



Mental measurement and the introspective privilege

Why extrospective methods can prevail over introspective reports.

Michael Pauen¹

Accepted: 24 August 2023
© The Author(s) 2023

Abstract

According to a long-standing belief, introspection provides privileged access to the mind, while objective methods, which we denote as “extrospection”, suffer from basic epistemic deficits. Here we will argue that neither an introspective privilege exists nor does extrospection suffer from such deficits.

We will focus on two entailments of an introspective privilege: first, such a privilege would require that introspective evidence prevails in cases of conflict with extrospective information. However, we will show that this is not the case: extrospective claims can prevail in cases of conflict. These conflicts are resolved by an inference to the explanation that best accounts for the available evidence. This explanation may speak in favor of extrospection.

Second, given an introspective privilege, the development of extrospective measurement techniques should be restricted by the accuracy of introspective reports. We will argue that this problem is part of a more general issue that comes up with the establishment of measurement in the natural sciences. We will identify three strategies that have proven successful in dealing with the problem in the natural sciences. It will turn out that all these strategies are available for extrospective measurement as well. Consequently, the insufficiencies of introspective reports do not impose a limit on the accuracy of extrospective measurement methods.

We conclude that neither an introspective privilege nor basic extrospective deficits exist. This does not mean that extrospection will ever replace introspection. Rather, both approaches provide independent and indispensable forms of epistemic access to the mind.

Keywords Mental measurement · Extrospection · Introspection · Epistemology of the mind · First-person privilege · Pain measurement · Neurophenomenology

✉ Michael Pauen
michael.pauen@philosophie.hu-berlin.de

¹ Department of Philosophy, Berlin School of Mind and Brain, Humboldt-Universität zu Berlin, Unter den Linden 6, 10099 Berlin, Germany

1 Introduction

1.1 The epistemological shift

One of the major long-term developments in the philosophy of mind is a move from ontological to epistemological questions regarding the mind. As many philosophers have come to believe that mental states are, in fact, physical states, the ontological version of the mind body problem has lost center stage. The pressing question today is *whether* and, if so, *how* we can explain and understand the mind with objective methods – methods similar to those that have proven successful in other fields of science. In fact, the mind is still seen as special by many, but not anymore because it is made of some special stuff. It is now seen as special because it poses specific – and perhaps unsurmountable – epistemic challenges.

Objective methods targeting conscious experience – which we will denote as “extrospective methods” – seem to suffer from basic epistemic deficits. The only adequate method of knowledge acquisition about the mental seems to be provided by introspection. In fact, introspection has long been taken to provide privileged access to the mind (Alston, 1971; Goldman, 1997).

1.2 The introspective privilege

Here we will argue that there is no introspective privilege, and that the mind is not “special” in an epistemological sense, compared to typical non-mental phenomena. The mind is just awfully complex which makes scientific investigation a particular challenge.

Taken by themselves, the two claims about the introspective privilege and the corresponding deficits of extrospection are extremely general, such that it would be almost impossible to prove them true or false. In order to make them tractable, we will focus on the weakest version of the introspective privilege according to Alston (1971), namely on the *incorrigibility* of introspective claims: According to this understanding, introspective claims may in fact be wrong or give us reasons for doubt at times, but still, extrospective evidence will never be able to justify a revision of introspective claims. We will argue that even this weakest version of the introspective privilege should be rejected.

To be even more specific, we will zoom in on two assumptions that follow from the incorrigibility version of the introspective privilege. According to the first assumption, introspection will (almost always) prevail over extrospective evidence in cases of conflict – otherwise extrospective evidence would correct introspective claims. The second assumption is that the accuracy of introspective reports sets a limit for the accuracy of extrospective methods, particularly of extrospective measurement techniques – otherwise extrospective claims would be taken to be superior to introspective ones.

In our discussion of the first assumption in Section 2, we will look at conflicts between introspective and extrospective evidence. We will show that extrospection

can trump subjective reports because, often enough, we have multiple sources of evidence available, among them behavioral, physiological, neuroscientific, neurological, and etiological data in addition to subjective reports. Conflicts between different sources of evidence can be resolved with an inference to the best explanation – at times in favor of an extrospective claim.

In Section 3, we will discuss the second assumption, according to which extrospective methods are unable, in principle, to exceed the accuracy of introspective reports. In order to show that this is not the case, we will refer to similar difficulties that arise in measurement techniques in the natural sciences, which we will denote as “physical measurement techniques”. We will identify three strategies scientists have developed to deal successfully with those difficulties.

In Section 4 we will argue that these strategies can be used for the improvement of extrospective techniques as well. We will conclude that no introspective privilege exists, not even in its weakest version of mere incorrigibility. The mind is not epistemically special, neither does extrospection suffer from basic epistemic deficits.

As we will show in Section 5, this does not mean that objective techniques will replace subjective reports. While extrospective measurement can overcome the limitations of introspective reports particularly regarding quantitative accuracy, introspection has clear advantages if we want to capture the full breadth and diversity of first-person experience. Thus, introspection will remain an essential and indispensable source of knowledge about the mind.

2 Conflicts between intro- and extrospective evidence

2.1 Conceptual clarifications: introspection, extrospection, and conscious experience

Before we start, a clarification of a few essential concepts seems in place.

We understand *introspection* as a form of knowledge-acquisition directed at one’s own current mental states. So if you try to figure out whether the aversive feeling in your leg is a pain experience or just a tactile sensation, or whether you are really feeling comfortable with your new job, you are introspecting. While introspection does not need to result in a verbal report, it does require what Loar (1997) has called “recognitional concepts.” Recognitional concepts are ‘grounded in dispositions to discriminate by way of perceptual classification,’ thus enabling us to recognize our present experience as “one of that kind”, e.g., pain (Loar, 1997, 600).

Extrospection is a form of knowledge-acquisition as well, but unlike introspection, it uses objective evidence in order to figure out others’ mental states. So if you scrutinize the behavior and the verbal responses of your new employee in order to find out whether they feel comfortable with their new job, or if you do an fMRI scan of your patient’s brain in order to find out whether they are feeling pain, then you are extrospecting.

Conscious experience is the subject of both introspective and extrospective knowledge-acquisition but, taken by itself, it is devoid of epistemic implications.

Thus, from a purely conceptual point of view, you may have a pain experience without acquiring introspective knowledge about it.

Note that this conceptual distinction between introspection and experience does not beg the question against theories of consciousness that assume that conscious experience comes always with related introspective knowledge. Quite the contrary, the conceptual distinction is necessary in order to make this very empirical claim (see Section 5.3 for a discussion). Moreover, the distinction enables us to ascribe pain experience to animals and young babies who may be unable to acquire introspective knowledge.

2.2 Introspection and the introspective privilege

Introspection has long since been an important philosophical issue, mainly because it was thought to provide direct and, therefore, privileged access to the mental. Extrospection, by contrast, was thought to suffer from a basic epistemic deficit because it lacks direct access. Instead, it has to make do with fallible inferences based on indirect behavioral or physiological cues.

The idea of an introspective privilege can be traced back to ancient philosophers like Plato, Plotinus, or St. Augustine (Cary, 2000) who argued that directness saves first-person knowledge from the deceptive character of the senses. Much later, with the introduction of the concept of “consciousness” in the seventeenth and eighteenth centuries, the introspective privilege took center stage. The early proponents of the concept of consciousness tended to understand conscious experience itself as a form of knowledge acquisition (“science”, “scientia”) going along with (“co-“, “con-“) the mental state the knowledge is about: co-scientia, con-consciousness, conscience, co-scienza, Be-wusst-sein. Obviously, this intimate relationship between, e.g., an experience of pain and introspective knowledge *about* this pain state, results in a substantial privilege for introspective knowledge, which has been regarded as incorrigible, indubitable, infallible, or even omniscient by numerous philosophers (Alston, 1971).

The recent discussion has seen a further differentiation both regarding the metaphysics (Smithies & Stoljar, 2012) and the epistemology of introspection (Jack & Roepstorff, 2003), the latter aspect being central for the present paper. While the debate on the epistemological merits is far from reaching any consensus, it seems that skepticism regarding strong versions of the privilege has grown since the publication of Nisbett and Wilson’s (Nisbett & Wilson, 1977) study about the failures of introspective reports. Particularly the work of Eric Schwitzgebel has provided further empirical and theoretical evidence showing not only possible shortcomings of introspective reports (Schwitzgebel, 2002a, 2002b, 2008) but also explaining why we might have a tendency to underestimate these shortcomings. Schwitzgebel’s effort squares well with Emily Pronin’s work on the “Introspection Illusion” (Pronin, 2009; Pronin et al., 2007). Stronger skeptical claims regarding introspection have been made by Daniel Dennett and Elizabeth Irvine. While Dennett has long since argued that introspective beliefs have a fictional character (Dennett, 1991; Dennett

& Kinsbourne, 1992), Elizabeth Irvine (2019) expresses a “dark pessimism” regarding the relevance of introspective reports in consciousness science.

On the other hand, many scholars have defended the epistemic credits of introspection. Our discussion will pay special attention to Morten Overgaard’s justification of the introspective privilege (Overgaard & Sandberg, 2021; Overgaard, 2010, 2015; Sandberg et al. 2010; Overgaard n.d.). The reason is not only that his defense is particularly strong, but also that he focuses on conflicts between intro- and extrospective evidence that are pivotal for our argument as well. According to Overgaard, it is the dependence of extrospective methods on introspective reports that justifies the introspective privilege. E.g., if you want to use a specific sort of brain activity as an objective proxy for pain experience, you need introspective reports to make sure that this activity is a correlate of pain experience. Overgaard (n.d.) concludes that, given this dependence, extrospective methods cannot prevail over the sort of subjective evidence they depend on.

Goldman (Goldman, 1997) and Piccinini (Piccinini, 2003) have argued that introspection is a reliable source of evidence about the mind, but while Goldman denies that introspective claims can be validated, Piccinini thinks they can because there is no difference between introspection and third-person evidence in this respect. In a similar vein, Jack and Roepstorff (Jack & Roepstorff, 2002) have insisted that introspective reports are an extremely valuable, but often neglected, source of information about the mind, particularly for brain science.

Interestingly, there are still efforts to defend a kind of introspective infallibility, even if it might be restricted to specific circumstances. E.g., David Chalmers (2003) and Brie Gertler (2012) have argued that judgments directly based on phenomenal beliefs or direct acquaintance with phenomenal states may be infallible or quasi-infallible. The reason is not only the directness of the relationship but, even more so, the indexical character of judgments like ‘This [phenomenal] property is instantiated (in me, now)’ (Gertler, 2012). Importantly, the indexicals serve as placeholders for the experience itself, such that the statement is true, whatever the experience might be. While this raises triviality concerns, additional questions result from the restriction of this claim to very specific and rare cases of direct acquaintance, which makes it irrelevant for most introspective reports particularly in scientific experiments where more concrete forms of subjective report are required.

This extremely broad spectrum of positions calls for a clarification of the epistemic credits of introspection; a discussion of the introspective privilege can make an essential contribution here.

2.3 Conflict resolution

Here we will focus on a particularly weak version of the introspective privilege: are introspective reports privileged insofar as they prevail in cases of conflict with extrospective evidence?

One assumption underlying the opposite idea that introspective reports trump extrospective data seems to be, that these conflicts are he-said-she-said-cases, that is, standoffs between just two sources of evidence, introspective and extrospective.

Thus, a resolution of these conflicts, if it is possible at all, requires the application of general principles which assign epistemic superiority to one of these sources – like the introspective privilege.

Here we will try to show that this assumption is false: in a significant number of these conflicts, there are multiple sources of evidence available regarding a person's present mental state. This allows for a meaningful and empirically grounded resolution of such conflicts, typically by an inference to the explanation that best accounts for the available evidence. Importantly, this inference may, at times, support extrospective against introspective claims, depending on the evidence at hand.

2.3.1 Question begging?

But don't we beg the question against the introspective privilege if we assume that objective evidence may trump subjective reports? Let's first note briefly that there are very good reasons to reject the strongest version of the introspective privilege, namely the claim that subjective reports are infallible or even omniscient. Subjective reports are based on cognitive processes that help to memorize, interpret, and classify information about a given mental state, e.g., to recognize your present experience *as* an experience of pain. For all we know, cognitive processes fail at times, so it would appear unreasonable to assume that introspective claims are infallible. Moreover, the infallibility claim would immediately raise questions whether or not a given statement can be categorized as an introspective report and therefore as infallible. We think that these are reasons enough to reject infallibility.

2.3.2 Case study: Anton's syndrome

So let's come back to the question whether introspective evidence prevails in cases of conflict with extrospective data and let's turn to a case study in order to show that even this weak version of the privilege should be rejected.

Anton's syndrome is a rare neurological disorder (Anton 1899; Maddula, Maddula et al., 2009; Othman et al., 2019). Patients suffering from this syndrome are taken to be blind, although they insist, they can see. Thus, we have a conflict between patients' subjective reports and objective evidence provided by scientists and clinicians, and it is an inference to the best explanation that is taken to resolve the conflict in favor of the objective evidence.

Most importantly, the case shows how the inference is supported by quite a number of different sources of evidence. First, Anton's patients fail simple behavioral tests for visual perception: they do not realize when the psychiatrist in front of them extends their left hand when asking them for a handshake, so they extend their right hand; they claim that the psychiatrist wears a tie, although he doesn't; they maintain that the examination room has a window even if this is not the case (Othman et al., 2019; Maddula et al., 2009). Second, neurological evidence shows that Anton's patients have severe lesions in their visual cortices, which makes visual perception virtually impossible, and third, interviews show that the reports about their alleged visual experience are highly stereotypical: rather than reflecting actual perception of their environment or hallucinatory experiences like patients suffering

from Charles Bonnet's Syndrome (Eperjesi & Akbarali, 2004) do, Anton's patients seem to convey just generic knowledge about what can normally be expected in the present context, e.g., when one is asked for a handshake or sits in an examination room. Finally, there is neuroscientific evidence explaining why, due to their lesions, Anton's patients may confabulate experiences they don't have (Carjaval et al., 2012; Das & Naqvi, 2020).

Taken together, the best explanation for all these data is that Anton's patients are blind, and that their reports about visual perception are confabulations based on generic knowledge about their environment – even if the patients may believe in their reports. This line of reasoning may of course be challenged, e.g., opponents might insist that Anton's patients have hallucinatory experience. We would then have to look for further evidence. E.g., one could look for neurological markers that help to identify hallucinatory experience. The important point here is that there is a chance to find additional evidence even in the case of sophisticated objections.

The above discussion illustrates another issue that is relevant for the resolution of such conflicts. Resolving a conflict means showing that at least one piece of conflicting evidence is erroneous or misleading, while others are trustworthy. Ideally, the resolution would include an explanation of how this error came about. So, if we reject a subjective report in favor of objective evidence, we should be able to explain why subjects came to make false claims under the given conditions – as it is possible with respect to the confabulations of Anton's patients. The same goes when we reject objective evidence in favor of subjective reports. E.g., subjects under high stress may deny feeling pain although they are severely wounded (Hardcastle, 1997, 132; Melzack & Wall, 1983; see below). In this case there are strong reasons to assume that subjects are right because there are neural mechanisms in the spinal cord which block the afferent pain signals from reaching the brain and thus from producing pain experience in stressful situations.

But what about Overgaard's justification of the introspective privilege? Let's get back to Anton's syndrome in order to discuss his view. If Overgaard is right that the dependence of objective methods on subjective reports justifies the introspective privilege, this dependence should affect the role of, e.g., objective neurological data in our assessment of the Anton's patients' visual experience. In this case we have objective evidence (widespread lesions in the patients' primary visual cortex) on the one hand. Based on a substantial number of subjective reports, this evidence is thought to show that patients lack visual perception. On the other hand, we have patients' claims that they *do* visually perceive.

But if we reject subjective reports in the case of Anton's patients, how can we accept reports about the effect of lesions in the visual cortex? The reason is simple: as it is common practice in scientific research, we do not apply theoretical principles like the introspective privilege in our treatment of the evidence, rather we take all the evidence available and then look for a conclusion that best explains all the evidence. In the case at hand, our rejection of Anton's patients' subjective reports is motivated by very specific and well-justified doubts regarding these particular circumstances. These very specific doubts can go along with the – as we think well-founded – assumption that, in general, introspective reports are an important and reliable source of evidence. This holds in particular if these reports are confirmed

by a host of other evidence – as it is the case with our knowledge about the consequences of lesions in the primary visual cortex. We conclude that even if extrospective methods depend on subjective reports, this does not exclude the rejection of individual pieces of subjective evidence under specific circumstances.

The present case study can also demonstrate how the resolution of conflicts between first- and third person evidence can help to make progress in our understanding of the mind and, eventually, in the development of measurement methods. The reason is that these resolutions may help to identify systematic errors in our methods and our beliefs about the mind whose removal, in turn, may improve our interpretation of the relevant data: once we have understood why subjects who are severely wounded do not feel pain in stressful situations, we will be less likely to draw wrong conclusions from severe tissue damages in stressful situations. It would follow then that conflict resolution may be an important driver of progress in our endeavor to understand and measure the mind.

So, let's conclude that, in general, there is neither a need nor a justification for the use of a priori principles like the introspective privilege if we want to resolve conflicts between subjective reports and objective data in extrospective research. Due to the availability of multiple sources of evidence, inferences to the best explanation can help to resolve these conflicts pretty much like they can help to draw reasonable conclusions from imperfect and maybe inconsistent evidence in other fields of science. Still, introspection can be regarded as a particularly important and reliable source of evidence, even if it does not have any specific privilege over objective information and may be erroneous at times.

3 Subjective observation and the development of measurement techniques in the natural sciences

In this section we will start discussing the second question following from introspective incorrigibility which relates to extrospective measurement. Measurement techniques are essential to any scientific endeavor, they are well understood, and the standards for good measurement techniques are quite uncontroversial as well (BIPM, 2019).

If the mind was epistemically special, such that introspection was privileged while extrospection suffered from a basic epistemic deficit, we would expect that this deficit affects extrospective measurement as well. The extreme differences between the accuracy of physical and extrospective measurement seem to show that this is indeed the case. And as introspection is our only method for directly accessing the mind, the obvious limitations of introspection seem to impose a ceiling on extrospective measurement techniques as well – as Overgaard (n.d.) has argued.

In what follows we will denote this limitation of extrospective methods through the shortcomings of introspection as the *Ceiling Problem*. While we agree with Overgaard that introspection plays a pivotal role for the establishment and calibration of extrospective measurement, we will argue that the limitations of introspective reports particularly regarding accuracy do not impose a ceiling on extrospective techniques. The problem is part of a larger issue that affects the establishment

and calibration of physical measurement techniques as well. Here, in Section 3, we will describe three strategies that have been developed to deal with this problem in physical measurement. Below, in Section 4, we will show that the same strategies can be employed successfully for the improvement of extrospective measurement techniques.

Historically, the development of measurement techniques, such as thermometry or photometry, has typically taken its departure from subjective reports (Chang, 2007), but before long, scientists established objective methods for measuring temperature (Chang, 2007), brightness (Chen, 2005), and electricity (NIST, 2021b). But how did they push these objective methods to the extreme levels of accuracy and precision they enjoy today, given that all they could initially rely on was subjective observation with all its shortcomings, including limited discriminatory abilities regarding, say, temperature, weight, or brightness? Moreover, how did these scientists overcome the problems that came up later whenever they tried to establish a new measurement technique, or improve an existing one, but only had the old and fallible methods by which to calibrate the new or improved techniques? We will show that there exist at least three strategies for improving measurement techniques or introducing new methods that are partially independent from already existing methods in the field, including subjective observation.

3.1 Strategy I: Improving existing techniques

Let's first talk about some of the methods that helped to improve given measurement techniques and to make them more powerful than existing methods.

3.1.1 Improving specificity and sensitivity

One important approach is to improve a given measurement technique by increasing its specificity or sensitivity. A lack of specificity can result from a confounding variable. Measuring rods, for example, expand and contract when temperature changes, as do glass thermometers.

One way to deal with this problem is to control for the confounding variable in question. E.g., length measurements that aimed at high accuracy were conducted in closed rooms with a constant temperature; likewise, the standard meter was kept at 0 °C (Klein, 1988, 185). In other cases, the technique itself was modified to compensate for the confound: the scale of glass thermometers can be adjusted in order to account for the expansion of glass at higher temperatures, and Huygens avoided variations of the pendulum frequency brought about by changing amplitudes by forcing the pendulum to follow a cycloidal path (Bell, 1941).

The decisive point here is that these responses and their effects do not depend on the accuracy of existing measurement methods in the relevant field. Huygens' introduction of the cycloidal pendulum was independent of existing measurement techniques because it was based on theoretical considerations (Huygens, 2007 (1673); Bell, 1941). Likewise, the gridiron pendulum automatically compensated for the

expansion of metal at rising temperatures, even without calling for additional evidence (Andrewes, 2002).

3.1.2 Leveraging

Second, the results of existing measurement techniques can be leveraged in order to confirm improvements that exceed the accuracy of existing techniques. So if you want to adjust the scale of a thermometer to compensate for the expansion of the glass body at higher temperatures, you need length- and temperature measurement in order to quantify the effect. Importantly however, the gain in accuracy for the new technique does not depend on the accuracy of the existing method: you can make up for the potentially low accuracy of existing thermometric techniques by measuring the expansion between a very high and an extremely low temperature and then interpolating the missing values between them.

3.2 Strategy II: Using and manipulating measurement standards

The second strategy employs measurement standards as an independent source of information, which can be used for both the calibration of new methods and their comparison with extant techniques.

3.2.1 Measurement standards

Measurement standards are objects or procedures that realize a specific measurand, or a quantity thereof. Paradigmatic standards are the International Prototype Meter, or the Kilogramme des Archives. Measurement standards reverse the relationship between measurand and measurement technique. Normally, we use a measurement technique, e.g., a thermometer, to provide information about the measurand, e.g., its temperature. By contrast, when using a standard like the boiling point of water in a calibration process for a thermometer, the measurand (i.e., the standard) provides information about the measurement technique, e.g., whether the thermometer correctly displays 99.97 °C (211.9 °F).

3.2.2 Manipulation

Another reason why measurement standards are important sources of evidence is that we can manipulate them in systematic ways and then examine whether a measurement technique correctly reflects these manipulations. So, if you have a measurement standard of one kilogram on a weighing scale, and then add a second standard of one kilogram, your measurement device should display two kilograms.

3.2.3 Indirect standards

Standards for physical measurement may well be indirect. For example, the measurement standard for an ampere, the base unit of electric current, was once defined

as a certain amount of silver produced by an electrolytic process in a silver voltameter (Suplee et al., 2018). Later, scientists agreed on a definition in terms of a specific attractive force between two wires (NIST, 2021b). In both cases, the measurement standard is not the phenomenon itself, i.e., the current, but one of its indirect causal consequences. This indirectness, however, does not necessarily result in an epistemic deficit: if the causal mechanisms are well understood and easy to control, indirect procedures may even be superior to direct ones. That is why a direct standard like the platinum-iridium bar of the international prototype meter has been replaced by an indirect method that determines “the length of the path traveled by light in a vacuum in $1/299,792,458$ of a second” (NIST, 2021c).

3.3 Strategy III: Improving measurement standards

The accuracy with which a given measurement standard realizes the related unit may differ vastly. In general, the improvement of standards proceeds very much like the improvement of measurement techniques. Thus, one way of improving measurement standards is to identify potential confounds. Once a confound has been identified, the definition of a measurement standard can be qualified accordingly, the realization can be updated, or a new standard can be introduced. E.g., prior to 1948, the standard for the unit of a candela was a spermaceti candle. Although carefully defined, the standard left room for considerable variations between individual realizations. But once scientists began to suspect that certain variables might affect the brightness of the flame, they could test this suspicion and, if the test was positive, control for these factors, modify the standard, or introduce a new standard altogether, e.g., a laser beam with a precisely defined power and frequency.

3.4 Conclusion

Let us conclude, that scientists had three strategies at their disposal to address the Ceiling Problem: First, they could improve existing techniques by removing confounds and leveraging available quantitative methods; second, they used measurement standards to calibrate new techniques and compare them with existing ones; and, third, they could enhance the accuracy of measurement standards, particularly indirect standards, in pretty much the same way as they could improve measurement techniques, that is, by removing confounds. These three strategies can explain why physical measurement techniques have reached the unprecedented levels of accuracy that we see today.

4 Developing mental measurement

But what do these considerations about physical measurement techniques tell us about extrospective measurement? Here we will present theoretical considerations and case studies in order to show that basically all the strategies that were

successfully used in the development of physical measurement techniques can be employed for the improvement of extrospective methods as well. That is, we can.

- improve extrospective measurement techniques in ways that are largely independent from existing methods (Strategy I),
- use objective measurement standards to confirm that newly developed extrospective methods are superior to existing ones (Strategy II), and
- improve the accuracy of objective standards for the measurement of subjective experience as well (Strategy III).

As case studies play an important role in our argument below, it is important for us to stress that we use these studies as proofs of principle only. They are intended to illustrate that a certain method is applicable in principle, even if you may doubt that the application in the specific case in question was successful and the results obtained are correct.

Many of our examples come from pain measurement. So far, pain measurement has typically relied on standardized questionnaires, such as the McGill Pain Questionnaire (Melzack, 1987). More recently, however, scientists and engineers have started to develop objective measurement techniques for pain, particularly in light of the increasing availability of machine learning techniques (Mouraux & Iannetti, 2018; Woo et al., 2017).

Here we will focus on one of the more recent methods for measuring pain, namely the Neurologic Pain Signature (NPS) which was developed by Tor Wager and his group (Wager et al., 2013). The NPS is based on a machine learning algorithm that uses fMRI data of activity distributions in pain-relevant areas of the brain. The NPS can measure up to six degrees of actual physical pain induced by noxious thermal stimuli, and can distinguish these from similar experiences, such as social pain, pain anticipation, or pain recall, with a specificity and sensitivity of well above 90%.

Obviously, subjective reports were needed for the establishment and calibration of this extrospective measurement technique. But the point we want to make in what follows is, that, this dependence notwithstanding, the improvement of the NPS is not limited by potential shortcomings of subjective reports. So unlike Overgaard's claims, the NPS *can* exceed the accuracy of the subjective reports which were used for establishing this measurement technique (compare Overgaard, n.d.; Overgaard & Sandberg, 2021).

4.1 Strategy I: Improving existing techniques

Our first point will be that extrospective measurement techniques can be improved in basically the same way as physical measurement techniques can. The recent development of the NPS illustrates how this is possible – that is, how Strategy I can be applied to extrospective measurement. First, a variable affecting the specificity or sensitivity of the measuring mechanism is identified. Then, the variable can either be controlled or it can be accounted for by the measuring mechanism (in this case, by the NPS).

4.1.1 Case Study: Psychological factors and the neurologic pain signature

In its original form, the NPS is sensitive only to the activity in nociceptive brain areas. Woo et al. (2017) demonstrated that psychological factors like pain expectancy and perceived control do affect pain intensity as well, even if they are not accounted for by the NPS. That is why, in a second step, Woo et al. reanalyzed pain-related fMRI data in order to identify the residual variance that remains after removing those variations that the NPS in its present form was already sensitive for. In a third step, Woo et al. developed another classifier, the SIIPS1, that was able to account for psychological factors. Combined with the NPS, the SIIPS1 lead to more accurate results.

Woo et al.'s study provides a powerful illustration of how the accuracy of extrospective measurement techniques can be improved if we account for variables that used to be ignored so far. First, we have to establish that a specific factor (e.g., pain expectation) does affect the outcome variable (reported pain intensity) but is *not* accounted for by the existing measurement technique (the NPS) – as Woo et al. did in the first step of their study. Once we have established the relevance of the factor in question, there are two ways to move on. First, we can control for this factor; second, we can improve the measurement mechanism, such that it accounts for the factor – as Woo et al. did by adding the SIIPS1 to the NPS-based measurement technique. As already indicated above, we do not take Woo et al.'s study at face value; rather, we regard it as a proof of principle that Strategy I can be employed for the development of extrospective methods.

Note that, conversely, basically the same strategy would allow us to address a lack of specificity. In this way, we could improve a (hypothetical) NPS-style measurement technique that was affected by pain expectancy and perceived control, although these factors have turned out to be irrelevant for pain experience. Again, we could either control for the confounding psychological variables in the design of the measurement process, or improve the NPS, e.g., by excluding activity from non-nociceptive areas from the dataset in the first place.

But as the entire procedure described above used first-person reports as a measure for pain intensity, why can't Overgaard and other skeptics insist that extrospective results are affected by the limitations attendant to subjective reports? Due to the essential role that subjective reports play in step 1 of Woo et al.'s study, the limited sensitivity of these reports should still impose a ceiling on the overall sensitivity of the entire pain measurement system (NPS & SIIPS1), according to Overgaard's view, even if these limitations did not become visible in Woo's study.

In order to see that this is not the case, imagine that subjective reports were indeed severely limited, such as to be insensitive to subtle changes in pain experience, as they occur under real-world conditions. As a consequence, Woo et al. would not have been able to link these variables to changes in pain experience in the first step of their study. But even then, Woo could have leveraged subjective reports by using only extremely salient stimuli, whose effects on pain experience *are* introspectively accessible. By pursuing this strategy, Woo et al. should have been able to establish the connection between these particularly salient psychological stimuli and pain experience in step 1. In step 2a, the SIIPS1 could then be trained by means

of these particularly salient stimuli. Note that this does not rule out that the SIIPS1 can distinguish even weaker contrasts between more subtle real-life stimuli as well. If this were the case, it would give us good reasons to conclude that the SIIPS1 *does* measure a real difference in experience, even if this difference escapes subjective report.

To sum up, our considerations above allow for two important conclusions: first, we can improve extrospective measurement methods by identifying variables affecting the sensitivity or the specificity of these methods, very much like we do when improving physical measurement techniques. Second, we can deal with these variables in two ways. We can control for them, or, as Woo et al. have shown, we can improve the sensitivity of extrospective measurement techniques in a way that is not limited by the shortcomings of first-person reports. We have also demonstrated that there are reasons to believe, that the same strategy would work for a lack of specificity as well. All this shows that Strategy I should be available as a ceiling breaker also for extrospective measurement.

4.2 Strategy II: Using and manipulating measurement standards

Our second question will be whether extrospective methods can use and manipulate measurement standards as described in Strategy II. According to our earlier definition, measurement standards are objects or procedures that realize a specific measurand, or a quantity thereof like the International Prototype Meter. In the case of extrospection, the standard could be the experience of a specific shade of red or a certain degree of pain intensity.

Admittedly, our general claim that extrospective measurement can utilize basically the same strategies as physical measurement may appear less obvious in this case: extrospective and physical measurement standards seem to differ in kind and not only in degree. In physical measurement, we have standardized methods to make sure that, e.g., a given sample of water has a temperature of 100 °C and thus is a correct realization of a specific quantity of heat, due to our direct access to this sample. In extrospective measurement, however, it seems that we have to ask for a subjective report, if we want to make sure that the experience in question is a correct realization of the property we want to measure (e.g., a specific intensity of pain). Thus, subjective reports that we tried to throw out through the front door by using an objective measurement standard, seem to return through the back door once we try to establish these standards. This is problematic not only because extrospective standards turn out to be indirect. It is a problem also because these standards seem to inherit the shortcomings of subjective reports – as it has been stated by Overgaard. So, against our hypothesis, it would seem that extrospective measurement *does* suffer from basic epistemic deficits.

In order to assess these issues, let's first focus on the problem of indirectness. As we have already seen above in Section 3, indirect measurement standards play an important role in physical measurement as well. In fact, not only is the measurement standard for (e.g.) electrical current indirect, but even paradigmatic *direct* standards, like the International Prototype Meter and Kilogram, have now been replaced by

indirect standards, because these indirect standards are more accurate than the previous direct ones. However, indirect standards or procedures are available for mental states as well. For example, a noxious thermal stimulus with a specific temperature can produce a pain experience of a certain intensity, and thus can be used as a standard for calibrating a computer-based measurement technique like the Neurologic Pain Signature (NPS).

Taken by itself, however, the mere existence of indirect standards for extrospective measurement doesn't show very much, even if these standards are common in physical measurement as well. Rather, we have to demonstrate that these indirect standards can play the same role for the development of extrospective measurement as indirect standards did in the improvement of physical measurement techniques. This means, first, that we will have to see whether objective measurement standards, like noxious thermal stimuli, can be used to calibrate new and improved mental measurement techniques independently of the limitations of subjective reports, even if these reports have been instrumental in establishing such standards. This is what we will explore here in three case studies on Strategy II. Second, we will have to ask whether we can improve the relationship between indirect extrospective standards and the actual measurands in a manner that is basically analogous to the way this is done in physical measurement. This question will be addressed below in Section 4.3 which will be devoted to Strategy III.

4.2.1 Case Study: Optokinetic nystagmus

Our first example is taken from research on binocular rivalry. In a typical binocular rivalry experiment, two different stimuli are projected simultaneously to the right and to the left eyes of an experimental subject, e.g., a green grating moving to the right is projected to the right eye, and a red grating moving to the left is projected to the left eye. However, the visual system switches between these two stimuli such that, at any given point in time, the subject will experience only *one* stimulus. In standard experiments, subjects are asked to report when these switches between the stimuli occur.

Recently however, Frässle et al. (2014) and Naber et al. (2011) used the Optokinetic Nystagmus as an objective measure for the stimulus switch instead. The Optokinetic Nystagmus is a swift eye motion that follows a moving stimulus, as it might occur when a person is standing on a platform and watches a train passing by. Interestingly, the Optokinetic Nystagmus can also reveal the dominant stimulus in contexts of binocular rivalry: the eyes will move to the right if the green stimulus moving to the right is dominant, and they will move to the left if the red stimulus moving to the left is dominant.

The authors claim that this technique is superior to subjective reports – just as photodetectors turned out to be superior to subjective photometry in the early twentieth century. But how can this claim be justified, given that any statement about subjective experience seems to require subjective reports at some point? And if so, how can subjective reports confirm a level of accuracy that goes beyond the accuracy of these very reports? This is, exactly, Overgaard's question.

The solution here is that Frässle et al. made use of an independent source of information, namely an objective measurement standard and its manipulation. In the so-called “Replay Condition”, the experimenters produced the stimulus switches themselves, rather than relying on the visual system to do so. That is, they always projected just *one* stimulus to both eyes of their experimental subjects, switched to the other stimulus, and then back to the first one from time to time. As a consequence, they knew exactly when the switches would occur. This provided the authors with a standard for a comparison between subjective reports and the Optokinetic Nystagmus, and it was the latter that turned out to be more accurate.

Note that the standard is indirect as the stimulus switch is taken to be the cause of the measurand, namely the related switch of the experience. This assumption might of course be challenged, and, in the case of serious objections, would have to be defended with additional evidence – or dismissed, if no such evidence could be found. Importantly, this case study thus shows how extrospective methods can utilize independent, objective information provided by measurement standards and their manipulation to show that a new measurement technique like the Optokinetic Nystagmus is superior to subjective reports.

4.2.2 Case Study: A thought experiment

It might be objected, though, that Frässle et al.’s study presents a somewhat unusual case of extrospective measurement. After all, their study does not quantify any property of the experience in question, it just indicates whether or not a person has a specific visual experience at a given point in time. So, is there any reason to assume that quantitative extrospective measurement can take advantage of measurement standards, as it does in physical measurement?

In order to see that there are such reasons, consider the following thought experiment. Imagine, first, that a machine learning algorithm like the NPS can distinguish between six different degrees of pain intensity (1, 2, 3, 4, 5, 6), while subjective reports distinguish between three degrees only (A, B, C). Let’s call these degrees *weak*, *medium*, and *strong*. Finally, assume that we have noxious thermal stimuli with six different temperatures (i, ii, iii, iv, v, vi) which are used as measurement standards, thereby providing a third source of evidence. The decisive question, now, is whether the lower accuracy of the first-person reports imposes a limit on the accuracy or resolution of the NPS.

Let’s further posit that introspective reports confirm that subjectively experienced pain intensity (as indicated by A-C) increases continuously with the temperature of the noxious stimuli (i-vi). Even then we should expect that, averaged over the entire group, objective stimuli i, ii, iii, iv, v, and vi lead to a continuous increase of subjectively reported pain intensity from A over B to C, because it is more likely for, e.g., stimulus ii to be associated with experience B than for stimulus i. The NPS, by contrast, tracks the six different stimulus categories in each individual subject, again in increasing order (i:1, ii:2, iii:3..., vi:6).

We think that, in the context of the available objective evidence, subjective reports in this – hypothetical but possible – scenario would give us very strong reasons to assume that the NPS classifier does distinguish not only the six *objective*

stimulus categories, but also the six degrees of *subjective* pain intensity. If this is true, an objective measurement technique can exceed the accuracy of subjective reports. And, again, it is the manipulation of objective measurement standards, that is, of the temperature of noxious thermal stimuli, that provides an important piece of evidence in support of this conclusion.

4.2.3 Case Study: Citalopram

In fact, this is not just a hypothetical scenario; rather, Ma et al. (2016) have actually conducted an experiment along these lines. The authors systematically manipulated a standard, namely pain intensity caused by a noxious thermal stimulus, to compare subjective reports and the NPS. The manipulation was effected by means of Citalopram, a selective serotonin reuptake inhibitor (SSRI) that can also be used as a pain killer. The results of this study support our claim that the accuracy of the current version of the NPS can overcome the limitations of subjective reports – even if only by a small margin.

Importantly, experienced pain intensity was measured in two ways: by subjective reports and by the NPS. While there was a trend towards pain reduction by Citalopram in both the self-report and the NPS group, the effect of the pain killer reached significance only in the NPS group, indicating that the NPS might be more sensitive to pain experience than subjective reports, even if the trend in subjective reports confirms the NPS results. As in the thought experiment above, it was the manipulation of a measurement standard, i.e., heat-induced pain experience, by the administration of a pain killer, that provided the evidence for the comparison between self-report and objective method; and it was the latter that turned out to be superior in this comparison.

To sum up, these three case studies show that indirect measurement standards can be instrumental in providing evidence that allows for a comparison between subjective and objective measurement methods, thus supporting our claim that Strategy II applies to mental measurement as well.

4.3 Strategy III: Improving measurement standards

The conclusions drawn above are based on an assumption that may be false, namely that there is a systematic covariance between the indirect measurement standard and subjective experience, such that, e.g., the intensity of a pain experience covaries with the temperature of a thermal stimulus. As we have seen above, the accuracy of standards in physical measurement can be challenged and improved as well, and it was Strategy III that allowed scientists to identify inaccuracies and to address them.

Here we want to show that Strategy III does apply to extrospective measurement: we can identify and address inaccuracies of extrospective measurement standards as well. Due to principal similarities between Strategies I and III we can largely refer to what we have said about Strategy I above. As in Strategy I, the control of confounds is an important tool for achieving these improvements. There are many examples for confounds that may affect extrospective standards, some of which have already

been mentioned above. E.g., noxious thermal stimuli may be affected by psychological variables like pain expectancy and perceived control, as the study by Woo et al. (2017), described above, has shown. Likewise, pain experience may be attenuated by stress, e.g., in combat, or sports competition (Hardcastle, 1997, 132; Melzack & Wall, 1983), and perceived temperature may be affected by previous experiences of heat or cold.

However, as we have also demonstrated above in Section 4.1, it is possible to control for or to avoid these confounds once they have been identified. In order to do so, we can often use the leveraging technique mentioned above, that is, we can use particularly salient stimuli in order to detect confounds with experimental methods as shown above in Section 4.1.1 regarding pain expectancy. This method can help us to identify confounds that might go undetected in subjective reports under real-world conditions. Once these confounds are identified, we can make sure that experimental subjects are not affected by the psychological confounds, that they do not experience stress, or that all subjects in an experiment are adapted to the same temperature before applying a thermal stimulus.

Of course, the number of potential confounds in extrospective measurement is significantly higher than in most fields of physical measurement. But this is not because the mind is “special” due to its first-person character; it is because the mind is so complicated, or more precisely, because it is sensitive to so many factors. While this is an obvious challenge for extrospective measurement, it is a challenge that psychology has learned to deal with, e.g., by using statistical methods.

These problems notwithstanding, the above examples and considerations show that we can improve the accuracy of our indirect extrospective measurement standards in basically the same way as we can improve indirect standards in physical measurement. This means that all the three strategies that have proven essential for the solution of the Ceiling Problem in physical measurement, can be employed in extrospective measurement as well. That is, even if new methods depend on quantitative knowledge provided by old and well-established techniques including subjective reports, this does not mean that the limitations of the old methods restrict the accuracy of new ones. The main reason is that we can utilize various techniques and sources of evidence that are independent of existing measurement methods in a given field. Most importantly, we can remove confounds from existing techniques, we can use measuring standards for comparison and calibration, and we can improve these standards in much the same ways that we can improve measurement techniques – in both physical and in extrospective measurement. And this means that there are good chances to overcome the Ceiling Problem which, in turn, provides further evidence that the mind is not epistemically special, and no introspective privilege exists.

5 Consequences and objections

5.1 Will extrospective measurement replace introspective reports?

In the last section of this paper, we want to discuss some of the most obvious objections and possible consequences. Our first example is both a possible consequence and an objection at the same time: if mental measurement can provide us with objective standardized data about almost any experience, thus overcoming the limitations of introspection particularly regarding accuracy, wouldn't that mean that it will replace introspective reports sooner or later?

In order to answer this question, let's look briefly at the inherent strengths and weaknesses of both kinds of knowledge acquisition. While it is true that extrospective measurement techniques provide objective, standardized data, they do so only regarding one small aspect of first-person experience, e.g., pain intensity. While extrospective measurement may exceed the accuracy of introspective reports, it seems impossible that it will ever be able to parallel introspection in another important respect: the wealth of qualitative information that introspection is able to provide. The reason is that introspection has a much larger scope; it may cover almost the entire spectrum of experience of a person, including their emotions, thoughts, and desires, and, as far as pain is concerned, it covers various aspects of pain, beyond mere intensity – even if introspective reports lack objectivity, standardization, and quantification.

So what we have here are two substantially different methods of knowledge acquisition about the mind, each of which is suitable for specific situations and less useful for others: If you are a doctor who needs detailed and standardized information about a person's pain intensity, then use a pain measurement device, but if you are a psychiatrist who needs the broad picture regarding your patient's thoughts, feelings and desires, just ask them for an introspective report. Of course, the difference is not completely written in stone, as e.g., certain introspective methods in psychology can help to make progress regarding standardization of introspective reports and there are even ways to combine both methods as we show in Section 5.2 below. Still, we think that the basic difference is difficult to overcome.

Given this basic difference, it is very unlikely that extrospective measurement will displace introspective reports: the information each of them provides is just too different. We cannot give up on introspective report without ending up with a severely impoverished understanding of subjective experience. What we can expect is the usual relationship between different sources of empirical evidence, each with its own strengths and weaknesses which make it sometimes prevail over and sometimes succumb to the other.

Another reason why a replacement of introspective reports by extrospective measurement is highly unlikely has already been mentioned above in Section 2.2: the establishment and calibration of extrospective measurement techniques requires introspective information. If you take a certain neuroscientific activity as a proxy for pain, then you have to ask the participants for subjective reports of pain experience in the first place.

5.2 Extrospective measurement, quality spaces, and phenomenology

But if subjective and objective methods are so different – wouldn't it make sense to combine them? This is indeed the case. One example are so-called quality spaces that map the different aspects of one type of experience like color sensation (Rosenthal, 2010) or pain (Coninx, 2022) to a multidimensional space. Color spaces employ methods from psychophysics including subjective reports and map hue, saturation, and lightness to a three-dimensional space, thus allowing for an objective measurement of the most essential aspects of color experience.

Classical phenomenology provides a particularly important example since it does offer a systematic way to investigate subjective experience which can help to make progress with objective methods as well. As Geniusas (2016) has shown, Scheler's phenomenology of pain (Scheler, 1954a, 1954b) already distinguishes between sensory and affective pain in the early twentieth century, that is, more than fifty years before this distinction was established in empirical pain research. While the opportunity to use phenomenological insights for neuroscientific research was missed in Scheler's case, Neurophenomenology, as initiated by Francisco Varela combines phenomenology with objective neuroscientific methods (Berkovich-Ohana et al., 2020; Gallagher, 2015). While the original idea was to make progress on the hard problem of consciousness (Thompson et al., 2005), more recent approaches try to capture the full phenomenology of subjective experience, e.g., of awe and wonder, to identify the neural correlates of this experience (Gallagher, 2015).

The examples underline that subjective and objective methods are complements, both of which are needed to get an adequate idea of the full complexity of phenomenal experience and rigorous objective methods for experimental work.

5.3 Does the development of extrospective measurement lead to substantial progress?

But can we really expect that the development of mental measurement as outlined above will have any significant consequences for future research, and, if so, what might those consequences be?

One reason to assume that they will have such consequences is that improved extrospective measurement techniques will put neuroscientists into a much better position to measure relevant aspects of first-person experience. No matter whether we want to know how a certain drug, a behavioral training, or the activity in a specific neural areal affect pain intensity: the better our abilities to measure this aspect of pain experience, the better are our chances to investigate and understand the neural mechanisms underlying this effect. Moreover, an improvement of pain measurement will have substantial relevance for clinical practice, as current pain-assessment by questionnaires, even if highly useful (Ngamkham 2012), has inherent limitations, particularly in the case of chronic pain, for patients with dementia, and for small children.

The debate about cognitive and non-cognitive theories of consciousness (Michel & Morales, 2020; Overgaard & Grünbaum, 2012) provides a concrete example for the effect of progress in extrospective measurement on neuroscientific research.

Cognitive theories, like Baars's and Dehaene's Global Neural Workspace Theory (Baars, 1996, 1997; Dehaene & Naccache, 2001) claim that cognitive processes including introspection and the underlying activities in central areas of the brain are constitutive for consciousness. As a consequence, introspective reports are the method of choice if evidence for conscious experience is needed. By contrast, non-cognitive theories like Tononi's and Koch's Integrated Information Theory (Tononi et al., 2016) deny this; in their view activities in peripheral brain areas can be sufficient for conscious experience and introspective reports come in addition to conscious experience (Michel, 2017).

It could be argued that these are just conceptual questions regarding the semantics of "consciousness", questions that cannot be decided empirically. Fortunately however, so-called no-report paradigms like the Optokinetic Nystagmus mentioned above in Section 4.2.1 speak to this debate. Frässle et al. (2014) did an experiment where they used subjective reports in one condition and the Optokinetic Nystagmus in the other condition to determine the dominant stimulus in conscious experience. It turned out that central brain activity was present only in the subjective report condition, indicating that central activity is *not* a necessary ingredient of conscious experience – as predicted by non-cognitive theories of consciousness. While the interpretation of this study is still under debate (Michel & Morales, 2020; Tsuchiya et al., 2015), the experiment shows how progress in extrospective measurement, in this case the availability of no-report measurement techniques, can have tangible consequences for empirical research already today.

6 Conclusion

In this paper we argued against a widely shared view according to which the mind is epistemologically "special", such that it cannot be investigated and explained with objective methods, as it is done elsewhere in the natural sciences. According to this view, the method of choice for acquiring knowledge about the mind is introspection which is thought to have privileged access due to its directness. By contrast, objective third-person methods, which prevail in any other field of scientific research, are thought to suffer from substantial epistemic deficits when it comes to the mind.

In order to make these somewhat general and diffuse claims tractable, we focused on two specific issues regarding extrospective measurement that would follow from an introspective privilege. These issues have been raised more recently by Morten Overgaard who defends an empirically and epistemologically grounded version of the introspective privilege. Such a privilege would imply, first, that, in cases of conflict, introspective evidence prevails over extrospective data. Second, it would follow that the accuracy of extrospective measurement techniques is limited by the shortcomings of introspective methods, given the essential dependence of extrospective methods on introspection.

In Section 2, we demonstrated against the first assumption that extrospective evidence can prevail over introspective data. Cases of conflict can be decided with an inference to the best explanation based on multiple sources of empirical evidence. At times, the best explanation can favor extrospective over introspective evidence.

Putting an emphasis on measurement methods in Section 3, we demonstrated that the supposed limitation of extrospective measurement methods is just one instance of a more general problem that affects physical measurement as well: in both cases, progress requires replacing older and less accurate methods with newer and more precise techniques, even if the latter may depend on the former. However, scientists have successfully used particularly three strategies to improve existing measurement methods and to establish more accurate new ones.

In Section 4, we demonstrated that these strategies can be applied to extrospective measurement as well, thus showing that the second assumption has to be rejected as well. That is, (1) we can improve existing extrospective measurement techniques largely independently of the limitations of subjective reports, e.g., by removing confounds and leveraging subjective reports, (2) we can use objective standards in extrospective measurement as an independent source of evidence for the calibration of new measurement techniques; and (3) we can improve the accuracy of indirect objective standards as well.

In light of these considerations, we conclude in Section 5 that introspection does not enjoy a specific epistemic privilege, neither does extrospection suffer from basic epistemic deficits. Rather, both forms of epistemic access are important sources of information about the mind. More particularly, extrospective methods will not displace introspective ones; rather, both methods are needed in order to avoid an impoverished picture of the mind, which can be investigated and explained in basically the same ways as other highly complex systems.

Acknowledgements The author would like to thank Morten Overgaard, Elisa Filevich, Sascha Benjamin Fink, John-Dylan Haynes, Franz Hildebrandt-Harangozó, Dimitri Coelho Mollo, Cosima Pauen, Jens Schneider-Mergener, Sera Schwarz, Ali Yasar, the members of the Research Training Group “Extrospection” at the Berlin School of Mind and Brain, and two anonymous reviewers for helpful suggestions and important comments.

Author’s contribution The paper was entirely written and conceived by Michael Pauen.

Funding Open Access funding enabled and organized by Projekt DEAL. The author’s work has been funded by the German Research Foundation,

Deutsche Forschungsgemeinschaft,
grant GSC 86/1 Berlin School of Mind and Brain,
grant 337619223/RTG 2386 Extrospection and
grant 390523135/EXC 2002/1 Science of Intelligence.

There was no involvement whatsoever of the funding organization in the conceptualization, writing, and editing of this project.

Data availability Not applicable.

Declarations

Ethical approval Not applicable.

Statement regarding research involving human participants and/or animals Not applicable.

Informed consent Not applicable.

Competing interests The author declares that he has no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Alston, W. (1971). Varieties of Privileged Access. *American Philosophical Quarterly*, 8, 223–241.
- Andrewes, W. J. (2002). A chronicle of timekeeping. *Scientific American*, 287(3), 76–85. <https://doi.org/10.1038/scientificamerican0902-76>
- Baars, B. J. (1996). Understanding Subjectivity: Global Workspace Theory and the Resurrection of the Observing Self. *Journal of Consciousness Studies*, 3, 211–216.
- Baars, B. J. (1997). In the Theatre of Consciousness. Global Workspace Theory, A Rigorous Scientific Theory of Consciousness. *Journal of Consciousness Studies*, 4, 292–309.
- Bell, A. E. (1941). The "Horologium Oscillatorium" of Christian Huygens. *Nature*, 148(3748), 245–248.
- Berkovich-Ohana, Aviva, Yair Dor-Ziderman, Fynn-Mathis Trautwein, Yoav Schweitzer, Ohad Nave, Stephen Fulder, and Yochai Ataria. 2020. "The Hitchhiker's Guide to Neurophenomenology – The Case of Studying Self Boundaries With Meditators." *Frontiers in Psychology* 11. <https://doi.org/10.3389/fpsyg.2020.01680>.
- BIPM. (2019). *The International System of Units* (9th ed.). BIPM.
- Carjaval, J. J., Romero, A. A., Cárdenas, A., Pazmiño, G. Z., & Herrera, P. A. (2012). Visual Anosognosia (Anton-Babinski Syndrome): Report of Two Cases Associated with Ischemic Cerebrovascular Disease. *Journal of Behavioral and Brain Science*, 2, 394–398.
- Cary, P. (2000). *Augustine's invention of the inner self: The legacy of a Christian Platonist*. Oxford University Press.
- Chalmers, D J 2003. "The Content and Epistemology of Phenomenal Beliefs." In *Consciousness: New Philosophical Perspectives*, edited by Q. Smith and A. Jokic. Oxford.
- Chang, H. 2007. *Inventing temperature: measurement and scientific progress*, *Oxford studies in philosophy of science*. Oxford ; New York: Oxford University Press.
- Chen, X. 2005. "Visual Photometry in the Early 19th Century: A "Good" Science with "Wrong" Measurements." In *Wrong for the Right Reasons*, edited by Jed Z. Buchwald and Allan Franklin, 161–183. Springer.
- Coninx, S. (2022). A multidimensional phenomenal space for pain: Structure, primitiveness, and utility. *Phenomenology and the Cognitive Sciences*, 21(1), 223–243. <https://doi.org/10.1007/s11097-021-09727-0>
- Das, J M., & I A. Naqvi. (2023). Anton Syndrome. [Updated 2023 April 3]. StatPearls [Internet]. Treasure Island (FL): StatPearls Publishing. <https://www.ncbi.nlm.nih.gov/books/NBK538155/#article-17722.s1>. Accessed 13 Sept 2023.
- Dehaene, S., & Naccache, L. (2001). Towards a cognitive neuroscience of consciousness: Basic evidence and a workspace framework. *Cognition*, 79(1–2), 1–37.
- Dennett, D. C. (1991). *Consciousness Explained*. Backbay Books.
- Dennett, D. C., & Kinsbourne, M. (1992). Time and the observer: The where and when of consciousness in the brain. *The Behavioral and Brain Sciences*, 15, 183–247.
- Eperjesi, F., & Akbarali, N. (2004). Rehabilitation in Charles Bonnet syndrome: A review of treatment options. *Clinical & Experimental Optometry*, 87(3), 149–152. <https://doi.org/10.1111/j.1444-0938.2004.tb03167.x>

- Frässle, S., Sommer, J., Jansen, A., Naber, M., & Einhäuser, W. (2014). Binocular Rivalry: Frontal Activity Relates to Introspection and Action But Not to Perception. *The Journal of Neuroscience*, *34*, 1738–1747. <https://doi.org/10.1523/JNEUROSCI.4403-13.2014>
- Gallagher, S. (2015). *A neurophenomenology of awe and wonder : Towards a non-reductionist cognitive science, New directions in philosophy and cognitive science*. Palgrave Macmillan.
- Geninas, S. (2016). Max Scheler's Phenomenology of Pain. *Frontiers of Philosophy in China*, *11*(3), 358–376.
- Gertler, B. (2012). Renewed Acquaintance. In D. Smithies & D. Stoljar (Eds.), *Introspection and consciousness* (pp. 93–128). Oxford University Press.
- Goldman, A. I. (1997). Science, publicity, and consciousness. *Philosophy of Science*, *64*(4), 525–545.
- Hardcastle, V. G. (1997). When a Pain is Not. *The Journal of Philosophy*, *44*, 381–409.
- Huygens, C. 2007 (1673). *Horologium oscillatorium sive de motu pendularium ad horologium aptato demonstrationes geometricae*. Translated and annotated by Ian Bruce. Paris: F. Muguet.
- Irvine, Elizabeth. 2019. "Developing Dark Pessimism Towards the Justificatory Role of Introspective Reports." *Erkenntnis*<https://doi.org/10.1007/s10670-019-00156-9>
- Jack, A, and A Roepstorff. 2003. *Trusting the subject? The use of introspective evidence in cognitive science*. Exeter, UK ; Charlottesville, VA: Imprint Academic.
- Jack, A., & Roepstorff, A. (2002). Introspection and cognitive brain mapping: From stimulus-response to script-report. *Trends in Cognitive Sciences*, *6*, 333–339. [https://doi.org/10.1016/S1364-6613\(02\)01941-1](https://doi.org/10.1016/S1364-6613(02)01941-1)
- Klein, H A. 1988. *The Science of Measurement. A Historical Survey*. New York: Dover Publications.
- Loar, B. 1997. Phenomenal States. In *The Nature of Consciousness. Philosophical Debates*, edited by Ned Block, Owen Flanagan, and Güven Güzeldere, 597–616. Cambridge MA: MIT Press.
- Ma, Y., Wang, C., Luo, S., Li, B., Wager, T. D., Zhang, W., Rao, Yi., & Han, S. (2016). Serotonin transporter polymorphism alters citralopram effects on human pain responses to physical pain. *NeuroImage*, *135*, 186–196. <https://doi.org/10.1016/j.neuroimage.2016.04.064>
- Maddula, Mohana, Stuart Lutton, and Breffni Keegan. 2009. "Anton's syndrome due to cerebrovascular disease: a case report." *Journal of Medical Case Reports* *3* (9028).
- Melzack, R. (1987). The short-form McGill Pain Questionnaire. *Pain*, *30*(2), 191–197. [https://doi.org/10.1016/0304-3959\(87\)91074-8](https://doi.org/10.1016/0304-3959(87)91074-8)
- Melzack, R., & Wall, P. D. (1983). *The Challenge of Pain. Completely* (rev). Basic Books.
- Michel, M. (2017). Methodological Artefacts in Consciousness Science. *Journal of Consciousness Studies*, *24*(11–12), 94–117.
- Michel, M, and J Morales. 2020. "Minority Reports: Consciousness and the Prefrontal Cortex." *Mind and Language*:1–21.
- Mouraux, A., & Iannetti, G. D. (2018). The search for pain biomarkers in the human brain. *Brain*, *141*(12), 3290–3307. <https://doi.org/10.1093/brain/awy281>
- Naber, M., S. Frässle, and W. Einhäuser. 2011. "Perceptual Rivalry: Reflexes Reveal the Gradual Nature of Visual Awareness." *PLOS ONE* *6* (e20910).
- Nisbett, R. E., & Wilson, T. D. (1977). Telling More Than We Can Know: Verbal Reports on Mental Processes. *Psychological Review*, *84*, 231–259.
- NIST. (2021c). SI Redefinition: Meter. Accessed, 05(01), 2022.
- NIST. 2021b. SI Redefinition. Ampere: History. Accessed 20.09.2021b.
- Ngamkham, S., Vincent, C., Finnegan, L., Holden, J. E., Wang, Z. J., & Wilkie, D. J. (2012). The McGill Pain Questionnaire as a multidimensional measure in people with cancer: An integrative review. *Pain Management Nursing : Official Journal of the American Society of Pain Management Nurses*, *13*(1), 27–51. <https://doi.org/10.1016/j.pmn.2010.12.00>
- Othman, Bayan Al, Andrew Go Lee, and Ashwini Kini. (2019). "Anton Syndrome." American Academy of Ophthalmology, Last Modified January 8, 2020. https://eyewiki.aaopt.org/Anton_Syndrome.
- Overgaard, M. (2010). How consciousness will change our view on neuroscience. *Cognitive Neuroscience*, *1*, 224–225. <https://doi.org/10.1080/17588928.2010.497585>
- Overgaard, M. (2015). The Challenge of Measuring Consciousness. In M. Overgaard (Ed.), *Behavioral Methods in Consciousness Research* (pp. 7–20). Oxford University Press.
- Overgaard, M., & Grünbaum, T. (2012). Cognitive and non-cognitive conceptions of consciousness. *Trends in Cognitive Science*, *16*(3), 137. <https://doi.org/10.1016/j.tics.2011.12.006>
- Overgaard, Morten, and Kristian Sandberg. 2021. "The Perceptual Awareness Scale—recent controversies and debates." *Neuroscience of Consciousness* 2021. <https://doi.org/10.1093/nc/niab044>.

- Piccinini, G. (2003). Data from Introspective Reports: Upgrading from Common Sense to Science. *Journal of Consciousness Studies*, 10, 141–156.
- Pronin, E. (2009). The Introspection Illusion. In M. P. Zanna (Ed.), *Advances in Experimental Social Psychology* (pp. 1–67). Academic Press.
- Pronin, E., Berger, J., & Molouki, S. (2007). Alone in a crowd of sheep: Asymmetric perceptions of conformity and their roots in an introspection illusion. *Journal of Personality and Social Psychology*, 92(4), 585–595. <https://doi.org/10.1037/0022-3514.92.4.585>
- Rosenthal, D. (2010). How to Think about Mental Qualities. *Philosophical Issues*, 20, 363–393.
- Sandberg, K., Timmermans, B., Overgaard, M., & Cleeremans, A. (2010). Measuring consciousness: is one measure better than the other? *Conscious and Cognition*, 19(4), 1069–1078. <https://doi.org/10.1016/j.concog.2009.12.013>
- Scheler, Max. (1954). Vom Sinn des Leidens. In Maria Scheler & Manfred S. Frings (Eds.), *Gesammelte Werke* (pp. 36–72). Bern: Francke.
- Scheler, Max. (1954). Zu einer philosophischen Lehre von Schmerz und Leiden. In Maria Scheler & Manfred S. Frings (Eds.), *Gesammelte Werke* (pp. 331–333). Bern: Francke.
- Schwitzgebel, E. (2002a). How Well Do We Know Our Own Conscious Experience. The Case of Visual Imagery. *Journal of Consciousness Studies*, 9(5–6), 35–53.
- Schwitzgebel, E. (2002b). Why Did We Think We Dreamed in Black and White? *Studies in History and Philosophy of Science*, 33, 649–660.
- Schwitzgebel, E. (2008). The Unreliability of Naive Introspection. *The Philosophical Review*, 117(2), 245–273.
- Smithies, D., & Stoljar, D. (2012). *Introspection and consciousness*, *Philosophy of mind*. Oxford University Press.
- Overgaard, M. (n.d.). Methodological reductionism or methodological dualism? The continued debate of the relation between subjective and objective measures of consciousness.
- Suplee, Curt, Jennifer Lauren Lee, Alison Gillespie, Gail Porter, Ben P. Stein, Bill Phillips, Jon Pratt, Mark Esser, and Chad Boutin. 2018. "A Turning Point for Humanity: Redefining the World's Measurement System." *National Institute for Standards and Technology*, accessed 09/20/21.
- Thompson, Evan, Antoine Lutz, and Diego Cosmelli. 2005. "Neurophenomenology: An Introduction for Neurophilosophers." In *Cognition and the brain: The philosophy and neuroscience movement.*, 40–97. New York, NY, US: Cambridge University Press.
- Tononi, G., Boly, M., Massimini, M., & Koch, C. (2016). Integrated information theory: From consciousness to its physical substrate. *Nature Reviews Neuroscience*, 17, 450–461.
- Tsuchiya, Naotsugu, Melanie Wilke, Stefan Frässle, and Victor Lamme. 2015. "No-Report Paradigms: Extracting the True Neural Correlates of Consciousness." *Trends in Cognitive Sciences* 19. <https://doi.org/10.1016/j.tics.2015.10.002>.
- Wager, T. D., Atlas, L. Y., Lindquist, M. A., Roy, M., Woo, C. W., & Kross, E. (2013). An fMRI-based neurologic signature of physical pain. *New England Journal of Medicine*, 368(15), 1388–1397. <https://doi.org/10.1056/NEJMoal204471>
- Woo, C.-W., Chang, L. J., Lindquist, M. A., & Wager, T. D. (2017). Building better biomarkers: Brain models in translational neuroimaging. *Nature Neuroscience*, 20(3), 365–377. <https://doi.org/10.1038/nn.4478>