



Market Failures in Science

Milan Cvitkovic

Market failure: when the allocation of resources by a market doesn't produce as much overall value as it could.

Science: the part of society focused on understanding the natural world and creating new technologies.

I. Information Asymmetries

i. Knowledge monopolies

Scientists are incentivized to withhold as much information as possible.

Scientists are incentivized to, and often do, [withhold as much information as possible](#) about their innovations in their publications to maintain a monopoly over future innovations. This slows the overall progress of science.

For example, a chemist who synthesizes a new molecule will publish that they have done so in order to be rewarded for their work with a publication. But in the publication they will describe their synthesis method in as minimal detail as they can while still making it through peer review. This forces other scientists to invest time and effort to reproduce their work, giving the original author a head-start in developing the next, better synthesis.

This is not necessarily malicious. Effective communication of ideas requires effort that the scientist may rather spend doing something else. Most scientists will divulge details of their experiments when asked in personal communication, but their incentive is still to respond as slowly and cryptically as possible.

If there were a way to reward scientists in proportion to the ease of replicating their results that would be better, but it would have to outweigh the reward of a monopoly on future publications.

ii. Mathwashing

Scientists are incentivized to make their work as inscrutable as possible.

Scientists often use overly complex mathematical or statistical methods to obscure problems from reviewers. This works because most reviewers will not admit to not understanding details of work they are reviewing. The harder it is to spot an error, the



less foolish you are for having failed to spot it. Unreadable academic prose can also serve a similar obfuscatory purpose, though this seems to be a bigger problem in [the humanities](#).

II. Poor Credit Assignment

Credit is the primary reward for a scientist. Credit accrues into influence, which is valued in its own right, and is also convertible into funding.

Credit is given in science by priority: if you publish an idea or discovery first, you get the credit.

This system successfully incentivizes the creation of public goods, which is impressive. But it's not perfect.

i. Scooping and duplicate work

Science wastes a lot of effort by not coordinating between groups.

When scientists don't coordinate it can lead to duplicate work. If two groups are interested in the same project, ideally they would divide up the relevant work between their members and share the credit. This often doesn't happen because groups typically only publicize what they're working on after they've finished and published the work. This leads multiple groups to do nearly identical work simultaneously. The first group to publish, "scooping" the others, usually gets all or most of the credit. Even [Darwin](#) thought this was a bad system.

Obviously we want some work in science to be duplicated: the lack of duplicate work is a cause of the replication crisis. But non-observational work like theory or method development gains almost nothing from groups scooping each other.

Why don't scientists coordinate more? Partly because broadcasting what one is working on would let competitors free-ride on your ingenuity. They could even claim they were working on the same idea when they weren't. Another reason is that it's hard to evenly divide labor between groups. Plus the more authors are on a paper, the less credit each receives. If you publish before your competitor, you don't have to share credit with them, though of course you might lose the race. At some career stages, having solo or majority credit for an idea matters a lot, like when working on a job-market paper. This can even prevent people from the same lab from collaborating with each other.

One way to prevent coordination failures are turf norms. Some fields like geology or parts of biology have norms about not working in areas (literally, for the geologists) that other people study. But forcing everyone to find a separate niche is inefficient: if everyone in a field thinks a particular idea is the most promising thing to work on, we want them all to be able to coordinate on doing so!

Based on an analysis of PDB submissions, [Hill and Stein \(2019\)](#) suggest that duplicate



work doesn't have a huge impact on the careers of scientists themselves, despite the waste of scientific effort. The same authors went on to study the impacts of competition on research quality, again using PDB data, with structure resolution as a metric for quality ([Hill and Stein, 2021](#)). It's a great idea, and they present a very detailed model, but PDB crystallography is also an unusual "minigame" in science where groups are all trying to produce the exact same thing. In most of science, when you get scooped you pivot (often in a direction you don't really care about), or dress your work up to look unique, or try to tack more on for novelty. People almost never publish purely duplicate work, unless the projects are so contemporaneous that they make it into the same "release cycle". I don't know how one could quantify that kind of inefficiency.

One solution to this problem is working in public, but that opens up opportunities for flag-planting, and researchers often find working in public embarrassing since flaws aren't hidden. Another option would be a trusted neutral party to whom everyone could tell what they're working on (along with appropriate proof of investment in the idea) and groups duplicating each other's work could be put in touch.

ii. Author lists and contributions

A meritocracy is only as good as its ability to assign merit, and science is awful at this.

Every field has different conventions, but generally authors on scientific papers are listed from greatest to least contribution. The first author did the most work, and the last author (usually the principal investigator(s) supervising the project) did the least work. People who only helped a little don't get authorship, but are thanked in an Acknowledgements section. First authorship matters more than any other position in the author list because papers are typically cited as "Lastname *et al.*" in text, so only the first author's name is seen by many more readers.

This is not a great way to apportion credit. Partly because an ordering doesn't tell you how much more work one author did than the next. But mainly because credit can't be reduced to a scalar quantity. Contributions of different kinds aren't comparable. What's the exchange rate between hours writing prose and hours building apparatus? Even if one measured contribution purely in hours spent on a specific project, how does one weigh preparatory work done on a similar previous project, or an entire career of learning?

The market failure in all this is that a meritocracy is only as good as its ability to assign merit. Science is inefficient when talent is misjudged. This inefficiency is not hypothetical — every scientist has at least one colleague whose scientific skills are dwarfed by their ability to get their name on papers, and colleagues who are perennially undercredited.

Some authors now include a section detailing precisely what each author contributed to a paper. This probably helps apportion credit in narrow groups of peers. Some authors take an absurdist approach and [declare that their paper has 19 co-first authors](#). But I



suspect that as long as there is a top-line author list, most readers will apportion credit inaccurately. In [this paper](#), even though the authors listed their names randomly, everyone in machine learning would agree that the first author has gained vastly more name recognition than the others.

iii. Citation pile-ons

Scientists have little incentive to cite accurately or fairly.

What do you do when you're writing a paper and need to cite an idea whose provenance you don't know? You see whatever reference someone else cited and cite it too.

Sometimes this is the right citation, but often an idea is developed by multiple people across multiple publications. Citing all of them is tedious, so a winner-take-all reward accrues to just one of these publications. Often the "winning" publication is the one that coins a term rather than the discovery of an idea, or more often is the publication with the most famous author.

The worst case is when people cite references that don't even contain the information they're citing. Rekdal (2014) "[Academic Citation Practice: A Sinking Sheep?](#)" is a delightful paper on the subject.

There is a small incentive for finding the right citation: if you cite an older paper than the one readers are expecting, you look smart. But this rarely outweighs the ease of citing whom everyone else cited.

III. Public Goods

Work everyone agrees someone should do but which nobody does, because it's inadequately rewarded in money or reputation.

- [Better animal and in vitro models for disease](#) (credit:effort ratio is too low for academics to work on it, too hard to protect IP for industry to work on)
- "Unsexy" studies like exhaustive parameter sweeps, calibration curves, etc.
- Open-source scientific software (a lot still gets made, but it is undersupported given its importance)
- Open-access publication, at least pre-Sci-Hub
- Open-access dataset collection/curation, esp. highly paywalled ones like SciFinder and Reaxys
- Insufficiently profitable drugs
 - [vaccines](#), at least pre-COVID
 - antimicrobials, at least until the Hyper-SARS Pandemic of 2028
 - therapies for rare diseases
- Shared fabrication facilities, [MOSIS](#) being an ideal example



IV. Organizational inflexibility

Scientific career paths, organizational structures, and funding sources are too standardized.

It's a bit suspicious that all PhD programs should take about the same amount of time, regardless of field. Or that it's so hard to fund anyone who isn't on the PhD-to-postdoc-to-prof path, like the [high-level independent contributors common in the tech world](#). Or that the entirety of a professor's career should be [planned years in advance](#) to increase the odds of getting an NIH R01 grant. Or that universities should prevent scientists from hiring engineers to build critical infrastructure by setting low salary caps.

Someone invented the existing ways we do science, and we can invent better ways. And people are. The [Overedge Catalog](#) is a great collection of new types of research organizations breaking these molds.

V. Bureaucratic Bottlenecks

It takes too long to get permission and money.

“[Drug lag](#)” is the delay between discovery of a drug and its availability to patients. The largest contributor to drug lag is the time required to obtain regulatory approval to market the drug. The justification for regulatory control of drugs is that many drugs are credence goods, but the current system is certainly not at the Pareto frontier between minimizing drug lag and maintaining customer confidence in drugs.

Similar delays exist across biomedical science. Like delays for giving scientists the Investigational Device Exemptions (IDE) they need to do research on new medical devices. This contributes to the [“1- to 3-year delays in the introduction of new device technologies into general clinical practice within the United States as compared with Europe.”](#) Or delays for determining whether a new clinical trial protocol will convince the FDA of efficacy. And in general, as [Fast Grants](#) has shown, grantmaking is also much slower than it needs to be.

Changing from a pre-approval to an escrow-until-approval or post-monitoring regulatory system could help. More explicit, quantitative regulatory standards could also help. e.g. rather than new clinical trial protocols being approved by qualitative judgment of the FDA's statisticians, have the FDA state a false-positive rate cutoff that new protocols can prove they are below (in the absence of other issues).



VI. Taxes

i. Administrative Burden

Scientists only spend 1/2 their time on active research.

A study of 11,167 principal investigators with active federal grants in 2018 found that only 55.7% of their time was spent on active research ([2018 Faculty Workload Survey](#), Schneider 2020). The rest of it went to things like grant administration and satisfying [Institutional Review Boards](#) (IRBs) and other compliance requirements. Worse, an institution's IRB is incentivized to disallow any risky (i.e. potentially valuable) research lest some lawsuit be raised.

This statistic is less worrying than it might seem, since most research is done by graduate students and postdocs who spend (much) more time doing active research than PIs. But it's clearly a deadweight loss to science.

ii. Institutional Fees

Universities get paid by the amount of grant money, not scientific output.

Universities get paid overhead when their scientists win grants, taking [around 1/3](#) of what comes in the door. In return scientists get facilities, admin help, a parking spot, etc. Unsurprisingly, this leads to [perverse outcomes](#) like incentivizing universities to grant tenure based on grant size rather than scientific merit.

iii. Course Teaching Requirements

Teaching is its own skill, distinct from research, and most scientists are bad at it. We shouldn't force them to do it.

Most training relevant to actually doing research occurs in the lab, where students learn from more senior students and occasionally from their advisors. But many researchers are forced to spend inordinate amounts of time teaching subjects to general student populations.

Pushing the frontiers of a field of science does not, contrary to what many seem to think, make one better informed about or better at teaching the basics of that field. While many researchers enjoy teaching, many more are bad at it (and often these are the same people). Information dissemination is obviously a good thing, but forcing research-oriented professors to disseminate basic information through the medium of lecture-based courses is wildly inefficient.

Traditional lecture-based courses are themselves, of course, far from the least efficient way of imparting knowledge.



VII. Small Market for Research Tools

Experimental apparatus for science is mostly awful since there's little money to be made selling it.

Most experimental science relies on specialty equipment. Some equipment like petri dishes or high-performance computers have large enough customer bases for them to be available reliably and cheaply to scientists. But newer and less-popular niches of science require equipment provided by a few companies, if any at all.

Most small businesses providing scientific equipment do it as a labor of love, since the markets involved are too small for grandeur. But small markets mean these companies can't invest in product development. Bad products, bad service, and slow innovation follow. And occasionally, among less generous companies, oligopolistic price hiking ensues.

I'm not sure what solutions there are besides helping research equipment companies capture more value from work done with their tools. Giving research equipment companies equity in developments done with the tools might help. Research equipment (a.k.a. "platform") companies in biotech often do this. [Advance market commitments](#) might also help.

VII. Switching Costs

When scientists publish in new areas, their work is less impactful (see "[Adaptability and the Pivot Penalty](#)").

To the extent it's because their work is worse, it's bad for the adaptability of science. To the extent it's because their work is undervalued, it's bad for scientific progress.

IX. Graduate-student Illiquidity

Once they pick an advisor, grad students are mostly stuck with them. Obvious problems arise.

It's extremely hard for graduate students to change advisors or work without an advisor. This labor market illiquidity is due to the advisor's monopoly over graduate students' funding and future career prospects. Graduate students are funded by their advisor, except for a small minority of students who receive personal fellowships like the NSF's [GRFP](#). And a graduate student's (or postdoc's) advisor's reference letter is by far the most important factor in future academic employability.

The negative effects of this illiquidity include advisors taking excess credit for their students' work, advisors forcing students to work on projects they don't care about or do jobs that benefit the advisor but not the student, and students being trapped into working under abusive advisors.



Funding students directly rather than funding advisors could ameliorate some of this problem.

X. Bias by Funding Source

Scientists are incentivized to produce results their funders will like.

For example, tobacco companies that sponsor research on how smoking isn't unhealthy. Or [Berkeley's controversial agreement with Novartis](#). These days scientists have to disclose their funding sources, but it's not hard for funders to hide themselves.

XI. Bad motivation

The cooler science becomes, the more people become scientists for the wrong reasons.

Science relies on an obsession with truth-seeking to hold scientists to norms of good conduct. But the increasing status of scientists in society may have increased the number of scientists who are motivated by status-seeking rather than truth-seeking.

It's not clear how much fraud, p-hacking, strategic citation, authorship-jockeying, academic politics, etc. occurred when science was less cool. But I certainly know a lot of people who do science for status and not truth. Maybe it's worth making science less cool. And yes, it's a stretch to call this a market failure.

Useful References

- [Microeconomics of Market Failures](#)
- [How Economics Shapes Science](#)
- [New Things Under the Sun](#)
- [Ben Reinhardt's List of Academic Constraints](#)
- [José's reminder that we don't know if fixing any of the above will help](#)

Gardener Comments

Ted Wade:

This paper's strengths are how the market failures are smartly categorized, cleverly and concisely stated along with concrete examples, and likely to resonate with most scientists except the big winners. These virtues are perhaps enough to justify accepting it.

It also references other similar critiques and meta-science studies. There appear to be lots of diagnoses of sciences' ills available now, as well as bottom-up attempts to cure it.



The framework used in the paper is the common one from economics that blames mis-aligned incentives for pretty much anything and everything that goes wrong in the world. The paper makes some suggestions of ways that alignment can be improved, and cites some examples of attempts in that direction.

There is some creative use of punctuation and incomplete sentences. These contribute to the easy, informal tone of the paper, but might be tightened somewhat.

I disagree with “Someone invented the existing way we do science.” That’s too casually stated. It was very many someones, in a broader cultural and historical context.

I would like to know more about “escrow-until-approval.” Also, how one finds a “false-positive rate cutoff that new protocols can *prove* they are below.” That sounds like an invitation to wishful thinking, or maybe it’s that I am 40 years behind in inferential statistics. Maybe a reference would do?

Eloy Parra-Barrero:

This piece reads well and does a good job at pointing out issues with the way science is done in academia.

It occurred to me while reading the section on credit assignment, that a solution to the problem could perhaps be to have no credit assignment at all: people doing science and collaborating with one another openly simply because they like doing so, not because they are trying to increase their h-index. But of course this sounds utopic, and it would make good resource allocation very difficult.

Anonymous1:

The manuscript offers a manifesto of market failures in science. It addresses a number of issues known to everyone in academia, but have never (to my understanding) been explicitly mentioned. In general, the writing is quite clear, but I noticed a number of typos along the way. I like the topic of the paper, but I think that in some sections, it needs justification.

Some thoughts I can share:-

. On author lists and contributions, I think you might mention that early-career researchers might also be willing to do anything to get a publication - at times even sacrificing lead authorship, or accepting relegation in the authorship list (I think this would be somewhat tangential to the Switching Costs section).

On Bad Motivation, I think your second paragraph was very neatly addressed by the OpenScience Framework (OSF). Perhaps it would be more relevant to stick to the trade-off between status-seeking and truth-seeking.

Antman:

This was a round-up of some of the problems facing the scientific endeavor at the moment. While this makes a nice blog post, there is nothing really new here. More



problematic is that issues are addressed, but very little effort has gone into suggesting well thought out solutions. Usually, no solution is offered at all. In essence, this becomes an exercise in moaning - fine for a coffee break or on twitter, but not really constructive. Many big issues are waded into without nuance, for example the 'drug lag' bit. I agree that this is an issue, especially in the USA, but the solution suggested had its own major ethical issues.

All in all, I found this not to be a very scholarly work - it's more of a moan than anything, and doesn't have enough intellectual content to support publication.

Dan James:

This is an excellent summary and analysis of what has been seen as a creeping marketisation of science (Meagan Day- 'Capitalism is Ruining Science' 2022), that has potentially negative effects on the way science is conducted. The analysis here is persuasive, making a compelling case for failures of a market-place approach to science by highlighting specific examples. However there was one statement that I feel could possibly benefit from some revision before publication– the notion that the first publication of an idea in science guarantees credit (or even fame!). In fact the issue is far more nuanced than this – as shown by Stigler's 'Law of Eponymy' (Stigler 1980), There have been a great many instances where even some of the most famous discoveries in science have not been attributed to the original proposer. It seems that the crucial factor in attribution relies far more on sociological factors- that it is the scientist who convincingly presents the idea to the scientific community who gets the credit.

Whilst the paper does not go far in suggesting remedies to address 'marketisation' problems, it does act as an interesting discussion piece for further work. Splitting the topic down into individual sections made the whole paper clear, easy to read and quite 'punchy', though I felt that a broader introduction and the addition of a conclusion would have been beneficial.

What was not 'teased' out of the argument in the paper and perhaps could be mentioned in a recommended conclusion, was the idea that the problems identified are not necessarily solvable by better methods and more stringent procedures (although these would surely help), rather that they are best addressed by a closer adherence to good values and ethics, moreover that it is above all the attitude of the scientist that builds integrity in any scientific community, as so eloquently expressed by Lee McIntyre in 'The Scientific Attitude'

Roger's Bacon:

While all of the issues are known to greater or lesser degrees, I think there is considerable value in compiling and explaining them in one article as the author has done here.

"A Problem Well Stated is Half Solved"



I imagine some people will want a more extensive introduction but I enjoyed the minimalist aesthetic choice and didn't find it necessary anyways.

"The cooler science becomes, the more people become scientists for the wrong reasons. Science relies on an obsession with truth-seeking to hold scientists to norms of good conduct. But the increasing status of scientists in society may have increased the number of scientists who are motivated by status-seeking rather than truth-seeking.

It's not clear how much fraud, p-hacking, strategic citation, authorship-jockeying, academic politics, etc. occurred when science was less cool. But I certainly know a lot of people who do science for status and not truth. Maybe it's worth making science less cool. And yes, it's a stretch to call this a market failure."

You put this at the end for obvious reasons but this might be the most important part of the paper. Is science too cool or not cool enough? In contemporary western culture science is seen as much cooler than religion (I'm sure some would dispute this); is that a good thing or a bad thing?

How many of these problems could be solved if science were fully anonymous? What new problems would we have?

Jared:

format comments

- article needs a more detailed explanation of 'market failures' as an introduction, or a clear example. My understanding is that SoS readers will be from general backgrounds so maybe worth the explanatory effort

content comments

- Interesting framing that the reward signal for scientists should be based purely on the reproducibility of their research. At the margin this is probably true, although in the degenerate case it seems to lead to a world where scientists get max-reward for doing simple, easily reproducible experiments decoupled from the main goal of science. One could argue that individual researchers have enough pride to *want* to do good research, but then this seems to be saying that we're relying on science working *despite* the incentive system, which seems to be our current situation anyway.

- Empirically, how common is mathwashing? Difficulty of error spotting reduces to a principle-agent problem, so maybe some research done by the economists/game theorists on that could also point the way here?

- Public goods - a question worth asking is why a relatively large amount of open-source academic software does get made. Is it just because the materials costs are low



(innovating in bits-vs-atoms) or are there other (reproducible) things we can learn from this?

- Advance market commitments for certain unsexy or small market-cap research tools are intriguing.

- On the last point, I don't think the coolness of science is the problem, although I'll admit I hadn't thought of it like that before. Maybe at the margin it's worth thinking about the Scientists as a precious resource of uniquely altruistic people, but I expect that making 10% progress on even half of the market failures listed here would shift incentives to a new equilibrium which might make the 'coolness' hypothesis seem implausible.

Overall enjoyable compilation of ideas/examples - it would also work as a good blog post (not meant as an insult, I love blogs, and this is the kind of post I like to read)

Umar Khan:

Fascinating read! Succinctly and clearly breaks down the incentive issues in science.

Sergey Samsonau:

The manuscript provides a collection of opinions aimed to support the "market failure in science" statement, claiming that an alternative allocation of resources might produce more outcome/value. While the current system is not ideal, it does not mean automatically that there is an alternative that can be more efficient. An alternative may be worse in various expected and unexpected ways, or may be better. I believe the content of the manuscript fails to support the main statement.

Jan Kirchner:

In the article "Market Failures in Science" the author(s) explore(s) societal shortcomings in how we incentivise research. They identify eleven factors that contribute to the inadequacy of the scientific apparatus, describe how these factors play out, and sometimes highlight possible improvements. The article is well-written and delivers its punches well.

I enjoyed reading the article and think it presents a valuable resource of analyses and ideas for others to build on. The only point in which the article falls short is in the (non-existent) introduction. I happen to already be convinced that market failures in science are real and important, but other readers might need additional context and arguments to benefit from the article. I don't think this shortcoming stands in the way of publication though.

Partha Ghosh:

This is very well written. The issues are logically established. The potential solutions are not always clearly presented - perhaps because we do not know them today, and more reasons why this study should be expanded.