

# Collective epistemic vice in science: Lessons from the credibility crisis

Duygu Uygun Tunç, Duncan Pritchard

## Abstract

We investigate the explanatory role of epistemic virtue in accounting for the success (or failure) of science as a social institution that is characterized by predominantly epistemic ends. Several structural explanations of the epistemic success of science that commonly rule out virtue attributions to scientists are explored in reference to a case of collective epistemic vice; namely, the credibility crisis in the social and behavioral sciences. These accounts underline the social structure of science as the chief explanatory factor in its collective success, and endorse a common conclusion, namely that individual virtue is neither necessary nor sufficient for science to be successful. While acknowledging that divergent motivations and behaviors might also serve the collective goals of science, our analysis of the credibility crisis shows that the presence of a significant proportion of epistemically virtuous scientists in a scientific community is a necessary condition for collective epistemic success in a scientific community.

## Introduction

Do scientists have to be epistemically virtuous in order for science to be successful? This question as to the role of epistemic virtue in science can be embedded in a broader debate on how the goals of epistemically virtuous communities and epistemically virtuous individuals are related. We can identify two positions as marking the extreme ends of this debate. At one end we can place the view that individual epistemic virtue is necessary and sufficient for collective epistemic virtue. Call this position *virtue traditionalism*. At the other end, we can place the recently much more popular position that individual virtue is neither necessary nor sufficient for collective virtue, and that individual vice can possibly contribute to the collective success of science. Call this position *virtue radicalism*. This latter position has been defended by many in the philosophy of science to explain the success of science in terms of its social structure, most notably the arguments from the invisible hand and

division of cognitive labor offered by Hull (1978; 1988) and Kitcher (1990; 1993), respectively. These arguments maintain that the success of science is primarily due to the incentive structures of its credit economy, most importantly the priority rule in science (which we will explain below). Thanks to the resulting social structure of science, scientists can serve the collective epistemic good by pursuing non-epistemic goals like esteem and recognition instead of genuinely epistemic aims such as truth or knowledge. The same conclusion is defended in the sociology of science in reference to scientific norms, most notably by Merton (1973). The sociological argument maintains that scientists behave in a way that serves the institutional goals of science not because they are epistemically virtuous but because they comply with the rules and norms enjoined by the scientific institution, which are thus the real cause behind its success.

While acknowledging (*contra* virtue traditionalism) that divergent motivations and behaviors might also serve the collective goals of science, in this paper we argue (*contra* virtue radicalism) that the presence of a significant proportion of epistemically virtuous scientists in a scientific community is a necessary condition for collective epistemic virtue. This is because neither incentive structures nor institutional norms can be established, maintained, or reformed in a way that reliably serves collective epistemic ends of science without there being a significant proportion of epistemically virtuous scientists. We arrive at this conclusion indirectly, through an analysis of a set of major arguments for virtue radicalism, and offer objections in reference to the credibility crisis in the social, behavioral and life sciences that has been going on for more than a decade now.

The structure of the paper is as follows. In §1 we make some preliminary remarks on how we use the terms epistemic virtue and vice in the following argument. In §2 we present the major economical arguments for virtue radicalism and offer objections in reference to meta-scientific research and theorizing on the credibility crisis. In §2a we present the argument from invisible hand and offer four objections. First, invisible hand arguments are formulated as causal explanations for the success of science, but the same invisible hand mechanisms correlate with collective epistemic vice as well as virtue. Second, scientists respond to incentives or sanctions differently depending on both individual and contextual factors. In some cases, the most salient explanation of the heterogeneous behavior of scientists in relation to situational factors are features pertaining to their epistemic characters. Third, the social structure of science is not a sufficient explanation of the collective epistemic virtue of the scientific community and epistemic virtue has an indispensable explanatory role. Fourth, many central aspects of the success of science cannot be explained without

reference to individual epistemic virtue. In §2b we present the sociobiological version of the economic argument. In response, we point out that an adequate sociobiological model should also include group level factors, which may undermine the virtue radicalist conclusion. In §2c we present the argument from the division of cognitive labor. In response, we point out an inherent normative tension within scientific inquiry that stems from the differences in the contexts of discovery and justification. Consequently, we argue that no single social mechanism, such as the priority rule in science, can be normatively adequate. Lastly, in §3 we present the sociological argument for virtue radicalism and argue, in response, that horizontal mechanisms for norm-compliance are equally important as the top-down enforcement of institutional norms and that individual epistemic virtue is one of the most important determinants of these horizontal mechanisms.

### **1. Epistemic virtue and vice**

Epistemic virtue and vice have been defined variously in the literature.<sup>1</sup> For the purposes of the present paper, we take epistemic virtue to consist broadly in a set of relevant cognitive competences that are acquired, exercised and maintained with proper epistemic motivations and reliably produce good epistemic outputs such as true or justified beliefs. By proper epistemic motivations we mean those that generally manifest a responsible degree of “concern with truth” (or other epistemic goods such as accuracy or validity).

Epistemic virtue in science comprises, most broadly, scientific competency and responsibility. An epistemically virtuous scientist can accordingly be construed as someone who competently, and responsibly, contributes to the advancement of scientific knowledge (however big or small the contribution). The particular content of competency and responsibility is generally dependent on the scientific field and the typical research questions in it. Generally, scientific competency comprises the relevant field knowledge and cognitive skills such as perceptiveness, correct reasoning, adequate mathematical knowledge or technical know-how, and responsibility comprises a motivational state conducive to reliable, rigorous and credible scientific inquiry. Responsible conduct of scientific inquiry involves actively aiming at truth (or some other scientific good such as credible empirical findings, well-corroborated theories, or empirically adequate models), avoiding negligence (e.g., avoiding disregard of relevant evidence, considering the total and not the immediately available evidence), reasonable sensitivity to epistemic risk (e.g., suspending judgment below a reasonable threshold of uncertainty), being alert to error and self-serving biases, choosing the best methods of

inquiry, and interpreting and communicating research findings honestly. Scientific responsibility in the sense we use it here is thus not a merely behavioral notion, having to do with adherence to the prevailing norms and standards of research, but involves an intrinsic concern with advancing scientific knowledge.

Competency and responsible are not simply two independent factors, but bring about epistemic virtue by their proper interaction: An epistemically virtuous scientist is one who cultivates and uses the relevant scientific skills with a view to advance scientific knowledge. Hence, a lack of relevant, necessary competences may amount to epistemic irresponsibility (an insufficient concern with epistemic goods), or improper motivations may lead excellent scientific skills astray and undermine their value (for instance, in the case of research misconduct).

An epistemically virtuous scientist puts epistemic goals and values above all else in their approach to inquiry. Thus, they could diverge from mainstream scientific practices if they deem these insufficiently justified. Epistemically virtuous scientists need not have exclusively epistemic motivations, though. Non-epistemic goals such as recognition, productivity or career advancement are potentially compatible with epistemic virtue, as long as these do not take precedence over epistemic motivations in making choices regarding research design or the interpretation and communication of findings. We also do not envision epistemically virtuous scientists to be excellent individuals in all respects. From the particular normative angle we adopt, epistemic virtue is perfectly compatible with moral vice or other failings of character. In some circumstances pursuing epistemic goals at all costs might itself amount to a weakness in practical reasoning. What matters from the perspective of epistemic virtue is that the strictly cognitive aspects of science are governed by purely epistemic concerns.

We identify epistemic vices in reference to epistemic virtues, as commonly done. Thus, an epistemic vice is a disposition or motivational state that consistently produces bad scientific outputs or systematically stands in the way of proper scientific inquiry. An epistemically vicious scientist is consistently negligent towards counterevidence or norms of proper inquiry, is culpably incompetent, reports false findings, makes unwarranted scientific inferences, engages in bad methodological practices, irrationally clings to bad theories and models, or in the worst case falsifies or fabricates data.

In analogy to the individual level, we take an epistemically virtuous scientific community as one that is capable of reliably advancing scientific knowledge. An epistemically vicious community, on the other hand, is one that not only lacks a progressive character but also consistently produces bad epistemic outcomes, such as a high proportion of erroneous or false findings in the literature, widely accepted but wrong or useless theories and models, and ineffective outputs such as treatments, interventions or scientific advice for social policy. Let us leave stagnating or immature fields aside, as they are difficult to class as either virtuous or vicious.

Since we identified individual epistemic virtue as comprising a competence component and an epistemically proper motivational state, by analogy we should be able to identify appropriate collective counterparts to these. We can define collective epistemic competence of a community as its capacity to reliably produce good epistemic effects. In most general terms, the collective epistemic competence of a scientific community is not only a product of the individual competences of the scientists working in it, but also of its theoretical maturity as a field or discipline, the accuracy and reliability of its measurement techniques, the division of cognitive labor, or the efficiency of collective research foci and strategies. So, collective epistemic virtue is not straightforwardly summative and may involve other impersonal determinants too.

The collective counterpart of an epistemically proper motivational state, or a “collective concern for epistemic goods,” does not have to be a group motivation in the literal sense, which might be too restrictive to characterize whole scientific communities (and better suit smaller groups). A better candidate for a scientific community can be found in the shared values, goals, and standards that incorporate a concern for truth. These are commonly conceived as the “epistemic ethos” of a community (see e.g., Fricker, 2020). Ethos is not a straightforwardly summative notion either. The epistemic ethos of a scientific field or discipline is not only a product of shared motivational dispositions, but also of its particular social organization, institutional norms and sanctions, or scientific culture (see e.g., Knorr-Cetina, 1999). In a good scientific ethos, epistemic values such as truth, accuracy, rigor, or objectivity are cherished and effectively promoted. In a bad scientific ethos, we can expect low deterrence or even incentivization of bad scientific practices.

### **Traditionalism vs radicalism about epistemic virtue**

Let us go back to our main question as to the nature of the relationship between individual and collective virtue in science. The virtue traditionalist view would say that the goals of epistemically

virtuous scientific communities and epistemically virtuous scientists overlap, because collective epistemic virtue is made up of individual epistemic virtue. According to this, a scientific community is capable of reliably advancing scientific knowledge (chance discoveries aside) if and only if it consists of epistemically virtuous scientists. This is in line with the commonplace impression of scientists as competent truth-seekers, as high-minded inquirers who care to advance humanity's collective knowledge. The counterpart of this view would be that an epistemically vicious scientific community (which consistently fails to deliver) implies epistemically vicious scientists and vice versa. Thus, individual epistemic virtue is necessary and sufficient for collective epistemic virtue.

*Virtue Traditionalism:* A scientific community manifests epistemic virtue by reliably reaching its collective epistemic goals if and only if all (or most) of its members manifest epistemic virtue in their individual contributions to those collective goals.

In contrast, there are several relatively recent economical and sociological arguments for the thesis that what is essential to epistemically successful science is its social structure. These arguments maintain that individual epistemic virtue is neither necessary nor sufficient for collective epistemic virtue. We thus get virtue radicalism:

*Virtue Radicalism:* A scientific community can manifest epistemic virtue by reliably reaching its collective epistemic goals even without a single member manifesting epistemic virtue in their individual contributions those collective goals.

We agree with these arguments (and thus disagree with virtue traditionalism) to the extent that individual epistemic virtue may not be the only salient factor in explaining the collective success of science. We thus do not reduce collective success as a scientific field or discipline to the competences or motivations of its individual members. As we also acknowledged in the preceding, the success of science may also depend on other factors, such as variables relating to the theoretical and methodological state of the individual fields, institutional research goals and strategies, or the norms and values promoted in the community. Thus, the social structure of science can furnish a partial explanation of its success (or failure).

Moreover, we also agree that individual epistemic virtue is not necessary for making a contribution to science. The history of science features numerous chance discoveries, unintended breakthroughs, or eventual vindication of theories that for a long time were thought to be utterly false and irrational

to endorse. Thus, consisting entirely of epistemically virtuous scientists might be neither sufficient nor necessary for a scientific community to be successful, as the traditionalist position holds.

What we claim, instead, is that the presence of a significant proportion of epistemically virtuous scientists in a scientific community is necessary for its collective scientific success in the longer-term. Thus, facts about the social structure of science cannot completely or exhaustively explain its collective success but furnish at most a partial explanation. Collective epistemic success (and failure) is not purely systemic, but both systemic and individual at the same time.<sup>2</sup>

Lastly, while we argue that the presence of a significant proportion of epistemically virtuous scientists in a scientific community is necessary for it to achieve and maintain a progressive course, what proportion would count as significant is variable. For instance, it can be argued that theoretically and methodologically less mature fields require a much higher proportion of epistemically virtuous scientists, who have a distinctive concern and competence for advancing scientific knowledge, in order to develop reliable research methods and objective evaluation criteria for research outputs. In the absence of these, epistemic vice is also a much more serious concern than it is for more mature fields, because non-epistemic factors may be more influential in determining how widely research outputs are accepted by other scientists. In mature fields ordinary skills and merely practical motivations may allow one to make a significant contribution and epistemic vice is potentially less detrimental, thanks to the objective reliability of the prevailing research procedures and evaluative standards. But even in the most mature fields, epistemically virtuous scientists have a crucial role to play, although a more limited number would be enough. Scientific consensus on the superiority of certain theories and methodological procedures (Kuhnian normal science) exerts a conservative influence on scientists. They consequently become less prone to detect anomalies to the dominant theory or to identify the potential weaknesses in their standard procedures. This prevalent conservatism can become a hindrance for further scientific progress. In such a context it is for the benefit of science to have some members in the scientific community who are characteristically less prone to be tempted by social influences than others. When the demands of epistemic rationality are divergent from the collective preferences of the scientific field, epistemically virtuous scientists play a vital role by pointing this out to the others.

Below we analyze several arguments for virtue radicalism and offer objections in reference to the credibility crisis of the numerous fields of the social, behavioral and life sciences. We formulate the

objections in an accumulative manner, so that the subsequent ones will presume and build upon the previous ones.

## **2. The economic argument for virtue radicalism**

### **2a. Argument from the invisible hand**

What we call the economic argument constitutes a bold attempt towards a “disenchantment” of science. Before roughly the second half of the 20<sup>th</sup> century, the commonplace conception of the ideal scientist was a champion of epistemic virtue. This picture was slowly replaced by a picture of the scientist as an entrepreneur. Scientists do not contribute to the advancement of scientific knowledge because they intrinsically endorse certain epistemic norms and values more than other people, but because doing so is in their personal practical interest. Strevens (2011) puts the motive for the disenchantment in the economic argument quite explicitly when he says: “It is hardly possible to believe any more, if it ever was, that scientists are made of [...] better stuff than the rest of us... They are intellectual entrepreneurs.”

We see one of the most common forms of the economic argument for virtue-radicalism in the philosophy of science in the various invisible hand arguments for the success of science (e.g., Hull 1978; 1988; 1997; Kitcher 1993). The general structure of invisible hand arguments is that a decentralized social process results in an outcome that is not intended by the participating agents (Ylikoski 1995). Quoting from Wray (2000), an invisible hand argument describes a particular outcome “as an unintended consequence of the intentional behavior of a number of individuals. The individuals have one end in mind, and act accordingly; but their concerted efforts give rise to a consequence that was no part of their intentions.” Typically, an invisible hand explanation describes a case where the pursuit of selfish interests produces common good. Applied to science, an invisible hand argument for the success of science purports that scientists primarily seek credit, not knowledge, and in doing so they advance scientific knowledge. In Hull’s formulation, the credit economy of science is so organized that the collective outcome of each scientist just pursuing their individual interest is that the aims of science are served. “Scientists want credit and [...] science is so structured that this desire leads to increased knowledge of the empirical world” (Hull 1988, 357), so that “some of the behavior that appears to be most improper actually facilitates the manifest goals of science” (*ibid.*, 32). Thus, the credit economy of science gives rise to collective virtue as if it is intentionally designed so as to promote the advancement of knowledge.



The social structure of science that creates this invisible hand mechanism is characterized as a “winner-take-all” credit market (Hull 1997), which is due to what is commonly called the *priority rule in science* (Merton 1957). Strevens (2003) describes the priority rule in terms of two main characteristics: “rewards to scientists are allocated solely on the basis of actual achievement, rather than [...] on the basis of effort or talent invested” and “no discovery of a fact or a procedure but the first counts as an actual achievement.”

Another characteristic of the social structure of science, and for many a result of the priority rule (Hull 1997; Heesen 2017), is the system of mutual citation. The scientific institution incentivizes originality, speed and impact, since credit goes only and completely to those who are the first to discover something, and to the extent that they can make others accept and use their research. Scientists want their research to be accepted in order to get credit for it, and they can maximize the chances of their being accepted by using well-established work of others, thus by citing them. The competition for credit is also the very mechanism by which science polices itself or corrects its own record. Having to use the findings of others also makes the scientists interested in their reliability, because it affects their prospects for credit. As a result, scientists are incentivized to be sensitive for any potential errors in others’ research, thereby realizing scientific self-correction.

The crux of these arguments is that thanks to the social structure of science, in this case the particulars about the credit economy, epistemic virtue and selfish credit-seeking are *behaviorally or consequentially indistinguishable*. The social structure of science is miraculously successful in producing collective epistemic goods in the same way as a truth-seeking motivation would do. Hull (1997) says: “Scientists need not be saints to contribute to science...the really neat thing about the reward system in science is that it is so organized that, by and large, more self-serving motivations tend to have the same effect as more altruistic motivations.”<sup>3</sup>

The invisible hand argument rules out epistemic virtue as a viable explanation of why scientists do what they typically do. While it does not advance the empirical claim that scientists generally do not have epistemic motivations but only practical motivations, it does advance the theoretical claim that the seemingly virtuous practices of scientists can be explained with practical motivations only and thus any reference to epistemic motivations is an unnecessary commitment for an explanatory account. We think that scientists are truth-seekers by looking at their behavior, but the most parsimonious account of their behavior is that they pursue their own practical interests just as much as anybody else. Thus, the social structure of science is the chief explanatory factor in accounting for

the success of science and individual epistemic virtue is redundant. This suggests that the practical motivations of scientists are sufficient to explain the success of science.

There is a slightly different and stronger conclusion that is conjoined with a similar set of premises, and it is that regardless of whether they actually have epistemic motivations, scientists *must* have a prominent credit-motivation, at least some significant proportion of them, for science to be successful. Hull (1988) maintains that if science consisted only of high-minded scientists (who do not care about credit) it would not be as successful as it is. This is because the success of science is primarily due to the incentive structures of its credit economy and if all or most scientists lacked a prominent credit motivation this structure would fall apart. For instance, the dissemination of scientific knowledge and scientific self-correction are explained in terms of the system of mutual citation, where it is the credit-seeking motive of the scientists does the causal heavy lifting. Scientists who lack such a motivation would be supposedly less willing to share their findings or to criticize other scientists' work. This suggests the practical motivations of scientists are necessary to explain the success of science.

### **An anomaly for the model: the credibility crisis**

The invisible hand mechanism is invoked as a causal model to explain the success of science. But a serious flaw in its justification is that its proponents only look at science at its most successful and abstract out some social interaction patterns to explain this success. A good test for the model would involve, however, also looking at less impressive cases to see whether they feature the same or different incentive structures.

The credibility crisis (Vazire 2019; IJzerman et al. 2020) and increased attention on the effect of the present incentive structures on research quality and reliability reveals a critical anomaly for the model. The same reward system not only fails to work as the proponents of the invisible hand describe, but also correlates with highly undesirable epistemic outcomes such as unreplicable findings, low methodological standards, publication bias, and consequently an unacceptably high rate of false or misleading scientific claims in the literature (Ionnadis 2005; Camerer et al. 2018; Open Science Collaboration 2015; Button et al 2013). The root causes of this credibility crisis have been widely traced to a range of established but faulty research practices (Bakker, Van Dijk & Wicherts 2012; Fraser et. al. 2018; John, Loewenstein & Prelec 2012; Simmons, Nelson & Simonsohn 2011). These 'questionable research practices' (QRPs), as they are commonly called, make up a wide spectrum, some clearly bordering on research misconduct.<sup>4</sup> In general terms, QRPs refer to the use

of wrong methods, wrong use of right methods, misinterpretation of results, selective reporting of results or studies, and drawing of unjustified conclusions (Altman 1994). The use of QRPs lead to seriously bad epistemic effects, because published results generate false impressions about the reliability and importance of empirical findings.

The invisible hand idea, or at least some of its major premises, consequently came under attack in the aftermath of the credibility crisis. In the meta-science literature there are several arguments underlining the damage the existing incentive structures do to science. The causes of the widespread use of QRPs have widely been traced back to the institutional incentives that favor them (Chambers 2017; Nosek et. al. 2012; Munafo et. al. 2017; Simmons, Nelson & Simonsohn 2011; Smaldino & McElreath 2016). According to these analyses, the institutional incentive structures are not well-aligned with the aims of science but actually undermine them. Competition for priority leads to unreliable research processes (Tiokhin et al. 2021a). Playing the credit-maximizing game strategically requires collecting publishable results most efficiently, which consists in doing research in a way that inflates effects and increases rates of false positives in the literature (Bakker et al. 2012). From a broader angle, valuation of high competition, emphasis on novelty, productivity and originality actually prevent the realization of science's aims because they make the proliferation of bad methodological practices inevitable. Smaldino and McEreath (2016) call this the "natural selection of bad science." Among philosophers, Heesen (2018) demonstrated through a rational choice model that the incentive structures of the credit economy of science praised by the proponents of the invisible hand idea makes low replicability of findings inevitable, because it "incentivizes scientists to focus on speed and impact" at the expense of replicability. Romero (2017) argues resonantly that the priority-based credit system of science creates a tension between the values of novelty and replicability. Bright (2021) adds that the credit incentive even abets and encourages fraud.

Bad research habits and other norm-deviant behavior are not only very common, but they also resist numerous calls for amelioration. There are several candidate explanations for the persistence and prevalence of QRPs and misconduct (see Gopalakrishna et al., 2022; Haven & van Woudenberg, 2021; Holtfreter et al., 2020)) such as publication pressure, problems in training and mentoring, the academic culture, peer norms, and low deterrence of misbehavior, all of which have been variously linked to the priority-based credit economy of science and the highly competitive institutional culture it creates. The quality control mechanisms of science seem to be dramatically inadequate to prevent the continuous introduction of error and misinformation into the scientific literature and the

eventual canonization of falsehoods as knowledge. Moreover, the system of mutual citation seems neither to consist in the social mechanisms nor serve the functions that the invisible hand explanation takes it to do. Citation networks are shown to be able to create unjustified epistemic authority through distortions (Greenberg, 2009). Let alone being a more or less reliable indicator of scientific quality, citations do not seem to be correlate with scientific quality in any useful way. For instance, citations cannot be used to distinguish between studies that replicated and those that failed to replicate (Arslan & Eleftheriadou 2019; Schafmeister 2021), not even between articles that are retracted and those still in the literature (Candal-Pedreira et al., 2020). As the interesting case of a phantom reference shows, sometimes a high number citations do not even indicate that the article, the journal or the authors exist (Harzing & Kroonenberg 2017).

Most meta-scientific research and theorizing also calls for some form of change in the scientific institution, ranging from virtuous interventions in the credit economy to adjust the rules to increased governance of science by strengthening institutional gate-keeping and sanctions to better enforce scientific norms. Obviously, even setting aside an emphasis on stronger norms and sanctions, the idea of engineering incentives makes an uneasy bedfellow with that of invisible hands. Some go still further. For Tiokhin et al. (2021b), the whole idea of achieving collective ends through the pursuit of individual interests might be mistaken, because it “creates a conflict between what is best for scientists’ careers and what is best for science,” and we should organize science around with a focus on collective epistemic aims.

### **Epistemic virtue and credit-seeking are behaviorally distinct**

Behavior that manifests epistemic virtue and behavior that is motivated by credit market incentives are indeed behaviorally or consequentially distinguishable under the same incentive structures. Arguably the reason why they seemed indistinguishable to the proponents of the invisible hand argument is their disregard for circumstantial variables due to their exclusive attention to the most successful examples of science.<sup>5</sup> A virtuous epistemic character would diverge from others regarding behavioral and epistemic outcomes under different conditions affecting the social-institutional context of research.

Let’s imagine three kinds of scientist, which we can label the *lawful good*, the *chaotic good* and the *chaotic neutral*. The lawful good scientist has a motivation for unveiling nature’s secrets, advancing human knowledge, or developing the best models to represent the world. This scientist has a relatively disinterested concern for epistemic goods, thus maximizing credit is not part of the motivational

explanation of why they choose method X over method Y, how they interpret the implications of the results for the hypotheses under test, how they go about error-control and use researcher degrees of freedom (Simons, Nelson & Simonsohn 2011), how they determine which part of the scientific literature to build on, how they respond to the critique of peers and criticize other published work, why they decide to publish or file-drawer a study, and so on. They try to maximize the epistemic rationality of their actions, and this concern guides how they assess own competence, methodological strategy, and the epistemic risks involved in the situation. Thus, if they doubt their own competence in a given method, they refrain from using it or collaborate with more skillful others. If all available methods lack sufficient reliability (e.g., construct validity or measurement invariance), they suspend judgment on certain research questions until further methodological advances are made or maybe even desire to engage in the development of new ones. They do not avoid an honest discussion of their own work or hesitate to point out flaws in others' work. The list goes on.

The chaotic good scientist is interested in getting credible empirical results, uses methods that are at least standard in that field, avoids publishing false or unjustified scientific claims and so on. But this scientist does all that also, or ultimately, for prestige, reputation, academic promotion or some other such aim. Thus, truth is a significant concern for them but not the primary concern. They might polish and oversell their results, but they do not fake research. They might go for relatively high-impact fields and questions, but would try to get it reasonably right in answering them. They have a higher risk of self-serving bias in scientific reasoning. They may tend to deal with ambiguities and degrees of freedom in the research process in way that increases publishability, while possibly also convincing themselves that the strategy that maximizes publishability also maximizes epistemic reliability. Their choice of methods is not solely guided by reliability. They may opt for less reliable or riskier methods, theories or research questions in order increase their chances of standing out if the competition in the field is fierce. They would increase the transparency and standards of rigor in their research if the costs involved are relatively little and there is some recognition to gain for it. They are less inclined to do replications, search for error or other flaws in published research or do rigorous peer-review. They do not put extraordinary effort in developing measures or infrastructure, or in own further training. They are less inclined to collaborate with others (as the potential credit is distributed) or invest in auxiliary or second-order competences, which are not often or directly

rewarded (e.g., highly specialized skills, skills for error-control, knowledge of measurement theory, calibration know-how).

Different circumstances can place different demands on the intellectual character of scientists, in order for them to reliably and responsibly contribute to the advancement of scientific knowledge. It is easier to be epistemically virtuous in some circumstances than others, and in some circumstances epistemic virtue requires extraordinary qualities. Just as some kinds of knowledge are relatively easy to acquire with a common level of epistemic competence and some kinds of knowledge require exceptional cognitive skills, what kind and degree of epistemic motivations are sufficient for responsible scientific inquiry differs depending on the circumstances. If a scientific field already has well-established and objectively reliable scientific standards, protocols and methods, has reached a certain theoretical maturity, is on a progressive track, and has plenty of resources, there will not be any significant difference in the epistemic quality of the inquiry procedures and outputs of the lawful good and chaotic good scientists. But if the prevailing norms and standards are weak, the dominant theory has started to accumulate anomalies, or there is a scarcity of research funding and academic jobs, there can emerge a gap between the demands of epistemic rationality and practical rationality. Then we can expect the lawful and chaotic good scientists to diverge behaviorally. In a context such as the credibility crisis, on top of dramatically low scientific standards, we also see serious contextual pressure that creates incentives for epistemically bad research practices such as high competition for scarce funding, too few employment or promotion options together with temptations such as a high demand for flashy results combined with low deterrence for questionable research practices or scientific fraud. In such circumstances, making a reliable and credible contribution to the advancement of scientific knowledge requires a higher level of epistemic virtue. Thus, we might expect to see the chaotic good scientist succumb to pressure and start cutting corners, thereby producing potentially misleading or at least insufficiently rigorous research outputs.

Lastly, the chaotic neutral scientist does not have any significant epistemic motivation, no concern for truth whatsoever, but only a credit-motivation. They pursue good epistemic ends only if doing so brings more credit. When there are low-hanging fruits (low-risk high-gain research) they pick them up, but they would be more reluctant to pursue conservative research strategies or to contribute to incrementally progressing fields or projects. They are not responsive to norms of scientific inquiry if these are inadequately enforced through scientific quality control mechanisms and sanctions. They respond only to a gain-loss calculus involving the prospects for credit, the costs of producing

publishable and citable research outcomes, and the risks of not passing through scientific quality-control and gate-keeping. The chaotic neutral scientist has a much higher risk of engaging in QRP or even outright fraud and falsification. If these become even slightly more profitable than conforming to the standards of research, such as when the prospects of getting flawed research published and potentially cited are reasonably high (due to scientific quality and error control mechanisms being ineffective), and the risks of sanctions for misbehavior are low (due to the vagueness of standards or difficulty of establishing misbehavior), the chaotic neutral scientist easily becomes a free-rider.

The invisible hand argument has it that the chaotic neutral scientist and the lawful good scientist behave in the same way and produce similar quality results, because the scientific practices of the lawful good scientist are also those that are the most rational in economic terms. The outcome of the incentive structures of science, however, are dynamic and context-dependent. This analysis could hold only in rather rare and transient periods of scientific success in which the chances of making scientific discoveries are so high that the scientific endeavor resembles a gold rush. When the gold veins are depleted and the scientific community must proceed with more caution, criticality, and diligence, the attitudes and behavioral tendencies of chaotic neutral scientists turn into an epistemic liability.

### **The social structure of science is not a sufficient explanation of its success**

Wray (2000) offered a criticism of Hull's virtue radicalism that nonetheless shares most of the premises of the invisible hand argument. He maintains that an invisible hand explanation for the success of science has to be combined with a hidden hand explanation. The scientific institutions are intentionally designed in a way that scientists' pursuit of credit contributes (unintentionally) to scientific knowledge generation. Thus, we cannot completely deny an explanatory role to individual epistemic virtue, and we have to significantly qualify the virtue radicalist conclusion. Wray assumes as Hull does that scientists are primarily motivated by credit, such that the success of science is not due to their intentional pursuit of truth. Nonetheless, a concern for truth is not completely out of the picture, because "many of the practices and institutions constitutive of science are intentionally designed by scientists with an eye to realizing the very goals that Hull believes need to be explained by reference to an invisible hand mechanism" (Wray 2000). A hidden hand with properly epistemic intentions designs the rules of the game and the subsequent outcomes of the game can be explained purely in reference to the rules and the reward-seeking behavior of the players. It follows that as long as the institutional structure of science and the standards of scientific practice are initially

designed competently and out of a concern for truth, collective epistemic virtue can be achieved even with only chaotic neutral scientists.

Wray's argument offers a significant counterpoint to the idea of the invisible hand by maintaining that the incentive structures do not self-organize into epistemically virtuous configurations; they are designed that way by epistemically virtuous people. But we can go much further. There is not only an initial design of scientific institutions with an eye to the advancement of scientific knowledge required, but regular "virtuous interventions" to the credit economy are also necessary in order to maintain the progressive character of a scientific community. This takes us still further away from the idea of an invisible hand.

The current reward system of science incentivizes QRPs and research misconduct and leads to the proliferation of bad methods as metascientific studies amply demonstrated. External reward systems (credit as well as promotions, grants, awards) by their nature incentivize gaming or free-riding (in the sense of getting credit for epistemic contribution without making any real contribution or getting more credit than the research deserves), and they do so to the extent that getting away with flawed or fake research is possible. Scientific quality control is far from being fail-safe. There are sanctions for misbehavior, but they are too low compared to those for misbehavior in other institutions, and misbehavior is too difficult to clearly identify. The process of peer-review which stands at the heart of scientific quality control, for instance, falls dramatically short of being a reliable social mechanism for detecting error and fraud (Anderson et al. 2013; Shatz 2004; Smith 2006). The deficiencies of the peer-review system and the need to reform it have been variously noted (see also Gibson 2007; Nosek & Bar-Anan 2012). With sufficient ill-intent and skill one can possibly pass their research reports through any quality control process. This epistemic vulnerability is not a problem when most scientists are indeed trustworthy in their intentions and most scientific error is honest error (Resnik & Stewart 2012). If there were no external rewards in science, it would probably attract only scientists with such intentions. But since external reward systems create incentives to exploit the epistemic vulnerability of the system, they also invite divergent motives. The scientific institutions are also not originally designed to fit the challenges of the industrial scale science we have today, which amplify the epistemic vulnerability of the system. So, the current combination of external reward systems with inherently vulnerable filtering and gatekeeping mechanisms makes it necessary that the rules of the game (e.g., publication standards) are constantly monitored, modified or calibrated with a view to decreasing incentives for free-riding and increasing incentives for research



quality and rigor in the credit economy. On top of this, science also goes through constant social and technological changes, such as the transition from individual research to large collaborations or advances in research infrastructure, which requires appropriate modifications to the norms and standards of credit allocation. All this implies that the practices and institutions of science cannot be designed once and for all in a way that effectuates collective epistemic success, as Wray contends—the incentive structures of science must change and evolve in ways that cannot be pre-determined by initial design. Thus, a system which centrally features external rewards needs virtuous interventions by people who would in effect act outside of the credit economy in terms of their motivations, who would act as guardians of scientific goals and values. Let us dwell a little further on this last point.

Epistemically problematic incentive structures do not only proliferate bad methodological practices, they also select certain intellectual characters. Scientists who have a strong motivational concern for truth, who place epistemic values above all others in their attitude towards research, cannot easily survive within such a community, while those do not have the slightest concern for truth can potentially thrive. The force of incentives applies the most to untenured scientists (Fanelli et al. 2015), but less so to established ones, and not at all to those who fulfil mainly administrative roles as heads of scientific institutions. These scientists are supposed to act as ‘guardians’ of scientific values. However, there is no additional, specifically designed credit allocation mechanism to create a strategic or instrumental interest in truth in them, so they must have a non-instrumental motivation for the truth. It is clearly possible to have a collectively virtuous scientific community that involves many chaotic neutral scientists, as the hidden hand argument has it, as long as the individual interests of these scientists are aligned with implementing epistemically good scientific practices. But maintaining a critical ratio of lawful good scientists is necessary in order to uphold scientific values and preserve the epistemic reliability of scientific institutions. The dominance of chaotic neutrals in the population of scientists could increasingly corrupt scientific institutions and reduce their capacity to safeguard collective scientific goals and values. In the worst case, we can end up with a scientific community that cannot even get to change its perverse incentive structures because there are not enough scientists who non-instrumentally value the aims of science among the ranks of executives at scientific institutions and science policy makers.

### **The explanatory role of epistemic virtue**

Many practices that contribute to the success of science need to be explained in terms of individual epistemic virtue.

The invisible hand mechanism can never promote some goods that are clearly in the collective epistemic interest, such as research transparency and methodological rigor. Well-conducted peer-review contributes to the success of science by filtering flawed research, but it is never in anyone's individual interest to do it well. Sharing one's data and code allows others to assess the reliability of research, but it is in a researcher's best interest not to do so (and most do not). Increasing methodological rigor often comes with serious extra costs (Forscher et al. 2020). Formulating intersubjectively testable and falsifiable hypothesis is in the interests of science, but not necessarily in the interests of individual scientists. The structure of the scientific credit economy is not enough to explain why scientists would do these things. Let us expand on these points a little more.

Some defenses of the current credit system of science remarkably conflate its benefits and risks. For instance, Heesen (2017) and Bright (2021) argue that the priority rule can even explain the origin and persistence of the so-called Mertonian norm of communism (1973), which is the common ownership of scientific knowledge. According to this reasoning, scientists are incentivized by the priority rule to share their results with the widest possible audience as soon as possible because this way they can get maximum credit for their discoveries. Anybody who does research involving data knows, however, that sharing one's results is not the same as sharing one's data and analysis code. While sharing one's results is in one's interest given the priority rule, sharing one's data and code is not. As long as this is not mandatory, nobody is incentivized to do so because it only diminishes one's prospects of getting credit as anybody can have at least honest errors or questionable choices in one's analysis strategy. If and when people are transparent about their research, this is much better explained by norm-following or (when the norm is weakly enforced) by intellectual honesty. One can be reminded of incentives such as preregistration or open data badges, but it is easy to see that the chaotic neutral type of scientist will always be engaged in an effort to go around, or game, whatever measures or metrics we use to evaluate transparency. One could also object that any tinkering with the data or bad choices in data handling or analysis lowers its replicability and thus one would not be credited for unreliable results. One could be so credited by a conceivable system of quality control, because it is in the clear disinterest of science to conduct each reported study from scratch by independent institutions or peers due to sheer inefficiency. One could also object that at least the researchers who want to make use of the results would know that they are unreliable and thus would not cite the study. At this point it is pertinent to invoke another closely related

claim, namely that the priority rule and the resulting mechanism of mutual citation ensures the self-correction of the scientific literature (and incentivizes reliability and rigor). Hull (1997) says that

“the chief credit in science, the currency that really matters, is use. Scientists do not use the results of other scientists idly. In adopting the views of other scientists or using their data, scientists are voting with their career. [...] With the possibility of credit comes the possibility of discredit. One effect of uncovering mistakes in the work of other scientists is increased care in using the work of these scientists in the future. In fact, often one egregious instance is enough for a group of aggrieved scientists to cease using the work of those scientists found guilty [... *Scientists*] get credit for discovering mistakes in the work of others, especially if this research is taking place in one of the “hot” areas of science. The rush to publish, when properly constrained, increases the pace of science. The monitor on this pace is the punishment meted out to those scientists who produce unreliable results. [...] Erroneous results hurt the careers of anyone who uses them. Thus, [...] science can be viewed as a self-policing system of mutual exploitation or, if you prefer, cooperation.”

While neat as a philosophical analysis, we know that scientific findings that are of no use epistemically (be the cause a false positive, fraudulent research, or that they are unreproducible) seem to be ubiquitous in the literature, which suggests wide-scale QRPs like p-hacking, publication bias, fraud, or a terrifying mixture of these (Bakker et al. 2012). Misinformation can also find use for itself in practice. Let alone replication failures, even retractions do not prevent researchers from widely citing these papers. Moreover, retracted papers seem to perform better than others in science communication and dissemination (Peng et al. 2022). These cannot be easily explained away, since the incentive system is pretty much in force as described in all these cases. The most efficient strategy given the current rules of the credit-seeking game is to get publishable results out there as fast as one can following methods that lead to unreliable or misleading results, and with the least degree of transparency that is allowed (Bakker et al 2012; Simons, Nelson & Simonsohn 2011). Consequently, researchers who make all the potentially relevant aspects of their research procedure publicly available and reproducible for others to freely use, who make the extra effort to maximize the accessibility and usability of their data and code by proper labeling, indexing and so forth, are not behaving in a maximally rational manner if modelled as credit-seekers. Thus, the priority rule does not explain (away) the communist norm in science, neither does it explain mutual criticism or error-correction, which are part of the explanation of its success.

## **2b. The evolutionary/sociobiological argument**

Hull (1978) also defended another version of this argument for the success of science, similarly radicalist about epistemic virtue, from an evolutionary rather than strictly economic perspective. The evolutionary version of virtue radicalism is advanced to explain the altruistic behavior of scientists, such as disinterested inquiry, freely sharing knowledge, policing themselves for error and flaws, and giving credit when it is due. Hull follows the general methodological principle that sufficiently selfish reasons must be found to explain a case of seemingly selfless behavior. He argues that what appears to be epistemically virtuous (i.e., altruistic) behavior in scientists is actually self-interested behavior aimed at recognition and credit, in the form of appropriate citations, awards, appointments and prestige. As organisms compete to transmit their genes, scientists compete with each other “to have one’s own ideas incorporated into the generally accepted body of scientific knowledge” (Hull 1978). But the pursuit of this selfish goal by individual scientists is well-calculated to realize the manifest goal of science as an institution given the social organization of science. In analogy with theories of kin selection and parental manipulation that maintain that altruistic behavior can be an optimal strategy for the transmission of own genes, Hull argues that scientists conduct and share research in the benefit of the scientific community and the public because this enables them to get their ideas accepted by others most reliably and efficiently in the form of citations. Hull maintains that on the whole scientific worth correlates with citations (if not merely quantitatively but also qualitatively assessed), which in turn correlates with other forms of recognition. This argument concludes that the success of science is due to its social structure that enables the institutional goals of science to coincide with the scientists’ individual interests by facilitating cooperation among competitors.

### **The rationality of virtuous behavior**

Applying insights from evolutionary theory to the social organization of science does not require that we model competition and selection exclusively at the individual level (see also Tiokhin et al. 2021b). As scientific knowledge is collectively produced, it is often the case that scientists have supraordinate goals. Even when scientists do not organize in teams or larger research collaborations, they typically strive to answer a common set of research questions. While we can say that scientists would prefer that they be the one to answer a question, they also care that the question is successfully answered even if by others, because scientific and technological progress is in the interest of everyone. In a reliably progressing field, it is not surprising to find many scientists who would unhesitatingly prefer an epistemically better idea by another scientist to be canonized instead

of their own, even if they would have no prospect of building on it later in their research. Some behaviors of scientists are truly altruistic, in the sense that they aim to maximize the chances for collective rather than individual success. Many scientists are more rigorous in their research than required by publication standards, engage in error detection in the published literature when the topic is not closely related to their primary activity and with no prospect for personal credit (sometimes even with high prospects of reproach), conduct replication or self-replication studies when it is neither required nor rewarded, do quality peer-review, engage in teaching and mentoring and so on. This is true even in the fields suffering from a credibility crisis. Doing so is rational from the collective viewpoint, because it increases the chances that the “jackpot” hits at least one individual or team. If everybody invests their time and effort to maximize their own chances for credit, there is a real possibility of group failure in the sense that we cannot reliably answer some research questions that we set before ourselves. Thus, the motivation to obtain reliable findings, in own research as well as in the literature, cannot be reduced to a motivation to achieve acceptance or credit.

### **2c. Argument from the division of cognitive labor (or the systemic argument)**

The first two arguments asserted that thanks to the incentive structures of the credit system of science, epistemically virtuous inquiry is behaviorally or consequentially indistinguishable from inquiry motivated only by maximizing credit. Formulated differently, the miracle of the social structure of science is that what is instrumentally rational for scientists to do in order to maximize credit is also the epistemically most rational choice. The third argument differs from these in asserting that thanks to the incentive structure of its credit economy science can be successful even when scientists clearly deviate from what epistemic rationality would require, such as choosing the better method or the better theory over alternatives. Kitcher, who is one of the earliest and the most prominent proponents of this argument, deems it “possible that there should be a mismatch between the demands of individual rationality and those of collective (or community) rationality (1990).” This is because there is no guaranteed success in inquiry, only different prospects for success. An epistemically irrational inquiry strategy might eventually turn out to be the successful one while the rational strategy fails. From the perspective of collective epistemic rationality, the aim should be an optimum division of cognitive labor. If all individuals in a scientific community behaved in accordance with epistemic rationality, the scientific community would put all its eggs into one basket. A better distribution of cognitive labor that maximizes the chances of success at the

level of the scientific community or field would include a diversity of strategies. However, epistemically virtuous scientists do not make any epistemically irrational choices. Thankfully there are other norms in science that incentivize divergent inquiry strategies. The credit incentive in science serves science better than purely epistemic norms given an appropriate arrangement of the credit economy. Kitcher (1990, 1993), and following him Strevens (2003, 2011), praise the credit economy of science organized around the priority rule for serving exactly this function: The priority rule motivates scientists to follow divergent research strategies which eventually leads to a better division of cognitive labor in the scientific community than the case where all scientists behave in accordance with epistemic rationality. Strevens maintains further that the priority rule brings about the best reward system for science because it leads to the most efficient allocation of resources for maximum benefit to society (2003). Given this reward system of science, a scientific community where the choices of scientists are motivated by recognition and prestige can perform even better epistemically than one consisting of high-minded inquirers. Kitcher (1990) explicitly maintains that a “Hobbesian” scientific community might work much better (in making a discovery) than a community of the high-minded spirits who would fail to divide cognitive labor between epistemically reliable and unreliable methods. Thus, “the thirst for fame and fortune [...] might actually play a constructive role in our community epistemic projects [...] social institutions within science might take advantage of our personal foibles to channel our efforts toward community goals rather than toward the epistemic ends that we might set for ourselves as individuals.” This argument thus agrees with the first argument in its conclusion that that epistemic virtue is neither necessary nor sufficient for the success of science, as the structure of the scientific credit economy is adequate for collective epistemic virtue.

### **The normative inadequacy of discovery-oriented values**

Kitcher’s argument focuses on the context of discovery and disregards the context of justification within scientific inquiry. In the context of discovery (i.e., formulation of novel hypotheses, choice of which theory to pursue given insufficient evidence, developing methods to detect conjectured phenomena etc.), the choice of the most epistemically rational act is characteristically underdetermined (weaker options also have non-zero success chance). So, epistemically irrational behavior that is motivated by instrumental rationality can possibly contribute to the production of collective epistemic goods under specific circumstances. The existence of even completely chance discoveries in the history of science shows further that the road to discovery can be most variable.

This analysis does not work in the context of justification (theory/ hypothesis testing). In the context of justification, we aim to minimize researcher degrees of freedom, rather than increase them, because flexibility in the interpretation of results invites bias, noise and error. We also aim to maximize intersubjective testability and verifiability; thus we desire qualities such as replicability, robustness, validity, and invariance. Here we cannot make room for epistemically irrational acts, as they are certain to undermine the collective epistemic good. Thus, while it is reasonable that a competitive system of external rewards such as the current scientific credit economy serves the interests of science in the context of discovery, while it has the opposite effect in the context of justification. As Romero (2017) similarly argues, the current priority-based reward system of science is not “normatively adequate,” because it creates a tension between the values of novelty and replicability.

The contexts of discovery and justification are inseparable aspects of inquiry, but the context of discovery demands a different value system than the context of justification, thus there is a normative conflict within scientific inquiry. Qualities like ambitiousness, epistemic risk-taking, prestige-seeking, self-interestedness can contribute to the proliferation of novel ideas, hence can be desirable for the sake of collective epistemic good in the context of discovery.<sup>6</sup> In the context of justification however, the desirable qualities are epistemic competence, intellectual diligence, critical thinking, intellectual humility, intellectual carefulness, disinterestedness, impartiality, and so on, which are all intellectual virtues, broadly conceived. Thus, the ideal type of scientist for the purpose of discovery is the well-awarded scientist, but the ideal type for rigorous, trustworthy research is the intellectually virtuous scientist.

We cannot maximize both productivity (i.e., discovery) and reliability with a single value system. There is always a trade-off. Thus, for a scientific community to be epistemically virtuous, a set of norms which are based on purely epistemic values and promote epistemically virtuous research is necessary besides a system of external rewards with its associated values. This brings us to the topic of scientific norms, which is the pivot of the sociological argument.

### **3. The sociological argument for virtue radicalism**

According to the famous sociologist of science Merton, the success of science should be explained not in reference to psychological/cognitive variables pertaining to scientists but to the norms enjoined by scientific institutions. Primarily, these are the “four sets of institutional imperatives

taken to comprise the ethos of modern science [...] communism, universalism, disinterestedness, and organized skepticism” (1973). This is a particular application of a more general methodological heuristic in sociology, that differences in behavior across institutional contexts should be explained in terms of the social structure of those institutions (Hull, 1978). Merton (1973) says that it is not “distinctive motives” but “rather a distinctive pattern of institutional control of a wide range of motives which characterizes the behavior of scientists.” If scientists behave in ways that are more epistemically virtuous than other professions, this is due to the rules, standards and norms that are enforced by scientific institutions. Scientists may comply with them for fear of penalty or they may also internalize them to mitigate psychological/cognitive dissonance. This is contingent and inconsequential from the perspective of collective outcomes. What matters is that it is in their best interest to comply with the norms.

Bridgstock (1982) describes how scientific institutions enjoin epistemic norms through reward and punishment:

“Such values are reinforced according to a rewards system within science that utilizes positive and negative sanctions to influence behavior. Positive rewards such as job security and promotion, citation, research grants, and honorific awards are given to those that adhere to the values of science. In contrast, negative sanctions such as dismissal, cessation of research, and suspension of grant writing are wielded against those subverting the common set of scientific values.” (Quoted in Sovacool, 2008).

The sociological argument resonates with the invisible hand argument and the sociobiological argument in maintaining that scientific success does correlate necessarily with virtuous inquiry. Scientists’ seemingly virtuous behavior is actually conformity to norms and sanctions, and the best (or the most parsimonious) way to explain this conformity is not an intellectual endorsement of the values promoted by those norms and sanctions but a motivation for seeking reward and avoiding punishment.

### **The top-down and horizontal mechanisms for normative behavior**

The sociological argument for virtue-radicalism focuses on the top-down enforcement of norms and ignores the horizontal mechanisms for norm-compliance, which in turn have systemic or collective consequences. One such horizontal mechanism is peer normative pressure, another is the influence of exemplars of epistemic virtue.



Since epistemic virtue has an inherent motivational element, such as sufficient concern for truth, collective virtue has an analogous counterpart. The collective counterpart of a proper epistemic motivation can be found in shared values, standards and goals that incorporate a concern for truth. We can conceive the normative constitution of a scientific community in terms of its ‘epistemic ethos’, as Merton also does. But an epistemic ethos can be understood more comprehensively than a few institutional imperatives. It has to do with the normative constitution of a scientific community, which involves the collective endorsement of a large set of values, goals and standards.<sup>7</sup> This normative constitution has two aspects: the content of the values, goals and standards themselves (by which they can be evaluated for their truth-conduciveness), and the normative pressure on individual actions that results from their collective endorsement. A good scientific ethos involves collectively shared norms and standards of knowledge-conducive or reliable scientific practices. If any of these aspects are deficient, we can end up with a bad epistemic ethos and thus collective epistemic vice. First, (even if the Mertonian institutional imperatives are in place in the abstract) it is possible that a scientific community adopts the wrong values, standards or goals, in the sense that non-epistemic values take precedence over epistemic ones (such as academic productivity or profitability instead of credibility, rigor or progress) or there is a mistaken idea about intermediate or ultimate epistemic goods (such as the best methodological practices or the right disciplinary goals). Or there might be nothing wrong with the values, goals or standards adopted but these may be nonetheless insufficiently endorsed and thus bereft of normative force. By virtue of collectively endorsing certain scientific values, goals and standards, the members of a scientific community gain a standing to criticize those who do not comply with or implement them. In a bad scientific ethos, deviant practices would generate insufficient criticism and epistemically justified practices could meet normative resistance.

The credibility crisis showed epistemic deficiencies in both. Lundh (2019) maintains that it involves a normativity crisis: The norms according to which the research community functions, those that govern research and those that govern the credit economy, are not conducive to scientific progress. Widespread QRPs, weak publication standards, and the incentivization of publication count over scientific contribution are the most commonplace examples, which feature a mixture of adopting wrong norms and false beliefs about what is epistemically good or bad. For instance, quite many scientists either do not know why QRPs are bad or underestimate how bad they are. Simmons, Nelson and Simonsson (201) write that many researchers who p-hacked shared a false belief about

how serious a vice it is: “We knew many researchers—including ourselves—who readily admitted to dropping dependent variables, conditions, or participants to achieve significance. Everyone knew it was wrong, but they thought it was wrong the way it is wrong to jaywalk [...] simulations revealed that it was wrong the way it is wrong to rob a bank.”

Moreover, norm-deviant practices meet little interpersonal criticism and epistemically justified practices meet normative resistance. Engagement in QRPs is widely seen as “normal misbehavior” (De Vries et al., 2006). Most scientists who subscribe to scientific norms against QRPs nonetheless see them as the peer norm or descriptive norm (Gopalakrishna et al., 2021), and this indicates a conflict between what the scientists commonly approve and what they commonly do—their injunctive and descriptive norms (Cialdini et al. 1990). This is worrisome because a conflict or disparity between injunctive and descriptive norms undermines the desired effect of injunctive norms by reducing people’s motivation to engage in norm-compliant behavior (Smith et al. 2012; Kallgren et al. 2000). Moreover, vocalizing suspicion of QRPs is difficult. For instance, a decade ago it was widely counter-normative to call out a replication failure. While many researchers working on a particular research question might have unsuccessfully tried to replicate a famous finding, they would confess this to others off the record, so the belief in the unreplicability of the finding would not turn into a collective belief (or eliminate the collective belief in the finding). Any occasional public announcement of replication failures used to meet with rebuke. On a more general level, the scientific reform movement in response to the replication crisis experienced normative clashes everywhere: it was counter-normative to go through the methods and results in the published literature and call out a serious suspicion of QRPs or data fraud. People who did so have even been labelled “methodological terrorists.”<sup>8</sup> Accordingly, people who identify errors in the scientific literature did not have the *standing* to criticize other’s scientific claims although they were epistemically *justified* to do so (Weatherall and Gilbert 2015). Here we have a more serious problem than the occasional unwarranted collective belief in certain scientific claims (which might result in epistemically irresponsible insistence on holding onto them), a collective resistance to forms or methods of epistemically justified criticism (such as doing replication studies), which in effect significantly restricts the venues for critical discussion –thus, a lack of a critical culture.

Coming to the role of exemplars, the existence of epistemically virtuous scientists itself is a normative force, especially in the training of younger generations of scientists. Exemplars are important because intellectual habits are also acquired by social learning, and they contribute

significantly to increasing the normative force of scientific values. Exemplarism is one of the central ideas in traditional virtue theory. The guiding idea behind it is that the most natural development of virtuous character is by observing and emulating people that one admires, which allows one to gain a much clearer understanding of appropriate conduct than a study of norms in the abstract as well as gives one a stronger motivation to act in the right way (Croce & Pritchard 2022). Many good scientists indeed say that they learned the most about proper scientific inquiry from their mentors or academic heroes, rather than textbooks. De Ridder (2022) observes a similar influence in research teams: “when a group member sets an example by an impressive display of virtuous behavior, others will want to live up to that ideal and strive to improve their own behavior in the image of that ideal. Especially when stories about an exemplar are often repeated or when little ritual-like practices are formed around it, they have a lasting influence.”

Both of these points have negative implications for the sociological analysis of how scientific norms acquire their normative force. Norms cannot acquire and maintain their normative force only through reward and punishment. First, externally motivated behavior tends to undermine or go around the norms when opportunity arises as we argued for in the preceding. We can assume that opportunity will constantly arise, because norm-deviant behavior and misbehavior detection mechanisms are always in an arms race: scientists who intend to get credit for lesser or no accomplishment will invent more ingenious ways to pass through detection and the detection mechanisms will have to keep advancing to keep up with those innovations. Tightening institutional control is the intuitive solution but it comes with a serious downside, because it harms scientific autonomy, productivity and efficiency. Thus, we would want high norm subscription instead of mere punishment-avoidance. Findings also suggest that norm subscription is an important predictor of refraining from QRPs and misconduct (Gopalakrishna et al. 2022). Another relevant finding indicates that aiming for the appearance of epistemic virtue has different behavioral consequences than aiming for virtue itself. Janke et al. (2019) found that the goal of achieving an ‘appearance’ of epistemic competence positively predicts engagement in QRPs, while the goal of achieving actual epistemic competence (e.g., striving for skill development) is a negative predictor. Moreover, they found that “the more researchers favored publishing over scientific rigor, the stronger the association between appearance-oriented goals and engagement in QRPs.” On the other hand, it is quite optimistic to assume that institutions can actually reward epistemic virtue or adherence to scientific values. Rewards do not incentivize virtue but its appearance. When we set out to

normatively evaluate epistemic performance, we resort almost exclusively to academic performance indicators, which actually are not measures of *epistemic* performance. But more importantly, most of these indicators now turned into false signs that do not indicate any other quality of researchers' than their ability to produce them. Because, as Donald Campbell (1976) says, the “more any quantitative social indicator is used for social decision-making, the more subject it will be to corruption pressures and the more apt it will be to distort and corrupt the social processes it is intended to monitor.” Or as summarized in Goodhart’s law: when a measure becomes a target, it ceases to be a good measure. When scientists do not intellectually endorse scientific norms, we might easily end up with allocating rewards in a way that does not track the best epistemic practice.

A bad epistemic ethos is not a purely systemic problem. Accordingly, a good epistemic ethos cannot be credited to the enforcement of scientific norms and sanctions alone: Collective epistemic virtue is a product of the interplay of individual and collective goals and values.

## **Conclusion**

The structure of the credit economy arguably serves to increase rate and speed of scientific discovery, and to steer the attention of scientists to high-impact ideas. Without a competitive reward system, it is possible that scientists would pursue only whatever they deem worthy to know, and as efficiently as they deem appropriate. They could also refrain from making their gained knowledge public until they are perfectly sure or unless they also desire to advance collective knowledge. All these factors have a role in scientific progress, but it is a mistake to think that these serve collective epistemic interests unconditionally. As efficient and speedy scientific progress is in the collective epistemic interest, as is reliable knowledge generation and error elimination—but these are not the kind of aims that can be achieved via external rewards. On the contrary, any competitive reward system is by itself a risk factor from the perspective of reliability and credibility because it incentivizes unreliable research practices and misconduct. Thus, it is wrong to think that the same social mechanisms can facilitate fast, impactful and reliable outputs at the same time. What we can at most expect from an external reward system is not to create a substantial tension between the epistemic and practical motives of scientists so that it in practice incentivizes epistemic vice and creates a selection pressure on virtuous scientists.

An epistemically reliable scientific ethos is vital to incentivize epistemic virtue and deter epistemic vice. A necessary feature of a reliable scientific ethos is the institutional imperatives and sanctions, as

the sociological argument maintains. But it is a mistake to think that institutional imperatives and sanctions are sufficient by themselves. First, sanctions can be imposed on major deviations from scientific norms but not to unreliable or suboptimal (but legitimate) research strategies that substantially undermine epistemic outcomes at the collective level or in the long run, which constitute the main bulk of the metascientific literature on the credibility crisis. It is of course possible to change rules and standards to raise the bar on scientific rigor, as many propose to do with various scientific reform proposals that are targeted at particular problematic research practices. But if the primary motivation of researchers is to satisfy the standards but not to achieve good epistemic outcomes, adjusting or reforming particular standards easily turns into a kind of arms race between policy makers and researchers because it is always possible to create a false appearance of good research however stringent the standards are. Second, major deviations from norms are difficult to prove. It is difficult to detect QRPs at the individual level and close to impossible to establish fraud. Third, scientific quality control and error detection is extremely costly and inefficient to do at the level of each scientific output. Thus, if norm-following were done merely in fear of sanctions the scientific institution would turn into a cumbersome bureaucracy. Last, and most importantly, an epistemically reliable scientific ethos requires active stewardship. Without there being scientists who are motivated by epistemic reasons above all else, there is no mechanism to safeguard scientific institutions from epistemic corruption. This can happen if the scientific norms are reduced to a mere nominal status and institutional decisions are taken in accordance with other concerns. In that case we can also expect a gap to emerge between injunctive and descriptive norms, as it is actually the case in several fields at the moment. When that happens, it is impossible to rehabilitate a degenerated ethos without active efforts by epistemically virtuous scientists. Furthermore, scientific norms and standards are not set in stone. They must also be re-interpreted, evaluated and updated in the face of novel challenges to scientific progress, because they are not valued in themselves but for their instrumental role in bringing about good collective epistemic outcomes. Epistemically virtuous scientists are necessary to navigate situations of axiological uncertainty or crisis, where the scientific community has to collectively reason to redefine rules, standards, goals and values in accordance with epistemic rationality. For all these reasons, the scientific ethos like any other normative system needs a significant proportion of individuals who endorse those norms autonomously.

Since we have maintained that epistemically virtuous scientists, especially the lawful good characters, might have difficulty in competing in a cut-throat credit economy, we might need to create and maintain niches for them. These can be in the form of diversifying the recruitment and promotion criteria or creating academic positions dedicated to other contributions to science than first-order scientific outputs such as science policy development, design or coordination of large collaborative projects, or systematically conducting replication studies and meta-analyses. Of particular importance is broadening our evaluative criteria for scientific production to include performances that contribute the advancement of scientific knowledge in indirect ways, such as engaging in rigorous peer review, serving as good teachers and mentors, facilitating interdisciplinary communication and coordination, undertaking replication research, or detecting fraud and errors in the scientific literature. Tiokhin et al. (2021b) call such contributions ‘indirect effects’, which is meant to denote the causal effect of a scientist’s epistemic performance on a measurable scientific outcome that goes through (or mediated by the epistemic performances of) other scientists. Ignoring such contributions to collective knowledge production in the credit economy of science not only disregards these and similar actions but creates the side-effect of a negative selection pressure on scientists who engage in them.

## References

- Altman, D. G. (1994). The scandal of poor medical research. *BMJ*, *308*(6924), 283-284.
- Anderson, M. S., Shaw, M. A., Steneck, N. H., Konkle, E. and Kamata, T. (2013). Research Integrity and Misconduct in the Academic Profession. In M. B. Paulsen (ed.), *Higher Education: Handbook of Theory and Research*, Vol. 28, Dordrecht: Springer, pp. 217–61.
- Arslan, R. C., Eleftheriadou I. I. (2019). ‘Revised: Are studies that replicate cited more?’, Accessed: 28.08.2022. Available at: <https://rubenarslan.github.io/posts/2019-01-02-are-studies-that-replicate-cited-more/>
- Bakker, M., Van Dijk, A., & Wicherts, J. M. (2012). The rules of the game called psychological science. *Perspectives on Psychological Science*, *7*, 543-554.
- Bridgstock, M. (1982). A sociological approach to fraud in science. *Journal of Sociology (Melbourne, Vic.)*, *18*(3), 364–384. doi:10.1177/144078338201800305.
- Bright, L. K. (2021). Why Do Scientists Lie? *Royal Institute of Philosophy Supplements*, *89*, 117–129. <https://doi.org/10.1017/S1358246121000102>
- Button, K. S., Ioannidis, J. P., Mokrysz, C., Nosek, B. A., Flint, J., Robinson, E. S., & Munafò, M. R. (2013). Power failure: why small sample size undermines the reliability of neuroscience. *Nature Reviews Neuroscience*, *14*, 365-376. <https://doi.org/10.1038/nrn3475>
- Camerer, C. F., Dreber, A., Forsell, E., Ho, T. H., Huber, J., Johannesson, M., ... & Wu, H. (2016). Evaluating replicability of laboratory experiments in economics. *Science*, *351*, 1433-1436. <https://doi.org/10.1126/science.aaf0918>
- Campbell, D. T. (1976). *Assessing the impact of planned social change*. Hanover, NH: The Public Affairs Center, Dartmouth College. <http://goo.gl/AAaPx>.
- Candal-Pedreira, C., Ruano-Ravina, A., Fernández, E., Ramos, J., Campos-Varela, I., & Pérez-Ríos, M. (2020). ‘Does retraction after misconduct have an impact on citations? A pre–post study.’ *BMJ Global Health*, *5*(11), e003719. <https://doi.org/10.1136/bmjgh-2020-003719>
- Chambers, C. (2017). *The seven deadly sins of psychology: A manifesto for reforming the culture of scientific practice*. Princeton, NJ: Princeton University Press.
- Cialdini, R. B, Reno, R. R., & Kallgren, C. A. (1990). ‘A focus theory of normative conduct: Recycling the concept of norms to reduce littering in public places.’ *Journal of Personality and Social Psychology* *58*, 1015-1026.

- Croce, M., & Pritchard, D. (2022). Education as the social cultivation of intellectual virtue. In *Social Virtue Epistemology* (pp. 583-601). Routledge.
- De Ridder, J. (2022). 'Three Models for Collective Intellectual Virtues.' In *Social Virtue Epistemology*, (eds.) M. Alfano, C. Klein & J. de Ridder, 367-85. London: Routledge.
- De Vries, R., Anderson, M. S., & Martinson, B. C. (2006). Normal Misbehavior: Scientists Talk about the Ethics of Research. *Journal of Empirical Research on Human Research Ethics*, 1(1), 43–50. <https://doi.org/10.1525/jer.2006.1.1.43>
- Fanelli, D., Costas, R., & Larivière, V. (2015). Misconduct Policies, Academic Culture and Career Stage, Not Gender or Pressures to Publish, Affect Scientific Integrity. *PLOS ONE*, 10(6), e0127556. <https://doi.org/10.1371/journal.pone.0127556>
- Fricker, M. (2020). Institutional epistemic vices: The case of inferential inertia. In *Vice Epistemology*. Routledge.
- Gibson, T. A. (2007). 'Post-publication Review Could Aid Skills and Quality', *Nature* 448, 408.
- Gopalakrishna, G., Ter Riet, G., Vink, G., Stoop, I., Wicherts, J. M., & Bouter, L. M. (2022). Prevalence of questionable research practices, research misconduct and their potential explanatory factors: A survey among academic researchers in The Netherlands. *PloS one*, 17(2), e0263023.
- Greenberg, S. A. (2009). 'How citation distortions create unfounded authority: analysis of a citation network.' *BMJ*, 339, b2680.
- Hardwig, J. (1985). 'Epistemic Dependence.' *Journal of Philosophy* 82(7), 335-49. <https://doi.org/10.2307/2026523>
- Harzing, A-W., & Kroonenberg, P. (2017). 'The mystery of the phantom reference.' Accessed: 28.08.2022. Available at: <https://harzing.com/publications/white-papers/the-mystery-of-the-phantom-reference>
- Haven, T., & van Woudenberg, R. (2021). 'Explanations of Research Misconduct, and How They Hang Together.' *Journal for General Philosophy of Science* 52(4), 543-61. <https://doi.org/10.1007/s10838-021-09555-5>
- Heesen, R. (2017). Communism and the Incentive to Share in Science. *Philosophy of Science*, 84(4), 698–716. <https://doi.org/10.1086/693875>



Heesen, R. (2018). Why the Reward Structure of Science Makes Reproducibility Problems

Inevitable: *The Journal of Philosophy*, 115(12), 661–674.

<https://doi.org/10.5840/jphil20181151239>

Holtfreter, K., Reisig, M. D., Pratt, T. C., & Mays, R. D. (2020). The perceived causes of research misconduct among faculty members in the natural, social, and applied sciences. *Studies in Higher Education*, 45(11), 2162–2174. <https://doi.org/10.1080/03075079.2019.1593352>

Hull, D. L. (1978). Altruism in science: A sociobiological model of co-operative behaviour among

scientists. *Animal Behaviour*, 26, 685–697. [https://doi.org/10.1016/0003-3472\(78\)90135-5](https://doi.org/10.1016/0003-3472(78)90135-5)

Hull, D. L. (1988). *Science as a Process: An Evolutionary Account of the Social and Conceptual Development of Science*. Chicago, IL: University of Chicago Press.

Hull, D. L. (1997). What's Wrong with Invisible-Hand Explanations? *Philosophy of Science*, 6, S117–S126. <https://doi.org/10.1086/392592>

Ijzerman, H., Lewis, N. A., Przybylski, A. K., Weinstein, N., DeBruine, L., Ritchie, S. J., ... & Anvari, F. (2020). Use caution when applying behavioural science to policy. *Nature Human Behaviour*, 4(11), 1092-1094.

Ioannidis, J. P. (2005). Why most published research findings are false. *PLoS Medicine*, 2, e124.

<https://doi.org/10.1371/journal.pmed.0020124>

Janke, S., Daumiller, M., & Rudert, S. C. (2019). 'Dark Pathways to Achievement in Science: Researchers' Achievement Goals Predict Engagement in Questionable Research Practices.' *Social Psychological and Personality Science* 10(6), 783–791.

<https://doi.org/10.1177/1948550618790227>

John, L. K., Loewenstein, G., & Prelec, D. (2012). Measuring the prevalence of questionable research practices with incentives for truth telling. *Psychological Science*, 23, 524-532.

<https://doi.org/10.1177/0956797611430953>

Kallgren, C. A., Reno, R. R., & Cialdini, R. B. (2000). 'A Focus Theory of Normative Conduct: When Norms Do and Do not Affect Behavior.' *Personality and Social Psychology Bulletin*, 26(8), 1002–12. <https://doi.org/10.1177/01461672002610009>

Kitcher, P. (1990). The Division of Cognitive Labor. *Journal of Philosophy*, 87(1), 5–22.

<https://doi.org/10.2307/2026796>

- Kitcher, P. (1993). *The advancement of science: Science without legend, objectivity without illusions*. Oxford University Press.
- Knorr-Cetina, K. (1999). *Epistemic cultures: How the sciences make knowledge*. Harvard University Press.
- Lundh, L. G. (2019). The crisis in psychological science and the need for a person-oriented approach. In *Social philosophy of science for the social sciences* (pp. 203-223). Springer, Cham.
- Merton, R. K. (1957). Priorities in scientific discovery: a chapter in the sociology of science. *American sociological review*, 22(6), 635-659.
- Merton, R. K. (1973) [1942]. The Normative Structure of Science. In Merton, Robert K. (ed.), *The Sociology of Science: Theoretical and Empirical Investigations*, Chicago: University of Chicago Press, pp. 267–278.
- Montmarquet, J. A. (1987). 'Epistemic Virtue.' *Mind* 96(384), 482-97.
- Munafò, M. R., Nosek, B. A., Bishop, D. V. M., Button, K. S., Chambers, C. D., Percie du Sert, N., . . . Ioannidis, J. P. A. (2017). A manifesto for reproducible science. *Nature Human Behaviour*, 1, Article 0021. <https://doi.org/10.1038/s41562-016-0021>
- Nosek, B. A. and Bar-Anan, Y. (2012). Scientific Utopia: I. Opening Scientific Communication. *Psychological Inquiry*, 23, 217-243, DOI: 10.1080/1047840X.2012.692215
- Nosek, B. A., Spies, J. R., & Motyl, M. (2012). Scientific Utopia: II. Restructuring Incentives and Practices to Promote Truth Over Publishability. *Perspectives on Psychological Science*, 7(6), 615–631. <https://doi.org/10.1177/1745691612459058>
- Open Science Collaboration. (2015). Estimating the reproducibility of psychological science. *Science*, 349. doi:10.1126/science.aac4716
- Peng, H., Romero, D. M., & Horvát, E.-Á. (2022). Dynamics of cross-platform attention to retracted papers. *Proceedings of the National Academy of Sciences*, 119(25), e2119086119. <https://doi.org/10.1073/pnas.2119086119>
- Romero, F. (2017). Novelty versus Replicability: Virtues and Vices in the Reward System of Science. *Philosophy of Science*, 84(5), 1031–1043. <https://doi.org/10.1086/694005>
- Shatz, D. (2004). *Peer review: A critical inquiry*. Tottowa, NJ: Rowman & Littlefield.
- Simmons, J. P., Nelson, L. D., & Simonsohn, U. (2011). False-positive psychology: Undisclosed flexibility in data collection and analysis allows presenting anything as significant. *Psychological Science*, 22, 1359-1366. <https://doi.org/10.1177/0956797611417632>
- Smith R. (2006). 'Peer review: a flawed process at the heart of science and journals.' *J R Soc Med.* , 99(4), 178-82. doi: 10.1258/jrsm.99.4.178.

- Smith, J. R., Louis, W. R., Terry, D. J., Greenaway, K. H., Clarke, M. R., & Cheng, X. (2012). 'Congruent or conflicted? The impact of injunctive and descriptive norms on environmental intentions.' *Journal of environmental psychology* 32(4), 353-361.
- Sosa, E. (2007). *A Virtue Epistemology: Apt Belief and Reflective Knowledge*. Oxford University Press.
- Sovacool, B. K. Exploring Scientific Misconduct: Isolated Individuals, Impure Institutions, or an Inevitable Idiom of Modern Science?. *Bioethical Inquiry* 5, 271–282 (2008).  
<https://doi.org/10.1007/s11673-008-9113-6>
- Strevens, M. (2011). Economic Approaches to Understanding Scientific Norms. *Episteme*, 8(2), 184–200. <https://doi.org/10.3366/epi.2011.0015>
- Strevens, M. (2003). The Role of the Priority Rule in Science. *The Journal of Philosophy*, 100(2), 55–79.  
<https://doi.org/10.5840/jphil2003100224>
- Resnik, D. B., & Stewart Jr, C. N. (2012). 'Misconduct versus honest error and scientific disagreement.' *Accountability in research* 19(1), 56-63.
- Smaldino, P. E., & McElreath, R. (2016). The natural selection of bad science. *Royal Society Open Science*, 3(9), 160384. <https://doi.org/10.1098/rsos.160384>
- Tiokhin, L., Yan, M., & Morgan, T. J. (2021a). 'Competition for priority harms the reliability of science, but reforms can help.' *Nature Human Behaviour* 5(7), 857-867.
- Tiokhin, L., Panchanathan, K., Smaldino, P. E., & Lakens, D. (2021b). 'Shifting the level of selection in science.' <https://doi.org/10.31222/osf.io/juwck>
- Vazire, S. (2019). Do we want to be credible or incredible? *APS Observer*, 33.
- Weatherall, James Owen and Gilbert, Margaret (2015) Collective Belief, Kuhn, and the String Theory Community. Available at: <http://philsci-archive.pitt.edu/11413/>
- Wray, K. B. (2000). Invisible Hands and the Success of Science. *Philosophy of Science*, 67(1), 163–175.  
<https://doi.org/10.1086/392767>
- Zagzebski, L. T. (1996). *Virtues of the mind: An inquiry into the nature of virtue and the ethical foundations of knowledge*. Cambridge University Press.

## NOTES

<sup>1</sup> For two major accounts, see Zagzebski (1996) and Sosa (2007).

<sup>2</sup> See also Haven & van Woudenberg (2021) for the limitations of the structural explanations of prevalent research misconduct.

<sup>3</sup> Hull and others view a proper concern with truth or knowledge as an altruistic motive. Although selfishness-altruism are not epistemic categories, proper epistemic motivations reflect an altruistic motivation from the perspective of social scientists, because the latter regard scientific knowledge as a common good and only credit or material rewards as an individual good. Purely epistemic (non-instrumental) aims are not regarded as valid motivational factors to explain behavior by themselves.

<sup>4</sup> Here are some examples of serious QRPs:

1. Selectively reporting results or studies / publication bias: This is when the outcome of an experiment or research study influences the decision whether to publish or otherwise distribute it. For instance, to report only some portion of the results or some of a group of studies that are undertaken to investigate a particular hypothesis, or to publish only studies that report a statistically significant finding. Studies with negative or inconclusive results about the same research question rarely see the press, thus the literature as a whole creates a false impression of repeated corroboration with no disconfirmation.

2. P-hacking. This is when researchers collect or select data or statistical analyses until nonsignificant results become significant, thereby inflating the false positive rate (by conducting analyses midway through experiments to decide whether to continue collecting data, recording many response variables and deciding which to report post hoc, deciding whether to include or drop outliers post hoc, excluding, combining, or splitting treatment groups post hoc, and stopping data exploration if an analysis yields a significant p-value).

3. Harking. This refers to presenting a post hoc hypothesis (i.e., one based on or informed by one's results) in one's research report as if it were, in fact, an a priori hypothesis.

4. Overselling. This is glossing over some known issues/problems with a study in its presentation.

<sup>5</sup> It is important to note here that the notion of a credit incentive involves that something else than knowledge or purely epistemic performance is rewarded, typically an outcome or product that may correlate (or not) with knowledge. For instance, in the current reward system of science the knowledge that A does not influence B, or the second-order knowledge that an experiment worked or did not work properly rarely deserves credit, while an accidental discovery by a researcher who lacks the associated knowledge deserves credit nonetheless.

<sup>6</sup> As the economic argument has it, an exclusive community of high-minded, disinterested inquirers could possibly stifle scientific productivity. A community of lawful good type of scientists may end up studying subjects that are deemed insignificant for the public or other scientists, be reluctant to share own research with others until they are confident enough in the outcomes, may be much less productive than others, or have a much lesser chance of making unexpected discoveries by following riskier methods. It is moreover difficult to steer or nudge their research strategies in directions desired by science policy makers by engineering incentives, since they do not respond to external rewards as easily.

<sup>7</sup> In the same vein, Fricker (2020) proposes a richer conception of an epistemic ethos as “the set of shared motivational dispositions and evaluative attitudes geared towards ultimate or intermediate epistemic goods.” De Ridder (2022) elaborates on a similar but more delimited idea in terms of a research group's ethos, as embodied in their practices: “are questions welcomed; are junior team members mentored; is there organizational support for learning and development; are there opportunities for creativity and out-of-the-box thinking; do team members (especially those in hierarchical relations) welcome feedback and criticism; do team members give each other credit; do group members take pride in being part of the group; do people experience the organization's overall goals as worth caring about; who are the group's role models; are work hours and compensation in proportion to the tasks and results that are expected; are successes celebrated; etc.”

<sup>8</sup> <https://www.thecut.com/2016/10/inside-psychologys-methodological-terror-debate.html>

<sup>9</sup> Presumably the effect of social organization of science interacts with the effect of exemplars in determining the normative force of one's epistemic values on the practices of others. In general, the values promoted by more senior scientists are more influential on the community than those endorsed by junior scientists. If a scientific field is organized around research teams, the structure of the research teams becomes an important factor. In hierarchical teams the values cherished by those in leading positions would have more normative force on the others. This makes it all the more important that senior scientists and research team leaders are epistemically virtuous. Ironically, the institutional ethos generated by the competitive external reward system of science favors strategic, career-focused individuals to those who are passionate for science, and those that make their way up to senior academic positions through all the selection mechanisms are often of the first kind. This gives the institutional ethos a rather conservative tendency, and makes it all the more difficult for younger scientists to effect a change in it.