PESSIMISTIC META-INDUCTION AND THE EXPONENTIAL GROWTH OF SCIENCE[1]

LUDWIG FAHRBACH
University of Konstanz

1 SCIENTIFIC REALISM

This paper presents the outlines of a defense of scientific realism against the argument of pessimistic meta-induction (PMI for short). I will understand the position of scientific realism to consist of the claim that our current empirically successful scientific theories are probably approximately true. Examples of such theories are the atomic theory of matter, the theory of evolution or claims about the role of viruses and bacteria in infectious diseases. In what follows, I omit "probably" and "approximately" (as in "probably approximately true") and simply use "true." Furthermore, I use the term "theory" in a rather generous sense, so that it also denotes laws of nature, theoretical statements, sets of theoretical statements, and so on.

A theory is defined as being *empirically successful* at some point in time if there are sufficiently many cases of fit and no serious cases of non-fit between the known observational consequences of the theory and the observations gathered by scientists until that time. In contrast, if a consequence of a theory conflicts with some observations and scientists cannot find any other source of the error, e.g., they cannot blame an auxiliary statement, the theory is refuted and does not count as successful.

---

2. WEAK PMI AND STRONG PMI

There are several different versions of PMI. I only deal with some of them here. I will use the expression "PMI" in a generic way to denote the common idea of the different forms of pessimistic meta-induction discussed here. Stated generally, PMI starts from the premise that the history of science is full of theories that were accepted for some time but were later refuted and replaced by other theories, where these changes in theories occurred even though the refuted theories were empirically successful while they were accepted. This premise has a strong and a weak reading. The strong reading is that *most* successful theories in the history of science were later refuted. The weak reading is that at least *a significant fraction* of all successful theories accepted in the history of science were later refuted. The qualification "at least" is used here in order to make the weak premise logically weaker than the strong premise, so that it is easier to provide support for it from the history of science. To illustrate, "most" may be taken to mean "90 percent" and "a significant fraction" may be taken to mean "20 percent". Instead of the cumbersome term "significant fraction", I will also use the term "some". Thus, the weak premise states that *some* successful theories accepted in the history of science were later refuted.

In accordance with the two premises, there are two versions of PMI; a strong version and a weak version.[2] The strong version invites us to infer from the *strong* premise that *most or all* of our current successful theories will be refuted at some time in the future. The weak version of PMI invites us to infer from the *weak* premise that at least *some* of our current successful theories will be refuted at some time in the future. The weak PMI deserves its name, because its conclusion is compatible with some or even the majority of our current successful theories' being true.

The two PMIs are naturally associated with two forms of antirealism, one strong and one weak. The strong form uses the strong PMI to predict that most or all current successful theories will be refuted and therefore recommends believing that they are false (compare Ladyman 2002, p.231). An example of the strong form is Kuhn's account of science of changing paradigms. One plausible interpretation of his perspective is that he believed that "all paradigms are doomed to fail eventually" for the reason that "the world is just so *complicated* that our theories will always run into trouble in the end." (Peter Godfrey-Smith 2003, p.177, his empha-

---

[2] The distinction between strong and weak PMI was put forward to me by Gerhard Schurz.

96

sis). The weak form of anti-realism uses the weak PMI to predict that at least a significant fraction of our current best theories will be refuted. It recommends agnosticisms about these theories offering as reason that even if some of them are true, we do not know which ones are the true ones.

As stated, the two PMIs are cases of enumerative induction. They can be understood as being instances of the argument schema that projects the relative frequency of As among *observed B*s to the relative frequency of As among *unobserved B*s. (The As are the refuted theories and the Bs are the successful theories.) However, if we examine them more closely, we quickly see that in several respects, they are more complicated affairs and more fraught with difficulties than it seems at first. So, let us examine them more closely with the help of some idealizing assumptions. This will also help us to clarify what both arguments are actually stating. Let us assume for the moment that scientific fields are defined in such a way that in each one, exactly one theory is accepted at any time, and that scientific fields neither merge nor split. Let the number of scientific fields be N, and the number of successful but refuted theories from science's past be R. N is also the number of theories that are held today. Then the total number of all successful theories in the past is $R + N$. The ratio of refuted theories to all successful theories of the past is represented by the number $R/(R + N)$. The strong PMI states that if this ratio is not far from one, then most or all currently successful theories will be refuted in the future. Its premise can be taken to mean that on average there were several theory changes per scientific field in the past. For example, assuming that there are 100 scientific fields, and that on average every field experienced three theory changes, then $R/(R + N)$ is $300/400 = ¾$. The conclusion of the strong PMI can then be taken to be that most or all scientific fields will experience further theory changes in the future. Given the premise, this seems to be a reasonable conclusion.

The weak PMI states that if the fraction $R/(R + N)$ is "significant", then we should expect at least "some" changes among our current best theories. The premise can be taken to mean that at least "some" scientific fields have experienced one or a few theory changes in the past, and its conclusion can be taken to say the same thing about the future. However, it is not so clear how to make the inference more precise. One may wonder whether one should project the fraction $R/(R + N)$, the relative frequency of refuted theories among all successful theories of the past, or the fraction R/N, the average number of refuted theories per scientific field, as a lower bound for the expected fraction of refutations among our current best theo-

ries. Fortunately, for small R, the two ratios do not differ much. For example, if 33% of all fields experienced one theory change in the past, while 67% did not experience any theory changes, then R/N is 33%, while $R/(R + N)$ is 25% (because in that case of all successful theories of the history of science, 25% were refuted).[3]

Both PMIs are sensitive to additional considerations. For example, assume again that 33% of all fields experienced one theory change in the past, while 67% did not experience any theory changes, i.e., R/N is 33%. On the one hand, one might think that the fact that 33% of theories have been refuted shows that some of the fields with non-refuted theories probably also have unstable theories, only it did not show until now; as a result, this line of interpretation would continue, we should expect that more than 33% of current theories will be refuted. On the other hand, one might think that since the 33% figure implies that the theories of the majority of disciplines have been stable, this shows that stable theories will probably remain stable and, what is more, that some of the current theories in the bad disciplines will also remain stable in the future, because scientists have hit on the true theories, so less than 33% of current theories will be refuted. Which one of these projections is more plausible depends on one's background assumptions. These are only some of the problems and complications that one encounters when attempting to make meta-induction more precise. Further problems quickly crop up, e.g., how to individuate theories, how to determine the reference class (i.e., the set of all successful theories in the history of science), and how to do justice to the fact that theories often have parts that one may want to deal with separately. To my mind, all of this shows that it makes little sense to try to formulate PMIs with any higher degree of precision than that of the words "most" and "some".

The premises of both PMIs require evidence. Larry Laudan (1981) famously presented a long list of theories (which I will not repeat here), all of which were once successful, and all of which are now considered to have been refuted. An example discussed especially often in the literature is the sequence of theories of light in the 18th, 19th and early 20th centuries, see Figure 1. Many philosophers have considered Laudan's examples to be impressive evidence for the premise of PMI where the premise is usually understood in the strong way. If the strong PMI is correct, scientific realism is refuted, as the latter holds that our current successful theories are probably true. However, scientific realism is also refuted by the weak PMI,

---

[3] The numbers $a = R/N$ and $b = R/(R+N)$ are related by $a = b/(1–b)$ and $b = a/(a+1)$.

which concludes that at least some of our current successful theories are false. This is significant, because it is considerably easier to provide evidence for the premise of the weak PMI than the premise of the strong PMI.
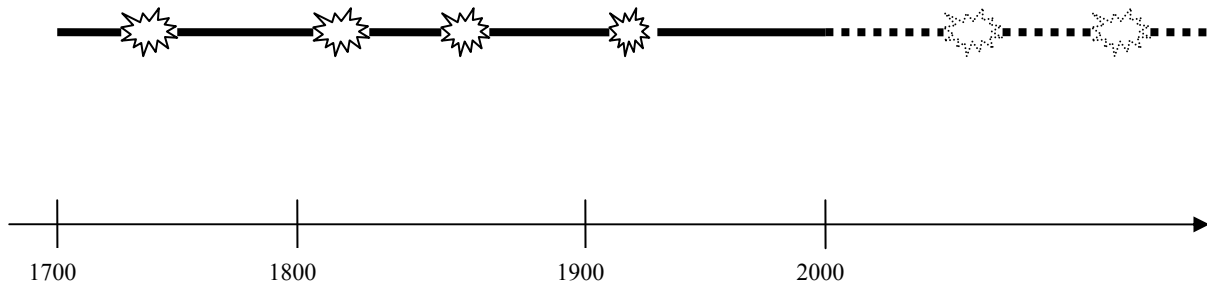

Figure 1: The sequence of theory changes in the case of theories of light and the projection of that sequence into the future

3 REFINEMENTS OF PMI

There are further ways to refine PMI. The premises of both PMIs average over all scientific fields, in the sense that all theory changes of all scientific fields are pooled disregarding any possible differences between the fields. To do this is not unreasonable, as Laudan's examples of refuted theories come from many different scientific disciplines such as physics, chemistry, biology, medicine, etc. Still, theory changes may not be distributed so evenly among scientific fields, as some fields may suffer from more theory changes than others do. This is the case, for example, if the premise of the weak PMI but not the premise of the strong PMI is true – i.e., if theory changes occurred in some, but not all, scientific disciplines or fields. Furthermore, even if theory changes are distributed evenly, the future stability or instability of a theory from a certain scientific field may be taken to depend on the specifics of that field, e.g., its subject matter or its methods, while the past stability or instability of theories of *other scientific fields* may not be taken as indicative of the future fate of theories *in that field*. Hence, one may want to refine PMI in such a way that the projection of theory changes is made relative to specific scientific fields. Every such PMI is then a localized "contextualized" inference. Some fields may then give rise to weak PMI, while others give rise to strong PMI. These field-relative forms of PMI can be made more precise in different ways, e.g., by

99

specifying how broadly the respective scientific field is to be understood. At one extreme, a field may be defined such that it consists of only one theory at every point in time. At the other extreme a field may be defined as encompassing an entire scientific discipline such as biology or chemistry, or it may even be understood as encompassing all of natural science.

As an example of the application of a field-relative form of PMI consider the field of clinical studies. In a recent meta-study, 49 highly cited (i.e., cited more than 1000 times) original clinical studies claiming that a drug or other treatment worked were examined. It turned out that subsequent studies of comparable or larger sample size and with similarly or better-controlled designs contradicted the results of 16% of the earlier studies and reported weaker results for another 16%. This means that nearly one-third of the original results did not hold up.[4] For example, the refuted studies had seemingly shown that hormone pills protect menopausal women from heart disease and that vitamin A supplements reduce the risk of breast cancer. An argument very similar to a field-relative form of PMI may then be used to infer that at least the same percentage of currently non-refuted clinical studies will not hold up in the future. Such a projection is akin to a weak form of PMI, because what is projected is not that *most or all* of these studies are false, but only that *some* of them are (where we do not know which ones). Finally, relativizing PMI to the field of clinical studies makes it possible to enrich and improve PMI by taking into account additional factors, such as the quality and size of the clinical studies.

Field-relative forms of PMI provide the basis of some forms of anti-realism. These forms of anti-realism state that many successful scientific fields have experienced theory changes in the past, and that if a field has experienced theory changes in the past, then we should, because of the respective strong or weak field-relative PMI, disbelieve or be agnostic about the current successful theories of that field. Such a field-restricted form of anti-realism may recommend, for example, that we not believe our current best theory of heat (roughly, that heat is mean kinetic energy) because several incompatible theories of heat were successful and accepted at different times in the past. Kyle Stanford's (2006, Ch.8) form of anti-realism seems to be a version of this field-restricted form of anti-realism. He essentially recommends that although we may believe *some* predictions of the best theories of unsteady fields, we should not believe many other predictions of those theories, especially those about new kinds of phenomena, again for the reason that the respective field-restricted PMIs should make us ex-

---

[4] Paraphrased from John Ioannidis (2005) and Lindey Tanner (2005)

pect that some of those theories will eventually fail empirically at some future time.

Figure 1 suggests a further idea regarding how to understand PMI. Assume once again for a moment that the notion of a scientific field is defined such that scientists of that field accept just one theory at any given moment in time. If the theory changes in such a field show a regular pattern, we can extrapolate that pattern along the timeline into the future, similar in kind to the extrapolation of a time-dependent regular curve of a certain form along the timeline into the future.[5] In the same vein, we can extrapolate *some chosen feature* of the pattern of past theory changes in such a field into the future. For example, a possible feature for projection suggested to me by Gerhard Schurz is to the "mean survival time" of theories of the respective scientific field. Let us call this form of PMI "the dynamic PMI". It is obviously a version of field-restricted PMI. Furthermore, we may understand the notion of scientific field in broader ways, e.g. as sets of theories. Thus understood, one possibility is to extrapolate how the frequency of theory changes per time developed over time where the frequencies are obtained from the whole scientific field: this frequency may have increased, decreased, or stayed the same, and this development is projected into the future. Then earlier theory changes have less of a bearing on the extrapolation than later theory changes do. Such extrapolations that are based on sets of patterns of theory change in broader scientific fields will also count as dynamic PMI. In general, the premise of any dynamic PMI is a statement about a feature of the set of time-dependent patterns of theory change in some scientific field in the past, and its conclusion is the statement that the same field will exhibit that feature in the future. I will call this focus on the pattern of theory changes "the dynamic understanding of the history of science". The dynamic PMI is also an instance of enumerative induction, although of a more complicated sort.

In discussions of PMI, the dynamic understanding is usually not mentioned explicitly, although it is lurking in the background and adds to the intuitive plausibility of PMI. The difference between the strong/weak PMI and the dynamic understanding is that the former completely abstracts from the timing of the theory refutations and theory changes, while for the

---

[5] The step-function which takes time points t as arguments and natural numbers as values where the value of the step-function at t is the number of theory changes up to t in the given scientific field encodes the pattern of theory changes in that field. If it has some regular form, that form can be extrapolated into the future.

latter, the distribution of the points of time of the theory changes are relevant.

What could the rationale for the dynamic PMI be? One might think that it has no rationale, because the points in time of the theory refutations are irrelevant for the projection of theory changes into the future; the timing of the refutations and their distribution seem to have no epistemic significance. Still, a preference for the nearer past over the more distant past may be justified by claiming that science is constantly changing for the better in several respects – methodologically, for example (Devitt 2005, p.787) – but most importantly, in the amount and quality of evidence that scientists have accumulated (compare Gerald Doppelt 2007). To this, one might reply that the increase in evidence and the increase in success of the theories do not really make much of a difference, as this is the same kind of success and nothing qualitatively new (compare Stanford 2006, Section 1.2). However, we will see in a moment that the dynamic understanding of the history of science, if developed further in the right direction, has more to it than meets the eye.

4 REFUTATION OF PMI

Let us now turn to developing the argument against PMI. Consider the amount of scientific work done by scientists in different periods of time, and how that amount increased over time. Here, "scientific work" means such things as making observations, performing experiments, constructing and testing theories, etc. It is plausible to assume that the amount of scientific work done by scientists during some period can be measured by the number of journal articles published in that period. Over the last few centuries, this number has grown exponentially. The doubling rate of the number of journal articles published every year has been 15–22 years over the last 300 years (see Figure 2). This is a very strong pattern of growth. As is shown in the appendix, these doubling rates imply that at least 95% of all scientific work ever done has been done since 1915, and that at least 80% of all scientific work ever done has been done since 1950.
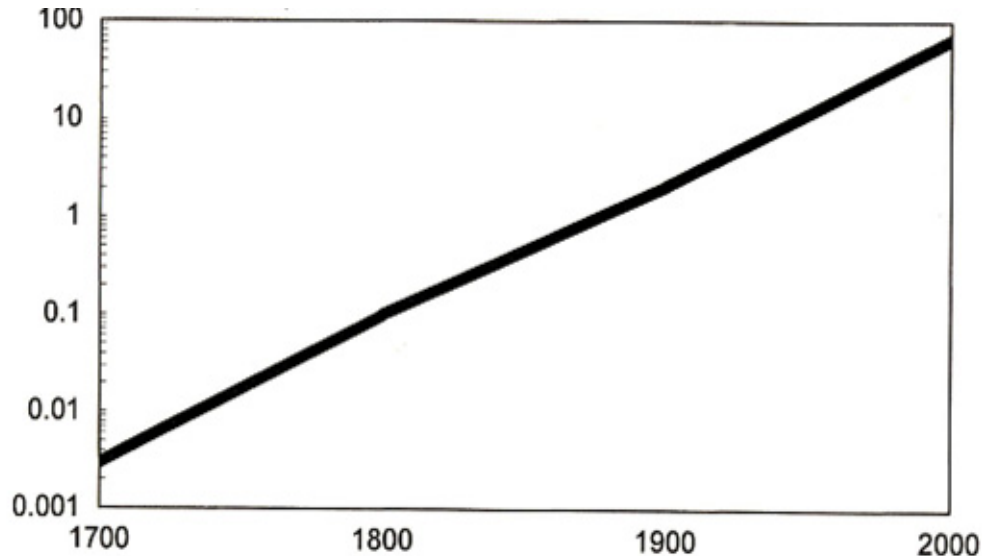
Figure 2. The growth in the number of scientific journal articles over the last 300 years. The vertical axis shows the cumulative number of scientific journal articles measured in millions. It has a logarithmic scale; hence, the straight line represents exponential growth. Note the small bend in the curve. Note also its thickness, demonstrating the uncertainty of the data. (From Brian C. Vickery 2000, p.xxii).

Let us examine how the exponential growth of science affects the different forms of PMI. Let us first consider the premise of PMI in its intuitive form: the history of science is full of theories that were once successful and accepted but were later refuted. As we saw earlier, proponents of PMI support it by offering numerous examples of such theories. Now, however, given the exponential growth of science, we have to recheck whether these examples are really evidence for the premise of PMI. If we do so, we get a very different idea of the matter. Inspecting Laudan's list, we see that all entries on that list are theories that were abandoned more than 100 years ago. This means that all corresponding theory changes occurred during the time of the first 5% of all scientific work ever done by scientists. As regards the example of theories of light, all changes in those theories occurred before the 1930s, whereas 80% of all scientific work ever done has been done since 1950. The same holds for practically all examples of theory changes offered in the philosophical literature. Thus, it seems that the set of examples offered by proponents of PMI is not representative and cannot be used to support the premise of PMI. If this is right, the premise lacks support and PMI does not work.

To this argument, one might object that the intuitive understanding of the premise is misleading. We should use the premise as it is understood in

the weak and strong PMI. Both of these state only relative frequencies (namely, of the appearance of refuted theories among those that were successful in the past). They abstract from the periods of time at which the theories were accepted and the points in time at which they were refuted, and rightly so (according to this objection), because neither has any epistemic significance. The only thing that is relevant for PMI is that the refuted theories were successful. Hence, Laudan's list is a representative sample of refuted theories after all, and it can serve to support the premises of the strong or weak PMI.

In response to this objection, I want to show that it is actually the premises of relative frequency that are not adequate as premises for meta-induction. I will do so by elaborating on the dynamic PMI. The dynamic PMI states that certain features of the pattern of refutations and stability in the past of science are to be projected into the future. But how exactly is the projection to be accomplished? Consider Figure 1, where the x-axis is weighted in a linear fashion such that equal lengths of time are represented by intervals of equal length on the x-axis. With the exponential growth of science in mind, a second weighting suggests itself: the x-axis could be weighted in such a way that the length of any interval on the x-axis is proportional to the amount of scientific work done in that interval; see Figure 3. I will call these two ways of weighting the x-axis the linear weighting and the exponential weighting.
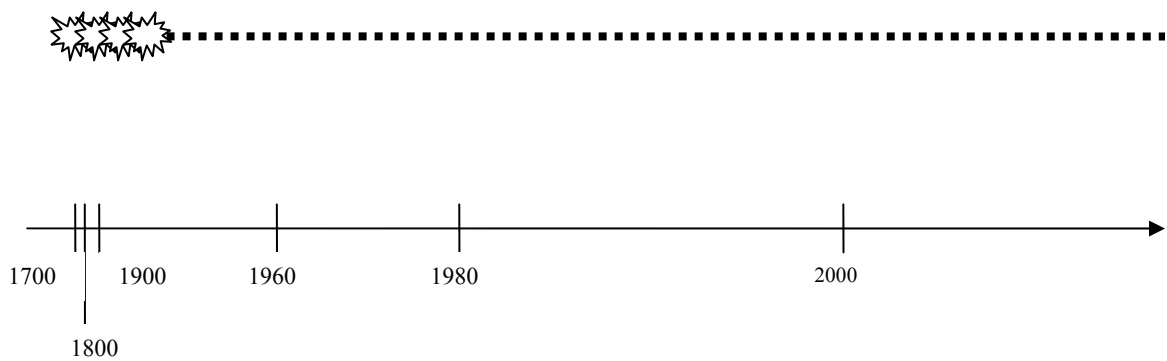
Figure 3: Exponential weighting of the x-axis and the sequence of theories of light.

Both weightings could be used in the premises of the dynamic understanding of PMI. Which one should be used? If we want to project the past development of science into the future, which weighting is the right one? The exponential weighting is more plausible for the following reasons. If we want to determine how stable or unstable the best scientific theories

104

have been in the past, we should look at the amount of scientific work done by scientists, because the amount of scientific work can be expected to be very roughly proportional to the amount and quality of empirical evidence compiled by scientists to confirm or disprove their theories. More concretely, but still on a very general level, more scientific work results in the discovery of more phenomena and observations, which, in turn, can be used for more varied and better empirical tests of theories. More varied and better empirical tests of theories, if passed, result in greater empirical success for theories. Although it is certainly not plausible that *all* scientific disciplines profited from the increase in scientific work in this way, it is even less plausible that *no* scientific disciplines and *no* scientific theories profited from the increase in scientific work in this way, and it is obviously the latter–the disciplines and theories that did profit–that realists want to focus on.[6] This consideration is a good reason to adopt the exponential weighting of the x-axis in the premise of the dynamic PMI.

Although this consideration offers a prima facie strong case for the exponential weighting, it clearly is in need of further elaboration. However, in order to develop it more fully, we would need a better worked-out notion of the degree of success of scientific theories. Such a notion has to be developed in future work. The main tasks are then, first, to establish a sufficiently strong connection between the amount of scientific work (as measured by the number of journal articles) and the degree of success of the best scientific theories, and second, to show that such a connection can be exploited for a more fully developed argument against PMI. From now on, I will proceed using the assumption that such a connection can be established.

The assumption implies that the exponential weighting of the x-axis is the correct weighing. It also implies that the premises of the strong and the weak PMI, which make statements exclusively about frequencies, are inadequate: in projecting the past development of science into the future, it is not appropriate to abstract from the times at which the theories were accepted or changed. Time matters because different periods differ very strongly in terms of the amount of scientific work done in them, and because the amount of scientific work is linked with the degree of success of the best theories. Hence, the two premises leave out relevant information and should therefore not be used in an inductive argument; it follows that neither the weak PMI nor the strong PMI are cogent arguments.

---

[6] This focus is, of course, a form of field-restriction.

Let us proceed by examining what the exponential weighting implies for the dynamic PMI. Our observation from the beginning of this section, that all examples of theory changes discussed in the philosophical literature are rather old, shows that this set of examples is not representative and therefore cannot support the premise of the dynamic PMI. Based on this set of examples, nothing can be inferred about the future change or future stability of scientific theories. This is illustrated in Figure 3. Thus, we need to come up with a more representative sample set. We should examine the last 50 to 80 years. Only then can we decide whether the premise of the dynamic PMI is plausible. So, let us look at this period. Moreover, as we just observed, we should focus on the *best* (i.e., most successful) scientific theories. If we do so, it quickly becomes clear that virtually all of our best scientific theories have been entirely stable in the last 50 to 80 years. Despite the very strong rise in the amount of scientific work, refutations among them have basically not occurred. Here are some examples of such theories (remember that the realist endorses the *approximate* truth of those theories):

- The theory of evolution
- There are 92 natural chemical elements
- The conservation of mass-energy
- Infectious diseases are caused by bacteria or viruses
- $E = mc^2$
- The brain is a net of neurons
- There are billions of galaxies in the universe.
- Sound consists of air waves
- In the Earth's past, there were many ice ages
- And so on

Proponents of PMI will have a hard time finding even one or two convincing examples of similarly successful theories that were accepted for some time during the last 50 to 80 years but were later abandoned (and one or two counterexamples could be tolerated because, after all, we are dealing with inductive inference here). This does not mean that there were no theory changes in the last 50 to 80 years, which there clearly were; of course the large amount of scientific work done in the recent past has also brought with it a lot of refutations. It only means that there were practically no theory changes *among our best (i.e., most successful)* theories. For example, the highly cited clinical studies mentioned earlier, which were

partly or wholly refuted later on, constitute theory changes; still, every such study is just *one single study*, which is clearly insufficient for one to count its result as belonging to the most successful theories.

At this point, one might object that the notion of "the most successful theory" is intolerably vague; it can neither be used to delineate a set of theories, nor for the statement that the most successful theories have been stable. A satisfactory reply to this objection would have to rely on a better elaboration of the notion of success, a task that, as noted earlier, has to await another occasion. In this paper, I have to appeal to our intuitive understanding of it and trust that it is not too vague to serve my purposes here. However, as a preliminary reply to this objection, I want to offer the following argument.

The argument is meant to defend that even if the notion is rather vague, the claim of the stability of our current most successful theories is quite plausible. Consider Laudan's list once again. Although the set of *theory changes* on Laudan's list is not representative for projection, it is plausible to assume that the set of *scientific fields* of those theories form a representative set of scientific fields for projection. These fields belong to many different scientific disciplines: astronomy, chemistry, biology, medicine, geology, etc. When we examine the currently accepted descendant theories of the theories on Laudan's list, that is to say, those theories that are about the same subject matter as Laudan's theories and that are accepted today, then we observe that most of them can be regarded as being among our current most successful theories. At this point, the objection was that the notion of a "most successful theory" is dangerously vague and should not be used to describe the past of science. However, it turns out that it does not matter that this notion is rather vague, because *all* of the descendent theories have been entirely stable for the last few decades (and most of them for a far longer time). This is the case for the descendant theories of phlogiston theory, the crystalline spheres of ancient and medieval astronomy, the effluvial theory of static electricity, both theories of heat, and the theory of circular inertia. (For some theories on the list, it is not clear which current theories to count as descendant theories, because they are too general or too unspecific – for example, the humoral theory of medicine and the vital force theories of physiology.) Thus, for a representative set of theories, the vagueness of the notion "most successful" does not matter.

What follows from all this for the dynamic understanding of PMI? As we saw, the dynamic understanding of the history of science is the right

understanding of the history of science for any instance of meta-induction, but it also turned out that the premise of this version of the dynamic PMI was false (if restricted to the set of our current best theories). We already concluded above that neither the weak nor the strong PMI is a cogent argument. Thus, no form of PMI considered in this paper, neither the weak PMI nor the strong PMI nor the dynamic PMI, is a cogent argument. Their conclusions that some, most or all of our current best scientific theories will fail empirically in the future is therefore not supported by the history of science. Hence, PMI in any of these forms is refuted.

The refutation of PMI has two consequences. First, field-relative forms of PMI are not cogent for fields involving our current best theories. Hence, Stanford's form of anti-realism is not tenable for those theories (although it may be tenable for theories of less successful scientific fields such as the field of clinical studies). Second, scientific realism, the claim that our current most successful theories are true, is saved from being undermined by PMI.

5 APPENDIX

1. In the main text, I use the number of scientific journal articles published in some period as a statistical quantity to measure scientific work in that period. As is the case for any statistical quantity (and, actually, any scientific quantity generally), we can distinguish two main tasks, the task of definition and the task of determination, where both tasks come with their specific difficulties. First, we have the conceptual task of providing sufficiently precise and fruitful definitions of the respective statistical quantities, in this paper the notion of a scientific journal article. The accompanying difficulty is that usually, several different ways to define any quantity are possible, and often, no definition is clearly superior to all other definitions. Furner writes, "If we are to count 'publications,' … should we count monographs, or serials, or both? If we are counting serial publications, should we count yearbooks and technical reports, in-house organs and newspapers, as well as journals? Should we count journal titles, or journal articles, or pages, or words? How, moreover, should we define 'scientific'?" (Furner, 2003, p.9) Second, we have the empirical task of determining values of the quantity, i.e., the number of journal articles per year. Here, the difficulty is that whatever definition of "journal article" we choose their number at any point in time is not easy to ascertain. Especially

for centuries earlier than the 20[th] century, Bibliometricians can only yield estimates with limited accuracy. In any case, neither kind of difficulty threatens our project here because all we need are rough estimates; for this, what Bibliometricians offer is entirely sufficient. Different definitions of "journal article" and different estimates of their numbers by Bibliometricians do not lead to significantly different results as far as the aims of the paper are concerned.

2. We have to show that of all scientific work ever done, at least 95% has been done since 1915, and at least 80% has been done since 1950. Consider the following list of the number of scientific journal articles during different periods[7]:

1600 till 1900	2 million
1901–1915	1 million
1950s	4 million
1960s	600 000–900 000 annually
1600 till 1970	20–25 million
1600 till 2010	60–80 million

As for the first claim, adding the first two entries of the table and comparing the sum with the last entry, we observe that at most 5% of all scientific work ever done was done before 1915. As for the second claim, the upper bounds of the last two entries show that at most 25 million out of the total 80 million articles ever published were published before 1970. Around 6–9 million articles were published in the 1960s, and at least 4 million were published in the 1950s, leaving at most 15 million articles before 1950, which is less than 20% of 80 million. We obtain a similar number if we use the lower bounds of the last two entries. Because the last two entries in the table are not independent from each other, we need not consider the combination of the upper bound of the second to last entry with the lower bound of the last entry. Apart from the last assumption, my calculations are mostly rounded in such a way that if they err, they err to my disadvantage.

We obtain the same results when we use the doubling rate (the length of time in which the number of journal articles doubles). Because we only

---

[7] The numbers are taken from Vickery (1990, 2000, p.xxii, see figure 2), Meadows (2000, p.89) and Furner (2003). See also Mabe/Amin (2001, pp.145–5). Numbers before 1600 can be neglected. This is especially obvious for the time before the invention of movable type printing by Gutenberg around 1450.

need a low level of precision, we can assume that the number of publications per year follows an exponential function $f(t) = a \cdot \exp(\lambda t)$, where a is a normalization factor (which disappears in the results), and λ is the growth rate. Let d be the doubling rate. The doubling rate has been between 15 and 22 years over the last 300 years. λ and d are related by $\lambda d = \ln 2 = 0{,}7$ (because $f(t+d) = a \cdot \exp(\lambda(t+d)) = a \cdot \exp(\lambda t) \cdot \exp(\lambda d)$ and $f(t+d) = 2 \cdot f(t) = 2a \cdot \exp(\lambda t)$). For example, a doubling rate of 20 years corresponds to a growth rate of 3.5% per year. The number of publications between two points in time is given by the integral of $f(t)$ between those two points in time. For example, the number of publications from the beginning of science until time T is given by

$$\int_{-\infty}^{T} f(t)\,dt = a \cdot 1/\lambda \cdot \exp(\lambda T).$$

This integral is a measure of overall scientific work done in that period.

We are interested in how many of all publications ever published were published after a certain year T. Let r denote this ratio. Let S denote the present time, and n be such that $n = (S - T)/d$ (n measures how often the number of publications doubles between T and S.). Then

$$
\begin{aligned}
r \quad &= \quad \int_{S-nd}^{S} f(t)\,dt \ / \ \int_{-\infty}^{S} f(t)\,dt \\
&= \quad a \cdot 1/\lambda \, (e^{\lambda S} - e^{\lambda S - \lambda nd}) \ / \ a \cdot 1/\lambda \, e^{\lambda S} \\
&= \quad 1 - e^{-\lambda nd} \\
&= \quad 1 - 1/2^{n}
\end{aligned}
$$

For example, assume that we want to know what fraction of all scientific articles ever published were published between 1970 and 2010. Let us assume a doubling rate of 20 years. Then $n = 2$ and $r = 1 - \frac{1}{4} = 75\%$. Hence, the answer is that of all articles ever published 75% were published between 1970 and 2010. Here are some values of r for different years and for doubling rates of 20 and 22 years. (Judging from the literature, a doubling rate of 22 years is too high an estimate, and here serves merely as an upper bound.)

T	1810	1873	1918	1953
r for $d = 20$ years	99,9%	99,1%	95,9%	86,1%
r for $d = 22$ years	99,8%	98,7%	94,9%	83,4%

We see once again that around 95% of all journal articles ever published have been published since 1915, and more than 80% have been published since 1950.

REFERENCES

Devitt, M. (2005), "Scientific Realism", In *The Oxford Handbook of Contemporary Analytic Philosophy, Frank Jackson and Michael Smith*, eds. Oxford: Oxford University Press, pp.767–791.

Doppelt, Gerald (2007). "Reconstructing Scientific Realism to Rebut the Pessimistic Meta-induction", *Philosophy of Science*, Vol.74, pp.96–118.

Furner, J. (2003). "Little Book, Big Book: Before and After Little Science, Big Science: A Review Article, Part II", Journal of Librarianship and Information Science, Vol.35, No.3, pp.189–201.

Godfrey-Smith, P. (2003). *Theory and Reality: An Introduction to the Philosophy of Science*, The University of Chicago Press, Chicago.

Ioannidis, John P. A. (2005). "Contradicted and Initially Stronger Effects in Highly Cited Clinical Research" *Journal of the American Medical Association*, July 13, 2005; 294: pp.218–228.

Ladyman, James, (2002). *Understanding Philosophy of Science*. London: Routledge

Laudan, L. (1981) "A Refutation of Convergent Realism", Philosophy of Science, 48 March, pp.19–49.

Mabe, M. and M. Amin, (2001). "Growth Dynamics of Scholarly and Scientific Journals", Scientometrics. Vol.51, No.1, pp.147–162.

Meadows, J. (2000). "The Growth of Journal Literature: A Historical Perspective" in: The Web of Knowledge: A Festschrift in honor of Eugene Garfield Edited by Blaise Cronin and Helen Barsky Atkins ASIS&T Monograph Series, Information Today, Inc., Medford, New Jersey.

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

Tanner, Lindey (2005). "Review of Medical Research Turns Up Contradictory Results", The Associate Press, July 13, 2005.

Vickery, B.C. (1990). "The Growth of Scientific Literature, 1660–1970", in: The Information Environment: A World View, Studies in Honour of Professor A.I. Mikhailov, Elsevier Science Publishers B.V. North-Holland.

Vickery, B.C. (2000). *Scientific Communication in History*, The Scarecrow Press, Lanham, Maryland.