

HARNESSING THE WISDOM OF CROWDS: The New Contours of Intellectual Authority

**Remarks to The Fields Institute for Research in Mathematical Sciences
Annual General Meeting
Toronto, 15 June 2006**

by

Dr. Peter J. Nicholson

President, Council of Canadian Academics
180 Elgin Street, Suite 1401, Ottawa, ON, K2P 2K3
613-567-5000; peter.nicholson@scienceadvice.ca

OUTLINE

- **The Decline of Deference**
- **Role of the Media**
- **The Dilemma of the Information Age**
- **Massively Distributed Collaboration**
- **Harnessing the Wisdom of Crowds**
- **Implications for Mathematics**

HARNESSING THE WISDOM OF CROWDS: The New Contours of Intellectual Authority

I will argue that what qualifies as *intellectual authority* in contemporary societies – who and what to believe – is changing fundamentally. I will speculate as to the reasons, and draw out some of the implications for intellectual work in the future. I will even venture to speculate on possible implications for the conduct of mathematics research, though the main object of my message this evening concerns the broad sweep of intellectual endeavour and not any specific domain.

The thesis in a nutshell is this. People today are much less prepared to defer to the experts. But at the same time, we are being swamped with data and information – a glut that cries out for analysis and summary. So there's a dilemma. Who to turn to? Increasingly the answer is – Well, to ourselves of course, as individuals empowered by a world wide web that has rapidly evolved into a *social* medium. More specifically, it is a medium that today supports *massively distributed collaboration* on a global scale that – we can only hope – will help us make sense of it all.

This phenomenon of massively distributed collaboration is perhaps best exemplified by Wikipedia – the wildly popular, on-line, user-created encyclopedia. I will have a lot more to say about Wikipedia later in these remarks. But the underlying notion has more prosaic origins, many of which are described in James Surowiecki's fascinating book – *The Wisdom of Crowds*, [1] the title of which is an ironic inversion of Charles MacKay's 1841 classic, *Extraordinary Popular Delusions and the Madness of Crowds*.

The *wisdom* of crowds, as Surowiecki describes, was encountered – perhaps to his dismay – by the great statistician and elitist, Francis Galton, during a visit to a country fair in England a hundred years ago. There, his interest was drawn by a contest to guess the weight of a particular ox. He later consulted the record of the roughly 800 guesses and discovered to his amazement that their mean value was essentially spot-on, differing from the ox's actual weight of 1,198 pounds by just a single pound! The collective guess of the crowd – notwithstanding its motley composition – was much more accurate than that of any individual despite the presence

of many contestants who, as farmers and butchers, could be considered experts in the field of “ox knowledge.”

So here at least we have a counterexample to the elite convictions of such great intellects of that time as Freidrich Nietzsche who wrote that “madness is the exception in individuals, but the rule in groups.” Or, as Thomas Carlyle haughtily put it: “I do not believe in the collective wisdom of individual ignorance.”

A lot has changed in the hundred and twenty-five years since Carlyle’s death. My purpose in these remarks is to offer a perspective on what seem to me to be the deepest and most pervasive changes that are now shaping the contours of virtually all forms of intellectual authority, transforming its landscape in ways that stand Carlyle’s elitist conviction on its head.

Let me say at the outset that I am not particularly comfortable with the future I foresee. I am, after all, a charter member of the ‘old guard’ and will never really belong to the new. But I am also an optimist and a realist. The world has changed – and so must we.

The Decline of Deference

Let me begin with some very general remarks on contemporary attitudes toward hierarchical authority generally – of which intellectual authority is but one instance.

As President of the new Council of Canadian Academies – an organization that will oversee expert studies of the science underlying important public questions – I am in the business of brokering intellectual authority. I admit to being a traditionalist in the sense that I believe intellectual authority should have a close correlation with expertise. And it should flow from the tried and true, though never infallible, processes of peer review and other forms of elite consensus building.

More than that, I am comfortable with hierarchies that are based on merit. And I am quite willing to defer to the well-established institutions in today’s society since, on balance, I

believe that their power is adequately constrained by the legal, economic and political structures of modern democracy.

But I am also convinced that the values that have shaped my world view – and that of my demographic peers throughout the industrialized world – are being eclipsed by a new paradigm. This new framework is shaped by technology – primarily information and communications technology; by globalization; by post-industrial affluence; and by a culture which, as never before, celebrates and empowers the *individual*.

One of the most significant symptoms of this pervasive shift is the decline of deference to virtually all forms of traditional authority – whether the church, the school teacher, the family doctor, the business executive, the union leader, the politician, and not least, the intellectual. In short – out there on main street, mistrust and scepticism reign.

While all this is widely recognized, the truly fundamental reasons for the decline of deference seem not to be generally understood in broad sociological terms. The explanations we do see typically cite the public revulsion that stems from *specific* cases – for example, scandals in the Catholic Church; or in businesses like Enron; or in politics – Watergate; our own “sponsorship affair”; the failure to find WMDs in Iraq; or to warn the British public of BSE. Take your pick.

The key point is this. The decline in trust of – and therefore deference to – traditional sources of authority is a nearly universal feature of advanced societies. It transcends every specific, local instance. And it didn't just happen yesterday. Deference to hierarchical authority has been declining for at least the past 50 years. Clearly, therefore, we are witnessing a socio-cultural change whose roots run deep in the character of economically advanced societies.

From whence does it spring? The best account I have read is by U of T political scientist, Neil Nevitte. His 1996 masterpiece, *The Decline of Deference*, draws on a rich vein of multi-country time-series data – the World Values Survey – to establish convincingly that “the new citizens are less likely than their predecessors to be satisfied with any form of

authoritarianism. . . Citizens, cut from the newer cloth, are more attracted to formations that are bottom-up." [2]

Thus societies formerly based on deference to authority, community loyalty, and the struggle for the material basics of life have given way to societies, the affluence of which has engendered a generational shift toward the so-called "post-materialist" values of self-esteem, quality of life, and the search for personal fulfillment.

When these objectives are combined with the empowering tools of universal education, a rights-oriented political culture, and the Google search engine, we should not be surprised that people – and particularly younger people – regard *ex cathedra* expert authority with scepticism, if not outright hostility.

The paradox is that expert opinion is being sought and cited more than ever. But increasingly, it is individuals themselves who weigh the various authorities and come to their own conclusion. Just ask doctors about their web-savvy patients.

Role of the Media

Let me open an important parenthesis here on the role played by the media in shaping broader public attitudes toward intellectual authority. The prevailing ethic in journalism is that "fairness" requires that all views on an issue be presented, often without regard for the relative weight of authority of various sources being quoted. The objective is simply to report point, and counterpoint, with an emphasis increasingly on sensationalism, official screw-ups, and conflict – i.e. those things that can attract at least fleeting attention, and advertising dollars, in an information environment that has become super-saturated.

The net effect is to create in the public mind an impression that experts can never agree; and expert authority is thereby diluted. Fortunately, I don't see this as representing much of a threat to the authority of expert mathematicians – but then again, math controversies don't get a lot of front page ink anyway!

The same certainly cannot be said for medical journalism where the daily *reported* advice does make the front page. And the advice in the mass media on how to stay healthy keeps flip-flopping, whereas the full text of the journal articles would reveal the provisional nature of findings, statistical caveats, and so forth. The bottom line is that superficial media treatment of scientific and technical issues reinforces the prevailing scepticism as to the consistency and trustworthiness of expert authority.

The Dilemma of the Information Age

Coming back to the main line of my argument, we find that while expert-based authority is being challenged, the volume of information and the economic significance of knowledge are exploding. Information technology itself – whose capacity continues its four-decade exponential improvement – is clearly a key part of the reason. But so too is the huge global expansion of knowledge-generating capacity, the more so as China and India and other giants plug into the economic and research networks of the industrialized world. These societies are adding tens, and soon hundreds, of millions of trained knowledge workers. They will bring not only new sophistication and motivation, but also cultural and intellectual perspectives that are quite different from those of the West. We can therefore expect an unprecedented surge of innovation, and new impetus to the information glut, as the two worlds meet, finally on equal terms.

So we are confronted with a dilemma.

On the one hand, the whole world is struggling to cope with an information explosion that shows no sign of letting up – quite the contrary. We need somehow to transform a data torrent into useful information and knowledge that can power economic progress and human fulfillment.

But on the other hand, the agents we have relied upon traditionally to filter and manage information, and to broker formal knowledge – agents like research universities, the serious media, and highly trained experts of all kinds – are less trusted as intermediaries than they once were. And even if that were not the case, we might doubt that these expert resources are really

up to the task of managing the information glut anyway. Just ask journal editors and referees, or researchers in any dynamic field, how well they are keeping up. Ask yourselves.

Part of the response, of course, has been to deploy the same computer technology that is facilitating the information explosion in the first place, to help cope with its management. In other words, the offence is also the defence. That's why Google Inc. today has a total stock market value of more than US \$115 billion – over four times the combined worth of Ford and GM. And it's also why "Google" has become a verb. (Who remembers when it was merely 'one' followed by a hundred zeros!)

But Google and its ilk notwithstanding, the sheer volume of information, its global origins, and especially the dynamic, real-time nature of information today is simply overwhelming our traditional, centralized institutions of information screening and management – whether research libraries, book and journal publishers, or newspapers and other mass media.

The *infosphere*, if I could use that term, therefore needs new and decentralized mechanisms of self-regulation and self-organization, much like a complex economy which, as Adam Smith realized, needs the guidance of an invisible hand.

Massively Distributed Collaboration

I believe that the outlines of just such a mechanism are already emerging in the multifaceted development of what cyber-prophet, Mitch Kapor, recently dubbed "massively distributed collaboration." [3] Probably the single best example, as I mentioned at the outset, is Wikipedia, the user-edited encyclopaedia that in just over five years has become one of the most-visited sites on the web.

In fact, something much broader is going on. The world wide web has already morphed into a *social* medium – what some are calling Web 2.0 – a global many-to-many meeting place, very unlike the one-to-many connections of radio, TV, books and newspapers. The latter media are inherently hierarchical – a communicator of one to an audience of many. The social web,

on the other hand – like Thomas Friedman’s new world – is *flat*. It is in tune with today’s ethos. Just consider some of the manifestations:

- 20 million blogs – and counting;
- Self-expression portals like *My Space* and *Facebook*, growing explosively – indeed a new cultural phenomenon, tellingly dubbed “Me Media;”
- Massive multiplayer games like *Ever Quest* and *Second Life* where the players themselves shape the dynamic environment;
- The Linux operating system, flagship of the open source software movement, and maintained by a worldwide network of volunteers;
- eBay – the many-to-many model implemented as a phenomenally successful digital marketplace;
- Amazon, and countless other “collaborative filtering” sites that tally and report user satisfaction;
- And Google itself, which indirectly exploits massively distributed collaboration via its page-rank technology to aggregate the behaviour of millions of users into an index of relevance.

In summary – and this is my key message – we are witnessing in these examples the *convergence* and *mutual reinforcement* of two of the great defining movements of the past half-century – one cultural, the other technological – i.e. the ascendancy of the ordinary individual together with the empowering technology of the computer, now enormously amplified by global networking – creating essentially a “cyber nervous system” for the entire planet.

This is an epochal development that will not be reversed. The job for all of us beyond a certain age, but still hankering to be part of the action, is to figure out how to be a constructive part of it.

Harnessing the Wisdom of Crowds

In the remainder of these remarks I want to take a closer look at one important example of massively distributed collaboration (MDC) – specifically, the on-line encyclopaedia movement, since this illustrates most directly how MDC is already harnessing the wisdom of crowds and thereby reshaping the contours of intellectual authority.

The flagship example is Wikipedia, founded only in January 2001, but already the site of nearly four *million* entries in almost 200 languages. There are more than 1.1 million articles in English, growing by about 1,500 a day. [4] The Encyclopaedia Britannica, by contrast, has merely 65,000 articles in the print edition and 75,000 on-line. [5]

What is most amazing is that Wikipedia is doing all this on an annual budget of just \$1.3 million, 60 per cent of which goes for the cost of computer hardware, leaving only about \$500,000 to cover everything else! [6] How can that possibly be? With apologies to those who are already very familiar, here's what's going on.

For starters, the articles are written, and re-written, by volunteers. The website is equipped with so-called "wiki" software that allows anyone with a browser to edit virtually any article at the push of a button. In the flat culture of Wikipedians, experts and dunderheads are equally welcome. The main editorial principle is that articles should reflect *a neutral point of view*. This is not a site for cranks and propagandists. Acts of deliberate vandalism are not tolerated and are usually corrected very quickly. On the other hand, decisions as to what is deemed to be unjustified bias are taken consensually, and this can be excruciatingly drawn-out in contentious areas.

At first blush, it admittedly sounds a lot like "monkeys with typewriters." But in fact it's not. In a widely publicized and controversial head-to-head test with Britannica, reported last December in the journal, *Nature*, expert reviewers determined that Wikipedia articles, on average, contained "only" a third more inaccuracies than their Britannica counterparts. More to the point, only eight serious errors were reported in the sample of 42 topics with an equal number, four, attributed to each source. [7]

Having read the full text of the debate between *Nature* and Britannica over the methodology of the comparison, I would grant many of Britannica's objections, but would still conclude that the essence of *Nature's* findings remains intact. Wikipedia is surprisingly good, especially for a five-year old; and even a source as well-researched as Britannica still contains a significant number of inaccuracies. So much for any presumption of expert infallibility.

The real bottom line, of course, is that notwithstanding doubts about its reliability, Wikipedia has taken off like a rocket. We need to understand why.

Obviously, being instantly accessible and *free* – no ads at all – is a big plus. But the real power of Wikipedia is that it's in perfect synch with web culture – which mirrors today's attitudes, and even more so tomorrow's. Wikipedia is also in synch with globalization – 200 languages represented with much of the content original to each language, not simply translated. And Wikipedia – like Google, and blogs, and open source software – operates in synch with the *rhythm* of the web, incorporating new information continuously in real time, 24/7.

This last point is important, and is part of a much larger story. I can only summarize. The “half-life” of *active* information has been getting shorter and shorter due primarily to the sheer rate of information generation. There is more and more to process, but not more hours in the day, and not more raw individual brain power to apply. So we graze, or we gulp, and then we move on.

The half-life is also shrinking due to the very nature of electronic technology which makes overwrite so easy and natural. We are all becoming addicted to the “refresh” button. Documents of every kind – certainly in my experience in business and government – are being revised continuously until the moment they become virtually obsolete. And as the shelf-life of any particular information product gets shorter – whether it's an e-mail or a position paper – basic principles of economics dictate that fewer resources of time and money can be put into its creation. The ubiquitous deck of bullet points is the iconic example.

The result is a dumbing down of written communication. We can decry it – and I do – but it reflects a probably necessary trade-off in favour of easier and quicker absorption, unfortunately at the expense of nuance and rigor.

This has profound implications for how good is “good enough” when it comes to authoritative information. Of course, complete accuracy still matters as much as ever where lives or fortunes (or mathematical theorems) depend on it. But for most everything else, the tradeoff point is moving toward faster, not deeper.

This is a context in which massively distributed collaboration systems like Wikipedia excel. But the advocates of MDC claim much more, and believe that it can be both faster *and* deeper. They may have a point based on the old adage that two heads are better than one – and thousands or millions of heads are incomparably better. And so we come full circle to the wisdom of crowds and to the validating belief of the open source software movement, summed up in the motto – *Given enough eyeballs, all bugs are shallow*. [8]

Maybe. But in the case of specialized subjects where quality criteria are more judgemental (unlike software bugs), or where relevant expertise is spread very thinly – and perhaps nowhere so thinly as at the frontier of mathematics research – the “crowd” is unlikely to be sufficiently wise. So there will always be a secure niche for expertise in the traditional sense. Indeed, that conviction led Wikipedia’s co-founder, Larry Sanger, to leave what he had created out of despair over the hostility toward expert authority that dominates Wikipedian corporate culture. Sanger is now creating a new on-line authority, Digital Universe, that seeks to provide both expertly-created as well as collaboratively-developed content. [9]

We should stay tuned, because the puzzle that the Larry Sangers of this world are trying to solve is to *evaluate* and *integrate* very different methods of ascertaining intellectual authority – ranging from the continuously-flowing, collaboratively-determined “truth” of Wikipedia and its ilk, to the timeless records of solitary genius.

Indeed we should be thinking of the infosphere as an ecosystem where different “species” are adapted to specific niches. Google, for example, delivers fantastic volume but the measure of relevance is still pretty crude. Blogs give you an up-to-the-minute read on what’s hot. Wikipedia provides a great first cut at coherently organized material plus a good set of relevant links. And if reliability is a critical objective, then refereed journals and original documents become progressively more important. But the contemporary niches in the information ecosystem are neither stable nor secure. Instead they are shifting continuously in response to technological and cultural changes, which, as I have argued this evening, are reshaping fundamentally the contours of intellectual authority.

Implications for Mathematics

What are we to make of all this? As always, the one key question for each of us is – What does it mean for me? For most of the audience in this room that question translates to – “Does any of this relate to the world of mathematics research?”

One might reasonably think not, arguing that mathematical creativity of the highest rank has always been a solitary and esoteric activity, the talent for which is extremely rare, much like artistic and literary genius. Proving theorems about non-Abelian groups is not a bit like averaging guesses about the weight of an ox.

True – but on the other hand, writing the software to implement the Linux operating system is – in terms of sheer intellectual complexity – comparable to resolving some of the knottiest problems in pure mathematics. Yet the Linux challenge – and a growing number like it – are being met through a massively decentralized collaborative effort by thousands of volunteer hackers motivated both by the sheer challenge, and by the psychic and social rewards of belonging to a very special club of peers.

Might not similar conditions apply to important areas of mathematics research?

When I googled “collaborative mathematics,” one of the first entries was something called the *Flyspeck Project* which is dedicated to the collaborative development of a purely formal proof of the ancient Kepler Conjecture on the maximum packing density of spheres. This famously difficult problem was apparently solved by traditional methods by Hales and Ferguson in 1998. At that time, a panel of 12 referees was assigned to verify the proof, but after *four years* the panel could only conclude that it was 99 per cent certain of its correctness. [10]

The Flyspeck Project hopes to use specialized computer software and an army of volunteers to develop a “formal” proof of Kepler, the correctness of which will be virtually assured by the method of its construction. The Flyspeck web site includes the following line:

“We are looking for mathematicians, from the advanced undergraduate level up, who are computer literate and who are interested in transforming the way that mathematics is done.”

An isolated case, you say. Perhaps. And certainly the traditional skills of mathematicians will continue to be relevant in collaborative projects just as world-class software skills are table stakes for the open source hackers who keep improving the Linux operating system.

But I would wager that mathematics will not remain isolated from the deep and pervasive changes we have been discussing this evening. And fundamentally, that's simply because the coming generations of mathematicians will be children of the web culture, globally-networked, equipped with unimaginable information-processing power, and devoid of deference to hierarchical authority, intellectual or otherwise.

Can we believe that they won't do things very differently?

Source Notes

1. James Surowiecki; *The Wisdom of Crowds*; 2004.
2. Neil Nevitte; *The Decline of Deference*; 1996.
3. Presentation by Mitch Kapor at UC Berkely; 9 November, 2005
(<http://www.sims.berkeley.edu/about/events/dls11092005>) Kapor states: "...The sudden and unexpected importance of Wikipedia...represents a radical new modality of content creation by *massively distributed collaborations*. This talk ...will examine the intriguing prospects for application of these methods to a broad spectrum of intellectual endeavours."
4. *Internet Encyclopedias go head to head*: Nature 438; 15 December, 2005.
5. *Who Knows?* The Guardian; 26 October, 2006.
6. Cited in article on *Wikipedia* in Wikipedia. It is reported that 4th Quarter, 2005 costs were \$321,000 (\$1.3 million at annual rate) with hardware making up almost 60%.
7. Nature 438; *op.cit.*
8. Eric Raymond; *The Cathedral and the Bazaar*; First Monday; 2 March, 1998.
9. Larry Sanger; *Why Wikipedia Must Jettison its Anti-elitism*; Kuro5hin.org; 31 December, 2004.
10. <http://www.math.pitt.edu/~Ethaes/flyspeck/>