are starting with those parts of emotional life that they can count and model (the activities that make them economists). But once they're in this terrain, they'll surely find that the processes that make up the inner life are not amenable to the methodologies of social science. The moral and social yearnings of fully realized human beings are not reducible to universal laws and cannot be studied like physics.

In other words: human nature being what it is, there will always be a little madness mixed in with the wisdom. And that has implications for the limits on how far we can go with these ideas, as we'll see now.

The Wisdom of Corporations?

As I said earlier in this talk, people have been trying to harness the wisdom of the crowds in many different fields, and the field of business and finance is no exception. In fact, Surowiecki's book *The Wisdom of Crowds* was a *New York Times* bestseller not in the *general* category, but in the *business* category; Surowiecki's hope and intention was to inspire businesspeople to find ways to harness the ideas in the book. For example, he outlined ways in which the ideas of collective intelligence could be used to produce superior corporate decision-making processes.

I had the occasion in Rome exactly one month ago [10 March 2010] to meet a person who is working in this area right now. Professor Michele Missikoff is a director of the Laboratory for Enterprise Knowledge and Systems at the Italian National Research Centers in Rome. He has been trying to raise awareness in Europe of the possibilities for harnessing the power of the crowd in the enterprise. He is promoting a view that each person in the enterprise can be considered as a valuable carrier of knowledge – however partial and incomplete – that contributes to a collective, enterprise-wide "wisdom" that is more powerful than anything any individuals could put together on their own. This could lead to whole new ways of managing knowledge in the enterprise.

"Considering the importance of the enterprise in our economy, it's amazing that more hasn't been done to try to make it function better," he told me. "The current economic crisis has made authorities in the European Community more aware of the need to look at enterprise improvement from an economic point of view rather than just a technical point of view."

"So what's holding back progress?" I asked.

"The enterprise is the single most complex structure ever made by humans," he explained. "The problem is that it's not just made *by* humans, it's made *of* humans. And human nature being what it is … for example, it's not always easy to convince a manager that he should relinquish his hard-earned right to make decisions to the collective intelligence of some kind of group." He shrugged his shoulders and smiled, "That's just the way we humans are."

III. The Future

I mentioned at the beginning of this talk that I was originally led to this topic not by reading a book, but rather by what I was seeing happen all around me. Examples of

trying to harness collective intelligence are everywhere now, especially on the Internet. You don't have to look any farther than your home page on the Web for a great example of this – that is, if your home page is Google, the famous search engine.

Did you ever wonder how Google actually finds the things you look for on the web? Among all the web pages out there, how does it find the ones that are relevant to your search, sort them and rank them so that you get the best ones? Simple: it uses a technique for ranking pages imaginatively named … PageRank.

Now, you could be forgiven if you thought that the name stood for "ranking pages." But you'd be wrong. The "Page" actually refers to Larry Page, one of the co-founders of Google. He developed this technique while he was a student at Stanford University (by the way, the work of an Italian, Massimo Marchiori at the University of Padua, is cited as being one of the sources of inspiration for this technique). Another interesting fact is that, although the *name* PageRank is a trademark of Google, the *patent* on the technique is actually owned by Stanford University, and they made a bundle of money licensing the exclusive rights to Google.

So how does this page-ranking technique work? Let's hear it from Google themselves:

> PageRank relies on the uniquely **democratic** nature of the web by using its vast link structure as an indicator of an individual page's value. In essence, Google interprets a link from page A to page B as a **vote**, by page A, for page B. … Votes cast by pages that are themselves "important" weigh more heavily and help to make other pages "important".

Democratic. *Vote*. Those words again! Yes, that's right; Google essentially harnesses the wisdom of crowds to let the Internet itself tell you which pages are relevant to your search.

Perhaps that's not so surprising, because the Internet is the most fertile ground you could possibly imagine for harnessing the wisdom of crowds. And even less surprising is that where it's all happening right now is in the *social networks* – the ultimate aggregators of crowds.

There probably aren't three people in this room who aren't registered on Facebook (Italy has one of the world's most enthusiastic Facebook communities). Last year I introduced you to Twitter, who taught us that the burning question on everybody's mind was "What are you doing right now?"

This year, we have gone one step further. Writing a couple of weeks ago [21 March 2010] in the *New York Times*, David Carr described the unnerving experience of sitting in a bar in Austin, Texas, where suddenly about seventy people, all sitting in different parts of the bar, got up out of their seats and rushed over to another bar. What made them all suddenly move together as a group, like a swarm of bees?

> At large events, people have always moved in groups to the next big thing. But … the ubiquity of so-called **ubiquitous presence** — location-based services like Foursquare and Gowalla — meant **the hive suddenly knew what it was collectively doing**. … It was striking to see the digital location effect in the wild, with people reacting to an unseen dog whistle and moving en masse, on command.

Ubiquitous presence? Location-based services? What on earth is he talking about? This is the new frontier in social networking. The burning question isn't just "What are you doing right now?" any more, but also "*Where* are you right now?" *Social* networking is now morphing into *geosocial* networking. In geosocial networking, friends not only tell each other *what* they are doing but *where* they are doing it.

But how do they do this? It has a lot to do with the concept of *ubiquitous presence* mentioned above. This is a hot topic in the research programs of the European Union right now, as a matter of fact. In its so-called Seventh Framework Programme of Research, the term "ubiquitous presence" is cited often, and refers to the idea of everything being interconnected, from your refrigerator to your clothes to you yourself. (Another name commonly used in association with the concept is the *Internet of Things*.) But how do you connect *you* to others? Simple: your mobile phone.

Foursquare, one of the "location-based services" mentioned above, works like this: when a person enters a place (like a popular bar) he can *check-in* using his smartphone. In that way, his friends all know where he is. Of course, Foursquare *also* finds out where he is, and privacy issues immediately come to mind. Carr writes:

> To someone not in their 20s whose location generally isn't that interesting to others … the idea of handing over your privacy with both hands to strap on a digital ankle bracelet sounds profoundly unattractive. … But to a younger cohort that lives on the grid, the location of people you know and care about is vital information, the coin of the realm.

Like in a beehive, we are beginning to see new forms of emergent behavior in social networks, like groups suddenly standing up in a bar and moving to a different place. Like those cars coordinating with each other on the highways, think of this as a kind of "human platooning," where everybody knows where everybody else is. This kind of collective, emergent behavior is also encouraged by services like Foursquare, who have introduced game-like facilities and rewards into the service. For example, you can earn badges by checking into places. And if you check in the most times to a popular place (like a bar), you can be named "mayor" of that place. These kinds of awards are reminiscent of the whiskey prizes using in finding the submarine *Scorpion*, and in the same way, they encourage the crowd to interact in all kinds of unexpected ways.

Furthermore, location-awareness is not destined to remain in the purely physical world. As Carr writes:

> … What if location became not just a physical place, but a digital one? The possibilities for old and new media could be significant. "The check-in is bigger than location," said Yancey Strickler of Kickstarter, a Web site that helps with fund-raising for media products. "Think of media: Checking into watching 'Lost,' being declared the mayor of 'The Brothers Karamazov' or earning a badge for braving free jazz."

Recently, much larger social networking companies like Facebook and Twitter have also announced that they will soon provide location-aware facilities. This means that crowds will arise more and more often, with their members getting together and acting in collective, emergent ways that we can't begin to predict. Where will all this end? Once again I find myself at the end of a talk having to use the same old, tired line: Only time will tell.

So I guess I didn't do a very good job of answering the question I posed at the beginning of this talk: are we collectively smart or stupid? As usual, I ended up equivocating: we'll always be a little wise, and we'll always be a little mad. But most of all, more and more we'll be *we*. Unlike the individual, royal "We" of Francis Galton's beloved Queen Victoria, the "we" of the future will be ever more connected, interacting, self-aware, the ultimate universal swarm of humanity.

Resources

James Surowiecki, *The Wisdom of Crowds*, Anchor (August 16, 2005)

Charles Mackay, *Extraordinary Popular Delusions and the Madness of Crowds*, with a foreword by Andrew Tobias (1841; New York: Harmony Books, 1980). ISBN 0-517-53919-5. Also available free on the Internet, e.g. www.manybooks.net.

Toby Segaran, *Programming Collective Intelligence: Building Smart Web 2.0 Applications*, O'Reilly Media; illustrated edition (August 16, 2007)

Don Tapscott, Anthony D. Williams, *Wikinomics: How Mass Collaboration Changes Everything*, Portfolio Hardcover; Expanded edition (April 17, 2008)

Jeff Howe, *Crowdsourcing: Why the Power of the Crowd Is Driving the Future of Business*, Three Rivers Press; unedited edition (September 15, 2009)

Richard Ogle, *Smart World: Breakthrough Creativity And the New Science of Ideas*, Harvard Business School Press; 1 edition (June 5, 2007)

Keith Sawyer, *Group Genius: The Creative Power of Collaboration*, Basic Books (June 4, 2007)

Steven Johnson, *Emergence: the connected lives of ants, brains, cities and software* (2002) Scribner, ISBN 0-684-86876-8.

Mitchel Resnick, *Turtles, Termites, and Traffic Jams: Explorations in Massively Parallel Microworlds* (1997), MIT Press, ISBN 0262680939.

Peter Bernstein, *Against the Gods: The Remarkable Story of Risk* (1998), Wiley & Sons.

Peter Bernstein, *Capital Ideas: The Improbable Origins of Modern Wall Street* (1992), Wiley & Sons.