Metascience2019

The inherent inefficiency of grant proposal competitions and the possible benefits of lotteries in allocating research funding

Carl T. Bergstrom

Department of Biology

University of Washington

Kevin Gross

Department of Statistics

North Carolina State University

@ct_bergstrom

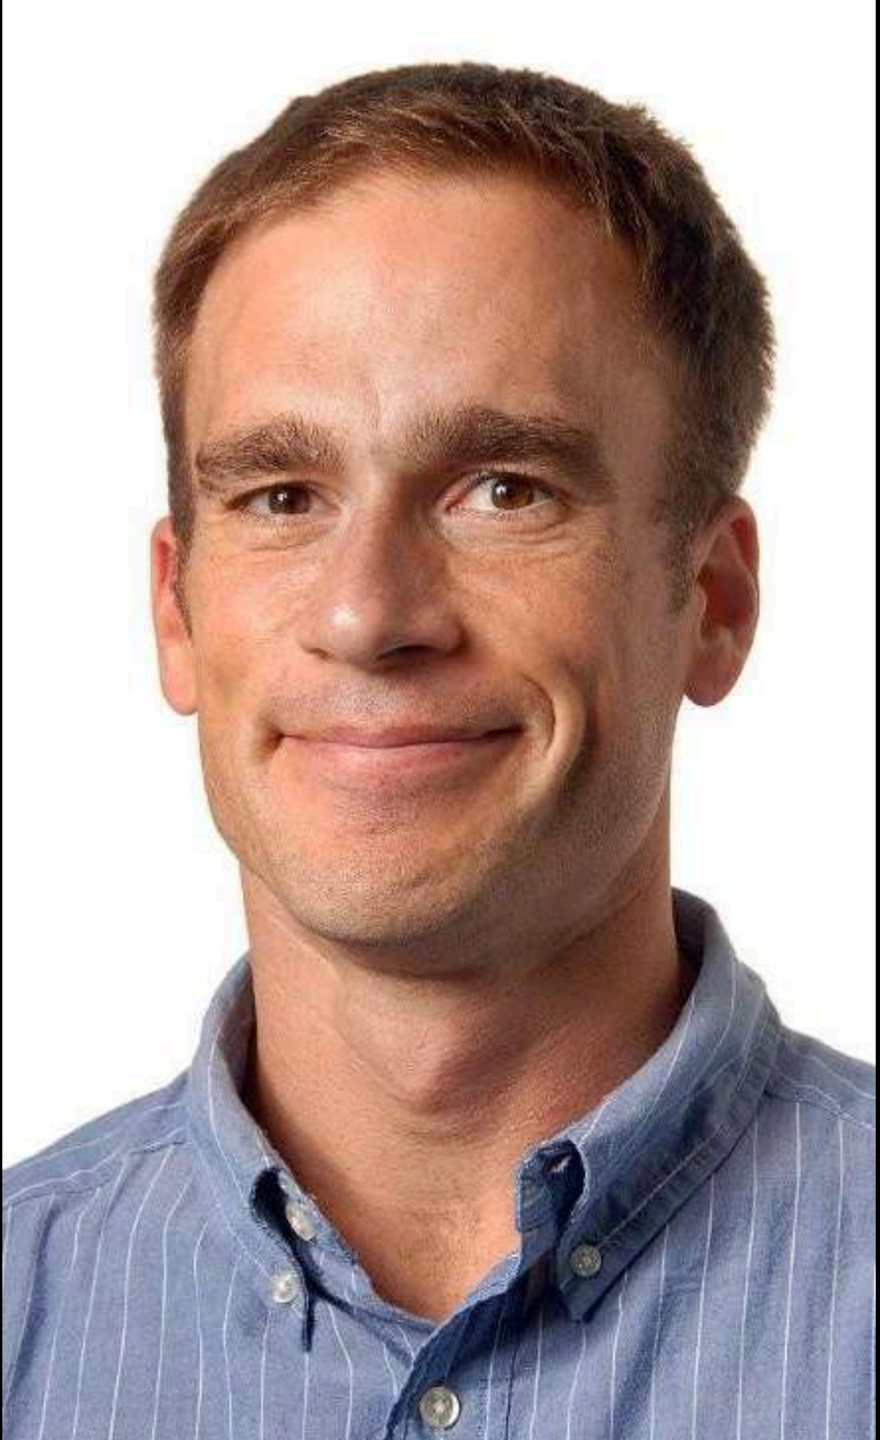

Metascience2019

The inherent inefficiency of grant proposal competitions and the possible benefits of lotteries in allocating research funding

Carl T. Bergstrom
Department of Biology
University of Washington

Kevin Gross
Department of Statistics
North Carolina State University

@ct_bergstrom

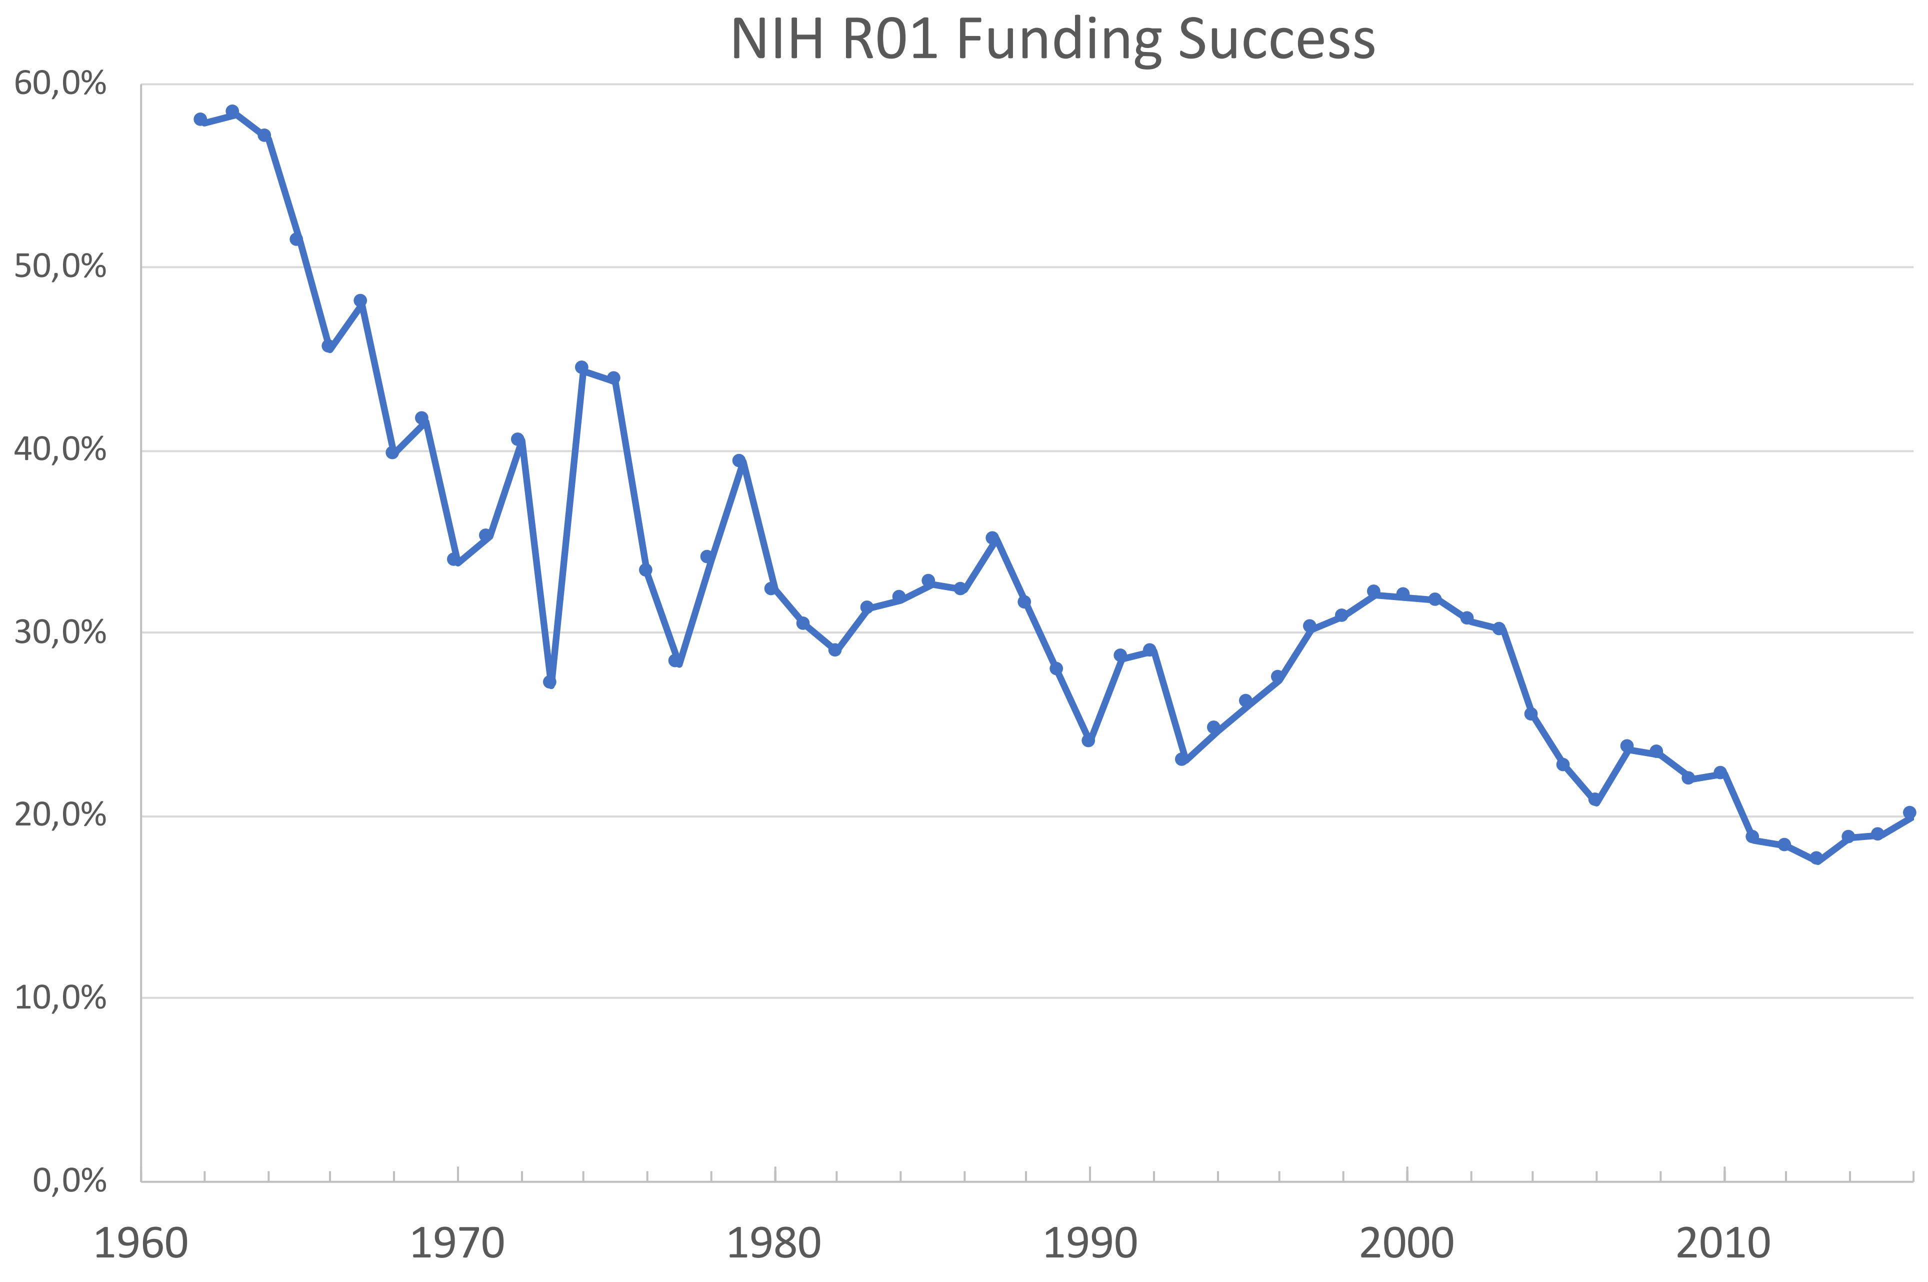

NIH R01 Funding Success
60,0%
50,0%
40,0%
30,0%
20,0%
10,0%
0,0%
1960
1970
1980
1990
2000
2010

Faculty effort toward writing grant at research universities is approximately:

~10% of total time, and

~20% of research time,

In medical schools:

up to 50% of total time

Geard and Nobel 2010; Link et al 2008; Siliciano 2007

All this stress. All this effort. All of this uncertainty.

- Is it worth it?
- How have falling paylines changed the calculus?
- Is there a better way to do things?

All the science
made possibly by
grant funding.

All the science
that doesn’t get done
because we are
busy writing grants

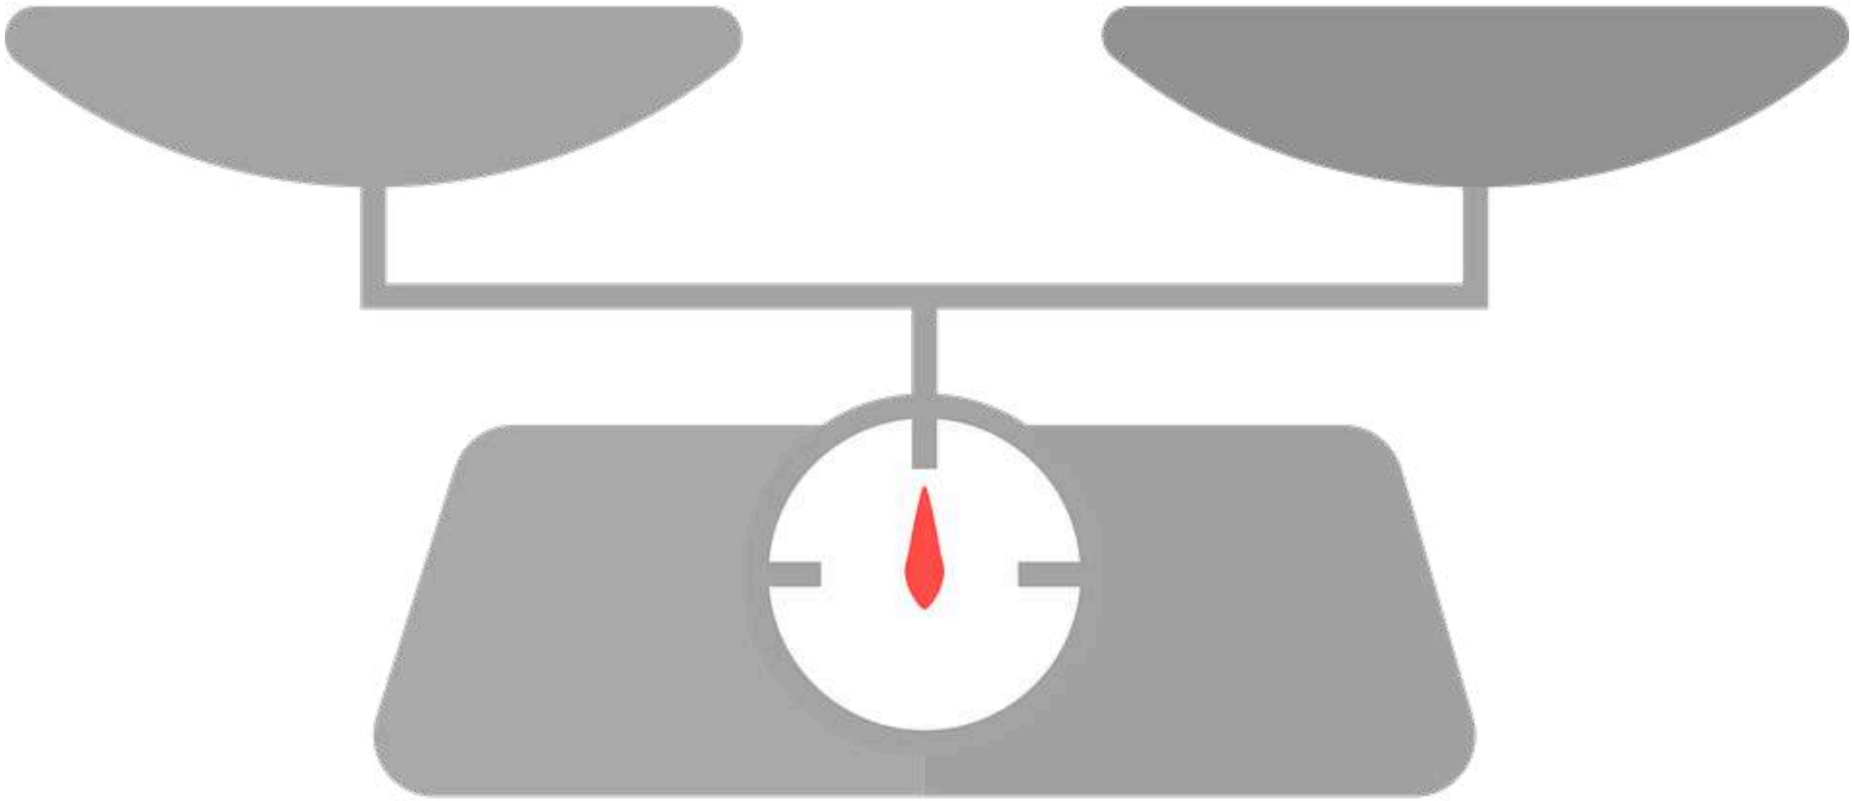

Contest theory — a branch of game theory

- Agents (contestants) $N=\{1,2,\ldots n\}$
- Possible actions (moves) G_1, G_2 … G_n for each contestant
- Cost of moves and value of prize(s)
- Contest success function $V: G^n \rightarrow P^n$

NETFLIX

Netflix Prize

COMPLETED

Home | Rules | Leaderboard | Update

Leaderboard

Showing Test Score. Click here to show quiz score

Display top 20 **leaders.**

Rank	Team Name	Best Test Score	% Improvement	Best Submit Time
Grand Prize - RMSE = 0.8567 - Winning Team: BellKor's Pragmatic Chaos				
1	BellKor's Pragmatic Chaos	0.8567	10.06	2009-07-26 18:18:28
2	The Ensemble	0.8567	10.06	2009-07-26 18:38:22
3	Grand Prize Team	0.8582	9.90	2009-07-10 21:24:40
4	Opera Solutions and Vandelay United	0.8588	9.84	2009-07-10 01:12:31
5	Vandelay Industries !	0.8591	9.81	2009-07-10 00:32:20
6	PragmaticTheory	0.8594	9.77	2009-06-24 12:06:56
7	BellKor in BigChaos	0.8601	9.70	2009-05-13 08:14:09
8	Dace_	0.8612	9.59	2009-07-24 17:18:43
9	Feeds2	0.8622	9.48	2009-07-12 13:11:51
10	BigChaos	0.8623	9.47	2009-04-07 12:33:59
11	Opera Solutions	0.8623	9.47	2009-07-24 00:34:07
12	BellKor	0.8624	9.46	2009-07-26 17:19:11
Progress Prize 2008 - RMSE = 0.8627 - Winning Team: BellKor in BigChaos				
13	xiangliang	0.8642	9.27	2009-07-15 14:53:22

Welcome to Kaggle Competitions

Challenge yourself with real-world machine learning problems

Submit »

New to Data Science?
Get started with a tutorial on our most popular competition for beginners, Titanic: Machine Learning from Disaster.

Build a Model
Get the data & use whatever tools or methods you prefer to make predictions.

Make a Submission
Upload your prediction file for real-time scoring & a spot on the leaderboard.

Learn more | InClass

✖ Dismiss

General | InClass

Sort by Grouped

All Categories | Search competitions

17 Active Competitions

2018 Data Science Bowl Find the nuclei in divergent images to advance medical discovery Featured · a month to go · biology	$100,000 2,015 teams
Google Cloud & NCAA® ML Competition 2018-Men's Apply Machine Learning to NCAA® March Madness® Featured · a month to go ·	$50,000 421 teams

salesrewards.com

Principal-agent framework

- A sort of game-design approach to game theory.

- How can the *principal* design the rules of the game so that the *agents* do the desired / socially beneficial thing?

- Typically under information asymmetry.

Principal | **Agent**

Offers prizes → Provides work

Reveals (something about) type

Principal **Agent**

Offers prizes → Reveals type

Does (largely useless) work

NIH

National Institutes of Health

NSF

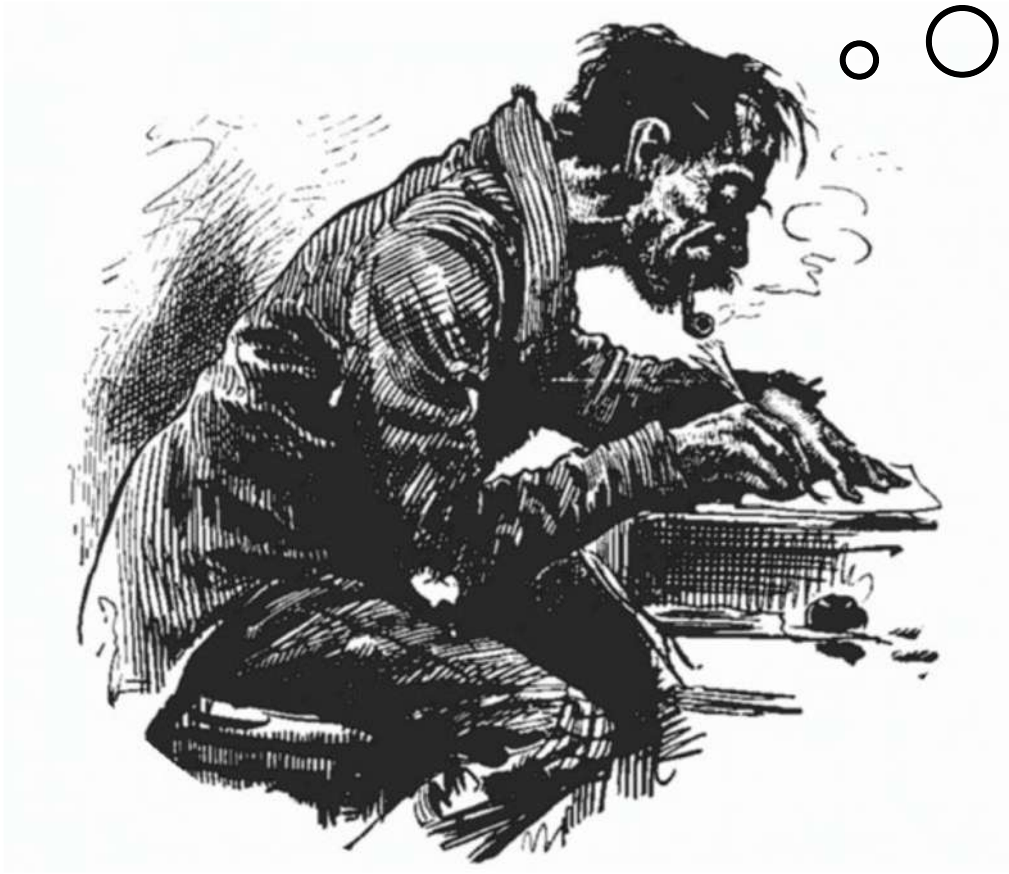

I have a research idea and I have an idea of how good it is.

How much effort should put into writing a proposal?

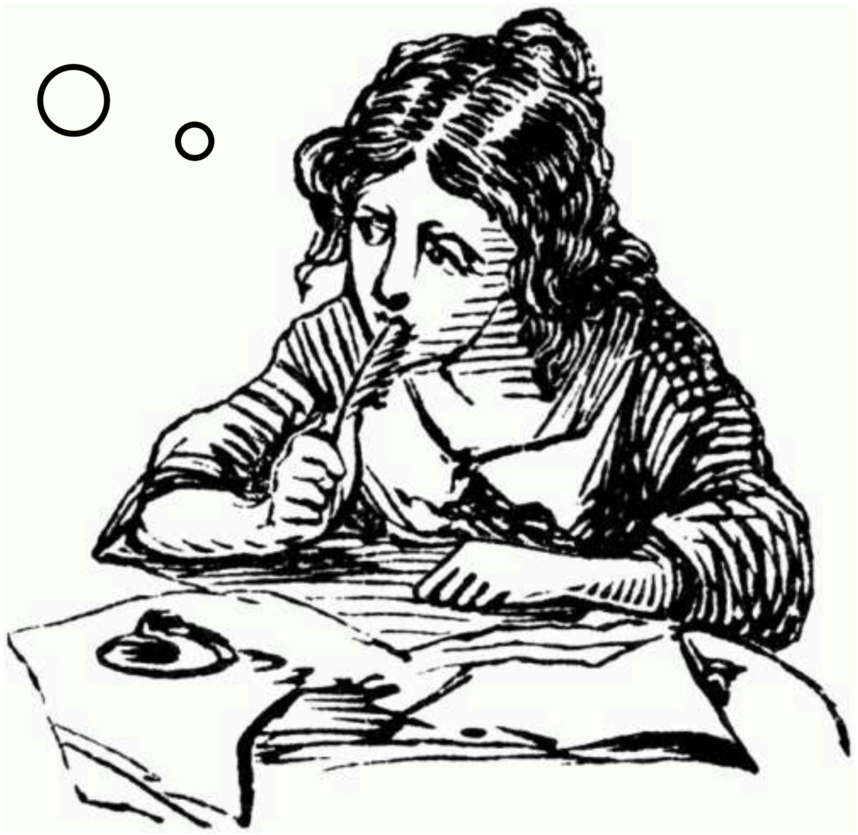

Stronger proposals are more likely to be funded, but take more effort to write.

At least it is easier to write a strong proposal for a good idea than a weak one.

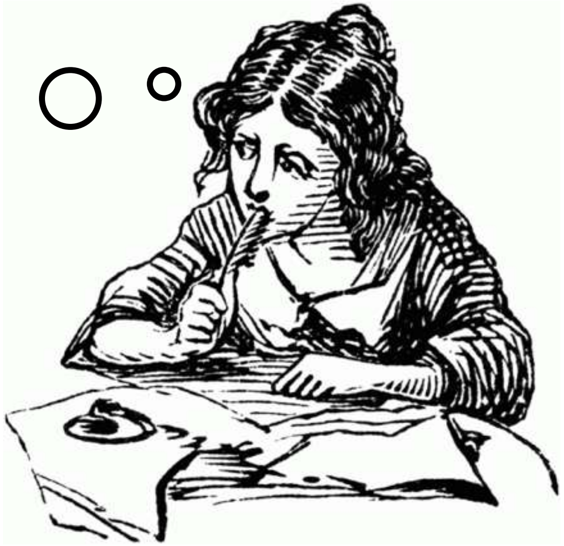

An idea has a *scientific value* v, both to researcher and to funder.

A researcher writes a proposal of *strength* x.

The *cost* of writing a proposal of strength x for an idea of value v is $c(v,x) = g(v)h(x)$,

where $g'(v) < 0$ and $h'(x) > 0$

A fraction *k* of that cost is *recaptured*.

The *grant panel* iis more likely to fund a strong grant than a weak one.

It chooses to fund a proposal of strength *x* with probability

$$\eta(x) \text{ where } \frac{d\eta}{dx} \geq 0$$

Investigator with idea v wants to write a proposal of strength *x* to maximize

$$v\,\eta(x) - (1-k)c(v,x)$$

Benefit × chance of funding

Cost of writing discounted by “recovery”

How strong (*x*) of proposal should I write?

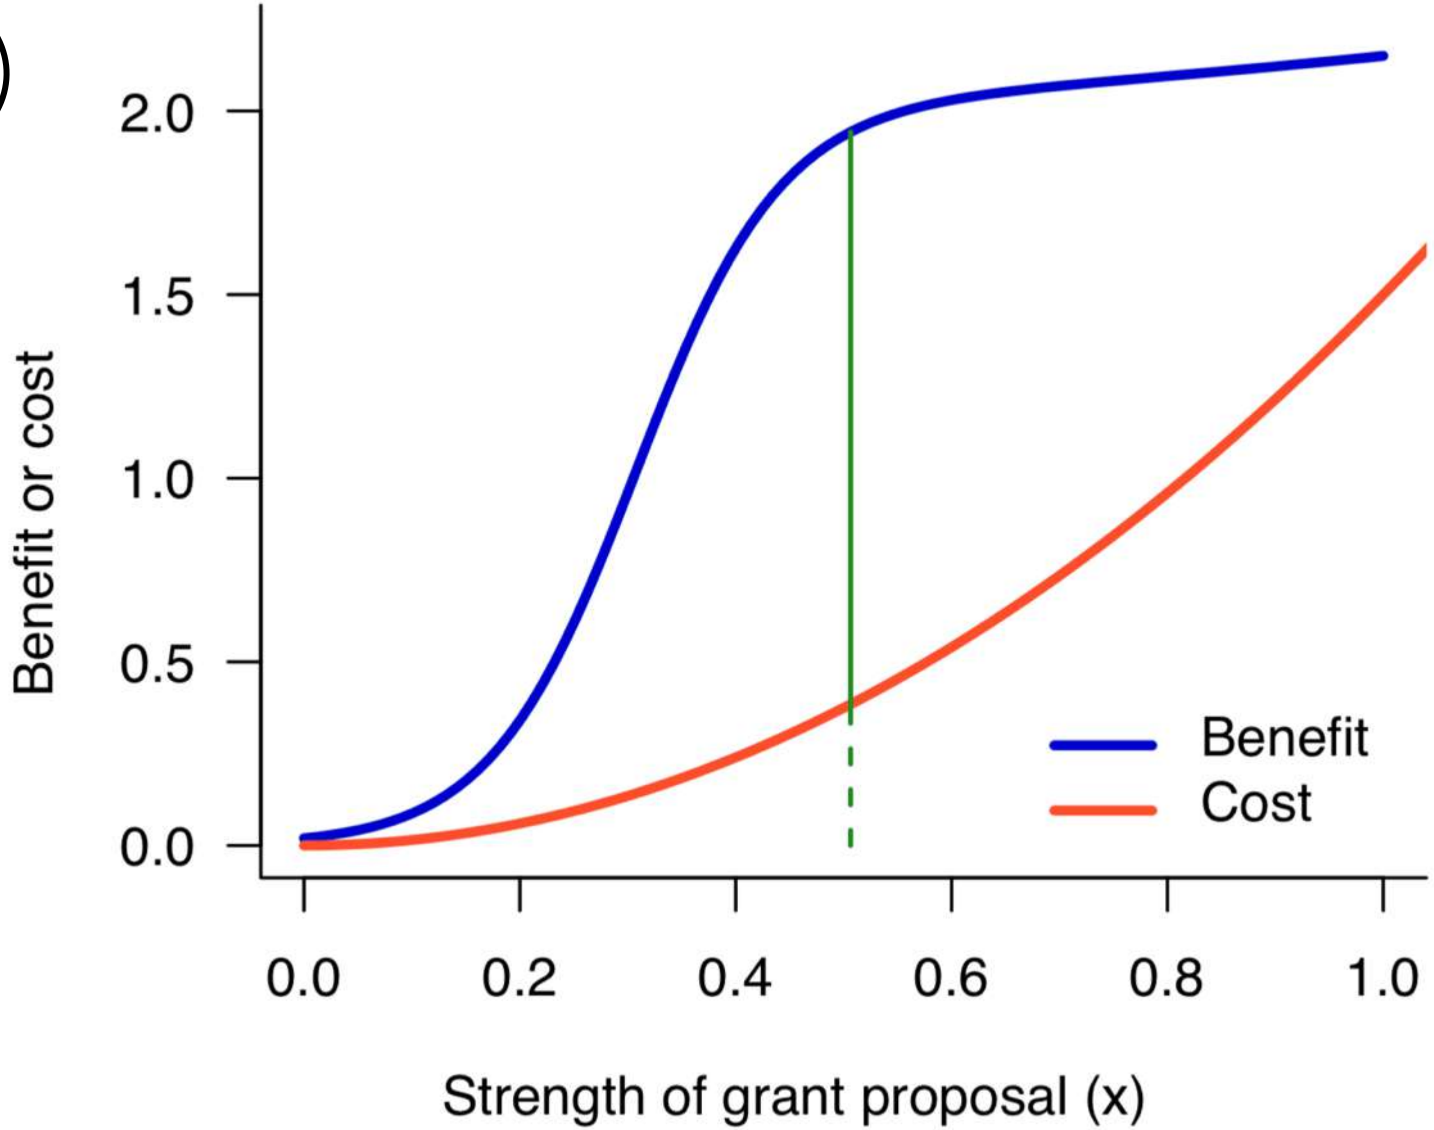

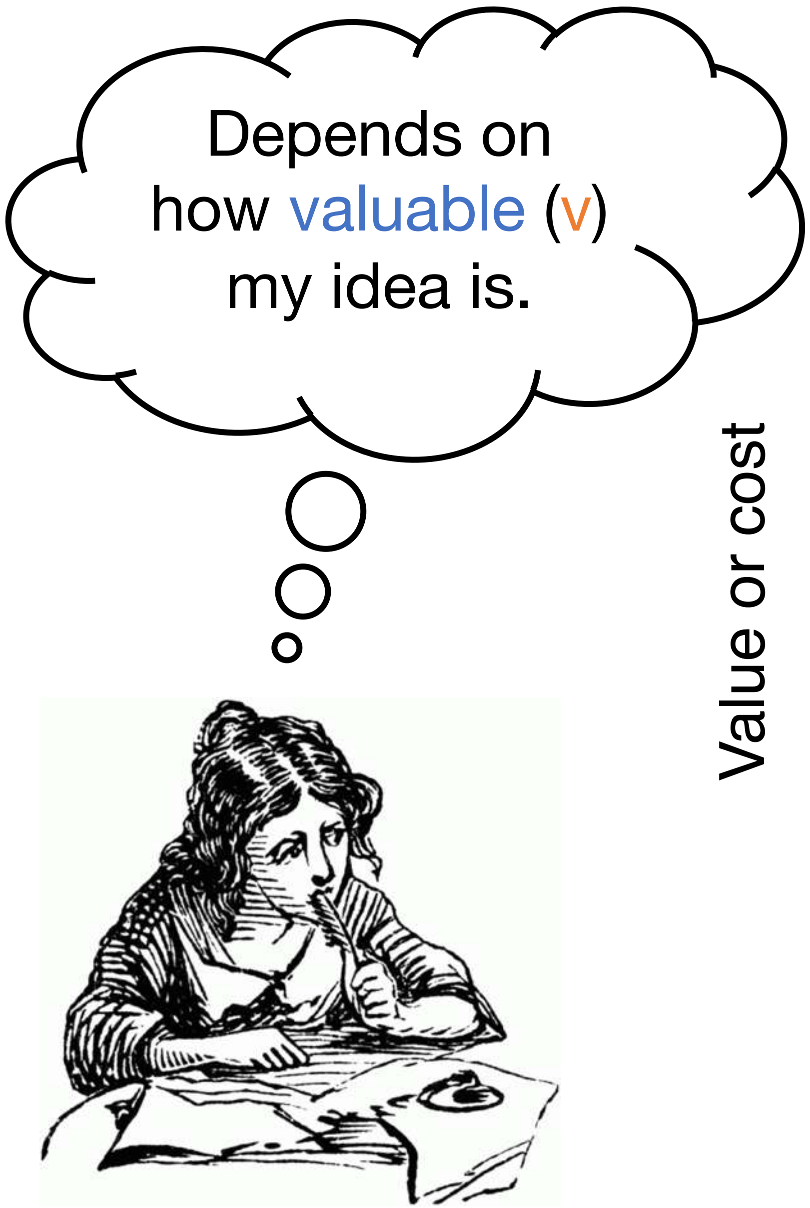

Depends on
how valuable (v)
my idea is.

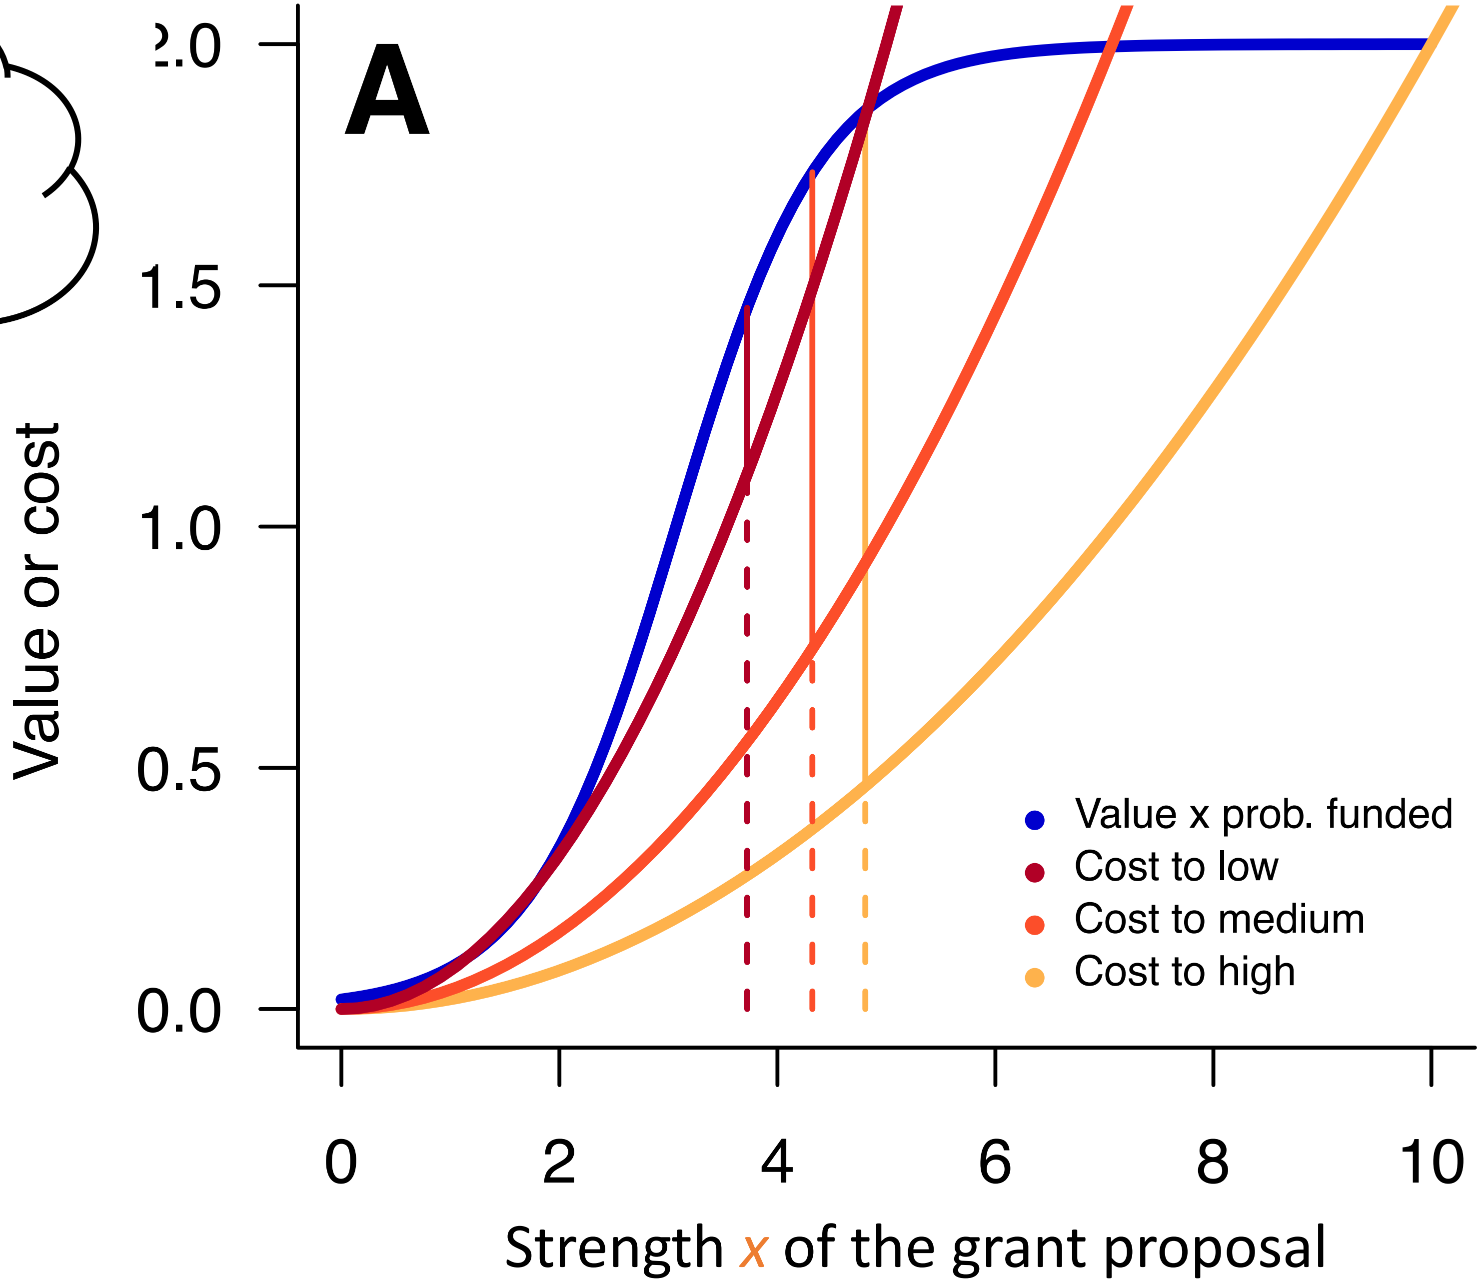

A
2.0
1.5
1.0
0.5
0.0
Value or cost
0
2
4
6
8
10
Strength x of the grant proposal
Value x prob. funded
Cost to low
Cost to medium
Cost to high

Mate Selection—A Selection for a Handicap

AMOTZ ZAHAVI

1975

JOB MARKET SIGNALING *

MICHAEL SPENCE

1973

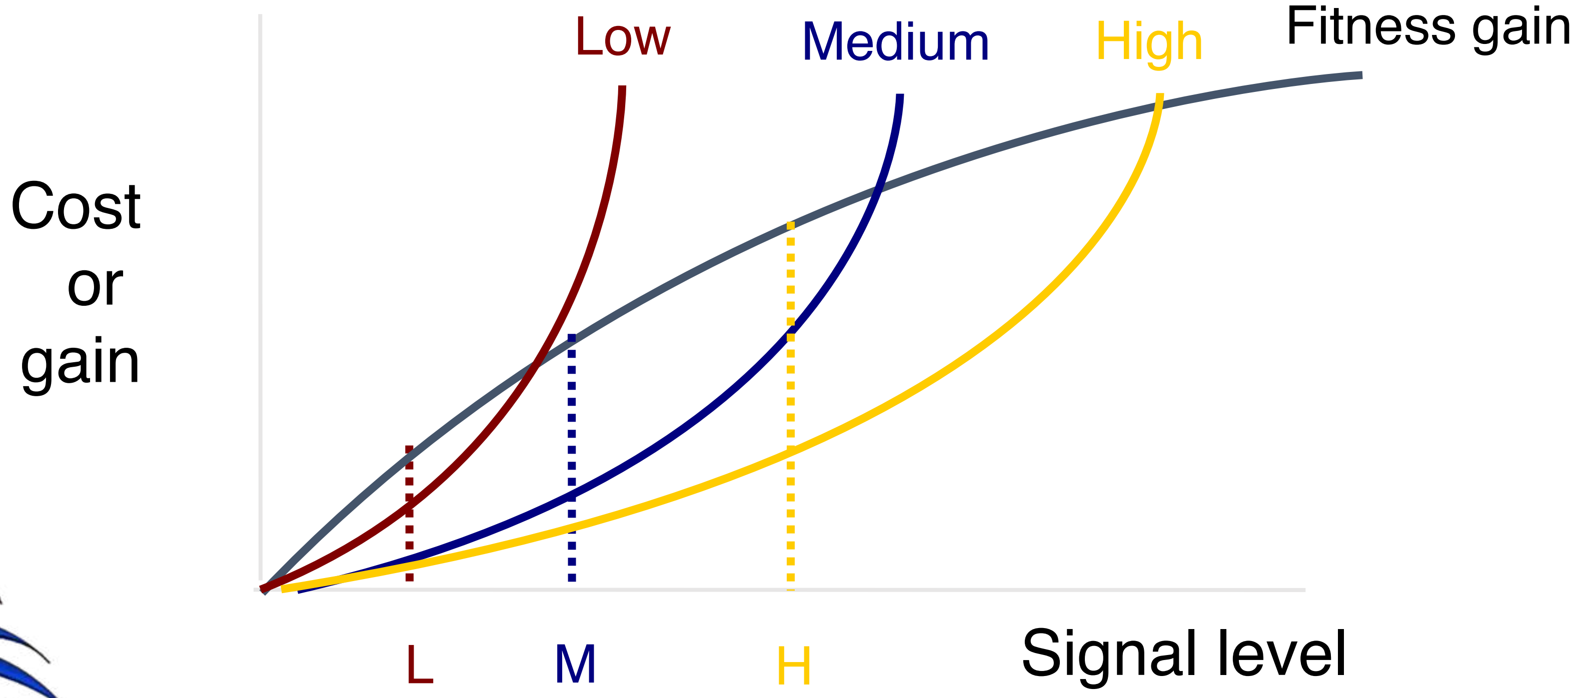

Lachmann, Szamado, and Bergstrom (2001)

But where does the probability of being funded, $\eta(x)$, come from?

- Who the other players are.
- What they decide to do.
- How accurate the panel is at assessing quality.

Copulas: joint probability distributions with uniform marginals.

The joint distribution of actual and assessed quantile is a *copula*.

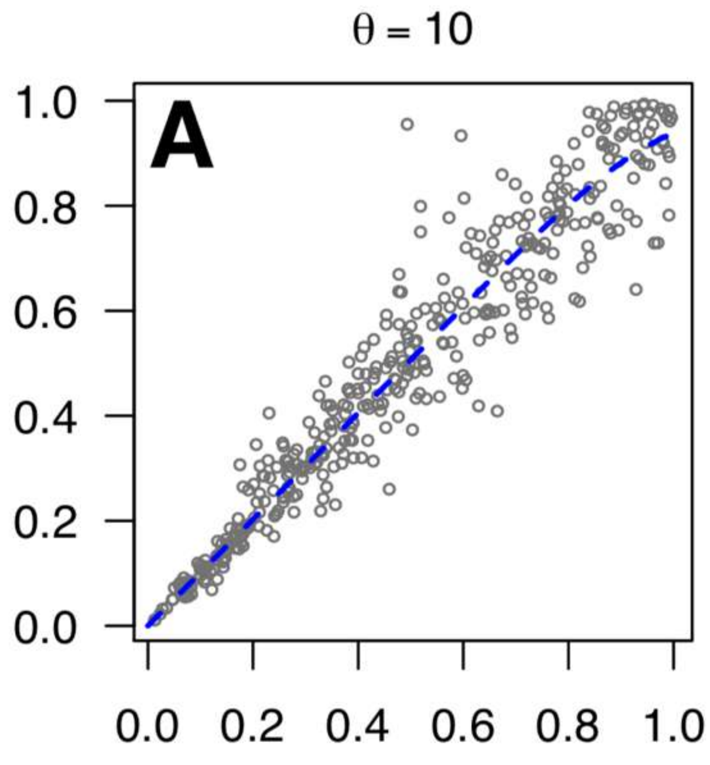

Actual quantile

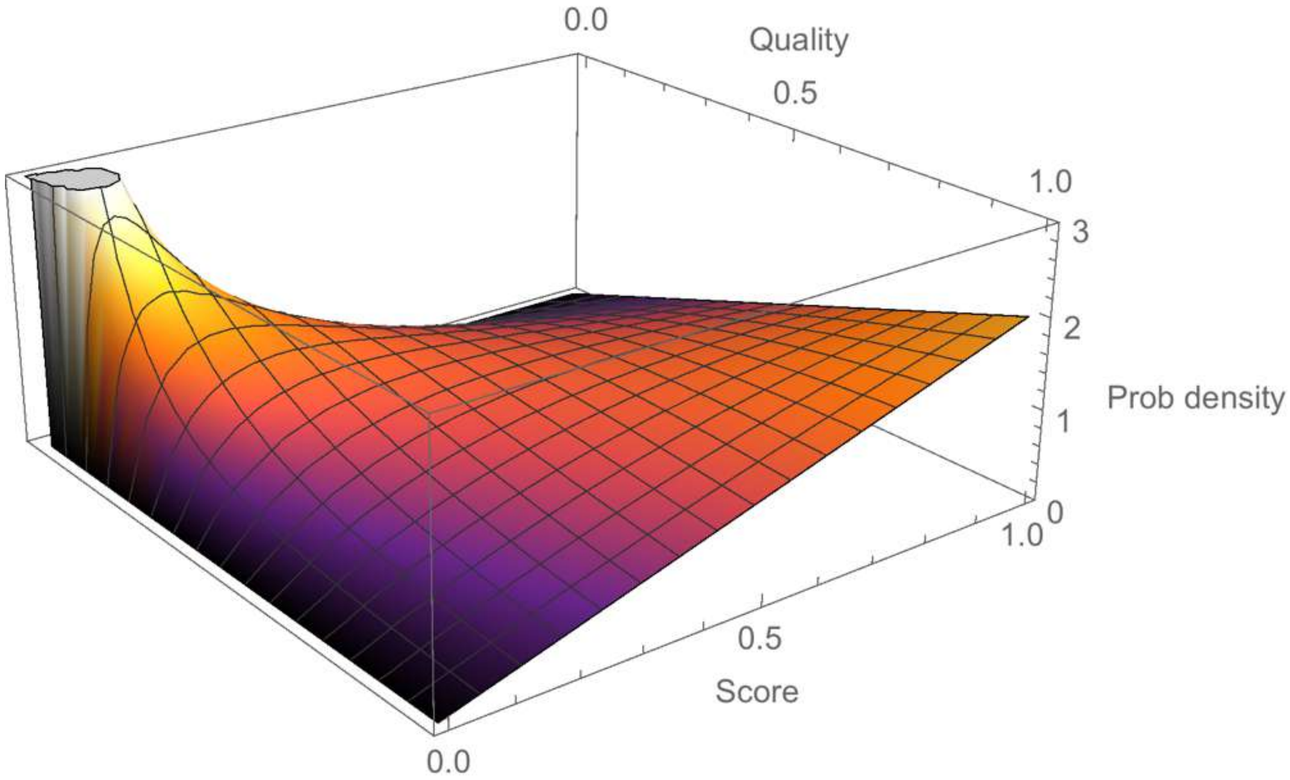

How strong of a proposal should I write?

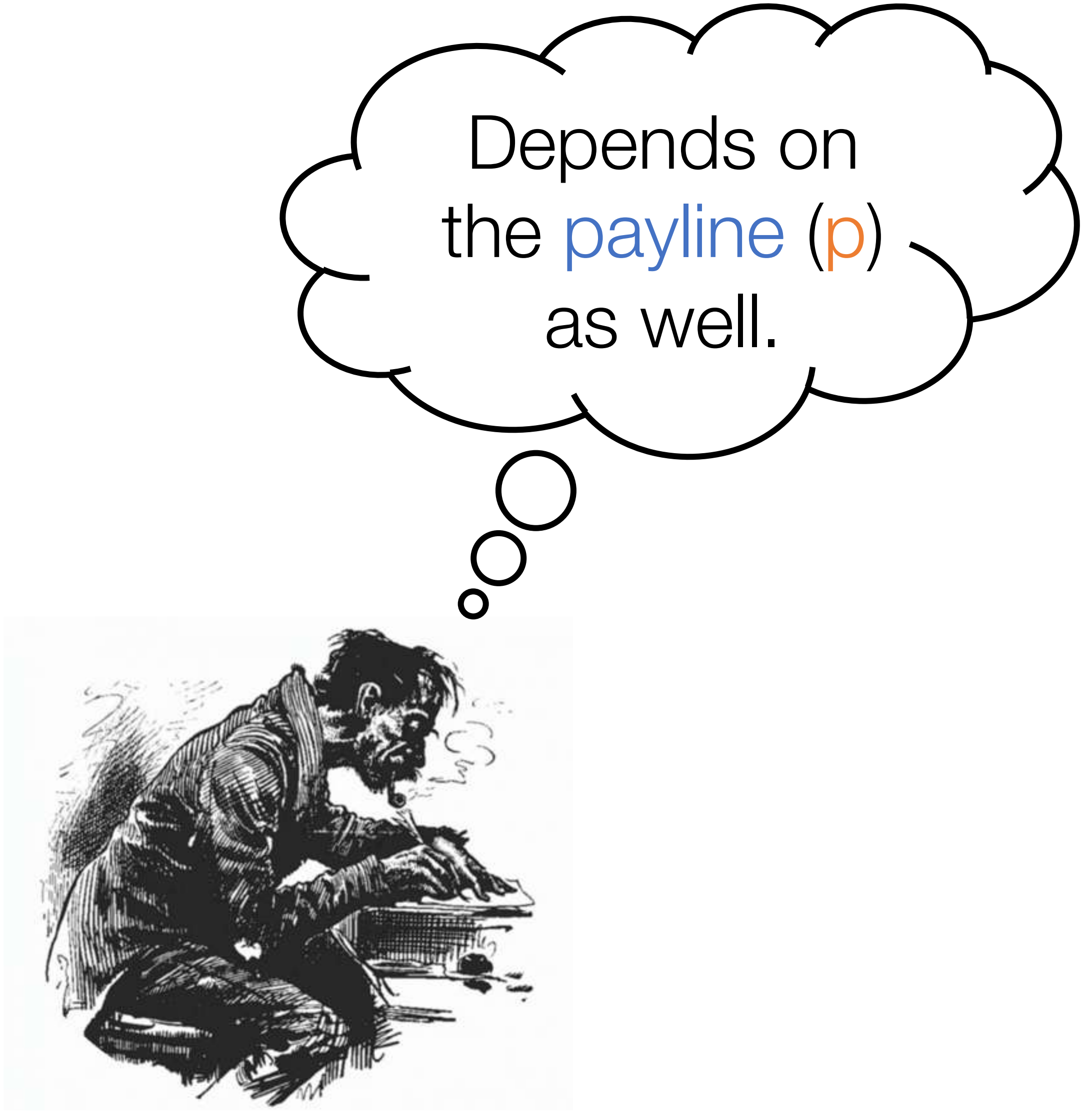

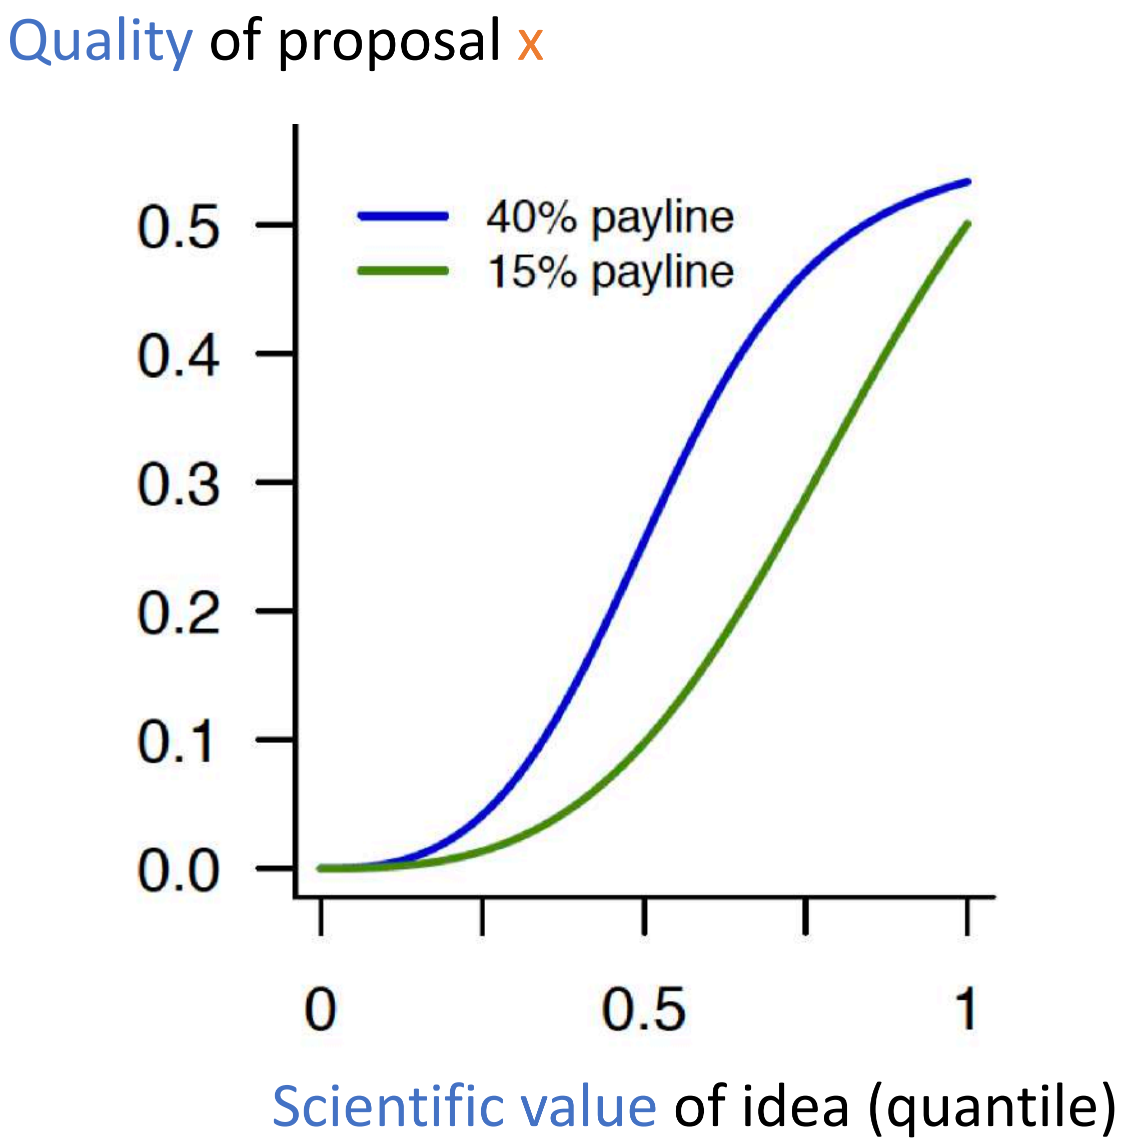

Funders might be concerned with how efficiently their money generates a scientific surplus.

So they might try to maximize something like

Return per Dollar: value v / cost c

But this ignores the cost of the competition in terms of lost scientific output.

The proper way to measure ROI is to include cost.

Return on Investment (ROI) = (v – c) / c

This is the *net gain* per dollar invested.

So what does ROI look like?

ROI increases with scientific value—but more so for high paylines where investigators don't have to work so hard to write proposals.

Return on investment

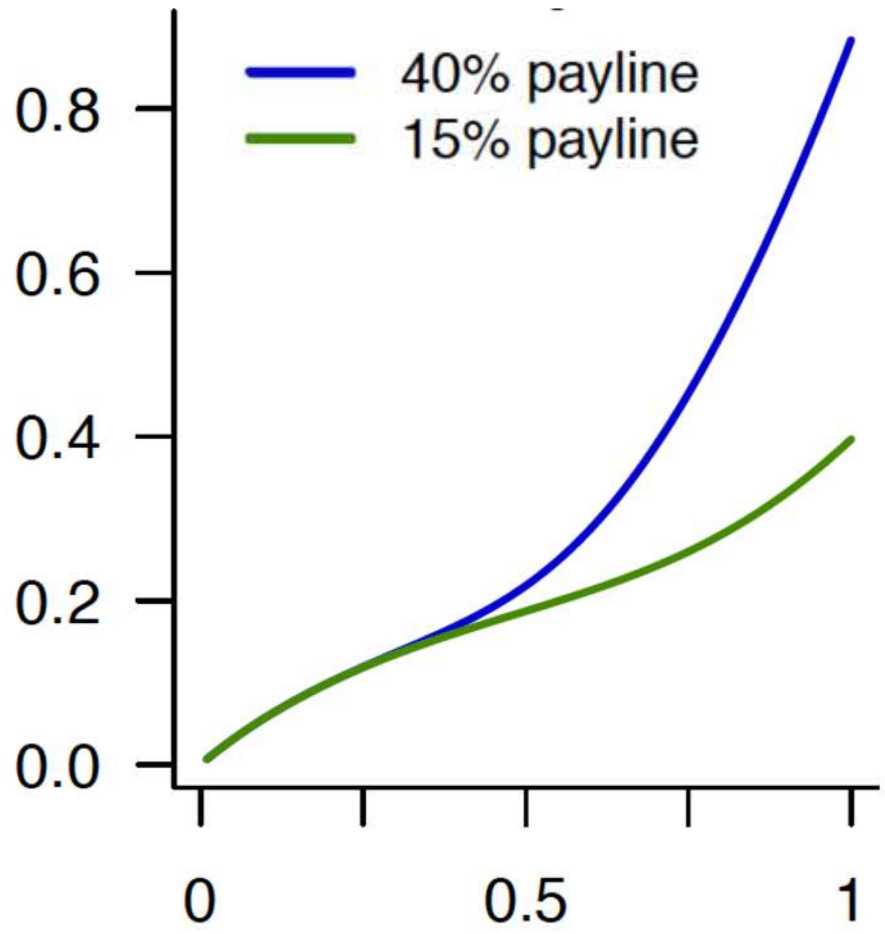

Scientific value of idea (quantile)

As paylines drop*:

- Average return to investigator decreases.
- Average scientific value of funded proposals increases.
- Total scientific ROI eventually decreases.

*So long as panels prefer better proposals to worse ones. Empirical evidence is rumored to be mixed on this point.

Grants aren’t just about the money

Hiring	Promotion
Tenure	Salary
Space	Status
Power	etc.

1968

Scientists get grants to do research.

2018

Scientists do research to get grants.

Investigator with idea *v* now wants to write a proposal of strength *x* to maximize

$$(v_0 + v)\eta(x) - (1 - k)c(v, x)$$

Private benefit

Public benefit

Private benefits of grant funding

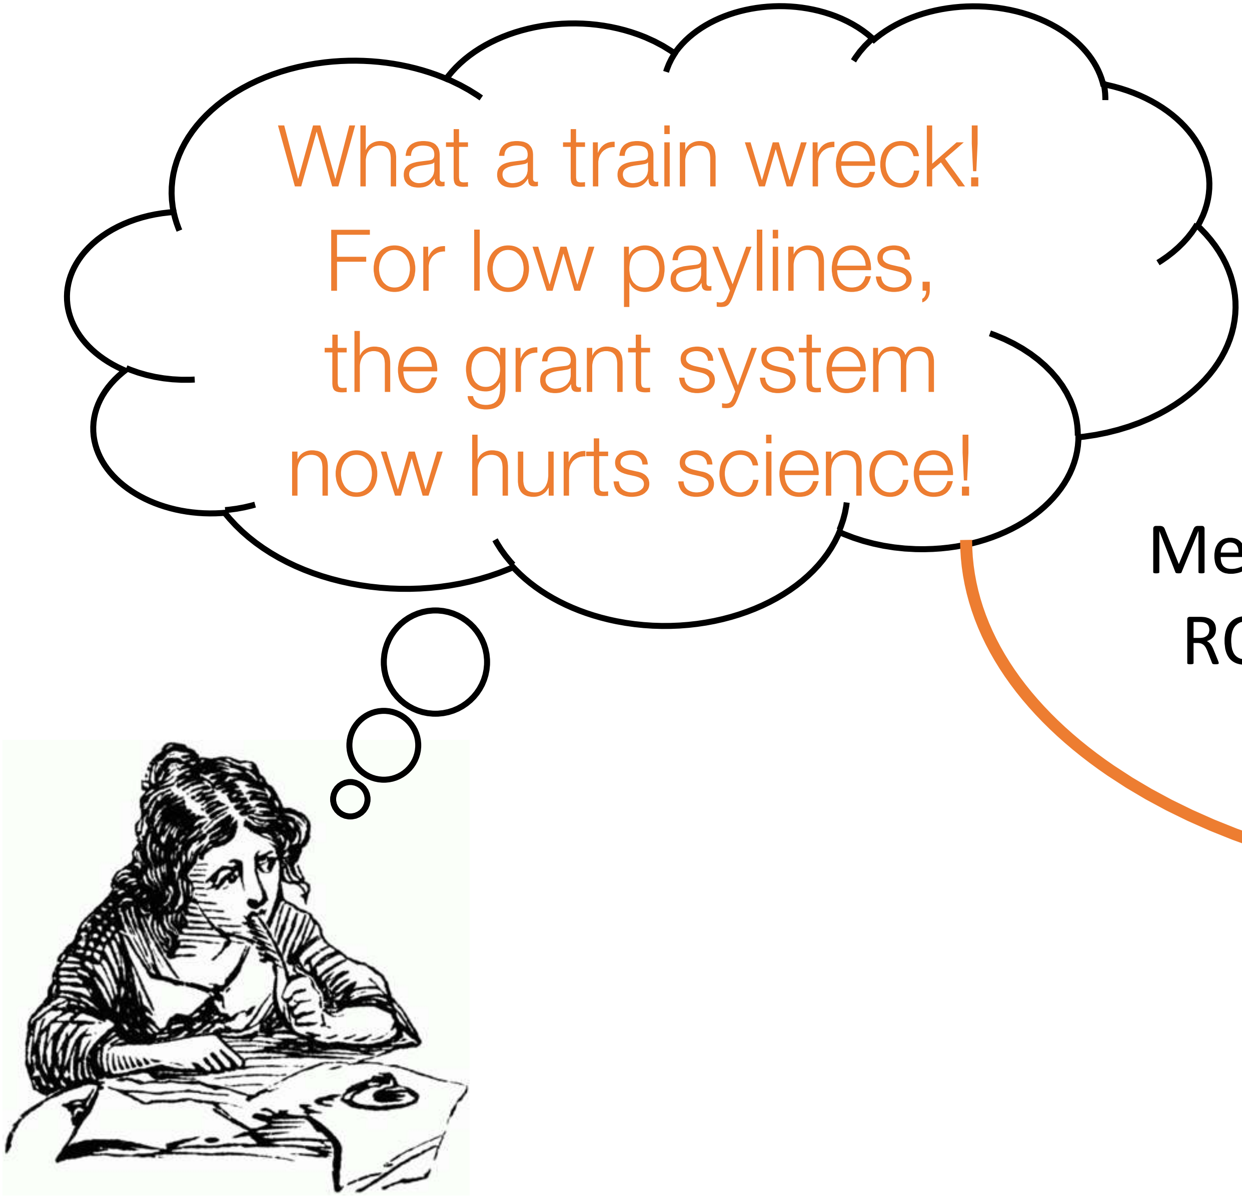

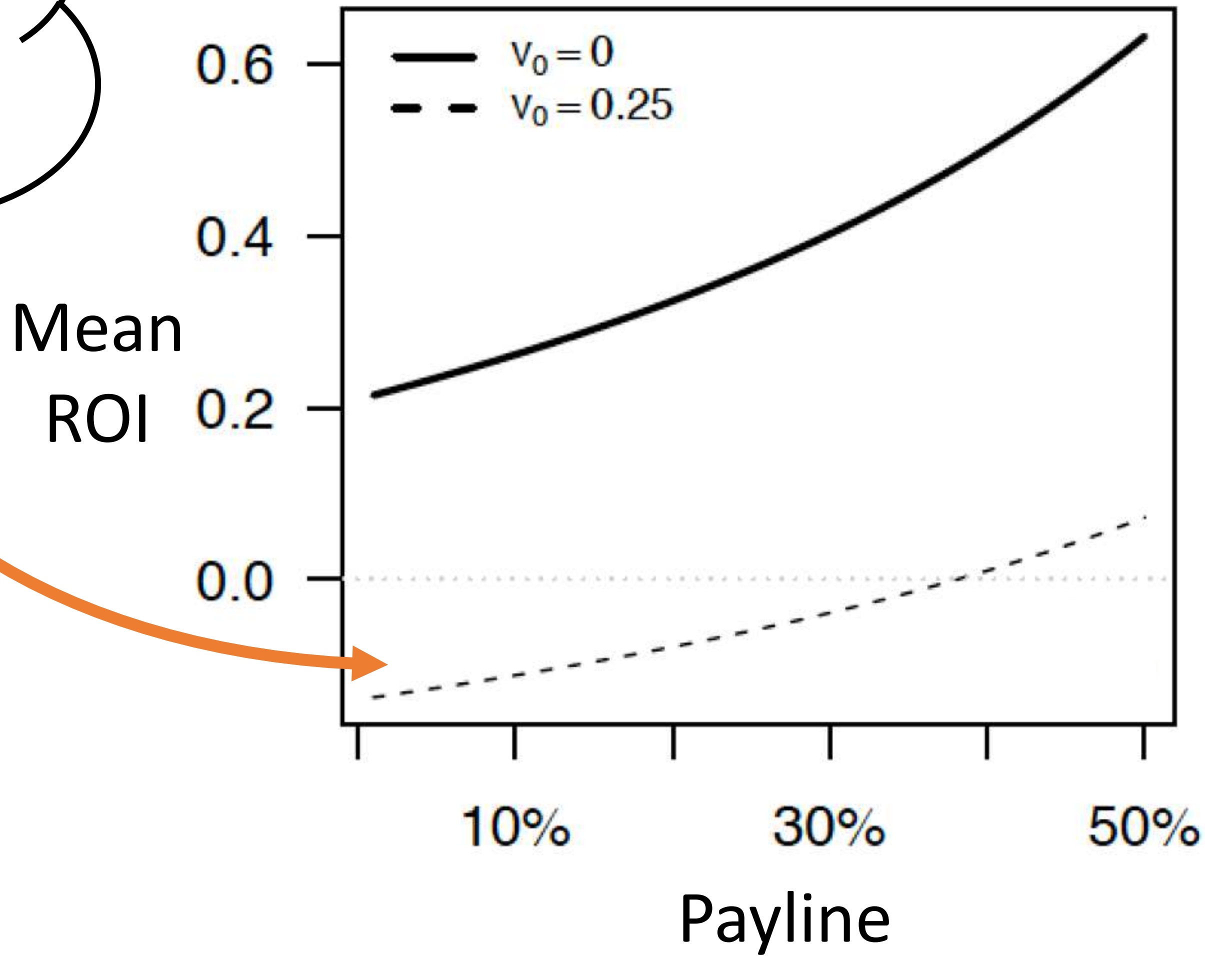

The Wall Street Journal, April 14, 2014

COMMENTARY

Taking the Powerball Approach to Funding Medical Research

Winning a government grant is already a crapshoot. Making it official by running a lottery would be an improvement.

By Ferric C. FangAnd Arturo Casadevall

April 14, 2014 7:08 p.m. ET

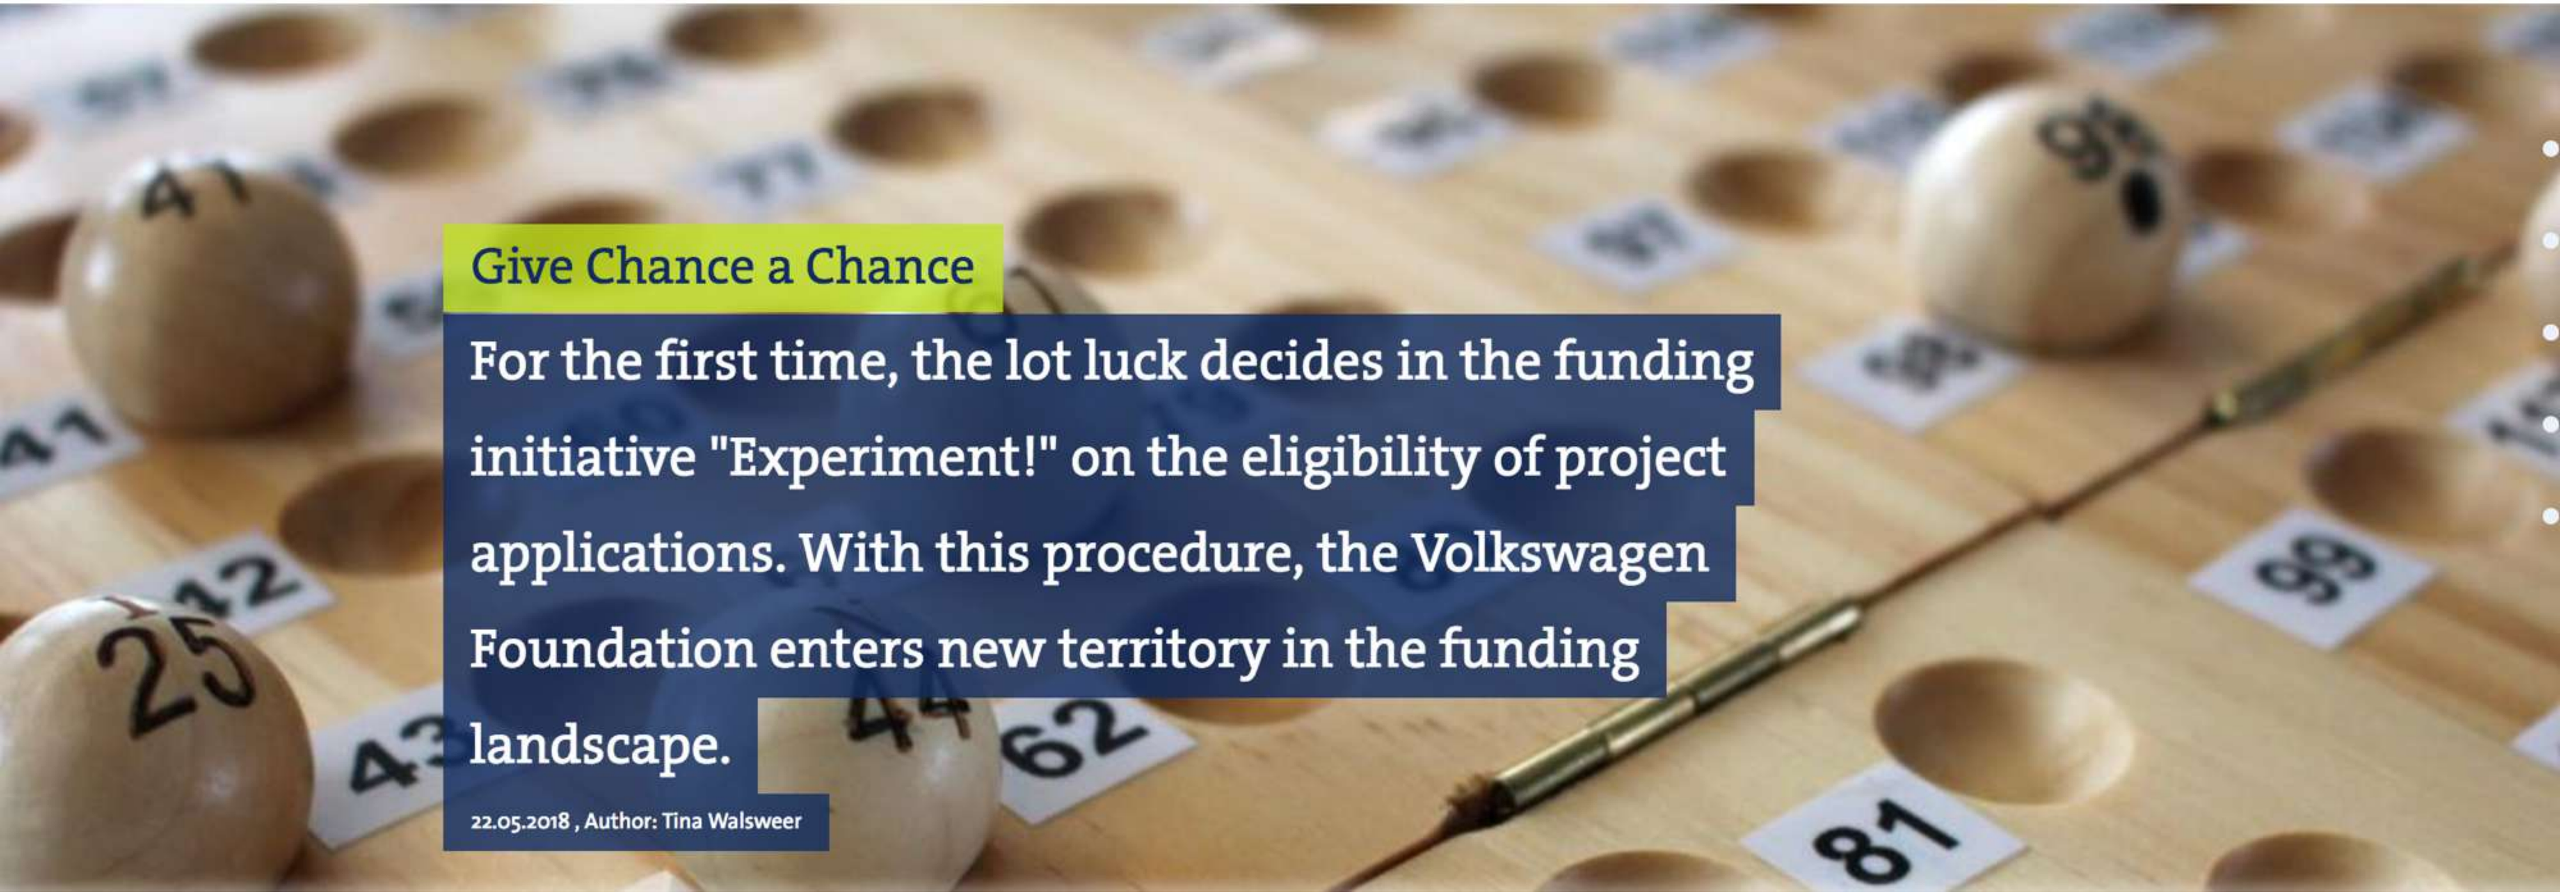

Give Chance a Chance

For the first time, the lot luck decides in the funding initiative "Experiment!" on the eligibility of project applications. With this procedure, the Volkswagen Foundation enters new territory in the funding landscape.

22.05.2018 , Author: Tina Walsweer

Partial lottery

Researchers submit proposals as before.
Proposals are scored as before.

The top L percent of applicants receive not a grant, but a lottery ticket for a possible grant.

We call L the *lottery line*.

Benefits of lotteries

AMERICAN SOCIETY FOR MICROBIOLOGY | mBio®

EDITORIAL

Research Funding: the Case for a Modified Lottery

Ferric C. Fang,[a] Editor in Chief, *Infection and Immunity*, **Arturo Casadevall,**[b] Founding Editor in Chief, *mBio*

Departments of Laboratory Medicine and Microbiology, University of Washington School of Medicine, Seattle, Washington, USA[a]; Department of Molecular Microbiology and Immunology, Johns Hopkins Bloomberg School of Public Health, Baltimore, Maryland, USA[b]

ABSTRACT **The time-honored mechanism of allocating funds based on ranking of proposals by scientific peer review is no longer effective, because review panels cannot accurately stratify proposals to identify the most meritorious ones. Bias has a major influence on funding decisions, and the impact of reviewer bias is magnified by low funding paylines. Despite more than a decade of funding crisis, there has been no fundamental reform in the mechanism for funding research. This essay explores the idea of awarding research funds on the basis of a modified lottery in which peer review is used to identify the most meritorious proposals, from which funded applications are selected by lottery. We suggest that a modified lottery for research fund allocation would have many advantages over the current system, including reducing bias and improving grantee diversity with regard to seniority, race, and gender.**

> **The lottery is in the business of selling people hope, and they do a great job of that.** —John Oliver (1)

The American research establishment has been facing the most prolonged funding crisis in its history. After a doubling in funding at the turn of the 20th century, the budget of the National Institutes of Health (NIH) was flat from 2003 to 2015, translating into a 25% reduction in actual buying power after taking inflation and the increasing costs of research into account (2). Although the increased NIH support in the 2016 spending bill is welcome news (3), this does not alter long-term uncertainty regarding the federal commitment to scientific research. The research funding crisis has been paralleled by other problems in science, including concerns about the reliability of the scientific literature, demographic imbalances, and various antiscience campaigns that question evolutionary theory, the usefulness of vaccines, human impact on climate change, and even the occurrence of the moon landings. What is perhaps most remarkable in this time of crisis and change is how little scientific leaders and governmental officials have done to combat these trends. Although each of these problems merits its own essay, we focus here on the allocation of U.S. biomedical research funds by the NIH. Specifically, we provide a detailed jus-

the overwhelming majority of the NIH budget, is allocated by a mechanism of prospective peer review in which scientists must write grant proposals detailing future work that are reviewed and criticized by a panel of experts known as a study section. The difference in funding mechanisms used by the intramural and extramural programs is significant because it shows that there is already some flexibility in the approach used by the NIH to distribute its research dollars. In this essay, we will focus on the prospective peer review mechanism used to allocate funds to extramural investigators. The fundamentals of NIH extramural peer review have not changed in a half-century. The process involves writing a proposal that is reviewed by a panel of "peers" and assigned a priority score that is converted to a percentile ranking. The NIH then funds proposals depending on the amount of money available, with the payline being that percentile ranking up to which funding is possible. At the time that the system was designed, paylines exceeded 50% of the grant applications received. However, in recent decades there has been a precipitous drop in the proportion of grants that are funded. Today's paylines and success rates are at historically low levels, hovering at around 10% in some institutes. Despite a drastic reduction in the likelihood of funding success, the essential features of NIH peer review and

- Reduce bias; increase diversity
- Reduce nepotism
- Fund more high-risk research

nature
International journal of science

RESEARCH GRANTS

Europe's top funder shows high-risk research pays off

European Research Council publishes third impact assessment of the projects it supports.

The independent review, undertaken in 2017, assessed 223 completed ERC projects that had ended by mid-2015. It deemed that 79% of them achieved a major scientific advance, 19% of which were considered fundamental breakthroughs. That proportion rose to 27% for ERC Advanced Grants, which are awarded to experienced researchers. Only 1% of the total were judged to have made no appreciable scientific contribution. The review was published on 31 May.

Benefits of lotteries

AMERICAN SOCIETY FOR MICROBIOLOGY | mBio®

EDITORIAL

Research Funding: the Case for a Modified Lottery

Ferric C. Fang,[a] Editor in Chief, *Infection and Immunity*, **Arturo Casadevall,**[b] Founding Editor in Chief, *mBio*

Departments of Laboratory Medicine and Microbiology, University of Washington School of Medicine, Seattle, Washington, USA[a]; Department of Molecular Microbiology and Immunology, Johns Hopkins Bloomberg School of Public Health, Baltimore, Maryland, USA[b]

ABSTRACT **The time-honored mechanism of allocating funds based on ranking of proposals by scientific peer review is no longer effective, because review panels cannot accurately stratify proposals to identify the most meritorious ones. Bias has a major influence on funding decisions, and the impact of reviewer bias is magnified by low funding paylines. Despite more than a decade of funding crisis, there has been no fundamental reform in the mechanism for funding research. This essay explores the idea of awarding research funds on the basis of a modified lottery in which peer review is used to identify the most meritorious proposals, from which funded applications are selected by lottery. We suggest that a modified lottery for research fund allocation would have many advantages over the current system, including reducing bias and improving grantee diversity with regard to seniority, race, and gender.**

> **The lottery is in the business of selling people hope, and they do a great job of that.** —John Oliver (1)

The American research establishment has been facing the most prolonged funding crisis in its history. After a doubling in funding at the turn of the 20th century, the budget of the National Institutes of Health (NIH) was flat from 2003 to 2015, translating into a 25% reduction in actual buying power after taking inflation and the increasing costs of research into account (2). Although the increased NIH support in the 2016 spending bill is welcome news (3), this does not alter long-term uncertainty regarding the federal commitment to scientific research. The research funding crisis has been paralleled by other problems in science, including concerns about the reliability of the scientific literature, demographic imbalances, and various antiscience campaigns that question evolutionary theory, the usefulness of vaccines, human impact on climate change, and even the occurrence of the moon landings. What is perhaps most remarkable in this time of crisis and change is how little scientific leaders and governmental officials have done to combat these trends. Although each of these problems merits its own essay, we focus here on the allocation of U.S. biomedical research funds by the NIH. Specifically, we provide a detailed jus-

the overwhelming majority of the NIH budget, is allocated by a mechanism of prospective peer review in which scientists must write grant proposals detailing future work that are reviewed and criticized by a panel of experts known as a study section. The difference in funding mechanisms used by the intramural and extramural programs is significant because it shows that there is already some flexibility in the approach used by the NIH to distribute its research dollars. In this essay, we will focus on the prospective peer review mechanism used to allocate funds to extramural investigators. The fundamentals of NIH extramural peer review have not changed in a half-century. The process involves writing a proposal that is reviewed by a panel of "peers" and assigned a priority score that is converted to a percentile ranking. The NIH then funds proposals depending on the amount of money available, with the payline being that percentile ranking up to which funding is possible. At the time that the system was designed, paylines exceeded 50% of the grant applications received. However, in recent decades there has been a precipitous drop in the proportion of grants that are funded. Today's paylines and success rates are at historically low levels, hovering at around 10% in some institutes. Despite a drastic reduction in the likelihood of funding success, the essential features of NIH peer review and

- Reduce bias; increase diversity
- Reduce nepotism
- Fund more high-risk research
- Reduce peer review effort
- Make underfunding transparent

- Inter-rater reliability is low anyway
- And predictive ability is poor.

- Reduce effort in proposal preparation.

Proposition: In a lottery, the return to the investigator and the return to the community are set by the lottery line, and are independent of the payline.

Effect of the lottery line

Average scientific ROI

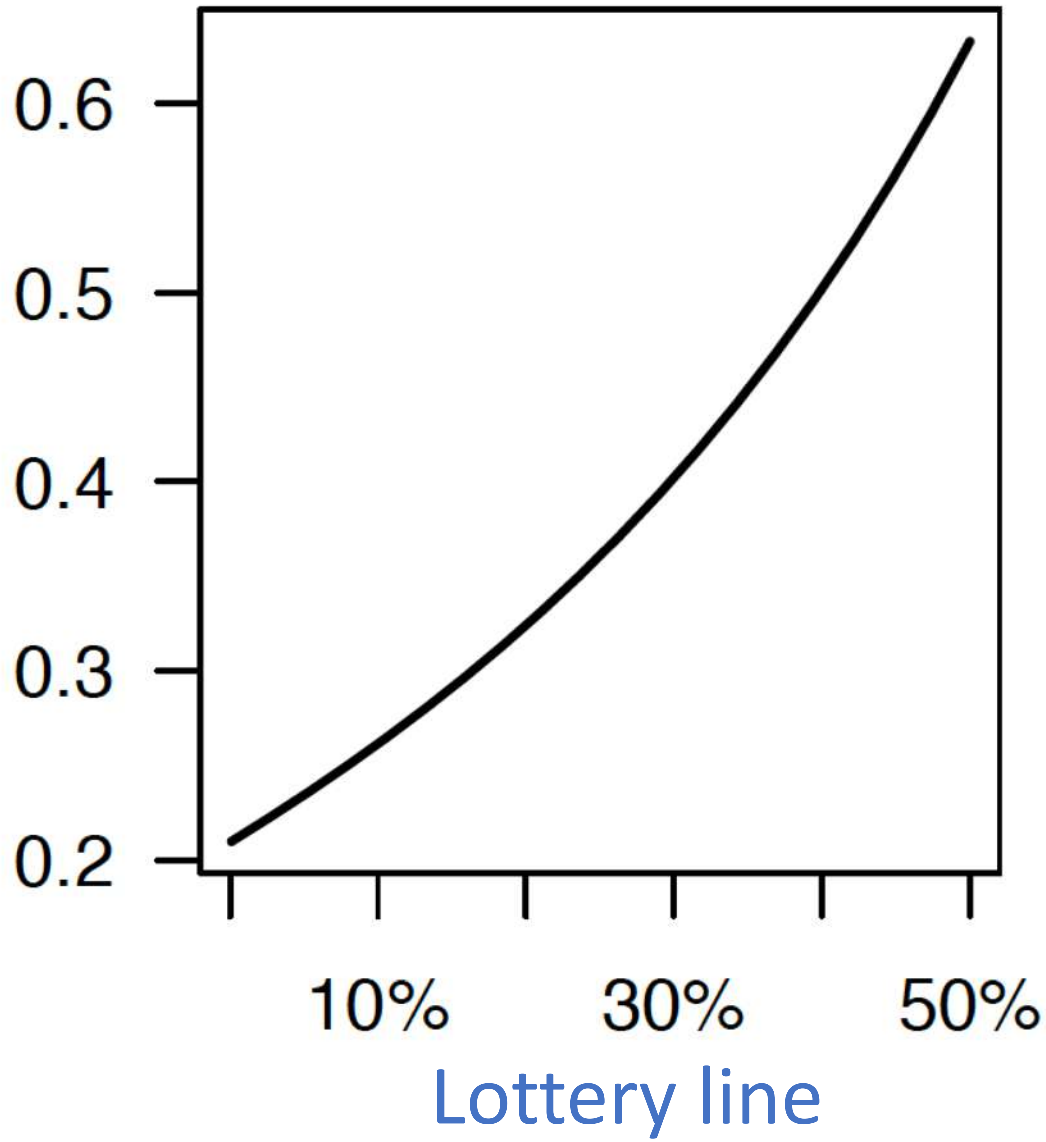

We can capture the efficiency benefits of a high payline by replacing a sure payout with a lottery.

Additionally this weakens the value of the grand award for assessment purposes, reducing the overall investment in grant preparation due to private benefits.

Lotteries may be politically untenable

Switching to lotteries may drive investigators to prepare more grants.

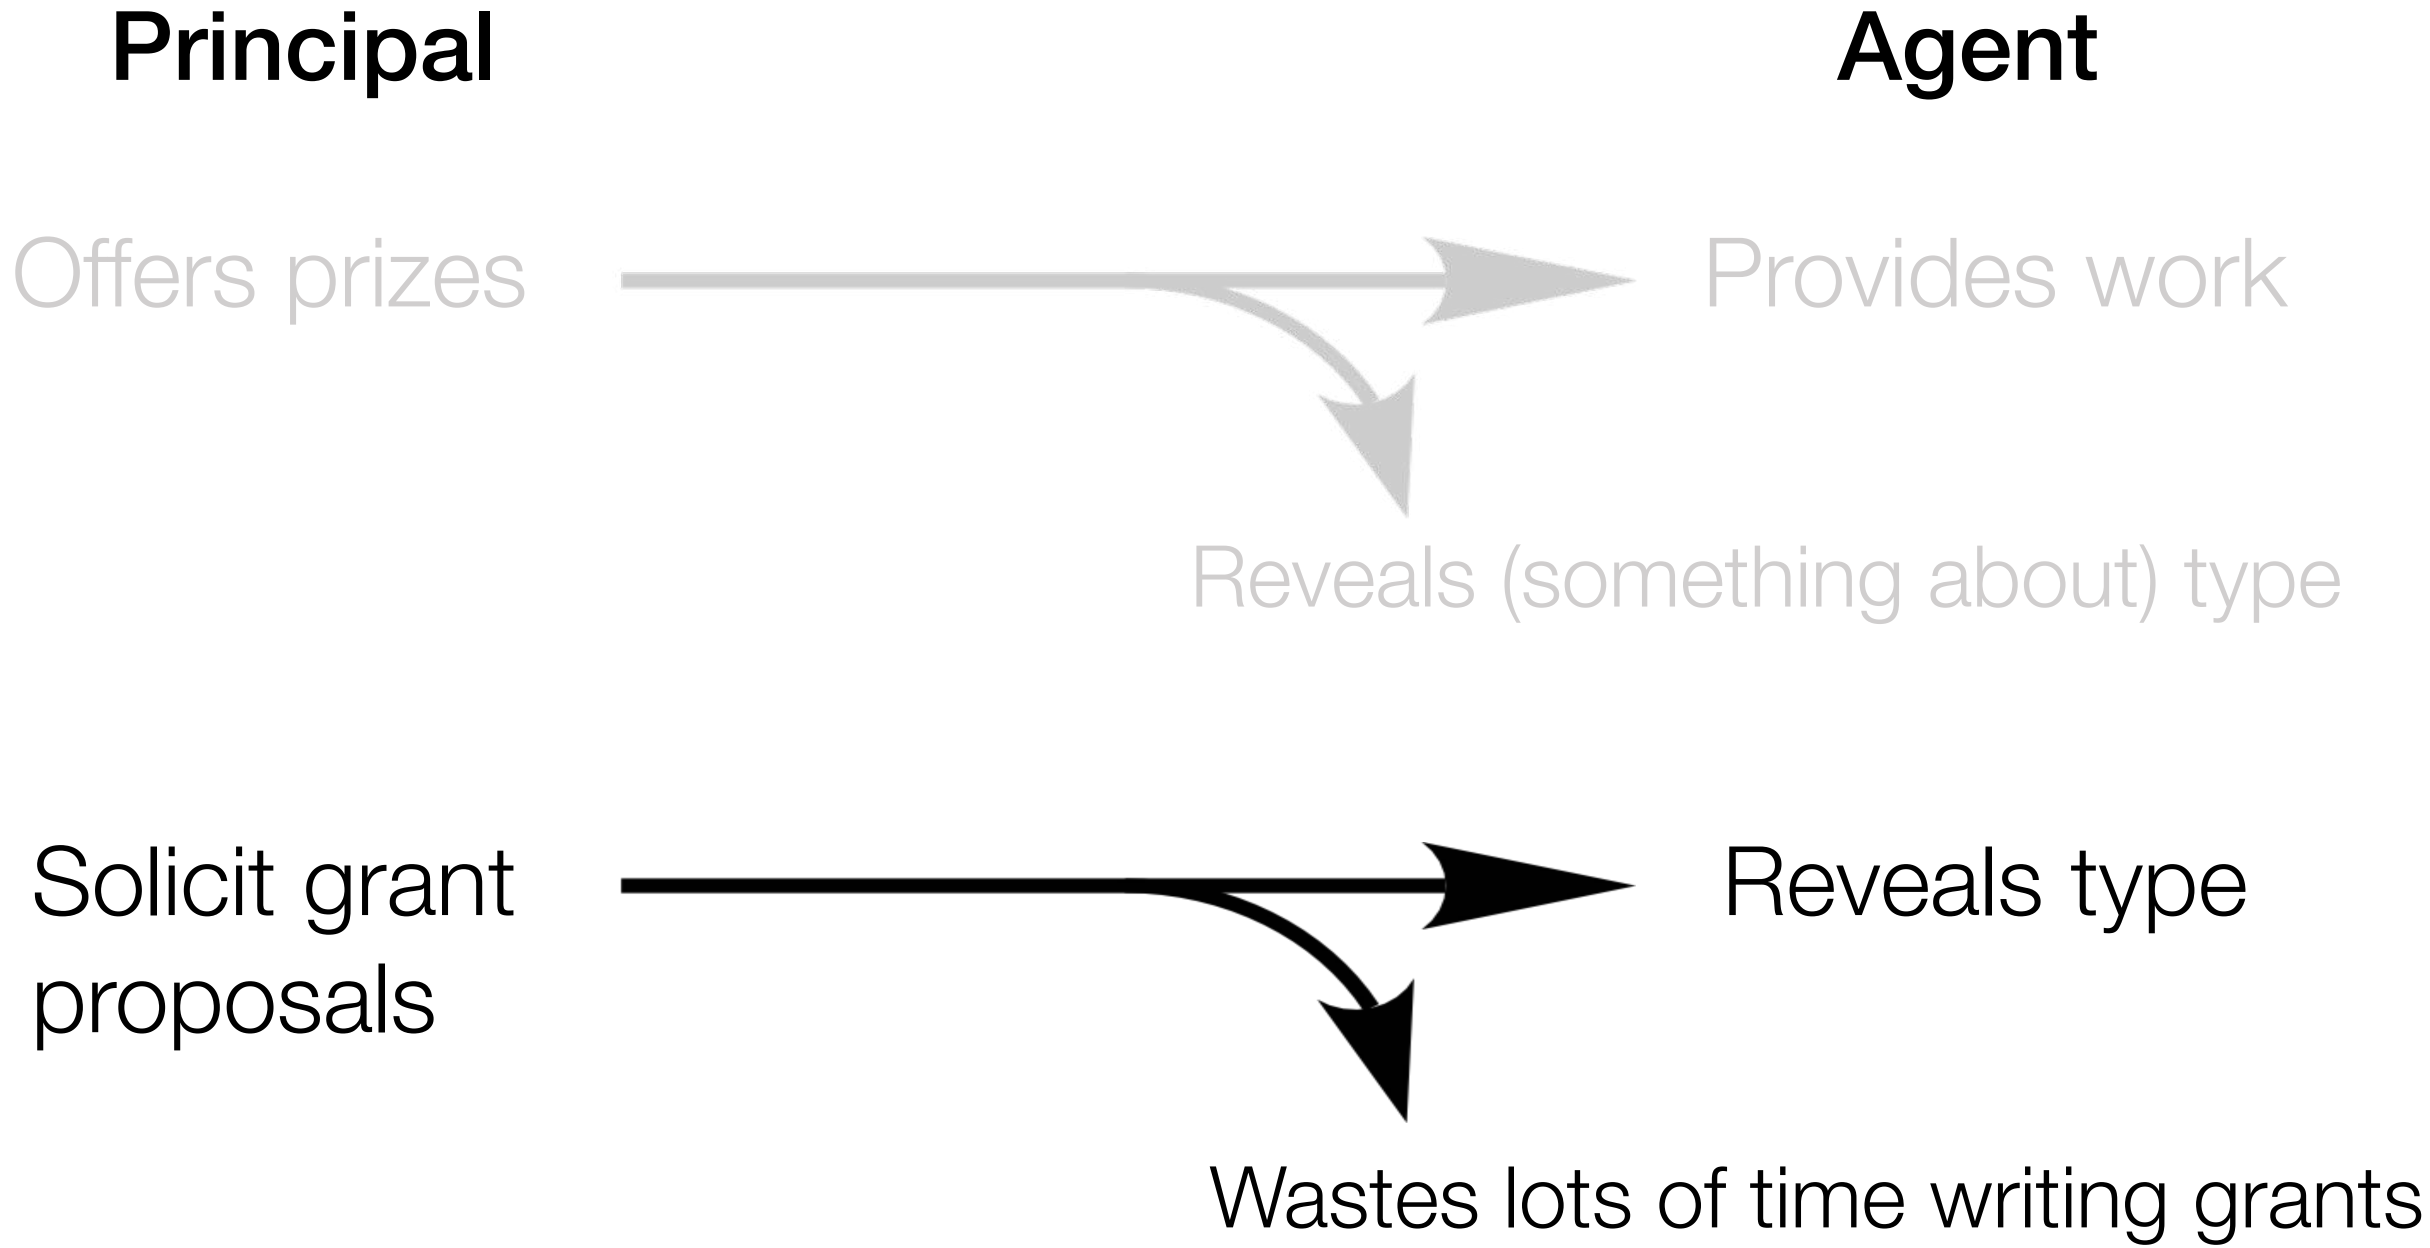
Principal
Agent
Offers prizes
Provides work
Reveals (something about) type
Solicit grant proposals
Reveals type
Wastes lots of time writing grants

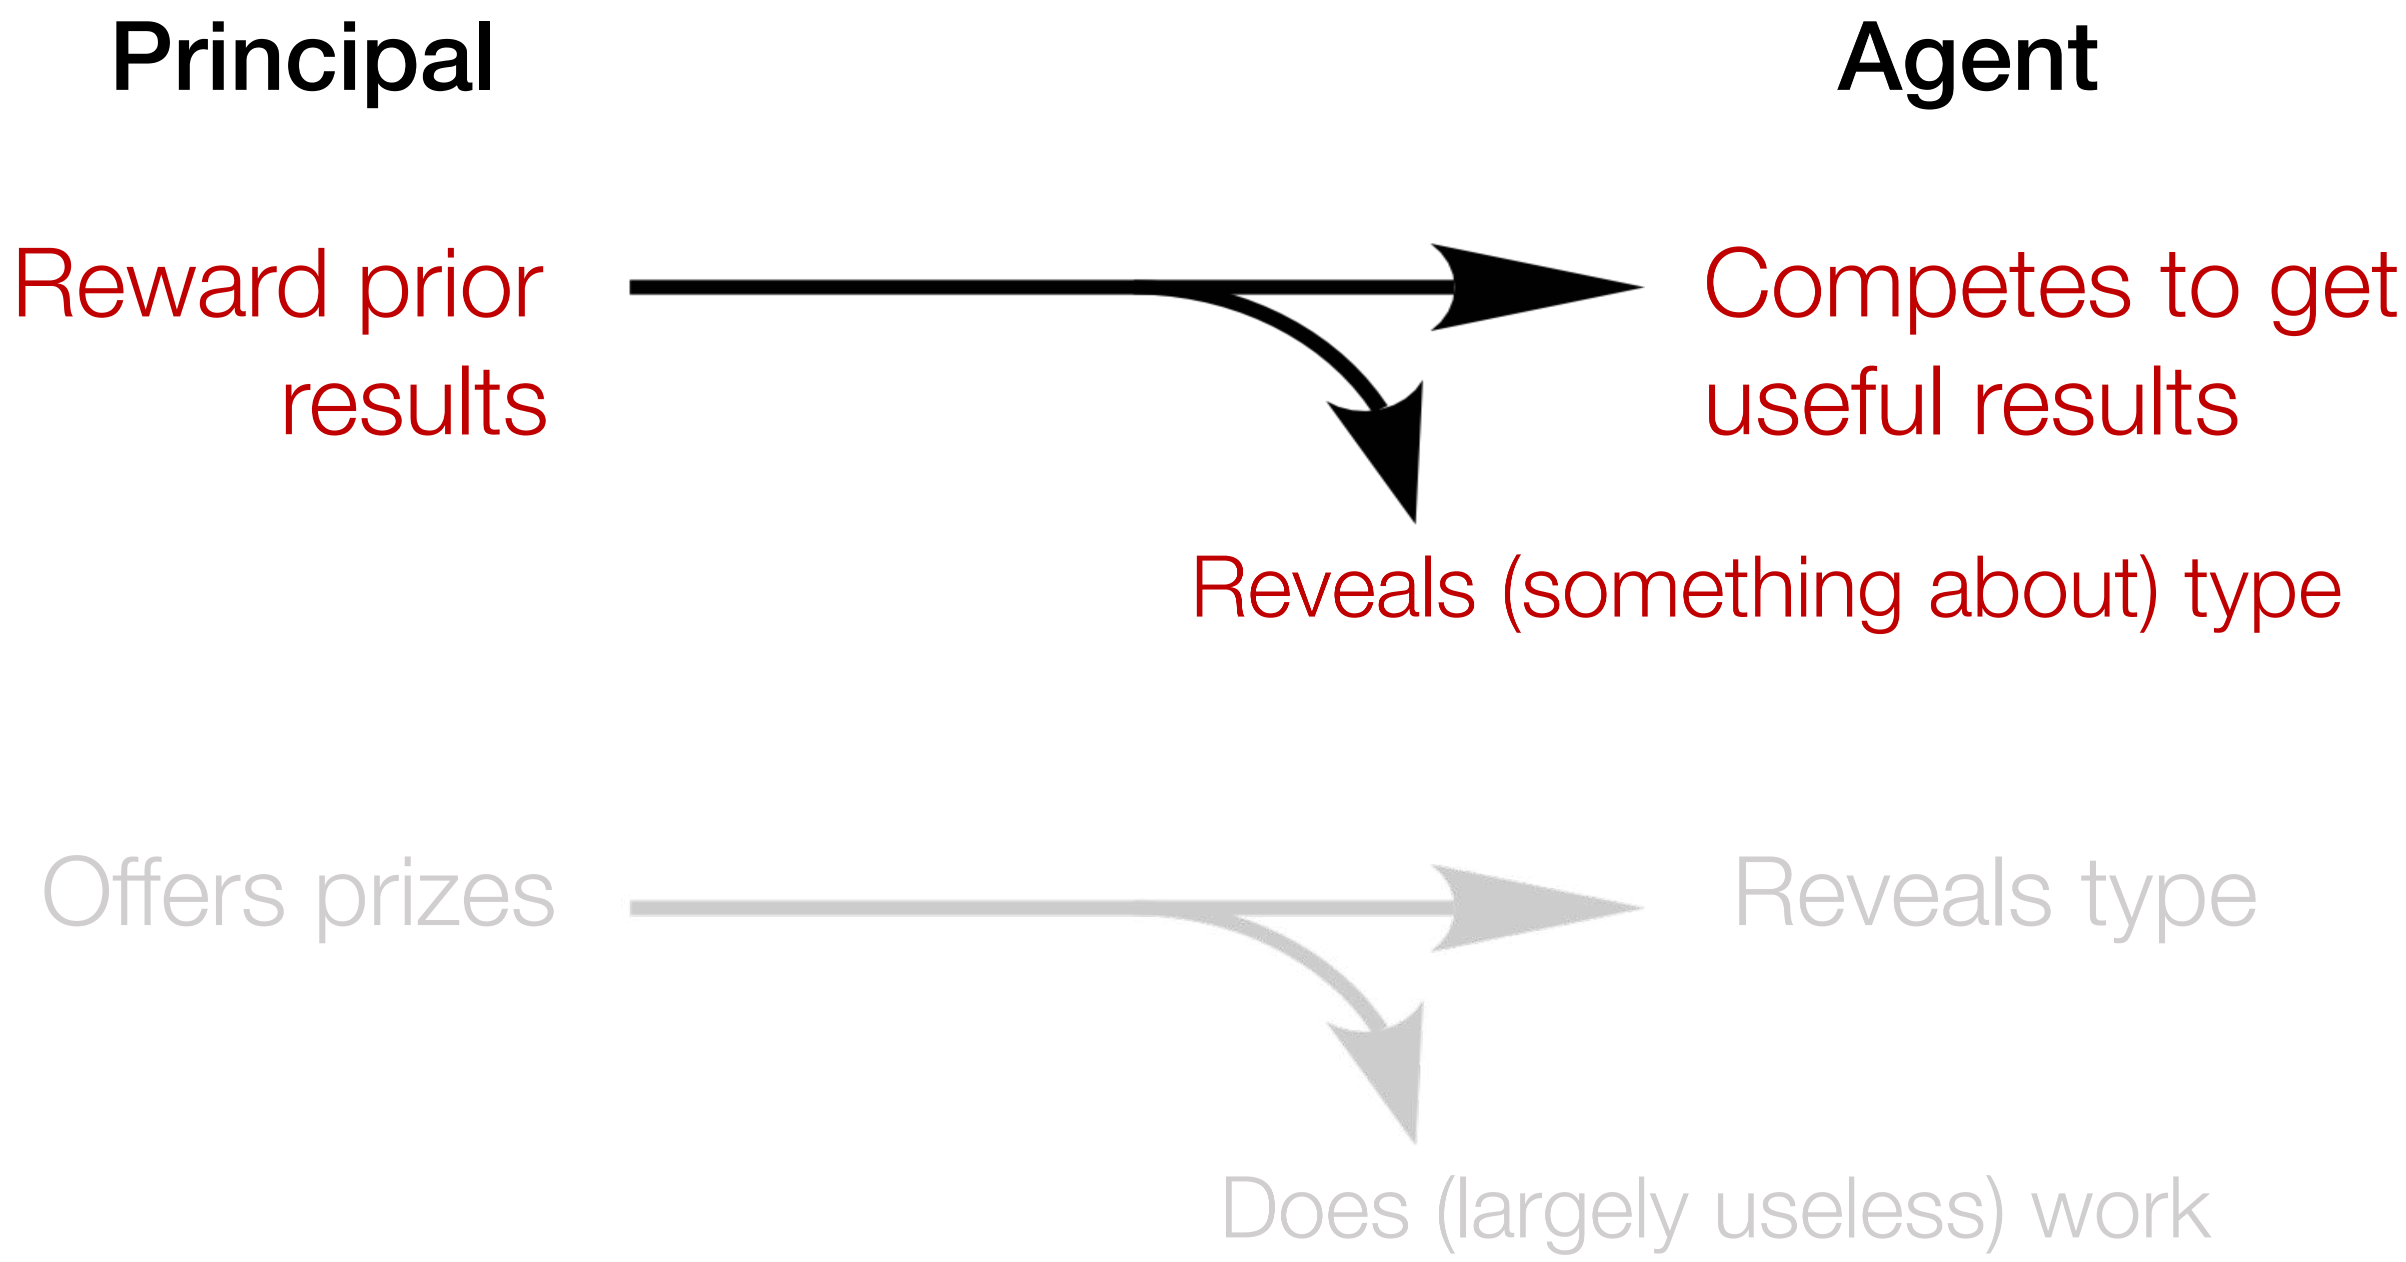
Principal
Agent
Reward prior results
Competes to get useful results
Reveals (something about) type
Offers prizes
Reveals type
Does (largely useless) work

So what does all this tell us?

- Using a (mostly useless) contest as screening mechanism is inefficient.
- Private benefits to grants make the problem worse. Some funding programs could be net negatives for science.
- We illustrate the mathematical logic behind using a partial lottery system to reduce costs of grant preparation.
- If we want to rationally design a proposal-based funding system, it is useful to think within the framework of contest theory.

Grant proposal contests as all-pay auctions

Revenue equivalence

In a single-object, private-value auction, all auction designs that

> allocate object to the highest bidder, and
>
> allow individuals to not bid at all

generate the same expected revenue. This includes first and second price auctions, but also all-pay auctions.

Revenue equivalence

A standard English auction sells to the top bidder at the second bidder’s price.

As the number of bidders gets large (for reasonable distributions), the difference between the first and second bidders’ values goes to zero. This means there is no consumer surplus.

Translated to grants, this means as the number of bidders gets large and the payline gets small, scientific cost approaches scientific value and there is no net gain even in the absence of private benefits to getting funded.