C.R.AIPS 2

Proceedings of *the Académie Internationale de Philosophie des Sciences*

Comptes Rendus de l'Académie Internationale de Philosophie des Sciences

Tome II

Justification, Creativity, and Discoverability in Science

Éditeur

Lorenzo Magnani

Second-level evidence for future-proof science?

Mario Alai

Università degli Studi di Urbino Carlo Bo, via Curiel 41, 47521 Cesena, Italy

1 Can we identify securely true hypotheses in current science?

According to the "pessimistic meta-induction" none of our current scientific theories, hypotheses or assumptions are true and will be preserved in the future. For some hyper-optimistic outlooks (e.g., Doppelt 2007, 2011), practically all current science, unlike past science, is true and (save minor adjustments) it will stay forever. The latter view runs against fallibilism, and the former against the hope that through science we can reach at least some truths. However, the current debates on scientific realism abundantly show that both extremes are wrong: most realists are nowadays selective realists, but also many antirealists are only *selectively* so. This means that we now possess at least some truths. Moreover, if science is progressive, or at least not badly regressive, it follows that many, or at least some, of our true scientific claims are future-proof, i.e., they will never be refuted.

The question, is, however: Can we identify which ones? Clearly, not all of those we hold now, for that would mean that now we are infallible. Moreover, it would seem that to distinguish which of our claims are true and enduring, and which ones are false, we should be able to anticipate future scientific progress, which is impossible. Yet, this question is becoming crucial today, not only for philosophers or historians of science but also for policymakers and the general public: the COVID-19 pandemic has shown how important it is for our very safety that even individual laypersons become able to distinguish between mere scientific *opinions* and established scientific *facts*.

Some, like Alberto Cordero (2017a, 2017b), maintain that for certain claims we now have such overwhelming evidence that it would be nonsensical to imagine that they can be rejected in the future. For instance, despite wide disagreements on the interpretations of quantum mechanics, physicists of all leanings and allegiances agree on certain basic tenets (Cordero 2001, p. 307). But who is to say when evidence is enough to warrant that we are in front of an indisputable *fact*? Earlier on, Rescher (1987, Ch. 5) distinguished between "forefront science", which is precise but mostly false, and "schoolbook science", which though vague and imprecise includes the true core of forefront science. But this distinction is vague; moreover, it doesn't seem to capture the distinction between reversible and perduring claims, because some of today's forefront science will still be true in the

Justification, Creativity, and Discoverability in Science, edited by Lorenzo Magnani.
Comptes Rendus de l'Académie Internationale de Philosophie des Sciences 2 (2024).
M. Alai, Second-level evidence for future-proof science?, pp. 1–13.

future, while much science that was in the schoolbooks of the past has subsequently been rejected.

Deployment realists (notably, Psillos 1999) have convincingly used the "no miracle" argument from novel predictions (henceforth 'NMA') to argue that when a hypothesis was essentially deployed in a novel and risky prediction, we can be practically certain of its truth; but from this it is a short step to conclude that (save unfortunate but unlikely scientific regressions) it is future-proof. There is a problem, however: to begin with, we must distinguish between claims that are *completely* true (i.e., true *tout court*) and others that are only *partly* true (i.e., false but with some true content, i.e., consequences). For instance,

(SW) All swans are white

is false, for some Australians swans are black. However, it entails the true statements that all European swans are white, all American swans are white, all Asian swans are white, and all African swans are white. These latter four claims are completely true, while (SW) is only partly true. This is what we mean by saying that (SW) is *approximately* true, and it explains the remarkable empirical and practical success that (SW) provides to its holders (Musgrave 2006–7). If we trust in the progressive nature of science, we can expect a completely true claim to be future-proof, while a partly true claim, in the long run, will (hopefully) be rejected and substituted by its completely true parts (i.e., consequences).

Now, when a hypothesis H is used to derive a novel risky prediction, there are two possible cases: either it has been deployed *essentially*, hence most probably it is completely true, or it was deployed *inessentially*, hence it is only partially true (Alai 2014a, § 7; 2021). H is deployed inessentially in a prediction when the latter was actually deduced from it, but it might equally well have been deduced just from a part of it. In that case, only the essential part of H is certainly true. For instance, (SW) may be employed inessentially in deriving the prediction

(PR) Any swan I see in Urbino will be white,

for the same prediction may also be derived from some of its parts, like 'All European swans are white', or 'All Italian swans are white', etc. Still, even the latter statements would be inessential to that prediction, while *in practice* only something like

(UrSW) All swans in Urbino are white

might count as essential.[1] For instance, as argued by Psillos (1999, 121), we need not be committed to the caloric hypothesis, although it was actually

[1]It may be remarked that, from a purely logical point of view, only (PR) itself is essential to deriving (PR). However, by hypothesis the fact described by a prediction (PR)

deployed in Laplace's prediction of the speed of sound in the air because it was not deployed essentially, hence that prediction "did not depend on this hypothesis". The same goes for the existence of the ether and other false assumptions that were deployed in successful novel predictions, albeit inessentially (Alai 2014a, § 7, 2021, §§ 9.4, 9.5). Therefore, only the assumptions that have played an *essential* role in novel scientific predictions can be trusted to be *completely* (not just *partially*) true, hence destined for preservation in the long run.

However, whenever a novel prediction *np* has been derived from a hypothesis H, it is *practically* impossible to tell whether H was employed essentially or not, and if not, which part of H was essential. One reason is that which is the minimal assumption a scientist needs to predict *np* depends on her background knowledge. Only in hindsight, if H is found to be false by subsequent research, it is shown that H had not been essential (Alai 2021, § 9.5). It follows that we cannot exactly circumscribe future-proof claims: the most we can learn from the NMA is that a hypothesis is at least partly true, and this doesn't guarantee that it will be preserved in the long run. Of course, this is to say that we do have some completely true (hence future-proof) beliefs (those that were actually essential in deriving the relevant predictions), but we don't know which ones.

Moreover, in his book *Identifying Future-Proof Science* (2022), Peter Vickers maintains that the question whether a claim is actually new and risky seldom, if ever, allows for a clear-cut answer, and it can be decided only through interdisciplinary competencies. In the past, I too argued that novelty, riskiness, and inessentiality are gradual properties (2014b, §§ 3.4, 4). Summing up, in many cases the NMA can be of little help in identifying future-proof claims.

More generally, Vickers points out that the first-level empirical evidence one would need to assess in order to decide whether a claim is future-proof is so vast and requires such specialized competence in different disciplines that not only no philosopher or layperson, but not even an individual scientist could master all of it. Even if a scientist could, over many years, study all that material, he would still see it from his individual and potentially biased perspective.

2 Vickers' criterion for future-proof scientific claims

Although neither the NMA nor the direct assessment of any other first-level empirical evidence allows us to decide whether a claim is future-proof, Vickers argues that this can be decided by a second-level criterion:

---

is unknown to us, and we must derive it from some assumption (H) we know. Moreover, if (PR) is to be deduced from (H), (H) must be stronger than (PR). Yet, there are problems in isolating essential assumptions: see the next paragraph.

(C) If the community of scientists competent concerning a claim C is sufficiently large and representative of different perspectives and sociological groups, and at least 95% of its members believe that C describes an established scientific fact (i.e., beyond reasonable doubt, save philosophical skepticism), then C is future-proof.

Since checking whether a given claim complies with (C) is in principle possible even for philosophers or laypersons, (C) may allow us to recognize future-proof claims even for the needs of philosophy and practical life.[2] Moreover, (C) is only a sufficient condition of truth and permanence, not a necessary one: unbeknown to us, even claims that don't command a 95% consensus, or do not yet, might be future-proof.

Granted, (C) runs against the current wisdom of both laypersons and philosophers: as famously argued by Kuhn (1962), scientific consensus may be achieved even for sociological (non-epistemic) reasons, and many theories or hypotheses that in the past were accepted as a matter of course have subsequently been rejected by the "scientific revolutions" and are now considered straightforwardly false. Yet, Vickers holds this criterion is borne out by the history of science: no claim fulfilling the requirements of (C) has ever been rejected: all the once largely accepted claims that have subsequently been rejected were still debated by at least some 6% of specialized scientists. Moreover, Vickers lists 30 statements or whole bodies of knowledge which met criterion (C) a long time ago, (pp. 12–18, 220).

However, one might grant that (C) has been confirmed so far, but ask whether this is only a contingent fact, or we can be assured that also present and future claims fulfilling the conditions of (C)'s antecedent will be future-proof. In other words, one may ask whether (C) is predictively reliable, and an affirmative answer may be provided only by *justifying* (C): *why* can we trust that claims fulfilling (C)'s antecedent are future-proof?

To begin with, Vickers remarks that in some cases almost unanimous consensus is reached when, thanks to technical progress, the entities, structures, or behaviors that were originally unobservable become directly observable by means of appropriate instruments, as happened with continental drift and the SARS-CoV-2 virus. But, he adds, scientific debates on the reliability of instruments, like the question of whether we see through optical microscopes, were resolved many years ago (pp. 198–200, 221). Therefore, claims that are

[2] In practice this may be more difficult: as Vickers remarked in discussion, some scientists are not convinced that global warming is anthropogenic, but they are very few. More or less than 5%? Of course, there aren't exact figures, moreover, it depends on which scientists we count as competent on this matter. Therefore, some have objected that in practice checking whether a claim satisfies (C) can be so hard, that it isn't actually helpful. Yet, while in some cases it may be difficult, in other cases it is certainly easier. Besides, even if it may be difficult to give an *exact* answer, in many cases it will be possible to give an approximate one, and for practical purposes that may be enough.

unanimously accepted on such bases are certainly true and destined to be preserved in the future.

As to the claims about entities that are not even instrumentally observable, it must be considered that science is essentially a critical activity, with strong epistemic and sociological premiums on criticism and nonconformism; therefore, if any doubts about a given claim were still possible, at least a substantial minority of scientists (well over 5%) would have raised them. "Any solid international scientific consensus is so hard-won that the evidence base has to be truly enormous to achieve it" (p. 112). As remarked by Kuhn himself, "History suggests that the road to a firm research consensus is extraordinarily arduous" (1962, p. 15). This is why if at least 95% of scientists have no doubts about a claim, it must be future-proof.

Still, it may be objected that various sociological and epistemic forces in scientific practice push to conformity, and Vickers is ready to grant that these forces may be even stronger today than in Kuhn's times. Therefore, when a claim gains acceptance, the "bandwagon" effect is a possible threat (p. 220). Thus, in the end, Vickers gives up on providing a full principled explanation of *how* scientists may reach a 95% consensus, and *why* it is such a reliable indicator of truth and stability; instead, he settles for just maintaining *that* it is, by induction from its confirmation in the history of science so far.

3 Why is Vickers' criterion reliable?

However, if a 95% consensus could be reached for merely sociological reasons (or for any epistemically irrelevant reason), the fact that claims passing the (C) standard have been preserved *so far* wouldn't be a sufficient reason to assume that they are true and that they will always be accepted in the future, i.e., the inductive support for (C) would become weaker, or even irrelevant. One reason to suspect that even claims that now command a 95% consensus might not be really future-proof is Stanford's (2006) idea that at any time, hence even now, there are alternatives to accepted theories that escape us, one of which might be the true answer to our questions. Another reason is the abovementioned fact that hypotheses which now seem to have been essential in deriving certain novel predictions (hence completely true) may actually have been inessential. Therefore, the questions of how the 95% consensus is reached and why it is reliable remain compelling.

Thus, in order to establish the reliability of criterion (C) one should show

(A) which *epistemically* relevant reasons (i.e., relevant to the truth of a claim) may produce a 95% consensus,

and

(B) why epistemically irrelevant reasons cannot produce such a consensus.

As mentioned earlier, a clear and satisfactory answer to question (A) is provided by Vickers for the claims about facts that have become observable thanks to new sophisticated instruments. Here I shall try to explore some further steps toward answering this question in the general case, and in the case of non-observable facts.

It is usually assumed that the probability conferred by some empirical evidence e to a hypothesis H is given by Bayes' theorem:

$$\text{(i)}\quad p(H|e) = \frac{p(e|H) \cdot p(H)}{p(e)}$$

or, in its extended formulation,

$$\text{(ii)}\quad p(H|e) = \frac{p(e|H) \cdot p(e)}{[p(e|H) \cdot p(H)] + [p(e|\neg H) \cdot p(\neg H)]}.$$

However, according to the NMA, the success of novel predictions also provides some second-level evidence for H, and this is why: all true hypotheses have only true consequences, while not all the consequences of false hypotheses are true: in fact, most false hypotheses have mostly false consequences. For instance, if it is true that

H Tomorrow is Tuesday,

all the consequences of H are also true:

C_1 Today is Monday.

C_2 Yesterday was Sunday.

C_3 The day after tomorrow is Wednesday.

C_n Etc.

C_{n+1} Today is a weekday.

C_{n+2} Tomorrow's name's initial is T.

C_{n+3} Etc.

Instead, consider the six false hypotheses about which day tomorrow is: none have the true consequences C_1 to C_n; besides, only five of them have C_{n+1}, only one has C_{n+2}, etc.

Now, theoreticians look for true hypotheses, and they try to build them in order to entail certain known evidence e, e', e'', etc. Hence their hypotheses, even if false, will entail e, e', e''. However, it is only by chance that they build a false hypothesis with a true *novel* and unlikely consequence ne. Therefore, we may ask what the probability is that by chance one picks a false hypothesis H with a true consequence ne. Clearly, it is equal to the ratio of false hypotheses that entail ne to all the false hypotheses on that subject.

In turn, that ratio is inversely proportional to the logical content of ne: if ne is tautological all hypotheses entail it, while if it is contradictory none do; between these two extremes, the greater the logical content of ne, the fewer the hypotheses that entail it. We may measure the logical content of ne as the ratio of the number of equiprobable cases it excludes to the number of all equiprobable cases (where equiprobable cases are those which we have no reason to believe have different probabilities).

Conversely, we may call the "logical probability" of ne the ratio of the number of equiprobable cases it allows to the number of equiprobable cases: if ne is tautological (allowing all cases) its logical probability is 1, and if it is contradictory (excluding all cases) it is 0. Hence, the probability that by chance one picks a false hypothesis H entailing ne is equal to the logical probability of ne.

Determining which are the equiprobable cases is not always possible, and it involves a principle of indifference which depends on our background assumptions. However, in several actual scientifically relevant cases one can figure out in a substantially clear way what the possible alternative cases are, hence what the logical probability of ne is.

For instance, when Adams and Leverrier predicted the mass and position of a new planet (later called Neptune), the equiprobable cases were the other positions and masses that a new planet might possibly have had. As for the mass, it is difficult to tell what the possible equiprobable alternatives were (e.g., those not obviously too large or too small). However, their prediction of the position missed the mark by just $1°$ over $360°$, so their true prediction was $n° \pm 1°$, hence its logical probability was $2/360$, or $1/180$. This is to say that only one in 180 logically possible (groups of) hypotheses entailing predictions of Neptune's position entailed the right one. Therefore, the logical probability that a theoretician picked a hypothesis entailing the right prediction *by chance* was also $1/180$.

Now, only by chance a false hypothesis entails a true consequence. Hence, the probability that Adams and Leverrier predicted Neptune's position using false hypotheses (i.e., that the hypotheses of Newton's gravitation theory essentially used by them were false) is $1/180$. On the opposite, the probability $p(H/ne)$ that they made this prediction using true hypotheses (i.e., that the hypotheses of Newton's gravitation theory essentially used by them were true) is fairly high. Granted, if the antecedent probability of Newton choosing false hypotheses $p(H)$ were very low, also $p(H/ne)$ would be low. But in general, this is not the case (Alai, 2023).

Since the probability of deriving a novel prediction ne from a false hypothesis H is proportional to ne's logical probability lp(ne), when ne is very risky (i.e., when lp(ne) is very low) it would be a "miracle" if H were false.

Besides, many contemporary novel predictions are much riskier and much more approximate than the prediction of Neptune's position, hence the probability that the hypotheses used to derive them are true is proportionally much larger: for instance, quantum electrodynamics predicted the magnetic moment of the electron to be $1159652359 \times 10^{-12}$, while experiments found $1159652410 \times 10^{-12}$: hence John Wright (2002, pp. 143–144) figured that the probability to get such accuracy *by chance*, i.e., through a false theory, is as low as 5×10^{-8}.

Thus, in the ideal case of NMA (when the logical probability of ne can be assessed with sufficient confidence) we might recognize that H is (at least partly) true from just one piece of evidence (the confirmation of ne). Something similar, perhaps, happened with Eddington's confirmation of Einstein's prediction of the bending of light in the 1921 solar eclipse. Hence, this is a kind of epistemically relevant consideration which may underly, explain and warrant a 95% consensus.

It might be objected that, just because there are infinitely many false hypotheses for any true one, the probability of finding a true hypothesis entailing ne is still lower than that of finding a false one entailing ne. Indeed, this would be the case if all hypotheses were picked by chance. But in that case, even the probability of picking a false hypothesis entailing a true risky prediction would be minimal, definitely too low to explain the many extraordinary successful predictions of our scientific theories. Therefore, we must conclude that scientists do not pick hypotheses randomly, but seek true hypotheses (which necessarily entail true consequences) and sometimes find them, and this happens because they employ reliable heuristics. This is why hypotheses that license novel risky predictions are most probably true (White 2003; Alai 2014c, § 6).

Bayes' theorem establishes the conditional probability of H on the basis of the antecedent probability of H, the antecedent probability of e, and the logical relationship of H to e. In general, the antecedent probabilities of H and e are different from their logical probabilities, for they are based on certain prior empirical or theoretical evidence we have for them. Only when we have no prior information, do their antecedent probabilities reduce to their logical probabilities, but this is seldom or practically never the case. Thus, Bayes' theorem assesses the conditional probability of H with first-level evidence.

Instead, the NMA assesses the probability that H is true only by considering that it has licensed a novel risky prediction ne, irrespective of what H

and *ne* say and of the first-level (empirical or theoretical) evidence we have for them. Thus, it provides second-level evidence for H.

Of course, in the overall assessment of H's probability one should also take into consideration the first-level evidence for it, i.e., H's probability conditional not just on *ne*, but on all relevant pieces of evidence e, e', e'', etc. However, as pointed out by Vickers, it is practically impossible to establish that a claim is future-proof just by assessing first-level evidence. In fact, especially because of the empirical underdetermination of theories, H's probability based only on first-level evidence cannot become sufficiently high. Instead, in ideal cases the NMA by itself may produce a reliable 95% consensus that a hypothesis describes a scientific fact.

There are problems, however, and not all cases are ideal. First, as explained above, in this way we cannot establish whether H was used essentially, hence whether it is completely true or not, and so whether it is actually future-proof. Alternatively, we might say that we can establish that it is future-proof *in the weaker sense* that *a part of it* is completely true, hence properly speaking future-proof. Second, the novel prediction *ne* may be not very risky, and the probability of H decreases as the logical probability of *ne* increases. Third, a NMA may be less than ideal also because it is unclear whether the prediction or predictions involved are actually novel, and even more because (as suggested by the literature on novel predictions), we should probably consider novelty as coming in degrees (Alai 2014b, p. 312).

Often, however, a hypothesis H is employed in more than one (more or less risky) novel predictions $ne_1, \ldots, ne_j$. For instance, Adams and Leverrier predicted both the position and the mass of Neptune and Mendeleev predicted various properties of different new elements. Of course, the conjunctive probability of many predictions gets smaller and smaller as the number j of predictions increases. Therefore, it can be very low, hence the probability that the hypothesis H from which those predictions were derived is true can be very high, even if each of those predictions is not *very* risky.

Instead, if it is unclear whether the predictions are actually novel, or if they are only partially novel, the probability of H will be proportionally lower.

Even this smaller probability, however, may raise the degree of confirmation conferred to H by first-level evidence.

Finally, the fact that H accounts for *old* evidence may also constitute additional second-level evidence for it. In fact, we may ask what the probability is of finding a hypothesis (whether true or false) which accommodates a known datum e. Apparently, the answer is that if the theoretician was minimally skillful, the probability was 1, for this was just a puzzle-solving exercise. But since there is an infinite number of false hypotheses for each true one, the probability that H is false is practically 1. However, things are seldom so simple. To begin with, H must usually account for *many* data $e_1, \ldots, e_m$, and the more they are, the harder the task becomes. Moreover, the data are practically never entailed *just* by H, but by H in conjunction with several collateral assumptions $A_1, \ldots, A_n$, which in turn must be derived from, or at least be consistent with, a number of independently accepted theories $T_1, \ldots, T_n$. So, one must also find the right $A_1, \ldots, A_n$, and H must not only entail $e_1, \ldots, e_m$, but be compatible with $T_1, \ldots, T_n$.

Therefore, finding a hypothesis entailing $e_1, \ldots, e_m$ (whether true or false) may become impossible for a minimally skilled theoretician, and very difficult even for a truly gifted one: the probability that H is false decreases with the number of accommodated data and of the collateral assumptions needed. If these numbers become very high, it is no longer plausible that H has been found by puzzle-solving skill alone, and another hypothesis becomes more plausible: that the theoretician was not just trying to accommodate $e_1, \ldots, e_m$, but, more importantly, looking for a true hypothesis using a reliable heuristic, so she actually found one.

This is why I once suggested that certain confirming instances which are apparently different from the confirmation provided by novel predictions are actually of a similar nature: for instance, the convergence of independent theories, the convergence of measurements by different experimental procedures based on independent theories, Keynes' distinction of confirming instances, and non-*ad hoc* explanations (Alai 2014b, § 4). In other words, this argument from the complexity of the theoretician's task can be turned into an argument from the improbability that just by chance independent theories converge in accounting for a large number of disparate data. Even this argument, however, may show at most that H is partly true, for even a partly false hypothesis may entail $e_1, \ldots, e_m$, and so be employed (inessentially) to account for them.

4 Conclusion

According to Vickers, a certain kind of second-level evidence (i.e., a 95% consensus) may show *that* certain claims are future-proof. Here I have suggested that certain different kinds of second-level evidence (provided by NMA-like considerations) may *justify* the achievement of such a consensus. That is, they can explain *why* and *how* it was reached for sound epistemological reasons rather than just for non-epistemic sociological drives, and hence why it can be a reliable indicator of future-proof claims. In other words, when at least 95% of specialists in a field take a claim as describing a scientific *fact*, they are probably right, at least in the *weak* sense that at least certain (possibly unidentified) parts of the claim are future-proof. A 95% consensus is justified especially when those claims yield various novel and risky predictions, but possibly also when they display great systematicity and unifying power (by accounting for a very large number of known data $e_1, \ldots, e_m$) and great plausibility and coherence with accepted theories and assumptions (by accounting for $e_1, \ldots, e_m$ in full coherence with assumptions $A_1, \ldots, A_n$ and theories $T_1, \ldots, T_n$).

Vickers grants that a claim fulfilling criterion (C) may be only *approximately* true, hence it may be future-proof *modulo-minor adjustments*. This qualification, however is a potential threat for his enterprise, for approximation is a vague concept, thus there is a continuum ranging from being simply true to being approximately true, half-way between true and false, more false than true, or, finally, completely false. So, if a 95% consensus can be achieved by claims that are only approximately true, why couldn't it be achieved by claims that approximate the truth less and less, or are even very distant from it? He suggests that in practice we can clearly distinguish when a claim is substantially true from when it is not. Perhaps we can, but only retrospectively: *ideally*, and with some approximation, we might suppose that 100% consensus shows that H is 100% true, 95% consensus shows that it is 95% true, and so on:

> The further a claim is from 100% true, the less likely it is that a truly solid consensus will be reachable. So, a claim only halfway true would probably never reach a 95% consensus. Other things being equal, there will be barriers to a halfway-true claim reaching 95% solid scientific consensuses that are not present for a claim that is true. E.g., it will likely be less thoroughly tested, and the scientific community will know that.[3]

However, things may be far from ideal, I am afraid there might be a very wide consensus on claims that are far from *completely* true (various examples are provided by the history of science). Again, therefore, it might be safer to assume simply that Vickers' criterion warrants that a claim is at least *partly* true, hence *weakly* future-proof.

Acknowledgment. I thank Peter Vickers for helpful comments on the first draft of this paper.

References

Alai, M. (2014a). "Defending Deployment Realism against Alleged Counterexamples", in G. Bonino, G. Jesson, J. Cumpa (eds.) *Defending Realism. Ontological and Epistemological Investigations*, De Gruyter, Boston, Berlin, Munich, pp. 265–290.

Alai, M. (2014b). "Novel Predictions and the No Miracle Argument", *Erkenntnis* 79, n. 2, pp. 297–326.

[3]Vickers in discussion.

Alai, M. (2014c). “Why Antirealists Can’t Explain Success” in F. Bacchini, S. Caputo and M. Dell’Utri (eds.) *Metaphysics and Ontology Without Myths*, Cambridge Scholars Publishing, Newcastle upon Tyne pp. 48–66.

Alai, M. (2021). “The Historical Challenge to Realism and Essential Deployment”, in *Contemporary Scientific Realism. The Challenge from the History of Science*, Lyons, T. and Vickers P. eds., Oxford University Press, New York 2021, pp. 183–215.

Alai, M. (2023). “Bayesian ‘No Miracles Argument’ and the Priors of Truth”, preprint, Pittsburgh PhilSci Archive, https://philsci-archive.pitt.edu/23321/.

Cordero, A. (2001). “Realism and Underdetermination: Some Clues from the Practices-Up”, *Philosophy of Science* LXVIII, pp. 301–312.

Cordero, A. (2017a). “Retention, truth-content and selective realism”, in E. Agazzi (ed.) *Varieties of Scientific Realism*, Cham: Springer, pp. 245–256.

Cordero, A. (2017b) “Making Content Selective Realism the Best Realist Game in Town”, *Quo Vadis Selective Scientific Realism?* Conference, Durham, UK, 5–7 August 2017.

Doppelt, G. (2007). “Reconstructing Scientific Realism to Rebut the Pessimistic Meta-induction”. *Philosophy of Science*, 74: 96–118.

Doppelt, G. (2011). “From Standard Scientific Realism and Structural Realism to Best Current Theory Realism”. *Journal for General Philosophy of Science*, 42: 295–316.

Kuhn, T. (1962). *The Structure of Scientific Revolutions*, The University of Chicago Press, Chicago, Ill., London.

Musgrave, A. (2006–7). “The ‘Miracle Argument’ for Scientific Realism”. *The Rutherford Journal* 2. Article 020108.

Psillos, S. (1999). *Scientific Realism. How Science Tracks Truth*, Routledge, London.

Rescher, N. (1987). *Scientific Realism: A Critical Reappraisal.* Dordrecht: Reidel.

Stanford, P. K. (2006). *Exceeding Our Grasp: Science, History, and the Problem of Unconceived Alternatives*, Oxford: Oxford University Press.

Vickers, P. (2022). *Identifying Future-Proof Science*, Oxford University Press, Oxford.

White, R. (2003). "The epistemic advantage of prediction over accommodation". *Mind*, 112(448): 653–683.