

What are we optimizing: short-term sum or long-term max?

Daniel Liberzon*

December 2, 2014

Optimizing some utility function is at the heart of many human activities and processes occurring in nature. Engineering designs often aim to achieve or approximate some optimal performance, while figuring out the right utility function is a major first step towards explaining the behavior of a complex system. Here I would like to apply this principle to try to accomplish a goal that is both philosophical and practical: to better understand how we do research and go about related professional tasks.

The current academic environment, at least in engineering departments in the United States, encourages faculty members to publish many papers, raise a lot of research funds, graduate many students, and in general to be “active,” “productive,” “industrious,” “prolific,” etc. The system rewards such individuals by granting them promotion and tenure, giving them salary raises and awards, and placing them in positions of prominence at their institutions and professional societies. If we try to imagine a utility function that this kind of behavior optimizes, we realize that it has two important features. First, this is *short-term* optimization. Indeed, salary raises are typically based on one’s performance over the past year; the time horizon for promotion and tenure decisions is only slightly larger, typically around 5 years or less. Second, since each additional paper published, research grant awarded, and dissertation supervised counts towards the overall goal, this short-term utility function takes the form of a *sum*. With some effort, it can be made into a weighted sum where, for example, publications in top-tier journals are considered more important. Still, such utility functions unfortunately favor short-term gains for the faculty members themselves (and their institutions) over long-term benefits that their work brings to the research community and the society at large.

On the other hand, if we examine long-term impact over several decades or an entire career, a very different picture emerges. When we assess someone’s lifetime achievements, things like the total amount of research funding that this person has raised will probably not even cross our mind. If the person has published several hundred papers, we might be impressed, but most likely only a handful of these papers will remain relevant today. The most important questions that we will ask are: What mark has this person left on the field? Has he/she written a paper or a book, or developed a concept or a technique, that still shapes the way people think? In other words, only the most significant contributions count in the long run. In our optimization setting, this means that *the long-term utility function computes a maximum (instead of a sum)*. There have been some true luminaries who managed to make several contributions of lasting value. However, most of us would be lucky to leave behind one piece of work that will outlive us. The sad truth is that a vast majority of things that we are so preoccupied with will soon be forgotten.

The two forms of utility functions described above—the short-term sum and the long-term max—lead to very different, even conflicting, optimal strategies. The path to increasing the sum is clear: work harder, grow your research group, produce more results. The best way to increase the max is somewhat less obvious, especially since it is difficult to predict which ideas will end up having the most lasting impact; however, it almost certainly involves focusing on fewer things and doing them better, and in general holding ourselves and our students to a higher standard. Of course, in reality we need to work with some convex combination of the two utility functions. Few of us can afford to completely ignore day-to-day pressures or annual performance evaluations. At the same time, I want to believe that most of us do care about the long-term value of the work we are doing as researchers and educators. I think it would be good for all of us to think carefully about the weights that we assign to each of these two components in the overall utility function. In my own opinion, we should not be afraid to increase the weight of our own long-term component, and we should also look for ways to encourage other people to do the same. If we do this, we will all be rewarded in the long run.

*Coordinated Science Laboratory, University of Illinois at Urbana-Champaign, Urbana, IL 61821, U.S.A., liberzon@uiuc.edu.