

This is a draft of a paper to appear in a collection edited by Jonathan Matheson and Aaron Creller
Please do not quote without permission
Comments welcome
Last Update: 06 January 2025
Word Count: 8,864

The Science Contract: Scientific Inquiry, Public Trust in Science, and the Division of Zetetic Labor

GABRIELE CONTESSA

Department of Philosophy
Carleton University

gabriele.contessa@carleton.ca

ABSTRACT. What can we, as a society, legitimately expect from science? And what, if anything, can science legitimately expect from society? This paper argues that the relationship between science and society is governed by a science contract. I first introduce the notion of an expertise contract—a social contract that governs the relationship between experts and non-experts, bestows on experts certain fiduciary duties towards non-experts, and enables the division of epistemic labor in society. I then argue that the science contract cannot be simply identified with the expertise contract as it applies to scientific experts but requires at least two modifications in order to reflect the global nature of science and the role of scientists as researchers. I conclude with some remarks about public trust in science and a non-ideal theory of the scientific contract.

1. THE TRANSACTIONAL CONCEPTION OF THE SCIENCE CONTRACT AND ITS LIMITS

What can we, as a society, legitimately expect from science? And what, if anything, can science legitimately expect from society? This paper sketches a potential answer to these two questions according to which the relationship between science and society is governed by a science contract. At first, the idea of a (more or less tacit) agreement between science and society might not appear to be particularly original.¹ Over the last century, it has been popular to think of the relationship between science and society in roughly transactional terms. What we might call the *transactional conception* of the science contract sees the relationship between science and society as, essentially, a *quid pro quo*. Society provides science with the material resources required to produce scientific knowledge and, in return, science provides society with scientific knowledge valuable to society. While, to my knowledge the transactional conception has never been explicitly embraced (let alone defended), something along

¹ For an enlightening discussion of the notion that the relationship between science and society is governed by a social contract, see (Douglas and Branch 2024). While, for reasons of space, I won't be able to discuss Douglas and Branch's views here, it is worth mentioning that I do not see the notion of the science contract developed here as an alternative to the one embraced by Douglas and Branch but rather as largely complementary to it (as, for one thing, the two notions operate at different levels of abstraction).

these lines seems to underpin Vannevar Bush’s influential report *Science: The Endless Frontier*, which shaped American science policy for most of the twentieth century.²

The transactional conception, however, faces a number of problems. Let me briefly mention four here. The first problem is that it seems to imply that society can only legitimately expect something from science as long as society foots its bills. This seems to be particularly concerning at a time when scientific research is increasingly reliant on private funding and, especially, funding by corporations that, ultimately, aim to profit from that research.³

The second problem is that there is no guarantee that scientific research will automatically produce knowledge that is valuable to society just because it is publicly funded. After all, publicly funded science is funded by governments (or their agencies) and, arguably, governments (and their agencies) don’t always act in the best interest of the public. For example, much of the research funded by the US government in the twentieth century was intended to have military applications and it is at best unclear that this is the best interest of the American public.⁴

The third problem is that, insofar as governments do fund science, they do so at a local, regional, or national level, but scientists seem to have duties that extend beyond any of those borders. For example, it would seem to be at best parochial (and at worst callous) to maintain that the scientists who contributed to the development of the atomic bomb needed not to consider the effects of their research would have on its potential (non-American) victims because they only had duties towards the American government or the American public.

The fourth problem is that the transactional conception leaves society vulnerable to being taken advantage of. For example, as we have seen, society tries to fund scientific research that is likely to produce knowledge that is valuable to it. However, it is difficult for society to independently assess which research programs are genuinely promising. One potential problem is *scientific hype*—i.e., when a field or subfield misrepresents what it might be able to deliver, such as in the case of stem cell therapy or genetic medicine in an attempt to garner attention or funding.⁵

While these remarks are admittedly brief, they seem to give us reasons to believe that the transactional conception is inadequate (at least, on its own). However, I do not think that the most appropriate response to the apparent inadequacy of the transactional conception is to abandon the notion of a science contract altogether. Rather, the best response is to develop a better conception of the scientific contract. In this chapter, I aim to sketch a stronger conception, which I call the fiduciary conception.

Since that the road that will take us there is somewhat tortuous, let me briefly sketch the route here. In §2, I distinguish a number of different notions of the division of epistemic labor and discuss their relationships with one another. In §3, I argue against a transactional conception of the everyday relationship between ordinary people and individual experts. In §4, I introduce the idea that non-experts’ dependence on experts generates certain fiduciary duties on the part of experts towards non-experts. In §5, I offer a contractualist argument for the existence of an “expertise contract,” which underpins the fiduciary duties of experts and enables the efficient division of epistemic labor. In §6, I discuss (the fiduciary conception of) the science contract as an expansion of the expertise contract. In

² It is significant in this respect that a recent reprint of Bush’s report is accompanied by an essay by Rush D. Holt, the CEO Emeritus of the American Association for the Advancement of Science, that is entitled ‘The Science Bargain’ (Holt 2021)

³ See, e.g., (Conner 2020). Notably, this is even true of research that is carried out by universities (see, e.g., (Babina et al. 2020)).

⁴ See, e.g., (Conner 2020). The fact that much of that research eventually found civilian applications (see, e.g., (Mazzucato 2015)) is beside the point here, as that was not the intended purpose for which the US government funded it.

⁵ See, e.g., (Master and Resnik 2013) and (Intemann 2022).

§7, I connect this to the notion of public trust in science and, in §8, I apply the ideal theory developed in the previous section to our very non-ideal context.

2. GNOSTIC LABOR AND ZETETIC LABOR

Nowadays, it is not uncommon to talk about the division of epistemic labor.⁶ For our purposes, however, it is convenient to distinguish two forms of the division of epistemic labor. The first, which we might call *the division of gnostic labor*, is the one that has received most attention from social epistemologists so far. Gnostic labor is the labor that goes into acquiring or retaining (existing) knowledge. It is a form of labor because acquiring or retaining knowledge requires using one's limited resources (such as time or memory) and, as such, it has an opportunity cost—the resources that one devotes to acquiring or retaining one piece of knowledge are resources that cannot be devoted to other purposes (such as acquiring or retaining some other, unrelated piece of knowledge).

Contemporary societies rely extensively on the division of gnostic labor. While, in any given society, there are pieces of general knowledge that almost every adult is expected to have (e.g., the sort of stuff most people in our society are taught in grade school), much of our more specialized knowledge is only possessed by select groups of individuals who are experts in the relevant domains; the more specialized the knowledge, the smaller the group of individuals who possess it.⁷ When the rest of us (individually or collectively) need to access one of these more specialized bits of knowledge, we must rely (directly or indirectly) on the relevant experts.

The main reason why societies rely so extensively on the division of gnostic labor is, presumably, that it is efficient. The division of gnostic labor enables a society to collectively know much more than any of its individual members could possibly ever know. If I am not feeling well, it would be wildly inefficient for me to go to medical school in an attempt to acquire enough medical knowledge to proficiently diagnose my own condition and identify potential treatments for it. It is extraordinarily more efficient to outsource those tasks to a medical doctor, who already has acquired the relevant knowledge (or can refer me to a specialist who has more specialized knowledge). As this example illustrates, the division of gnostic labor is efficient because it allows us to outsource our gnostic labor to others. The only alternative (a radical gnostic self-reliance) would not only be wildly inefficient (as it unnecessarily duplicates the gnostic labor of those who already have acquired the requisite knowledge) but beyond the capacity of any human (given our limited resources and the number and variety of issues each of us is likely to face in our lifetime).

The second form of the division of epistemic labor, which we might call *the division of zetetic labor*, on the other hand has received more attention from philosophers of science than from social epistemologists.⁸ Zetetic labor is the labor that goes into producing (new) knowledge (as opposed to acquiring (existing) knowledge). Zetetic labor is a form of labor because producing knowledge also requires us to use our limited resources (such as time or material resources) and, as such, it also has an opportunity cost—the resources devoted investigating one question cannot be devoted to any other purposes (including that of investigating any unrelated questions).

Over the last few centuries, societies have come to increasingly rely on the division of zetetic labor. In particular, the division of zetetic labor is one of the factors underlying the remarkable expansion of

⁶ See, e.g., (Hardwig 1985) and (Goldberg 2011).

⁷ See, e.g., (Millgram 2015).

⁸ See, e.g., (Longino 1990), (Hardwig 1991), and (Kitcher 1993).

our scientific knowledge over that period of time.⁹ Scientific research is a collective effort, and this collective effort is based on a massive division of zetetic labor.¹⁰ This is not only true in the sense that the global community now outsources much of its zetetic labor to professional scientific researchers. It is also true in the sense that scientific researchers divide their zetetic labor among themselves. At the most general level, scientific researchers divide their zetetic labor in the sense that different scientific researchers operating in different disciplines, fields, or subfields pursue different research questions about different aspects of our (natural or social) world from one another; scientific community, different scientists within the same (sub)field explore the same (or similar) research questions using different assumptions and methodologies; at the level of research teams, different members often are responsible for different tasks and have different (and often complementary) competences and skills.

The main reason why we have come to rely so extensively on the division of zetetic labor is also, presumably, that it is efficient. The division of zetetic labor is what has enabled scientific knowledge to grow at the pace it has.¹¹ In fact, it is arguably because of the massive amount of scientific research that is carried out and the knowledge that it has produced that we must rely on a division of gnostic labor in the first place. If the idea of going to medical school to diagnose my own medical condition seemed wildly implausible and extremely inefficient, the idea of me doing my own original research on my own medical condition seems to be even more ludicrous. Where could any one of us ever find the resources to start their own lab or run their own clinical trials? But the same applies to my family doctor (or the specialist to whom she referred me). They too must rely on a body of knowledge that is the product of the zetetic labor of a large number of scientific researchers across the world.

Let me conclude this section with two relatively trivial points that go often unmentioned. While it is common for social epistemologist to employ the analogy between the division of gnostic labor and the division of labor (*simpliciter*) (see, e.g., (Matheson 2023)), the relationship between the two forms of the division of epistemic labor I discussed here and the general division of labor is more than a mere analogy. This is because the division of gnostic labor seems to be (in part) a consequence of the general division of labor, and the division of zetetic labor seems to be a form of the general division of labor. Let me briefly explain each point. The first point is often associated with Adam Smith's example of the pin factory (*The Wealth of Nations* I.I.3) (Smith 1776, sec. I.I.3)—the general division of labor tends to both generate (and require) specialized knowledge. Becoming an electrician, a mechanic, a surgeon, a lawyer (or, for that matter, a scientist) is a learning process and, as we have seen, learning is a kind of gnostic labor. The second point is that, now that scientific research is fully professionalized, being a scientific researcher has become a job like many others.

3. BEYOND CONTRACTUAL DUTIES

The two points I made at the end of the previous section help my argument because it is easier to conceive of the science contract (as I propose to do here) as a special case of a more general contract that I will call *the expertise contract*. Let me start from a simple example in which a layperson relies on a (non-scientific) expert for a given task.¹² Suppose that Aman's car makes a strange clicking sound

⁹ While this is not the only form of the division of zetetic labor on which modern societies rely (e.g., we outsource much of our zetetic labor in other areas of inquiry to journalists or law enforcement agencies), it is the only one I will focus on here.

¹⁰ See, e.g., (Hull 1988), (Longino 1990), and (Kitcher 1993).

¹¹ For a rough estimate, see, e.g., (Fahrback 2011).

¹² For an overview of different accounts of expertise, see (Watson 2021). Which of the many accounts of expertise is most compatible with the fiduciary conception of the social contract will become clearer in §7 below.

sometimes and, since he does not know much about cars, he decides to take it to a mechanic. The mechanic, Beth, tells him that the left CV axle is damaged, and, if it is not replaced, it might fail completely or damage the transmission. While Aman is not sure what a CV axle is or what it does, he is nevertheless alarmed and, therefore, he authorizes Beth to carry out the repair. Once Beth is done, she returns his car to him, and he pays her for the job. Much to Aman's relief, the clicking sound is now gone.

When the relationship between Aman and Beth is described this way, one might feel inclined to adopt what we might call a *transactional conception* of their relationship (which, as the name suggests, is analogous to the transactional view of the science contract discussed in §1). According to the transactional conception, the relationship between Aman and Beth is purely transactional and it is defined by the legal contract between them—i.e., the contract according to which Beth agrees to replace the CV axle and Aman agrees to pay her a given amount of money in return.

However, while it is true that, when non-experts rely on experts, their relationship often has a transactional component (as in the case of Aman and Beth), it is important not to mistake the whole relationship for one of its (inessential) components. The easiest way to see this point is to consider that, if the relationship between experts and non-experts was purely transactional, then the only duties each party would have toward the other would be their contractual duties (i.e., the duties that the legal contract between them bestows on each of them).¹³ However, this does not seem to be true. I want to argue that, whenever non-experts depend on experts, the experts acquire certain duties towards the non-experts (irrespective of any contractual duties towards them). Here, I offer two brief arguments for the existence of such non-contractual duties.

The first argument is that experts (such as plumbers, lawyers, or doctors) have certain duties towards non-experts who depend on them even when there isn't (and there will never be) any legal contract between them. For example, mechanics, lawyers, and doctors sometimes might informally give information or advice to family, friends, or acquaintances. Or, to pick another example, many experts provide potential customers with free estimates. Suppose, for example, that the noise Aman heard is not due to a damaged CV axle but to his tire rubbing against a loose piece of trim and that Beth is simply trying to take advantage of Aman. Suppose, however, that on this version of the scenario, Aman is suspicious of Beth's advice and declines to have the CV axle replaced. Since there was no legal contract between them (and there will never be one), then Beth cannot have violated any contractual duties towards Aman; and yet, intuitively, it would still seem that she did something wrong. This seems to suggest that Beth has duties towards Aman that go beyond any strictly contractual duties she might have towards him.¹⁴

The second argument is that, even in cases in which there is a legal contract between the two parties, its content is often shaped by the information or the advice that the expert has given to the layperson. For example, suppose that, on this scenario, Beth has managed to convince Aman to replace his CV axle (even if there is nothing wrong with it). In this scenario, it would seem that all it would take for Beth to discharge her contractual duties towards Aman would be to perform the unnecessary replacement adequately. After all, that's all that the agreement between her and Aman requires her to do! However, it seems that Beth should not have advised Aman to replace the CV axle

¹³ Here, I understand contractual duties in a strict sense, as the duties explicitly bestowed by the contract on each party. Many legal systems, however, seem to presuppose a more expansive conception of contractual duties (by including, e.g., notions such as those of an unconscionable contracts (see, e.g., (Wertheimer 1992))). However, I suspect that these more expansive notion of contractual duties tacitly presupposes the existence of fiduciary duties between the parties of the sort for which I argue below.

¹⁴ Of course, supporters of the transactional conception could adopt a broader notion of what a transaction is. I suspect that doing so could lead in a similar direction to the fiduciary conception (see, e.g., (Hardin 2002)).

in the first place. Once again, Beth would seem to have duties towards Aman beyond her contractual duties towards him (which, we can assume, she discharged adequately).

Before concluding this section, let me consider a possible reply to these arguments. The reply is that even those who are sympathetic to the transactional conception typically acknowledge that people have duties that extend beyond their contractual duties. If these duties include a (defeasible) general moral duty not to lie to (or mislead) others, then this would take care of both of the arguments above. However, even admitting that this solution works in the two scenarios discussed above, it does not provide the transactional conception with a general solution to the problem. To see why, consider, for examples, two new scenarios. In the first scenario, Carla, a 40-year-old woman with no family history of breast cancer, is considering whether she should start undergoing regular breast cancer screenings. While US guidelines used to recommend that women like Carla undergo biannual screening, they now recommend that they discuss their situation with their doctor. This is partly due to the fact that there is now less confidence in the effectiveness of routine screenings for those patients and a better appreciation of potential downsides, such as overdiagnosis and overtreatment.¹⁵ Carla's doctor, Dr Dodd, however, advises her to undergo the screening without taking Carla's own preferences into consideration or discussing the benefits and risks of either option with her because he is more concerned about his own potential legal liability than Carla's best interest.

In the second scenario, another patient of Dr Dodd, Erica, shares with him her concerns about childhood vaccines for her young child. Dr Dodd, who is overworked and behind schedule, makes short shrift of her concerns and simply tells her that there is absolutely no evidence that vaccines are unsafe. His brisk dismissal, however, not only does not persuade Laura, but it makes her feel like he does not really have the best interest of her child at heart.

While Dr Dodd's dereliction of his duties towards his patients might not be as flagrant as the dereliction of Beth's duties towards Aman in the scenarios above, it does not seem to be any less serious. In fact, if anything, it seems to be more serious, as it might have more harmful consequences for his patients than Beth's dereliction of duty has for Aman. But what sort of duty would have Dr Dodd violated? After all, he does not seem to have violated any general moral duty (along the lines of the general moral duty not to lie to (or otherwise mislead) others). I claim that Dr Dodd has violated his fiduciary duties towards Carla and Erica, which is not a general moral duty but a specific moral duty that arises from the fact that his patients rely on him as an expert.¹⁶ In the next section, I briefly outline the account of fiduciary duties I adopt here and, in the following section, I offer an argument for thinking that experts have fiduciary duties towards the non-experts who rely on them.

4. FIDUCIARY DUTIES

In this section, I outline the general account of fiduciary duties that I adopt in this paper. According to it, what distinguishes fiduciary relationships (i.e., relationships in which it would be appropriate for one party to trust or mistrust the other) from other relationships of reliance is that, in a fiduciary relationships (unlike in other relationships of reliance), the trustor's reliance on the trustee gives rise to certain fiduciary duties towards the trustor on the part of the trustee. But where do these fiduciary duties come from? In what we might call consent-based fiduciary relationships, the fiduciary duties

¹⁵ For a discussion of these issues that is both philosophically sophisticated and empirically informed, see (Plutynski 2012) and (Plutynski 2017).

¹⁶ Of course, this is not the only possible answer. An alternative answer is that Dr Dodd has a duty of care towards his patients. However, at least in the case of experts, it is not clear to me that these are two genuinely different answers (as opposed to terminological variants of the same answer).

might arise simply from the fact that the trustee has consented to the trustor's relying on them.¹⁷ If you agreed to pick me up from the airport, your having consented to picking me up from the airport gives rise to certain fiduciary duties towards me on your part (such as the duty to pick me up and the duty to let me know if extraordinary circumstances prevent you from doing so).

However, not all fiduciary relationships seem to be consent-based. If you come across a crying child who is lost in a park, you seem to have certain fiduciary duties towards the child simply in virtue of the fact that the child now depends on you. It would seem to be impermissible for you to just walk away claiming that you have never agreed to helping the child. In a case such as this, the fiduciary duties are role-based (rather than consent-based)—they arise from your role as an adult in relation to a lost child who has come to depend on you.¹⁸

According to the account I adopt here, the relationship between experts and the non-experts who depend on them is also a role-based fiduciary relationship. The experts have fiduciary duties towards their trustors in virtue of the fact that their trustors have come to depend on them as experts. Let me call this sort of role-based fiduciary relationship an *eidemonic relationship* (from the Greek word for 'expert')¹⁹ and the fiduciary duties that arise from them *eidemonic duties*. In the next section, I offer a contractarian argument for the existence of this sort of fiduciary duties.

5. THE EXPERTISE CONTRACT

So far, I have claimed that experts have special duties towards those who depend on them as experts. However, I am yet to offer any substantial argument in support of this claim. In this section, I offer such an argument. More specifically, I offer a contractarian argument for the conclusion that eidemonic duties stem from what I call *the expertise contract*—an hypothetical agreement between all members of society according to which experts have eidemonic duties towards those who depend on them as experts (irrespective of whether they have any other duties towards them), which might include the duty not to misinform their trustors, the duty to give them as much information as the trustors need in a manner that is as much as possible comprehensible to them and useful for them, the duty to give advice that is in the best interest of their trustors, and the duty to inform the trustor(s) of any threats that might be unaware of (in view of their lack of expertise).

I should note immediately that my goal here is not to offer the only possible argument for the existence of eidemonic duties but to offer an argument for eidemonic duties that relies on as weak a set of assumptions as possible. In particular, the contractarian argument does not presuppose the existence of any general moral duties from which eidemonic duties can be derived (although it does necessarily deny that eidemonic duties might be also derivable from a sufficiently robust set of general moral duties (such as the aforementioned moral duty not to lie to (or mislead) others)); nor does it presuppose that anyone is motivated by anything other than self-interest (although it does not necessarily deny that real people have a much more complex and rich set of motivations). What the

¹⁷ There are many ways for a trustee to consent to the trustor's relying on them, which include the trustee's promising to do something for the trustor, the trustee's accepting a request by the trustor to do something for the trustor, the trustee's offering to do something for the trustor, or the trustee agreeing to do something for the trustor. For a possible account of how these are related to one another, see (de Kenessey 2023).

¹⁸ While, for lack of space, I cannot articulate the distinction properly here, I take dependence to be distinct from (mere) reliance in that reliance can be sometimes impermissible, but dependence is never impermissible. For example, the person who takes advantage of the generosity of their friends out of stinginess is (impermissibly) relying on them but does not depend on them.

¹⁹ I would like to thank Nikolaos Drosos for suggesting this label to me.

contractarian argument tries to establish is that not even a purely self-interested expert could reasonably reject the expertise contract

The first step of the contractarian argument is that, if, as the argument assumes, experts are purely self-interested, then, whenever their self-interest diverges from that of their trustors (as it is often the case in the real world),²⁰ they will try to take advantage of their trustors whenever they believe that they can get away with it.

The second step is that, in eidemonic relationships, trustors are often more vulnerable to betrayal than in most ordinary fiduciary relationships. If, for example, you forget to pick me up from the airport, it is going to be very obvious to me that you have failed to meet your fiduciary duties towards me. However, an expert's failure to meet their duties is usually not as obvious to their non-expert trustors. We have already seen a few of examples of this above. However, the point can be made more simply by noting that, in eidemonic relationships, the trustor is rarely in a good epistemic position to accurately assess the expert's performance due to the epistemic asymmetry between them and, therefore, the trustor is often unable to determine whether or not expert has met their eidemonic duties towards them.²¹

The third step is that, as a matter of fact, no one (not even the experts) can avoid entering in eidemonic relationships as a trustor. This is simply a consequence of the fact that no one can acquire enough expertise in each domain in which they will need specialized knowledge or expert advice throughout their lives, and, therefore, even people who are experts in one domain have to frequently rely on those who have expertise in other domains. The mechanic and the doctor might each be expert in their own respective domains but each needs to rely on the other when they need to access the specialized knowledge they do not possess.

These three steps allow us to reach the intermediate conclusion of the contractarian argument, which is that everyone (including the experts) has self-interested reasons to accept the expertise contract. Let me illustrate this point with an example. Suppose that Beth is purely self-interested and, therefore, has no non-self-interested reasons to accept the expertise contract. In fact, she arguably has self-interested reasons to reject it. After all, rejecting the expertise contract would allow her to claim that she has no other duty towards Aman (or any of her other customers) than her contractual duties. However, Beth also has self-interested reasons to accept the expertise contract because, in many circumstances, she is the one who finds herself in the trustor role in an eidemonic relationship. If she is not feeling well, if she is trying to invest her money, or if a disgruntled customer sues her, then she will be the one who will have to turn to the relevant experts and, in those cases, she will have self-interested reasons to accept the expertise contract because she does not want to be taken advantage by these experts.

However, this is not enough. In order for the contractarian argument to be successful, it must also show that Beth's self-interested reasons for accepting the expertise contract outweigh her self-interested reasons for rejecting it. The fourth and final step of the argument is meant to fill this gap. What would a society of purely self-interested individuals who reject the expertise contract look like? In this hypothetical epistemic "state of nature," it would be permissible for Beth to take full advantage of her customers in all sorts of ways (as far as he can get away with it) and, since she has self-interested

²⁰ See, e.g., (Guerrero 2016).

²¹ See, e.g., (Nguyen 2020). It is worth noting that this is true even of those domains that are deceptively opaque to non-experts. For example, when our symptoms disappear after starting the new treatment recommended by a healthcare practitioner, we might be inclined to conclude that the treatment was effective. However, if people could reliably assess the effectiveness of a treatment (or the trustworthiness of a healthcare practitioner) from anecdotal evidence (such as the first-hand evidence we gather about our condition), there would be no need to run costly randomized control trials or complex meta-analyses. For all we know, the relief of our symptoms could be due to a coincidence or to the placebo effect. It is worth noting that

reasons to do so (e.g., she can earn more money or work less by taking advantage of her customers) and her customers will often be unaware of her deceptions, she might have reasons to do it most of the time. Moreover, her customers do not have any real alternative. In this epistemic “state of nature,” other mechanics are no less deceptive than Beth and, as we have seen, radical epistemic self-reliance is impracticable.

So, at first, it might seem that Beth’s self-interested reasons for rejecting the expertise contract outweigh her reasons for accepting it, as it might seem to be a world in which, thanks to her deceptions, Beth has more free time and more money than in a world in which she’s bound to the expertise contract. However, even if this is the case, there seems to be very little point for Beth to have more free time or more money if she does not do anything with her extra resources. For example, she might want to use some of her money to buy a fancy house and save the rest for retirement. However, in the absence of an expertise contract, it would also be permissible for her real estate agent and her investment advisor to take advantage of her and, since they have self-interested reasons to do so and are likely to get away with it, then it is likely that they will. Beth, then, is likely to end up with a crumbling house and no savings, which would ultimately defeat the point of having all the extra money that she would make by taking advantage of her own customers. In other words, while Beth might be perfectly happy to live in a world in which she is the only expert exempted from the expertise contract, not even she would be happy to live in a world in which no expert abides by the expertise contract because, in that world, her only two options would be to rely on untrustworthy experts or on her own limited expertise in all fields. The conclusion is therefore that not even Beth can reasonably reject the expertise contract because even for someone like her the disadvantages of rejecting the contract outweigh the potential advantages of doing so.

The contractarian argument can be reformulated more abstractly and concisely as follows. All of the members of any society that relies as heavily on the division of epistemic labor as most contemporary societies do have self-interested reasons to accept the expertise contract and those reasons outweigh any self-interested reasons that even the experts themselves might have to reject it because, unless all members of society can assume that (something akin to) the expertise contract governs the relationship between experts and non-experts in a society, society would be unable to divide its epistemic labor efficiently. Its members would live in something akin to an epistemic “state of nature.”

Let me conclude this section with two brief (and somewhat related) remarks. First, the expertise contract needs to be carefully distinguished from any of the specific legal contracts that often bind experts to non-experts. According to the contractarian account, the expertise contract should be understood as a social contract in the spirit of the social contract tradition (i.e., as roughly an implicit and hypothetical contract among all members of society) rather than as a legal contract between any specific parties. Crucially, the contractarian argument maintains that the expertise contract holds even in the absence of any legal contract between an expert and a non-expert and that this is why experts often have duties towards non-experts even under those circumstances.

Second, the one between the expertise contract and the social contract is more than a mere analogy. One of the main functions attributed to the social contract is to support cooperation among the members of society. If the function of the expertise contract is to enable the division of epistemic labor in society, then it can be understood as a section of (or an addendum to) a broader social contract that applies to societies that, like ours, rely extensively on the division of epistemic labor. More specifically, the expertise contract can be seen as the part of the social contract that enables *epistemic* cooperation among the members of the relevant society. Without an expertise contract, a society would be condemned to live in an epistemic “state of nature,” where specialized knowledge is limited, hard to acquire, and even harder to produce.

6. THE SCIENCE CONTRACT: A FIDUCIARY CONCEPTION

We are now in a better position to see an alternative conception of the science contract, which we might call the *fiduciary conception*. To a first approximation, the fiduciary conception of the science contract maintains that the science contract is simply the expertise contract insofar as it applies to a proper subset of all experts—i.e., scientific experts. On this preliminary understanding, scientific experts have the same sorts of eidemonic duties as all other experts. However, this preliminary understanding has a number of limitations. Here, I focus on two of them.

The first limitation is that, while the eidemonic duties of scientists in their capacity as scientific experts might coincide with those of all experts, it would seem that, in their capacity as scientific researchers, scientists also have a distinct set of fiduciary duties. In fact, we can think of the science contract as an expanded version of the expertise contract that, in addition to the eidemonic duties of scientists as experts, also includes their eidemonic duties as researchers. But what are these extra eidemonic duties? For reasons of space, I focus on one example here. In the case of biomedical research, it is relatively uncontroversial who the biomedical researchers' primary trustors are—they are the patients who are the ultimate users of any knowledge produced by their research because it is the patients who ultimately depend on their research. According to the fiduciary conception, this means that, when researchers do research that is relevant to patients affected by a given condition, they have certain fiduciary duties towards those patients. In particular, when the researchers are facing any decisions in the course of their research, they have a fiduciary duty to act in the best interest of those patients and that these fiduciary duties override any contractual duties they might have towards those who sponsor the research. So, for example, if the pharmaceutical company that sponsors research on a new drug expects the researchers to cherry-pick the studies that appear to show that the drug is effective, the researchers have an eidemonic duty not to do so (irrespective of any contractual duty they might have towards the company). More generally, Heather Douglas has persuasively argued that scientists have a the duty to take into account the potential practical consequences of their decisions in their decision-making process (see, in particular, (Douglas 2009)). While Douglas seems to think of these duties as deriving from general moral duties of individual scientists (such as a duty not to be negligent), on the fiduciary conceptions, these duties should be understood as fiduciary duties deriving from the science contract.

The second limitation of the preliminary understanding of the fiduciary conception is that the traditional understanding of the social contract is that the social contract applies at a relatively local level (e.g., the national level) and this might even be true of the expertise contract. However, it does not seem to be true of the science contract. Today, science is more than ever a global enterprise and, in order to reflect that, the science contract must be understood as applying at a global level. So understood, the science contract can overcome many of the problems that were faced by the transactional conception. In particular, researchers have fiduciary duties towards the global public irrespective of who funds their research (whether it is governments or corporations). This allows the fiduciary conception to solve in one fell swoop all of the problems faced by the transactional conception.

7. THE RECIPROCAL DUTIES OF NON-EXPERTS

Contracts are usually understood as requiring some reciprocity between the parties. Typically, all parties to a contract acquire certain duties towards one another in virtue of their mutual agreement. This seems to be true of the social contract as well. However, as I have described them so far, the

eidemonic and the science contracts do not seem to require anything of non-experts. This might lead one to wonder whether it is correct to describe them as contracts in the first place. While I suspect that this worry relies on a number of questionable premises, I think that the fiduciary conception has a way to address it directly, which is to claim that the science contract does require something of the rest of society—it requires that, if the experts hold up their side of the science/expertise contract by fulfilling their fiduciary duties towards non-experts, then the non-experts should reciprocate by trusting (and epistemically deferring to) the experts on the relevant matters.

But why would this be the case? I think there are a number of possible answers at different levels of abstraction. Let me briefly mention two here. The first answer is that it would seem appropriate for a trustworthy trustee to resent the (inexcusable) mistrust of their trustors.²² For example, suppose that Fritz is a trustworthy lawyer who is working *pro bono* to overturn Gail's wrongful conviction. At the beginning, he might excuse Gail's mistrust of him, as, due to her previous experiences, Gail's generalized mistrust of lawyers might not seem to be unwarranted. However, if Gail keeps being mistrustful of Fritz despite all the hard work that he has put into her case and all the evidence she has that he is trustworthy, Fritz might eventually come to (appropriately) resent her mistrust. But why should Fritz expect Gail to trust him? One possible answer is that fiduciary relationships give rise to a reciprocal duty in the trustors to trust the trustee if the trustee upholds their fiduciary duties, and, since Fritz has been fully trustworthy towards Gail, she should reciprocate his trustworthiness with her trust (despite her past negative experiences with lawyers).

The second answer is more abstract. The point of the expertise/science contracts is to enable an efficient division of epistemic labor. However, since an unwarranted mistrust of experts would undermine the efficient division of epistemic labor, the expertise/science contract can only perform its function if it also requires that non-experts trust (and epistemically defer to) the experts on the relevant matters insofar as the experts are trustworthy.

This approach allows us to reconceive the meaning of the oft-used expression 'epistemic authorities' to refer to the experts in a certain domain. The science/expertise contract framework allows us to understand the relationship between epistemic and practical authority (such as political authority) better than other conceptions available. For example, according to the prominent view of epistemic authority developed by Linda Zagzebski (2012), what characterizes an epistemic authority is that, if an epistemic authority believes something (within their domain of authority), this provides everyone else with a preemptive reason to believe it as well. However, while such an account might work in the case of practical authorities (see, in particular, (Raz 1986)), it does not seem to work for the case of epistemic authorities. For one thing, when epistemic authorities disagree with one another, we are left under an impossible obligation—i.e., that of having contradictory beliefs. This problem does not typically arise for practical authorities, as distinct practical authorities do not usually have overlapping jurisdictions. The expertise/science contract framework, however, provides us with a different approach to understanding the duty to trust epistemic authorities. Epistemic authorities arise from the division of epistemic labor and any duty we have to trust (or defer to) them on the relevant matters is not an epistemic duty *per se* but a social duty that is conditional on the authorities fulfilling their social role and the (fiduciary) duties towards us that come with it.²³

²² For a much more in-depth discussion of what is owed to the testifier in cases of testimony, see (Goldberg 2022).

²³ As these remarks suggest, the conception of expertise that is closest to the one adopted by the fiduciary conception is the one developed by Christian Quast (2018).

8. THE SCIENCE CONTRACT IN A NON-IDEAL WORLD

Much of this paper has been concerned with ideal theory (in the Rawlsian sense of a normative theory that assumes full compliance).²⁴ However, it would be wildly unrealistic to assume full compliance with the principles of the science contract in the real world. Over the last couple of decades, much attention has been devoted to the non-compliance on the part of the general public. A growing number of academics and commentators have lamented the existence of pockets of unwarrantedly low level of trust within the general public.²⁵ The two standard examples of the general public's failure to live up to its side of the science contract are the discrepancy of opinions between scientific experts and the general public with respect to the issues of the reality of anthropogenic climate change and that of the safety of childhood vaccines.

However, once we start thinking of these issues along the lines of the fiduciary conception, it is harder to see how this assessment of the situation can be plausibly defended. After all, its proponents surely acknowledge that we are under no duty to trust the lawyers or the mechanics on whose services we have to rely. So, why should the situation be any different in the case of scientific experts and scientific researchers?

As far as I can see, the proponents of this view are likely to answer that in the case of scientists we have no reason to assume that their interests diverge from ours as often as in the case of mechanics and lawyers. However, on closer scrutiny, it is unclear that this is true. In fact, I already mentioned a couple of reasons to doubt that it is. It does not seem to be true in the case of scientific hype (briefly discussed in §1), nor in the case of research that is beholden to its corporate (or military) sponsors,²⁶ nor in the case of social groups that have been (and often still are) subjected to discriminatory or harmful research practices (such as women or racial minorities).²⁷

While this list could be extended indefinitely, the problem seems to run even deeper, as the misalignment often stems from science's internal system incentives (rather than on any undue influences of "external" factors (such as external funding sources or social prejudices) on scientific research).²⁸ While scientific researchers might have all sorts of purely epistemic or even altruistic motives, they also have an interest in furthering their own career. After all, as I have mentioned earlier, being a researcher is a job like many others and publications (and citations) are the currency of the scientific profession. Moreover, researchers nowadays operate in what is often described as a "publish or perish" environment, which means that, everything else being equal, those who publish more are more likely to secure a larger share of valuable resources (such as more citations, more funding, better jobs, better graduate students or post-docs, etc.) than researchers who publish less. While, in theory, this system of incentives could be used to harness the professional interests of scientists in the service of science (or of society), in practice, it also tends to generate perverse incentives for the researchers, which contribute to the misalignment between their interests and those of the general public. One of the aspects of this problem that has received most attention lately is the so-called replication crisis, which seems to be largely a result of the incentive to publish at all costs and which underlies phenomena such as publication bias, *p*-hacking, and even scientific fraud.²⁹

²⁴ See, e.g., (Rawls 1971, 216). For a helpful discussion of how that sense of 'ideal theory' relates to other senses, see (McKenna 2023, chap. 2).

²⁵ See, e.g., (Braun 1999), (Gauchat 2012), (Gawande 2016), and (Dobson 2022).

²⁶ For an overview of these issues, see, e.g., (Conner 2020).

²⁷ See, e.g., (Washington 2007) or (Cleghorn 2021)

²⁸ See, e.g., (McBrayer 2024).

²⁹ See, e.g., (Bright 2017) and (Romero 2019).

Elsewhere (Contessa 2023), I have argued that the prevalent individualistic approach to public trust in science is misguided in its focus on trying to persuade individuals to trust science more and that, instead, we should embrace a social approach that focuses on trying to improve our dysfunctional socio-epistemic infrastructure—i.e., the system of institutions, norms, practices, and procedures that promotes the production, transmission, reception, and application of information and prevent the and the spread of misinformation. I hope that seeing these issues through the lens of (the fiduciary conception of) the science contract might help us see why this is the case. Once we think of the division of zetetic labor as a form of the division of labor, it becomes evident that some mistrust might not be entirely unwarranted. For example, rethinking the way in which scientific research is currently funded, organized, and regulated might go a long way towards bringing the interests of researchers into better alignment with those of the rest of society, which would improve the trustworthiness of science, and, in turn, might contribute to fostering public trust in science.

ACKNOWLEDGEMENTS

A draft of this paper was presented at the Inquiry Conference at the University of North Florida. I am grateful to the audience for their insightful feedback. I would also like to thank Aaron Creller and Jonathan Matheson for their helpful comments on an earlier draft, as well as T.Y. Branch and Heather Douglas for engaging discussions on the topic of this paper. This paper is based on research supported by the Social Sciences and Humanities Research Council.

REFERENCES

- Babina, Tania, Alex Xi He, Sabrina T. Howell, Elisabeth Ruth Perlman, and Joseph Staudt. 2020. “The Color of Money: Federal vs. Industry Funding of University Research.” Working Paper. Working Paper Series. National Bureau of Economic Research. <https://doi.org/10.3386/w28160>.
- Braun, Richard. 1999. “The Public’s Growing Distrust of Science?” *Nature Biotechnology* 17 (5): 14–14. <https://doi.org/10.1038/70359>.
- Bright, Liam Kofi. 2017. “On Fraud.” *Philosophical Studies* 174 (2): 291–310. <https://doi.org/10.1007/s11098-016-0682-7>.
- Cleghorn, Elinor. 2021. *Unwell Women: Misdiagnosis and Myth in a Man-Made World*. Dutton.
- Conner, Clifford D. 2020. *The Tragedy of American Science: From Truman to Trump*. Chicago: Haymarket Books.
- Contessa, Gabriele. 2023. “It Takes a Village to Trust Science: Towards a (Thoroughly) Social Approach to Public Trust in Science.” *Erkenntnis* 88 (7): 2941–66. <https://doi.org/10.1007/s10670-021-00485-8>.
- Dobson, Geoffrey P. 2022. “Wired to Doubt: Why People Fear Vaccines and Climate Change and Mistrust Science.” *Frontiers in Medicine* 8. <https://www.frontiersin.org/articles/10.3389/fmed.2021.809395>.
- Douglas, Heather. 2009. *Science, Policy, and the Value-Free Ideal*. Pittsburgh, PA: University of Pittsburgh Press.
- Douglas, Heather, and T. Y. Branch. 2024. “The Social Contract for Science and the Value-Free Ideal.” *Synthese* 203 (2): 1–19. <https://doi.org/10.1007/s11229-023-04477-9>.

- Fahrbach, Ludwig. 2011. "How the Growth of Science Ends Theory Change." *Synthese* 180 (2): 139–55. <https://doi.org/10.1007/s11229-009-9602-0>.
- Gauchat, Gordon. 2012. "Politicization of Science in the Public Sphere: A Study of Public Trust in the United States, 1974 to 2010." *American Sociological Review* 77 (2): 167–87. <https://doi.org/10.1177/0003122412438225>.
- Gawande, Atul. 2016. "The Mistrust of Science." *The New Yorker*, June 10, 2016. <https://www.newyorker.com/news/news-desk/the-mistrust-of-science>.
- Goldberg, Sanford C. 2011. "The Division of Epistemic Labor." *Episteme* 8 (1): 112–25. <https://doi.org/10.3366/epi.2011.0010>.
- . 2022. "What Is a Speaker Owed?" *Philosophy and Public Affairs* 50 (3): 375–407. <https://doi.org/10.1111/papa.12219>.
- Guerrero, Alexander A. 2016. "Living with Ignorance in a World of Experts." In *Perspectives on Ignorance from Moral and Social Philosophy*, edited by Rik Peels, 156–85. New York, NY: Routledge. <https://doi.org/10.4324/9781315671246-15>.
- Hardin, Russell. 2002. *Trust and Trustworthiness*. The Russell Sage Foundation Series on Trust, v. 4. New York: Russell Sage Foundation.
- Hardwig, John. 1985. "Epistemic Dependence." *Journal of Philosophy* 82 (7): 335–49.
- . 1991. "The Role of Trust in Knowledge." *The Journal of Philosophy* 88 (12): 693–708. <https://doi.org/10.2307/2027007>.
- Holt, Rush D. 2021. "The Science Bargain." In *Science, the Endless Frontier*, by Vannevar Bush. Princeton, NJ: Princeton University Press. <https://press.princeton.edu/books/hardcover/9780691186627/science-the-endless-frontier>.
- Hull, David L. 1988. *Science as a Process: An Evolutionary Account of the Social and Conceptual Development of Science*. Chicago, IL: University of Chicago Press.
- Intemann, Kristen. 2022. "Understanding the Problem of 'Hype': Exaggeration, Values, and Trust in Science." *Canadian Journal of Philosophy* 52 (3): 279–94. <https://doi.org/10.1017/can.2020.45>.
- Kenessey, Brendan de. 2023. "Promises, Offers, Requests, Agreements." *Ergo an Open Access Journal of Philosophy* 9 (0). <https://doi.org/10.3998/ergo.2623>.
- Kitcher, Philip. 1993. *The Advancement of Science: Science without Legend, Objectivity without Illusions*. Oxford: Oxford University Press.
- Longino, Helen E. 1990. *Science as Social Knowledge: Values and Objectivity in Scientific Inquiry*. Princeton University Press.
- Master, Zubin, and David B. Resnik. 2013. "Hype and Public Trust in Science." *Science and Engineering Ethics* 19 (2): 321–35. <https://doi.org/10.1007/s11948-011-9327-6>.
- Matheson, Jonathan. 2023. *Why It's OK Not to Think for Yourself*. Routledge.
- Mazzucato, Mariana. 2015. *The Entrepreneurial State: Debunking Public vs. Private Sector Myths*. Revised Edition. New York, NY: Public Affairs.
- McBrayer, Justin P. 2024. "Caveat Auditor: Epistemic Trust and Conflicts of Interest." *Social Epistemology* 38 (3): 290–301. <https://doi.org/10.1080/02691728.2022.2078250>.
- McKenna, Dr Robin. 2023. *Non-Ideal Epistemology*. Oxford: Oxford University Press.
- Millgram, Elijha. 2015. *The Great Endarkenment: Philosophy for an Age of Hyperspecialization*. Oxford: Oxford University Press.
- Nguyen, C. Thi. 2020. "Cognitive Islands and Runaway Echo Chambers: Problems for Epistemic Dependence on Experts." *Synthese* 197 (7): 2803–21. <https://doi.org/10.1007/s11229-018-1692-0>.
- Plutynski, Anya. 2012. "Ethical and Scientific Issues in Cancer Screening and Prevention." *Journal of Medicine and Philosophy* 37 (3): 310–23. <https://doi.org/10.1093/jmp/jhs017>.

- . 2017. “Safe or Sorry? Cancer Screening and Inductive Risk.” In *Exploring Inductive Risk: Case Studies of Values in Science*, edited by Kevin C. Elliott and Ted Richards, 149–70. Oxford: Oxford University Press.
- Quast, Christian. 2018. “Expertise: A Practical Explication.” *Topoi* 37 (1): 11–27. <https://doi.org/10.1007/s11245-016-9411-2>.
- Rawls, John. 1971. *A Theory of Justice*. Cambridge, MA: Harvard University Press.
- Raz, Joseph. 1986. *The Morality of Freedom*. Oxford: Clarendon Press.
- Romero, Felipe. 2019. “Philosophy of Science and the Replicability Crisis.” *Philosophy Compass* 14 (11): 1–14. <https://doi.org/10.1111/phc3.12633>.
- Smith, Adam. 1776. *An Inquiry into the Nature and Causes of the Wealth of Nations*. 2 vols. London: Strathan & Cadell.
- Washington, Harriet A. 2007. *Medical Apartheid: The Dark History of Medical Experimentation on Black Americans from Colonial Times to the Present*. 1st edition. New York: Doubleday.
- Watson, Jamie Carlin. 2021. *Expertise: A Philosophical Introduction*. London: Bloomsbury.
- Wertheimer, Alan. 1992. “Unconscionability and Contracts.” *Business Ethics Quarterly* 2 (4): 479–96. <https://doi.org/10.2307/3857584>.
- Zagzebski, Linda Trinkaus. 2012. *Epistemic Authority: A Theory of Trust, Authority, and Autonomy in Belief*. Oxford University Press. <https://doi.org/10.1093/acprof:oso/9780199936472.001.0001>.