Why Hiring the 'Best' People Produces the Least Creative Results

While in graduate school in mathematics at the University of Wisconsin-Madison, I took a logic course from David Griffeath. The class was fun. Griffeath brought a playfulness and openness to problems. Much to my delight, about a decade later, I ran into him at a conference on traffic models. During a presentation on computational models of traffic jams, his hand went up. I wondered what Griffeath – a mathematical logician – would have to say about traffic jams. He did not disappoint. Without even a hint of excitement in his voice, he said: 'If you are modelling a traffic jam, you should just keep track of the non-cars.'

The collective response followed the familiar pattern when someone drops an unexpected, but once stated, obvious idea: a puzzled silence, giving way to a roomful of nodding heads and smiles. Nothing else needed to be said.

Griffeath had made a brilliant observation. During a traffic jam, most of the spaces on the road are filled with cars. Modelling each car takes up an enormous amount of memory. Keeping track of the empty spaces instead would use less memory – in fact almost none. Furthermore, the dynamics of the non-cars might be more amenable to analysis.

Versions of this story occur routinely at academic conferences, in research laboratories or policy meetings, within design groups, and in strategic brainstorming sessions. They share three characteristics. First, the problems are complex: they concern high-dimensional contexts that are difficult to explain, engineer, evolve or predict. Second, the breakthrough ideas do not arise by magic, nor are they constructed anew from whole cloth. They take an existing idea, insight, trick or rule, and apply it in a novel way, or they combine ideas – like Apple's breakthrough repurposing of the touchscreen technology. In Griffeath's case, he applied a concept from information theory: *minimum description length.* Fewer words are required to say 'No-L' than to list 'ABCDEFGHIJKMNOPQRSTUVWXYZ'. I should add that these new ideas typically produce modest gains. But, collectively, they can have large effects. Progress occurs as much through sequences of small steps as through giant leaps.

Third, these ideas are birthed in group settings. One person presents her perspective on a problem, describes an approach to finding a solution or identifies a sticking point, and a second person makes a suggestion or knows a workaround. The late computer scientist John Holland commonly asked: 'Have you thought about this as a Markov process, with a set of states and transition between those states?' That query would force the presenter to define states. That simple act would often lead to an insight.

The burgeoning of teams – most academic research is now done in teams, as is most investing and even most songwriting (at least for the good songs) – tracks the growing complexity of our world. We used to build roads from A to B. Now we construct transportation infrastructure with environmental, social, economic and political impacts.

The complexity of modern problems often precludes any one person from fully understanding them. Factors contributing to rising obesity levels, for example, include transportation systems and infrastructure, media, convenience foods, changing social norms, human biology and psychological factors. Designing an aircraft carrier, to take another example, requires knowledge of nuclear engineering, naval architecture, metallurgy, hydrodynamics, information systems, military protocols, the exercise of modern warfare and, given the long building time, the ability to predict trends in weapon systems.

The multidimensional or layered character of complex problems also undermines the principle of meritocracy: the idea that the ‘best person’ should be hired. There is no best person. When putting together an oncological research team, a biotech company such as Gilead or Genentech would not construct a multiple-choice test and hire the top scorers, or hire people whose resumes score highest according to some performance criteria. Instead, they would seek diversity. They would build a team of people who bring diverse knowledge bases, tools and analytic skills. That team would more likely than not include mathematicians (though not logicians such as Griffeath). And the mathematicians would likely study dynamical systems and differential equations.

Believers in a meritocracy might grant that teams ought to be diverse but then argue that meritocratic principles should apply within each category. Thus the team should consist of the ‘best’ mathematicians, the ‘best’ oncologists, and the ‘best’ biostatisticians from within the pool.

That position suffers from a similar flaw. Even with a knowledge domain, no test or criteria applied to individuals will produce the best team. Each of these domains possesses such depth and breadth, that no test can exist. Consider the field of neuroscience. Upwards of 50,000 papers were published last year covering various techniques, domains of enquiry and levels of analysis, ranging from molecules and synapses up through networks of neurons. Given that complexity, any attempt to rank a collection of neuroscientists from best to worst, as if they were competitors in the 50-metre butterfly, must fail. What could be true is that given a specific task and the composition of a particular team, one scientist would be more likely to contribute than another. Optimal hiring depends on context. Optimal teams will be diverse.

Evidence for this claim can be seen in the way that papers and patents that combine diverse ideas tend to rank as high-impact. It can also be found in the structure of the so-called random decision forest, a state-of-the-art machine-learning algorithm. Random forests consist of ensembles of decision trees. If classifying pictures, each tree makes a vote: is that a picture of a fox or a dog? A weighted majority rules. Random forests can serve many ends. They can identify bank fraud and diseases, recommend ceiling fans and predict online dating behaviour.

When building a forest, you do not select the best trees as they tend to make similar classifications. You want diversity. Programmers achieve that diversity by training each tree on different data, a technique known as *bagging.* They also *boost* the forest 'cognitively' by training trees on the hardest cases – those that the current forest gets wrong. This ensures even more diversity and accurate forests.

Yet the fallacy of meritocracy persists. Corporations, non-profits, governments, universities and even preschools test, score and hire the 'best'. This all but guarantees not creating the best team. Ranking people by common criteria produces homogeneity. And when biases creep in, it results in people who look like those making the decisions. That's not likely to lead to breakthroughs. As Astro Teller, CEO of X, the 'moonshoot factory' at Alphabet, Google's parent company, has said: 'Having people who have different mental perspectives is what's important. If you want to explore things you haven't explored, having people who look just like you and think just like you is not the best way.' We must see the forest.

—