Why scientists chase big problems
(and why it matters)

Modeling Science, Technology, and Innovation
National Academy of Sciences

May 17-18, 2016
Washington, DC

Carl T. Bergstrom
University of Washington

Jacob Foster
Department of Sociology
UCLA

Yangbo Song
School of Management & Economics
Hong Kong University of China

John Templeton Foundation

The University of Chicago Knowledge Lab

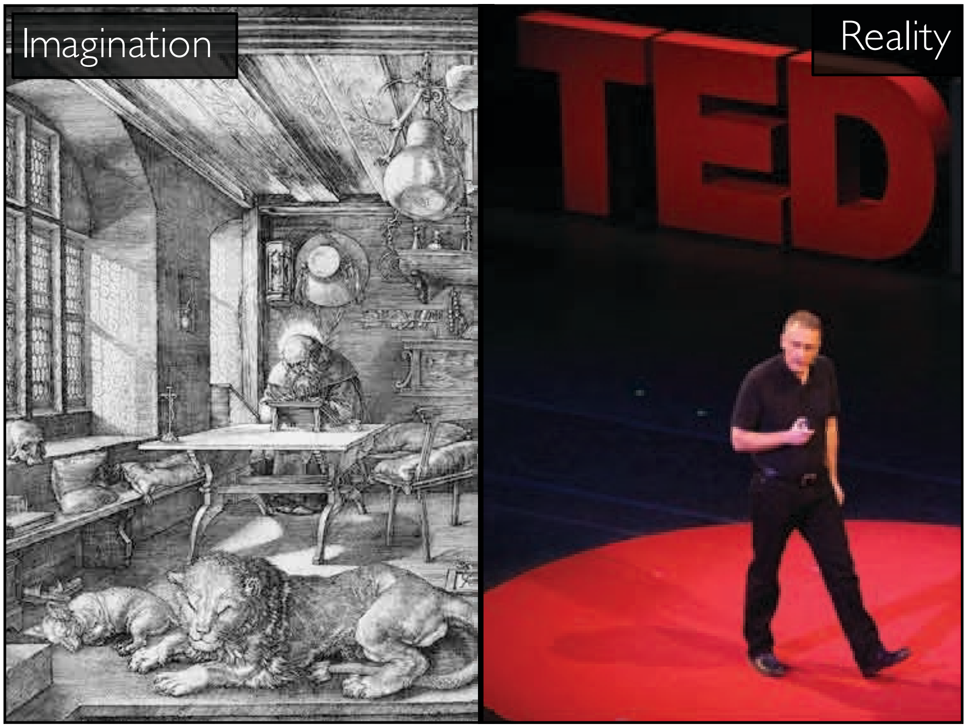

These days, scientists do research to get grants, rather than vice versa.

Given that scientists are epistemically sullied, how do the incentives created by contemporary scientific institutions lead scientists to allocate research effort across problems?

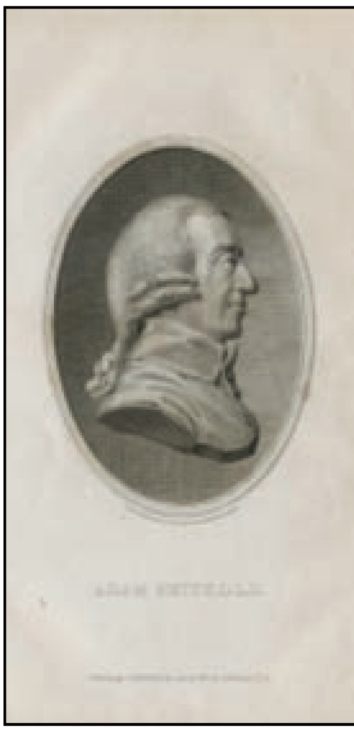

Are scientist, like Adam Smith's economic actors,

"led by an **invisible hand** to promote an end which was no part of [their] intention" ?

A theoretical framework for Science of Science Policy

Hot problems: examples

- How does DNA encode amino acid sequence?
- What is the etiological agent of AIDS?
- How do we explain the Lamb shift?
- Does Zika virus cause microcephaly and if so, how?

Hot problems: properties

- Well-defined in scope
- Widely agreed to be important within the scientific community
- More rewarding than day-to-day work.
- Not a "grind" – unclear how to solve.
- May be immediately useful to society at large, but need not be.
- Pull scientists away from other ongoing research

Model setup: the credit race game

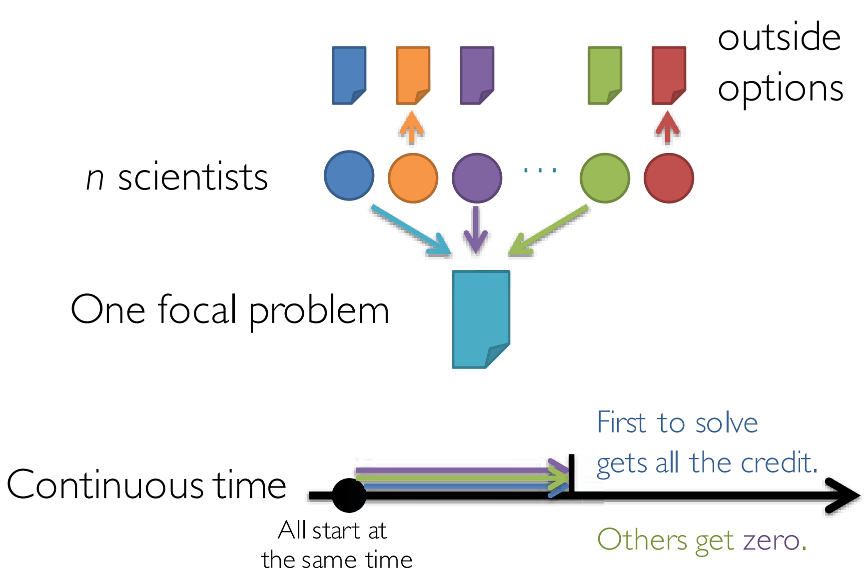

Model setup

Scientists aim to maximize the credit they receive.

Scientists vary in the amount of scientific capital h_i they can bring to bear upon a problem.

Outside option pays off in credit at rate $h_i x$

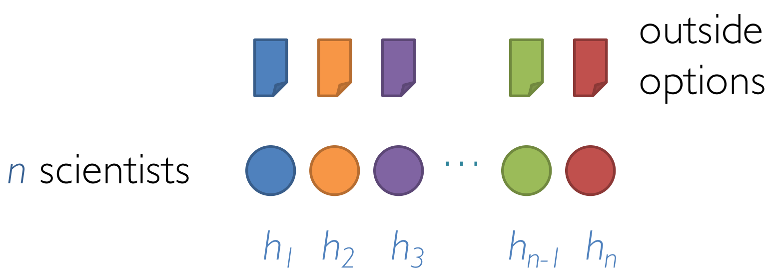

Model setup

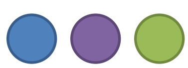

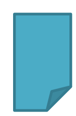

Scientists who pursue focal problem pay a startup cost F.

Constant instantaneous probability of solution h_i/d where d is the difficulty of the problem.

The expected time to solution τ by *someone* is then

$$\mathrm{Prob}(\tau \leq t) = 1 - e^{-\frac{\Sigma h_i}{d} t}.$$

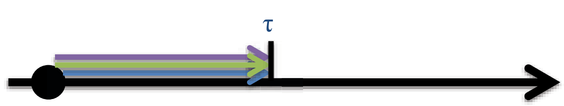

Probability that scientist j wins the race is $h_j / \Sigma h_i$.

Solver gets payoff V. Everyone else gets 0.

Model setup

Where I' is the set of scientists who pursue the problem, expected payoff U_i to scientist i of pursuing the problem is

$$U_i(I', h_i, h_{I'-i}) = V \frac{h_i}{h_{I'}} - h_i \, x \frac{d}{h_{I'}} - F$$

Reward — Chance of winning — Opportunity cost — Expected duration — Cost of entry

Model setup

A Nash equilibrium of the credit race game is a subset I^* of scientists such that those in I^* and no others do better to pursue the focal problem rather than their outside options.

Every participant does better to stay in

$$U_i(I^*, h_i, h_{I^*-i}) \geq 0 \text{ for every } i \in I^*$$

$$U_j(I^* + j, h_j, h_{I^*}) \leq 0 \text{ for every } j \notin I^*$$

Every non-participant does better to stay out

Proposition 1. A Nash equilibrium always exists in the credit race game.

Example 1

Mysterious symbols are found in Paleolithic cave art. They seem to be some kind of code.

Paleolinguists Alice (F.R.S), Bob (assistant professor), and Carol (Miller fellow) could drop everything and try to decode them.

Image: Alexander Putney, *Sanskrit*

Paleolithic code example

Alice has capital 10
Bob has capital 5 | Value V=20 | Difficulty d=5
Carol has capital 4 | Fixed cost F=4 | Opportunity cost x=1

Alice has capital 10		
Bob has capital 5	Value V=20	Difficulty d=5
Carol has capital 4	Fixed cost F=4	Opportunity cost x=1

Nash equilibria, and associated payoffs, are as follows:

	U(Alice)	U(Bob)	U(Carol)
{Alice, Bob}	6	1	(-0.84)
{Alice, Carol}	6.71	(-0.05)	0.29
but not			
{Bob, Carol}	(3.89)	4.33	2.67

With multiple equilibria, how can we derive predictive value from our model?

We need some method for equilibrium selection.

Risk dominance

The risk dominant Nash equilibrium is the one with the largest basin of attraction.

It is the one for which the cost of making the wrong move – or wrongly behaving as if you were at a different equilibrium – is the largest.

Proposition 2. A unique risk dominant equilibrium exists in the game.

At this equilibrium, the individuals with the highest scientific capital pursue the problem and all others opt out.

Paleolithic code example

Unique risk dominant equilibrium is {Alice, Bob}

	U(Alice)	U(Bob)	U(Carol)
{Alice, Bob}	6	1	(-0.84)
{Alice, Carol}	6.71	(-0.05)	0.29

High costs of switching

Low costs of switching

Top researchers chase hot topics

Extension

Scientists vary not only in their scientific capital, but also in the productivity of their normal work.
That is, the value of outside options is an increasing function of scientific capital $x_i=\chi(h_i)$ with $\chi'(h_i)>0$.

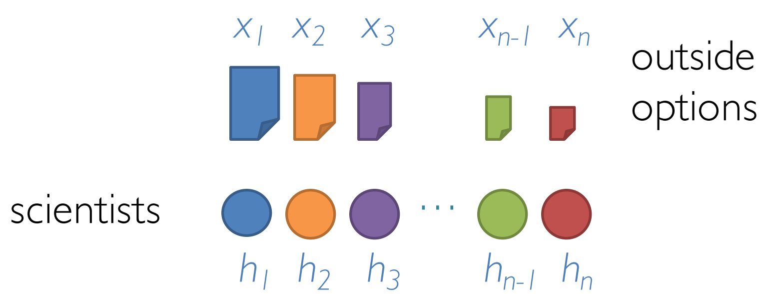

Proposition 3. When outside options differ, the equilibrium I^* composed of scientists with highest capital is the unique risk dominant equilibrium provided that $h_i\chi(h_i)$ is concave and

$$\frac{h_{i^*-1}\,\chi(h_{i^*-1}) - h_{i^*}\,\chi(h_{i^*})}{h_{i^*-1} - h_{i^*}} < \frac{V}{d}.$$

Publishing partial results

The huge difference between academic credit races and patent races is that in academia we are credited for being cited, not for some final product. Thus it can often be in our interest to publish partial results.

Publishing partial results

The focal problem now comes in stages that must be solved sequentially. Scientists have the option of publishing after completing any stage, or they may wait until solving the entire problem. If they publish intermediate results, they are guaranteed credit for that result but others can then move immediately to the next stage.

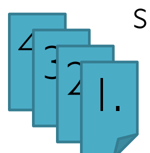

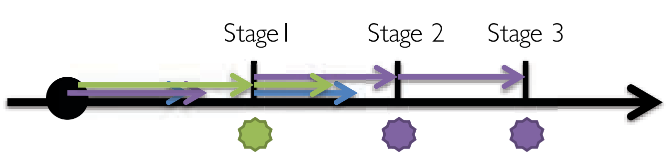

Publishing partial results

A public sharing equilibrium (PSE) is an equilibrium in which all participants publish immediately upon solving any stage.

For stage m of a problem, let I^*_m be the set $\{1,2,\ldots i^*_m\}$ of participants in that stage:

$$i^*_m = \max\{i : (V\frac{h_i}{d} - h_i\, x)\frac{d_m}{\sum_{j=1}^{i} h_j} - F_m \geq 0\}$$

Publishing partial results

For stage m of a problem, let I^*_m be the set $\{1,2,\ldots i^*_m\}$ of participants in that stage:

$$i^*_m = \max\{i : (V\frac{h_i}{d} - h_i\, x)\frac{d_m}{\sum_{j=1}^{i} h_j} - F_m \geq 0\}$$

Proposition 4. The unique risk dominant equilibrium is the PSE in which scientists in I^*_m participate in stage m, provided that for each consecutive pair of stages m and m',

$$\text{For any } i,\ \frac{V_m h_{I^*_m - i}}{d_m} - \frac{h_i}{d_{m'}}\left(\frac{V_{m'} h_{I^*_{m'} - i}}{h_i + h_{I^*_{m'} - i}} + \frac{x d_{m'} h_i}{h_i + h_{I^*_{m'} - i}}\right) \geq 0.$$

Publishing partial results

A researcher will publish partial results when

- When a stage is relatively valuable ($V_m > V_{m'}$) or easy ($d_m < d_{m'}$).
- When there are many competitors and/or competitors with high scientific capital
- When she has low scientific capital
- When the opportunity cost is low

Example 2

An unknown disease strikes a dozen people in Louisiana, many of them golfers. It causes temporary paralysis and long-term joint pain. To take preventative action, we must solve a two-stage problem:

1) What is causing this disease?
2) And where is it coming from?

Two teams, one from CDC, and one from LSU are well-poised to solve this mystery.

Louisiana golfer example

CDC has capital 12 Values $V_1=V_2=30$ Difficulty $D_1=D_2=15$
LSU has capital 5 Fixed costs $F_1=4$, $F_2=8$ Opp. costs $x=1$

Only public sharing equilibrium in this game:

Both CDC and LSU attempt stage 1. Whichever solves it first publishes immediately. Regardless of who solved stage 1, only CDC attempts stage 2 (and obviously publishes immediately.)

Louisiana golfer example

CDC has capital 12 Values $V_1=V_2=30$ Difficulty $D_1=D_2=15$
LSU has capital 5 Startup costs $F_1=4$, $F_2=8$ Opp. costs $x=1$

If researchers were not allowed to publish partial results:

The only equilibrium has the CDC attempting the problem and LSU opting out. Here we see a partial publication norm recruiting increased effort to the problem.

Social welfare
(the "why it matters" part)

In our model (as in life, largely) scientists are looking out for their own interests. How does this align with society's interests?

Because we model scientists' outside options, we are in a position to consider this question explicitly.

Social welfare

Research creates knowledge, generating a (discounted) flow of value indefinitely onward into the future.

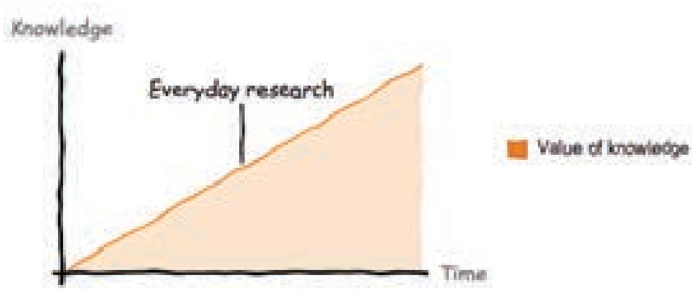

Where $K(t)$ is knowledge at time t, social welfare W is

$$W = \int_{t=0}^{\infty} \frac{K(t)}{(1+r)^t}$$

Social welfare

Working on a hot problem interrupts daily progress, in exchange for a larger subsequent jump in knowledge when the hot problem is solved.

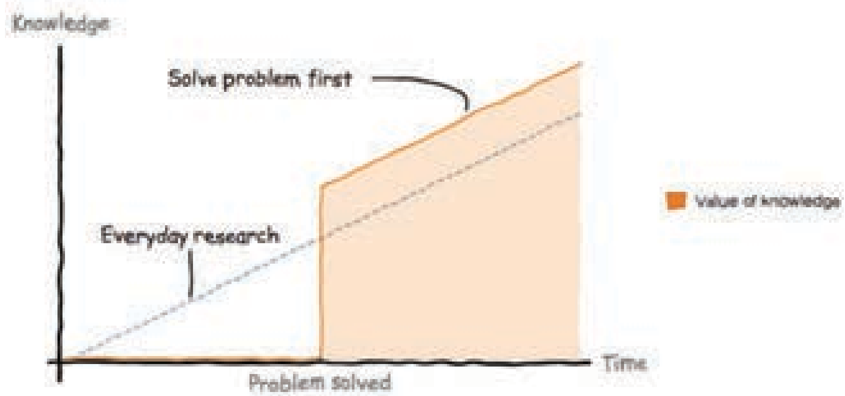

Social welfare

Where $\hat{V}$ is the social value, $\hat{x}$ is the social opportunity cost, and $\hat{F}$ is the social cost of startup, the social welfare when scientists {1,2,…*i*} with total capital H_i work on the hot problem is

$$W(i) = \frac{\hat{V}}{r}\frac{H_i}{H_i + r\,d} - \hat{x}\,d\frac{H_i}{H_i + r\,d} - i\,\hat{F}$$

Discounted Reward

Discounted opportunity cost

Cost of entry for *i* scientists

Proposition 5. There is a unique social optimum, in which the i_e scientists with the highest scientific capital work on the hot problem and the rest opt out.

Social welfare results

How does the risk dominant equilibrium (what scientists will do) compare with the socially optimum (what we would like them to do)?

Both under-participation and over-participation are possible. Roughly (ignoring some subtle discounting terms) we see under-participation when

$$\frac{\frac{\hat{V}}{r} - \hat{x}\, d}{\hat{F}} > \frac{V - x\, d}{F}$$

Example 3

Halibut fisheries in Alaska abruptly collapse. This poses a two-stage problem:

1) What is causing the collapse? Climate change? Disease? Overfishing?

2) What can we do about it?

Teams from NMFS and UW are able to attack the problem.

Image (cc) flickr: Steve Johnson

Halibut fishery example

NMFS has capital 10 Values $V_1=60$ $V_2=30$ Diff. $D_1=D_2=15$
UW has capital 6 Startup costs $F_1=1$, $F_2=6$ Opp. costs $x=1$
Value to society $\hat{V}=100$.

Only public sharing equilibrium in this game:

NMFS and UW attempt stage 1. Whichever solves it first publishes immediately. Regardless of who solved stage 1, only NMFS attempts stage 2 (and obviously publishes immediately.) Social welfare is $W=122.94$

Halibut fishery example

NMFS has capital 10 Values $V_1=60$ $V_2=30$ Diff. $D_1=D_2=15$
UW has capital 6 Startup costs $F_1=1$, $F_2=6$ Opp. costs $x=1$
Value to society $\hat{V}=100$.

If we don't allow partial progress sharing, both teams attempt the entire problem. In this case, the expected time to solution is shorter, and social welfare is higher $W=126.98$!

Contrary to conventional wisdom,

Proposition 6. Allowing partial progress sharing does not necessarily accelerate the rate at which a hot problem is solved.

Why not?

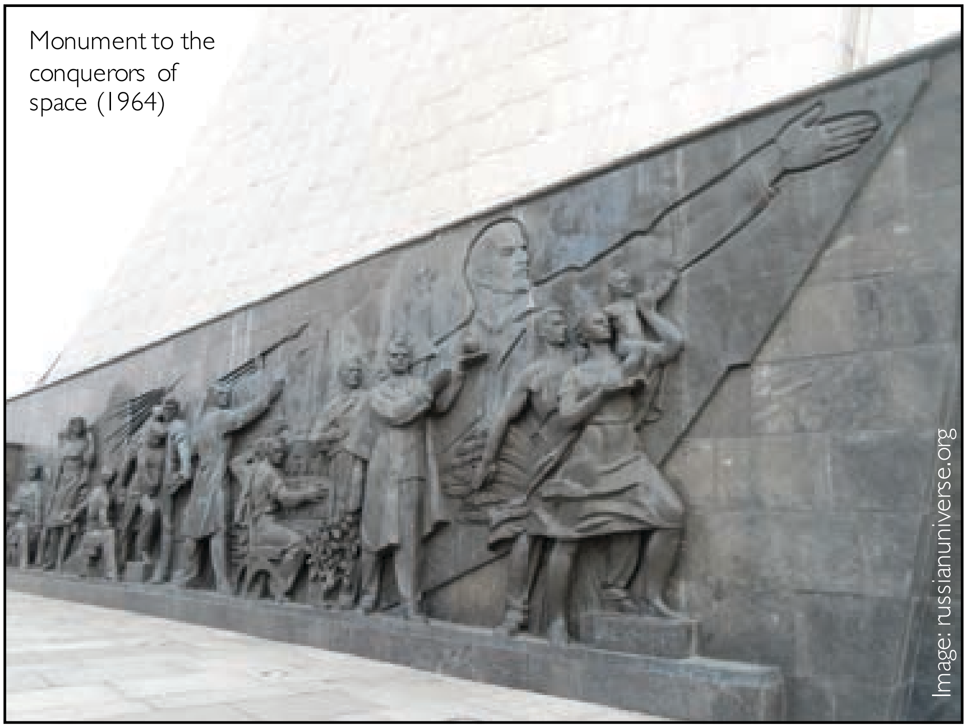

Monument to the conquerors of space (1964)

Future directions

Problems of unknown difficulty

- Why scientists give up on a problem
- Strategic informational issues around non-publication.

Reproducibility

- How do our institutions shape individual incentives in ways that contribute to or ameliorate the "reproducibility crisis."

Policy implications and suggestions

Principal-agent framework

–How do alternative forms of credit allocation influence scientists' behavior?

–What can a govt. agency do to shift scientists' efforts toward the socially optimal allocation?

–How does the publishing system contribute to irreproducibility and inefficiency?

–How do within-university institutions accelerate or retard the progress of science?

Graphical view of expected payoff