Further Properties of the Priority Rule

Michael Strevens

Draft of July 2003

Abstract

In Strevens (2003), I showed that science's priority system for distributing credit promotes an allocation of labor among research programs that maximizes science's payoff to society. The present paper extends this result by describing two ways in which the incentives induced by the priority rule change with the structure of the equation determining the payoff to society so as to maintain the optimality of the payoff even as the parameters determining the nature of the payoff vary. Two parameters are considered: the speed with which a research program is likely to realize its goal, and the correlation between the success of a given program and the successes of its competitors. This paper assumes familiarity with Strevens (2003).

1. Introduction

"The Role of the Priority Rule in Science" (Strevens 2003) argues that science's priority-driven system for distributing credit or prestige tends to attract scientists to different research programs in such a way that they arrange themselves so as to optimize, or nearly optimize, the contribution of the scientific enterprise to society.

The aim of the present paper, which assumes familiarity with the earlier paper, here referred to simply as *Priority*, is to investigate some additional advantages of the priority rule in achieving an optimal allocation of labor resources among scientific research programs.

As in the final parts of *Priority*, I will not aim for a perfect result, that is, I will not argue that the priority rule will always motivate a distribution of labor among research programs that is the very best possible for society. Rather, I will argue that as important parameters of the social benefit function change, the priority rule will automatically adjust the incentive scheme in the right sort of way. By this I mean

1. The adjustment will be in the right direction, so that when a change in a parameter requires, for society's sake, a shift in the allocation of labor from one program A to another B, the same change will result in the priority rule's making B relatively more attractive to scientists than before.

2. The adjustment will be roughly in proportion to that required, at least insofar as the increase in the attractiveness of B will become greater as the shift in allocation required for maximum social benefit becomes greater.

When an incentive scheme responds to a change in a given parameter of a payoff function in this way, say that the scheme is *well-attuned* to the parameter from the perspective of the payoff's recipients.

My aim, then, is to show that the priority rule is, from society's perspective, well-attuned to certain parameters, namely, the degree of correlation between different programs' successes and the speed with which different programs can be expected to succeed. It will not follow that the priority rule always finds the allocation of labor that optimizes benefit to society, but the reader will nevertheless be impressed, I think, by the degree of sensitivity inherent in the rule.

2. Correlation in Success Probabilities

In *Priority*, I assumed for the sake of the argument that the successes of two competing research programs were stochastically independent, while noting that my qualitative claims held whatever the degree of correlation between the programs' successes.[1] I now want to consider the effect of varying the degree of correlation between the success of one program and the success of its competitor.

2.1 Setup

As in *Priority*, I focus on the winner-contributes-all case, the case in which second and subsequent attainments of the same goal are worthless to society. Society's benefit is maximized in the winner-contributes-all case when the probability that at least one program succeeds is maximized. In a two-program race, this probability is equal to

$$s_1 + s_2 - s_{12}$$

where the terms are, respectively, the two programs' individual probabilities of success and the probability of joint success.

---

[1] The exception being perfect anti-correlation, i.e., mutual exclusivity.

I assumed in *Priority* the stochastic independence of the successes, so that the quantity to be maximized is:

$$s_1 + s_2 - s_1 s_2.$$

So as to relax this assumption, I now introduce a parameter r to represent the degree of correlation between successes, defining r so that:

$$s_{12} = r s_1 s_2.$$

Varying r will not affect a program's probability of success, but will affect the proportion of that probability that corresponds to s_{12}. When r is zero, the programs' successes are mutually exclusive; when r is one, they are stochastically independent; and when r takes on its maximum possible value, the lower intrinsic potential program cannot succeed without the higher potential program also succeeding. (Of two competing research programs, one has a higher intrinsic potential than the other if, for any fixed quantity of resources, the one has a higher probability of attaining its goal given those resources than does the other.)[2]

2.2 Optimality

Consider the effect that varying the correlation between successes has on the optimal allocation of labor resources. The probability to be maximized is:

$$s_1 + s_2 - r s_1 s_2.$$

---

[2] The maximum possible value is the value at which $r s_1 s_2$ is equal to the success probability for the lower intrinsic potential program.

The optimal distribution of N worker-hours, then, allocates n worker-hours to the first research program in such a way as to maximize

$$s_1(n) + s_2(N - n) - rs_1(n)s_2(N - n)$$

where $s_1(\cdot)$ and $s_2(\cdot)$ are the two programs' respective probabilities of success as functions of the resources allocated them.

Define the marginal return function $m(\cdot)$ for a program as in *Priority*, so that $m(n)$ is the increase in the probability of success due to an additional worker-hour being invested on top of n hours already invested. Then the optimal value of n will be the value for which

$$m_1(n)(1 - rs_2(N - n)) = m_2(N - n)(1 - rs_1(n)).$$

The resources allotted to the first program by the optimal distribution of labor will increase as the "adjusting factor" $1 - rs_2(N - n)$ increases, thus as $rs_2(N - n)$ decreases. Similarly, the resources allotted to the second program will increase as $rs_1(n)$ decreases. Assuming that the first program has higher potential than the second, the degree to which the optimal allocation of labor favors the first program over the second will vary with $r(s_1(n) - s_2(N - n))$. Thus, as r is increased, the optimal allocation of resources increasingly favors the higher potential research program. In short, the higher the correlation between the successes of the two programs, the more resources ought to be allocated to the higher potential program to achieve the greatest good for society.

2.3 Rewards

The priority rule is societally well-attuned to success correlation if it increasingly favors the higher potential program as the correlation r increases. I will now show that this is so.

In the context of a priority race, a program's success probability may be divided into two mutually exclusive, exhaustive parts, the probability that the program succeeds and its rival fails, and the probability that both programs succeed. The effect of increasing r is to decrease the ratio of the former to the latter, while keeping the total probability constant. Thus as r increases, the probability that the program succeeds but its rival fails decreases, while the probability that both the program and its rival succeed increases.

Because, in the event that both programs succeed, only the members of the first to do so are rewarded, varying r has the following two effects on the expected rewards of the scientists in the two programs:

1. As r increases, all scientists' expected rewards decrease, and
2. The decrease is in proportion to a program's scientists' probability of *losing* a priority race. This is the probability, then, that the rival program succeeds first, given that both programs succeed.

As r increases, then, programs that are relatively more likely to win a priority race will become relatively more attractive.

In my original paper, I assumed that higher potential programs are more likely to win a priority race than their rivals. I now invoke this assumption again, to conclude that as r increases, the priority rule will attract relatively more scientists to relatively high potential programs. Thus, the priority rule is well-attuned to the degree of correlation between successes, as desired.

2.4 More on Correlation

While on the topic of correlation, I ought to point out that, all other things being equal, it is to everyone's advantage to have as low an r as possible: society's because the lower r, the higher the probability of at least one success, and individual scientists' because everyone's expected reward decreases as r increases.

This reflects another respect in which the priority rule is well-attuned to r. Insofar as it is possible to choose between research strategies that are correlated and those that are not, there is a harmony between the interests of society and the individual.[3]

3. Success, Speed, and Priority

In *Priority*, I assumed that the higher a program's potential relative to a competing program, the more likely it was to win a priority race, in the event that both programs realized their goal. I now want to investigate the effects of relaxing that

---

[3] Dasgupta and Maskin (1987) discuss a model in which correlation-avoidance, though present, is tempered by other factors in such a way that the allocation of resources achieved by maximizers of personal reward will not be optimal for society. This result, based on a somewhat different model than mine, is consistent with my qualitative claim in the main text that the priority rule is well-attuned to correlation.

assumption.

Consider first the following simple case. Two programs, *A* and *B*, are competing to make the same discovery. They have the same potential. Thus, by my assumption about what might be called the programs' speediness, they are equally likely to win a race, given that they both succeed in their common goal. In this case the priority rule makes both programs equally attractive, resulting in an equal distribution of labor, which will indeed be optimal for society.

Now suppose that *A* is the "speedier" of the two, in the sense that if both programs realize their goal, *A* is more likely to do so first. Because the priority rule rewards only the first program to succeed, under a priority regime, *A*'s speediness enhances its attraction. More scientists will join *A*. Furthermore, the speedier *A* is relative to the *B*, the more attractive it becomes, and so the greater the number of scientists who will allocate themselves to *A*.

If the priority rule is to count as societally well-attuned to speediness, it must be that when *A* is speedier, society benefits more from a relatively greater allocation of labor to *A*. If only the probability of success is taken into account, this is not so, since speediness makes a difference only if success is assured. But it seems reasonable to suppose that, all other things being equal, it would be better for a discovery to made sooner than later. If *A* and *B* are equivalent in all respects, then, but *A* represents a speedier approach to realizing the goal, then more resources should be devoted to *A* than to *B*. What's more, the speedier *A* is relative to *B*, the more the allocation should favor *A*, if society is to receive the maximum expected benefit. This is precisely what the priority rule achieves.

Now consider two competing programs of differing intrinsic potential. In *Priority*, I assumed that the higher intrinsic potential program would be speedier, that is, more likely to finish first in the event of both programs succeeding. Suppose that this is not so: the lower potential program is speedier.

The effect of the new assumption is to make the lower potential program more attractive under a priority regime than it would be if it were less speedy than its competitor. This is as it ought to be: because an earlier discovery is more desirable, a speedy low potential program is worth more investment than a sluggish version of the same. (To see this easily, just imagine the speed of the discovery factored into the value of the payoff; as can be seen from the models in *Priority*, the higher the value of a program's payoff, the greater the optimal allocation of resources to that program.) Note, however, that a slow high potential program will still likely receive more resources than a speedy low potential program, since the probability of success plays a relatively greater role in determining expected rewards than does the probability of winning a priority race in the event of two successes.

I have explained why more resources ought to be allocated to a relatively speedy low potential program, and I have shown that the priority rule follows this maxim. But does the priority rule ensure just the right increase in the resource allocation? That is a question I cannot answer here; as the reader will see immediately, the answer will depend on further parameters quantifying, in particular, the advantage to society of speed, and so another round of model-building, which will have to be reserved for a longer paper.

This last comment of mine exposes a shortcoming in the priority rule's handling of the question of the speed: the advantage accorded by the rule to a speedy program is not sensitive to the premium society puts on speed. The additional attraction of a speedy program depends only on its relative speediness, and not at all on how much that speed is worth, given the goal in question, to society. The priority rule will not distinguish between a race to understand some abstruse question in philosophy—to solve the grue problem, say—and a race to develop a serum for a plague sweeping across the globe. Of course, labor allocation in a crisis as dramatic as the coming plague will not likely be left to the priority rule, but still, it appears that the priority rule, though well-attuned to speediness itself, is not well-attuned—not attuned at all, in fact—to the societal benefits of speed.

A way out of this difficulty lies in the interpretation of the priority rule's notion of the *actual contribution* made by a discovery. If the plague is wiping out thousands every day, and the contribution made by the discoverers of the serum is proportional to the lives saved, then a rapidly developed serum makes a far greater contribution to society than a late arrival. Rewarding actual contributions, then, will make a speedy program far more attractive than a sluggish program, under these circumstances.

I suggest—this is based on observation and intuition rather than any formal results—that it is only when time is very much of the essence that timing effects are factored into the assessment of the contribution made. Otherwise, speediness affects incentives only in the way described above, that is, as a result of speedy programs' improved chances in a priority race. Thus, when timing makes relatively little difference to society, the priority rule is insensitive to to the precise amount of difference made.

I will discuss one further issue: the integration of the two forms of well-attunedness discussed in this paper. Consider the following problem. In my treatment of correlation, I assumed that high potential and speediness go together. In my treatment of speed, however, I assumed that they may not go together, and asserted that the priority rule would nevertheless create the right sort of incentive effects. It seems, then, that in the presence of speedy low potential programs, the priority rule is well-attuned to speed but not to correlations between programs.

Or at least, this is true unless there is some reason why speed matters more when success is correlated. Perhaps there is: perhaps in these circumstances, speed matters *relatively* more, since on matters other than speed there is less to choose from between the two programs. I will, however, put this issue to one side.

4. Conclusion

Nothing is perfect, but the priority rule is extremely sensitive, in the right sort of way, to a number of factors affecting the optimal allocation of resources. This is quite remarkable, given that

1. The rule is very simple,
2. The rule is computationally very tractable, and so puts a very light cognitive load on scientists (in their capacities both as recipients and bestowers of prestige), and

3. Implementation of the rule requires relatively little information, and what information it does require is relatively accessible, so that the rule puts a very light epistemic load on scientists.

It is hard to believe that there is another reward scheme that is both as easy to use and as well-attuned, from society's perspective, to so many different factors as the priority rule. Whatever the dynamics of the process that led to its adoption, it was surely a peerless attractor.

References

Dasgupta, P. and E. Maskin. (1987). The simple economics of research portfolios. *The Economic Journal* 97:581–595.

Strevens, M. (2003). The role of the priority rule in science. *Journal of Philosophy* 100:55–79.