Exceptional service in the national interest

Sandia National Laboratories

http://www.smithsonianmag.com/science-nature/The-Top-Ten-Daily-Consequences-of-Having-Evolved.html

Engineering the Ultimate Dynamical Social System:

what we know and don't know about how scientists do science

Jeff Tsao

2012 December 5

U.S. DEPARTMENT OF ENERGY · NNSA National Nuclear Security Administration

Sandia National Laboratories is a multi-program laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000.

- HELLO. Good afternoon. I'd like to talk today about what I sometimes think of as the ultimate dynamical social system: the system that produces knowledge, the ultimate form of which we call science.
- WHAT IS SCIENCE? Science, of course, is first and foremost the ideas that humanity uses to understand and control the world around it. But science, secondly, is the human social community in which those ideas exist and which has, at least for the time being, conquered the earth. Perhaps most importantly, science is widely viewed as key to humanity's future: how we fix our energy and environment problems, and how we innovate to mitigate various "gathering storms." Indeed, as some of you may know, the National Science Foundation is in the fifth year of a major program aimed at creating a science of science policy and ultimately at engineering an improved system of science.
- SCIENCE OF THE SYSTEM OF SCIENCE VS SCIENCE OF HOW INDIVIDUALS DO SCIENCE. In fact, the National Science Foundation might be missing an important piece. In order to engineer an improved system of science, it's not only important to understand in a holistic way the system of science, but to understand in a reductionist way how individual scientists within the system of science practice the art of science. Why do I think this? I think this because it's difficult to intelligently engineer a system unless you know something about how its parts behave. It'd be like trying to intelligently engineer a system of education without knowing how people learn. Or like trying to intelligently engineer a system of highway transportation without knowing how people drive.
- OPPORTUNITY IN THE REDUCTIONIST SCIENCE OF HOW SCIENCE IS DONE. Moreover, I believe there is a lot of opportunity in digging more deeply into the reductionist science of how science is done. If we can understand better how science is done, perhaps we can better do science, even without re-engineering the system.
- TODAY'S TALK. So what I'd like to talk about today is what we know and don't know about how individual scientists actually "do" science, with an emphasis on those aspects that we don't understand but if we did could improve how those individual scientists do science.

Science (and knowledge in general)

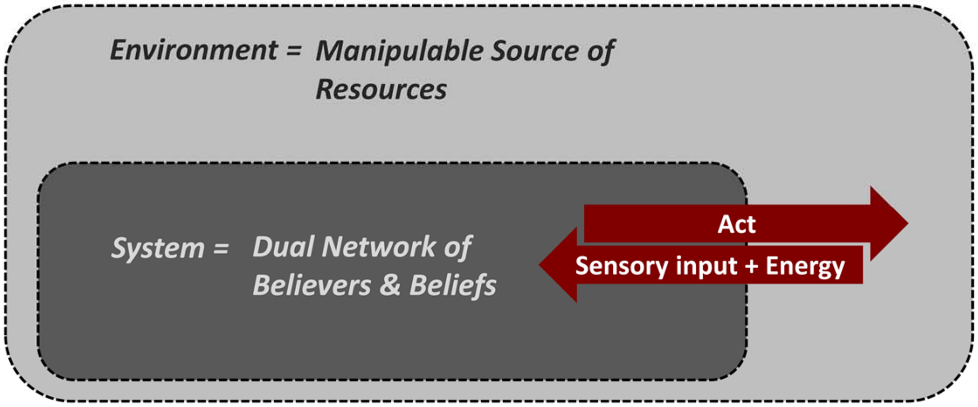

Birth: of new believers and beliefs

Growth: of influence of believers and beliefs

Death: by shrinkage of influence of believers and beliefs

- SYSTEM FRAMEWORK OF SCIENCE. Let's start by asking: what is the overall system framework within which science is done? Here I illustrate one emerging view.
- BELIEVERS AND THEIR BELIEFS. In this view, science is like knowledge in general, it is basically a dual network of believers and their beliefs, adapting to their environment. Believers, mediated by their beliefs, act on their environment; the environment, via sensory input and energy responses, provides a selection pressure on the believers and their beliefs.
- BIRTH, GROWTH, DEATH. This selection pressure mediates life-cycles of birth, growth and death for the believers and their beliefs. New believers and new beliefs are born. Believers and beliefs grow in influence, and ultimately shrink in influence and die.

Dual network of believers & their beliefs

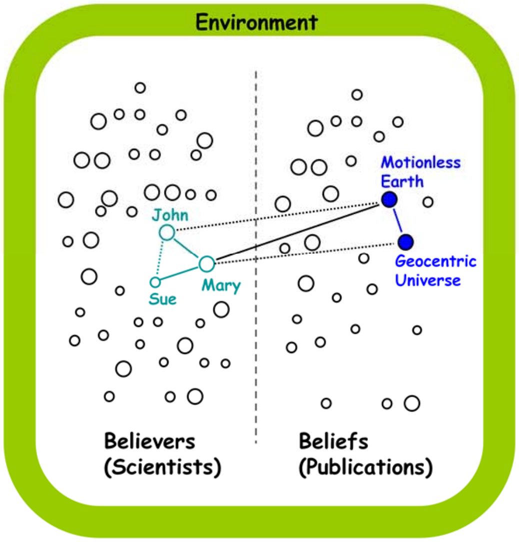

- **Each network "uses" the other**
  - **Believers interact in a social marketplace, and mediate the success or failure of beliefs**
  - **Beliefs interact in an intellectual marketplace, and mediate the success or failure of believers**
- **Links are within and across networks**
  - **Influence (social network) of believers**
  - **Influence (logical network) of beliefs**
  - **Footprints (beliefs onto believers, believers onto beliefs)**
- **The game the individual scientist plays**
  - **Give birth to new ideas: 1 choose good puzzles; 2 solve puzzles convincingly**
  - **Use ideas to build influence: 3 build influence for ideas; 4 build his or her own influence**

- NETWORKS OF BELIEVERS AND BELIEFS. To make this more concrete, we can think of the believers and their beliefs as the two interconnected networks of believers and their beliefs sketched here. The two networks co-evolve, each making use of the other in its co-evolution.
- BELIEVERS NETWORK. The first network, the one on the left, is that of believers interacting in a complex social marketplace, with the hierarchies and spheres of influence that social networks are characterized by. And, as the believers interact, they help mediate the success or failure of beliefs. In other words, believers can be viewed as tools that beliefs use to become successful – the certifiers of beliefs, so to speak.
- BELIEF NETWORK. The second network, the one on the right, is that of beliefs interacting in a complex intellectual marketplace, with the hierarchies and spheres of logical and observational consistency that beliefs are characterized by. And, as the beliefs interact, they help mediate the success or failure of believers. In other words, beliefs can be viewed as tools that believers use to become successful – the certifiers of believers, so to speak.
- LINKS. Importantly, there are links within and across the two networks.
  - There are interactions within the network of believers in the form of a social influence network.
  - There are interactions within the network of beliefs in the form of intellectual interdependencies across beliefs.
  - Then there are interactions across the two networks. The beliefs all have footprints onto the network of believers – each belief has a set of believers that believe in that belief. And the believers all have footprints onto the network of beliefs – each believer has a set of beliefs that the believer believes in.
- MICROSCOPIC PROCESSES. You can see that there are a lot of possible dynamical processes that can occur in this dual network, as believers and beliefs are born, strengthen, weaken, and die, and links between believers and beliefs are born, strengthen, weaken, and die. In fact, there are so many processes that in this talk I'd like to consider just a subset of them – the subset that has to do with how an individual scientist, already born, already trained, actively does science. From the perspective of that scientist, we can think of the scientist as playing a game, a game in which the scientist has to take four steps in series in order to score.
- 1&2: BIRTH OF IDEAS. The first two steps have to do with giving birth to new ideas. The scientist has to choose good puzzles, and then he or she has to solve them convincingly. Kind of like John giving birth to an idea like the motionless earth.
- 3&4: DIFFUSION OF IDEAS AND GROWTH OF INFLUENCE. The second two steps have to do with using those new ideas to build influence both within the world of ideas and within the world of scientists. That'd be like John's idea of a motionless earth giving rise to or supporting a new and subsequent idea of a geocentric universe, as well as diffusing to Mary and Sue. Finally, the scientist uses the influence of those ideas to build his or her own influence in the world of scientists, with John now having more influence over Mary and Sue than he had before he had developed his new idea of a motionless earth.
- THE BUSINESS OF SCIENCE. Now, I don't want to imply that every scientist actively does all four of these steps. In fact, at some level it seems a bit crass, like the only thing that scientists care about is fame. Of course, many scientists care just as much or more about the search for truth, just as athletes playing a sport might care just as much or more about the beauty of their performance as about actually winning. But, in the business that science has turned into, scientists must pay attention to winning and not just to truth just as, in sports, athletes must pay attention to winning and not just to the beauty of their performance.
- REMAINDER OF TALK. So, in the rest of this talk, I'd like to spend a few minutes discussing each of these steps that are associated with "winning" in the business of science. I'll discuss some of the things we know about these steps, but discuss much more the things we don't know but if we did would enable scientists to improve their game.

1/4: Choose good puzzles

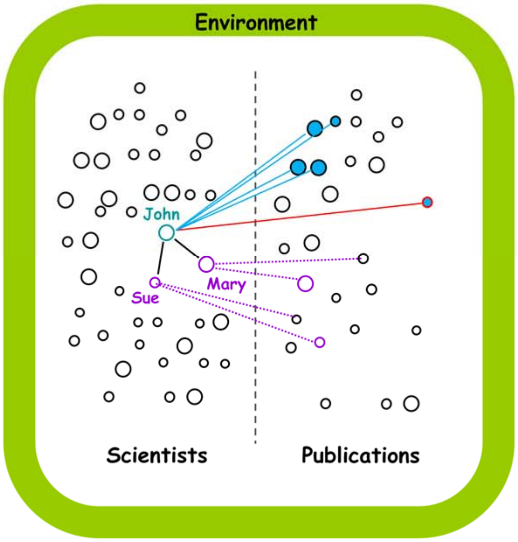

- **What is a good puzzle?**
  1. **Solvable (avoid Merton's trap)**
  2. **Adds verisimilitude not just truth (Popper)**
  3. **Useful and puzzle-generating (Kuhn)**

A Coming up with puzzles

- **Is this part blind (Campbell's "nothing about creativity can be comprehended except in the light of blind-variation and selective retention" and Simonton's "quality stems from quantity")?**
- **Can it be accelerated through cognitive or social tools?**

B Selectively retaining the good puzzles

- **Can goodness be quantified?**
- **Or is it a matter of scientific "taste"?**

- CHOOSE GOOD PUZZLES. Let's start with the first of these four steps: choosing good puzzles. The idea here is that not every puzzle is the same: some puzzles are better than other puzzles. Indeed, a few of the characteristics of good puzzles that sociologists and philosophers of science have talked about are listed here.
  - IS SOLVEABLE. First, the puzzle should be solvable using the tools and knowledge that the scientist has available to him or her. In other words, the puzzle should avoid what I like to call Merton's trap – puzzles that might be very interesting and important if they *could* be solved, but actually can't be solved and hence aren't useful for the advance of science at all.
  - ADDS VERISIMILITUDE. Second, the puzzle should be one that, if solved, adds something that one might call verisimilitude. This is Karl Popper's choice of phrase, and it means that the solution to the puzzle shouldn't just lead to predictions that are true, but to predictions that are specific and unlikely.
  - IS USEFUL. Third, the puzzle should, if solved, be useful in that it generates more puzzles for other scientists to tackle. As Thomas Kuhn would say, it is better to open than to close doors for others.
- COMING UP WITH GOOD PUZZLES. Of course, it's easy to say that a good scientist comes up with good puzzles. But how does he or she do that? Well, one way of thinking about the process is to say that it has two parts: A coming up with puzzles, good or bad, then B selectively retaining the good ones.
- A COMING UP WITH PUZZLES.
  - And for part A, coming up with puzzles, is it possible that this part of the process is blind? As Donald Campbell says, "nothing about creativity can be comprehended except in the light of blind-variation and selective retention," with Dean Simonton's corollary that, since you can't tell which puzzles will be good or not, you just have to come up with a lot of them – in other words, quality ultimately stems from quantity.
  - But, even if it is blind, can it perhaps be accelerated? For example, if John has a number of ideas in his head already (like these blue circles), can he use cognitive tools (like analogy, metaphor and visualization) to come up with new ideas (like this red circle)? Or, e.g., if John's colleagues Mary and Sue have additional ideas in their heads, can John use social tools (like group brainstorming and crowd sourcing) to blend those ideas with his to come up with new ideas?
- B SELECTIVELY RETAINING THE GOOD PUZZLES. And for part B, how do we, after we've come up with a puzzle, then selectively retain just the good ones? Can the goodness of a puzzle be quantified through bibliometric, lexical, or Bayesian measures? Or is the goodness of a puzzle simply a matter of scientific taste, cultivated through experience?

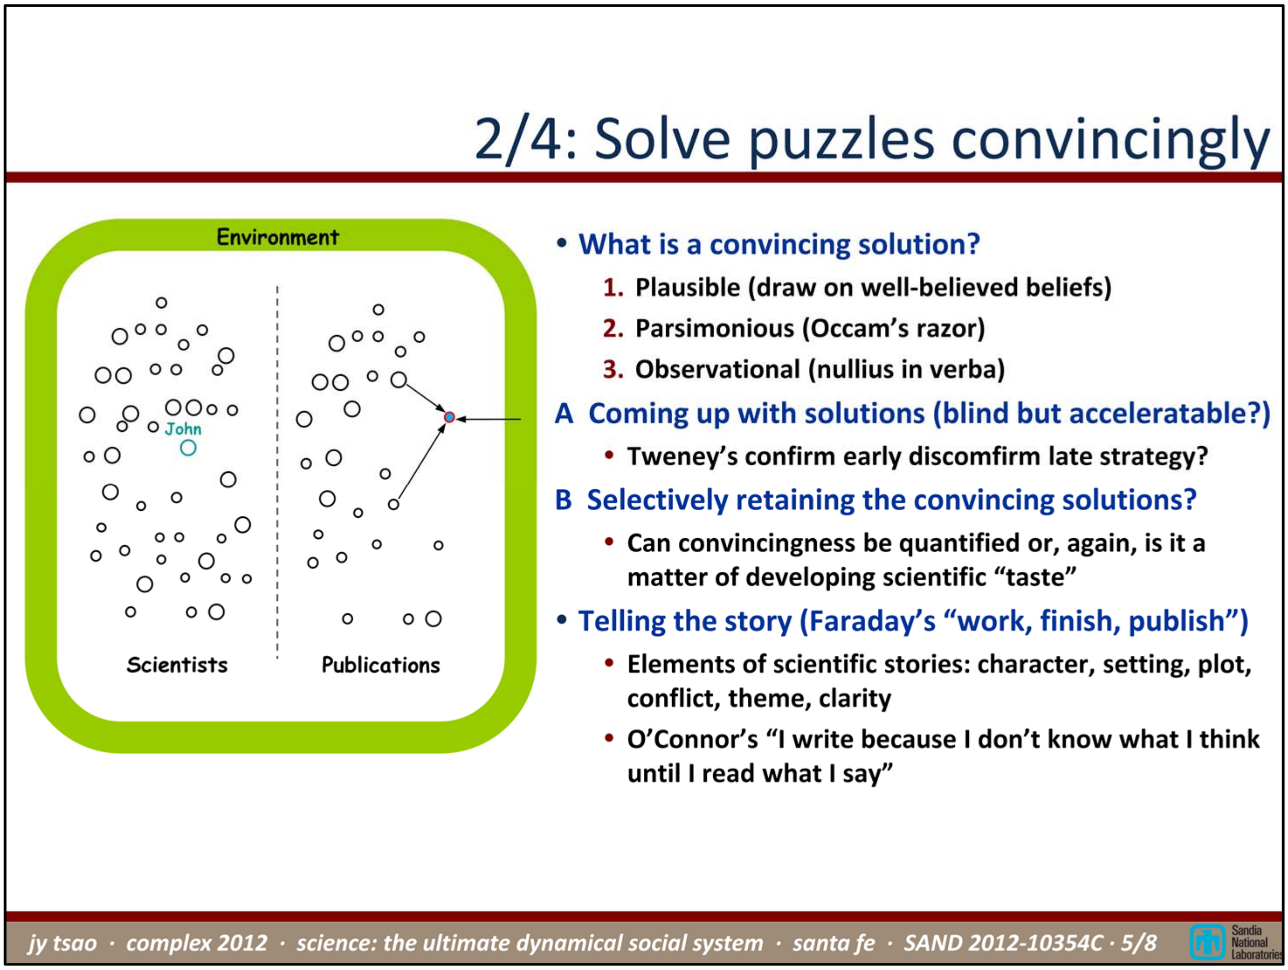

- SOLVE PUZZLES CONVINCINGLY. OK, now suppose you've taken the first step, you've come up with a good puzzle. The second step the scientist has to take is to solve the puzzle convincingly. In other words, it's one thing to come up with a good puzzle, it's another thing to solve it convincingly.
- CONVINCING SOLUTIONS. But what do we mean by a convincing solution? Well, a few possible characteristics are listed here.
  - PLAUSIBILITY. One characteristic is that the solution be plausible, that is, that it be based on ideas that are already accepted by the scientific community. In other words, this new idea here in the red circle should follow logically from other ideas, like these two blue circles, and the more accepted those blue ideas are the more plausible this new red idea will be.
  - PARSIMONIOUS. Another characteristic is that the solution be parsimonious, as in Occam's razor – the simpler the solution the more convincing it will be.
  - OBSERVATIONAL. And yet another characterstic is that the solution be as much as possible based on experiment or observation, as this blue idea is through this link from the environment. The tighter the tie to observation the better. This is of course the famous Royal Society admonition "nullius in verba" – don't trust words. It's also the famous Richard Feynman admonition "the test of all knowledge is experiment."
- COMING UP WITH CONVINCING SOLUTIONS. Just like it was easy in step 1 to say that good scientists come up with good puzzles, it's of course easy in step 2 to say that good scientists also come up with convincing solutions. But, again, how does he or she do that? Can we say again that there are two parts to the process: A coming up with solutions, then B selectively retaining the most convincing ones?
  - COMING UP WITH PUZZLES. And can we say that part A, coming up with solutions, is also blind, but perhaps can also be accelerated? For example, knowing that we as humans suffer from positive confirmation bias, are there ways of thinking that can help us either use this bias when it is useful, or overcome it when it is not useful? For example, in the early stages of idea exploration there are many ways in which experiments haven't yet been controlled properly and may disagree with an idea because the experiment wasn't done correctly, not because the idea is incorrect, so in these early stages one might want to enhance positive confirmation bias. But, later, after experiments have come under control and now we are probably overly wedded to our idea, might one want to enhance a disconfirmation bias?
  - SELECTIVELY RETAINING CONVINCING SOLUTIONS. And after we've come up with solutions, how do we judge and then selectively retain just the most convincing ones? Just as with the goodness of the puzzle, can the convincingeness of the solution to a puzzle be quantified? Or is the convincingness of the solution to a puzzle simply a matter of scientific taste?
- PUBLISH. Finally, once the puzzle has been solved, for the scientist playing the game of science, publishing is where the rubber hits the road. As Michael Faraday said, "work, finish, publish," not just "work, finish."
  - THE STORY. And here is where there's perhaps so much room for improvement in the art. Here is where all the elements that we know make a difference in ordinary story-telling (character, setting, plot, conflict, theme, clarity) could make a difference in scientific story-telling, but rarely do.
  - WRITING AS PART OF THE PUZZLE-SOLVING PROCESS. And writing of course plays another role – it can be part of the puzzle-solving process itself. As Flannery O'Connor said, "I write because I don't know what I think until I read what I say." It's important to begin writing before the puzzle is solved, so that the writing can help solve the puzzle, but by how much before?

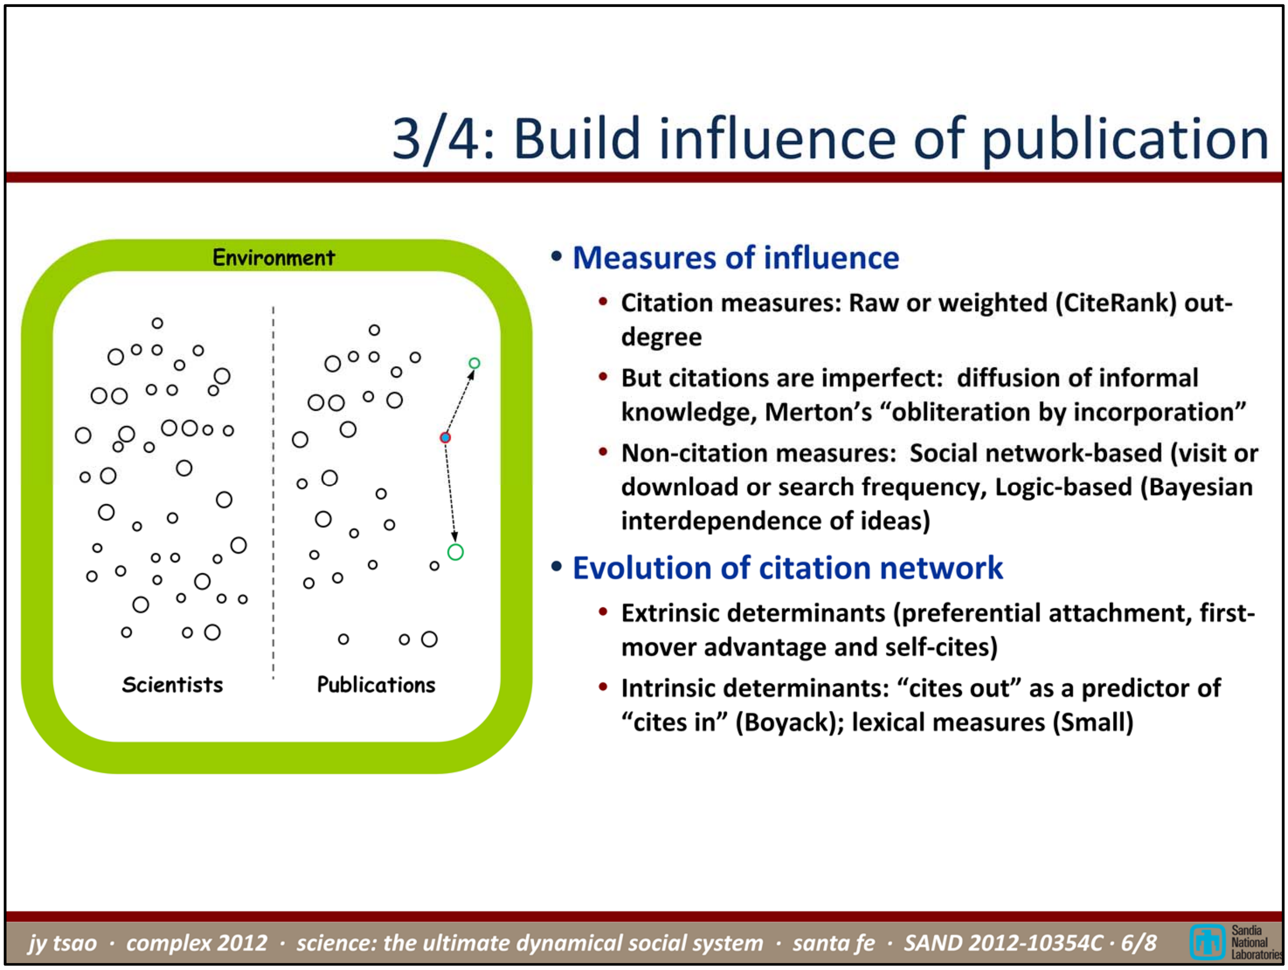

- BUILD INFLUENCE OF IDEAS. OK, now suppose you've, first, come up with a good puzzle, and, second, you've solved it convincingly and you've published your scientific story. Now you've got a new idea that no one, or hardly anyone, knows about yet. But only if this idea of yours spawns other ideas, like this red circle has spawned these two green circles, and becomes embedded in the world of ideas, can one say that the idea has been successful. So the third step is to build influence for the idea.
- INFLUENCE. But what do we mean by influence? Are there ways to measure influence?
  - CITATION MEASURES. Well, there is of course the simple-minded way: citations. Assuming you've published the idea as a paper, if a subsequent paper with its own ideas cites your paper, than certainly your idea has had some influence. And one could of course go beyond simple citation counts and move to a PageRank or CiteRank method where when one counts citations one weights the count by the influence of the papers doing the citing.
  - CITATIONS ARE IMPERFECT. But the main problem is that scientists don't always cite accurately. Ideas are not always learned about through papers, sometimes they are learned about through informal conversations, sometimes they are learned about through textbooks and have become so much a part of common knowledge that they are no longer cited at all – e.g., Robert Merton's "obliteration by incorporation" phenomenon.
  - NON-CITATION MEASURES. So one would really like measures that take these kinds of non-citation-based influences into account. E.g., could one make use of social networks to understand how an idea is influencing thought even without formal citations? Or could one come up with a logic or Bayesian based measure of how important one idea is to other ideas' abilities to make scientific predictions?
- EVOLUTION OF CITATION NETWORK. Now, suppose one does develop some measure of influence. How do we understand how that influence evolves within the network of ideas?
  - EXTRINSIC DETERMINANTS. We know that there are extrinsic determinants of how influence evolves. If a paper has already been cited a lot, more people will come into contact with the paper and will subsequently cite it in later papers. In other words, preferential attachment, as in all social networks, works. In fact, in some microscopic mechanisms that give rise to preferential attachment, like copying mechanisms, scientists don't even need to read the original paper, it is sufficient that the paper has been cited. And, regardless of the microscopic mechanism, clearly in order to get a first-mover advantage one wants to get as many cites as possible as soon as the paper is published, and of course self-citations are one obvious way of doing this. So maybe these extrinsic determinants of evolution can be manipulated to the scientist's advantage, if he or she wants to.
  - INTRINSIC DETERMINANTS. But, even more important, are there intrinsic determinants of the evolution of influence? That is, shouldn't it matter how good the puzzle is and how convincing the solution is, to how the citations of the final published paper evolves? Every scientist thinks it should, but is there a way to measure this, so that one could give earlier feedback to the steps that give birth to the idea, rather than waiting years or even decades to find out? For example, are the number and type of cites that a paper itself gives correlated with the number and type of cites that it eventually gets (Kevin Boyack, e.g., has some ideas along these lines)? Or are there lexical measures, as in the type of language that a paper uses in its abstract and introduction, that are tip-offs to papers that will eventually become exceptionally influential (Henry Small, e.g., has some ideas along these lines)?

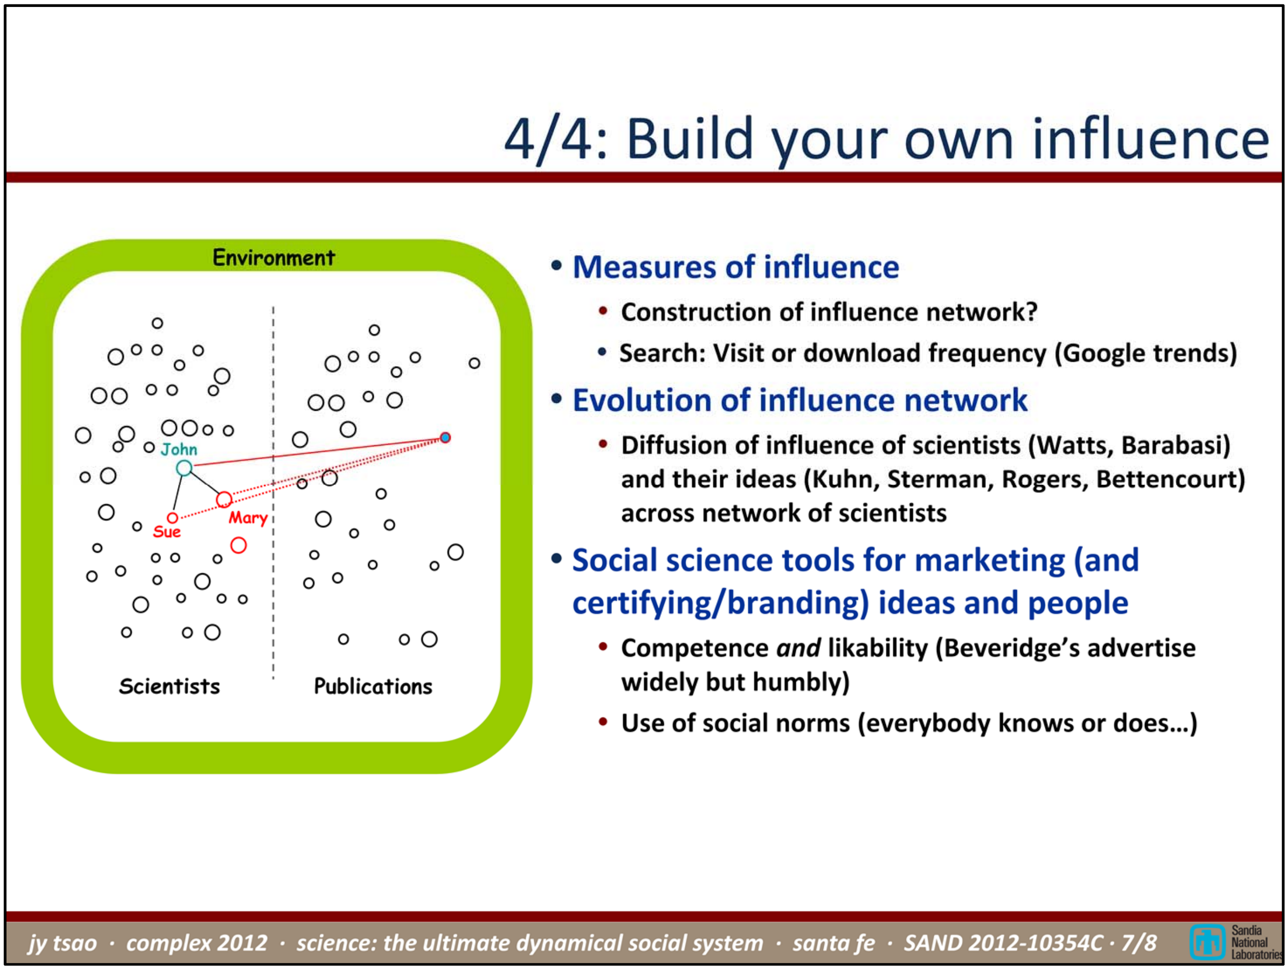

- BUILD YOUR OWN INFLUENCE. OK, now suppose you're on a good track. You've, first, come up with a good puzzle; you've, second solved it convincingly; and you've, third, built up the idea's influence in the world of ideas. Now you have one last step to do, which is to use the idea and its newfound influence in the world of ideas, to build up your *own* influence in the world of other scientists.
- MEASURES OF INFLUENCE. In particular, the measure of a scientist's social influence is not the same as his or her citation influence in the world of ideas. Maybe it's correlated, but we all know scientists with a ton of citations who aren't particularly influential and we all know scientists without a ton of citations who are very influential. The social influence network is influenced by the citation network but it isn't the same as the citation network. So the big question is whether one could construct a quantitative social network that somehow reflects actual influence? It should make use for sure of the citation network, but somehow has to go beyond it. It has to include measures like invited talks, service to professional societies, advisory councils, etc.
- EVOLUTION OF INFLUENCE NETWORK. And then all of the machinery that is being developed for the evolution of citation and other networks could perhaps be applied to the evolution of this scientific influence network. But of course it is more complicated than just the evolution of influence within one network. It is simultaneously the diffusion of the influence of the idea across the network of scientists, as this red idea is diffusing from John to Mary and Sue, and the increase in the influence of the scientist him or herself within the network of scientists, as John has now gained influence over Mary and Sue.
- INFLUENCING THE EVOLUTION OF THE INFLUENCE NETWORK. And then, if we understood better how these influence networks evolve, maybe we could develop social-science-based tools for influencing that evolution. Tools for certifying and then branding ideas and people, some of the tools, e.g., that sophisticated politicians use to build followers.
  - E.g., in politics, one often hears that it is both competence *and* likability that is important. Is it the same for scientists? Or is likability not a factor, it's really just competence?
  - Or, e.g., how susceptible are scientists to social norms and to wanting to "fit in" to their community, so if one can artifically create some buzz around a scientist can that buzz cause other scientists to get excited by that scientist, and then to catapult that scientist into a larger scale popularity? In fact, is this something that could be more pronounced in influence networks composed of scientists than of those composed of non-scientists, because scientists are much more divided into disciplines? If a scientist in one discipline knows that a scientist in another discipline is respected in that discipline, is he or she more likely to simply accept that respect as justified, because he or she is very unlikely to be able to verify that scientist's competence directly?

http://www.smithsonianmag.com/science-nature/The-Top-Ten-Daily-Consequences-of-Having-Evolved.html

Art & Science of the Individual Scientist's Game:

-- Steps 1&2 = Give birth to new ideas

-- Steps 3&4 = Use ideas to build influence

Towards Engineering an Improved System of Science

- THE INDIVIDUAL SCIENTIST'S GAME. So let me close here. I spent most of this talk talking about the art and science of the individual scientist's game. The first two steps have to do with giving birth to new ideas. The second two steps have to do with using those new ideas to build influence.
- WHY IS THIS IMPORTANT? Why do I think understanding these microscopic steps are so important? Because in order to engineer an improved system of science, my belief is that it's important to gain a reductionist understanding of how individual scientists do science. Moreover, the more we understand how scientists do science, the more we can improve how they do science, even if we didn't re-engineer the system of science.
- STAKES ARE HIGH. And of course the stakes are pretty high, if science is indeed key to solving the many problems that humanity faces now and in the future.
- THANK-YOU. Thank-you for your attention.