A Method for Making "God" Decisions
Ventilator Allocations and
Prestigious College Admissions

Is this a solution looking for a problem? Perhaps, but in struggling with it, we may move toward a better understanding of the problem. In this essay I propose to introduce the reader to a lesser-known decision assistance method.

Types of decisions

Decision theory has excellent methodologies for maximizing utility among alternatives. The decision maker charts choices and possible events, assigns values to endpoint outcomes, assesses the probabilities for each event, does a little math and, presto, the best set of choices glows in the dark. The outcomes have values, usually monetary, but, sometimes what theorists call utility, the dollar equivalent of something non-monetary. When an endpoint has such a bad dollar impact on your organization that you will lose your job if it occurs, we need to give that possible outcome a negative utility value. Our theorists tell us that we can value it in a way that expresses how much we dislike the event. Three years of our salary does not seem unreasonable to subtract from that outcome's dollar value to create its utility value.

There are also courageous choices that have nothing to do with decision theory. The enemy machine gun nest is cutting your men to pieces. You grab a grenade, yell, "Cover me!" and jump out of your foxhole, getting just close enough to pull the pin and heave as you are cut in half. You had to do it to save your men. You chose the courageous alternative and did not live to get your medal. The other alternative was cowardice.

Courageous choices do not often have easily evaluated monetary outcomes. They are choices between personal sacrifice and personal safety. The negative outcomes might be physical, mental, monetary, or emotional. You lose in the courageous choice, but others are protected or advantaged.

Third, there are organizational decisions that require extensive social interaction. In universities the setting of tuition and faculty and staff pay increase guidelines may be of this type. These decisions require long timeframes, specialized analyses, and a certain level of consent from constituents. While a person, often the president, is seen as the decision maker, the role can be symbolic with the president acting more as a ratifier of what the organization has settled on. Phrases associated with this type of decision are "running an idea up the flagpole," "getting a detailed impact analysis," and "awaiting the committee's recommendations."

A fourth type of decision is a "god" decision. These decisions require that someone allocate a scarce and valuable resource to fewer people than require it. Generally, the decision maker does not need the resource, and there are no clear rules for the utility of the receipt of the resource. Ventilator allocation when the demand is greater than the supply during a Covid crisis is one example. Admissions to a prestigious university is another.

Decision theory would require that we value each dying patient's life to create the necessary utility that we would want to maximize. Just tell me the value of each person's life, and decision theory has your answer. If I were to ask a doctor to do that, my life wouldn't be worth two cents, as they say.

Characteristics of "god" decisions

The January/February 2021 issue of *The Atlantic* in "The Committee on Life and Death" by Jordan Kisner notes, "A wide spectrum of approaches and values exists within the bioethicist community…" (p. 42). The lack of agreement on a set of ethical criteria are characteristic of these decisions. Kantians (deontologists) might value saving the most lives, which sounds good, unless it means neglecting those who have been poorly served by health care in the past. People of color may again lose out under this criterion for ventilator allocation. As the article says, "'Equity still matters,' even in a crisis situation," (p. 42, quoting Matthew Wynia, the director of the Center for Bioethics and Humanities at the University of Colorado).

The number of value functions for university admissions may be even greater than the number of those possible for ventilator allocation. Test scores are popular as supposed indicators of future academic success but may also measure the quality of the preparatory schools (and, hence, the income level of the neighborhood), as well as the family's ability to purchase test preparation. The applicant's potential benefit from education is another possible choice dimension for admission. Interestingly, potential benefit may be inversely proportional to test scores: the lower the starting point, the greater the potential. The contribution to the campus is another worthwhile dimension, looking for talented people and leaders. Diversity may be valued as a way to correct historical injustice, but it can also provide a range of cultures and viewpoints that promise growth among students. One former Harvard student notes that she was evaluated on whether she would make a good roommate (https://www.youtube.com/watch?v=4VED5gzUNIQ). We cannot forget that loyalty, often defined as having parents who are big alumni contributors, is also valued on many campuses.

There have been few satisfactory methods for allocating scarce hospital resources. Governors and hospital administrators have been unwilling to give guidance. When bioethicists are available, they hesitate and patch together a process, like, treat the person most likely to recover first, unless a person who has historically been denied adequate health care has a likelihood of recovery that is close to that of the most-likely-to-recover person. This is not a terrible method, but it shifts the decision from the doctor to the bioethicist. Hence, they usually hedge their instructions even more than the imprecise method described above.

There are many other possible dimensions to consider. Should some preference be given to younger people who have more life in front of them? What about those with families who depend on their earnings?

The method of searching for the first solution that seems to satisfy most concerns was called satisficing by Cyert and March (Cyert, Richard and James G. March, *A Behavioral Theory of the*

Firm, 2nd ed., Wiley-Blackwell, 1992.) Cyert and March characterized a firm as a collection of groups with different stakes in various solution options with no single group having the power to "optimize" a solution for them alone. So, a meander through "solution space" that led to the first proposal that no one was willing to veto was the normal way of solving problems. Cyert and March noted that satisficing a solution was not always optimal for the firm as a whole (whatever that was), but that satisficing was descriptive. Satisficing appears reasonably descriptive, as well, of the method used by bioethicists to allocate ventilators as described above.

The *Atlantic* article seemed to imply that governors and hospital administrators are insufficiently courageous to endorse a value function for making critical health care allocation decisions. Nevertheless, one can see that, while it does take courage to decide that A will live and B will die, it takes the same amount of courage to decide that B will live and A will die. These are not courageous decisions requiring self-sacrifice. These are decisions where it is unclear what values should be maximized.

University admissions directors take another approach. They view their work as an art. They don't admit individuals; they sculpt a freshman class. They have minimum requirements for all values, but above that, they look for people with a set of characteristics that would make a good addition to the class. They choose people who are "solid," some who "sparkle," and others who "bring a rare quality."

Nevertheless, Harvard was challenged (and has so far survived the challenge) that their admissions folks harbor biases that illegally discriminate against a class of people. That's the problem with covering your decisions with the "art-form cloak." It's hard not appear that you are perpetuating yourself and your ideals.

Approaching "god" decisions

I submit that the reason governors don't want to specify the value functions, which would implicate them in the choices that were made, and the reason why the decisions of admissions officers are suspect, is that they *should not be making these decisions.* The most striking characteristic of a "god" decision is that there is no right answer, only wrong ones. Humans resist being forced into making decisions that are always wrong. No one wants to be accused of playing the role of a god.

I've been reading a few of the ancient Greek works lately. The gods are, in general, badly behaved. They have favorites, grab illicit lovers, and take sides in wars. Still, when Fate has determined that Achilles will die under the walls of Troy, no god (or other god, depending on the era) can interfere. So, what can we do to avoid trying to be Zeus, yet help Fate take a satisfactory course?

Interestingly, "god" decisions resemble multi-person value function decisions. That is, the same headache of multiple sets of criteria that we find in "god" decisions, we also find in decisions that are to be made by a group of people who have strong and differing ways of assessing

outcomes. The most benign example would be a group of six stock advisors who have investigated similar sets of companies, looking for differing kinds of strengths and values and must decide on only one stock to buy.

Fortunately, a founder of the field of management science, Merrill Flood, has attacked this problem ("Effectiveness of Dynamic Value Voting," Human Systems Management, vol. 4, no. 3, pp. 152-162, 1984). In 1977 I was a research assistant in the Stanford Associate Provost for Budget's office and had a cubbyhole outside the office of my boss, Raymond Bacchetti. One day Ray brought over an older gentleman who Ray said would like to share my cubbyhole. He thought we would get along. The man was Merrill Flood, who could have retired years before but was then working on value voting. He was kind, quiet, humble, and brilliant. He explained value voting to me and was trying it out on Stanford's investment folks.

He asked each analyst to distribute one hundred percentage points among a set of stocks under consideration. At this point the system sounds like distributive voting. An analyst might vote 100% for her favorite stock or might distribute her votes among several preferences. So, analyst X might have a score of 30% for stock A, while analyst Y might have a score of 20% for stock A.

Flood then thought he wanted to add information about the skill of the analyst. Perhaps analyst X was the more senior and had an excellent record for picking stocks. Flood would then weight analyst X's scores. The weighted scores for each stock were then added together and normalized, that is, the sums were divided by the total of the sums. Thus, we might find that stock A had a weighted average score of 33.6% and stock B 14.2%, etc., such that the sum of all these overall analyst scores added to 100%.

Flood then refused to take the easy way out. He did not say, "Buy the stock with the highest score," as in distributive voting. People are not infallible. The sum of all these considerations about the future, that's what stock picking is about, is still greatly uncertain. He turned to Fate. He generated a random number and used the random number to select a stock. The scores of each stock determined the probability that a stock might be picked by the random number. Imagine a dartboard divided according to the scores. Imagine a blindfolded dart thrower hitting the board. While the stock with the highest score was *most likely* to be picked by the random number or hit by the dart, it was not *certain* to be picked or hit. Any stock with a score greater than zero could be picked.

Applying Value Voting to the Allocation of Ventilators: Five Patients, One Ventilator

In this example, we do not have multiple decision makers with different value assessments, we have different ethical value systems that might be used to choose the patient who will be given the one available ventilator.

Ventilator Decision: Five people, one ventilator

								Switch, on=1, off=0	1
								Random number	0.4007314
	Probability of Living	Previous Healthcare Access (10=none, 1=good)	Normalized Living	Normalized Access	Weighted Living	Weighted Access	Raw Sum	Normalized Sum	Selected
Person A	14%	4	6.7%	14.8%	10.10%	14.81%	24.91%	10%	-
Person B	33%	7	15.9%	25.9%	23.80%	25.93%	49.72%	20%	-
Person C	64%	8	30.8%	29.6%	46.15%	29.63%	75.78%	30%	**Person C**
Person D	7%	2	3.4%	7.4%	5.05%	7.41%	12.46%	5%	-
Person E	90%	6	43.3%	22.2%	64.90%	22.22%	87.13%	35%	-
			100.0%	100.0%				100.00%	
Additional Weight	50%	0%							

Figure One: Ventilator Decision Spreadsheet

The yellow cells are for input. There are two value dimensions in this example: the probability of living (given access to the ventilator) and previous healthcare access (a justice dimension). I do not advocate these dimensions. I use them only to demonstrate the method. As noted above, many other dimensions might be used or added.

The spreadsheet could easily be designed to handle a larger number of patients, but the redesign for a larger number of ventilators is a bit more challenging (but will be demonstrated in the next example for university admissions).

The spreadsheet must be created ahead of time, although the input can be done at the last minute. This means that the criteria and the weights must be decided ahead of time. For this, a bioethicist would be helpful. Also, rules for not considering a patient should be created in advance. Certainly, a person with no chance of being helped should be ruled out. We also know that others, say, the country's president, would automatically get the ventilator.

In this example, two criteria are in place and the "probability of living" criterion has been given a 50% higher weight than the "justice" criterion simply to demonstrate the use of weights. Note also that one criterion is a percentage, and the other is a numeric scale. The criteria must be structured such that higher numbers indicate a greater likelihood of access to the ventilator. Thus, people who have been directly or indirectly discriminated against in their previous attempts to access health care receive a higher score.

Once the scores of the patients are input, the calculations are visible. (It is best to keep the random number off until all data is entered so no false hopes are raised.) First, the scores are normalized so that the score for each dimension adds to 100% across all patients.

These normalized scores are then weighted. Next, the weighted scores are summed, and, finally, these sums are re-normalized. The final scores now add to 100%. By typing a one to turn on the random number generator, a person is chosen.

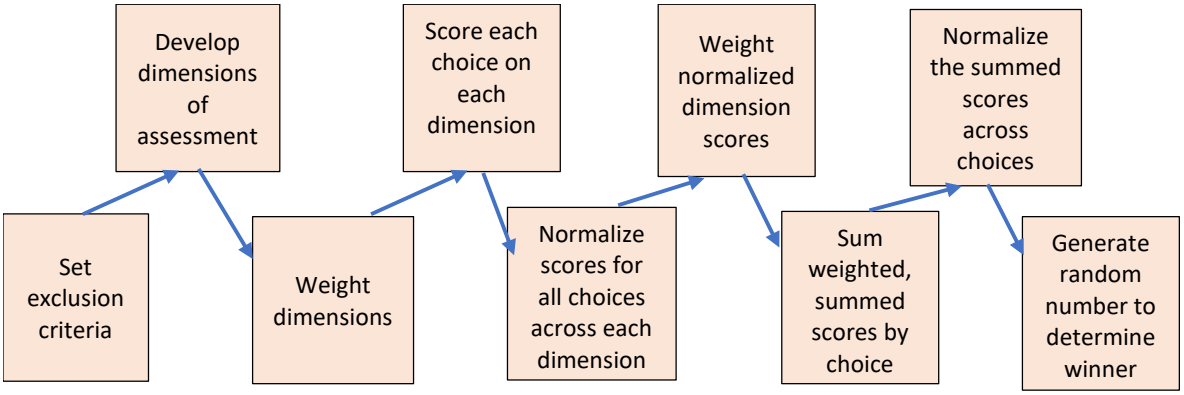

Figure 3: The principles as steps

The random number in the example is about .40 or 40%. Person A would have been chosen, if the number had been 10% or less. Person B would have been chosen, if the number were greater than 10% and less than or equal to 30%. Person C was chosen because 40% was between 30% and 60%. Note that Person E had a higher probability of being chosen than Person C, but Fate chose Person C.

How do we feel about this? We have chosen a person with only a 64% chance of living after the intervention and with only two points poorer access to previous healthcare, than a person with a 90% chance of living.

First, I counsel you to lose your worship of numbers. How accurate do you think those 64% and 90% estimates are, even if made by a well-trained physician, who, by the way, has been working a twelve-hour shift in horribly uncomfortable PPE, facing five dying people?

Second, doctor, suppose you gave the ventilator to Person E, and she died anyway. How would you feel? While this is less likely than for Person C, it is still possible.

Third, did we include all the necessary information? What about the worth to society of these people or the number of people who depend on them day-to-day? Achilles knows he's not Zeus. We aren't either. I propose leaving the decision to Fate, even if it's only the small god of a laptop.

Note, unfortunately, that five percent of the time (or, more accurately, there is a five percent chance) the small god of the computer will choose Person D with only a 7% chance of survival. Perhaps we should have set our inclusion minimum at a 10% chance of survival level to prevent choices outside a reasonable range. Because of these possibilities, one should first use the spreadsheet as a game. One learns about one's self and the difficult positions that any choice system creates.

Finally, the numbers I have used in the example are, in fact, much more optimistic than reality. Ventilators seem only to prolong a sort of life. As noted above, there are no "correct" answers. Doctors are overburdened, suggesting that we should leave the choice to the computer. Nevertheless, I can understand that doctors will be relieved just to have the calculation portion without the random choice. They may prefer to say, "We considered X, Y, and Z and your spouse was a poor choice," rather than, "Your spouse was not granted Fate's favor."

Applying Value Voting to the Selection of Freshmen at a Prestigious College: 96 Viable Applicants, 10 Slots

This is a more complicated problem, but the principles are the same as those shown in Figure 2. The greatest change in the spreadsheet comes from having more than one person to select. In this case, the spreadsheet must be programmed to fill the ten slots sequentially. After each selection, that person's scores must be removed and the scores of the remaining people re-normalized.

Farherd University Freshman Selection			Applicant 1	Applicant 40	Applicant 59	Applicant 93	
96 apps, 10 selected			Applicant 19	Applicant 42	Applicant 80		
Run=1, off=0	1		Applicant 29	Applicant 49	Applicant 83		
Criterion	Test Score	Leadership	Adversity	Talent	Unique Contribution Reviewer A	Unique Contribution Reviewer B	Loyalty
Range, high	1600	10	20	25	10	10	10
Range, low	1000	5	0	5	2	2	0
Weight	3	1	1	2	1	1	0.5
Applicant 1	1,555	7	4	18	8	6	4
Applicant 2	1,251	7	3	11	6	5	6
Applicant 3	1,281	9	15	23	8	7	5
Applicant 4	1,330	6	5	19	5	4	7
Applicant 5	1,292	6	12	15	10	8	2
Applicant 6	1,542	6	6	18	6	6	7

Figure 3: Top-right corner of Freshman Selection Spreadsheet

Only six of the ninety-six applicants are shown in Figure 3. While the calculation columns are not visible, a run's results are. There are seven criteria. I make no claims on the adequacy of these criteria: not collectively and not individually. I simply wanted to have criteria with a variety of ranges. All scores were devised such that a higher score improved the probability of selection.

Note that I have specified the allowable range of each criterion. Applicants with a test score less than 1,000 were not considered. Those with leadership scores below five were not considered. The 96 are "viable" applicants.

Note also that I have taken a page from the stock selection example. I have two different people entering their estimates of unique contribution. Their views were weighted the same. This gives "Unique Contribution" an actual weight of two.

These weights are multipliers. Test score gets a triple weight, talent a double weight, and loyalty a half weight.

						Random #	0.431437414		
Norm/wtd Test Score	Norm/wtd Leadership	Norm/wtd Adversity	Norm/wtd Talent	Unique Contribution Reviewer A	Unique Contribution Reviewer B	Norm/wtd Loyalty	Score Round 1	Cumulative Scores	Winner
0.035143399	0.00994746	0.005238949	0.02387028	0.014065285	0.013181638	0.004538278	0.01115635	0.01115635	-
0.028267234	0.010091085	0.004101708	0.01428526	0.010174976	0.009974577	0.006560117	0.00878473	0.01994108	-
0.028934804	0.012758016	0.019687173	0.031703996	0.013593292	0.013976552	0.005321478	0.01326056	0.03320164	-
0.030053383	0.008715873	0.006900355	0.025438757	0.00825293	0.008916149	0.008169527	0.01015231	0.04335395	-
0.029198353	0.007755247	0.015830627	0.020174314	0.017008308	0.017413931	0.001998619	0.01151362	0.05486757	-

Figure 4: Calculation columns, applicants 1-5 shown, Round 1 columns

In this spreadsheet, I saved some column space by normalizing and weighting each criterion in individual columns and then adding and renormalizing them in a single column as "Score Round 1," as shown in Figure 4. I created a "Cumulative Scores" column to make finding the "Winner" or selected applicant easier. Since the first random number was about 0.431, the "winner" of this round doesn't show up in the first five applicants. Applicant 42's cumulative score was just above 0.431 at 0.436. Applicant 1 was selected in the sixth round with a random number of 0.0019.

Three columns were needed for each round. Each round had its own random number, and all the scores re-normalized after the winners of the previous rounds were removed. Ten more columns were needed to save the removed scores. Other rows and columns were added to re-format the list of winners into something easily viewed.

Applicants 87 through 96 had average test scores 100 points below the others. They also had adversity scores an average of five points higher than the others. If these folks represented one sort of diversity, the method, in general, but not always, admitted one or two of them in test runs.

I believe that admissions directors and their staffs are noble and gifted people. I think the art they practice is challenging and that they have raised the quality and importance of higher education. If the choice were merely between taking applicants with the highest test scores or this method, I would strongly advocate this method. If the choice is between the current art of admissions and this method, I would equivocate. I hope that something like the use of test scores only, as was mentioned in the Harvard litigation, is never forced on us.

Nevertheless, there are sufficient advantages of this method over the art of filling a class to consider it. First, it can be transparent. A college can publish the criteria and give examples of how they are scored. Previous minimum scores may be referenced, but current minimums would depend on application volume.

Second, as in the ventilator example, all criteria are inexact measures. Why make an exact choice based on inexact measures. Sometimes Fate is preferable to pseudoscience.

Third, the future is not well understood. How different is a hunch from a thrown dart? Many applicants who look good on paper fail to thrive. The measures are inexact and not perfectly predictive.

Fourth, most dimensions have both positive and negative aspects. Many autocrats are known as leaders. A person with countless leadership indications on her application may not necessarily be a good campus addition. Also, as another example, high test scores are negatively correlated with social adjustment in my experience.

Fifth, to the extent that the process is currently hidden, accusations of illegal discrimination are easily raised, if less easily proved. Some biases are visible in the matrix and thus, more easily avoided. Bias is also possible in the ratings along the dimensions. Nevertheless, if an applicant were to successfully challenge a rating, the random choice step is still unlikely to select the challenger.

More research is possible. How many applicants who apply to two prestigious institutions are accepted by both? My assumption is that very few are. If this is true, are the institutions selecting for different qualities, or, as I feel is more likely, is the process reasonably random once minimum qualifications are met? If the process is already random, then perhaps it is not necessary for admissions officers to expose themselves to criticism.

The workbook with spreadsheets for both examples in this article is available for download at www.dickmeyerconsulting.com/publications.

Nathan Dickmeyer, January 2021