Statistical Scapegoating & the Rise of Mimetic Science

Loud Experiments with Silent Evidence

According to Herodotus, Psamtik I, a Pharaoh of Egypt was the first to conduct a systematic study on the origin of language (of all professions). The pharaoh aimed to discover humanity's original language. For this he engineered a tightly controlled experiment in which a shepherd was tasked to raise two newborn infants in total isolation from spoken words. The hypothesis here being that their first utterance would reveal the universal root tongue.

Years passed, and one day the shepherd heard the infants say *'Beec'*. Excitedly, the Pharaoh's translators interpreted *'Beec'* as '*Becos*', the Phrygian word for bread, concluding Phrygian as the primordial language. Yet the background context, often more telling, is omitted from this apocryphal tale. The Pharaoh, and some of his inner circle wanted to have diplomacy with the neighbouring Phrygians, whom the Egyptians despised and assumed inferior to their culture. Such a revolutionary 'discovery' would make the bonds quicker to forge. The truth is, isolated from humans, the babies most likely mimicked their only companion, a goat. This overlooked explanation highlights a persistent human reasoning flaw: seeing only what we wish to confirm.

Today, we dismiss these attempts as primitive, yet modern science faces similar issues. Historically, the scientific method rested on an observable objective function. Modern scientific research heavily involves inferring unseen phenomena, clouding the clarity of a true objective function. We now depend on statistical interpretations within theoretical contexts more than ever.

This essay is about mankind's (unconscious) effort into domesticating statistics, giving our preconceived hypotheses dominion over it. As we'll see, if left unchecked in the present structure of the scientific system, this can trigger a scientific totalitarianism.

Taming Reality, the Bayesian way

Let's formalize the relationship between our ideas—hypotheses, theories, or models—and the reality they seek to explain. A good model explicitly states assumptions and produces predictions as though reality itself generated them. Yet, distinct models can yield similar outcomes. The central problem, then, is of model selection. Bayesian theory (Bernardo & Smith, 2008) offers a way to classify model selection in three distinct scenarios:

1. **M-Closed Problems**: These are problems where a true model exists and can be explicitly defined. However, there are multiple models to choose from, and the we as modelers must select the correct one. This situation arises in games, gambling, controlled experiments, simulations, and engineering. Here reality is well-defined and feedback immediate and precise.

2. **M-Completed Problems**: Nature does not like to announce its rules and mechanisms (say, through a podcast). In these problems, a true model exists, but it's not directly accessible. Instead, we must choose a set of surrogate models based on analytical tractability or communicative clarity. Most real-world decision-making and scientific inquiry fit this category. Here reality is incomplete but a decision has to be made with sparse and delayed feedback, generalizing into the future.

3. **M-Open Problems**: In scenarios like art, entertainment, politics, and culturally-generated realities, no single "true" model may meaningfully exist due to inherent volatility. Here reliable evidence is absent; thus, evidence is generated internally from 'proxy' models.

The difference between artistic and scientific approaches to reality further highlights these distinctions. Subject several artists to identical experiences, and each will produce unique interpretations reflecting subjective truths. Now consider the case of giving three scientists the same problem to vex over. Assuming intelligence, motivation, access and sustenance are equal, given enough time each of them would solve and have the same solution, often independently, even if they took different approaches.

Yet, even capable scientists fail if their prior beliefs or biases steer them away from the correct hypotheses. Historical delays in scientific breakthroughs often reflect the difficulty of unlearning entrenched misconceptions rather than the challenge of discovery itself. This struggle emerges from a fundamental tension between *Culture* and the *Scientific Method*, a previously stable but increasingly fragile relationship today.

The Imitation Game

The late philosopher René Girard proposed three foundational principles that explain how human culture emerges. Girard's theses can be used to coarsely predict how complex multi-agent systems with resource constraints (e.g., modern scientific enterprise) behave from seemingly simple principles.

The first principle is the concept of Mimesis. Put simply, it means most of our desires originate from seeing it in others. That is, we look for desires in others, and incorporate them as our own. Rather than creating models of reality ourselves, we are biased to adopt it from others. Humans being the most gifted imitators in nature doesn't help develop criticisms against this but mimesis is more than just raw imitation. It means a desire for similar valuable and thus scarce. This often leads to infighting, trend-chasing, fear of missing out and more. A doctoral student desiring their supervisor's achievements and a highschool student wanting to create the next AI company in 2025, might be mimetic.

While this drives ambition it also restricts diversity of thought. Science depends on diverse hypotheses to confront uncertainty, yet mimetic desire encourages conformity. Quantum physics, machine learning and neuroscience all had distinctly various schools of approach, each of which despite individually incomplete, was instrumental in piecing together the final jigsaw. Just as vagueness is the archenemy of scientific knowledge, mimesis is the archenemy of scientific hypotheses. Consider the situation where there are only a few sources of funding, spread out to all scientists. In a perfectly rational world, they would all have different approaches - one would be empirical heavy (clever experiments), one would be method heavy (sophisticated modelling) or measurement heavy (state-of-the-art machinery). But in a perfectly mimetic world, all of them would effectively be a copy of each other's approach (or the most famous) and become highly correlated. Nature, however, abhors undifferentiation.

This is aggravated by the overconnected nature of modern science. Massive science journalism, blurring of pop science and media, unholy mixing of science and politics via Twitter and other platforms, announcing finished/non-peer-reviewed works to great fanfare etc, might on the surface appear banal and harmless. But these deeply erode the very defense science had historically developed against our latent mimetic tendencies. Far fewer publications and even fewer journals were present until recently. The famed Philosophical Transactions journal by Royal Society was started to facilitate the then-assumed slow-moving correspondence between scientists, which was heavily one-to-one communication via letters. An overpopulation of journals continued the already prevalent aftermath of the publish-or-perish model. Massive connectivity via social media was not a factor present in the mix of scarce grants and over-publishing. But with it, our innate tendencies to mimic have taken a nasty detour.

This is not always gloomy however. I would also posit that honorable activities that were once hard to do as a group before, are now increasing due to Mimesis. Editors, as the guardians of any journal, should hold themselves to the highest ethics and clarity of mind possible ideally. An editor resigning over a dispute with the corporate suit is not news. An entire editorial board walking out is an act of rebellion, usually against the finance-first staffs in the publication house. In 2023 alone, 8 such rebellions took place, compared to 10 in the previous 5 years (*The Retraction Watch Mass Resignations List – Retraction Watch_, n.d.).* And it will only get more frequent now. Though short-lived, and temporarily taints the prestige of the journal, this can be seen as a positive effect of mimesis where righteous, but hard-to-do or speak-about activities occur more frequently.

The second principle is about Violence, and is a direct consequence of uncontrolled Mimesis in a society. We implicitly hate people that are (too) similar to ourselves. Girard posits complex societies inevitably become hostile toward individuals or groups perceived as overly similar, especially when competing for the same resources. Physical violence and war are natural, concrete forms to resort to for religion, politics and resources (indeed war is just a continuation of politics by forceful means). One can argue that science needs an internal selection mechanism, akin to 'intellectual' violence but more abstract and much more sophisticated, by eliminating the ineffectual theories.

But consider that, what was done usually by evidence can now be conducted by marshalling social capital, the central currency in Girardian mechanics. The recent debacle on Twitter of invalidating a famous neuroscience theory of consciousness (Integrated Information Theory) that resulted in an open letter (and article) penned by over 100 scientists (many of whom are esteemed) is a chilling example (*IIT-Concerned et al., 2024*). Later on, it was said that this stunt was done to dethrone its hegemony on grant acquisition for other less famous theories; they simply created shock value. While a commendable act, it must be said that it is essentially scientific violence and history will (ignominiously) look upon it as the watershed moment when consequences of Mimesis as the hand behind the lever. The irony is that it started with consciousness, the crown jewel of philosophy. Coincidentally enough, shortly afterwards an open letter was formed by another (cabal) in Artificial Intelligence, urging to stall progress because of the hidden (unknown) risks of that technology (*Pause Giant AI Experiments: An Open Letter - Future of Life Institute_, n.d.*). This was ineffectual obviously because capitalistic machines in Silicon Valley and hobbyists around the world have a different power structure and restraints, and cannot be stopped by such revolutions. But universities do not possess this and are completely in the stranglehold of social media and culture. A mark of impending violence in a culture is the complete disdain for social contracts and protocols. By usurping a theory (regardless of its scientific consistency) by force and in a very public manner, these researchers have marked the first act of (Girardian) violence in modern scientific society.

The third principle and the conclusive act is called Scapegoating, and is the natural defense developed by societies to halt Violence. Girard deduces that the only way for complex societies to sidestep violence is to put all the blame on a single, identifiable individual and thus 'wash off' the sins for collective relief. The origin of sacrifice in religion was deemed this manner, where it focuses on fringe 'beings' such as animals (goat in Abrahamic religions), children (Paganism), war prisoners (Mesoamerican societies) etc. The God-Man Christ was seen as the last great scapegoat (so far) after which a society is reborn with the sacrifice celebrated. This is extremely prevalent in corporate cultures where they term it the 'fall guy'. Almost every single example of corruption, scandal or bankruptcy is followed by a 'sacrifice' in the media of a named individual or two, while the true perpetrators move on to the next venture (or write memoirs).

But who are we scientists scapegoating then? The subtle answer is statistics.

Wrestling Reality

Jobber Metrics & Scapegoat Statistics

A scientifically trained mind infected with an ideology (political or scientific) will go further and devise careful hypothesis with intricate experiments to tease out data in such a way as to uphold an ideology. But surely cold statistical inference cannot be toppled by faith, an aspect the scientific method conquered centuries ago? Alas it can, by inserting a deceptive layer in between as a proxy of what represents the measured.

Enter the jobber metric.

"A Jobber (*jabroni* in pseudo-Italian) is a professional wrestling term (WWE/kayfabe) used to describe a wrestler whose job is to routinely get defeated by the high-profile wrestlers."

A jobber is a wrestler whose job, in that situation is to make the other person look good for entertainment and the overarching plot. Depending on the current fan trend and merchandise sold, the main wrestler ("face") will need to win and in a very demonstrable way. Their opponents will be these jobbers who are distinct from a villain ('heel'), whom the face has to defeat with difficulty. The narrative appeal of WWE, at its heart, is to see the current fan favourite win by making a fully detailed example out of jobbers. No other sports (let alone entertainment) company has tapped into this latent human desire the way this company did, to this day.

A jobber metric achieves the same thing by getting defeated by the hypothesis under the ideology. Most commonly, it looks like an undeniably bad correlation scatterplot used to make grand claims. But this is not usually between the actual reality and the model. Rather a metric is created ad-hoc, to act as a proxy for the reality to be explained 'strongly' by the hypothesis. This is different from statistical tests forged to distinguish false positive/negative conditions. This is different from blatant image tampering prevalent in cell biology studies. Here, a lineage of metrics, are specifically created for that experiment to defeat the alternate hypothesis, while appearing statistically significant. Recall that p-hacking and more generally data dredging, are all prevalent habits. However, it becomes much, much harder in the modern age where opaque machine learning algorithms come into the fray. These can include cherry-picked parameter tuned deep learning results imitating the brain, IQ studies and generally where the proxy metrics for something, keeps changing from paper to paper.

One might argue this doom song about confirmation bias in science is mitigated by the law of large numbers; assuming that individual biases naturally average out across the scientific community. If true, metrics such as *p-values*, the most common way (but certainly not the best) to assess hypothesis testing, should reflect this balance by exhibiting appropriate statistical distributions.

Alas, it is not so. An insightful paper, examined over 3000 such p values, and found a pronounced bump just below the widely-used cutoff of 0.05 (Masicampo & Lalande, 2012). The peculiar location of this strongly suggests widespread p-hacking, cherry picking, or Jobber Metrics.

Jobber Metrics, placed mendaciously between hordes of data and statistics, can deceive even the most seasoned intellectuals. Because it is neither altering/fabricating the data as some practitioners often clumsily do (and later appear in Retraction Watch), nor using wrong statistical assumptions, but instead introducing a weak metric that is designed to make the main hypothesis victorious. The mark of such a metric is its ephemeral nature. That is, they only appear in one or two works, making the key hypothesis win, and never to appear again (even in other similar works of the same authors). As such, a sequence of many different jobber metrics in various studies, all supporting the same ideologies, leads to an assured buildup and accumulated 'evidence', cementing the ideology instead.

Scientific culture always had selection mechanisms to prevent this despite the ineptness of the mainstream culture at any given time in history. But orbits are now colliding in ways never seen before and right now we are living in a culture space that never existed before.

A Culture of Confirmation

Our contemporary society can be described as Confirmation Culture, a pervasive consequence of one of humanity's oldest cognitive biases—confirmation bias.

This bias originates from how we were built to naturally handle M-Open problems, through much of humanity. Since we lack an infinite computational engine (even in Washington) we approximate probabilities to samples. That is, we mentally draw exemplars from different models of reality we learned, probabilistically and adjusting until some satisfactory threshold or tolerance is reached (Ratcliff & McKoon, 2008; Stewart et al., 2006; Vul et al., 2014). And then choose the model that fits the data best, rather than computing raw probabilities. One needs to survive and probability cannot cook one's food.

Beliefs carry with them practical consequences. Rejecting a potentially useful belief prematurely can impose real costs if it turns out to be a false negative. Another agent can follow the belief and profit in some permanent way (new group, new mates or new ways of foraging altogether). The competition between accepting and rejecting beliefs is not just about explaining away reality, but what helps in selecting the best actions upon it. Conversely, even incorrect beliefs can persist if they consistently lead to beneficial actions. An epistemologically wrong belief that helps in selecting the right action (often with ease), will always win over an epistemologically accurate belief that denies or even discommodes the right action. In the end, time will select the right actions, and the beliefs free-ride the survival. You may not need the data to act, but you do need the (belief in) hypothesis for it.

Historically, complex real-world (M-Completed) problems were approached as if they were open-ended (M-Open), making belief dominance more critical than belief accuracy. Confirmation Culture is a child of this latent bias and the modern overconnectedness. Many intellectually competent individuals seemingly fall prey to social theatre. Ironically this increases the confirmation for population-wide beliefs against science. Unlike the contrarian ideals of 1960s counterculture, Confirmation Culture prioritizes radical affirmation of one's existing beliefs, suppressing contradictory evidence. Disruption has evolved into consolidation.

The Sacred and the Significant

How would it look if one did not give into one's internal or social confirmatory biases?

Consider the Nobel awarded to John O'Keefe, for discovering place cells in the human Hippocampus; arguably one of the most brilliant discoveries of our times, demonstrating the mammal navigation system. He originally wanted to take on the ambitious task of single neuron recordings from Thalamus in freely moving rodents, an exceptional endeavour in the 1970s. A fruitless few years followed. Later on, the electrode coordinates accidentally moved laterally and he ended up recording from the hippocampus instead. And the rest is history (John O'Keefe – Biographical - NobelPrize.Org, n.d.). Imagine if the man satisfied his original bias using some jobber metric instead, only to be caught by someone else in posterity.

How would it look if Bias wears a labcoat?

In 2020, a celebrated paper examined ~90 million works and found that in most scientific disciplines, progress is not only remarkably slow but in turn consolidating existing ideas than exploring new ones(Chu & Evans, 2021). This was also seen in technology, aside from scientific discoveries, where there is a trend to have large groups to confirm beliefs and small teams are needed to disrupt (Wu et al., 2019).

We might think science always finds a way like in 1700s and 1900s where it prevailed against the Church and government. Alas we live in two distinct cultures to form similarity. Science was, is and always will be in need of heretics. Especially against itself.

But how will this phenomenon impact the trajectory of modern science?

Paths to Totalitarian Science

Here I outline three possible futures, and an anti-argument.

Retribution of Ideology

The government would infuse a fixed direction and like magnetic field aligns scientific hypothesis generation in a singular direction. Anything outside of this would be unfunded at best, and categorized as pseudo-science at worst.

Modern culture is increasingly becoming skeptical of science, and not in the expected way. Laypeople cherry-pick scientific data supporting their ideological beliefs, and form clusters around this, defending (and consolidating) their 'models' of the world, often violently. But this also opened the eyes of the 'laymen' in various hierarchies of the government, business, military and media (the modern 'monarchs') that, there are a significant portion of scientists who are, due to lack of funds, fame or faculty are willing to ally themselves with an ideology. People don't select beliefs, beliefs select people.

This is especially important since, all the structures of doing modern science - funding, publishing, machinery and administration - are a large part outsourced components driven by groups that are bedfellows with authority.

Mimesis of Corporations

A much less Orwellian, but strangely more realistic setting is the acceleration of Science cabals (the 'Big Science'). These possess larger than life ambition, overflowing with funding and influence, dictates how science should progress, and is a bonafide clone of corporate culture. A few giant 'science' corporations dominate the market, while fringe players play for scraps with a much lower ceiling to aim for. Unlike market dynamics, science differs fundamentally in the fact that scientists produce for other scientists, and not the masses directly. This lack of a selection mechanism would result in a sort of Big science consolidating scientific theories, models while ‘tiny’ science becoming largely fringe.

Return of Renaissance

A third but natural state of affairs in extreme uncertainty is to revert to a past state. Renaissance science was largely done by scholars who had their living made by other means (patronage, consulting, art) or inheritance (or priesthood) invalidating the need for labor.

A glimpse of this can be seen in modern independent 'artists' on YouTube and similar platforms. They find passionate patrons (with disposable income), and a consistent audience (that pays by attention) leading to advertisement and sponsorship deals, while they work on their craft and content. It's not hard to see dedicated scientist-as-a-craftsman working on concrete problems, whose output directly funds their pursuits in the near future.

What about AI?

A skeptical reader would, after entertaining, carefully refute most claims developed so far, and instead point solely towards the resurgence of modern AI as the healthy state of science. This is a misinference. Firstly, AI is largely a practitioner's calculus, much like the Steam Engine,Trading, Architecture and Surgery. We can generate images off it. We can get our answers to any factual question it is trained on. But nobody really knows how it works. While it has deep scientific and statistical foundations, the modern avatar as we know it was not a decidedly scientific prediction from the past masters, and more tellingly, nobody knows how to get to the promised land of AGI concretely. There would not be fields of Explainable AI, Interpretable AI, Aligned AI if we truly grasped it, say like astrophysics.

Modern AI belong to a class of fields where the implicit knowledge outshines and outgrows the explicit, and the rapid computing power available has increased the space to expand one's implicit knowledge. True Science, on the other hand, is all about the explicit, codifiable, and ultimately falsifiable knowledge.

Epilogue

I developed the argument that confirmation bias, effectively the oldest devil in any methodological inquiry is at its most rabid form in modern society. In the ancient world of limited data, what may have been a successful trait has corrupted into a force that has its growing claws on the scientific method itself. An abundance of data, and an overconnected social media, amplifying our Girardian tendencies, effectively scapegoating statistical approaches. The Confirmation Culture we find ourselves in, is extremely susceptible to sudden top-down regime changes making the scientific enterprise at mercy of this. While the path ahead may be treacherous and incredibly difficult reforms may be needed, our instincts always have a way of convincing us - not in my backyard. This isn't about preserving one's backyard; it's about saving our house. A house that has been built over centuries by minds far stronger and more formidable than ours, and a house that needs to stand for centuries for the minds that haven't formed yet, but need to live here.

And the house is on fire.

References

Bernardo, J. M., & Smith, A. F. M. (2008). Bayesian Theory. *Bayesian Theory*, 1–595. https://doi.org/10.1002/9780470316870

Chu, J. S. G., & Evans, J. A. (2021). Slowed canonical progress in large fields of science. *Proceedings of the National Academy of Sciences of the United States of America*, *118*(41), e2021636118.

IIT-Concerned, Fleming, S., Frith, C. D., Goodale, M., Lau, H., LeDoux, J. E., Lee, A. L. F., Michel, M., Owen, A. M., Peters, M. A. K., & Slagter, H. A. (2024). *The Integrated Information Theory of Consciousness as Pseudoscience*. https://doi.org/10.31234/OSF.IO/ZSR78

John O'Keefe – Biographical - NobelPrize.org. (n.d.). Retrieved January 31, 2024, from https://www.nobelprize.org/prizes/medicine/2014/okeefe/biographical/

Masicampo, E. J., & Lalande, D. R. (2012). A peculiar prevalence of p values just below .05. *Quarterly Journal of Experimental Psychology (2006)*, *65*(11), 2271–2279. https://doi.org/10.1080/17470218.2012.711335

Pause Giant AI Experiments: An Open Letter - Future of Life Institute. (n.d.). Retrieved January 31, 2024, from https://futureoflife.org/open-letter/pause-giant-ai-experiments/

Ratcliff, R., & McKoon, G. (2008). The Diffusion Decision Model: Theory and Data for Two-Choice Decision Tasks. *Neural Computation*, *20*(4), 873. https://doi.org/10.1162/NECO.2008.12-06-420

Stewart, N., Chater, N., & Brown, G. D. A. (2006). Decision by sampling. *Cognitive Psychology*, *53*(1), 1–26. https://doi.org/10.1016/J.COGPSYCH.2005.10.003

The Retraction Watch Mass Resignations List – Retraction Watch. (n.d.). Retrieved January 31, 2024, from https://retractionwatch.com/the-retraction-watch-mass-resignations-list/

Vul, E., Goodman, N., Griffiths, T. L., & Tenenbaum, J. B. (2014). One and Done? Optimal Decisions From Very Few Samples. *Cognitive Science*, *38*(4), 599–637. https://doi.org/10.1111/COGS.12101

Wu, L., Wang, D., & Evans, J. A. (2019). Large teams develop and small teams disrupt science and technology. *Nature 2019 566:7744*, *566*(7744), 378–382. https://doi.org/10.1038/s41586-019-0941-9