Science of Enabling Diversity: How to manage diversity to solve the most challenging problems

by
Dr. Norman L Johnson
Chief Scientist
Referentia Systems
norman@santafe.edu

Presented to:

Program Leadership Committee Workforce

**Invited talk for the NIH Workshop on:
"Who's in Your Portfolio?
Building A Stronger Program through Diversity"**

July 8, 2014

National Institute of Health
Bethesda Campus
Bethesda, Maryland

Table of Contents

Slide 1 – Title: Science of Enabling Diversity: How to manage diversity to solve the most challenging problems

I am honored to be here today. Particularly, I'd like to recognize the Program Officers in their mentoring new researchers and applicants. I've overseen staff funded by NIH and have attended your outreach events, and am deeply impressed at the support you provide to new researchers at a challenging time in their careers.

You do this better than any other agency, and to me, it clearly makes a difference in their careers, by helping their confidence and by initiating them on a solid path for later improvement. At other agencies, the approach is like throwing the baby in the water and letting them sink or swim. This mentorship of new voices is a prime example of what we need to do for diversity enablement in general – not for the sake of diversity, but for everyone. With this goal, I'm standing here today.

Slide 2 – Why Am I Here: "Expert"

I'm going to introduce the topic of "enabling diversity" by presenting three themes of why I'm here – to illustrate the three main lessons I want you to leave with today.

Theme #1: Listen to an expert, and they will solve your problems.
This is the common reason why a person stands at a podium before you: the argument is that the speaker is an expert, and therefore you should trust what they say as true and better than what you know. Indeed, providing a person's expert qualifications is almost mandatory–where the messenger is more important than the message. Often, their expertise isn't even in the field of interest. It's enough that they are successful and have many awards.

Use my career as an example: I've had a distinguished technical career in the physical sciences – polymers, physics, combustion, and fluids – always with an emphasis on modeling hard problems. And, what does that have to do with enabling diversity? I do have experience in technical diversity.

After 9/11, I radically switched my technical expertise to help the Los Alamos National Laboratory's biological threat reduction program, under Dr. Gray Resnick – someone you might know. Together, we managed the large biothreat reduction portfolio at Los Alamos. At the same time, I lead technical projects that were game-changers for the nation, as you can read in my bio. I'll highlight two of them later in this talk: the Epidemiological simulation project that was part of the NIH MIDAS program and the National Bio-Risk Assessment project for Presidential Bioterrorism Directive (HSPD-10).

In my last position at Los Alamos, I was acting Deputy Group Leader of the prestigious Theoretical Biology group, a major recipient of NIH funding and where Alan Perelson resides, famous for developing the field of computational immunology.

I also had the honor of working with Dr. David Franz (former USAMRIID Commander), co-chairing a National Academy study on advanced building protection from biological and chemical agents.

All of the above is intended to impress you and convince you that you should listen to me, even though most has nothing to do with what is the main goal of the workshop!

Slide 3 – Why Am I Here: "Expert in the failure of experts"

Theme #2: Collective intelligence. If you believe what Sharon just told you, there is more problem solving power in all of you than there is in the speaker before you. In Theme #2, my goal here is give you a radical understanding that can help you to achieve greater problem solving power in your program.

If you look under the hood of my hard-science career on the prior slide, in my free time I was pursuing a quest of understanding social dynamics and group problem solving. Here is a list of research that has given me the dubious title of "The Expert in the failure of experts." You can read about this adventure in a chapter from a recent book on Collective Intelligence – I reference it on my LinkedIn page.

How did I become the expert on the failure of experts? In the mid 1990s, a group of scientists at Los Alamos investigated the future of the internet, long before it was is was ubiquitous in our lives. We collectively believed that someday individuals using the internet for their own selfish needs would create an emergent capability that could solve the world's most challenging problem – we called it *Symbiotic Intelligence*. You can think of this as the modern version of Adam Smith's 'invisible hand of markets' where selfish acts lead to a global regulation that benefits all. Symbiotic Intelligence is an example of an *emergent property,* a key feature of complex systems.

Our effort was data poor at the time, so after I was lost in Amsterdam, I had a revelation and created a model simulation that showed how information from individuals solving their own problems can solve a problem greater than they can even understand. I'll introduce you to this research later. While I was successful in the demonstration, I didn't understand how it really worked. So like a good scientist, I started looking at what metrics correlated with the remarkable collective performance. What I discovered, very much to my surprise, was that the collective performance correlated with the group's diversity – when diversity is defined a specific way. What was remarkable and disturbing was that if the diversity was decreased in any way - such as selecting a collective of the best performers, the collective performance declined!

This finding was in contradiction to what we believe as the "survival of the fittest" approach to collective improvement: By selecting and reproducing the capabilities of highest performers and eliminating of the lower performers, the group performs better. Scott Page found these same results independently using a different type of model problem about the same time. When both Scott and I tried independently to publish these results in 1998, the papers were rejected. One reviewer said of my work, similar to the reason for rejection that Scott received: *"I don't see what's wrong, but it can't be right."* How do you respond to that? Scott didn't get his results published until 2002. I was luckier and published a summary of my results in a PNAS publication in 1999.

In the early 2000s, the finance world came to Los Alamos to seek me out. This was a major surprise – an excellent example of not knowing the importance of your self-serving research in another area. They brought a challenging problem: they observed in markets that a diverse collective of investors with lower skills and less access to information consistently outperformed the experts. Think about it: the baseline for evaluating the performance of an expert investor is whether or not they outperform the market as a whole. Imagine if your field was like this: the baseline of higher performance was the performance of all the researchers, not just the "best"! Many expert investors or firms can outperform the diverse collective for one year or a few years, but only a few individuals have ever done it for 10 years (e.g., Warren Buffett and Bill Miller). The financial visitors thought I had the answer: how a diverse collective could outperform the experts, and in a general way, I did. I know little about finance, yet in 2001 in front of investors that represented over $2 trillion in investments, I was voted the best of the presentations that featured many big names in finance and business.

I returned to my "normal" work, but a few years later, they visited again when the financial tide turned, because I researching how collectives respond to environmental changes and sometimes are prone to boom-and-bust cycles – you can read my paper and run the simple simulations that illustrated this collective dynamic. Again, the key to preventing the boom-bust cycle and the loss of robustness was retaining diversity: when collectives lose their diversity, in this case by replicating the best solutions and acting like a uniform herd, they lose robustness and are more likely to crash.

During this time I collaborated with an international expert on conflict resolution – Dr. Merle Lefkoff, in her quest to add science to her profession, with the goal of making it more effective. After a year of study, we concluded that social identity was the clarifying but missing understanding in her profession and in the social sciences in general, particularly as to how interacting groups create and sustain conflict. I'll share this with you shortly; it's key to managing diversity.

Recently, I've been studying the requirements for collective intelligence to be work – specifically by expanding the envelope to include collective intelligence of biased groups or individuals, an option the collective intelligence community has excluded.

Now you know why I am the "expert in the failure of experts" and a champion of diversity. While this theme may seem to the why I am standing here, isn't the real reason I'm here.

Slide 4 – Why Am I Here: "It's Tribal"

The previous theme would appear to be right on the mark for why I'm standing in front of you. But in truth, whether or not you're listening to me right now is all about the third theme. *Theme #3: Managing social identity is the key to progress in enabling diversity within and between collectives.*

In the two previous themes, I first appealed to your rational belief that you should listen to me because I'm an expert (theme #1) and then as a champion of diversity (theme #2) – both of these rational arguments. But the truth is that neither of these will make you listen to me, unless you first don't reject me as your part of your tribe or social identity group, your SIG. A fact: you'll summarily reject anything as false when said by an opposing tribe, and you'll accept as truth what your tribal members tell you.

So am I part of your tribe?   Here's my list of tribal attributes – each of these might be a reason to accept or reject me. Which do you react to, pro or con?

Here is an example. In the mid 2000s, after many attempts to share what I learned about collective intelligence, I was getting frustrated about the lack of acceptance of the indisputable facts, particularly by "high performers" in management. I realized that before decision makers can evaluate the truth, they first must be open to hearing the truth. I began to consider if there was a way to repackage the ideas in a way that would not trigger immediate rejection ("I don't see what's wrong, but it can't be right")– particularly from those that believe in the importance of leaders. To this end, I recast collective intelligence as a form of leadership and wrote a paper on the place of collective intelligence within traditional leadership models – and how collective intelligence was a necessary resource for solving the hardest problems.

I'll say more about how important it is to manage social identities and leadership to enable collective intelligence. (The photo is a "selfie", reflected off the mirrored egg in Chicago's Millennial Park.)

Slide 5 – An Example of Diversity Training

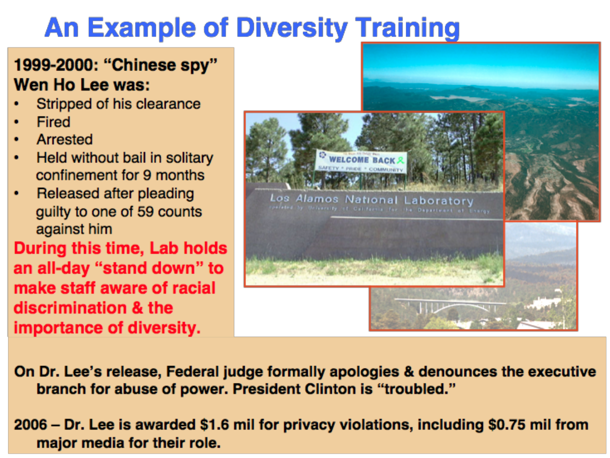

This is a personal example of the three themes. It is the first time I've publically shared this experience I had 14 years ago. It is a dark story on how we treat "others" when threatened, and particularly, how we villainize "others" as individuals, groups, countries and cultures. Most of you remember the persecution of Dr. Wen Ho Lee at Los Alamos.   I knew Wen for almost a decade before this incident happened. And when he lost his security clearance, he came to my group while he was being "investigated," and was there until he was arrested and held without bond for nine months in solitary confinement. My superiors publicly charged him in court with stealing the crown jewels of the Nation. You can read about his story and the final outcome on his Wikipedia page.

Because Dr. Lee was Chinese, the charge of being a spy caused racial discrimination against other Asian staff (cleared and uncleared), students, and visitors. Remember that Los Alamos is a weapons laboratory and practices legal, institutionalized discrimination - if you are foreign national from certain countries, you are restricted to where you can go, must report your location at all times, and who talks to you. This legitimate discrimination was acceptable to everyone, although often inconvenient and at times embarrassing. Hence, the discrimination caused by the spy charges was in addition to this background and escalated until the Laboratory decided to call an all-day stand-down to educate the 9,000 or so staff on illegal discrimination, and interestingly, on the importance of diversity.

Because I represented the three themes described earlier (scientist, diversity champion, and one of the "tribe"), I gave a presentation during the main event to the full laboratory staff – probably the largest live audience I'll ever address. The main point of my talk was the same as today: how diversity is required to solve the hardest problems and how views of diversity need to be changed. I was very nervous talking to my colleagues about research that wasn't "acceptable" or mainstream. Not surprisingly, I received many responses on my talk. For some, my talk was like fingernails on a chalkboard – unpleasant and impossible to accept. For me, their comments were the return of "I don't see what's wrong, but it can't be right." Then, I began to understand something I had been missing previously: tribal rejection.

The stand-down appeared to have worked: the complaints of discrimination declined. About a month after, I was in a review meeting in the Ombudsman office with Asian staff that had experienced discrimination. At first, most noted that things had improved. But as the group spoke, it was clear that the overt discrimination had shifted to being subtle and complex – from black-and-white to gray, particularly about one supervisor in common. I had an "ah-ha" moment: When a problem becomes more challenging, it takes a diverse collective looking at the problem from many different experiences to amplify the weak signals into an obvious truth – similar to what I observed in my research. And because the stand-down focused on changing the staff's behavior based on rational arguments (avoiding illegal actions) but missed the tribal causes of exclusion and discrimination, the discrimination continued.

To give a happy ending to the sad story of Dr. Lee, he was vindicated by an apology from the federal judge that was responsible for much of his extreme treatment, and he received a large cash award many years later for loss of privacy, including from news syndicates that amplified the villainization initiated by the government.

Slide 6 – Three Lessons in Retrospect

These are three themes for the rest of the talk – From my introduction of "why am I here" and from the Dr. Lee incident.

1) The "expert" theme recognizes there are limitations to our rational or objective approach to problems.

2) The "collective" theme introduces an alternative paradigm for solving hard problems.

3) And the "social identity" theme is the hidden key to enabling and managing diversity.

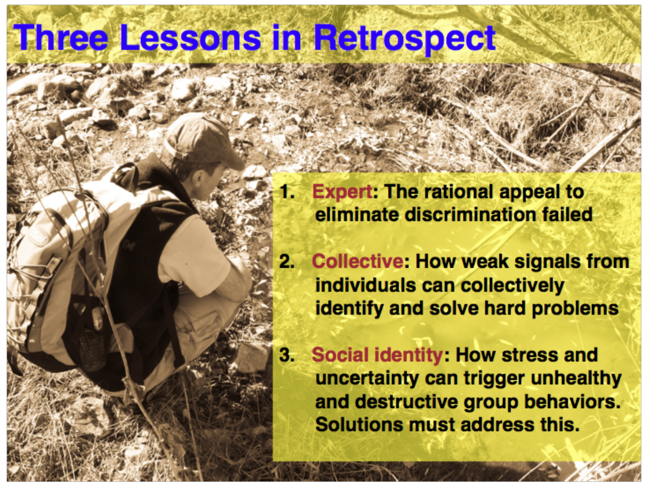

Slide 7– What is the average difficulty level of your program?

I'm now going to ask you questions to better understand why diversity should be of major interest to how you run your program. Let's first explore the *average* level of difficulty of the problems you are trying to solve. This can be technical problems of a program officer or the administrative problems as a manager.

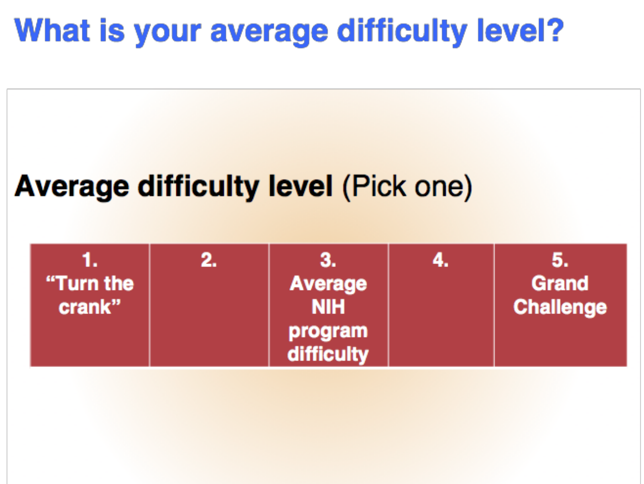

This slide shows the extremes of problem difficulty. The left side is where you know how to solve the problem easily, don't need an expert, and just need to turn the crank. The right side is are grand challenge problems: the risk of failure is high because you and the experts don't agree on how to solve it – for example, the complexity and challenge of obesity in the US. The middle (#3) is normalized to your perception of the average NIH program difficulty. Raise you hand where you think your average difficulty level lies.

The responses of the group: No one raised their hand for “turn the crank”. A few for #2. Maybe 75% picked #3. Maybe 20% picked #4. And two people picked #5 – both where in management (not program officers). I’m surprised: the program officer group doesn’t feel they are taking on the grand challenge problems, but are in the safe space of #3 and #4. Why isn’t NIH taking on grand challenges?

Slide 8– What are your solution approaches?

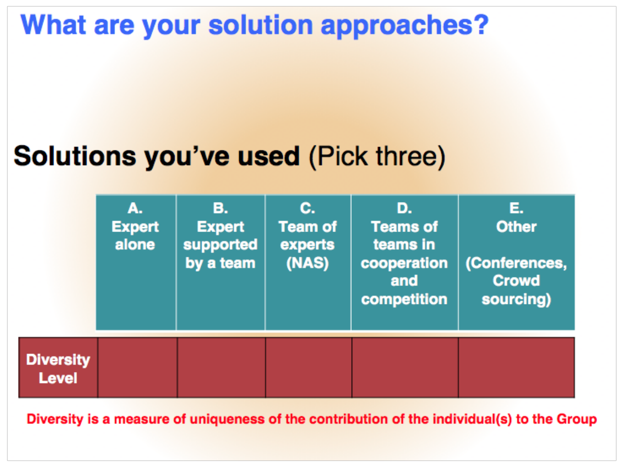

Now that we’ve assessed your average problem difficulty, let’s see what solution methods you use to solve your problems. You can pick three.

The options are:

A. **Funding (or hiring) a known expert to solve your problem** – this can include sole source or the use of a competitive approach where you solicit competing proposals from many experts.
B. **Funding a team lead by an expert,** as a professor with students.
C. **Utilizing a team of experts**, such as used in a typical National Academy study.
D. **Using interacting teams of teams that compete and cooperate** by sharing ideas and best practices.
    a. My personal example of this approach is the NIH NIGMS MIDAS program – with the goal to develop epidemiological simulation resources to support the development national policy and resources for contagious threats.
    b. I can attest to the success of this approach and program: I lead a team from Los Alamos that developed EpiCast. We were the first team worldwide to simulate an epidemic of the entire population of a country, in our case the USA with 300 million people – if you filled out a census form, then you were in our simulations. Never before could realistic, multi-level response strategies be evaluated.
    c. EpiCast and the MIDAS program was a game changer: Our work was used during the pandemic flu challenge to change the previous national epidemic response policy from a focus on quarantine to the rapid development of a vaccine, even a low efficacy one.
E. **Outside the box solutions**, such as using conferences to understand the scope of a problem and crowdsourcing.

A sample of hands from the audience showed that most of the solutions used are A and B as expected (maybe 90%) but a few (handful) use C-E. We could surmise that for the current problem difficulty of the NIH programs from the prior slide, the solution approaches of A and B using experts are optimal.

An interesting discussion would be to inquire from those that said they use the methods D-E (like the NIGMS MIDAS program) as to why they use these alternative methods. My contribution to that discussion would be that the complexity of the problem required the more diverse solutions.

To illustrate how you may already understand the importance of diversity in problem solving – particularly hard problems, let's see how you rank the level of diversity included in the different solution approaches.

Unfortunately, there is no standard definition of diversity – many say "I know it when I see it, but I can't define it." So before you can really answer this question, we need to be specific on a definition of diversity, otherwise we may not be using the same book, let alone be on the same page. Here is my definition of diversity and some observations:

- *A group's diversity is an aggregation of the expression of unique differences of individuals relative to the group.*
- This diversity measure of a group is zero if all individuals have the same information to contribute, even if everyone's information spans all possible options (what normally would be called diversity).
    - This point may seem subtle but it has huge implications when the "truth" is a weak signal held by a few that needs to be aggregated.
    - For example, the people that experienced discrimination at Los Alamos initially all agreed that the discrimination had gone away – they had the same information, but this changed when they started sharing their individual experiences of different activities by their common supervisor. They shifted from a unified herd viewpoint with no diversity to a unique individual perspective.
- The diversity measure is at a maximum if each individual contributes unique information that is not shared by any other individual.
    - This is also significant: Experts usually optimize their information, and in simple systems, experts will all have the same solution, so a group of these experts will have low diversity. Later I'll show you how a group of experts poorly solve harder problems.
    - Similarly, groups that share and coordinate information usually all have the same information, so they will also have low diversity.
- The definition makes a distinction between inherent individual diversity and the expression of the uniqueness of the individual. For example, herd thinking can cause the expression of uniqueness of individuals to be repressed and cause low diversity in the group, even if the inherent diversity is high.

(During the presentation, I didn't capture the audience responses, so the following is what I expected to hear.)

Here are your responses to the question: "Rank from 0 to 5, the degree that each of the methods capture diversity, ranging from zero to some maximum."

A. For a single expert, by definition, the diversity is 0. For competing experts, they have a higher diversity measure – say 1 or 2.
B. Because a team is lead by an expert, the diversity is still bounded, so the diversity is maybe 2-3 with the upper end being the case where the expert uses the diversity of the team to help solve the problem.
C. Evaluating the diversity of a team of experts is problematic:
    - A team of experts may have major differences, or may have minor differences that seem major because their differences are amplified by their posturing.
    - In most fields, particularly on teams of experts solicited by NAS, there usually is general agreement of common knowledge and minor disagreement on the fringe of knowledge. In this case the diversity would be low.
    - So the diversity range might be 1-3.
D. Everyone agreed that a team of teams has a high diversity: say 4-5.
E. Everyone also agreed that alternative solution methods, such as crowdsourcing that does not filter or select any contributions, has high diversity, say 4-5.

So we can conclude that the most used solution methods by NIH are in the low to moderate diversity range. And only a few programs or projects explore the upper extremes of diversity.

The above discussion of diversity may bring up different questions for you. For completeness, I'll share some observations I had to resolve from my research on diversity.

- This may surprise you, but it turns out there are very few quantitative or technical measures of diversity, particularly in the social sciences - maybe because a diversity measure is too controversial? We need to get past this barrier, in order to develop a real science of diversity.
- Diversity is a property of a group, not an individual:
    - An individual is only "diverse" relative to other individuals of a reference group.
    - We often say that someone is diverse, but what we really mean they are "diverse" relative to an implied reference group. This unspecified reference group could lead to problems if the group is different for the speaker and listener. Implied or assumed references often lead to misunderstandings and conflicts.

- There are many ways that a mathematical measure of diversity of a group can be constructed, so why use the diversity measure I defined, instead other possible ones?
    - I looked at many diversity measures (see my 1998 paper) and this was the one that best correlated to group performance for the model problem I'll show you shortly.
    - My diversity measure is similar to the diversity measure used by Scott Page.
    - Just as in any field of science, there needs to be research into diversity measures for many systems, leading to a standard for the discipline. We aren't even close to that place yet.
- Most people associate diversity with some characteristic that is checked on a hiring form. In my research, I don't directly associate diversity with ethnicity, religion, origin, etc.
    - I do believe that these social, cultural, and ethnic groups have unique ways of understanding and living in the world, which in turn produces diverse problem solving approaches and information resources that are not common to other groups. I believe this is not an accident, but necessary for our social evolution.
    - This viewpoint explains observed anomalies in social diversity, such as the large number of incompatible languages in India, much more than could be argued from geographical separation.
    - The creation of human social diversity may be the same reason that bee hives with populations above a threshold will divide to keep the hive population below a maximum, above which the hive can no longer function as a collective.

Because I'm arguing that you need to use diversity to solve harder problems, it is useful to state why the biomedical problems that NIH faces are more difficult than ever before:

- Problems NIH faces are no longer simply technical or medical, but the greater interconnectivity of society, biosciences and health care makes the problems much more complex.
- Our evaluation of the "best" in NIH's expert solution methodologies is complex and often failing. A gauge of this failure is that experts cannot reproduce their high performance or that experts in the same field don't agree, even on the basics.
- Solutions to our grand challenge problems (like obesity) require radically new problem solving methodologies and requirements of outcomes.
    - Ones that engage input and participation from diverse stakeholders.
    - Solution methods must be transparent, essential for diverse stakeholder buy-in and adapting and reusing solutions.
    - Outcomes must not only work, but also be robust and minimize unintended consequences.

The Bottom line:

- The problems facing NIH are becoming more difficult; probably many of ones that need to be solved are Grand Challenges, such as obesity.
- Yet, in this limited sample, NIH is not taking on the grand challenges.
- **Could the reason that grand challenges are not being addressed is that NIH is not embracing methods needed to address them?**

Slide 9– What are the failure modes of each?

Let's crowdsource the failure modes of the methods listed before. I think that intuitively you know a lot about the pros and cons of solution methods available to you, but you've probably never articulated this knowledge. Here's your chance.

What are the failure modes of each?

	Solution Approach				
Failure by:	A. Expert	B. Team lead by an expert	C. Team of experts	D. Teams of teams	E. Crowd sourcing
Isolation					
Herd thinking					
Internal Conflict					
Group Inefficiency					
Bounded ability*					

* Hitting a complexity barrier

What are the failure modes of the different solution methods?

- Note that to be specific for the last column, I've changed "other" to "Crowdsourcing" as the option where anyone can contribute - there are no restrictions for participation. Hence, this last column has the highest diversity.

I've listed five possible failure mechanisms for you to consider:

- **Isolation**: Failure due to not having access to sufficient information or skills to solve the problem. If you had these resources, then you could solve the problem.
- **Herd thinking**: Failure from the group dynamic that represses the expression of diversity of individuals, in the extreme, resulting in everyone having the same contributions. The outcome is the same as for isolation, but occurs from repressing group diversity that limits expression of individual resources, rather than access to diversity.
- **Internal conflict**: Failure to come to a conclusion because of the inability to overcome internal conflicts or disagreements, even though all the necessary resources are present to solve the problem.
- **Group inefficiency in the absence of internal conflicts**: Failure in the dynamics of group problem solving because a solution take too long relative to the time required for a solution, even though no internal

conflicts arise and all the resources are present. This is a common failure mode in meetings.

- **Bounded ability or complexity barrier**: Failure when the individuals or group hit a complexity barrier that can't be passed. Here the focus is on the problem difficulty, rather than the group state or dynamics. It is the flipside of isolation where the problem is less on lack of resources but rather on the difficulty of the problem.

Please raise your hand if you have experienced one of these failure modes for each of the solutions methods. We'll summarize the results in the next slide.

Slide 10– What are the failure modes of each? (My guess)

Here are my guesses at what was we just crowdsourced – Where the greater the number of symbols and darker the box indicate more hands were raised in agreement.

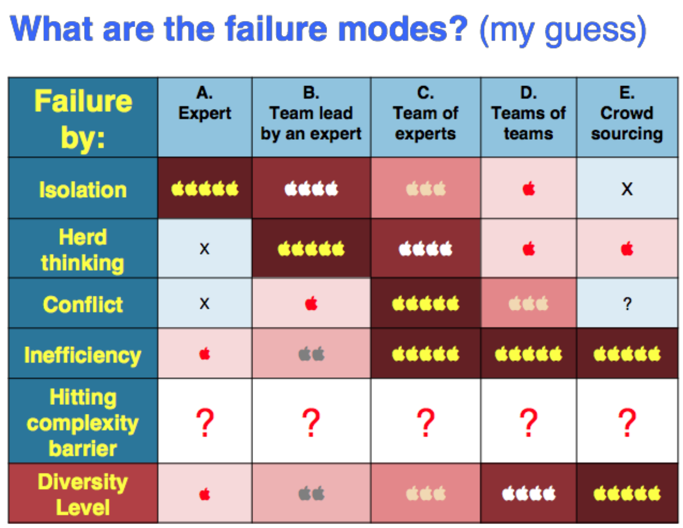

In writing up these notes after my talk, there were only a few surprises from the group that were different than what I guessed.

To the bottom of the figure, I've added the previously determined diversity level so we can correlate the results to diversity.

The reason for doing this exercise is to reinforce that you intuitively understand the tradeoffs between diversity and possible failure modes. Once we make these tradeoffs obvious, then you can use them in your programs.

There are trends worth noting in your responses. ***We note that there are two aspects of diversity being tested: Is diversity expressed and does the process able to handle the level of diversity?***

- **The greater the diversity, the less likely you think the isolation failure mode is activated**. Intuitively, we all understand that to address the failure mode of isolation, we must increase diversity. Relying on an expert to address their isolation is an exercise in failure in itself.
- Similarly, **herd thinking is less likely as diversity is increased**.

- **The number that experienced a failure mode is at a maximum for a team of experts** (many of groaned in agreement).
    - Maybe this reflects you choice of solution methods?
    - Clearly you think that a team of experts, while diverse, is deeply challenged to produce results.
    - It is interesting that the cooperating teams of teams and crowdsourcing methods are viewed to have fewer internal conflicts than teams of experts, presumably because teams moderate conflict internally, and typical crowdsourcing methods–particularly online resources–sidestep internal conflicts. This is the case in advance crowdsourcing methods like prediction markets, which may have major disagreements between the participants, but the process is not slowed by these disagreements. If the same people were put into a typical face-to-face group dynamic, progress would grind to a halt.
- Not surprising, most of you perceive that **as diversity increases, failure due to inefficiency also increases.**
    - This agrees with our experience of diverse groups particularly with traditional methods – they are less efficient because more coordination is required.
    - Crowdsourcing using online methods seem to break this trend, and is the major reason why the scientists at Los Alamos argued that the symbiotic intelligence of the entire planet's population was possible, where a traditional approach to world coordination is unthinkable.
    - The lesson is that new group methodologies can address traditional failure modes. Open Spaces and World Cafe are two excellent examples, which have been used to develop consensus over conflicting groups, for example a community planning effort in Colorado that involved 10,000 residents. Later, I will give you another example.
- We'll delay talking about the failure mode from hitting the complexity barrier, until after we look at how collective intelligence actually works.

Slide 11 – Needle in the Haystack Problem

Hopefully, by now, you are at least open to considering that there may be solution methods that use diversity that you haven't considered. Let's take a step back and talk what is collective intelligence from a technical viewpoint, not just some intuitive belief in the power of the group.

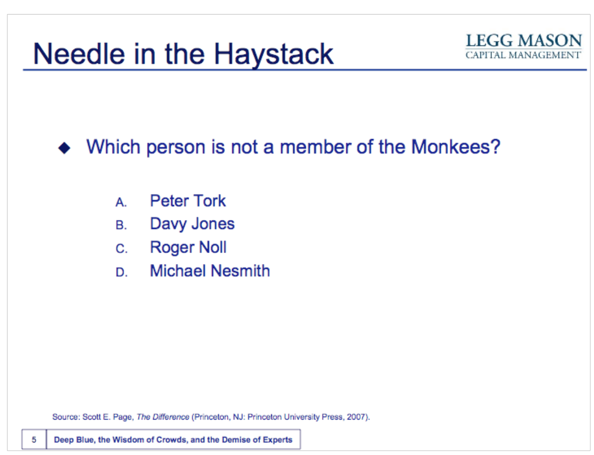

The next four slides are from my colleague and a champion of the wisdom–and failure–of the crowds in finance, Michael Mauboussin.

Question: Raise your hand at which name is the correct answer? If you don't know, then pick a name at random.

You are a VERY unusual crowd – most (70%) of you know of the Monkeys AND know the names of them. Honestly you don't look old enough to know - even I had to look it up to be sure!

Slide 12 – Needle in the Haystack (continued)

Let's suppose that the distribution of the groups knowledge on this question is as in the slide, where only a few of you know the right answer and many have to randomly guess.

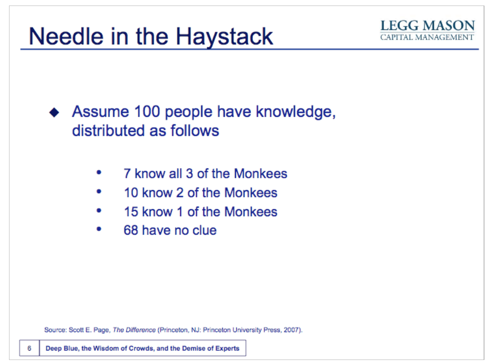

Slide 13 – Needle in the Haystack (continued)

If we repeatedly sampled these 100 people, we'd get that Noll on average would get 34 votes, 14 votes above the 25 votes if no one knew the answer and they picked from the four names randomly. This is remarkable given that only 7 people knew the right answer.

You can begin to **see how "Wisdom of the Crowds" works: A weak signal from the few individuals that knew all or part of the answer was amplified in the presence of random noise from the others to give the "truth".**

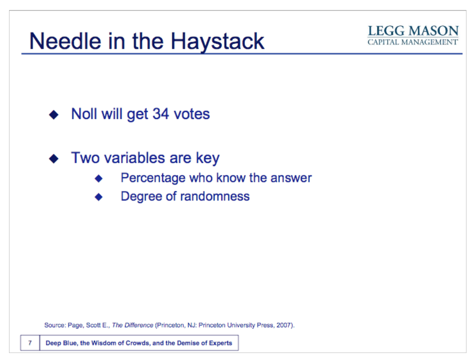

What is the important information here? The percentage that knew part or all of the answer. And the randomness of the other responses.

Notice that I told you to answer randomly if you didn't know. This is a good place to introduce how bias can make the crowdsourcing fail. For example, if you were one of the clueless and a couple of you liked the sound of one name over another, then this would bias the results.

Individual bias can cause the "random" answers not to be random, resulting in a false "weak signal" that can overwhelm the "truth". You can see from this example that if there is bias is towards an incorrect name that would influence 8 votes, then the aggregation will give the wrong result. But you can also see that if the biases were randomly distributed (say there are three groups that each prefer one of the incorrect names), then the wisdom of the crowd will still works.

From this understanding, I'll give you a sneak peak into my current research. I've concluded that individual or group biases are unavoidable and *necessary* in practical solutions of complex problems ("did he really stay that?"). They are unavoidable, because the capturing of "truths" in complex domains are interwoven with biases. They are necessary because the "truth" can not be separated from the bias without losing the truth. For example, objective methodologies that try to extract the truth from a biased source may fail because the weak signal (the truth) is lost in the extraction.

Another profound and disturbing observation is that sometimes unbiased or objective understandings are too complex to be useful, and biased understandings are more useful ("did he really say that?").

Here is an example of complex solutions that most of you know, but never realized the significance of it. There are two common extremes of geometry in use: Euclidian and non-Euclidian – each based on an assumption that is in contradiction with the other: lines remain parallel at infinity or lines intersect at infinity. Both types of geometry are useful in different applications. And a general geometry exists that includes both types of geometry (it's "objective"), but is almost never used because it's too complex.

The implications of this example is profound when working with complex problems: The scientific method is based an objective process, but what if the attempt to be objective is limiting our ability to solve hard problems? When I had this realization, my identity as a scientist was deeply challenged. Can you imagine a field of science that embraces non-objective subjectivity? I can, and I believe it's in our future. I'll let you reflect on this while you are falling asleep tonight.

Slide 14 – Diversity Prediction Theorem

As Sharon shared, Scott Page summarized many of the above conclusions in his *Diversity Prediction Theorem* (Scott is an applied mathematician). In it, you can see that the collective error (not getting Noll correct) decreases as diversity increases. And the baseline for the collective error is the average individual error: If the diversity is zero, then the collective error is the

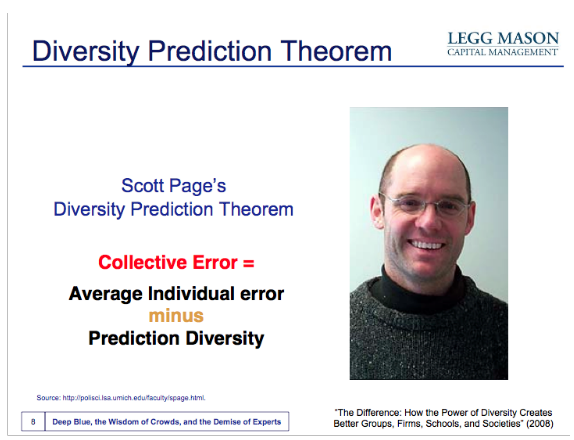

average individual error – as you'd expect.

When I asked this group the Monkey question, the individual error was low (most of you knew the answer), so diversity played a less important role than in the slide example. In harder problems, the individual error increases, and therefore, diversity becomes more important.

The types of problem that Scott studied are ones similar to the prior Monkey's question: finding a correct answer to a question (e.g., who is not a Monkey or how many beans in a jar) or predicting a single outcome (e.g., a political election or the winners of the Academy Awards).

By contrast, most of the world problems we want to solve are not single questions, but a series of interdependent questions to solve a larger problem. In the next couple of slides I'll show you the model problem I studied – a sequential decision problem – and show you that you get the same diversity theorem that Scott found. This is significant, because the "wisdom of the crowd" result appears to be independent of the type of problem! Generality of a research result is the dream of every scientist.

Slide 15 – Find the fewest steps in a complex maze

In my studies in the 1990s, I examined how the solutions of individuals in a maze (a sequential problem with many possible solutions) are combined to solve a global problem – I'll show you how in a moment.

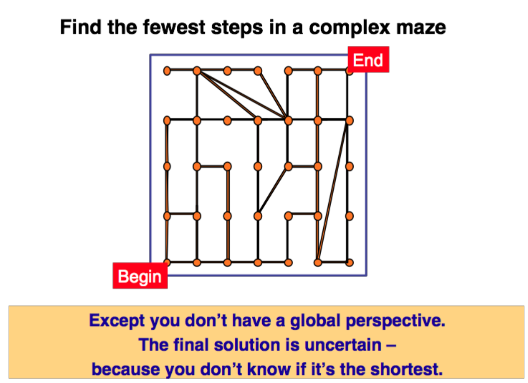

In this simple maze, individual preferences are which path to take at a junction, and all individuals see the same options–they have a common worldview. If we generalize the model problem, many types of information can be used to solve a hard problem: solution algorithms, scientific disciplines, personalities, environments, cultures, ethnicities, etc..

For my model problem, I examined how information learned by independent individuals (meaning they solved the maze alone) is aggregated as a collective at each of the decision points (the orange circles above). Hence, my diversity measure is calculated at each decision point for a group of individual at that point, each with different or similar experiences, and then summed over all decision points. You can see that my diversity metric includes decision points that aren't "on the main decision path." This is actually quite a profound observation: diverse knowledge outside the mainstream solution or paradigm, even in the extremes of knowledge, is important!

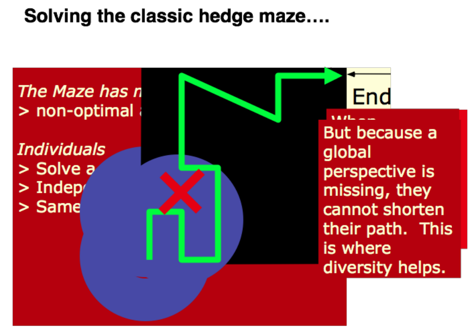

Slide 16 – Solving the classic garden hedge maze

This is a quick introduction to the maze problem I used to show how a diverse collective can solve a problem **both** better than any individuals AND discover a solution that the individual cannot even comprehend, in the lexicon of complexity science, this is an *emergent solution*.

To learn more, see my 1998 paper (the one that was rejected) for details.

I asked the question: *How can the information from a group of independent individuals solve a globally hard problem - without requiring coordination, cooperation or selection between the individuals?* By removing the cooperation stage, the focus is exclusively on how the information is combined, not how the individuals behave as a group. I've found it is a common mindset that the individuals must be cooperating to solve the global problem. They are not - keep telling yourself that.

What I did was to have many individuals first solve the maze - independently and with identical capability, but because they used random choices in a system with many options, they created a diversity of experiences. Therefore, the difference in performance is not due to different abilities, but different outcomes from random choices.

The individuals solve the maze without having a global perspective of the problem (they can't see the whole maze). They only see the local options at their current location. At any time in their exploration, they don't know how close to end they are–it could be around the corner or 100 steps away. Imagine the classic garden maze with high hedges.

(You will need to animate the PowerPoint slides to see the following.)

For example, at the beginning an individual has two choices - not knowing where the goal is, they randomly pick one path. Then at the next junction, they pick from 3 paths, and so on. They lay breadcrumbs on the way in order not to repeat paths that were dead ends or to select new paths. Until they find the goal.

While in the computer simulations I used specific heuristics for the individuals to solve the problem, I've also asked many people to solve the same maze before a public talk hosted by the Santa Fe Institute, and I found the same results that I found in the simulations. So the results are not dependent on the specific heuristics I used. In a separate study, I also examined individuals with different capabilities or heuristics – as another way to generate diversity – and found the same results.

The Red line is an example of a path of one individual from the start to the end. The Green line is the solution when an individual reapplies their learned information, because the individual will cut off unnecessary loops (remember the last time you drove to a new place again, you did not repeat going around the block as you did the first time; instead you optimized your solution based on your learned information).

At the Red X, we see that the individual might have shortened their path, but because individuals does not have a global perspective, they cannot see how to remove the unclosed loop. Filling this gap is the how the diversity of information from different individuals can improve a group's solution, simply based on information. I'll show a more detailed example of this shortly.

Slide 17 – Performance of Collectives

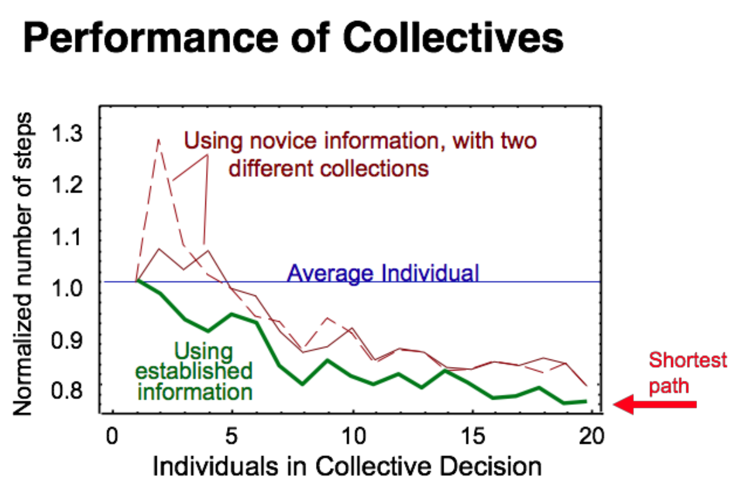

These are the main results from many simulations of the maze problem. Plotted are different the number of steps (normalized by the average individual performance) versus an increasing number of individuals in the collective.

The each curve shows the collective performance for randomly selected collections of individuals, up to 20 individuals. The normalized path length of the individual or group is the actual path length divided by the average normal performance of the all the individuals - 12.8 steps, compared to the shortest path length of 9 steps. For 2 individuals, the normalized number of steps is the collective path as found from the combination of information from the first individual plus a new individual. Then, a new individual is added to those two for a group of three individuals, and so on. The same individual heuristics (rules) are applied to the collective information, so the change in performance of the group over the individual is due only to the change of information, not rules.

Novice information (the top two curves) is when an individual fully contributes their learned information to the collective, including the extra loops that returned the individual to the same node. The *established information* (bottom green curve) is when an individual contributes only their optimal information that eliminates their closed loops. Note that collectives with a fewer individuals (2-4) do much worse than average if they contribute novice information, because the extraneous loops increase individual error, causing the collective error to increase in Scott's Diversity Theorem. But as the number of individuals increases, the collective error decreases because diversity increases, while the individual error stays constant. If the individual contributes their established information (bottom curve), then the diversity factor trumps the individual error from the beginning and the collective error decreases. One way to view the increasing collective performance is that as the number of individuals increase, the diversity increases while the individual error is constant, and the collective solution improves until it reaches the minimum path length!

The remarkable observation is that as more individuals are included in the problem, the collective solutions approaches an optimal, shortest-path solution, even though there is no concept of a shortest path in the individual (a concept that requires a global perspective). So, the collective robustly finds a global solution that is not even "understandable" by the individual. The ability of the collective to find an "emergent" solution (one that is not understandable by the individual) is what took me years to understand, after I first got these results. A mathematician colleague studied what types of maze (graphs) produce this result. His simple answer is a graph with multiple, interconnected paths, essentially very much like the path through life and a typical complex problem.

In Scott's Monkeys problem, we saw clearly how weak signals of individuals are aggregated to achieve a collective "truth". How does the collective improvement occur this maze, a much more complex problem?

Slide 18 – How collectives find the shortest path

This is a simple example of how the diverse collective improves the individual solution. Remember, because of the lack of global perspective, an individual can't eliminate unclosed loops in their path. But, as in this example, when all the paths are combined, the collective information can eliminate the unclosed excursions of other individuals. You can imagine that as the information of more and more individuals are combined, then all possible unclosed paths are

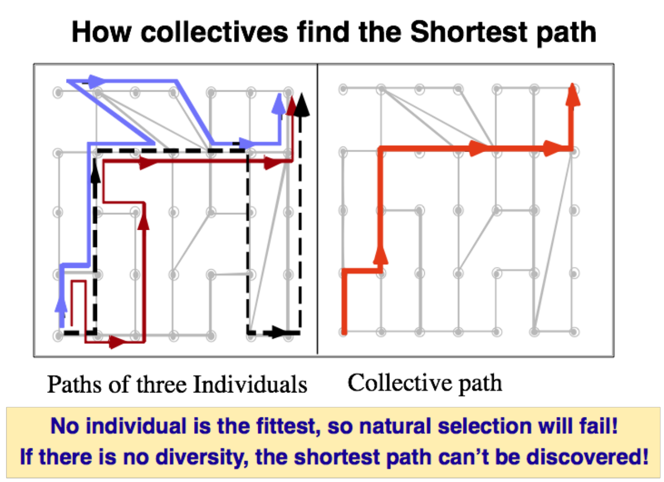

eliminated and the shortest path is found.

This result is remarkable because in this example no individual takes the shortest path. But the collective path is shorter than any individual path and is almost the shortest path. This might be very disturbing to you. It shows why "survival of the fittest" or "reward the best" is a failed concept for group optimization in complex problems. Because a shortest path is not embodied in any individual in this example, there is no one individual that is the best performer that can be selected to lead the others to improvement. In this example, it is the collective as a whole that is the optimal leader. In addition, we can also see in this example that if any one individual was eliminated from contributing, even the poorest performer, then the collective solution would be degraded!

Said another way, from a global perspective, what makes the worse performer have a long path doesn't affect the collective solution for larger numbers of individuals, but what makes the worse performer better (their weak signal of the truth) does improve the collective solution. This is another interpretation of Scott's Diversity Theorem: the random answers don't degrade the collective solution, but the correct, even weak, truths do. The only difference in the two types of problems with respect to Scott's Diversity Theorem is that randomness in the answers to single questions are replaced by unclosed extraneous loops in the maze problem.

These observations are based on a global perspective of the problem – we can see the shortest paths, and understand them. But there is a deeper consideration hinted at in the prior slide: If the individuals and collective don't have ability to understand a global viewpoint (nothing in the rules have a global perspective), the concept of a shortest path is not understandable. For example, if five individuals compared their shortest path and they found one to be shorter than everyone else, they don't know if it is the shortest path possible. Yet, the collective robustly finds the shortest path, even though that is not the expressed goal of the individual or collective! It is for this reason that the group of scientists as Los Alamos thought that a global collective might solve the hardest problems, even ones an individual can't understand. The full implications of this *emergent problem definition* and solutions in organizations and societies is yet to be fully understood and exploited, but you may intuitively appreciate the importance of such a insight in order to solving the hardest of the world's problems.

The other major observation in this example is that if all the individuals take the same path (zero diversity), then the collective can do no better than the individual performance (the collective error is identical to the individual error when the diversity is zero). This is how ants use diversity in path exploration to find the shortest path between your dropped potato salad and their nest, even if the path is remarkable complex and no single ant actually can see or understand the shortest path. You might pause a second and be open to the possibly that you are doing the same in society: you are currently contributing to a problem solving capability far beyond your understanding – this is what we called the *Symbiotic Intelligence*.

Slide 19 – Diversity Prediction Theorem (again)

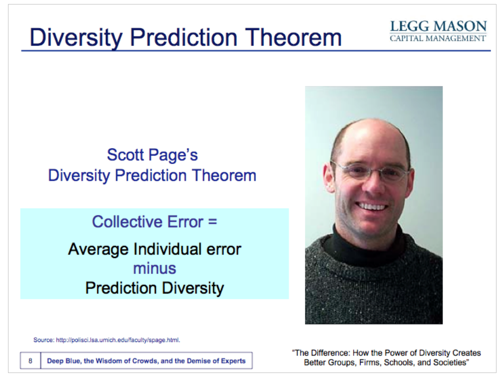

So what we find is that for a very different type of model problem and one that is more complex than a single prediction or question, we find that Scott's diversity theorem is valid as a general concept.

This is deep science: when you find the same "rules" in different types of systems, you know that what you found is a powerful explanation of the world around you.

Slide 20 – Utility of Performance versus Problem Complexity

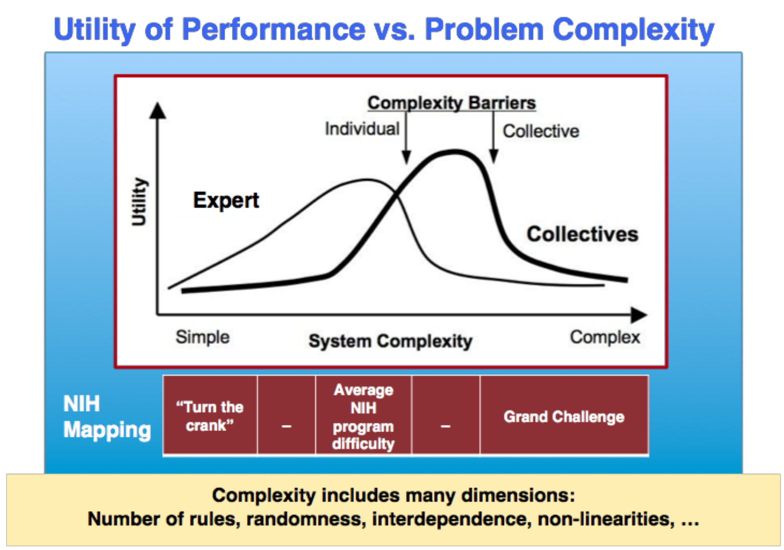

Let's recast the above results in a way that speaks to how we actually solve problems.

The figure above shows two generic plots of the utility (not of performance) of expert and the collective, plotted against problem complexity. Note that complexity can include many dimensions – I'll show you two dimensions in a bit. This graph is a generic (a cartoon), which can be made specific for a certain dimension of complexity, but all graphs will have the same trends.

What to observe from the expert utility curve:

- If everyone can solve the simple problem (left side), then the expert or high performer has no utility over everyone else, so the expert utility is low.
- But as the complexity increases, the expert has an advantage over other individuals, so the expert utility increases.

- At some point the expert performance must decline as the problem becomes too complex for them, at the *expert complexity barrier*.

What to observe from the collective utility curve:

- If the individual/expert can solve a moderately complex problem, then the collective has lower utility because the generally higher efficiency of the expert over the collective (no coordination issues).
- When the expert fails, the collective has the highest utility based on our earlier discussion, the Wisdom of the Crowds.
- As the complexity further increases, the collective begins to fail because the individual error increased faster than the effect of increasing diversity. Once the problem is too complex for the individual, resulting in high individual error, no amount of collective diversity will help. Stated alternatively, because there is no weak signal from the individuals for the collective to amplify, the collective cannot improve over the individual. I observe this collective decline in my simulations when I make the maze more complex or when the individuals only use clueless heuristics, such as random walk, to solve the problem. You can see examples of this in my 1998 paper.

I've added the NIH problem difficulty to the slide to get you to think about how you might start using diverse collectives to solve more challenging problems at NIH.

Slide 21– Robustness of Diverse Collective

Hopefully I've convinced you, or at least made you open to the possibility, that a diverse collective for certain problems can perform better than an expert. But what good is the promise of a high performing collective, if it regularly experience catastrophic failures, as in our financial markets?

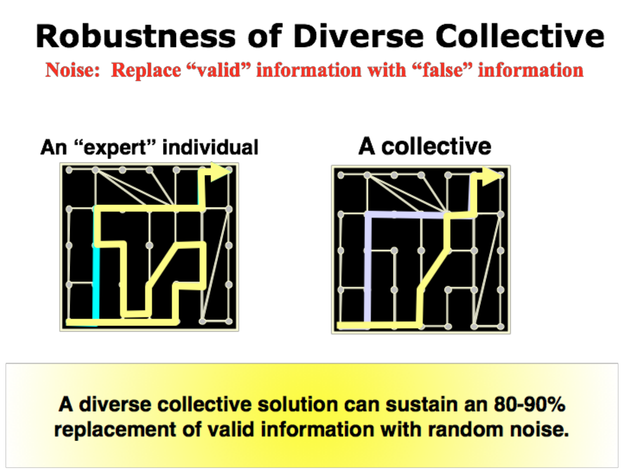

Here is an example of the results of a robustness study where I looked at the effect of noise on the performance of experts and collectives. To introduce noise, I replaced valid individual preferences with random information - wiping out real preferences and creating new preferences that lead nowhere, and then redid the collective problem solving with these noisy individual contributions. What I found was that for experts (left side of the figure), because they optimize their performance, they are very sensitive to noise in a complex domain (in a simple domain there noise doesn't a severe cause a problem, because options are limited). The impact to expert performance by noise is because they have no contingency experience to recover when the noise that sends them to unexplored portions of the maze. Consequently, they must go back to relearning the problem. It's not pretty when an expert stumbles.

By contrast, a diverse collective (right side) has lots of contingency information and doesn't even stumble when experiencing noise. In fact, I found that a large, diverse collective of individuals with 30% noise had little effect on the collective performance and with 70% could still find the shortest path with larger numbers of indivdiuals! This is an example in the diversity theorem where the increased individual error caused by the noise it trumped ultimately by the collective diversity.

A corollary to this robustness finding is for individuals with biases: if individual biases within a group are present but uncorrelated with each other (bias looks like noise), the diverse group will still find the shortest path. Biases are more than differences in preferences. Difference in preferences but not biases are where I like chocolate ice cream and you like vanilla, but we both like ice cream. Biases are a disagreement on options, where I don't think ice cream is an acceptable food but you do. Expressed in terms of the maze problem, a bias is where one person sees a doorway or path and another sees a blank wall–no option.

Hence, we conclude that when a diverse collective solves a hard problem, not only do they solve the problem better, but also the diverse collective solution is more robust than the expert! This is in contrast to most expert optimized solutions, where in order to get robustness, performance must be compromised, because optimal performance, such as an assembly line, leads to lack of robustness, typically because of loss of diversity. For example, companies that diversify their products to become more robust, also often reduce their performance because they have fewer bestsellers. In this example with diverse collectives, we can get both optimal performance and high robustness!

Slide 22 – Collectives in complex environments

Up until now, I've used a model problem where individuals start at the same beginning and end– basically representing an agreement on the problem definition and the goals.

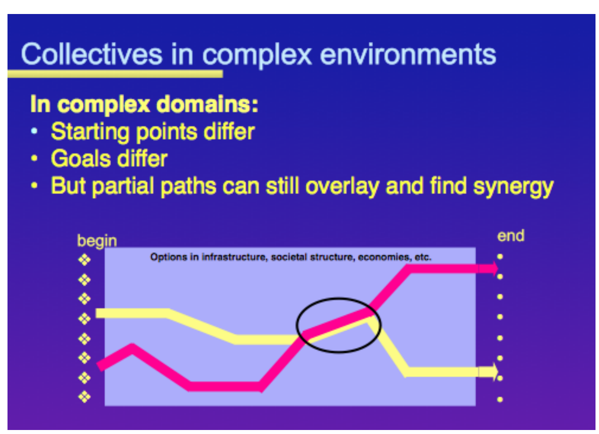

This type of problem is relevant to tight teams, but in the real world, particularly for complex environments, individuals have different beginnings and endings, even within organizations. This cartoon shows how individuals in these complex environments can still contribute locally to a global problem solution when their paths overlap. You might consider how this occurs in your expertise: your technical area might overlap with completely different discipline, because you both use cellular biology to understand your problems. Hence, a discovery in one discipline can benefit another discipline.

This figure also captures what we observed earlier in the mechanism for collective performance: An individual's "truth" or weak signal benefits the collective, but extraneous loops don't degrade the collective solution. In fact, in more complex problems with different goals, the extra loops for one goal may be the weak signal for another goal!

Slide 23 – Michael Mauboussin: What is the metric for identifying a Grand Challenge?

Because of a lively discussion at dinner last night with the PLC staff and speakers, I put this slide back in to show you today. The question being discussed was an important one: how do you know that a problem is a grand challenge, and worthy of a solution using diverse collectives.

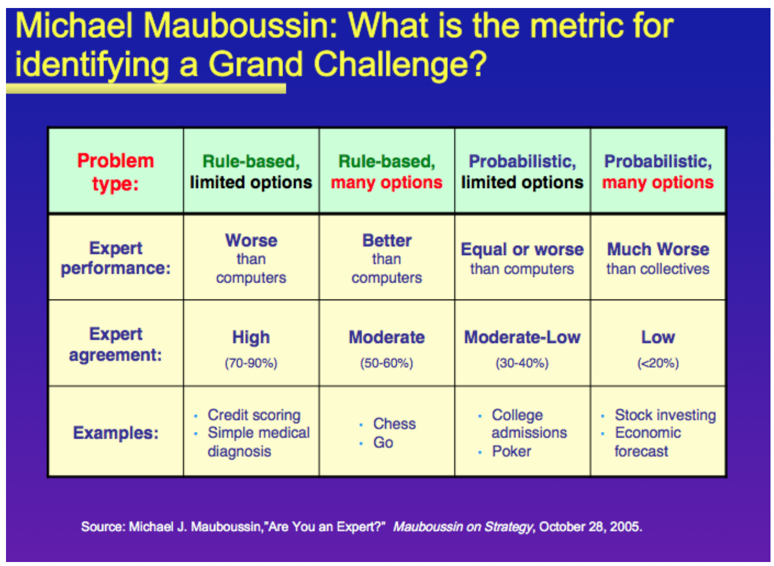
Michael Mauboussin: What is the metric for identifying a Grand Challenge?

Problem type:	Rule-based, limited options	Rule-based, many options	Probabilistic, limited options	Probabilistic, many options
Expert performance:	Worse than computers	Better than computers	Equal or worse than computers	Much Worse than collectives
Expert agreement:	High (70-90%)	Moderate (50-60%)	Moderate-Low (30-40%)	Low (<20%)
Examples:	• Credit scoring • Simple medical diagnosis	• Chess • Go	• College admissions • Poker	• Stock investing • Economic forecast

Source: Michael J. Mauboussin,"Are You an Expert?" *Mauboussin on Strategy*, October 28, 2005.

Michael Mauboussin (remember the collective intelligence champion from the finance world, I mentioned earlier) captured in this slide his analysis of how well experts perform relative to non-expert systems: computers and collectives, for different types of problems. He looked at two dimensions of complexity: how many rules and how much randomness. Both of these measures of complexity could be with few decision options or many. As you might expect from our prior discussions, experts do best in systems that are rule based but have many options, as the games of chess and go. Computers do well with rule-based problems, but limited options, such as simple medical diagnosis and many scientific calculations. And collectives do best for probabilistic problems with many options, as in stock investing and economic forecasting.

The main reason for including this slide is the third row that answers the questions: How do you tell if a problem is a grand challenge and appropriate for solution by a diverse collective? Michael's conclusion is that the more experts disagree, the more you need to enable a collective solution to address your grand challenge. For example, the experts I've talked to about obesity in the US have widely (wildly?) different understandings of and solutions to the problem. In my mind this is a grand challenge area. Can you think of examples within your area of expertise where expert disagreement occurs? Are they grand challenges?

Slide 24 – Collective Research Summary – information arguments alone

(The background photo is a new friend from Chicago that grew up in the city and never saw deer in the wild.)

This is a summary of what we've learned about collective solutions so far.

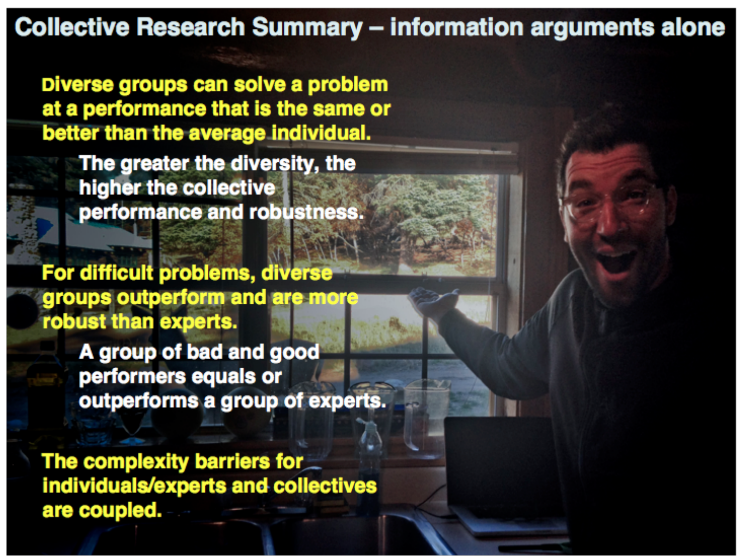

Diverse groups can solve a problem at a performance level that is the same or better than the average individual. And, the greater the diversity, the higher the collective performance and greater the robustness. For really difficult problems, diverse groups outperform and are more robust than experts. Finally the complexity barriers for experts and collectives are coupled: collectives cannot perform with "dumb" members.

I've added one result that both Scott and I independently found that may cause you cognitive dissidence: Diversity of individuals also includes differences in performance!

I found that if I selected the best individuals performers to contribute to a collective solution (similar to what we do all the time in teams), these expert collectives, from a purely informational basis, performed worse than a collective of all individuals, with good and bad performers. The reason for this is complex, but essentially groups of high performers have lower experiential diversity in harder problems (fewer unclosed loops), so the diversity theorem says that because they have lower diversity, they will perform worse!

The fact that both Scott and I observed lower performance of collective of experts in two very different types of problems again suggest this is a general feature of problem solving with diverse collectives. I note that because these results are based on information diversity alone, teams of experts also have challenges to perform well because of conflicts of ego and posturing (failure due to conflict), While these failure mode are not included in either results, differences in diverse groups could lead to conflict in many traditional group decision methods.

This last result may seem contradictory to what was stated earlier, where I noted that there is coupling between the collective and individual performance. If the individual doesn't perform well (high individual error), then the collective will not perform well. So why isn't the opposite be true: If I choose a team of higher performers with lower individual error, then the collective should perform better but Scott's diversity theorem. The resolution of this apparent contradiction is that the collective diversity decreases more than the individual error increases, so the net result is that the collective error increases.

Said another way, while the individuals must have some ability for the collective advantage to be expressed, in complex problems there is a greater need for exploring the domain (increasing diversity) than there is for higher individual performance. If the domain complexity is lower, then the group of experts will perform equal to the diverse performers.

Slide 25 – The conflict between the expert and diversity-enabled solutions?

If you are feeling cognitive dissidence at this point, you are not alone. In the decades before Scott's and my work, the prevailing view across every discipline that I examined is that when diversity is added to an already functioning group, the presence of increased diversity is destabilizing, and collective performance of the group is reduced. This even includes your brain, as observed in cognitive science!

Furthermore, our human resource and personnel management is also based on hiring the best, rewarding the high performers, and eliminating the poor performers. Now I just told you that you need lower performing individuals to make the collective perform better. It is no wonder that Scott's and my papers on diversity were summarily rejected in 1998. The reviewers could not see what was wrong, but the results couldn't be right – based on all of their prior experience, across many disciplines.

What I've come to understand is that while there is a technical argument to be made, the conflict described above is more than just whether or not a new idea is correct. The denial of the conclusions weren't based on rational reasoning, but rather because the results challenged the reviewers' paradigm on how the world works. Said another way, these results challenged the dominant belief on how diversity contributes to collective performance within all social organisms. And understanding of this challenge is key to enabling and managing diversity to solve the hardest problems.

(The photo is a panorama of a meeting where experienced and diverse participants rebelled against the leaders in the group, who thought they had to lead or they weren't doing their job. The body language in the photo is very telling.)

Slide 26 – What are your Social Identity Groups (SIGs)?

(This is the photo of a huge bear looking into my window in Santa Fe. It certainly triggered my defense of human social identity.)

To understand other people's social identity, you need to first discover your own social identities. You have many more than you are aware.

Use the definition of social identity: if someone in your identity group is attacked, then you feel attacked. Think of your family as an example. Take a minute to ask what are your different social identities applying this definition. Try groups of your ethnicity, gender, college, country, discipline, religion, sexual orientation, city, commuters, favorite sport, hobby, etc.

When I did this, I was surprised by some, for example a group I commuted with daily was very much one of my strong social identities, as I realized when one of my fellow commuters was attached while community, and how I felt violated. And I didn't even "know" the person.

How strong are your social identities? What would you to do defend them? Would you sacrifice your life for one? Some people would, e.g., a mother, a patriot, a suicide bomber, or a soldier.

Because we have many social identities and they are not expressed simultaneously, what situations or environments trigger the expression of each of your identities? For example, do you have an identity for home (a family) and another for work? Are they compatible? Or, are some identities exclusive of others?

The research on human social identity in economic transactions show that social identity easily forms and has a strong influence on our behavior:

> "... experiments show that competition is not necessary for group identification and even the most minimal group assignment can affect behavior. 'Groups' form by nothing more than random assignment of subjects to labels, such as even or odd."

Some of you might remember the experiments with children that formed strong social groups in a few days from just having a different color mark on their foreheads. And you treat you social identity group differently:

> "Subjects are more likely to give rewards to those with the same label than to those with 'other' labels, even when choices are anonymous and have no impact on their own payoffs. Subjects also have higher opinions of members of their own group."
>
> Akerlof, G. A. and R. E. Kranton (2000). "Economics and Identity." Quarterly Journal of Economics **115**, 715-753.

In the same children's study, children of one forehead color became protective of their own color and punished the "other" color.)

The bottom line is that when social identity is triggered ("I'll give up my life for my group"), it can override both rational ("I'll defend myself against threats") and habitual behavior ("I'll keep doing what I'm doing").

Slide 27 – How to manage diversity? Manage social identity in and between groups!

Earlier I defined diversity
as differences in

individual information that are contributed to a group when solving a common problem. *From a practical view in order for diversity to contribute to collective intelligence, diversity has to: 1) reside in the individual, 2) be expressed by the individual, and 3)used by the collective.*

What controls individual expression and group acceptance of diversity? What makes individuals within groups be similar and across groups be different? What are the foundations of intense group conflicts that defy reason and we all have experienced?

If you thought the answers to these questions are complex, Dr. Lefkoff and I concluded that the answer is much more straightforward than you'd guess, AND there are unexploited opportunities in this simple explanation. The key to answering all of these questions is simply *group social identity* and its dynamics within and between groups.

Social identity is what Merle and I identified as the missing key to the door of understanding and resolving conflicts within and between groups. I've now also come to understand that social identity of an individual or group also controls the expression of their diversity more than any other factor (e.g., personality or rational choice).

The major reason that social identity is a missing key is that we have great difficulty actually seeing others, ourselves and our behaviors – as in the photo above – we see distortions, instead of how we truly appear.

For our focus here, social identity is the key to managing diversity, because our social identities influence what we see, how we think, and how we treat others.

A social identity group (SIG) is in practice defined as: if someone does something to someone in your social identity group, you feel the same as if it was done to you. This is true for someone in your identity group winning a lottery or receiving a threat.

In the above slide, I list three core properties of social identities, to illustrate why they are important.

- Fact: Social identity is an innate property of all social organisms, from slime molds to social insects to social mammals to apes to humans. It is wired into all social organisms.
- Fact: Social organisms evolved social identities to coordinate and manage individual expressions of diversity for the collective good IN SIMPLE ENVIRONMENTS. Said another way: the primary function of social Identity is to repress individual diversity within the social identity group (SIG) in order to coordinate the SIG's activities to address a common challenge,

particularly from an opposing SIG. The classic example is stopping what you are doing to "circle up the wagons" to address a common attack.

- Fact: The acceptable expression of an individual's diversity is contained and managed within a social identity. The social identity (some would call culture when applied to a population) defines what is acceptable individual and group expression and what is not.

The above is for social identities for all social organisms *in simple environments*.

- In more complex environments, such as in our current society and in NIH programs, social identities may do more harm than good. Strong social identities can cause severe problems, specifically because when triggered, the repression of expressed diversity blocks the function of collective intelligence, needed for a group to function healthily in complex environments. The Wen Ho Lee case is an example. I'll show you why shortly.
- Fact: Many of our more intuitive problem-solving abilities are embodied or situated within a social identity (making use of definitions of embodied and situated as used in cognitive science and artificial intelligence). In the introduction, I gave a mathematics example of how a biased or subjective understanding can be more useful than an objective capability - Euclidian and non-Euclidian geometry. This is an excellent example of the utility of situated, non-objective knowledge that is based on a bias or has limited application. This observation puts into question the desirability of the universal application of the scientific method, an objective process from beginning to end.

To summarize, social identity is the mechanism in which diversity is coordinated at larger scales, within a self-consistent worldview. While ethnic or cultural identities come to mind as obvious social identities, this fact applies equally to scientific disciplines and even to seemingly casual social groups.

Main takeaways:

- Social identity is part of our innate programming and often determines the expression of diversity by individuals and by groups.
- If collective intelligence requires the expression of diversity, we must understand the dynamics of social identity, particularly in more complex domains where the expression of social identity can cause repression of diversity both inside the group and in accepting the diversity of outside group, thereby, preventing collective intelligence as an option for solving problems.

Slide 28 – Basics of Social Identity Dynamics

(This is a photo of the same bear – a meter high at the shoulders. My social identity as a human was definitely triggered, even though she/he was well behaved.)

Social identity isn't a static trait of social organisms. It is a heightened behavioral state of awareness and action that must be triggered and takes energy to maintain. When social identity or rationality, the two dominant forms of active problem solving, aren't being expressed, then the social organism makes decisions that are habitual, according to the current context or environment.

The basics of the dynamics of social identity in humans: how they are activated and what behavior they cause.

- Uncertainty or threats trigger the expression of social identity, otherwise an individual or group acts habitually or rationally.
- When social identity is triggered, individuals in the group collectively copy each other's actions, moving to a group coordinated response - they "circle up the wagons". This copying is not rational nor contemplated; it's impulsive.

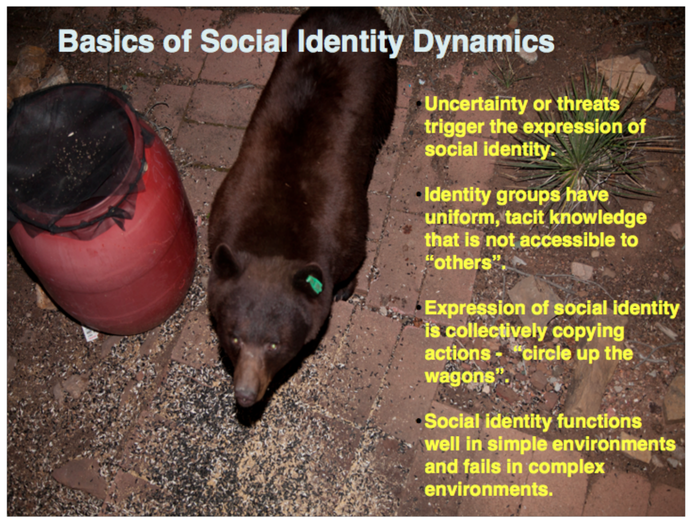

- The dynamics of social identity function well in simple environments and often fail in complex environments.
- In complex environments, interactions of identity group become more challenging because each group has uniform, tacit (hidden) knowledge or basis that is not accessible to "others". Hence, from the outside, another group's actions can seem incomprehensible or self-destructive, and can often appear as random or not rational. Within the group, the actions are obvious and acceptable.

Think of many recent actions of polarized groups that seem to contradict their core beliefs (suicide bombers or murder) and seem self-destructive from the outside, yet are uniformly supported from inside the group. For example, when 9/11 happened, America was faced with a threat that it did not understand. The rational U.S. approach would have been to seek out Americans that could explain why 9/11 happened and what could be done to prevent similar events. These subject matter experts were Muslim Americans. Instead, the America's patriotic social identity was triggered and approximately 80% of the population became hyper-patriotic as a group identity. As a result, instead of listening to Muslim Americans, the majority silenced them, isolated them, and even attacked them. The 20% that were not triggered were horrified at the 80% response. The U.S. now understands that the actions we took abroad only made the hated of the US stronger (such as physical searches of women by men), simply because we did not understand "other". In evolutionary theory, this type of innate behavior is called a *maladapted trait*: a trait that evolved to perform a useful function, but in different circumstances produces actions that do not serve the individual or group.

Slide 29 – Vegetarian Chicken Ham – Multiple SIGs

If the negative effects of social identities are a problem in complex problems, humans have the additional problem of having multiple social identities, unlike the slime mold. Remember back to the question about your different social identities.

This product produces cognitive dissonance in me, as a vegetarian.

Modern human societies have individuals with many distinct and often compartmentalized social identities. Some of these social identities are not compatible within an individual, meaning each is expressed within the presence or context of a social identity group. They are not "averaged" within the individual. This vegetarian ham loaf with chicken flavor illustrates how incompatible concepts create cognitive dissonance in our minds. So it is with our different social identities.

Vegetarian Chicken Ham

Complex social organisms have individuals with multiple, distinct social identities.

They often are not compatible within an individual, and each is expressed when within a Social Identity Group.

When you are at home, your family social identity is probably dominant. When you are at work, your discipline or work identity is dominant. But if an emergency from home intrudes on your work, you may have difficulties aligning the priorities of these family and work social identities.

- Fact: Every human has many social identities, each which is tied to certain context or environments, particularly social situations.
- The individual expression of different social identities may be in conflict because the same social identities are incompatible externally.
    - Said another way, diversity expression that is not compatible between different SIGs often is not compatible within an individual with the same social identities.
    - An example is a person of parents of two distinct ethnicities: The social identity of the child is not a mixture of parents' ethnic identities, but often the offspring retains both social identities and express them at different times. And if there is a situation where both need to be expressed, then there is cognitive dissonance.

For the above reasons, behavior of an individual in modern society can seem complex and inconsistent, but only because the social identities that generate the different behavior is unknown or hidden. If we recognize how social identities are dynamic and influence ours and others behavior, then actions and conflicts begin to be more understandable.

Hence, managing diversity requires understanding which social identities are expressed when.

Slide 30 – Application of Social Identity Groups (SIGs)

In the opening example of the case of Dr. Lee, four social identities were triggered by the incident, causing each of them to be defensive, reactive, irrational, and expressing herd behavior (acting as one):

1. Investigative branches of government felt threatened by a spy, and acted irrationally in aggressively targeting one individual and his ethnic group, largely at the exclusion of all other likely suspects.
2. Laboratory staff felt threatened by the public criticism of their loyalty, and some discriminated against an ethnic group, for no rational reason.
3. The Asian Laboratory staff felt threatened by the discrimination of Dr. Lee and because defensive and reactive as a group.
4. National Asian community felt threatened by the discrimination of Dr. Lee and the Asian laboratory staff, even though they were not in any way affected and created a powerful national campaign in protest.

Groups 1 and 2 were the ones with power and potentially would dominate the "other" groups – by discrimination and by causing loss of privacy.

Groups 3 and 4 were the ones without direct power and felt collectively discriminated against, even if an individual didn't directly experience the discrimination.

- Even though most of the individuals in the Asian staff and the National Asian community did not directly experience discrimination, they both felt personally affected, illustrating the definition of a social identity group given earlier.

For Group 1, the actions of the investigative and judicial branches of government were excessive and repressive – illustrating how a reactive and dominant social identity group, enabled by a powerful governmental system, can be self-justifying.

- The apology of the Federal judge, who at first identified with the patriotic social identity, was a remarkable public admission of the governmental excesses. Publicly, the judge did not explain why the government acted with such collective repressive actions. While the common media explanation was that the government was excessive in its attempt to find the "spy", based on the viewpoint of social identities, we can see that it more closely resembled the mass discrimination by a triggered social identity towards the "other".

The major lesson for me was that because the diversity stand-down presented a rational appeal to the staff – particularly the non-Asian staff, the effectiveness of the training was minimal. Group #2 marginalized their discriminatory actions to avoid repercussions, but the discrimination continued.

The second major lesson was that it took a diverse community to expose the new, more subtle discrimination, thereby, solving a challenging problem.

The final lesson was that if the awareness of social identity and its dynamics were understood and applied from the beginning, the entire event could have had been handled differently and with more positive outcomes.

For myself looking back 14 years, the event is an illustration of how powerful social identity can be when a group is threatened or uncertain, how the expression of diversity is severely repressed when social identity is triggered, and how in complex situations the behavior driven by social identity can cause negative consequences for all involved.

I'll give positive example shortly on how modern problem solving methods can both reduce triggering social identity and moderate existing conflicts between social identity groups – all with the goal of increasing collective performance when solving hard problems.

Slide 31 – What are the requirements for diversity-enabled solutions?

The question I receive most often when talking about collective intelligence–once a person has passed through the cognitive dissidence stage and accepts the concept–is: "What are the requirements for collective intelligence to work?" We need to know this, because we all have experienced the dreaded meetings that prove how *collective ignorance* works.

The list on this slide is a guide of requirements for collective intelligence to function, ordered from extreme at the top to minimal at the bottom.

This list of mine has evolved over 15 years – expanding outwards from the sweet spot in the center in blue, stating what I initially believed were the necessary requirements for collective intelligence to function. Now I see a spectrum of requirements, with a tradeoff between the ease of achieving collective intelligence with the ability to solve more difficult problems. What I find is that instead of a universal requirement, there is an optimal choice based on the difficulty and context of the problem we are trying to solve. This spectrum of options is similar to we saw with the different problem solving methods, in which you identified different failure modes, with the tradeoff between difficulty of the application of the method and difficulty of the problem that could be addressed.

The center requirement (in blue) captures the assumption that I made in my maze studies in 1998: where the individuals in the group have a *shared worldview*, meaning that they agree on the options, but not the preferences of the options. In the context of a maze, they all agree on the structure of the maze, but not which path is preferred. As I noted earlier, in more complex environments individuals can have different goals as well, and still contribute to collective intelligence. The reason why I call this a sweet spot is that there is less chance for disagreement on options that would lead to miscommunication or conflict. The other reason this is a sweet spot is that there are no restrictions that might limit the individual's creation and expression of diverse information, such as everyone having a common goal or social identity.

Moving up the list, I add the restriction that that the individuals in the group have shared goals, but not a shared social identity. This group state is similar to the maze problem I showed earlier, where all of the individuals had the same start and end points. Many researchers in the collective intelligence and organizational improvement community believe this more restricted state is the necessary requirement for collective intelligence. The challenge with having common goals is that it limits the complexity of problem that can be solved. Interestingly, organizations that try to energize their staff with a common vision or goals may be unnecessarily restricting the full expression of diversity of their staff and their ability to innovate as a group.

Moving further up the list, we add the restriction that the individuals are members of the same expressed social identity (I qualified identity with 'expressed', because individuals have multiple social identities but normally one identity is expressed for a given problem context). The advantage of having a common social identity is that miscommunication is unlikely, because all the individuals share common, information, vocabulary, information, goals and worldviews. But, this is also a limitation: having a common social identity can limit the diversity of the individual (they are all from one identity group) and limit expression of diversity (if the social identity is triggered).

So we can conclude that moving up the list from the sweet spot in the center leads to more efficient group communication and coordination, but at the expense of limiting diversity and therefore the complexity of the problems that can be solved. In my maze simulations I demonstrated this: by having individuals share information while exploring the maze - they "co-operated", the individuals converged to a shorter path faster, but with lower diversity and less robustness in the collective solution. Hence, while cooperation makes groups more efficient in finding a solution, it may be at the expense of diversity creation.

Moving down the list from the center sweet spot, the requirement of a common worldview is relaxed: the individuals can disagree on the options (acceptable paths), or alternatively, they can have biases that make them disagree about the options. Based on our prior discussion about biases, I add the restriction that the biases must be uncorrelated for collective intelligence to function.

My most recent research discovered that collective intelligence of a diverse group with uncorrelated biases can still discover the shortest path. It takes more individuals, but the collective intelligence still works even in the presence of biases.

Opening up collective intelligence to groups of biased individuals is a significant result for two reasons. The first is that the consensus requirement of the collective intelligence to function is the center requirement or above. These results using biased individuals shows that collective intelligence is more general, even if the collective might be challenged by biases.

An even deeper realization is that in complex problem domains, the weak signal of useful knowledge (truth) may be embedded in multiple individuals (as the name of the fake member of the Monkeys). In the same way that individuals in the maze problem can have weak signals of unknown significance, experts in complex domains can have incorrect conclusions that lead to expert disagreements, but also may have embedded knowledge (weak signals) that cannot be objectively extracted without loss of the weak signal. Said another way, the process of imposing objectivity may not be cognizant of the presence of a weak signal and may eliminate it. These weak signals do not become a recognized truth until they are combined with reinforcing expressions of the weak signal from other individuals, similar to the cutting off of unclosed loops in the prior explanation of how the collective solves hard problems.

This observation, combined with the result that the collective intelligence of uncorrelated, biased individuals still can function, allows us to conclude that collective intelligence methods are much more generally applicable and may be able to solve problems that are so complex that even the experts disagree – essentially the grand challenges. I'll give you a practical example of National importance of this next.

Finally at the bottom of the list is a group of groups with conflicting but diverse social identities. This is the equivalent of having multiple super-individuals (a super-individual being an identity group) with different goals and worldviews, but with all the diversity of a complex system captured by the different social identity groups. For example, one identity group in a complex social-political problem likely does not capture the full truth of the problem, such as in the challenging problems in the Middle East. Because this requirement captures the full-on representation of society, I put it at the bottom of the list as being the least restricted – essentially having no restrictions at all. While I have not demonstrated that collective intelligence still works for this system, my exploratory studies show that it is possible to for this bottom state to express collective intelligence. In justification of including this in the requirements list, I note there are some real world examples that have demonstrated this, e.g., the fall of the Berlin Wall, a world-changing example of a distributed, bottom-up solution to a complex social problem by multiple, distinct social identity groups.

In summary, what this list represents is the variety of social complexity that can arise in collective solutions. In the middle is a sweet spot where hard problems are solved with a compromise between high diversity and minimal failure modes. Going up the lists, reduces the possible diversity available and therefore the difficulty of problems that can be solved, but increases the ease of obtaining a collective solution.

The restrictions at the top of the list are what the social science literature imposes on a group, under the rubric of cooperation or coordination, in order for a group to function. Going down the list from the center increases the diversity and the complexity of the problems that can be solved, but at the challenge of increased failure modes, due to the conflicts of biases or even worse, conflicting social identities. Special solution methods are required to solve problem in these types of systems, as we will see next.

(The photo is of a birthday party with mostly 10 year olds – it illustrates how at this age group, there are unique individual expressions, interdependencies, and social competitiveness. I'd put them at the bottom of this list. This explains why it is so difficult to get a group of these raging egos to solve problems as a collective, unless cake and ice cream is used as an incentive. If only adults were so easily motivated.)

Slide 32 – Example: Comprehensive bio-threat risk assessment for the Nation

For the rest of the talk, let's turn to an example of the application of the above ideas to solve a Grand Challenge problem.

This is another first for me: I never used this example publicly or otherwise as an example on how to manage diversity. It was only in preparing for this NIH talk, ten years after the project ended, that I realized this project is an ideal example of how to solve a grand challenge with conflicting social identities (the bottom of the list in the prior slide). Maybe this achievement wasn't so surprising: I had a talented team and a national urgency, and we had no choice but find a solution that addressed the failure modes of a high complexity problem, one that addressed many of the challenges I've described above.

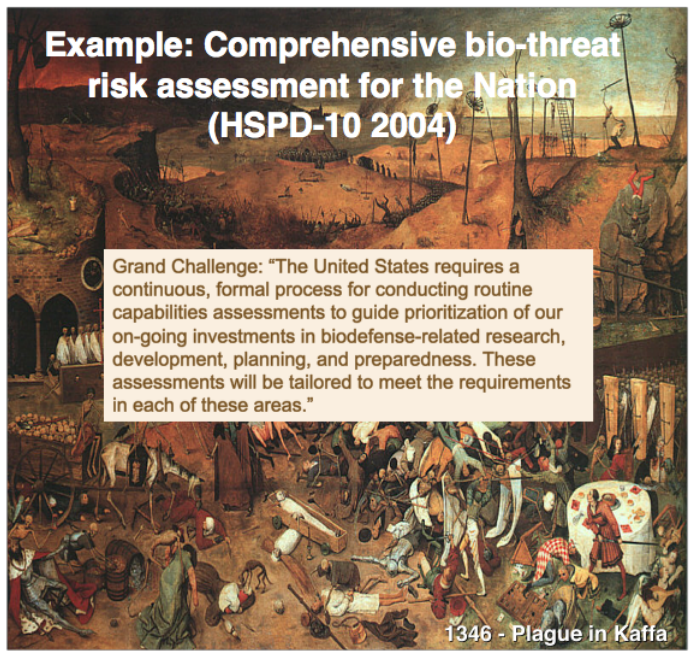

The background image of this slide captures one of the first historical bio-terrorist acts: when plague-infected bodies were catapulted into Kaffa, Ethiopia in 1346. Bioterrorism, intentional and accidental, is a growing threat in our increasingly crowded world and our society is deeply challenged to address it, for many reasons that I don't need to state for this group.

In 2004, President Bush released the *Biodefense Homeland Security Presidential Directive* (HSPD-10) to address the challenge of bioterrorism and natural threats. The Directive presented a grand challenge, excerpted in the quote, for the Nation to conduct a comprehensive risk assessment to guide biodefense investments across research, development, planning and preparedness. The successful completion of this challenge would change the direction of 100s of billion of dollars of federal funding, especially for HHS and specifically NIH.

To address this grand challenge, the Department of Homeland Security funded three parallel efforts of different sophistication: a dated approach, a current state-of-the-art approach, and an advanced approach. I lead the advanced approach at Los Alamos, which was considered to be the high-risk, high-payoff, high-expense option. At the peak of the one year effort, the staffing included about 30 full-time staff working and an equal number of on-call experts, aiding in

10,000s of SME elicitations, covering open and classified domains, across research to public health responses, as well as human, agriculture and animal threats. Nothing like this effort had ever been attempted before, because it was considered to be undoable. And because the outcome would likely impact budgets across the federal government, both public and military, the sponsor and performers expected there to be major pushback and resistance to the effort, even though we needed the stakeholders' input and eventual support to succeed.

Slide 33 – Enterprise Solution to Planning for National BioSecurity – the challenges and requirements

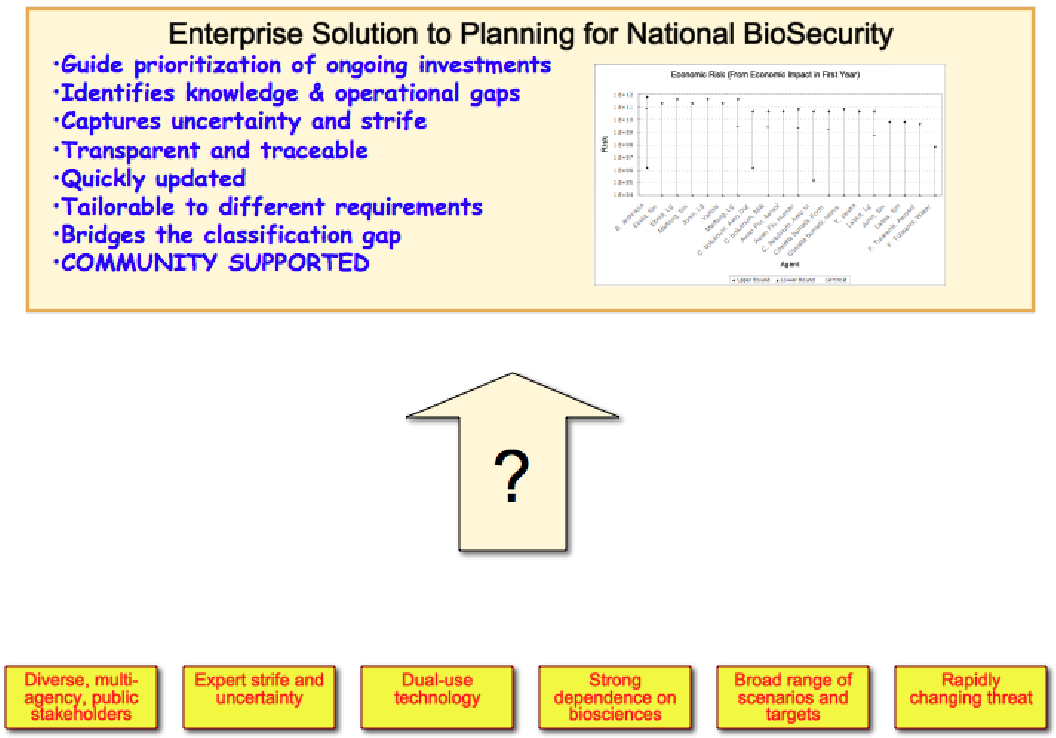

This slide highlights the nature of the grand challenge of the project.

The yellow boxes at the bottom represent the challenges – each of these was challenging enough to prevent this project from being attempted previously. In the broad scope of this project, all of them had be addressed simultaneously, as well as others I didn't include (such as complications of classification barriers and international issues, e.g., biothreat treaties).

The upper box has a list of DHS and HSPD-10 requirements for the final product – again, each of these was daunting, and we had to satisfy all of them. Addressing all of these challenges while satisfying the requirements made the effort a grand challenge at all levels: technical, operational, sociological, and political.

What made us stay awake at night?

- Stakeholders were extremely diverse with entrained social/technical identities (in beliefs, biases, and vocabulary) in both research and application practices and across multiple disciplines: from intelligence to health practitioners, from basic research to incident response, from cell biology to epidemiology.
- The most challenging stakeholders were upper management of federal programs that would be impacted, who had budgets to protect.
- The goal required huge amounts of work and SME input: do a comparative risk assessment on 20+ biological threat agents and thousands of scenarios in a transparent, comprehensive, threat-to-consequences process.

Each of the above challenges impacted how diversity was managed in the project. For example, of particular concern in executing the project was the need to get expert elicitation when there were major disagreements on technical issues –such as a mortality rate, a dose response (LD50), or an efficacy of a treatment regime. Because we were required to have quantitative results and needed quantitative input, how could we ever get consensus on parameters that different research groups had spent decades investigating, but still disagreed? This problem was compounded, because we needed consensus on tens of thousands of elicitations across the threat agents and scenarios – and getting consensus proved hopeless. Therefore, we had to use a methodology that was able capture but not resolve the conflicting inputs and aggregate them. Does this sound similar to how the collective solutions worked in the maze problem, but within a much more complex and detailed problem domain?

By the end, the project required over 40,000 expert elicitations from over 60 SMEs across all technical and operational domains - and that was just for human disease. Parallel efforts looked at plant and animal risk assessments (interestingly, these were much less contentious).

Slide 34 – Enterprise Solution to Planning for National BioSecurity – The Solution

How we accomplished the grand challenge was similar to the NIH MIDAS program, but on steroids. The following is a summary, but I'm happy to talk to any of you about methodology details and how to apply these to your grand challenges.

Here are some highlights of the project.

Before the full project, we did a pilot project for the sponsor with small team of five people to develop the main resource and to demonstrate the approach was feasible:

- We developed a *logic or inference tree* that describe the threat-to-consequence causality – similar to the maze but with more complex structure and pathways that opened and shut depending on the current state.
- The logic tree covered all (or almost all) *possibilities at the "right" level of resolution* (this is the real art of the methodology) and which all or most stakeholders could agree: mainly they could understand the logic, contribute their input, and see their interests were included in the bigger picture. We provided a process that could include everyone's contributions – even biases. This was essential for diverse and conflicting stakeholder buy-in.

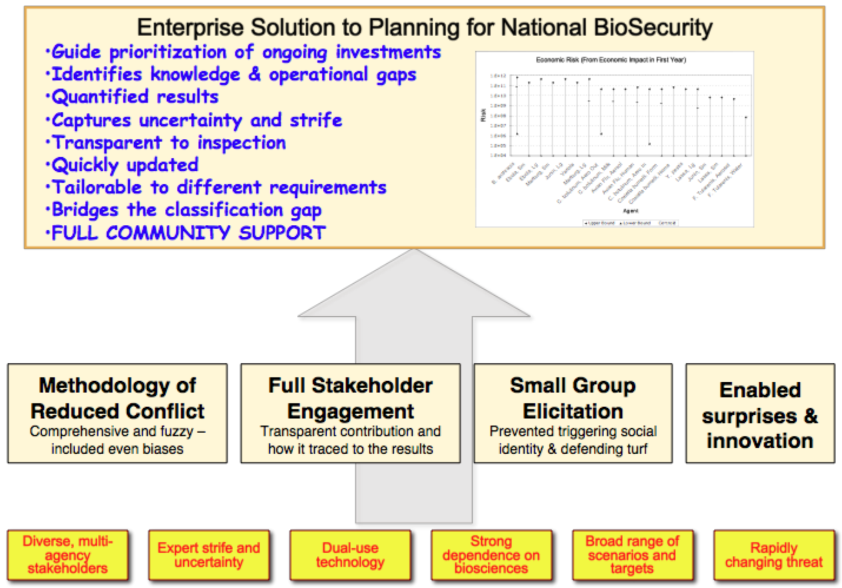

- The *fuzzy-set methodology and elicitation* essentially eliminated possible conflicts that would normally arise when having two or more experts provide information. No consensus was required in the fuzzy solicitation. Other forms of decision-making methods require agreement or consensus at each step in order to continue to the next elicitation. The importance of this difference in methodology cannot be overemphasized.
- The same fuzzy-set methodology also *captured the uncertainty* of the contribution of each expert. For example if an expert was uncertain about an LD50 value, a range can be recorded instead of a single value. Consequently experts did not have to limit their diversity by providing one answer at any point in the logic tree.
    - In my maze simulations, I found that when the individual could contribute their full range of experience, and not just their top preference, the collective solution was better.
    - Most traditional group solution methods consider a success to get one answer from everyone.
    - Therefore, enabling diversity requires contributions from all individuals plus getting each individual to contribute their fullest expression.

- Finally, because the process was quantitative from end-to-end, an SME or stakeholder could see how their contribution was included and influenced the outcome. The sensitivity of an expert's contribution to the outcome provided a context of the importance of their contribution (typically much less than they thought!).
- This sensitivity analysis also can provided feedback on how to improve the overall quality of the analysis by identifying what additional research reduced the uncertainty, satisfying a requirement of the project to identify and prioritize research gaps. This was one of the requirements of HSPD-10.
- This pilot methodology was tested on a few threat agents and scenarios, evaluated, and further refined.

Following the pilot demonstration, we expanded the team to do the risk assessment on all the agents and scenarios, and with the full stakeholder engagement. The center boxes show the main features of the methodology and execution.

- **The methodology of reduced conflict**. Because of the fuzzy-set methodology and because of the comprehensive expanse of the logic tree (from threat to vulnerability to consequence, both from an adversary and defender perspective), the different stakeholders could contribute and see how their contributions where captured and influenced the outcome.
- **Full stakeholder engagement with transparency**. By including all stakeholders in a transparent process, we achieved full stakeholder participation and support. We expected major push-back from the different stakeholders, but from the very beginning, just the opposite occurred: individuals stepped forward from multiple agencies to participate, because, as they expressed, the previous smoky back room that had determined their budgets promised to be replaced with a transparent process. This was a relief and an opportunity.
- **Small Group Elicitation**. All of the elicitation was done in small groups, targeting a specific part of the logic tree. Having small elicitation teams reduced or eliminated defending of turf and the triggering of conflicting social/technical identities.
- **Enable surprises and innovation**. The methodology enabled the capturing of surprises and innovation, addressing the most difficult challenge in risk assessment: determining "what you don't know that you don't know." Without this ability, our preparation for threats is constantly fighting the last battle. Because this feature is probably the outcome of greatest importance to programs that take on grand challenges, I'll spend more time with it.

First, having a comprehensive and quantitative assessment methodology enabled the identification of gaps in knowledge (part of the goal of HSPD-10 to prioritize future research areas). Sensitivity analysis identified what areas of information most affected the outcome, and consequently where a reduction in uncertainty would improve the accuracy of the results. Most of these gaps were already known, but some were surprises.

Additionally the logical inference of much of the bioscience-biodefense domain can be used to discover previously unknown vulnerabilities and new scenarios that exploit them (using large airplanes as WMDs to attack skyscrapers). For example, scenarios that required respirators for treatment (e.g., botulinus intoxication) identified a severe shortage of respirators (including manual ones) in almost all local public health facilities. This vulnerability is exacerbated by the lack of alternative treatment methods. While some public health officials knew of this resource gap for more than a few respiratory failures, there was no mechanism to communicate it to a policy level where it could be addressed. We could show that this vulnerability was so severe compared to other scenarios that it needed to be addressed at a national level.

Finally because of the bioscience foundation in the risk assessment, threat anticipation was possible to answer the question: where would an adversary invest to their efforts to maximize their impact – including developing new biological threat agents. This is possible because the different steps in developing a new threat–such as access to material, technical capability to accomplish the change, and detectability of the activity by a defender–are included in the risk assessment.

The threat anticipation capability is a game changer, because instead of the defender always responding to a past threat, they can look ahead to where the threat might most likely evolve. This threat anticipation is analogous to current studies of viral escape to better understand the unintended consequence of a vaccination in a population by forcing the evolution of the pathogen to higher mortality rates. This remarkable capability of innovation is only possible if the full diversity of stakeholders (including outliers) is captured across the hidden aspects of the problem domain. This is not too much different than Sharon's exercise on accidental discovery of new inventions – you have to be open to new possibilities, where, in this case, the discovery is enabled by the methodology.

Major lessons learned:

- *The quality of the outcome was directly a result of the diverse and comprehensive contributions, without selection and eliminating outliers, from the beginning.* We learned that a process which includes all stakeholder diversity lead to better solutions (higher performance) and are more robust and resilient (performed well with changes). Had the "best" experts been selected to contribute, the results would likely have been biased and with the outcome reproducing the current view of the risk landscape (this experiment could actually be done).
- *By using a process where all stakeholders participated from the beginning, the involvement in the process and acceptance of the outcome was high.* This

prevented a common failure of game-changing efforts due to lack of broad support and key players. This is a major lesson for complex programs at NIH, where a good idea or program can fail by not engaging the diverse stakeholders from the beginning.

- *Even if stakeholders didn't agree with the conclusions of the project, they could see how the results were obtained from a transparent process and could identify how their contribution was included.* This increased acceptance of the outcome and reduced conflicts, even when the results were contrary to a special interest or a paradigm.
- *And there were surprises: "what you don't know you don't know" was revealed.*

Slide 35 – First Year Economic Risk ($) for 28 agents and 52 general scenario-agent pairs

Here are preliminary results of the bio-risk assessment (the final results were never publicly released for obvious reasons), showing a range of short-term economic risk that is a quantitative combination of consequence, threat, and vulnerability with interdependencies. Similar plots show other outcomes, such as the medium-term and long-term costs or the number of deaths. Note that some scenarios have greater impact over long times, such as chemically induced cancer that would not show any symptoms for many years. Hence, it is not sufficient to make decisions based on immediate deaths or costs.

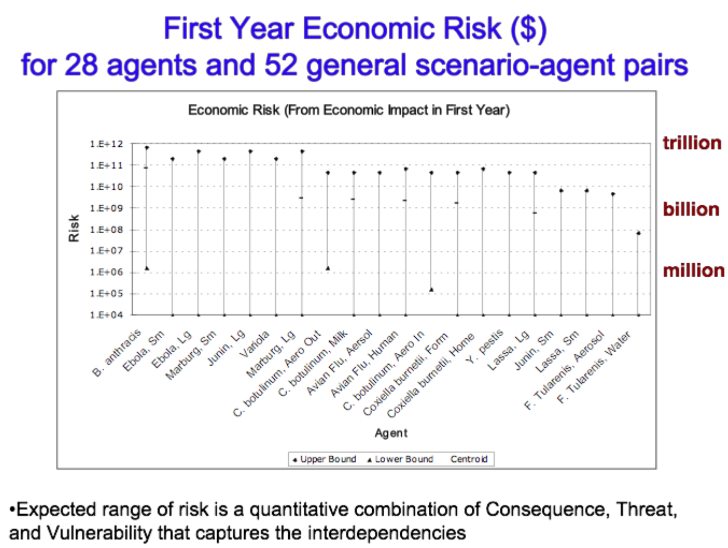

While these high risk results like these are often the ones that get the most attention (top five threats, oh my!), the loaded machinery of the comprehensive risk assessment can be a long term resource, one that can be continually improved. Hence the resource outlasts the immediate project requirements. For example, because the resource captures breadth and depth, it can be interrogated as a structured knowledge source that captures all mainstream and outlier knowledge - the true diversity of knowledge of the domain. Imagine if you had a similar resource for your discipline how you might use it.

Slide 36 – Leadership with Diverse Collectives

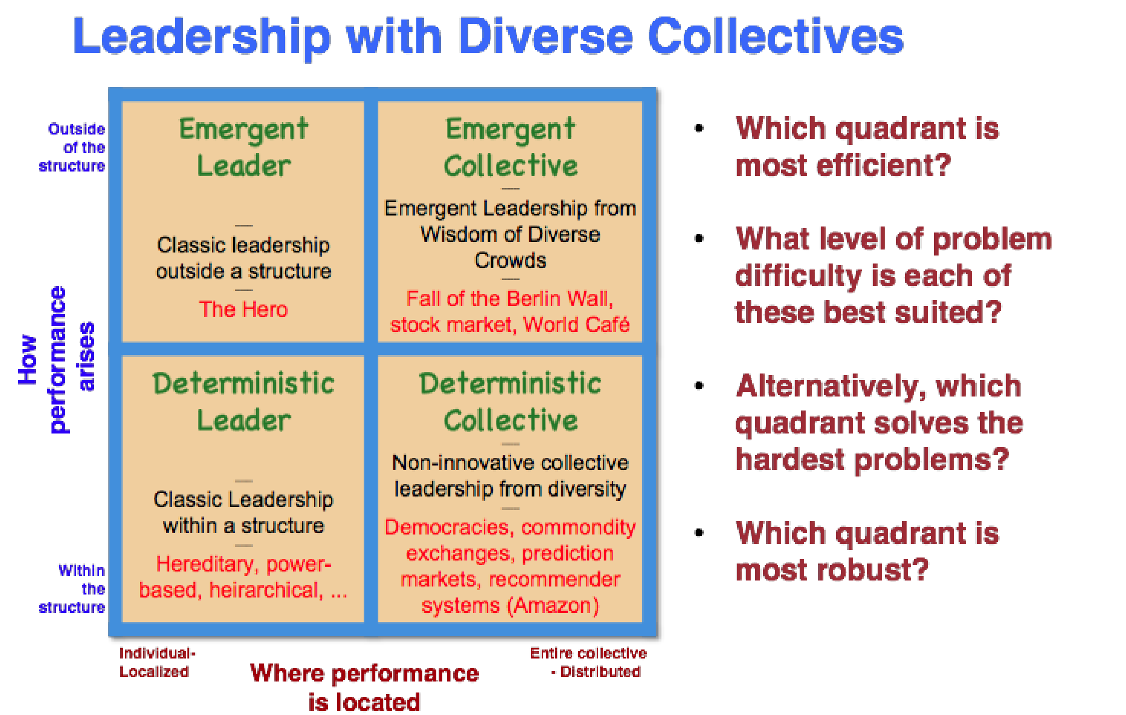

I didn't include this slide in the original presentation because of time limits. I'm including it as a resource to you when you encounter a traditional leader in order to enable them to accept using diverse collectives to solve their hard problems. It presents collective intelligence as a form of leadership. The figure is a landscape of Where and How Leadership Emerges (WHOLE) that the mechanisms for the different types of leadership and problem solving.

The horizontal axis captures where the performance is located, ranging from locally in the individual to distributed across the collective. The vertical axis captures how the performance arises, ranging from rules or structure – essentially traditional organizational methods – to *emerging* outside the structure. The landscape is divided for simplicity into four quadrants with titles in green, descriptions in black, and examples in red.

The lower-right quadrant captures the traditional individual or oligarchy leadership that is rule-base or structure-based. Above that, the upper-left quadrant captures how a an isolated leader arises outside of the main structure, essentially the hero. The two quadrants on the right are the ones we've been discussing today, the collective as a leader, either within a structure in the lower right, such as a democracy or crowd sourcing, or outside the structure in the less understood emergent collective resources in the upper right. Finding the shortest path by individuals that cannot comprehend a shortest path in my maze problem is an example of an emergent collective solution.

Based on the information discussed today, you should be able to answer the questions:

- *Which quadrant of leadership is most efficient?* (Assuming the leadership can solve the problem, that is, the expert complexity barrier of has not been hit):
  - The lower left – because an assembly line can be tuned to maximum efficiency, but is also potentially not robust, because failure in one component will cause failure in the entire assembly line.
- *Which quadrant solves the hardest problem?*
  - The right side overall and probably the upper right specifically.

The *deterministic collective* (lower right) is the diverse collective method currently in vogue – particularly in solving Michael's problems (slide 23) that are probabilistic and with many options. We are gaining a better understanding of how the deterministic collective methodologies work (see the reference at the end by my colleague, Jen Watkins: "Prediction Markets as an Aggregation Mechanism of Collective Intelligence"). This type of leadership is optimal when the problem posed is well understood but where maximum crowdsourcing (diversity) is required to get a better and more robust answer.

We didn't talk about emergent problem solving (upper right) except that in the example of how a group of individuals can solve a maze, but individually couldn't understand or apply the concept of a shortest path. Emergent problem solving also arose in the example of the bio-risk assessment, where an SME contributing their localized knowledge in the inference tree may not be able to (or even care to) understand the outcome of the project as a whole. The transparency of the bio-risk methodology in principle can be used to explain an emergent solution that was not initially apparent at a local level. Emergent collective leadership, just as how a hero arises in the upper-left quadrant, is a hot topic of study within the field of complexity science. I believe Emergent Collective leadership represents a major opportunity to solve the hardest problems facing humanity and is possibly the explanation of how societies reinvent themselves to address major challenges or opportunities.

If you want to better understand the issues and opportunities, I recommend reading:

Johnson, Norman Lee and Watkins, Jennifer H., "The Where-How of Leadership Emergence (WHOLE) Landscape: Charting Emergent Collective Leadership" (December 1, 2009). Available at SSRN – Social Sciences Research Network.

Slide 37 – Conclusions: How to solve Grand Challenges

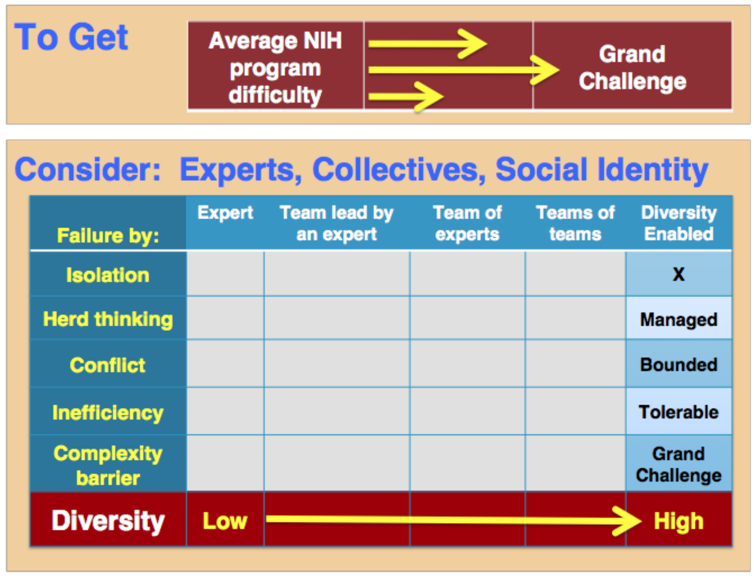

We have covered a broad range of topics – from why diverse collectives can outperform experts in hard problems, to different solution methodologies available to solve problems (the columns in this figure), to how to enable and manage diversity in methodologies that require collective solutions – specifically by being cognizant of social identities, and finally showing you a specific project that made use of enabling diversity and managing social identity to solve a grand challenge of interest to the Nation and health community.

The major takeaway for you is: if you want to include more grand challenges in your portfolio then you need to think differently about how to use and enable diversity.

Enabling diversity for NIH isn't about taking a correct or moral action. NIH has done this already better than any other federal agency. Now is the time to shift paradigms and use enabling diversity to achieve the core mission of NIH: to make the world a healthier place, and to do that will require solving grand challenges.

I hope you've seen that the diversity-enabled methodologies (rightmost column) can be implemented in a way that minimizes possible failure modes (conflict and inefficiencies), while breaking the complexity barrier of your current solution methodologies.

There is a new frontier for NIH grand challenges, all from enabling diversity.

Slide 38 – Collective Intelligence References

For other presentations by Norman Johnson on related topics, click here.

For a diverse view on Collective Intelligence (CI), the following book has summaries from many practitioners of Collective Intelligence: "COLLECTIVE INTELLIGENCE: Creating a Prosperous World at Peace." click here.

For Norm's myopic chapter that summarizes the development of this material: N. L. Johnson (2008). "*Science of Collective Intelligence*", Chapter in: Tovey, M. (Ed.) (2008). **Collective Intelligence: Creating a Prosperous World at Peace**. EIN Press, Oakton, VA. Click here.

Highly recommend the talks (both slides and text are available) of a science retreat for the Physicians' trade group: *Physician Accountability for Physician Competence*. See http://www.innovationlabs.com/summit/discovery1/ Norm helped put the program together, and the talks are a solid introduction to issues in any field of collective decision making and implementation in diverse and complex communities. Download the slides before you read the text and follow along transcribed talk. Norm gave two talks on "*Importance of diversity*" and "*Strategies in complex ecosystems*."

My out-of-date web site: http://CollectiveScience.com

Johnson, N. L. (1998). "*Collective Problem Solving: Functionality Beyond the Individual*." This is the research that started my investigation into diversity. http://collectivescience.com/Documents1.html

Johnson, N. L. (2002). "*The Development of Collective Structure and Its Response to Environmental Change*." S.E.E.D. Journal 2(3). This paper uses a model system to illustrate how collective intelligence behaves in response to different rates of change, particularly around robustness. The paper and the simulations are available here. Contact Norm for a presentation on the topic.

Johnson, Norman Lee and Watkins, Jennifer H., "*The Where-How of Leadership Emergence (WHOLE) Landscape: Charting Emergent Collective Leadership*" (December 1, 2009). Available at SSRN – Social Sciences Research Network.

N. L. Johnson and J. H. Watkins (2007). "*Interplay of Adaptive Selection and Synergistic Performance: As an example of natural selection and self-organization*" Peer reviewed paper for a presentation at Selection and Self-Organization Workshop, CSIRO-sponsored complex system science workshop, Katoomba, Australia. Available here.

Lichtenstein, Uhl-Bien, Marion, Seers, Orton and Schreiber. "*Complexity Leadership Theory: An interactive perspective on leading in complex adaptive systems*" Emergence: Complexity and Organization Volume 8, Number 4, 2006. A modern viewpoint on leadership in complex system.

Sawyer, R.K. (2006). **Social Emergence: Society as Complex Systems**. Good book as an introduction to complex systems and society.

Watkins, J.H. (2007). "*Prediction Markets as an Aggregation Mechanism of Collective Intelligence*." from http://public.lanl.gov/jhw. An excellent review of how prediction markets fit into the landscape of different decision support resources.

Watkins and Rodriguez (2007). "*A Survey of Web-based Collective Decision Making Systems*" from http://public.lanl.gov/jhw. This is a brilliant paper that surveys all the different collective decision processes – including democracy and markets and how they work.

Akerlof, G. A. and R. E. Kranton (2000). "*Economics and Identity*." Quarterly Journal of Economics **115**, 715-753. This is one of the few papers that presents social identity as an innate human trait, rather than a learned behavior, specific to economic theory and practice. A good summary of the limited social identity literature.

Avner Ben-Ner, Brian P. McCall, Massoud Stephane, and Hua Wang, "*Identity and Self-Other Differentiation in Work and Giving Behaviors: Experimental Evidence*" June 2005. Detailed research study available here. Paper here.