

Article

Mathematical Models of Consciousness

Johannes Kleiner 

Munich Center for Mathematical Philosophy, Ludwig Maximilian University of Munich, Geschwister-Scholl-Platz 1, D-80539 München, Germany; johannes.kleiner@lmu.de

Received: 21 March 2020; Accepted: 23 April 2020; Published: 30 May 2020

Abstract: In recent years, promising mathematical models have been proposed that aim to describe conscious experience and its relation to the physical domain. Whereas the axioms and metaphysical ideas of these theories have been carefully motivated, their mathematical formalism has not. In this article, we aim to remedy this situation. We give an account of what warrants mathematical representation of phenomenal experience, derive a general mathematical framework that takes into account consciousness' epistemic context, and study which mathematical structures some of the key characteristics of conscious experience imply, showing precisely where mathematical approaches allow to go beyond what the standard methodology can do. The result is a general mathematical framework for models of consciousness that can be employed in the theory-building process.

Keywords: models of consciousness; theories of consciousness; experience space; phenomenal space; mathematical consciousness science; mathematical phenomenology; phenomenal consciousness; epistemic asymmetry; non-collatability

1. Introduction

Conscious experience and its relation to the physical domain have been studied by philosophers, theologians, and scientists over many centuries [1]. In the previous three decades, there has been a resurgence of scientific investigations. Groundbreaking developments in neuroscience, cognitive psychology, and analytic philosophy lead to the emergence of a dedicated *science of consciousness*, whose aim is to develop a scientific account of conscious experience and its relation to the physical domain (e.g., brain processes).

A model of consciousness is a hypothetical theory about how conscious experience and the physical domain relate [2]. Examples include global workspace theories [3,4], multiple draft theory [5], higher order thought theories [6], or integrated information theory [7], among many others. Models of consciousness complement metaphysical theories of consciousness, such as the various forms of functionalism, identity theories, interactive dualisms, or neutral monisms. These theories are concerned primarily with ontological questions and address the general type of relation between consciousness and the physical domain.

1.1. The Rising Importance of Mathematics in Consciousness Studies

In recent years, many models of consciousness have been proposed that are mathematical in nature. Primary examples in neuroscience are the recent versions of integrated information theory [7–9], which aim to determine the quality and quantity of a system's conscious experience using a complex mathematical algorithm [10,11], or predictive processing theory [12], which can be interpreted as specifying the content of conscious experience of a system using an advanced minimization principle [13]. However, promising models have been proposed by other disciplines as well, including philosophy [14], physics [15,16], mathematics [17–19], or psychology [20], all based on various different metaphysical ideas about consciousness.

These developments point at an increasing relevance of mathematical tools and methods in the scientific study of consciousness, much like in other scientific disciplines throughout the last century, with promising new insights on the horizon. However, mathematization on its own does not have unique scientific merit. Valuable progress can only be made if the mathematization is based on and integrates previous theoretical, empirical, and conceptual work.

1.2. What Makes Consciousness a Problem

Consciousness is a phenomenon unlike any other studied by natural science. It is unique as an object of investigation both in its characteristic features and in its epistemic context. This is true in particular for its most relevant and mysterious connotation, which much of the science of consciousness and also this article is concerned with, namely phenomenal consciousness. Phenomenal consciousness refers to the way in which the world *appears* to us, i.e., the way in which we experience the world. This can roughly be paraphrased as “pure subjective experience” [21].

Much of philosophy of mind is concerned with analyzing in detail just what the characteristic features of phenomenal consciousness are and how precisely they relate to metaphysical ideas and efforts of scientific investigation. The arguably most crucial features attributed to phenomenal experience are its essential *subjective nature*, which sometimes is taken to mean that phenomenal consciousness embodies a particular point of view [22], but also that some of its parts or properties seem *ineffable* [23], *private*, or *unavailable to cognitive and linguistic processing or communication* [24].

Basic properties or simple constituents of phenomenal consciousness are called *qualia*, but this term is being used with many different connotations to date. Qualia are variously claimed to be *intrinsic* and *non-relational* or to have a *qualitative and non-quantifiable* nature. Phenomenal consciousness is also claimed to be *directly or immediately apprehensible*, to be *transparent* in the sense that it appears as if we are in direct contact with the content of our conscious experience or to be *homogeneous* [21]. Various different connotations of all of these notions exist, and different philosophers endorse various combinations thereof.

Complementing these characteristic features of consciousness is its unique epistemic context, which comes about from the fact that phenomenal consciousness per se is accessible only to the experiencing system itself. Thus, in any scientific approach, there are two fundamentally different methodological approaches that allow one to gather information, a first-person perspective and a third-person perspective. This is referred to as the *epistemic asymmetry* of consciousness [25].

1.3. The Need for a Mathematical Foundation

Any scientific analysis that strives to address and explain phenomenal consciousness needs to take these features of consciousness into account, at the very least as providing epistemic restrictions that constrain and shape the empirical access to conscious experience. Failing to do so at all amounts to ignoring what makes consciousness a problem in the first place, which no serious scientific investigation can afford.

To date, almost none of the existing formal models take any of these properties of consciousness into account. While mathematical structures are quickly associated with terms like “qualia”, “subjective experience”, or “act of consciousness”, contemporary models fall short of actually considering the conceptual meaning of these philosophical concepts.

What is necessary to mend this is a thorough foundation of mathematical models of consciousness that analyzes which implications the various characteristics of conscious experience have on the mathematical structure of these models and which provides a precise account of how the concepts developed in philosophy of mind relate to the mathematical structure of models of consciousness.

The goal of the work presented in this article is to provide this foundation for the case of ineffability, privateness, and cognitive, linguistic, and communicative unavailability.

1.4. A Framework for Formulating Models of Consciousness

The result of this work is a general mathematical *framework* in which models of consciousness can be formulated. Much like Lagrangian mechanics in theoretical physics, it does not provide any particular law or equation that constitutes a model of consciousness, but rather a general formal machinery. What this machinery achieves is to take into account properly that conscious experience has ineffable, private, or inaccessible aspects and that it exhibits an epistemic asymmetry. This framework provides a first mathematical foundation for models of consciousness and needs to be expanded in future work to take other key characteristics into account.

Crucially, the framework is independent of whether one considers any of these characteristic features to be ontological in origin or simply due to a system's particular design or cognitive functions [5]. What matters, from the perspective of this framework, are only the epistemic restrictions that arise from these features of conscious experience, i.e., that access to some parts of conscious experience is limited by consciousness' subjective nature and by ineffability, privateness, and inaccessibility in any type of experimental situation.

Thus, this framework can be dubbed *operational*. Much like quantum theory in its conventional formulation, it takes as its starting point the prototypical experimental situation in which a theory (of consciousness) is being tested, used, or inferred and then adds the particular epistemic context of consciousness, so as to arrive at a general operational description. Great care has been taken to keep the mathematical structure of this formalism as general as possible and to provide operational justifications of all essential definitions, so as to ensure that the framework is *compatible* with all types of mathematical structure one would want to use in modeling consciousness, including category theory [11,26], information theory [27], or complex system approaches [28], among many others.

1.5. An Axiomatic Conceptual Underpinning

In order to translate any concept into formalism, the concept itself needs to be rigorously defined. To date, neither the term qualia, nor the concepts of ineffability, privateness, and inaccessibility are defined in a rigorous enough manner to warrant thorough formalization.

Thus, in order to achieve our goal, it is necessary to represent the underlying philosophical concepts in an axiomatic form that is suitable for formalization. Since this whole program is operational in nature, it suffices in fact to provide an axiomatic definition of the *operational consequences* of these characteristic features of consciousness.

For the case of ineffability, privateness, and inaccessibility, this is possible at once by introducing the concept of *non-collatability*: A part, property, or feature of conscious experience is non-collatable if there do not exist any reasonable means to identify this part, property, or feature over several experiencing subjects in an experimental trial. Non-collatability is entailed by ineffability, privateness, and inaccessibility and arguably also by consciousness' subjective nature. As shown in detail throughout this work, it is precisely non-collatability that generates much of the epistemic difficulty in investigating conscious experience and that has substantial consequences for any empirically adequate model of consciousness.

The conceptual definitions we derived in order to found the mathematical structure of models of consciousness give rise to an axiomatic *grounding* of the scientific study of consciousness that is an alternative to and further development of the grounding that derives from David Chalmers' work. While our grounding was primarily intended to constitute an intermediate construction that links philosophical concepts and mathematical formalism, it may also have some conceptual worth in its own, providing an interim way to conceive of the task and methodology of the scientific study of consciousness from a more formal perspective.

1.6. A New Way of Consciousness Science

In summary, this article can be understood as taking seriously a new way of doing consciousness science that has been pioneered in eminent works such as [7,20], or [29]. Its central idea is to represent phenomenal consciousness in terms of mathematical spaces and to use these spaces to build theories of how conscious experience might relate to the physical domain. This facilitates a much richer and refined way of addressing conscious experience and offers promising tools to resolve some of the key issues that permeate contemporary consciousness studies.

What this article adds to this new methodology is the requirement that the essential features of conscious experience, studied in detail by philosophy of mind, are taken into account when building this mathematical representation. Doing so requires an account of how precisely mathematical spaces can be grounded in the phenomenology of experience and of which mathematical implication consciousness' fundamental epistemological context has. Answers to all of these questions are proposed here. The hope is that these considerations might provide a useful basis for further development of formal models of consciousness.

1.7. The Structure of this Article

This article is structured as follows. In order to make it accessible to readers without a formal background, we first summarize the results in Section 2, keeping mathematical details to a minimum. In this section, we also explain in detail the rationale and motivation of this work.

All subsequent sections aim for a concise presentation of definitions, explanations, and examples. In Section 3, we give the conceptual definitions on which our framework rests, making as few assumptions as necessary. This gives rise to an axiomatic grounding of the scientific study of consciousness. In Section 4, we show that there is an explanatory gap between qualia as defined here and natural science. Section 5 is devoted to deriving a general mathematical framework for formal models of consciousness, making use of a minimal set of ingredients of any formal theory and of consciousness' epistemic asymmetry. In Section 6, finally, we show how consciousness' characteristic features can be taken into account. We conclude this paper with a brief remark on a metaphysical question in Section 7 and various examples in Section 8.

In Appendix A, we review the grounding of the scientific study of consciousness that derives from David Chalmers' work in [30,31], emphasizing the logical relations among its parts. In Appendix B, we discuss problems that arise if one attempts to apply this grounding in a model-building process.

2. Summary of Results

Any research activity directed at conscious experience presupposes a conception of the phenomenon that is to be studied and a conception of a methodology that is suitable to do so. We call this a grounding of the scientific study of consciousness.¹

Definition 1. A grounding of the scientific study of consciousness *contains at least:*

- an explicit definition of what is to be studied.
- an explicit outline of the methodology.

Much of the research devoted to consciousness in the previous two decades has been guided by a grounding that derives from David Chalmers' work [30]. This grounding has played a pivotal role in the creation and consolidation of the field. However, it also exhibits several severe problems when being applied (Appendix B). In the first part of this paper, we introduce an alternative to Chalmers' grounding of the scientific study of consciousness. This alternative is built on a *thoroughly operational perspective*, which means that we define all notions relative to prototypical experimental investigations.

¹ The term "grounding" is one of several translations of the German word "Grundlegung".

Any experimental situation devoted to study conscious experience presupposes a preliminary choice of organisms that are considered to be conscious and whose conscious experience and physical state are probed during the experiment in order to gain information about how consciousness relates to the physical domain. We denote any such class of organisms by \mathcal{C} and call them *experiencing subjects*. Taking the operational perspective seriously, we consider \mathcal{C} to be a primitive notion. While it may be guided by theoretical insights and changed over the course of time, at any particular time, a class \mathcal{C} provides the basis for both inference and tests of theories about consciousness.

Having chosen our primitive notion, we can define experience relative to it. A promising choice is to use somewhat phenomenological terminology in defining the term conscious experience, referring to the *totality* of how experience “reveals itself” to an experiencing subject, how the experiencing subject finds itself experiencing, or how the “the world” appears to it. While this is what we have in mind, we have opted for a more approachable terminology and define the term *conscious experience* to denote the totality of impressions, feelings, thoughts, perceptions, etc., that an experiencing subject lives through at a particular instant of time (Definition 2). Experience so defined has various different *aspects*, where we define the term “aspect” to be a placeholder for any conception like “part”, “property”, or “element of” (Definition 3).

The key notion on which our grounding is built is that of non-collatability. An aspect of experience is *non-collatable* if there is no reasonable method to identify whether two or more experiencing subjects in an experiment experience this very aspect of experience; in other words, if the identity of this aspect over several different experiencing subjects in \mathcal{C} cannot be determined (Definition 5).

The distinction between collatable and non-collatable aspects of experience is what replaces the distinction between phenomenological and psychological concepts of mind in [30]. Whereas the latter distinction is defined in terms causal roles and spatiotemporal structure (cf. Appendix A), our distinction is defined axiomatically in terms of phenomenological or operational notions.

What is crucial is that non-collatability is implied by various essential characteristic features of conscious experience, e.g., any aspect that appears to be *ineffable* (i.e., that is *experienced* as ineffable) is also non-collatable in the above sense. The same is true for aspects of experience that are experienced as *private* or that are *not available to cognitive or linguistic processing*. All of these characteristic features destroy the possibility to identify an aspect under consideration over several experiencing subjects. Non-collatability is a necessary operational consequence of all of these characteristic features.

The same may be true of the subjectivity of conscious experience, if taken to warrant the claim that “there are facts that do not consist in the truth of propositions expressible in a human language” [22] (p. 441). In fact, one of the main claims in [22] is that there is at present no conception that allows one to establish the identity of a “what it is like” aspect of experience with a physical state. Our starting point, non-collatability, is closely related to this claim and may even be implied by it in reasonable cases.

Building on non-collatability, we define qualia as follows (cf. Definition 6).

Definition 6 We define the term *qualia* to refer to all non-collatable aspects of experience of an experiencing subject within the class \mathcal{C} .

This definition is warranted since it includes the paradigmatic examples of what qualia are claimed to be (Example 4), as well as aspects of experience referenced by the Nagelian “what it is like” conception (Example 5). It is furthermore axiomatic and replaces the concept of phenomenal consciousness as defined in Chalmers’ grounding. (We remark that for this and all other definitions, it does not matter whether non-collatability or any of the characteristic features that imply it are considered to be fundamental or merely the result of a system’s architecture. All that matters is that experience appears as such.)

The aspects of experience that satisfy Definition 6 are of special interest because the non-collatability induces a fundamental difficulty in any scientific approach: it implies that these aspects cannot be referenced intersubjectively, which in turn implies that they cannot be referenced

in a scientific model or empirical analysis. There is a fundamental explanatory gap (Section 4). The goal of this paper, when put in these terms, is to develop a mathematical framework that allows us to address both collatable and non-collatable aspects of experience, providing a formal methodology suitable to address this explanatory gap.

Next, we make use of the central idea that underlies many contemporary mathematical models of consciousness: to represent phenomenal consciousness as a mathematical space. In order to provide an accurate method to do so, we make use of two phenomenological axioms.

First, we make use of the fact that both collatable and non-collatable aspects of experience can be recognized to a certain extent (Phenomenological Axiom 2). Another way to say this is that aspects of experience may be experienced as identical. Following our operational perspective, this warrants the introduction of *labels* for both qualia and collatable aspects of experience, i.e., names relative to an experiencing subject.

Second, we make use of the fact that there are collatable relations between aspects of experience (Phenomenological Axiom 3). This might be considered obvious in the case of collatable aspects of experience. With respect to non-collatable aspects of experience, it corresponds to the observation that “structural features of perception might be more accessible to objective description, even though something would be left out” [22] (p. 449), or that “even if experiences are in some sense ‘ineffable,’ relations between experiences are not; we have no trouble discussing these relations, whether they be relations of similarity and difference, geometric relations, relations of intensity, and so on. As Schlick [32] pointed out, the *form* of experience seems to be straightforwardly communicable, even if the *content* (intrinsic quality) is not” [30] (p. 224).

Together, these two phenomenological observations allow us to define a mathematical space that *represents* conscious experience, which we call *experience space E*. The elements of this space are not experiences themselves, but *labels* that an experiencing subject may give for his/her aspects of experience, and the mathematical structure on this set of labels is induced by the collatable relations between aspects of experience. To conceive of *E* as a space of labels, rather than as a space of experiences, is of advantage because it prevents from the very beginning any implicit assumption of well-defined reference to aspects of experience. In contrast, working with a space whose elements are intended to express experiences themselves requires the introduction of a map that describes how these experiences can be inferred from reports (labels), e.g., as in [33]. The details of our introduction of experience space are explained in Section 3.2, and various examples are given in Section 3.5.

We remark that whereas our constructions are guided by conceiving of labels as something that an experiencing subject can express, which requires \mathcal{C} to comprise humans, this is not necessary. This is because the various principles that are used in experiments to date to infer the state of consciousness of some subject (e.g., button presses or behavioral indicators) are, in our terminology, in fact means to infer labels of aspects of experience. Whether a label is a recorded word or some other type of report, such as a particular movement, does not matter for our purposes. What is crucial about the terminology of labels is that one avoids from the very beginning any implicit assumption that there is an empirically well-defined method to refer to the qualia of an experiencing subject.

Non-collatability implies limitations on how aspects of experience can be referenced in a theory or empirical investigation. Whereas labels of collatable aspects of experience can be synchronized over all experiencing subjects in the class \mathcal{C} (because collatability holds iff there exist means to identify), labels of qualia cannot. In virtue of non-collatability, the definition of qualia implies that any label that one experiencing subject uses to denote a quale may denote another quale in a different experiencing subject. Any scientific investigation that aims to address qualia needs to take the resulting ambiguity into account. Ignoring it will lead to errors, such as the study of the wrong “information pathway” or confounding neural correlates of external signals with neural correlates of qualia.

In the next step of our construction, we quantify this ambiguity precisely. To this end, we make use of the mathematical representation *E* of experience constructed previously. As we explained in detail in Section 3.3, the conceptual and mathematical definitions imply that the ambiguity of any

reference to aspects of experience can be stated concisely in terms of the automorphism group $\text{Aut}(E)$ of E , i.e., the group of all transformations of E that change labels in such a way that the mathematical structure of E is left invariant.

We find that any statement about conscious experience that uses an individual label e could, in light of non-collatability, have equally well be formulated in terms of any label e' that is part of a subset $[e]$ of labels. This subset $[e]$ is called the equivalence class of the label e with respect to the automorphism group $\text{Aut}(E)$.

The crucial insight here is that these equivalence classes describe what is intersubjectively accessible or, in our terminology, empirically well-defined. Taken together, these equivalence classes describe what is amenable to the usual scientific methodology. Depending on the mathematical structure of E , this may well include some of the non-collatable aspects of experience.

This second step in our constructions may be sufficient for many investigations once experimental tools become advanced enough to proceed to the study of individual aspects of experience. It enables experimentalists to use structural features of phenomenal experience, as represented in E , to push the boundary of ineffability and all the other characteristic features that imply non-collatability back a little bit. However, as long as there are equivalence classes $[e]$ that contain more than one label, there are questions that evade the reach of the standard scientific methodology: why the subject had one of the corresponding experience rather than the other. Even though this cannot be expressed by intersubjective means, there is a fact to the matter for the experiencing subject, and hence a priori an open scientific question. The goal of the remaining part of the article is to develop tools that allow us to address this open question.

To develop these tools, we have to go beyond the mathematical representation of experience and in fact consider formal hypotheses about how conscious experience relates to the physical domain, i.e., formal models of consciousness. In order to remain as general as possible, as a first step in answering this question, we ask what the most general mathematical structure is that a model of consciousness needs to address.

In order to answer this question, we first give an account of the minimally sufficient formal structure of any scientific theory (Section 8). A theory needs to specify some dynamical variables d that describe what the theory intends to address, may contain some formal background structure, needs to have a parameter such as time that facilitates the description of variations of the dynamical variables, and finally, needs to contain some laws that pick out some variations of d from all possible variations.

In order to further fix the dynamical variables d , we make use of the epistemic asymmetry of conscious experience (Section 5.2). The epistemic asymmetry states that there are two fundamentally different ways of gathering knowledge about conscious experience, the first-person perspective and the third-person perspective. Thus, there are two *epistemically* different notions of state in any experimental situation, one that corresponds to first-person access and one that corresponds to third-person access. While any metaphysical theory of consciousness can ignore one of these states, a scientific model of consciousness cannot. The difference between a coherent idea and a scientific model of consciousness is precisely that the latter addresses both types of states, while the former need not.

Since the states that are accessible in the third-person perspective are in fact physical states (neural states, brain states, or similar) and the states that are accessible in the first-person perspective are aspects of experience (with “aspect” suitably defined, cf. above), this implies that the dynamical variables of a formal model of consciousness are in fact a subset of

$$d = E \times P,$$

where E denotes the mathematical representation of conscious experience we have introduced above and where P denotes the state space of some physical theory T_P .

Combining the above gives a general framework in which models of consciousness (Definition 10) can be formulated. It provides a reference to which models of consciousness *need to refer* in light of

consciousness' epistemic context, independently of how they are primarily defined and independently of which ontological ideas they express.

This general framework finally puts us into a position to investigate the implications of non-collatability in Section 6. First, in Section 6.1, we prove that in light of non-collatability, models of consciousness are only well defined if they carry a particular symmetry. This is comparable to physical theories. Much like general relativity carries a particular symmetry that ensures that the theory is well defined with respect to changes of coordinates, our results show that models of consciousness need to carry a particular symmetry that ensures that they are well defined with respect to changes of labels. The changes of labels in question are precisely those transformation that keep the equivalence classes $[e]$ introduced above constant, but transform individual members of these equivalence classes. The corresponding symmetry group is the automorphism group $\text{Aut}(E)$ introduced above.

What is crucial in our results of Section 6.1 is that the symmetry required to exist is not fixed uniquely. There is a freedom in its form that depends on the laws of a model of consciousness. This freedom describes how the transformations of labels relate to transformations of physical states. Sections 6.2 and 6.3 are devoted to proving that this remaining freedom is what allows formal models of consciousness to go beyond what the standard methodology can do.

In a nutshell, this is so because the standard methodology can only utilize intersubjectively well-defined references. Mathematically, this means that the standard methodology can only reference aspects of experience once it has imposed the group $\text{Aut}(E)$ in order to construct equivalence classes $[e]$. Formal models of consciousness, on the other hand, allow one to reverse this order. They allow one to relate individual elements of E and P prior to imposing the symmetry that ensures well definedness.

Since all our arguments, proofs, and derivations hold true also in the limiting cases where all aspects of experience are either collatable or non-collatable, we summarize all insights in a concise definition of what a model of consciousness is (Definition 12). This is the main result of our project.

“Many scientific discoveries have been delayed over the centuries for the lack of a mathematical language that can amplify ideas and let scientists communicate results.” [34] (p. 427)

3. Basic Definitions

In this section, we provide the basic definitions that underlie our constructions. In Section 3.1, we specify the notion of experience we consider, introduce some fundamental terminology, and give a definition of qualia in these terms. Subsequently, in Section 3.2, we discuss the mathematical representation of experience. In Section 3.3, we explain the implications of the defining characteristic of qualia for any reference to consciousness in an experiment or theory. As mentioned above, altogether, this can be taken to provide a grounding of the scientific study of consciousness, and Section 3.4 is devoted to summarizing the resulting picture. In Section 3.5, finally, we give several examples for the mathematical structure introduced in Sections 3.2 and 3.3.

3.1. Conscious Experience and Qualia

The starting point of every *scientific* activity related to consciousness is a preliminary choice of a class \mathcal{C} of experiencing subjects that are available for experimental investigations and that are targeted by theoretical models. The object of investigation of any empirical study, and what informs any model-building process, is the experience of these experiencing subjects in the following sense.

Definition 2. We use the term conscious experience (experience for short) to denote the totality of impressions, feelings, thoughts, perceptions, etc., that an experiencing subject lives through at a particular instant of time.²

The general idea underlying any conception of the scientific study of consciousness is to study experience and its relation to the physical domain by scientific means. Mostly, some part or feature of experience is under consideration. In order to emphasize that this part or feature may not be strictly separable from other parts of features, we use the term “aspects of experience”:

Definition 3. Aspects of experience denote specific or general features, parts, properties, or elements of a particular experience or of a set of experiences.

According to this definition, “aspect” is thus merely a placeholder for “feature”, “part”, “property”, or “element of”. Which of these notions is relevant is part of the specification of a model of consciousness.

Example 1. Aspects of experience range from individual visual, auditory, or tactile experiences to general characteristics, such as the experience of a first-person perspective, the unity of the conscious scene [2], or the structure and composition of experience [7].

A priori, every experiencing subject only has access to his/her own experience. However, systematic investigations of which aspects of experience are invariant over a large class of experiencing subjects are possible and have been carried out as part of the philosophical discipline of phenomenology.

Definition 4. A phenomenological axiom is a statement about aspects of experience that holds for all experiencing subjects in a class C .³

Phenomenological axioms serve as a starting point for any investigation in the scientific study of consciousness. In empirical studies, they are what can be correlated with physical states, e.g., to construct neural correlates of consciousness. When building models of consciousness, they are what informs the choice of mathematical structure. In simple terms, phenomenological axioms are statements about how experiencing subjects find themselves experiencing, or how “the world” appears to them. One could also say that they express invariant facts of “what experience is like” or of how experience “reveals itself”.

Next, we make use of three basic phenomenological axioms in order to examine in more detail which methodology may be used to study aspects of experience and their relation to the physical domain. In preparation, we define the concept of non-collatability.

Definition 5. An aspect of experience is non-collatable iff there does not exist a reasonable method to establish its identity over different experiencing subjects in the class C .⁴

² No special focus on subjectivity is intended when using the term “experiencing subject”. Alternatively, one could use the term “experiencer”. We also remark that the meaning of “instant” is to be fixed during the model-building process. It could refer to physical just as well as to experiential instants of time.

³ The restriction to a class C of experiencing subjects is necessary because a phenomenological analysis of invariants of experience is always restricted to experiencing subjects that are similar in some respects: “[O]ne person can know or say of another what the quality of the other’s experience is. [However, this] ascription of experience is possible only for someone sufficiently similar to the object of ascription to be able to adopt his point of view” [22] (p. 442). However, the choice of class C is not a constraint for models of consciousness, but rather a starting point, i.e., a preliminary choice that informs the model-building process. Models may eventually allow one to determine which organisms experience. We note also that the name “phenomenological axiom” is a tribute to phenomenology rather than an attempt to condense the phenomenological method into a simple definition.

⁴ I.e., an aspect experienced by subject S_1 is non-collatable iff there is a different experiencing subject S_2 such that there is no reasonable method to determine with which aspect e' of S_2 the aspect e of S_1 is identical. Put yet in different terms, this is

The non-collatability of an aspect of experience can be determined operationally in any experimental situation. Whenever there is no reasonable method to identify whether two experiencing subjects in an experiment experience the same aspect of experience, the aspect in question is non-collatable. However, the concept can also be applied in a more fundamental context, as part of a phenomenological axiom about how subjects experience the world. Several examples are given below.

Phenomenological Axiom 1. *Aspects of experience can be divided into two classes:*

- (a) *Aspects of experience that are non-collatable.*
- (b) *Aspects of experience that are collatable.*

The former class includes aspects of experience that are experienced as *ineffable* [23], *private*, or found to be *cognitively, linguistically, and communicatively inaccessible* [24]. However, they also include those that are referred to as having a *subjective character* [22] (p. 437) or connected to a *particular point of view* [22] (p. 441). Non-collatability is implied by, and hence a necessary condition of, all of these characteristic features of experience. The latter class includes those aspects that are experienced as accessible also *from other points of view* [22] (p. 443) or as having an *objective nature* [22] (p. 443).⁵

Example 2. *Consider, as a first example, experiences of awe. Subjects may report that they have an experience of awe and even give labels to various different such experiences, but there is, at present, no methodological procedure to establish whether any two experiences of awe of two different subjects are the same or not.*

Some [5,35] may hold that advanced neuroscientific theories may provide means to collate the experiences of awe eventually. However, as we will see below, collatability is a prerequisite of any theory that addresses specific aspects of conscious experience. If an aspect of experience is non-collatable, no theory can be empirically inferred or tested that addresses this aspect of experience.

Example 3. *Similarly, there is at present no possibility to ask meaningfully the question of whether color experiences of two experiencing subjects are the same or different. This simple but important fact is pointed at by the plain question of how two experiencing subjects might come to conclude that the experience of color that they have if they look at, e.g., the clear sky is the same. They may ensure that they use the same reference (“blue”) for the experience, that they see the same wavelength, and they might even be able to conclude that similar neuronal assemblies are active in both of their brains while having the experience in question. However, none of this is a priori related to the color aspects of their experiences (“what it is like to see blue”).*

Put differently, there is no reasonable way to assign truth values to statements of the form “my color experience e_1 is equal to your color experience e_2 ”; equality is not a well-defined concept when referencing experiences of two different experiencing subjects. Thus, color experiences are non-collatable aspects of experience in the sense of Phenomenological Axiom 1.

This non-collatability has consequences for any scientific account of color experience, e.g., any hypothesis that a particular neural activity occurs whenever a subject is experiencing the color “green” is not well defined, simply because there is no intersubjectively meaningful reference to “green”; the color experience one subject is having when presented a 510nm light source may be very different from the color experience another subject is having when presented the same light source. In other words, any intersubjective reference to color experiences carries a certain ambiguity, which has to be taken into account when constructing models or designing experiments related to color experience.

the case if there is no mapping from e to the aspects of experiences of other experiencing subjects that can reasonably be interpreted as establishing the identity of aspects.

⁵ When being presented with Phenomenological Axiom 1, scientists usually tend to think about how this can be derived from a theory of language. In our opinion, the more important task is to ground the underlying distinction in a thorough phenomenological analysis. We also remark that all formal constructions in this article are compatible with either of the classes in Phenomenological Axiom 1 being empty, even though this is most likely not the case.

The main point of this paper, argued for in detail below, is that non-collatable aspects of experience cannot be addressed by the usual scientific methodology. Since the term “qualia” is generally used to denote what is considered as essential in a particular analysis of experience, we introduce the following abbreviation.

Definition 6. We define the term qualia to refer to all non-collatable aspects of experience of an experiencing subject within the class C .

Example 4. According to Example 3, color experiences satisfy the condition of Definition 6. Thus, color experiences are qualia.⁶

Example 5. Example 4 is a special case of the aspects of experience referenced by Thomas Nagel in [22] when introducing his famous notion of “What is it like to be ... ?”:

“[F]undamentally an organism has conscious mental states if and only if there is something that it is like to be that organism—something it is like for that organism. We may call this the subjective character of experience.” (p. 436)

Nagel also uses the term “how it is for the subject himself” (p. 440) to point to these aspects of experience. Though [22] does not make the distinction of Phenomenological Axiom 1 central to his line of reasoning, one can find hints toward this distinction in [22]: e.g., he claims that “we do not possess the vocabulary to describe [what it is like to be us] adequately” (p. 440), there are “facts that do not consist in the truth of propositions expressible in a human language.” (p. 441)

3.2. Formal Representation of Experience

In order to define a formal representation of experience, we make use of two further basic phenomenological axioms. These are very general in nature, and it is plausible that they hold independently of the particular choice of class C . However, due to the restricted possibility of phenomenological analysis mentioned above, we generally assume C to comprise adult humans. The first phenomenological axiom expresses the observation that some qualia are *experienced* as identical, whereas others are not, or in other words, that one sometimes experiences a non-collatable aspect as identical to a non-collatable aspect one has experienced at another time.

Phenomenological Axiom 2. Qualia can be recognized to a certain extent: experiencing subjects can identify qualia that they have previously experienced.

Example 6. Phenomenological Axiom 2 states that experiencing subjects may perceive some aspects they experience at different times to be identical. For example, it could be the case that someone finds the taste aspect experienced when trying artificial strawberry flavor to be identical to the taste aspect experienced when eating an actual strawberry. This recognition of previously experienced aspects is simply a “subjective impression” of identity, so to speak.

Phenomenal Axiom 2 is important because it is the basis of the ability of an experiencing subject to introduce labels for his/her qualia, i.e., a name or reference for non-collatable aspects of his/her experience. Recognizability is presupposed in the notion of collatability, so that labels of collatable aspects of experience can be introduced by definition.

In what follows, we assume that labels are chosen such that different aspects of experience are associated with different labels and, using Phenomenological Axiom 2, that the same label is used to

⁶ We generally abbreviate “color aspects of experience” by “color experience”.

denote various occurrences of the same aspect.⁷ Furthermore, we assume that all experiencing subjects use the same set of labels, which we denote by E . For our purposes, E can be any set; which labels the set consists of does not matter in what follows.

The second phenomenological axiom expresses the observation that something can be said about how non-collatable aspects occur in, or constitute, experience. In [22], it corresponds to the observation that “structural features of perception might be more accessible to objective description, even though something would be left out” [22] (p. 449). In [30], it corresponds to the observation that “even if experiences are in some sense ‘ineffable,’ relations between experiences are not; we have no trouble discussing these relations, whether they be relations of similarity and difference, geometric relations, relations of intensity, and so on. As Schlick [32] pointed out, the *form* of experience seems to be straightforwardly communicable, even if the *content* (intrinsic quality) is not” [30] (p. 224).

Phenomenological Axiom 3. *Qualia have relations that can be collated within the class C .*

By Definition 4, this is a claim about experiences of all experiencing subjects in the class C . In simple terms, it expresses the fact that something can be said about non-collatable aspects of experience, something about how they appear in experience. The collatability of the relations implies that we may *represent the relations on the set of labels E* and assume (i.e., ask, cf. Footnote 7) labels to be chosen in such a way that the experienced relations between qualia are reflected in the relations represented on the labels. We assume that Phenomenological Axiom 3 also holds for collatable aspects of experience, so that they have relations, also, that can be represented on the set of labels.⁸

Example 7. *For qualia of the “what it is like to be” type (introduced in Example 5), these relations include:*

- ▶ Similarity: *Two qualia can be more or less similar;*
- ▶ Intensity: *A quale can occur in more or less intense versions;*

*among others.*⁹

Example 8. *Experiencing subjects typically experience some pairs of colors as similar to each other, whereas they experience others as not similar, e.g., small changes in hue usually result in colors that are perceived as similar, whereas large changes in hue result in colors that are not experienced as similar.*

What is crucial for our purposes is that one may (and in practice, often does) represent the experience of the similarity of colors on the set of color labels. Correspondingly, one may (and in practice, often does) ask experiencing subjects to choose labels for their color experiences in such a way that colors that are experienced

⁷ Note that throughout this section, assumptions are in fact conventions. E.g., this assumption can be satisfied by asking experiencing subjects (in an experiment, say) to choose labels as described. The assumptions can be made “without loss of generality”, so to speak.

⁸ Phenomenological Axiom 3 states that there are relations between qualia that are collatable. This expresses the observations in [22,30,32] that structural features of perception, relations between experiences, or the form of experience might be more accessible to communication or objective description. However, one might question whether this axiom is warranted and insist that relations between experiences are not (strictly, at least) collatable (thanks to an anonymous referee for pointing this out). The formalism developed here requires the collatability of relations, so that any non-collatable relation has to be ignored.

⁹ Similarity and intensity are simple examples of collatable relations between qualia. There may be many more collatable relations that express facts about how qualia appear in experience, some of which may only relate qualia of a particular type to each other. Further examples arguably include: *Composition*: Some qualia are experienced as a composition of two (or more) different qualia, i.e., the composed quale is but a combination (or simultaneous experience) of the composing qualia. *Inclusion*: Some qualia may be experienced as containing one (or more) other qualia. Here, the contained quale is but an aspect of the containing quale. Furthermore, the distinction between various *types* (visual, auditory, tactile, etc.) of non-collatable aspects of experience is a relation in the sense of Phenomenological Axiom 3.

as similar are similar according to the representation on the color space.¹⁰ We will study this in detail in Example 9 below.

Phenomenological Axiom 2 provides the possibility to introduce labels for non-collatable and collatable aspects of experience. What Phenomenological Axiom 3 adds to this is the possibility to represent relations between aspects of experience on the set of labels. Since any representation of a relation on a set is mathematical in nature, so are these representations. They give either relations on E in the mathematical sense of the word (i.e., a subset R of $E \times E$) or some more involved mathematical structure, which turns E into a mathematical space.¹¹

Thus, together, these two phenomenological axioms ground a representation of experience in terms of mathematical structure. We refer to the set of labels E together with its mathematical structure that represents relations between qualia as

$$\text{Experience space } E, \tag{1}$$

though it is important to keep in mind that this space does not describe experience per se, but only labels and the structural relations between aspects of experiences they represent. This space E is the *mathematical representation of experience* mentioned above. Every element $e \in E$ refers either to a collatable aspect of experience or to a quale. Several detailed examples are given in Section 3.5 below.

In order for the mathematics to come out right in what follows, we have to introduce an important mathematical convention with respect to collatable aspects of experience. By Definition 5, an aspect of experience is collatable if its identity over all experiencing subjects in the class \mathcal{C} can be established. This implies, in particular, that this aspect of experience can be referenced: by virtue of its collatability, it can be assigned a unique label used by all experiencing subjects in \mathcal{C} . Our convention, in what follows, is that this is represented in the mathematical structure of E .

Convention 1. We assume that for every collatable aspect of experience, the mathematical structure of E contains a unary collatable relation χ that allows one to select this aspect of experience uniquely.¹²

In practice, this means that for any $e \in E$ that is collatable, there is a subset $\chi_e \subset E$ that contains only e . This convention ensures that changes of labels, discussed next, can be represented conveniently using the automorphism group. It ensures that all the collatable information is represented in the relations between aspects, so that all aspects can be treated alike in the technical definitions that follow.

In summary, so far, we have constructed a space E whose elements denote aspects of experience (both collatable and non-collatable ones), e.g., phenomenal properties or elements of experience. Furthermore, this space carries relations or more advanced mathematical structures that express the structural features of experience, as well as information about which aspects of experience are collatable (by virtue of Convention 1). This allows us to give a concise account of references to aspects of experience that takes non-collatability into account, as we explain next.

3.3. References to Qualia

By virtue of non-collatability, any reference to qualia is ambiguous. In this section, we explain in detail why and, in doing so, develop formal tools that allow us to quantify this ambiguity precisely.

¹⁰ Note that this example is complicated by the fact that we calibrate color experiences in practice: we apply or learn rules on how to pick color labels related to external events such as wavelength impinging on the eye. This will be discussed in detail in Example 9 below. What is crucial is that a priori, individual labels so chosen do not correlate with color experience: two experiencing subjects may have a completely different color experience despite using the same label “blue”.

¹¹ Here, the term “mathematical space” is used to refer to a set that carries additional mathematical structure. Examples are metric spaces, topological spaces, vector spaces, differentiable manifolds, principal bundles, measurable spaces, and Hilbert spaces.

¹² A unary relation on E is simply a subset of E .

We proceed in two steps. First, we discuss the case where an experiencing subject uses labels to report on his/her experience without taking into account any of the collatable relations. This is a preparatory step whose purpose is to explain the following constructions in detail. Since it ignores the relations on E , i.e., structural features of experience, it is artificial and will give a pathological result. Subsequently, in the second step, we discuss the appropriate case that takes the mathematical structure of E into account.

3.3.1. Preparation: References that Ignore Relations

Let us assume that an experiencing subject uses labels to report on his/her experience without taking into account any of the collatable relations. In this case, the experiencing subject is free to choose any label to denote any aspect of experience, the only requirements being that different labels are used for different aspects and that the same label is being used for a recurrent aspect. We call a choice of labels of an experiencing subject to denote his/her experienced aspects a *labeling* and use the term *relabeling* to denote a change of labeling. In the present case, a relabeling is simply a map

$$s : E \rightarrow E, \quad e \mapsto s(e), \quad (2)$$

which determines which label $s(e)$ replaces the previous label e . Since different aspects are required to carry different labels, this map is injective. Since it furthermore has domain and codomain E , it is bijective. Since any composition of two relabelings of the form (2) yields another relabeling and since due to the bijectivity, each map (2) is invertible, all possible relabelings form a group: the group of all bijective maps from E to itself. This group is called the *symmetric group* of the set E .

The crucial insight here is that the group of relabelings allows us to quantify the ambiguity of any statement that refers to aspects of experience. Consider, e.g., the case where a statement only involves one label $e_1 \in E$. Since we are disregarding collatable relations at this point, this statement could just as well have been formulated with any other label $e_2 \in E$, simply because an experiencing subject may choose any label whatsoever to denote any quale. Mathematically, this is reflected by the fact that there is at least one relabeling s such that $s(e_1) = e_2$. The same reasoning can be applied to sequences (e_1, \dots, e_n) of labels, e.g., obtained by verbal reports at subsequent times. The ambiguity of a sequence (e_1, \dots, e_n) of labels is the set of all sequences (e'_1, \dots, e'_n) that can be obtained from the former by a relabeling s , i.e., the set of all sequences (e'_1, \dots, e'_n) for which there exists a relabeling s such that $(e'_1, \dots, e'_n) = (s(e_1), \dots, s(e_n))$.

These statements are in fact statements about equivalence classes. To see this, define two labels e_1 and e_2 to be equivalent, $e_1 \sim e_2$, if and only if there exists a relabeling s such that $s(e_1) = e_2$. The ambiguity of a label $e \in E$ is given precisely by the equivalence class of this label,

$$[e] := \{e' \mid e' \sim e\} = \{e' \mid \exists s : e' = s(e)\}, \quad (3)$$

because this class contains all labels that an experiencing subject could have chosen. The same is true for sequences: if we define two sequences to be equivalent, $(e'_1, \dots, e'_n) \sim (e_1, \dots, e_n)$, if and only if there exists a relabeling s such that $(e'_1, \dots, e'_n) = (s(e_1), \dots, s(e_n))$, the ambiguity of a sequence of labels is given precisely by the equivalence class of this sequence,

$$[(e_1, \dots, e_n)] := \{(e'_1, \dots, e'_n) \mid (e'_1, \dots, e'_n) \sim (e_1, \dots, e_n)\}, \quad (4)$$

because this class contains precisely all those descriptions of the sequence that an experiencing subject may give. Another way to put this is that the equivalence classes (3) and (4) are what is empirically well defined, not the labels themselves; these only have meaning for the experiencing subject him/herself once he/she has chosen a particular labeling.

This concludes the description of the case that ignores structural features of experience. Its artificial nature is reflected in the fact that the symmetric group allows one to map any choice (e_1, \dots, e_n) of

labels to *any* other choice (e'_1, \dots, e'_n) , provided that every label occurs at most once in each choice. Thus, there are very few equivalence classes (only one if $n = 1$). We now proceed to the discussion of the appropriate case.

3.3.2. Taking Relations into Account

Next, we take into account the collatable relations between aspects of experience as established in Phenomenological Axiom 3. To do so, we work with the *experience space* E introduced above: i.e., we assume that the relations between qualia have been represented on the set of labels¹³ and ask experiencing subjects to pick labels for the qualia they experience in accordance with this representation. As above, we refer to any such choice as *labeling*.

This constraint on how labels can be chosen implies that the freedom of every experiencing subject to choose labels is smaller than in the case above: the functions (2) only constitute relabelings if they preserve the collatable relations represented on the set E , i.e., if they preserve the structure of the space E . A bijective function from a space to itself that preserves the structure of this space is called an *automorphism* of the space. As above, automorphisms form a group. Thus, in the case where we take into account the collatable relations, *relabelings* are elements of the

$$\text{Automorphism group } \text{Aut}(E). \quad (5)$$

We summarize this by saying that the automorphism group $\text{Aut}(E)$ describes the *freedom of relabeling* of every experiencing subject.

It is here that Convention 1 is important. Since for every collatable aspect of experience, there is a unary relation that the automorphism group needs to leave invariant, this convention ensures that automorphisms do not change labels of collatable aspects of experience. As a result, the automorphism group allows us to quantify precisely the ambiguity inherent in any reference to qualia.

In order to identify the ambiguity of any statement that uses a sequence (e_1, \dots, e_n) of labels, we can argue exactly as in the simplified description above, replacing transformations (2) by automorphisms. The result is that the ambiguity is given precisely by the equivalence class

$$[(e_1, \dots, e_n)] := \{(e'_1, \dots, e'_n) \mid (e'_1, \dots, e'_n) \sim (e_1, \dots, e_n)\}, \quad (6)$$

where \sim denotes the equivalence relation defined as

$$(e'_1, \dots, e'_n) \sim (e_1, \dots, e_n) \quad \text{if and only if there is an } s \in \text{Aut}(E) \\ \text{such that } e'_i = s(e_i) \text{ for all } i = 1, \dots, n. \quad (7)$$

This class contains precisely all descriptions of the sequence of aspects of experiences that an experiencing subject might give: a description in every possible labeling. To obtain the ambiguity of individual labels, we simply set $n = 1$.

In summary, what this shows is that the empirically well-defined references to experience are given by elements of the quotient space

$$E^{\times n} / \sim, \quad (8)$$

where E is the space (1) whose structure represents collatable relations between aspects of experience, where \sim denotes the equivalence relation (7) and where $n \in \mathbb{N}$ is the length of a sequence.

¹³ For all practical purposes, one can obtain such a representation by simply asking one experiencing subject to pick a labeling and to report, in terms of this labeling, on his/her experienced relations. Other experiencing subjects are then required to choose labels according to this representation. For details, see Example 9. For explicit examples on how such a representation might look, cf. Examples 10–13.

Remark 1. *In practice, we typically establish labels by reference to particular “external” events, such as particular wavelengths emerging from a light source in the case of color experiences. Socially established labels of this sort are of course very useful in various circumstances, precisely because they correlate with external events. However, a priori, there is no reason to assume that qualia of different experiencing subjects that are denoted by the same label are the same, even if the labels correlate with the same external event.¹⁴ In fact, an assumption of this kind has no empirical meaning because the definition of qualia implies that neither the identity of qualia of different experiencing subjects with an external event, nor the equality of qualia of different experiencing subjects can be empirically tested. Statements of this sort can only be meaningful if formulated based on a scientific methodology that is compatible with the non-collatability of the aspects of experience under consideration.*

3.4. A Phenomenological Grounding of the Scientific Study of Consciousness

In the previous sections, we fixed the basic terminology, such as what we take the term experience to denote and how qualia are defined in these terms. We furthermore used phenomenological axioms to warrant the introduction of labels, which has in turn allowed us to ground a mathematical representation of experience. Finally, we analyzed the implications of non-collatability (and hence of ineffability, privateness, and inaccessibility) in terms of this formal representation. Together, this gives rise to a grounding of the scientific study of consciousness, i.e., allows us to specify what is to be studied and how.

First, concerning the task of the scientific study of consciousness, what is to be studied is simply experience as defined in Definition 2 and its relation to the physical domain. By Phenomenological Axiom 1, this includes collatable aspects of experience, as well as qualia. What is required to do so is a combination of the usual scientific methodology with some novel tools (developed in the remainder of this article). How these methodologies are combined is described by the formal representation of conscious experience in the experience space E .

The usual scientific methodology, e.g., the one in use today in the neuroscience of consciousness, can be applied to all intersubjectively well-defined references to experience, i.e., to the equivalence classes (6). When taken together, they constitute the quotient space (8), which provides a comprehensive description of all intersubjectively meaningful aspects of experience. This quotient space contains, in particular, all references to collatable aspects of experience.

However, the usual scientific methodology cannot be used to investigate individual elements of equivalence classes (6), if a class has more than one element, because the experiences labeled by these elements cannot be referenced intersubjectively in a meaningful way. These elements in fact generate an explanatory gap (Section 4). The study of these aspects of experience is, nevertheless, part of the task of the scientific study of consciousness. There is a fact as to which member of an equivalence class is experienced, and this fact cannot a priori be excluded from constituting a scientific explanandum.

The main achievement of this article in the following sections is to show that formal tools can be defined that allow us to go beyond a scientific analysis of the quotient space (8). Referring to these results, we can specify the grounding that arises from the previous definitions as follows. Due to the importance of phenomenological axioms in grounding the formal structure, we refer to this grounding as phenomenological grounding.

Definition 7. *What is to be studied by the scientific study of consciousness according to the phenomenological grounding is experience as defined in Definition 2 and its relation to the physical domain. This includes the study of intersubjectively well-defined aspects of experience using the quotient space (8) and standard scientific methodology, as well as the study of qualia proper, represented formally in the experience space (1) using the formal mathematical methodology derived in Section 6.*

¹⁴ One may even take this to be unlikely, given the difference of brain physiology and neuronal structure across individuals.

3.5. Examples

We close this section with several examples. First, in Example 9, we continue the discussion of color experience and show that color spaces, which are largely in use in commercial applications, constitute the experience spaces for color qualia as defined above. In Examples 10–13, we consider various possible mathematical structures of the experience space E , some of which have been proposed in the literature.

Example 9. To illustrate the meaning of the experience space E and the group $\text{Aut}(E)$, as well as Remark 1, we consider again color experiences. As we explained in Example 3, these satisfy the defining property of qualia.

We will generally denote the quale “what it is like to see light of wavelength λ ” as “experience of λ ”. For the purpose of this example, we will disregard the fact that color experience is highly sensitive to the geometry of the lighting of a scene and to the expected material properties of an object’s surface. We will use the symbol $\bar{\lambda}$ to denote a mixture of light of varying wavelength.

We start by fixing a particular human observer, the “standard observer” [36], and choose a set E_{cl} that is in one-to-one correspondence with all colors that this human can experience. As usual in color science, we assume that there is a large class C of humans that have the same set of possible color experiences as the standard observer. This assumption implies that every human in the class C can specify a one-to-one mapping between the set E_{cl} and his/her color experiences. The fact that color experiences are qualia as defined here is reflected in the fact that there is no unique one-to-one mapping. The set E_{cl} is thus a set of labels of color qualia as introduced after Phenomenological Axiom 2. It is also the basis of the definition of color spaces (cf. below).

The set E_{cl} can be calibrated: since color experiences arguably arise as a response to mixtures $\bar{\lambda}$ of light impinging on the retina, we may identify every element $e \in E_{cl}$ with a particular mixture $\bar{\lambda}$. The set of mixtures visible to the human eye can, in turn, be represented as a subset $S \subset \mathbb{R}^3$, roughly speaking by taking the three components of a vector $v \in S$ to represent the relative intensities of three reference wavelengths.¹⁵ Putting these two steps together, we may in fact choose the set E_{cl} to be the subset $S \subset \mathbb{R}^3$. In this case, every label $e \in E_{cl}$ is a three-tuple of real numbers that specifies which mixture $\bar{\lambda}$ of light has to be presented to a particular human to evoke the quale that he/she has denoted by that very label e .

This calibration may lead one to think that there is a unique way of referring to color qualia. However, this is not the case. To see this, assume that we fix some label/vector $e \in S = E_{cl}$, as well as two experiencing subjects A and B . Let us denote the mixture of light that corresponds to this vector as $\bar{\lambda}_v$. When we present this mixture $\bar{\lambda}_v$ to the two experiencing subjects, subject A has the color experience he/she has labeled as e , and so does subject B . However, this has nothing to say on whether the color experiences are the same or not: e.g., subject B might have the color experience subject A is having upon presentation of a completely different mixture $\bar{\lambda}_w \neq \bar{\lambda}_v$.

This illustrates the fundamental difficulty related to qualia as defined in Definition 6: if we “know” (e.g., as the result of some scientific investigation) that the presentation of the same color stimuli $\bar{\lambda}_v$ to various subjects results in them having the same color experience, we could meaningfully talk, or refer to, color experiences of different subjects in terms of stimuli. More generally, if statements of the type

$$\text{“subject } A \text{ will have color experience } X_1 \text{ once presented input } \bar{\lambda}\text{”} \quad (9)$$

were known, these statements would allow us to directly refer to A ’s color experiences, putting us into the position to do science as usual. However, the fundamental difficulty of the subject is that statements like (9) do

¹⁵ This is an experimental fact that is due to the biological details of the cone cells in the human eye. Since various mixtures $\bar{\lambda}$ evoke the same color experience, some conventions have to be made in order to fix the subset S uniquely (e.g., a choice of reference wavelengths). Furthermore, due to the particular responsivity curves of the cone cells, no finite set of wavelengths can be combined to achieve all colors that a human can experience. However, suitable experimental procedures exist so that all visible mixtures can be represented in \mathbb{R}^3 nevertheless [36].

not carry any intersubjective meaning at all: due to the impossibility of collating color experiences, Statement (9) cannot be distinguished (by anyone but subject A) from the statement

“subject A will have color experience X_2 once presented input $\bar{\lambda}$ ”,

where X_2 is any color experience of A with the same unary collatable relations (such as intensity). This problem exists independently of whether we consider the statement (9) to be a hypothesis or to be the result of some purported scientific investigation. Statements of this type do not have unambiguous intersubjective meaning.

As explained above, what has intersubjective meaning are the equivalence classes (6). They express facts about color experience that are invariant with respect to the labeling that an experiencing subject chooses. We now illustrate this in detail for color qualia.

First, we need to find the collatable relations between color experiences referred to in Phenomenological Axiom 3. Luckily, this has been on the agenda of color science for decades, so that we may simply turn to its results. Put in simple terms, there seem to be three types of collatable relations [36]: the continuity of the change of colors (whether some time-continuous sequence of color experiences is perceived as continuous or not), the behavior under mixtures of colors (whether a mixture of two color experiences is perceived as equal to another color experience or not), and (less well known) a notion of the distance of colors (whether two color experiences are perceived as more different to each other than another pair of color experiences).

Next, we need to translate these collatable relations into mathematical structures on the set E_{cl} . This yields the experience space (1) of color qualia. Again, color scientists have done the work for us: they have defined color spaces in order to formalize these collatable relations [36]. A color space is a closed subset S of \mathbb{R}^3 which is in a one-to-one correspondence with all colors humans may experience, chosen such that continuity is represented by the induced topology of \mathbb{R}^3 (a path of color experiences is continuous if the labels form a continuous path in the color space), mixture is represented by straight lines (equal mixing of two color experiences e_1 and e_2 yields the color experience that carries the label that is at the center of the straight line that connects e_1 and e_2), and finally, experienced distance of color qualia is represented by a metric on S .¹⁶ Thus, a color space is an experience space (1) for color qualia.

There are many subsets of \mathbb{R}^3 that satisfy these requirements: for any choice of subset S , there is a large class of transformation of \mathbb{R}^3 that, together with a corresponding transformation of the metric, yield another subset S' of \mathbb{R}^3 that equally represents color experiences, as well as their collatable relation. Color science uses the calibration described above to fix specific choices of subsets S , so that the coordinates of the elements of S can be translated into mixtures of wavelengths $\bar{\lambda}$. However, as explained above, for the study of color experience, calibrations do not have any relevance a priori, so that no particular choice of subset can be singled out.

In order to specify the group of relabelings for this example, we note that in more abstract terms, a color space is a smooth three-dimensional Riemannian manifold.¹⁷ Its topology represents the continuous changes of color experience, and its metric g specifies both the geodesics (generalized “straight lines”), which describe the mixture of color experiences, as well as a distance function, which describes the experience of distance between color qualia. The various choices of subsets S of \mathbb{R}^3 correspond precisely to the choices of the coordinates of this manifold. We summarize this as

$$E = (E_{cl}, g). \quad (10)$$

This is the actual form of the experience space (1) of color qualia. Its elements label the set of color experiences, and its structure represents the collatable relations between them. An experiencing subject can specify his/her color experiences by specifying points (in the case of individual color experiences) or curves (in the case of time-continuous color experiences) on this manifold. The freedom of choosing labels is described by

¹⁶ We take it that straight lines describe mixtures of color experiences, which have to be distinguished from the experience of mixtures of colors. Thanks to an anonymous referee for pointing this out.

¹⁷ Cf. [36]. However, note that a more axiomatic treatment may result in different mathematical spaces [37,38]. Furthermore, the assumption of smoothness may not be justified, and one might have to consider manifolds with corners.

the automorphism group of E . In the case (10) of a Riemannian manifold, this is the group of isometries, i.e., diffeomorphisms that leave the metric invariant

$$\text{Aut}(E) = \text{Iso}(E).$$

Thus, the ambiguity of any statement in terms of color labels (e_1, \dots, e_n) is given by the equivalence class

$$[(e_1, \dots, e_n)]$$

which is defined as in (6) with two sequences being equivalent if there is an isometry $s \in \text{Iso}(E)$ that transforms every element of the first sequence into the corresponding element of the second sequence.¹⁸ The actual form of the equivalence classes depends on the metric g , which can be determined experimentally. The current version of the distance function internationally in use was reviewed, e.g., in [39], a discussion of which however goes beyond the scope of this example.¹⁹

Putting everything together, we conclude that any statement, scientific or otherwise, that addresses color experiences *sensu stricto*, i.e., that addresses what it is like to experience colors, only makes sense if it is invariant with respect to $\text{Iso}(E)$ transformations. This is a consequence of the fact that qualia are non-collatable and of the corresponding freedom of every experiencing subject to choose names for the qualia he/she experiences.

The difference between labels of color experiences and color experiences (color experiences *de dicto* and color experiences *de re*, so to speak) can be crucial for scientific investigations. For example, if a study compares the calibrated label e that a subject reports with neural activity, it does not investigate the relation between neural activity and color experience, but rather the relation between neural activity and presentation of wavelengths $\bar{\lambda}$ to the retina. These two objects of investigation refer to completely different scientific agendas.

Example 10 (Pretopological structure on E). In the previous example, we relied on results from color science to provide the mathematical structure of the experience space E that represents the color aspects of experience. The goal of this example is to illustrate in more detail how the mathematical structure of E can be defined directly in terms of relations between qualia. To this end, we consider the relation of similarity of two qualia explained in Example 7, but understood in a binary way, i.e., for the purpose of this example, we make the simplifying assumption that any two non-collatable aspects of experience (of one experiencing subject) are experienced either as “similar” or as “not similar” and ignore the experience of varying degrees of similarity. While this restriction may not be warranted in practice, we take it to be justified for pedagogical purposes. When understood in this way, the similarity relation can be used to define a pretopological structure on E as described in [40] (with a slightly different goal in mind), whose presentation we now follow.²⁰

First, we define a binary relation $R_\circ \subset E \times E$ on E . If two qualia with labels e_1 and e_2 are perceived as similar by an experiencing subject, we define the corresponding labels to be related according to R_\circ , which we denote as $e_1 \circ e_2$ (i.e., $e_1 \circ e_2 \Leftrightarrow (e_1, e_2) \in R_\circ$, and similarly below). Thus, R_\circ is given directly by experience. We assume that $e \circ e$ for all $e \in E$. Second, based on the data of R_\circ , we define another relation R_\leq on E , called the “parthood relation” [40], as

$$e_1 \leq e_2 \quad \text{iff} \quad e \circ e_1 \Rightarrow e \circ e_2.$$

¹⁸ Since the ordering of distances between pairs of colors, rather than the numerical value of the distance itself, is collatable, one could make the point that the relabeling freedom is given by the group of diffeomorphisms that leave the metric invariant up to a conformal factor. Since the present example is, mainly, of a pedagogical interest, we do not explore this further at this point; cf. also Footnote 17.

¹⁹ We note that it is possible that some sequences (e_1, \dots, e_n) are not ambiguous, i.e., that $[(e_1, \dots, e_n)] = \{(e_1, \dots, e_n)\}$. This means that there is one unique sequence of color experiences that has the properties represented by the sequence (e_1, \dots, e_n) of labels, or put differently, that there is only one possible choice of labels for this sequence that takes into account the collatable relations as described. Sequences of this kind may be used to remove the ambiguity of the labels they contain and make these aspects of experience accessible to a proper scientific analysis.

²⁰ The following definitions and their relation to topology are intuitively accessible if one thinks about open balls in a metric space such as \mathbb{R}^3 , where \circ is defined as overlap. We remark, however, that the construction does not give rise to a topology, as claimed in [40], since the third Kuratowski closure axiom (idempotence) does not follow.

Thus, $e_1 \leq e_2$ holds iff all qualia that are similar to e_1 are also similar to e_2 . Third, we use the parthood relation R_{\leq} to define yet another relation R_{\sim} , called “connection”, as follows:

$$e_1 \sim e_2 \quad \text{iff} \quad \exists \tilde{e} \in E \text{ such that } \tilde{e} \circ e_1 \text{ and } \tilde{e} \circ e_2, \\ \text{as well as } e \leq \tilde{e} \Rightarrow e \circ e_1 \text{ or } e \circ e_2.$$

Note that $e_1 \leq e_2$ implies $e_1 \sim e_2$. We extend this notation to sets $A \subset E$ by defining

$$e \sim A \text{ iff } e \sim \tilde{e} \text{ for at least one } \tilde{e} \in A.$$

This allows us to define an operator pcl , which takes a subset $A \subset E$ to another subset $\text{pcl}(A)$ that contains all qualia that are connected to at least one of the qualia in A :

$$\text{pcl}(A) := \{e \mid e \sim A\}.$$

The operator pcl satisfies three of the four Kuratowski closure axioms (Section 3.2, [41]), but need not satisfy $\text{pcl}(\text{pcl}(A)) = \text{pcl}(A)$ for all $A \subset E$ (idempotence). Hence, it constitutes a preclosure operator, so that (E, pcl) constitute a pretopological space [42].

In order to define what constitutes a relabeling in this example, we note that a function f between two pretopological spaces (E, pcl) and (E', pcl') is defined to be continuous if

$$f(\text{pcl}(A)) \subseteq \text{pcl}'(f(A))$$

for all $A \subset E$. The automorphism group $\text{Aut}(E)$ of (E, pcl) is the set of all continuous invertible functions $f : E \rightarrow E$ whose inverse is also continuous, with the group operation given by function composition.

Thus, we see neatly how non-trivial mathematical properties of the experience space E can be defined directly in terms of experienced relations between qualia. The similarity relation established via Phenomenological Axiom 3 may, of course, not actually be binary: There seem to be various degrees, maybe even a continuum, of similarities of qualia.

Example 11 (Partial order on E). Our next example goes back to [29]. First, we observe that next to the two relations mentioned in Phenomenological Axiom 3, qualia may in fact have compositional relations that can be collated: an experiencing subject may find that the ineffable aspect of an experience he/she is having at a particular time includes an ineffable aspect he/she has had at another time. In this case, we may say that the former quale includes the latter quale. If e_1 is the label that the experiencing subject has chosen for the former quale and e_2 is the label he/she has chosen for the latter quale, we will denote this relation between the two qualia as $e_2 \leq e_1$.

By convention, we may put $e \leq e$ for all $e \in E$ (reflexivity). Furthermore, it is reasonable to hold that if both $e_1 \leq e_2$ and $e_2 \leq e_1$ for two labels $e_1, e_2 \in E$, these labels actually refer to the same quale, so that $e_1 = e_2$ (anti-symmetry). Finally, qualia seem to satisfy that $e_1 \leq e_2$ and $e_2 \leq e_3$ imply $e_1 \leq e_3$ (transitivity). Therefore, this actually constitutes a partial order on E and turns (E, \leq) into a partially ordered set. The automorphism group consists of bijective functions $f : E \rightarrow E$ that are order-embedding, i.e., that satisfy $e_1 \leq e_2$ if and only if $f(e_1) \leq f(e_2)$ for all $e_1, e_2 \in E$. Thus, one can see nicely that the automorphism group describes the freedom of relabeling: its elements represent changes of labels that preserve the inclusion relation between qualia.

Example 12 (Involutive semigroup structure on E). This example also goes back to [29]. In order to state it, note that in Definition 2, we defined the term “experience” with respect to instants of time. This implies that qualia (being aspects of experience) are associated with an instant of time as well, therefore excluding a sequence

of two qualia arising at two consecutive instants of time to constitute another quale.²¹ However, one might drop this restriction to instants of time and define qualia as aspects of experience in general. Following this line of thought, one could argue that for any two qualia e_1, e_2 , there is another quale e_3 that is the consecutive experience of the two qualia. One might denote e_3 as

$$e_3 = e_1 \& e_2,$$

where the “&” represents “and then” [29]. If one furthermore demands associativity, which does seem to be plausible, this defines a semigroup $(E, \&)$.

Next, one may consider an operation that reverses this temporal order of qualia. This may or may not have deep conceptual meaning: On the one hand, it may merely map any quale of the form $e_1 \& e_2$ to a quale of the form $e_2 \& e_1$, both of which have to exist due to the semigroup structure introduced above. On the other hand, it may express a deep fact about the reversal of psychological time [29]. In both cases, skipping over a few technical details, this gives rise to an involution [29], i.e., a map

$$\begin{array}{ll} * : E \rightarrow E & \text{such that} \\ e \mapsto e^* & (e^*)^* = e. \end{array}$$

In summary, the time composition relation of qualia may be represented on the space of labels in terms of an involutive semigroup structure.

Example 13 (Hilbert space structure on E). The last example is intended to evaluate how far the axioms of a Hilbert space can be grounded in the relations introduced in Phenomenal Axiom 3. The upshot is that whereas some of the axioms can be motivated based on Phenomenological Axiom 3, others cannot. Nevertheless, the example may prove valuable for constructing toy models of consciousness, which is why we include it here.

In what follows, we make several assumptions about the set of all experiences that an experiencing subject might have. These assumptions are phenomenological in flavor, yet some may ultimately not be justified.

(A1) We assume that with respect to any two qualia of one experiencing subject, the experiencing subject might have an experience that has exactly these qualia as ineffable aspects.

With respect to qualia of the “what it is like to be” type (Example 5), this assumption amounts to the following statement: if an experiencing subject has had an experience that included an ineffable “what it is like to be” aspect (quale) that he/she labels by e_1 and another experience that included an ineffable “what it is like to be” aspect (quale) that he/she labels by e_2 , then it is possible that he/she will have an experience that has exactly e_1 and e_2 as ineffable aspects. We will use the term “simultaneous experience of e_1 and e_2 ” as an abbreviation for the statement that the experiencing subject in question has an experience that includes both aspects e_1 and e_2 . To give an example, let e_1 refer to what it is like to taste cheese and e_2 refer to what it is like to smell wine. In this case, Assumption (A1) amounts to granting the possibility of the experiencing subject in question simultaneously experiencing what it is like to taste cheese and what it is like to smell wine. Whether this experience actually arises when the subject eats cheese and drinks wine is of no concern with respect to Assumption (A1). We take the combination of the same experience e with itself as denoting the experience of quale e , but twice as intense (cf. below). In order to motivate a group structure with respect to simultaneous experience, the following assumption is necessary:

(A2) We assume that there is a unique neutral quale, which we denote by “0”. Furthermore, we assume that for every quale e , there is a quale $-e$ such that an experience that includes both e and $-e$ is not distinguishable from (and hence equal to) the experience of the neutral quale.

It seems that this assumption is utterly beyond empirical justification, since it invokes something like “cancellation” of “what it is like to be” aspects of experiences, so that we may only be able to ground a semigroup

²¹ The term “instant of time” may refer to experiential instants of time or to instants of time as used in physics, i.e., points $t \in \mathbb{R}$.

structure of qualia with respect to combination (“simultaneous experience”). For the purpose of this example, we proceed nevertheless. We denote the simultaneous experience of two qualia e_1 and e_2 by \oplus , so that the ineffable aspect of the experience that comprises both qualia labeled as e_1 and as e_2 established by Assumption (A1) is labeled by $e_1 \oplus e_2$. Associativity and commutativity hold, so that we have:

- ▶ (A1) and (A2) imply that $\oplus : E \times E \rightarrow E$ is an abelian group.

Next, we model changes of intensity, as conceded in Phenomenological Axiom 3, by a positive real number in the following sense: if e_2 is the same quale as e_1 , but c times more intense, then we denote $e_2 = ce_1$, where $c \in \mathbb{R}^+$. For $c \in \mathbb{R}^-$, ce_1 is the opposite experience $-e_1$ introduced in (A2), but experienced $|c|$ times as intense as $-e_1$, where $|c|$ is the modulus of c . Finally, we assume that as intensity decreases, $c \rightarrow 0$, any experience goes over to the neutral quale, formally $\lim_{c \rightarrow 0} ce = 0$ for any $e \in E$, where 0 denotes the neutral quale introduced in (A2). Making the idealized assumption that a continuum of more and less intense versions of any experience is possible, we have:

- ▶ The intensity relation of Phenomenal Axiom 3 may be taken to give rise to a scalar multiplication $\odot : \mathbb{R} \times E \rightarrow E$.

As usual, we suppress the symbol \odot for scalar multiplication. We need to check whether the axioms of a vector space relating scalar multiplication and addition hold. Our interpretation implies that $1e = e$; hence, the neutrality of $1 \in \mathbb{R}$ holds. The two axioms of distributivity read

$$c(e \oplus e') = (ce) \oplus (ce') \quad \text{and} \quad (11)$$

$$(c + c')e = ce \oplus c'e \quad \text{for all } c, c' \in \mathbb{R} \text{ and } e, e' \in E \quad (12)$$

Axiom (11) says that a c times more intense simultaneous experience of e and e' arises as the combination of c times more intense experiences of e and e' , respectively, which we take as a plausible assumption in the context of this example. Axiom (12) states a compatibility of the addition of intensities with combinations of experience, e.g., it says that an experience e' that is the same as another experience e , but twice as intense, $e' = 2e$, can arise as the simultaneous experience of the combination of e with itself. We render this axiom at least somewhat plausible by defining the combination of an experience with itself to be the same experience experienced twice as intense. Finally, we note that the associativity axiom $(cc')e = c(c'e)$ is compatible with our interpretation of \oplus and \odot . We therefore have:

- ▶ (E, \oplus, \odot) satisfies the axioms of a vector space.

It remains to implement the relation of similarity between qualia. As before, we idealize and assume that there is a non-negative real number that specifies how similar two qualia e_1 and e_2 are. We denote this number by $\langle e_1, e_2 \rangle$. If e_1, e_2 are not similar at all, we set $\langle e_1, e_2 \rangle = 0$. If they are similar to some degree, we have $\langle e_1, e_2 \rangle > 0$, where a larger value implies more similarity. It seems natural to impose symmetry, $\langle e, e' \rangle = \langle e', e \rangle$, for all $e, e' \in E$. An inner product furthermore satisfies

$$\langle e, e \rangle = 0 \Leftrightarrow e = 0 \quad \text{(Definiteness)}$$

$$\langle e, ce' \rangle = c \langle e, e' \rangle \quad \text{(Linearity)}$$

$$\langle e, e' \oplus e'' \rangle = \langle e, e' \rangle + \langle e, e'' \rangle$$

for all $c \in \mathbb{R}$ and $e, e', e'' \in E$. Out of those three axioms, only the last one seems reasonable to some extent. It says that similarity is compatible with simultaneous experience: the similarity between a quale e and the simultaneous experience of qualia e' and e'' is given by the sum of the similarity of the quale e to each one of the qualia e' and e'' .

Definiteness says that the only quale that is not similar to itself is the neutral quale. This seems rather problematic if one chooses the interpretation of zero introduced in (A2). The first axiom of linearity says that the similarity between a quale e and a c times more intense version of a quale e' is given by c times the similarity between e and e' . As mentioned before, in order to have a nice and clear example, we will accept also these assumptions for now, so that in summary, we have:

- ▶ $(E, \oplus, \odot, \langle \cdot, \cdot \rangle)$ satisfies the axioms of an inner product space or pre-Hilbert space.

The inner product $\langle \cdot, \cdot \rangle$ introduces a norm on E as usual by $\|e\| = \sqrt{\langle e, e \rangle}$. This norm may be interpreted as the intensity of a quale e .

The inner product space $(E, \oplus, \odot, \langle \cdot, \cdot \rangle)$ may not be complete with respect to this norm, meaning that there are Cauchy sequences in E that do not converge to an element in E . In terms of qualia, this means that there are sequences of qualia whose elements become ever more similar to each other, but that do not converge to any quale in the topology specified by the similarity relation. In order to exclude such cases, we consider the completion of E with respect to the norm $\|\cdot\|$, which is unique up to isometric isomorphism. Alternatively, we may assume that there is a finite number of classes of non-similar qualia, so that completeness holds automatically. A complete inner product space is a Hilbert space. Denoting, as usual, completion by a line over the corresponding quantities, we have:

- ▶ The experience space E carries the structure of a real Hilbert space $\overline{(E, \oplus, \odot, \langle \cdot, \cdot \rangle)}$, which we denote by \mathcal{H}_E .

Note that this is an abstract Hilbert space: due to the ineffability of qualia, the elements of the Hilbert space do not have an intrinsic collatable nature (as, e.g., the case if one considers function spaces). The automorphism group $\text{Aut}(E)$ is the group $U(\mathcal{H}_E)$ of unitary operators.

4. Explanatory Gap

An “explanatory gap” [43] between a phenomenon²² and natural science occurs if the phenomenon has properties that render it incompatible with all notions of explanation used in natural science. This is in particular the case if the phenomenon violates a necessary condition for the application of any of these notions of explanation.

Explanatory gaps are taken by some to indicate or entail ontological gaps (cf. Chapter 5, Section 3.4, [31]). Whether this is legitimate or not is a question that we will not need to address here. What matters for us is that if there is an explanatory gap, a change of methodology is necessary if the phenomenon is to be addressed by scientific means.²³ This change may or may not be motivated by ontological considerations.

Whether there is an explanatory gap or not strongly hinges on what one takes scientific explanation to be, e.g., in [30], it was assumed that explanations in natural science can address “only structure and function, where the relevant structures are spatiotemporal structures, and the relevant functions are causal roles in the production of a system’s behavior” [31] (p. 105f.), which implies that phenomenal experience, defined in [30] to consist of precisely those aspects of experience that do not have a structure and function, cannot be explained by natural science (for details, cf. Appendix A). While this axiomatic derivation of an explanatory gap is undoubtedly important, the underlying notion of explanation is too narrow (Appendix B.3). This calls both the explanatory gap and the grounding built on it into question.

This is different for qualia as defined in Definition 6. Since the deductive-nomological model of explanation, the deductive-statistical model of explanation, the statistical relevance model of explanation, and the causal mechanical model of explanation all tacitly presuppose that descriptions of the explanandum can be collated [44], a thorough explanatory gap exists between any scientific explanation and qualia as defined here. No scientific methodology applied to date can be used to address non-collatable aspects of experience.

Put in simple terms, this comes about from the fact that all explanations used in natural science to date need to assume that the phenomenon under investigation is intersubjectively accessible. Since our

²² Here, by “phenomenon”, we mean anything that occurs or manifests itself in a general sense, including both scientifically observable “empirical phenomena” (such as the data of an experiment), as well as what is directly or indirectly perceived (experiences).

²³ This is not to say, of course, that every phenomenon *can* be addressed by scientific means. There may be phenomena to which the scientific method cannot be applied. However, it seems that the only way to establish whether this is the case for a particular phenomenon is to try to develop a suitable methodology and, if successful, to apply it.

definition of qualia comprises those aspects of experience that are not intersubjectively accessible, they cannot be addressed by the standard methodology. There is no possibility at present in natural science to explain why an experiencing subject experiences a particular quale over and above the explanation of the collatable relations between qualia.

Thus, we conclude that a change of methodology is necessary if qualia are to be addressed scientifically. The remainder of this article is devoted to developing this change in methodology. Our results show that mathematical tools can be devised that allow us to address it. The resulting methodology generalizes Chalmers' strategy (outlined in Appendix A) and constitutes a formal framework for models of consciousness.

5. The Mathematical Structure of Models of Consciousness

Models of consciousness are hypotheses about how conscious experience and the physical domain relate. In this section, we describe the general mathematical structure these models may take making use of the minimally sufficient mathematical structure of any formal scientific theory (Section 5.1) and of the epistemic asymmetry of conscious experience (Section 5.2). Finally, we introduce notation that will be used further below (Section 5.3).

5.1. The Mathematical Structure of Scientific Theories

There are various different accounts in the philosophy of science of what constitutes a scientific theory. Roughly, one may distinguish syntactic accounts, semantic accounts, and pragmatic accounts [45], which differ mainly in the role they attribute to mathematical formalization. Which account of scientific theories is most adequate for the scientific study of consciousness is yet to be seen.

The following list of formal ingredients is general enough to include any of the above-mentioned accounts of what constitutes a scientific theory. In preparation, we remark that a family $(d_t)_{t \in \mathcal{I}}$ is a function $f : t \mapsto d_t$, which we will call "trajectory". It describes the change of dynamical variables with respect to the parameter t . Needless to say, the following list is not intended to be sufficient.

Definition 8. *The mathematical structure of a scientific theory T is comprised of at least:*

- ▶ A set of dynamical variables d (Those quantities whose variation is determined by T to some extent.²⁴)
- ▶ Some background structure b . (Variables, or general mathematical structures, whose change is not determined by T . Background structure needs to be fixed in order to determine the variation of d in a particular application.) The variation or change of the dynamical variable of a theory can be expressed with respect to some parameter t that takes values in some set \mathcal{I} . Typically, the parameter is assumed continuous and interpreted as time. However, this is not necessary: the set \mathcal{I} may or may not carry some mathematical structure (such as a topology), and it may or may not be interpretable as time.
- ▶ A set of kinematically possible trajectories \mathcal{K} . Sometimes, this includes all possible trajectories, $\mathcal{K} = \{(d_t)_{t \in \mathcal{I}}\}$, but in many cases, trajectories need to satisfy certain mathematical requirements, such as differentiability with respect to the parameter t .
- ▶ Some laws \mathcal{L} . (Typically equations or variational principles, but \mathcal{L} may also include different formal ingredients (such as those provided by category theory) or even non-formal ingredients, as required by pragmatic accounts of scientific theories.)
- ▶ A set of dynamically possible trajectories \mathcal{D} , which we also call solutions of T . These are those kinematically possible trajectories ($\mathcal{D} \subset \mathcal{K}$) that are selected by the theory's laws in a particular application of the theory, given some choice of background structure and possibly taking into account some "nonformal patterns in theories" [46] (p. 55).

²⁴ We use the word "variable" in a general sense here: a variable may represent something as simple as a natural number just as well as an operator-valued field on some manifold.

In the next section, we will put these ingredients of a scientific theory into connection with the definitions introduced in Section 3.1. In doing so, we will have to distinguish between a general theory T and those theories that have been put forward (or are anticipated) by contemporary natural science. Similar to Chalmers' use of the term "physical domain" (Appendix A), we will refer to the latter as *physical theories*. We will use the symbol T_p to indicate one of these theories and denote its dynamical variables, background structure, kinematically possible trajectories, and solutions by d_p , b_p , \mathcal{K}_p , and \mathcal{D}_p . Finally, we will assume that the physical theories are formulated in terms of a *state space* P , which is chosen such that according to the laws of T_p , each $p \in P$ determines a unique trajectory in \mathcal{D}_p , i.e., there is a one-to-one correspondence between solutions $(p_t)_{t \in \mathcal{I}} \in \mathcal{D}_p$ and states $p \in P$.

We use the term *model* to denote a theory that is being proposed. This includes full-fledged theories that have not received the kind of empirical support usually required in science, but also "toy-models", which do not aim for a comprehensive account of some class of phenomena, but rather serve to study some specific aspect of it or to test a general idea of how the phenomena could be modeled.

Finally, for use in Section 6, we review the general definition of a symmetry group. Note that \mathcal{K} denotes the kinematically possible trajectories introduced above.

Definition 9. A group \mathcal{G} is a symmetry group of a theory T [47] (p. 43) if and only if the following conditions are satisfied:

- (a) There is an effective action $\mathcal{G} \times \mathcal{K} \rightarrow \mathcal{K}$ of \mathcal{G} on \mathcal{K} .²⁵
- (b) This action leaves the solutions \mathcal{D} of T invariant.

If ϕ is an action of \mathcal{G} on \mathcal{K} that satisfies the requirements (a) and (b), the pair (\mathcal{G}, ϕ) is a symmetry of T .

5.2. Models of Consciousness

We now apply Definition 8 to give a general account of what a model of consciousness is. To this end, we make use of the epistemic asymmetry of conscious experience.

Epistemic asymmetry is the name of the most fundamental epistemological problem associated with conscious experience, namely that there are "two fundamentally different methodological approaches that enable us to gather knowledge about consciousness: we can approach it from within and from without; from the first-person perspective and from the third-person perspective. Consciousness seems to distinguish itself by the privileged access that its bearer has to it" [21]. The epistemic asymmetry implies that there are two epistemically distinct notions of state, one associated with the third-person perspective and one with the first-person perspective.

Whereas metaphysical theories of mind may deal with only one of them and leave the relation to the other somewhat implicit, models of consciousness may not. Being *scientific* hypotheses about how experience relates to the physical domain, the relation of these two epistemically different states is exactly what models of consciousness need to address. Even if they take the third-person state to be fundamental (as in physicalist ontologies), they need to give a description of how the first-person state evolves in time, i.e., why conscious experience appears to be what it is. Even if they take the first-person state to be fundamental (as in idealist ontologies), they need to give a description of how the third-person state evolves, i.e., why the outside world appears to be what it is. The existence of these descriptions is what marks the difference between formal ideas and scientific models.

The mathematical representation of phenomenal consciousness developed in Section 3 is precisely what describes the first-person perspective in formal terms, with first-person states being elements e of the experience space E . A formal account of third-person states, on the other hand, is provided by natural science, which is devoted to the study of phenomena in the third-person perspective in the first place. Referring to theories of natural science by and large as "physical theories", the third-person states are thus the states utilized in physical theories, e.g., states of neural networks or

²⁵ An action is effective (\equiv faithful) if and only if no group element other than the identity fixes all elements of \mathcal{K} .

other descriptions of the human brain. Using the notation introduced in the last section, we denote physical theories by T_P and their state space by P .

In summary, the above shows that by virtue of consciousness' epistemic asymmetry, a model of consciousness needs to prescribe a relation between states of experience and physical states, independently of which ontology it seeks to express. In formal terms, this means that it needs to prescribe a relation between the experience space E and the state space P of a physical theory T_P . Applying the minimal formal ingredients of a scientific theory identified Definition 8, this means that the dynamical variables of a model of consciousness are given by $E \times P$ or in fact $E^n \times P$ with $n \geq 1$ in case a model of consciousness can prescribe more than one experiencing subject for a given physical state (as, e.g., the case with integrated information theory when making use of the exclusion postulate [7]).

This is summarized in the following definition.

Definition 10. Let T_P denote a physical theory. A pre-model of consciousness M is a theory as in Definition 8, where:

- (i) The dynamical variables are a Cartesian product of the physical state space P of T_P together with one copy of the experience space E for each experiencing subject,

$$d = \underbrace{E \times E \times \dots \times E}_{\text{experiencing subjects}} \times P. \tag{13}$$

- (ii) Kinematically possible trajectories \mathcal{K} are a subset of families of dynamical variables,

$$\mathcal{K} \subset \left\{ (e_t^1, e_t^2, \dots, e_t^n, p_t)_{t \in \mathcal{I}} \right\}, \tag{14}$$

where $e_t^i \in E$, $p_t \in P$, n is the number of experience spaces in (13) and \mathcal{I} is some parameter space.²⁶

We dubbed this structure a “pre-model of consciousness” since it does not yet take into account any of the characteristic features of conscious experience, so that the above mathematical structure may as well describe any other scientific theory that addresses two variables that are epistemically distinct. An improved definition that takes into account some of the characteristic features will be given in Section 6.2 below.

Note that by referencing a physical theory T_P , this definition takes into account that models of consciousness are built on and extend, or allow us to derive, physical theories. An example for the former is again integrated information theory (Section 8.1), and an example for the latter is conscious agent network theory (Section 8.3). Further examples are given in Section 8.

Implied by the above definition is that a model of consciousness M comes with laws \mathcal{L} that select from all kinematically possible trajectories \mathcal{K} a set of solutions \mathcal{D} . Each solution $(e_t^1, e_t^2, \dots, e_t^n, p_t)_{t \in \mathcal{I}} \in \mathcal{D}$ consists of families $(e_t^i)_{t \in \mathcal{I}}$ that describe changes of labels for every experiencing subject $i \in \{1, \dots, n\}$ and of a family $(p_t)_{t \in \mathcal{I}}$ that describes changes of the physical states. The solution thus realizes the mutual influence of conscious experience and physical states as described by the laws of the model M .

²⁶ The subset \mathcal{K} will typically be determined by demanding families $(e_t^1, e_t^2, \dots, e_t^n, p_t)_{t \in \mathcal{I}}$ to satisfy some mathematical properties, such as regularity, which are necessary for the laws \mathcal{L} of T to be well-defined. To exclude pathological cases, we assume that every label $e \in E$ is contained in at least one family $(e_t^1, e_t^2, \dots, e_t^n, p_t)_{t \in \mathcal{I}} \in \mathcal{K}$.

5.3. Notation

We conclude this section by providing a few abbreviations that will be of use further below. We will generally use the shorthand

$$(\bar{e}_t, p_t)_t := (e_t^1, e_t^2, \dots, e_t^n, p_t)_{t \in \mathcal{I}}, \quad (15)$$

where $\bar{e}_t = (e_t^1, e_t^2, \dots, e_t^n)$, to denote elements of \mathcal{K} . Furthermore, we denote by $\mathcal{D}|_P$ those trajectories in the physical state space P that are part of solutions \mathcal{D} of M ,

$$\mathcal{D}|_P := \{(p_t)_{t \in \mathcal{I}} \mid (\bar{e}_t, p_t)_{t \in \mathcal{I}} \in \mathcal{D}\}. \quad (16)$$

This set is not necessarily equal to the set \mathcal{D}_P of solutions of the contemporary physical theory T_P . Whether $\mathcal{D}_P = \mathcal{D}|_P$ or $\mathcal{D}_P \neq \mathcal{D}|_P$ is determined by the laws \mathcal{L} of the model M , cf. Section 7. Similarly, we define

$$\begin{aligned} \mathcal{K}|_P &:= \{(p_t)_{t \in \mathcal{I}} \mid (\bar{e}_t, p_t)_{t \in \mathcal{I}} \in \mathcal{K}\}, \\ \mathcal{K}|_E &:= \{(\bar{e}_t)_{t \in \mathcal{I}} \mid (\bar{e}_t, p_t)_{t \in \mathcal{I}} \in \mathcal{K}\}. \end{aligned} \quad (17)$$

Since the choice of subset \mathcal{K} in (14) is a technical condition prior to the application of any law \mathcal{L} , we may for simplicity assume that $\mathcal{K}|_P = \mathcal{K}_P$.

6. Taking Characteristic Features of Conscious Experience into Account

In the previous section, we derived the general mathematical structure of any model of consciousness. We showed that consciousness' fundamental epistemological feature has some implications regarding the formal structure of any scientific theory that seeks to address it. However, we have not yet taken into account any of the characteristic features of conscious experience.

The goal of this section is to do so. To this end, we work with the notion of non-collatability introduced in Section 3. Since non-collatability is implied by ineffability, privateness, and cognitive, linguistic and communicative inaccessibility, the mathematical structure identified here is in fact a consequence of all of these characteristic features.

In Section 6.1, we derive formal mathematical structures of models of consciousness that are necessary to account for non-collatability. In Section 6.2, we use the result to give an improved definition of what constitutes a model of consciousness and show by means of an example that these structures are also sufficient to address non-collatable aspects of experience. Finally, in Section 6.3, we compare the improved definition of a model of consciousness with the direct description of qualia that may otherwise be used and show that when it comes to non-collatable aspects of experience, mathematical models can achieve more than the direct description.

Together, these results show that the formal mathematical tools developed here can in fact address the explanatory gap between qualia and natural science identified in Section 4. Therefore, this section provides the methodology for the grounding of the scientific study of consciousness outlined in Section 3.4 when it comes to non-collatable aspects of experience.

6.1. Non-Collatability Implies Symmetry

In Section 3.1, we discussed intersubjectively meaningful references to qualia. We found that sequences of labels (e_1, \dots, e_n) are not empirically well defined and have shown that the empirically well-defined references to qualia are precisely the equivalence classes (8). We now repeat a similar analysis for pre-models of consciousness. We first introduce the necessary mathematical tools.

Let $s \in \text{Aut}(E)$ be an element of the automorphism group (5) that describes the freedom of an experiencing subject to choose labels for the qualia he/she experiences. Given a solution $(\bar{e}_t, p_t)_t \in \mathcal{D}$,

we may apply s to that experience space in (13), which is associated with the i th experiencing subject. This gives another trajectory

$$(e_t^1, \dots, s(e_t^i), \dots, e_t^n, p_t)_{t \in \mathcal{I}} \tag{18}$$

where $i \in \{1, \dots, n\}$. The map that takes $(\bar{e}_t, p_t)_t$ to (18) is an $\text{Aut}(E)$ -action ϕ_i on \mathcal{K} , defined as

$$\begin{aligned} \phi_i : \text{Aut}(E) \times \mathcal{K} &\longrightarrow \mathcal{K} \\ (s, (e_t^1, \dots, e_t^i, \dots, e_t^n, p_t)_t) &\mapsto (e_t^1, \dots, s(e_t^i), \dots, e_t^n, p_t)_t \end{aligned} \tag{19}$$

where the subscript i indicates on which experience space $\text{Aut}(E)$ acts. We may take into account the freedom of every experiencing subject to relabel his/her qualia by considering an action ϕ of

$$\text{Aut}(E)^n := \text{Aut}(E) \times \dots \times \text{Aut}(E) \tag{20}$$

on \mathcal{K} , defined as

$$\begin{aligned} \phi : \text{Aut}(E)^n \times \mathcal{K} &\longrightarrow \mathcal{K} \\ (s_1, \dots, s_n, (e_t^1, \dots, e_t^i, \dots, e_t^n, p_t)_t) &\mapsto (s_1(e_t^1), \dots, s_n(e_t^n), p_t)_t. \end{aligned} \tag{21}$$

This action corresponds to the transformations we considered in Section 2. However, in the context of models of consciousness, this is not the most general form of relabeling. The most general form is

$$\begin{aligned} \sigma : \text{Aut}(E)^n \times \mathcal{K} &\longrightarrow \mathcal{K} \\ (\bar{s}, (e_t^1, \dots, e_t^n, p_t)_t) &\mapsto (s_1(e_t^1), \dots, s_n(e_t^n), p'_t)_t, \end{aligned} \tag{22}$$

where similarly to (15), we used the shorthand $\bar{s} := (s_1, \dots, s_n)$ and where p'_t is given by an action $\tilde{\sigma}$ of $\text{Aut}(E)^n$ on $\mathcal{K}|_P$,

$$\begin{aligned} \tilde{\sigma} : \text{Aut}(E)^n \times \mathcal{K}|_P &\rightarrow \mathcal{K}|_P \\ (\bar{s}, (p_t)_t) &\mapsto (p'_t)_t. \end{aligned} \tag{23}$$

This action σ reduces to the action (21) if $\tilde{\sigma}$ is trivial. If $\tilde{\sigma}$ is non-trivial, σ specifies that the physical states are relabeled along with the qualia. We will see below (cf. Section 6.3 for details) that the possibility of a non-trivial $\tilde{\sigma}$ is what allows us to go beyond the standard methodology explained in Section 3.3.

Notation: In what follows, we will use the shorthand $\bar{s}(\bar{e}_t) := (s_1(e_t^1), \dots, s_n(e_t^n))$. As usual, we denote $\sigma(\bar{s}, (\bar{e}_t, p_t)_t)$ as $\sigma_{\bar{s}}((\bar{e}_t, p_t)_t)$. Furthermore, we use $k := (\bar{e}_t, p_t)_t \in \mathcal{K}$.

Remark 2. We remark that each action σ of the form (22) has two different meanings: On the one hand, they describe a relabeling of the trajectory k , i.e., $\sigma_{\bar{s}}(k)$ describes the same situation as k , but with respect to a different choice of labeling. This is the meaning we considered in Section 3.1. It is analogous to a change of reference frame in physics. On the other hand, $k' := \sigma_{\bar{s}}(k)$ is simply another trajectory in \mathcal{K} , which for $\bar{s} \neq \text{id} \in \text{Aut}(E)^n$ describes a scenario that is genuinely different from that of k . Whereas according to k , at time t , experiencing subject i experiences the quale he/she has labeled as e_t^i and physical state p_t pertains, according to k' , the experiencing subject experiences a quale he/she has labeled $s_i(e_t^i)$ and physical state p'_t pertains. This is reminiscent of the distinction between active and passive transformations in physics. Using this terminology, we have:

1. Passive meaning of σ : k and $\sigma_{\bar{s}}(k)$ are the same trajectory expressed in different labeling.
2. Active meaning of σ : k and $\sigma_{\bar{s}}(k)$ are different trajectories expressed in the same labeling.

The fact that active and passive transformation have an identical mathematical form is related to the fact that qualia by virtue of their non-collatability cannot be referenced intersubjectively.

Next, we use the fact that k and $\sigma_{\bar{s}}(k)$ describe the same trajectory with respect to two different choices of labels. Since a different choice of labels must not make a difference, it follows that if k is a solution of M , $\sigma_{\bar{s}}(k)$ needs to be a solution as well, for any choice of $\bar{s} \in \text{Aut}(E)^n$. This leads us to the following definition:

Definition 11. *A necessary condition for a pre-model of consciousness M to be empirically well defined is that there is an $\text{Aut}(E)^n$ action (22) on \mathcal{K} that maps solutions to solutions, i.e., that satisfies*

$$\sigma_{\bar{s}}(\mathcal{D}) = \mathcal{D} \quad (24)$$

for all $\bar{s} \in \text{Aut}(E)^n$.²⁷

Using Definition 9, this yields the following lemma.

Lemma 1. *A necessary condition for a pre-model of consciousness M to be empirically well defined is that $\text{Aut}(E)^n$ is a symmetry group of M whose action is of the form (22).*

Proof. According to Definition 9, $\text{Aut}(E)^n$ is a symmetry of the model M iff (22) is effective and leaves \mathcal{D} invariant. Invariance holds by Definition 11. Effectivity holds because for large enough \mathcal{K} (cf. Footnote 26), every action of the form (22) is effective: for any $\bar{s} \in \text{Aut}(E)^n$ with $\bar{s} \neq \text{id}$, there exists an $e_t^i \in E$ such that $s_i(e_t^i) \neq e_t^i$, as well as a trajectory $k \in \mathcal{K}$ that contains this label, so that $\sigma_{\bar{s}}(k) \neq k$. \square

If there are only collatable aspects of experience, the automorphism group $\text{Aut}(E)$ is trivial, so that (24) is satisfied. Therefore, in this case, all pre-models of consciousness are empirically well defined. If there are non-collatable aspects of experience, however, $\text{Aut}(E)$ is non-trivial, so that (24) poses a condition that needs to be satisfied, and Lemma 1 shows that the condition is in fact that there is an $\text{Aut}(E)^n$ symmetry.

Thus, Lemma 1 establishes the mathematical consequences of non-collatability: the need for an $\text{Aut}(E)^n$ symmetry in a model of consciousness. Since the existence of a symmetry is dependent on the dynamical trajectories of a model of consciousness, which are in turn determined by its laws \mathcal{L} , the condition posed by the lemma is in fact a requirement with respect to the model's laws.

6.2. The Mathematical Structure of Models of Consciousness Revisited

In the previous section, we found that the existence of an action (22) that constitutes a symmetry is a necessary condition for a pre-model of consciousness to be empirically well defined. We therefore need to include this requirement when specifying what constitutes a general model of consciousness. The result is given in the following definition. It specifies the necessary structure of any model of consciousness that is to address any non-collatable aspects of experience, e.g., ineffable, private, or inaccessible aspects. In particular, any model that aims to address any aspect of experience that is referenced by the Nagelian “what it is like” conception (Remark 5) necessarily needs to carry this mathematical structure.

Definition 12. *A model of consciousness is a pre-model of consciousness M as defined in Definition 10, which additionally carries an $\text{Aut}(E)^n$ symmetry of the form (22).*

We remark again that the additional requirement relates to the laws \mathcal{L} of a pre-model of consciousness M ; the laws need to be such that there is an action $\tilde{\sigma}$ that turns (22) into a symmetry.

²⁷ Here, \mathcal{D} is the set of solutions of M introduced above. Note that (24) states an identity of sets. It is equivalent to $\sigma_{\bar{s}}(k) \in \mathcal{D}$ for all $k \in \mathcal{D}$.

In Section 6.3, we will show that this framework indeed allows us to go beyond the limitations of the standard approach explained in Section 3.3. To furthermore make the point that this is a sufficient mathematical framework for the scientific study of consciousness (i.e., sufficient to study qualia proper, cf. Definition 7), we show in the following example that a typical class of ideas put forward in the neuroscience of consciousness can indeed be formalized in this framework: the idea that qualia are determined by physical states.

Example 14. We consider the hypothesis that qualia are determined by physical states, e.g., by neural activity in the brain. What exactly constitutes the determination does not matter in what follows. Examples are type identity theory, where qualia are types of physical states, or functionalism, where qualia are functional roles that are determined by physical states. Another example is non-reductive functionalism [48] where the physical states determine functional roles, which in turn determine qualia via non-reductive laws of nature.

The naive formalization of this idea would be to consider a function (in the mathematical sense) that specifies which quale an experiencing subject experiences for each physical state $p \in P$. As we saw in Section 3.1, the problem with this naive formalization is that qualia cannot be referenced intersubjectively, so that the specification of a function of this type is only meaningful up to the equivalence (7). In order to properly formalize this idea, we proceed as follows.

For simplicity, we consider the case of one experiencing subject ($n = 1$). We assume that a particular labeling has been fixed by the experiencing subject and assume that with respect to this labeling, a function

$$f : P \rightarrow E \quad p \mapsto f(p) \tag{25}$$

is given, where P is the state space of a physical theory T_P as above and E denotes the experience space. This function expresses in which way qualia are determined by physical states and could, ideally, be the result of experiments that include the experiencing subject in question. The state space P could, e.g., refer to neural activity.

Based on this function f , we can define a pre-model of consciousness M . To this end, we set $d = E \times P$, choose \mathcal{K} as the right-hand side of (14), and define the solutions of the model in terms of the solutions $(p_t)_t \in \mathcal{D}_P$ of the physical theory T_P as

$$\mathcal{D} = \{ (f(p_t), p_t)_t \mid (p_t)_t \in \mathcal{D}_P \}. \tag{26}$$

The solutions of this model are thus given by the solutions of the physical theory (e.g., brain dynamics) equipped with qualia as specified by f .

As it stands, this model is not invariant with respect to relabeling, e.g., if the choice of labels is being changed according to some $s \in \text{Aut}(E)$, the solution $(f(p_t), p_t)_t$ is being mapped to the solution $(s(f(p_t)), p_t)_t$, which in general will not be an element of \mathcal{D} as defined in (26). Thus, the theory is not empirically well defined.

In order to establish empirical well definedness, there are two choices: First, one could demand that $s(f(p)) = f(p)$ for all $s \in \text{Aut}(E)$ and all $p \in P$. This amounts to considering a function $f : P \rightarrow E \setminus \sim$, where $E \setminus \sim$ is as in (8), which does not achieve the task of Definition 7. The alternative is to specify an action $\tilde{\sigma}$ as in (23), as we now explain.

The action $\tilde{\sigma}$ describes how the physical state changes along with a change of qualia (active interpretation; cf. Remark 2). We observe that a definition of $\tilde{\sigma}$ as

$$\tilde{\sigma}_s((p_t)_t) := (p'_t)_t \quad \text{with} \quad p'_t := f^{-1}(s(f(p_t))), \tag{27}$$

where $f^{-1}(e)$ denotes any element of the pre-image $f^{-1}(e)$ of e , yields for (22)

$$\sigma_s((f(p_t), p_t)_t) = (s(f(p_t)), p'_t)_t = (f(p'_t), p'_t)_t,$$

where we used $f(p'_t) = f \circ f^{-1}(s(f(p_t))) = s(f(p_t))$. Thus, if $(p'_t)_t$ is a solution of T_P , the action (22) with $\tilde{\sigma}$ as defined in (27) is a symmetry of M , so that M is a model of consciousness, i.e., empirically well defined.

Thus, the idea that physical states determine qualia can indeed be formalized, even though qualia are defined to be non-collatable aspects of experience. The limitation of this approach is that the function f , being defined with respect to a particular choice of labeling of the experiencing subject, cannot be interpreted as specifying the quale that the experiencing subject experiences along with a particular physical state p . Nevertheless, the formalism allows us to treat the case that a quale, whichever one it is among the qualia in the equivalence class $[f(q)]$, is determined by the physical state p . Here, the equivalence class in question is the one defined in (7). A further analysis of the difference from a direct description will be given in Section 6.3.

We close this section by specifying the empirically well-defined part of the trajectories of a model of consciousness M . This specification is analogous to the specification of empirically well-defined sequences in (6).

As in Section 3.1, we define two trajectories $(\bar{e}_t, p_t)_t$ and $(\bar{e}'_t, p'_t)_t \in \mathcal{K}$ to be equivalent if one can be obtained from the other by relabeling the qualia of the experiencing subjects. In contrast to Section 3.1, relabeling is defined in terms of the action (22), which for non-trivial $\bar{\sigma}$ includes a relabeling of the physical states. We denote this equivalence by \sim_σ ,

$$(\bar{e}_t, p_t)_t \sim_\sigma (\bar{e}'_t, p'_t)_t \quad \text{if and only if there is an } \bar{s} \in \text{Aut}(E)^n \tag{28}$$

$$\text{such that } (\bar{e}'_t, p'_t)_t = \sigma_{\bar{s}}((\bar{e}_t, p_t)_t).$$

Note that $\bar{\sigma}$, and hence σ , depends on the laws of the model M under consideration. The empirically well-defined part of the trajectories is given by the quotient set

$$\mathcal{K} / \sim_\sigma, \tag{29}$$

i.e., by the space of equivalence classes of \sim_σ . This space describes the distinctions that remain once all trajectories are identified, which can be mapped to each other by relabeling the qualia of the experiencing subjects.

6.3. Comparison with Direct Reference

In the previous sections, we showed that non-collatability implies that models of consciousness need to carry a particular symmetry group in order to be well defined and that this allows us to address qualia proper, i.e., individual non-collatable aspects of experience. In this section, we compare the methodology so introduced with what may be called a “direct reference” of qualia: a description of experimental data or theoretical idea simply in terms of qualia’s labels, without invoking any of the mathematical details introduced in Definitions 10 and 12.

Mathematically, a direct reference of the qualia of n experiencing subjects is simply a family

$$(e_t^1, \dots, e_t^n)_{t \in \mathcal{I}}, \tag{30}$$

where \mathcal{I} is some parameter space as above. For example, this could be a time-series of reports of experiencing subjects. Whereas a direct description may ignore the mathematical details of Definitions 10 and 12, it cannot ignore the ambiguity induced by the non-collatability of labels, which is fundamental and independent of any of the mathematical tools introduced in Sections 5 and 6.

In Section 3.3, we studied this ambiguity and what it implies for references to qualia in detail and found that the corresponding well-defined statement is given by (6). In the present notation, this reads

$$(\bar{e}_t)_t \sim (\bar{e}'_t)_t \quad \text{if and only if there is an } \bar{s} \in \text{Aut}(E)^n \tag{31}$$

$$\text{such that } e_t'^i = s_i(e_t^i) \text{ for all } t \in \mathcal{I}$$

and

$$[(\bar{e}_t)_t] := \{(\bar{e}'_t)_t \mid (\bar{e}'_t)_t \sim (\bar{e}_t)_t\}. \quad (32)$$

In order to compare this with a model of consciousness, we assume that one wishes to relate (e.g., statistically analyze) the direct reference of qualia with some properties of a physical system, e.g., neural activation patterns. We assume that these properties are determined by states of the physical system and denote the state space as above by P . Thus, the data under consideration are of the form

$$(e_t^1, \dots, e_t^n, p_t)_{t \in \mathcal{I}}. \quad (33)$$

It could result, e.g., from verbal reports of experience and simultaneous fMRI scans or EEG recordings.

In (29), we found that the empirically well-defined trajectories of a model of consciousness are given by the quotient set

$$\mathcal{K} / \sim_{\sigma}. \quad (34)$$

The following lemma gives the corresponding result for the direct description.

Lemma 2. *The empirically well-defined trajectories of a direct description of qualia are given by the quotient set*

$$\mathcal{K} / \sim_{\phi}, \quad (35)$$

where ϕ is the action (21) and where \sim_{ϕ} is defined as in (28).

This lemma states what a direct description of qualia can reference in light of non-collatability. It expresses the epistemic constraints that ineffability, privateness, inaccessibility, and other characteristic features that imply non-collatability pose for any theoretical or experimental account of consciousness.

The difference between (34) and (35) is that in (34), one takes the quotient with respect to an action that generally acts non-trivially on the physical trajectories $(p_t)_t$, whereas in (35), one takes the quotient with respect to an action that acts trivially on the latter. We defer further discussion to after the proof.

Proof of Lemma 2. In terms of the notation (33), the equivalence (31) is given by

$$(\bar{e}_t, p_t)_t \sim (\bar{e}'_t, p'_t)_t \quad \text{if and only if there is an } \bar{s} \in \text{Aut}(E)^n \quad (36) \\ \text{such that } (\bar{e}'_t, p'_t)_t = \phi_{\bar{s}}((\bar{e}_t, p_t)_t).$$

According to (28), this is precisely the equivalence \sim_{ϕ} . Hence, the empirically well-defined trajectories (32) are elements of the quotient set (35). \square

The quotient space (35) contains the empirically well-defined trajectories that can be referenced by a direct description of qualia in light of non-collatability. The quotient space (34), on the other hand, gives the empirically well-defined trajectories of a model of consciousness. The difference between both is determined by the action $\bar{\sigma}$ defined in (23), which describes how the physical states transform if one relabels the qualia of an experiencing subject (passive meaning, cf. Remark 2), but also how the physical states change if the qualia of an experiencing subject change (active meaning in Remark 2). It is precisely the possibility of a non-trivial $\bar{\sigma}$ that allows the laws postulated by a model of consciousness to address individual qualia to a certain extent. The extent to which this is possible is limited by the requirement that (22) constitutes a symmetry of the model M .

Mathematically, this is reflected in the fact that elements of (35) are of the form

$$([(e_t^1)_t], \dots, [(e_t^n)_t], (p_t)_t), \quad (37)$$

where each equivalence class $[(e_t^i)_t]$ is as in (31) and (32). This coincides with our result in Section 2 on references to qualia (cf. quotient space (8)). The elements of (34), on the other hand, are not of this form: a non-trivial $\tilde{\sigma}$ results in equivalence classes in which the physical states are *mixed* with the labels of the various experiencing subjects. The empirically well-defined trajectories cannot be separated as in (37).

This means that if one chooses a direct description of qualia and investigates the relation to the physical domain, one first has to consider equivalence classes (32) (for only they are empirically well defined) and subsequently propose or analyze the relation to the physical domain. Using a model of consciousness, one may exchange this order: one may first postulate a relation of qualia and the physical domain and subsequently remove the arbitrariness of choosing labels by considering equivalence classes (28). The requirement of there being a symmetry σ simply makes sure that this relation (the law \mathcal{L} of a model of consciousness) is chosen in such a way that the second step (obtaining an empirically well-defined theory) is possible at all. A non-trivial $\tilde{\sigma}$ implies that after the second step, one does not end up with what one could have been obtained using the first procedure right away. This small detail can have large empirical consequences.

Thus, we showed that models of consciousness are more powerful than direct references in regard to the scientific study of non-collatable aspects of experience. Since collatable aspects are always contained in our formalization as a special case (trivial $\text{Aut}(E)$), we see that mathematical models of consciousness provide a suitable methodology for the scientific study of consciousness, which allows us to take the key characteristics of ineffability, privateness, and cognitive, linguistic, and communicative inaccessibility into account.

We conclude this section by remarking that Lemma 2 shows that if $\tilde{\sigma}$ is chosen as trivial in Definition 12, the result is a direct description of qualia as referenced here. Thus direct references constitute a special type of model of consciousness, and models of consciousness are in fact a genuine generalization of direct references.

7. Closure of the Physical

Great care has been taken in the previous sections to motivate all constructions in an operational, epistemological, or phenomenological way, making sure that they are independent of any metaphysical commitment. Metaphysical choices should only be made when constructing individual models of consciousness. In this section, we make a brief remark about a particularly important metaphysical choice, the closure of the physical.

The closure of the physical is often called “causal closure of the physical” [49]. It denotes the idea that “physical laws already form a closed system” [31] (p. 17) and is an important assumption that underlies many philosophical and scientific investigations of consciousness. We mention it here because, as explained in Appendix A, Chalmers’ grounding of the scientific study of consciousness needs to make use of the closure of the physical in its definitions, which limits the applicability of this grounding substantially (cf. Appendix B).

This is different for the grounding put forward in Section 3. Since none of the basic definitions or formal constructions refer to the closure of the physical in any way, this is in fact an independent assumption that one may or may not make when constructing models of consciousness. As a result, the conceptual and formal frameworks developed here are suitable also to construct models of consciousness that express metaphysical ideas such as dual aspect monism or property dualism that do not describe the physical as closed.

In terms of the formalism developed in Sections 5 and 6, the assumption of the closure of the physical can be stated in a particularly concise form.

Definition 13. *A model of consciousness M describes the physical as closed if and only if*

$$\mathcal{D}|_P = \mathcal{D}_P, \quad (38)$$

where $\mathcal{D}|_P$ is defined in (16) and where \mathcal{D}_P has been introduced in Section 5.1 to denote the solutions of the physical theory T_P that underlies the model M .

This definition says that a model of consciousness M describes the physical as closed if and only if the physical trajectories that are determined by the laws of M (as part of the solutions \mathcal{D}) are, as a set, equal to the solutions of the physical theory T_P on which M is based. Whether or not (38) is satisfied depends on the laws \mathcal{L} postulated by a particular model.

This concludes our brief excursion to metaphysics. In the next section, we review several examples and how they relate to the formalism introduced here. A conceptual point about the closure of the physical is made in Appendix B.1.

8. Examples

In this section, we review some models of consciousness that have been proposed in the literature and explain how they relate to the formalism introduced in Sections 5 and 6 and to the concepts introduced in Section 3.

8.1. Integrated Information Theory

Our first example is integrated information theory (IIT), which was proposed by Giulio Tononi in 2004 [27] and has since been developed considerably. The current version of that theory [7] consists of an algorithm whose input is a model of a physical system (together with a state of that system and including its dynamical laws) and whose outputs are formal quantities that give answers to the following three questions: 1. Which parts of the system are conscious? 2. What are they conscious of? 3. How conscious are they?

To answer the first question, the theory's algorithm identifies some (mutually disjoint) subsystems of the system. To answer the second question, for each such subsystem S , the algorithm specifies what is called a "maximally irreducible conceptual structure" (MICS). This is a mathematical object of the following kind: Let \mathcal{P}_S be the space of probability distributions (or probability measures) over the states of the subsystem S . A "concept" is an element of the space

$$\mathcal{P}_{CS} := \mathcal{P}_S \times \mathcal{P}_S \times \mathbb{R}_0^+.$$
²⁸

The maximally irreducible conceptual structure is an n -tuple of concepts, where n is determined dynamically by the theory and may vary from subsystem to subsystem, i.e., it is an element of the "qualia space" (graphical illustration in Figure 15, [7])

$$E_S := \mathcal{P}_{CS}^{\times n(S)}. \quad (39)$$

Finally, in order to answer the third question, the algorithm specifies the integrated conceptual information $\Phi^{\max}(S) \in \mathbb{R}_0^+$. In summary:

"[T]he central identity [of IIT] is the following: The maximally irreducible conceptual structure (MICS) generated by a [subsystem S] is identical to its experience. The constellation of concepts of the MICS completely specifies the quality of the experience (...). Its irreducibility Φ^{\max} specifies its quantity." [7] (p. 3).

The main papers of the theory remain somewhat silent about what exactly they take the terms "consciousness" or "quality and quantity of an experience" [7] to mean. The first observation is that ineffable aspects of conscious experience seem to have played at least a small role in the early

²⁸ In the terminology of [7], a concept consists of the maximally irreducible cause and effect repertoires of a mechanism M of S together with its integrated information $\varphi(M)$, provided that the latter is non-zero (Supplementary S1, p. 176, [8]).

development of the theory, e.g., in [27] (p. 229), Tononi notes that “[t]he notions just sketched aim at providing a framework for translating the *seemingly ineffable* qualitative properties of phenomenology into the language of mathematics” (our emphasis). As we explained in detail in Section 3.1, ineffable aspects of experience cannot be put in a one-to-one correspondence with mathematical objects, simply because two or more experiencing subjects have no means to ensure that they have associated the same ineffable aspect of experience with the same mathematical object. This was the reason for us to introduce experience spaces E via labels in Section 3.1 and what led to the requirement of there being a symmetry that describes relabeling. Following this path, we might take IIT’s “qualia space” E_S to constitute the experience space of qualia as defined in Definition 6.

This brings us to the question of how collatable relations between qualia so defined, such as the ones put forward in Phenomenological Axiom 1, are related to the mathematical structure of the space E_S . First, we note that the space E_S can be equipped with a metric: for any metric d on \mathcal{P}_S and using the usual metric on \mathbb{R}_0^+ , summation allows us to define both a metric on \mathcal{P}_{CS} and E_S . This metric can be taken to express the similarity relation in Phenomenological Axiom 1. Indeed, this may again have been a guiding idea in the development of the theory, “experiences are similar if their shape is similar” [27] (p. 228). Next, what has been called the “intensity” of an experience in Example 7 corresponds to the “quantity” of experience according to IIT. The corresponding mathematical structure is the \mathbb{R}_0^+ in which $\Phi^{\max}(S)$ takes values.

IIT arguably encodes another collatable relation between qualia: their *composition* in experience. This is usually formulated as an axiom of IIT, which states that “[c]onsciousness is compositional (structured): each experience consists of multiple aspects in various combinations” [7] (p. 2). The composition of the experience of a subsystem S in terms of more elementary experiences of the same subsystem is modeled by the Cartesian product structure of (39) in terms of the concept spaces \mathcal{P}_{CS} . One may interpret the \mathbb{R}_0^+ that constitutes the last factor of \mathcal{P}_{CS} as the intensity of the the more elementary experiences.

In summary, the basic definitions of IIT seem to fit quite well with the basic definition of the phenomenological grounding put forward in Section 3.4, and IIT can be taken to constitute a pre-model of consciousness as defined in Definition 10. However, in order to take into account the non-collatability of the corresponding aspects of experience, the symmetry (22) has to be implemented. This can be done by simply swapping the states of the physical system that give rise to particular labels $e \in E_S$: if $e_1, \dots, e_n \in E_S$ are the labels of the conscious subsystems of a system in state p_1 and $e'_1, \dots, e'_n \in E_S$ are the labels of conscious subsystems of the system in state p_2 , for any $\bar{s} \in \text{Aut}(E_S)$ that maps the former to the latter, we define the action (23) as $\bar{s}(p_1) = p_2$ (for details, cf. Example 14). Equipped with this symmetry, IIT constitutes a model of consciousness as defined in Definition 12.

We conclude this example with some conceptual remarks. First of all, we note that the phenomenological grounding approaches model-building differently than [7]. Whereas in the latter, phenomenological axioms were used to justify the definition of the algorithm, i.e., the dynamical equations of IIT, the phenomenological grounding uses phenomenological axioms to model the mathematical space associated with qualia. We saw above that an earlier version of IIT put forward in [27] was more aligned with this perspective. From the perspective of the phenomenological grounding, the main task would be to motivate the mathematical structure of (39) in more detail, e.g., one could ask why the elementary qualia are labeled by elements of the space $\mathcal{P}_S \times \mathcal{P}_S$ and not by a simpler metric space. Does the structure of the former have any phenomenological interpretation? Questions of this sort may have large consequences for the further development of the theory because the algorithm of IIT in its current form makes essential use of elements of $\mathcal{P}_S \times \mathcal{P}_S$.

Second, we remark that IIT describes the physical as closed: the dynamical evolution of the physical domain is not changed in any way by the theory. Thus, if one interprets the theory in Chalmers’

or the phenomenological grounding (or similar ones, in fact²⁹), the question arises of how the theory’s mathematical postulates (first and foremost the algorithm it specifies) can be evaluated experimentally at all, for it seems that all results one can hope to obtain from neuroscientific experiments that scan the brain (EEG, fMRI, etc.) are physical and therefore determined by the physical domain alone. Put in simple terms, one may ask what one actually learns when collecting experimental data that are in fact completely determined by the physical domain. This argument is outlined in more detail in Appendix B.2 and related to the transcendental argument against the closure of the physical given in Remark A1. In Section 8.1.1, we review a modification of IIT that avoids this problem.

8.1.1. Integrated Information-Induced Quantum Collapse

To overcome the problem mentioned in the last section, one needs to propose a model of consciousness that does not postulate the physical as closed. A first model of this kind based on integrated information theory (IIT) was given in [18]. It refers to an early version of IIT that only answers the third question of Section 8.1: How conscious is a physical system?

The physical theory on which this model is based is a quantum system with Hilbert space \mathcal{H} and Hamiltonian H . The states are density matrices $\rho \in L(\mathcal{H})$. Given a density matrix ρ , the *quantum integrated information* is defined as

$$\Phi(\rho) = \inf \left\{ S(\rho \parallel \otimes_{i=1}^N \text{Tr}_i \rho) \right\}, \tag{40}$$

where the infimum is taken over all decompositions of the Hilbert space \mathcal{H} , i.e., over all isomorphisms between \mathcal{H} and a Hilbert space of the form $\mathcal{H}_1 \otimes \dots \otimes \mathcal{H}_N$, where $\text{Tr}_i \rho$ denotes the reduced density matrix on the Hilbert space \mathcal{H}_i (i.e., Tr_i is a trace over all \mathcal{H}_j with $j \neq i$) and where, finally, $S(\rho \parallel \rho')$ denotes the quantum relative entropy defined as

$$S(\rho \parallel \rho') = \text{Tr} \rho \log \rho - \text{Tr} \rho \log \rho'.$$

Here, Tr denotes the trace over the whole Hilbert space \mathcal{H} . According to this model, (40) specifies how conscious the physical system is in the state ρ .

In order to specify how consciousness in turn influences the physical domain, the model modifies the time-evolution of the physical system; Whereas in quantum theory, the time-evolution of a closed system with Hamiltonian H is determined by

$$\frac{\partial \rho}{\partial t} = -\frac{i}{\hbar} [H, \rho],$$

this model proposes the evolution equation

$$\frac{\partial \rho}{\partial t} = -\frac{i}{\hbar} [H, \rho] + \sum_{n,m=1}^{N^2-1} h_{n,m}(\Phi(\rho)) \cdot (L_n \rho L_m^\dagger - \frac{1}{2} \rho L_m^\dagger L_n + \frac{1}{2} L_m^\dagger L_n \rho),$$

where $h_{n,m}(\Phi(\rho))$ are continuous functions of $\Phi(\rho)$, which vanish if $\Phi(\rho) = 0$ and where L_k are operators on H . This is a Lindblad evolution equation that describes, among other things, models of spontaneous wave function collapse. By choosing the functions $h_{n,m}$ suitably small, one can make sure that the model is compatible with physical experiments to date. The model is furthermore

²⁹ This problem seems to appear in any grounding that exhibits an explanatory gap. One could try to avoid the problem by interpreting (39) in terms of aspects of experience that do not exhibit an explanatory gap. This, however, would raise the question of why a novel law (the algorithm of IIT) should determine those aspects, as compared to some form of neural processing.

experimentally accessible in that it predicts a collapse rate that is dependent on $\Phi(\rho)$, rather than mass or the number of particles alone, as is the case in other spontaneous collapse models (cf. Section 5, [18]).

8.2. Global Neuronal Workspace Theory

The global neuronal workspace model (GNW) is, next to integrated information theory, the other model largely favored by neuroscientists. In contrast to the latter, however, it is usually stated directly in terms of brain physiology [3,50]. Even though this is sufficient to make some specific predictions [50], a more formal model would be desirable, not least to make a detailed comparison with IIT possible.

In what follows, we outline how a formal model could be constructed that takes as input any physical system (in a certain class of systems) and determines what the system is conscious of. To this end, we apply concepts from dynamical systems and nonbinary information processing whose connection with consciousness was recently suggested in [51]. While this attempt is *very* preliminary, the hope is that a genuine formal model can be developed along these lines in future work. A different goal was pursued in [52].

Let S be a physical system. We assume that it consists of a set N_v of components (“vertices”, representing neurons in a neuronal network), each of which is in a particular state $u_i(t)$. Here, $t \in \mathcal{I}$ denotes time, and $i \in N_v$ denotes the component in question. Furthermore, we assume that it consists of a set N_e of directed edges (representing axons, dendrites, synapses, etc., in a neuronal network), each of which may be in a state $w_l(t)$, $l \in N_e$ (e.g., representing the synaptic strength in a neuronal network). As usual, we define the parents Pa_i of the i th component to be those components from which a directed edge leads to i and assume $i \in \text{Pa}_i$. Finally, the dynamics of the system are specified component-wise by a set of “update-rules” $(f_i)_{i \in N_v}$, where f_i determines how the state $u_i(t)$ of the i th component depends on the states of its parents and the states of the edges coming from its parents at previous times.

This system has conscious representations, according to the GNW model, if two necessary conditions are satisfied. The first of these is that the system has “two main computational spaces, each characterized by a distinct pattern of connectivity” [50] (p. 56). The first computational space is a “processing network, composed of a set of parallel, distributed and functionally specialized processors or modular subsystems subsumed by topologically distinct (...) domains with highly specific local or medium-range connections” [ibid.]. The second computational space is a “a global neuronal workspace, consisting of a distributed set of (...) neurons characterized by their ability to receive from and send back to homologous neurons in other (...) areas horizontal projections through long-range excitatory axons” [50] (p. 56).

In order to construct a formal model of consciousness based on this hypothesis, a definition has to be given that specifies which structure in a physical system counts as a computational space of each kind, and which not; i.e., a definition of the necessary “patterns of connectivity” in terms of the mathematical structure of a physical system. In order to propose a such a definition, we combine the ideas of the GNW model with some of the ideas put forward in [51], using in particular the similarity between the processing network described by GNW and the “global directed network consisting of a large sparsely connected array of much smaller, irreducible subgraphs (ISGs), representing directed neuron-to-neuron connections” put forward in connection with consciousness in [51]. Here, an ISG is defined as follows.

[D1] A set $N_{\text{ISG}} \subset N_v$ of components of the physical system constitutes an irreducible subgraph (ISG) if there is a directed edge between any ordered pair of components in this set [51] (p. 25).

The similarity to the processing network of the GNW model comes about due to the major observations in [51] that each ISG “acts as an analog filter, a dynamical decision-maker (preferring one or another resonant mode), an amplifier, and a router” [51] (p. 27). In order to specify a necessary pattern for the global neuronal workspace, we simply require that this is a network with a directed edge going into and coming out of each ISG, noting that a further requirement on this network will be

added below. In summary, a proposal for the first necessary condition for the system S to be conscious may be put as follows.

[N1] The system S needs to contain two disjoint subsets $N_p, N_g \subset N_v$ of components: first, a set N_p of components whose induced subnetwork is a network of ISGs, where the inter-ISG-connections are feed-forward only; second, a set N_g of components with directed edges going from this set into all ISGs, and directed edges going to this set from all ISGs.

Clearly, this definition is preliminary and will have to be improved substantially to facilitate a full-fledged model to be defined.

In [51], the general properties of the state of ISGs were explained, which have been found in previous work. In particular, if the system satisfies a few conditions, including the feed-forward property mentioned in Definition [N1], the ISGs will typically carry out “successive pattern recognition tasks exploiting both remembered contextual information and prior expectations from past events (...) as well as the assumption of the structures (elements) that are identified at [a] previous level” [51] (p. 30). The result of this task is recorded by a dynamical attractor on the ISG’s components, which we denote by $m_k(t)$, where k indexes the ISGs in the system. These dynamical attractors represent “perceived sources/objects, (...) events, (...) narratives, (...) scenarios” (Figure 2, [51]). Combining these results with the idea that “[t]he entire workspace is globally interconnected in such a way that only one such conscious representation can be active at any given time” [50] (p. 58), we arrive at a proposal for the second necessary condition for the system to be conscious:

[N2] The induced subnetwork of N_g needs to be such that at any time t , its state “represents” only one of the ISGs’ dynamical attractors $m_k(t)$.

Here, one could, e.g., define the term “represent” to mean that at any time t , the state of the network is (essentially) equal to one of the states $m_k(t)$, but other more realistic choices might be possible.

If both necessary conditions [N1] and [N2] are satisfied at a particular time t , the GNW model claims that the system S is conscious of the “perceived object, event, narrative or scenario” $m_k(t)$ represented in the global workspace network N_g . Due to the directed edges from N_g to the ISGs, the state of the ISG k may be made “directly available in its original format to all other workspace processes” [53] (p. 15).

Clearly, this outline leaves open various questions; most notably, the question of how modes $m_k(t)$ of ISGs may relate to experience. Whereas IIT’s qualia space (39) has some structure that relates to phenomenology, it is highly questionable whether this can be asserted of the states of ISGs, which behave generically like a small number of “monotonically increasing phase variables” [51] (p. 25). This very question arises also, albeit in a more indistinct form, if GNW is formulated in terms of neuronal architecture: How does a “piece of information selected for its salience or relevance to current goals” [50] (p. 56), which is really just a state of some subset of the brain’s neurons, relate to experience? A proposal to this extent is presented in Section 8.4.

Whereas from a formal modeling perspective, there is some space for further development of the GNW model, it does seem to capture essential neuroscientific evidence in a simple and very plausible hypothesis: the global workspace. This idea might ultimately be combined with ideas of IIT or other models to give an explicit account of how the state of the global workspace relates to experience as we find it.

8.3. Conscious Agent Networks

A model that is based on idealistic metaphysics was developed in [20]. The underlying idea is that what exists are interacting conscious agents, each of which has a fundamental capacity to perceive, decide, and act and that the interaction between these conscious agents seems to each as if there is an external outside world. For simplicity, in what follows, we explain a slightly more general version of the model than presented in [20].

In order to explain a single conscious agent C , we first assume that there is a space W that is external to the conscious agent. One may think of this as states of some “world” that the agent perceives, but in fact, this space is constituted via interactions with other conscious agents, as explained below. Given this space W , a conscious agent is modeled as a five-tuple³⁰ $C := (X, G, P, D, A)$, where X and G are spaces and P, D, A are maps, interpreted as follows:

- ▶ X is a space that describes possible experiences of the conscious agent. Each element $x \in X$ represents a particular experience.
- ▶ G is a space that describes dispositions or intentions to act. Each element $g \in G$ corresponds to an action the agent has decided to carry out.
- ▶ $P : W \rightarrow X$ is a map that describes the agent’s “process of perception” [20] (p. 6). It specifies what the conscious agent experiences in response to the “world” being in a particular state $w \in W$.
- ▶ $D : X \rightarrow G$ is a map that models how the experience of the agent determines its disposition for an action, i.e., “the process of decision [in which] a conscious agent chooses what actions to take based on the conscious experiences it has.” (ibid.).
- ▶ $A : G \rightarrow W$ describes how the agent’s disposition for an action “is carried out”, i.e., how it affects the world: “In the process of action, the conscious agent interacts with the world in light of the decision it has taken, and affects the state of the world” (ibid.).

The structure of the spaces W and X , as well as the definitions of the maps P, G , and A are not fixed by the theory, but need to be chosen according to the application.³¹ Based on such a choice, the model specifies the dynamically possible trajectories $(x(t), g(t), w(t))_{t \in \mathcal{I}}$ as those trajectories that satisfy

$$(x(t+1), g(t+1), w(t+1)) = (Pw(t), Dg(t), Aw(t)),$$

where t is chosen as a discrete time parameter, i.e., $\mathcal{I} := \mathbb{Z}$.

The central hypothesis of this theory is called “conscious realism”: That “[t]he world W consists entirely of conscious agents” [20] (p. 7). This hypothesis is implemented via networks of conscious agents.

In order to describe a network of n conscious agents, we first assume that for every conscious agent, a space X_i of possible experiences and a space G_i of dispositions to act is given, as well as a “decision map” D_i as introduced above. The “external world” of the i th conscious agent is defined to be the product of the action spaces of all other conscious agents, i.e.,

$$W_i := G_1 \times \dots \times G_{i-1} \times G_{i+1} \times \dots \times G_n.$$

This choice is motivated by the idealistic idea that what exists are only experiences and dispositions to act and that the dispositions to act of some agents determines the experience of others, i.e., the process of the perception of the i th conscious agent is, in the case of a network of conscious agents, given by a map

$$P_i : W_i \rightarrow X_i.$$

This allows us to define the dynamically possible trajectories of the network of conscious agents via

$$x_i(t+1) = P_i(g_1(t), \dots, g_n(t)) \quad \text{and} \quad g_i(t+1) = D_i x_i(t).$$

³⁰ In [20], the specification furthermore included an integer N that counted perception-decision-action cycles and hence acted as a type of internal “psychological” time, which however we simply replace by the usual parameter $t \in \mathcal{I}$.

³¹ In [20], some general assumptions were made: the spaces W, X , and G were assumed to be measurable spaces, and the maps P, D , and A were chosen to be Markovian kernels, so that for every element of their domain, each map yielded a probability distribution on their co-domain.

If P_i is a partial function defined only for some $G_j \in W_i$, the i th agent is only able to perceive the dispositions to act of the corresponding other conscious agents. Various concrete proposals for how to choose P_i were discussed in [20] (p. 7ff.).

Due to the identification of the “outside worlds” W_i of each conscious agent with the dispositions to act of others, the action map $A_i : G_i \rightarrow W_i$ is not necessary to define the dynamics. In order to satisfy the definition of a conscious agent given above, one may define it formally as the map that takes $g_i(t)$ to $w_i(t+1) = (g_1(t+1), \dots, g_n(t+1))$. In simple cases (e.g., involving two conscious agents [20]), this definition can be fleshed out in terms of combinations of inverses of D and P . In general, it may require A_i to be time dependent.

In summary, the various objects the theory assumes in a particular application determine (possibly in a probabilistic manner) the dynamics of a network of conscious agents. The goal, then, is to specify plausible assumptions that allow us to deduce formally that “the perception of objects and space-time can emerge from such dynamics” [20] (p. 1) and to “explore [the model’s] theoretical implications in the normal scientific manner to see if they comport well with existing data and theories, and make predictions that are novel, interesting and testable” [20] (p. 7).

An early example of a result of this kind was given in [20] (p. 13ff.). In a nutshell, it is shown that if the state spaces X_i and G_i are finite, the dynamics of a network of two conscious agents can be described in terms of an object that bears some similarity to a quantum-mechanical wave function of a free particle.

From the perspective of models of consciousness as defined in Section 6, two crucial questions arise:

- (a) Whether the model would like to address aspects of experience that are non-collatable.
- (b) Whether the theory would (eventually or in principle) like to make predictions with respect to experiments that involve (reports of) conscious agents.

An affirmative answer to the first question might be indicated by the remark that the “qualia X of a conscious agent C are private, in the sense that no other conscious agent C_i can directly experience X ” [20] (p. 14). If this is indeed the case, the mathematical structure of the spaces X_i (and possibly also of G_i , if one holds that intentions to act are also experiences of some sort) could be defined based on a phenomenological analysis as explained in Section 3.1. This would, in particular, dismantle the objection that the “definition of conscious agents could equally well-apply to unconscious agents [so that the] theory says nothing about consciousness” [20] (p. 14).

More importantly, if the theory also answers affirmatively the second question, the results of Section 6 show that a further mathematical structure is necessary to ensure that the model is empirically well defined (Lemma 1).

8.4. Expected Float Entropy Minimization

One of the largest questions at present left open by the GNW model (Section 8.2) is how the state of the global neuronal workspace, ultimately a collection of states of individual neurons, relates to experience. Questions of this kind are addressed by the expected float entropy (EFE) model developed in [19]. In short, this is a proposal for how (probability distributions of) brain states determine relations among qualia.

In what follows, we review the definition of this model. Every brain state is assumed to consist of individual elements, each of which can be in a particular state. We denote the set of all elements (“nodes”) by S and the space of states of each node by V and assume both are a finite set. A brain state is thus a map

$$s : S \rightarrow V, \quad (41)$$

e.g., in a neural network, S is the set of neurons, and V is the set of possible states of each neuron. If applied to the GNW model as outlined above, S is the set N_g of nodes, and V denotes the corresponding space of states. In [19], s was called a “data element”, but we will refer to s simply as a “state”.

Let $\Omega_{S,V}$ denote the space of all states. We assume that a probability distribution p is given over $\Omega_{S,V}$. The probability $p(s)$ can be interpreted as the probability of the brain being in state s .

A weighted relation on a set S is a map $R : S \times S \rightarrow [0, 1]$. Given a set of states with the corresponding probability distribution, the theory developed in [19] allows one to determine two weighted relations R and U , where R is a weighted relation on the set S of nodes and where U is a weighted relation on the possible states V of each node. We will discuss the interpretation of R and U at the end of this example.

The theory determines both U and R as follows. For any state $s \in \Omega_{S,V}$, the composition $U(s(\cdot), s(\cdot))$ is a relation on S , which we denote as $U \circ s$. Define the float entropy fe and expected float entropy efe as

$$fe(R, U, s) = \log_2 |\{ \tilde{s} \in \Omega_{S,V} \mid d(R, U \circ \tilde{s}) \leq d(R, U \circ s) \}| \quad (42)$$

$$efe(R, U, p) = \sum_{s \in \Omega_{S,V}} p(s) fe(R, U, s) \quad (43)$$

where $s \in \Omega_{S,V}$, d is a distance function on the weighted relations on S , and where $|A|$ denotes the cardinality of a set A . The theory proposes “that a system (such as the brain and its subregions) will define U and R (up to a certain resolution) under the requirement that the efe is minimized.”, i.e., U and R are defined via

$$efe(R, U, p) = \min_{\bar{R}, \bar{U}} efe(\bar{R}, \bar{U}, p), \quad (44)$$

where the minimum is taken over all relations \bar{R} on S and all relations \bar{U} on V (the existence or uniqueness of minimizers was not discussed in [19]).

Concerning the interpretation of R and U , the theory proposes that if “a brain state is interpreted in the context of all these relations (...), the brain state acquires meaning in the form of the relational content of the experience”. If applied to the visual cortex, the theory aims to explain “perceived relationships between different colors, the perceived relationships between different brightnesses, and the perceived relationships between different points in a person’s field of view (giving geometry)”.

These interpretations were supported by several examples in [19], where the theory was applied to pictures, so that S was the set of all pixels and V described the possible color values at each pixel, which implied that U was a relation between color values and R a relation between pixels. The support for these interpretations becomes more difficult when the theory is being applied, e.g., to the visual cortex, for in this case, U is a relation on the states of the nodes where the nodes could be individual neurons or tuples of neurons in the visual cortex for example and R is a relation on the set of these nodes, making it somewhat unclear why in this case, a relation U might give an explanation of, e.g., why “blue appears similar to turquoise but different from red”.

One can, however, simply take the theory at face value by accepting that the relations of U and R , whichever mathematical form they take, are (describing) non-collatable aspects of experience and that U and R are (describing) the relations between them. Here, the non-collatability is essential, for otherwise, the identity of some collatable aspect of experience and elements of the set V or S would be questionable. In short, one may assume that the relations R and U correspond to the structure of the experience space E , which describes experience.

Several interesting questions are raised by this model. First of all, we note that since the model aims to explain the relations between aspects of experience, it is fully compatible with a direct description of qualia as discussed in Section 6.3 and does not aim for a description of qualia *sensu stricto*. This raises the question of whether this model is an alternative to, or rather a complement of, models that do intend to describe qualia *sensu stricto*, such as, e.g., integrated information theory.

One might conjecture that the relations among aspects of experience might in fact be nuanced enough to allow us to identify individual qualia by specification of the relations; in other words, that all orbits of the automorphism group (5) are trivial. Whether or not this is the case is a phenomenological question, which needs to be answered by a systematic account of the relations between qualia found in

experience and is a priori to any model-building process (just like general properties of an explanandum have to be fixed prior to an explanation). However, since the EFE model actually specifies the relations between aspects of experience, one can also study which answer the model itself gives to this question. The upshot of this analysis, which is presented in the next paragraph, is that if the probability distribution p is invariant with respect to a transformation (or permutation) of states, which is often the case, the model does in fact specify relations whose automorphism group has non-trivial orbits.

Consider a bijective transformation (permutation) of states $\sigma : \Omega_{S,V} \rightarrow \Omega_{S,V}$, which can be specified in terms of a bijective transformation $\sigma_S : S \rightarrow S$ of nodes and in terms of a bijective transformation $\sigma_V : V \rightarrow V$ of node-states, i.e., $\sigma(s) := \sigma_V \circ s \circ \sigma_S$. The probability distribution p is *invariant* with respect to this transformation if $p = p \circ \sigma$, i.e., if the transformation maps states s to states $\sigma(s)$ that have the same probability as the former, $p(\sigma(s)) = p(s)$. Defining the transformation of the relations U and R as

$$U'(\cdot, \cdot) := U(\sigma_V^{-1}(\cdot), \sigma_V^{-1}(\cdot)) \quad \text{and} \quad R'(\cdot, \cdot) := R(\sigma_S(\cdot), \sigma_S(\cdot)), \quad (45)$$

and using the fact that the metric d is chosen as one of the d_n metrics in [19], (p. 127) i.e., involves summation over all elements of $S \times S$, (42) yields that $\text{fe}(R, U, s) = \text{fe}(R', U', \sigma(s))$. Using the invariance of p and (43), this gives

$$\text{efe}(R, U, p) = \text{efe}(R', U', p).$$

This implies that for any minimizer R, U of (44), the pair R', U' is a minimizer of (44) as well. In other words, the theory only determines minimizers up to transformations (45). Assuming the uniqueness of minimizers, this in turn implies that the minimizing pair U, V satisfies

$$U(\cdot, \cdot) = U(\sigma_V(\cdot), \sigma_V(\cdot)) \quad \text{and} \quad R(\cdot, \cdot) = R(\sigma_S(\cdot), \sigma_S(\cdot)), \quad (46)$$

so that σ_V and σ_S are relation-preserving bijections, i.e., non-trivial elements of the automorphism group of the spaces (V, U) and (S, R) , respectively.

Another interesting question is which part of the brain generates those relations between aspects of experience that we find in experience. This is, to a large extent, a question that could be answered by simulations of the brain's neuronal network. If it turns out that these relations can be reproduced better by a distributed network, this model may actually be compatible, or even taken as support of, the global neuronal workspace hypothesis. The underlying challenge here is, of course, to identify the weighted relations R and U between (states of) neurons with the manifold relations between aspects of experience. This identification may also hinge on how the probability distributions $p(s)$, which is the only data that enter the definition of R and U , is interpreted when applied to the brain.

We conclude that this theory is an interesting approach to the mind-matter relation, which might complement more neuroscientific approaches such as the global neuronal workspace model. Depending on whether a phenomenological analysis confirms that there are qualia that cannot be distinguished by mere reference to collatable relations, the model may or may not have to be extended in some form or the other to talk about the hard problem of consciousness.

9. Conclusions and Outlook

Consciousness is in the focus of research projects around the globe. Empirical, as well as theoretical projects aim to investigate different aspects of experience, ranging from access consciousness or the unity of a conscious scene to phenomenal consciousness or the first-person-perspective [2]. The starting point of this paper was the observation that if an aspect of experience is under investigation that cannot be identified over several experiencing subjects (which cannot be *collated*), special care is necessary. Any reference to such aspects of experience, be it in a theoretical account or when giving reports, is

ambiguous, and this ambiguity may lead to ill-defined models, erroneous empirical predictions, and misinterpretation of experimental data.

A detailed summary of our results is given in Section 2. In order to develop a well-defined scientific methodology that can be applied to all aspects of experience (collatable or not), we used basic *phenomenological axioms* to specify how a formal representation of experience can be constructed. The result was a mathematical space that represents some parts of experience (such as visual experiences or auditory experiences) *completely*, including both the usual objects of investigation in cognitive neuroscience, as well as qualia. This avoids the usual hard cut between parts of experience that represent a difficulty for the contemporary scientific methodology and parts that do not. Both are interwoven in our formal representation, similarly to position and momentum being two aspects of a quantum state.

Subsequently, we showed that this mathematical representation of experience allows to quantify the ambiguity involved in any reference to conscious experience precisely. This is sufficient to avoid the problems mentioned above and yields a formal mathematical toolbox that could be applied in empirical or theoretical investigations of consciousness.

In the second part of the paper, we investigated how individual non-collatable aspects of experience (qualia “sensu stricto” [7]) can be studied scientifically. Since there is a fundamental explanatory gap, this question may be considered as equally relevant to the one addressed in the first step.

The main result of the second part was that formal models of consciousness can address individual non-collatable aspects of experience *if and only if* they carry a specific symmetry group related to the mathematical representation of experience referred to above. Because of the mathematical details of the action of this symmetry group, models of consciousness can be used to construct empirically well-defined theories of how qualia sensu stricto relate to the physical domain, *despite* the ambiguity inherent in any reference to them.

The results of this paper constitute a grounding of the scientific study of consciousness, which is an alternative to other groundings currently in use. It offers a thorough conceptual and mathematical framework in light of which existing models of consciousness can be interpreted and improved, and based on which new models can be constructed.

This constitutes a first step in developing a full-fledged conceptual and mathematical foundation for models of consciousness. Further work is necessary to investigate which mathematical structures are implied by other key characteristics of conscious experience, most notably the various connotations of subjectivity and intrinsicity.

Funding: This research received no external funding.

Acknowledgments: I am grateful for the questions and comments received during presentations of parts of this work at the the LPS Colloquium of the Munich Center for Mathematical Philosophy, the Mathematical Institute of the University of Göttingen, the Institute for Theoretical Physics of the University of Hanover, the Modelling Consciousness Workshop in Dorfgastein, the Models of Consciousness Conference in Oxford, the Online Seminar Progress and Visions in the Scientific Study of the Mind-Matter Relation, and the Conceptual Foundations of Science Workshop in Tegernsee. Most of this work was carried out while I was employed at the Institute for Theoretical Physics of the Leibniz University of Hanover, and I am very grateful for having had the opportunity to do so.

Conflicts of Interest: The author declares no conflict of interest.

Appendix A. Chalmers’ Grounding of the Scientific Study of Consciousness

The most prominent grounding of the scientific study of consciousness has been developed by David Chalmers in [30]. Since it is the blueprint of the grounding proposed in Section 3, we review its essential definitions. Note, however, that the following outline of Chalmers’ grounding is intended to highlight the relations among various constituents of his grounding and is not intended to be of an introductory nature. A good and short introduction to this topic is (Chapter 1, [31]).

First, we note that Chalmers' definition of 'physical domain' includes what is often called 'material' or 'physical' configurations, such as neurons or brain tissue, as well as more fundamental physical notions such as "mass, charge, and space-time" [31] (p. 17) or "atoms, electro-magnetic fields, and so on" [30] (p. 71). We thus define the term '*physical domain*' to refer to all those phenomena which are currently considered to be the subject of a natural science (physics, chemistry, earth science, biology, etc.). Chalmers assumes that:

(A1) "The physical domain is causally closed." [30] (p. 161)

"For every physical event, there is a physical sufficient cause." [30] (p. 125)

Central to Chalmers' grounding are the terms 'function' and 'structure'. "Here 'function' is not used in the narrow teleological sense of something that a system is designed to do but in the broader sense of any causal role in the production of behavior that a system might perform" [31] (p. 6). The term 'structure' is used in a spatiotemporal sense. Together, they constitute, according to Chalmers, the notion of explanation which is used throughout contemporary science: "One can argue that by the character of physical explanation, physical accounts explain *only* structure and function, where the relevant structures are spatiotemporal structures, and the relevant functions are causal roles in the production of a system's behavior." [31] (p. 105f.)

We denote this notion of explanation by (E1). Assuming some laws or theories relating to the physical domain as given (= accepted by the scientific community by and large) and referring to them as 'accepted theoretical notions', (E1) might be put as follows:

(E1) An explanation specifies the function and structure of an explanandum in terms of the the function and structure of accepted theoretical notions.

The crucial aspect of Chalmers' grounding is to establish, in a consistent and explicit way, that there are phenomena, related to consciousness, to which no function or structure (as defined above) can be associated. It follows that these phenomena cannot be explained according to (E1) and hence, if (E1) indeed captures all notions of explanations which are used throughout contemporary science, that they cannot be explained by contemporary science.—There is an "explanatory gap" [30,43]. Chalmers refers to these phenomena as "phenomenal concepts", "phenomenal qualities" or "qualia" [30].³² We refer to these phenomena as 'phenomenal aspects of consciousness':

(D1) *Phenomenal aspects of consciousness* are those aspects of conscious experience which do not have a function or structure, where 'function' and 'structure' are as defined above.

The key requirement for this definition of what is to be studied by a science of consciousness to make sense is to establish that there are aspects of experience which satisfy (D1), i.e., which neither have a spatio-temporal structure nor a causal role in the production of behavior. It is the second requirement with respect to which (A1) is crucial, for (A1) can be utilized to argue that nothing non-physical can have a causal influence on the physical domain. Therefore, all aspects of experience which do not have a spatio-temporal structure (e.g., in the Cartesian sense of being non-extended in space and space-time) automatically satisfy (D1). We will not review the various arguments which aim to prove the existence of phenomenal aspects of consciousness at this point.

Put in terms of Definition 1, what is to be studied in the scientific study of consciousness are, according to this grounding, phenomenal aspects of consciousness and their relation to the physical domain. Since these are, by definition, not accessible to the usual scientific methodology, Chalmers proposes that the task of a science of consciousness is to find what he calls "psychophysical laws" [30] (p. 127) which relate the physical domain to phenomenal aspects of consciousness. Due

³² In [31], he prefers to use the term 'experience': "Sometimes terms such as 'phenomenal consciousness' and 'qualia' are also used here, but I find it more natural to speak of 'conscious experience' or simply 'experience.'" [31] (p. 5).

to Assumption (A1) and an underlying stance on the nature of causality “[t]hese laws will not interfere with physical laws; physical laws already form a closed system. Instead, they will be *supervenience laws*, telling us how experience [= phenomenal aspects of consciousness] arises from physical processes” [30] (p. 127). In combination with (E1), this implicitly points at the major parts of the methodology to be used according to this grounding.

Chalmers’ grounding raises several questions related to the definition and ontological status of causality, to the validity of Assumption (A1), to the nature of experiments in his grounding and to the validity of the subsumed notion of explanation, which we discuss in Appendix B. The upshot is that there are severe conceptual problems which make it questionable whether a scientific research program based on this grounding can be carried out at all.

Furthermore, any scientific approach based on this grounding faces the question of which mathematical structure one is to use in order to describe phenomenal aspects of consciousness when formulating “psychophysical laws” [30] (p. 127). Whereas the physical domain comes with a clear-cut mathematical structure, Chalmers’ grounding merely asserts that the phenomenal aspects form a set and offers no systematic way of tying additional mathematical structure to the phenomenology of experience.

This strongly suggest the construction of other groundings of the scientific study of consciousness. In Section 3, we have introduced a possible alternative which avoids the above-mentioned problems. Whereas this grounding breaks with several of Chalmers’ main ideas, it retains the key idea of addressing an explanatory gap with mathematical tools.

Appendix B. Conceptual Problems of Chalmers’ Grounding

In this appendix, we briefly discuss several conceptual issues of Chalmers’ grounding. These issues are not motivated by metaphysical considerations and are not intended to have metaphysical implications; they simply arise if one wishes to carry out a scientific investigation of consciousness based on Chalmers’ grounding. Appendixes B.1 and B.2 are most crucial and might make it impossible to apply the grounding.

The abbreviations used below have been introduced in Appendix A. For reasons explained in Appendix B.4, we assume that Assumption (A1) is intended to express the fact that “physical laws already form a closed system” [30] (p. 127).

Appendix B.1. Closure of the Physical

Much has been written about Assumption (A1) both by David Chalmers himself (e.g. Chapter 5, [30] or Chapters 8 and 9, [31]) and by others (e.g., [54] or [49]). As noted in Appendix A, this assumption is crucial for Chalmers’ grounding in order to establish that there are aspects of experience which satisfy (D1).

To date, there is no valid argument which shows that Assumption (A1) is wrong, i.e., that the physical laws of nature cannot “form a closed system”. On the other hand, there also is no valid argument that shows that Assumption (A1) is right, i.e., that the physical laws of nature must form a closed system. (Note that no reasons are given in either [30] or [31] for why Assumption (A1) should hold true.) This assumption also cannot be backed by analyzing opinions or strategies of working physicists, for most physicists are prepared to accept, or even try to find, modifications of the known laws of physics due to yet unknown phenomena (e.g., related to dark matter, to quantum gravity or to dynamical collapse theories, to name just a few). They do not assume that the known physical laws form a closed system. “Physics itself does not imply its own causal closure nor is there any proof within physics of its own completeness, so CoP [causal closure of physics] must be a metaphysical principle” [49] (p. 45).

Based on this state of affairs, one might think that both Assumption (A1) as well as its opposite should be compatible with a scientific approach to consciousness. However, this is not the case, as the following remark shows. Despite the fact that the physical laws of nature may form a closed system, it

seems that Assumption (A1) is incompatible with a scientific approach to investigate consciousness because it violates a necessary condition for the possibility of the latter.

Remark A1. *The phenomenological grounding developed in Section 3 allows one to construct models of consciousness which postulate the physical as closed just as well as models which do not postulate the physical as closed (several examples of both are given in Section 8). However, it seems that in both Chalmers' and the phenomenological grounding of the scientific study of consciousness, it does not make sense to assume the closure of the physical because it violates a necessary condition of the possibility of the scientific study of consciousness itself.*

The goal of this remark is to explain in detail why this is so. To this end, we use the symbol Q to denote that which is to be studied according to the grounding at hand: In the case of Chalmers' grounding (CG), Q refers to qualia as defined in (D1) in Appendix A, whereas in the case of the phenomenological grounding (PG), Q refers to qualia as defined in Definition 6. In both cases, Q thus refers to aspects of conscious experience.

The above claim rests on two premises. First, that a scientific study of consciousness is possible only if scientists can communicate about Q at least to some extent. E.g., they need to be able to agree on Q 's definition and existence, need to be able to communicate certain general properties of Q (such as Phenomenological Axioms 1, 2 or 3 in the case of PG) or need to be able to record and exchange data related to Q . This is the necessary condition for the possibility of the scientific study of consciousness referred to above, which we abbreviate by NC.

The second premise is that communication is always mediated via communication channels C which are elements of the physical domain. To give some examples, consider verbal communication, which is mediated via sound waves, digital communication, which is mediated via electromagnetic signals, or printed texts, where communication is mediated via arrangements of molecules and electromagnetic fields.

Due to the second premise, an assumption concerning the closure of the physical (ACoP) has something to say about communication channels and therefore also about communication itself. If, in the grounding at hand, ACoP is fleshed out in such a way that it restricts the relation between Q and communication channels C to such an extent that communication about Q is impossible, the above claim holds: By the first premise, this implies a violation of a necessary condition of the possibility of a scientific study of consciousness.

Clearly, whether or not this is the case depends on what one takes to constitute 'communication' and which conditions one posits as necessary for something to count as 'communication about Q '. To find proper answers to these questions is of course the goal and task of various parts of philosophy. However, by restricting to a very simple situation, we may hope to work with a necessary requirement for 'communication about Q ' to be possible which is acceptable independently of which notion of communication one prefers.

The simple situation which we consider is the prototypical scenario of the mathematical theory of communication.³³ I.e., we consider a situation where one experiencing subject S_1 (the 'sender') formulates a message m_1 which expresses some properties of her experience of Q , such as which particular phenomenal quality she has experienced (in the case of CG) or whether two qualia are similar (in the case of PG). (If this is impossible, i.e., if the assumptions of a grounding are such that an experiencing subject cannot formulate a message which expresses some properties of her experience of Q , this grounding violates the necessary condition NC as claimed. This may be the case for CG, cf. (Chapter 9, [31]).) Subsequently, this message is being transferred via a communication channel C to another experiencing subject S_2 (the 'receiver'), who after decoding the channel's signals obtains a message m_2 . We abbreviate this scenario by MTCP ('p' for 'prototypical').

We denote properties of S_1 's experience of Q by q and states of the communication channel C by c . In what follows, we consider functional dependencies between the quantities q , c , m_1 and m_2 . In order to define what constitutes a functional dependency both mathematically and conceptually, we refer to the groundings under consideration:

³³ This is the original (and arguably more adequate [55]) name for 'information theory' [56].

Both CG and PG's specification of the task of the scientific study of consciousness includes the formulation of laws or theories concerning the relation of Q with the physical domain. Given enough further specifications (such as a model of the communication channel or more comprehensive physical laws), these laws or theories should be applicable to the MTCp setup. I.e., we may assume that both PG and CG allow one to construct (or even to deduce) mathematical models of the MTCp setup. The details of any such model of course depend on various factors, most importantly on which psychophysical laws (CG) or models of consciousness (PG) one considers. All that matters at this point is that given any such model, we may identify functional relationships between the quantities q , c , m_1 and m_2 :

(F1) A quantity $a \in \{q, c, m_1, m_2\}$ is functionally dependent on a quantity $b \in \{q, c, m_1, m_2\}$ according to some model of the MTCp setup iff according to this model, a is a non-constant function of b .³⁴

The reasons for focusing on functional dependency in order to argue for the main claim of this remark are threefold. The first reason is that in both CG and PG, ACoP implies a restriction of the functional dependencies which may hold between the quantities q , c , m_1 and m_2 . Consider first CG. Here, the various formulations of ACoP differ slightly depending on whether they utilize a notion of causality or not. However, it seems fair to say that they all intend to express the central claim that "physical laws already form a closed system" [30] (p. 127). Together with the second premise introduced above, this implies that in any model of the MTCp setup based on CG, c cannot functionally depend on q . In PG, ACoP implies that any state c of the communication channel is determined completely by the dynamics of the physical theory T_P , which does not include Q . Therefore, as is the case for CG, in PG ACoP also implies that c cannot functionally depend on q :

(A2) In both Chalmers' grounding (CG) and the phenomenological grounding (PG), the assumption of the closure of the physical (ACoP) implies that the states c of communication channels cannot be functionally dependent on q .³⁵

The second reason is that functional dependency also seems to allow us to formulate a fundamental necessary condition for 'communication about Q ' to be possible:

(C1) A necessary condition for communication between S_1 and S_2 about Q is that m_2 may depend functionally on q .

The third reason, finally, is that the MTCp is intended to express functional relationships in the first place. In particular, it can be taken to imply by definition that m_2 is functionally dependent only on c and acquires additional functional dependencies only via c 's functional dependencies.

This concludes the reasoning: A necessary condition of communication about Q in the MTCp setup is that m_2 is functionally dependent on q . By definition of the MTCp setup it can only be functionally dependent on q via c . CG and PG's ACoP however imply that c cannot be functionally dependent on q . Therefore, a necessary condition of communication about Q is violated, which by the first premise above is a violation of a necessary condition for the possibility of a scientific study of consciousness.

Clearly, this reasoning does not yet constitute a formal argument. Several of its suppositions have to be checked carefully for hidden assumptions, which goes beyond the scope of this remark.³⁶ Nevertheless, it is of

³⁴ Here, by 'constant function' we simply refer to functions which are formally dependent on b but whose value remains the same independently of which value b takes. E.g., $f(x, y) := x$ is a constant function of y .

³⁵ We emphasize again that the notion of 'functional dependence' is defined by the respective grounding under consideration. Thus it has a somewhat nomological flavor and does not express, e.g., simple covariation. The fact that both groundings contain notions of functional dependence is what allows the present argument to be stated in a comparably concise form.

³⁶ To give one example: As explained in Footnote 35, this argument rests on the notion of functional dependency contained in CG and PG in virtue of psychophysical laws or models of consciousness. In using these, we have avoided the difficult question of what a functional dependency actually expresses (i.e., how it is supposed to be defined and interpreted). E.g., when considering q , c , m_1 and m_2 as variables, which sort of possible worlds do they describe? Logically possible worlds, conceivable worlds, some sort of nomologically possible worlds? In what way can the assumptions of a grounding restrict these possible worlds and what effect does this have on functional relationships?

importance both with respect to Chalmers' grounding (where it raises a thorough problem) and with respect to the phenomenological grounding (where it is a basis for potential empirical predictions).

We close this remark by pointing out that arguments which try to prove that the closure of the physical cannot hold in light of empirical facts about our experience (most notably written or verbal statements which express some fact about conscious experiences, e.g., bafflement about why consciousness exists [54]) do not seem to be valid. The problem is simply that we may appear to be expressing facts about our conscious experience while in fact we are not. Similarly, we may appear to be communicating about consciousness while in fact we are not. This is the basis for Chalmers' efforts to develop a theoretical account of how judgments or statements about consciousness can be accounted for despite the closure of the physical, cf. (Chapter 5, [30]) and (Chapters 8 and 9, [31]).

In contrast, the claim proposed in this remark simply represents a transcendental argument: Independently of whether reality satisfies the closure of the physical or not, it does not make sense to engage in a scientific study of consciousness if one postulates the physical as closed, because the latter violates a necessary condition of the possibility of the former.

Appendix B.2. Experiments

An issue also arises with respect to experiments if one postulates that "physical laws already form a closed system" [31] (p. 17): Almost all experiments one might wish to perform are rendered meaningless. The reason is simply that most experimental data (fMRI scans, EEG signals, verbal reports, etc.) is stored on physical devices (hard drives, paper, sound waves, etc.) and hence subject to physical laws. If these are postulated to "form a closed system" it follows that the experimental data must be determined by these physical laws alone, independently of which "psychophysical law" [30] (p. 127) correctly describes how phenomenal properties depend on physical properties.

To see this in more detail, let us assume that two different psychophysical laws \mathcal{L} and \mathcal{L}' have been proposed. The idea behind Chalmers' and in fact any conception of the *scientific* study of consciousness is that experiments have to be carried out in order to evaluate which of the proposals better describes reality. Accordingly, assume that an experiment has been designed and carried out which purports to answer this question, e.g., by checking predictions based on the laws \mathcal{L} and \mathcal{L}' . Finally, denote by d the dataset produced by this experiment.

The term 'data' is applicable to any "putative fact regarding some difference or lack of uniformity within some context" [55], so that one might consider the case where d actually consists of non-physical quantities, e.g., of differences in one's own experience. However, as soon as the data is stored or processed as usual, e.g., on a hard drive in order to perform statistical analysis, the differences in question have been transformed into "difference or lack of uniformity" of physical quantities. Since almost all experiments, even when dealing with verbal reports or similar indications of conscious experience, perform some sort of statistical analysis, it seems that in almost all experiments, d eventually is a physical data set in this sense.³⁷ It is 'stored via' physical quantities.

If one assumes that "physical laws already form a closed system" [31] (p. 17), it follows that all physical quantities, as well the differences or lack of uniformity they exhibit, are determined by the laws of physics alone. Applied to the physical quantities on which d is stored, this statement literally says that the data d is determined by the laws of physics alone. Put differently, due to the fact that the experimental data d is stored on a physical device, closure of the physical implies that the data d is completely independent of whether \mathcal{L} or \mathcal{L}' or some completely different psychophysical law best describes how experience arises from physical processes.

³⁷ If one assumes that communication between two experiencing subjects is mediated via communication channels that are part of the physical domain (cf. Remark A1), it follows that every scientifically meaningful data needs to be transformed into physical data at some point.

Thus, in summary, the closure of the physical implies that whatever experiment one performs in order to evaluate psychophysical laws, if it yields data that is stored on physical devices, the result of the experiment is independent of how experience actually arises from physical processes, i.e., independent of that which it seeks to study.

This conclusion holds true even if we concede that every experiencing subject might interpret the physical dataset d in terms of his/her own experience, so as to give meaning to this set in a way that a philosophical zombie might not, simply because if d is independent of which law E best describes how phenomenal properties arise from physics, the meaning a scientist gives to d will generally be too.³⁸

Appendix B.3. Subsumed Notion of Explanation

Chalmers' grounding builds on, axiomatizes and extends the notion of an explanatory gap that has been introduced by Joseph Levine in [43]. To this end, Chalmers claims that a specific account of explanation covers all notions of explanation that are used throughout natural science: An account in terms of function and structure, cf. (E1) in Appendix A. He subsequently shows that there are aspects of experience which do not have any of these two properties, so cannot be explained in terms of natural science as usual. The gist of his grounding is that they may be addressed by a "new sort of explanation" [30] (p. 121) which consists of "new fundamental laws (...) specifying how phenomenal (or protophenomenal) properties depend on physical properties" [30] (p. 127).

The question of how scientific explanation is to be defined has occupied many philosophers throughout the 20th century [44]. To find a definition which is general enough to capture the various explanations in science, yet specific enough to exclude scenarios which are clearly not cases of scientific explanation turns out to be a very difficult task. Even basic questions such as whether or not causality is to feature in the definition of explanation (and if yes, which definition of causality), are still largely debated: "There is considerable disagreement among philosophers about whether all explanations in science and in ordinary life are causal and also disagreement about what the distinction (if any) between causal and non-causal explanations consists in." [44].

This sheds some doubt on Chalmers' notion of explanation, and the question arises whether (E1) really covers all, or even the most essential, uses of explanation throughout sciences. This is particularly so with respect to physics, whose notion of explanation seems to be a lot more formal than suggested by the terms 'function' and 'structure' as defined here. E.g., physics does seem to provide notions of explanation which can be applied to general dynamical quantities, whether they describe changes in the behavior of a system³⁹ or changes of a more general sort. (Chalmers might even reluctantly agree to this last observation when claiming that "throughout the higher-level sciences, reductive

³⁸ One may be able to avoid this last conclusion by insisting that the meaning attributed to d by any experiencing subject is dependent on the law E itself and if one furthermore argues that a conclusion about which law E best fits nature can be deduced from the meaning of d , despite d itself being determined independently of the former. At the present stage it seems quite unclear how such a deduction might work, let alone what role an experiment might play in this deduction in the first place.

³⁹ Recall that the term function refers to "any causal role in the production of behavior that a system might perform" [31] (p. 6). One could interpret this as referring to "any change in the behavior of a system" (cf. Appendix B.4). This could, in turn, be taken to mean "any change in the dynamical properties of a system", which would change the meaning of the claim that "physical accounts explain *only* structure and function" [31] (p. 105f.) to the following:

"Any account given in purely physical terms will suffer from the same problem. It will ultimately be given in terms of the structural and dynamical properties of physical processes, and no matter how sophisticated such an account is, it will yield only more structure and dynamics. While this is enough to handle most natural phenomena, the problem of consciousness goes beyond any problem about the explanation of structure and function [sic], so a new sort of explanation is needed." [30] (p. 121)

However, most or even all aspects of consciousness are dynamical in nature, which implies that the set of phenomenal aspects of consciousness (Definition (D1)) is, given this redefinition of the term 'function', either empty or trivial. Put differently, with this redefinition the grounding implies that all or almost all of conscious experience can be addressed by an "account given in purely physical terms". What is left out are only non-dynamical aspects of experience (if there are such aspects at all).

explanation works in just this [(E1)] way” [31] (p. 7), thus, in this quote, avoiding the claim that (E1) also applies to lower-level sciences, such as (presumably) physics.)

This is a problem because the legitimacy of proposing “new fundamental laws” which describe how phenomenal aspects of experience depend on physical properties [30] (p. 127), as compared to a reductive explanation in terms of physical accounts, is granted, in Chalmers’ grounding, by the existence of an explanatory gap between phenomenal aspects and contemporary scientific explanation. If scientific explanation is more powerful than Chalmers assumes, the justification of this explanatory gap breaks down and it becomes questionable whether this explanatory gap actually exists. “[A]n explanatory gap (...) cannot be made more precise than the notion of explanation itself” [43] (p. 358).

Appendix B.4. Causality

Finally, the question arises of what exactly one should take to constitute causality when applying Chalmers’ grounding. This is so because Assumption (A1) as well as Definitions (E1) and (D1) all relate to causality in an essential way (the latter via the definition of the term ‘function’, cf. Appendix A).

This question is widely debated both in physics and in the philosophy of causation [57]. It seems fair to say that consensus is missing on basically all aspects of a definition of causality, including basic questions such as which relata a causal relation is to refer to and how, given a choice of relata, causality is defined. Whereas this multitude of possible notions of causality may not matter much if one is concerned with philosophical investigations based on Chalmers’ grounding (one may just restrict to analyses that apply to every notion of causality), it does matter if one wishes to apply the grounding. In particular, if one wishes to model, let alone to identify, phenomenal aspects of experience, one does need to know what exactly the Definition (D1) amounts to. Since the term ‘function’ used in that definition refers exclusively to causality, the defining property of phenomenal aspects depends on what one takes causality to be.

Connected to questions of how to define causality is the question of the ontological status of causality. Does some definition of causality pertain to “reality” or the universe? (In physicists’ terms: Is causality “fundamental”? Does the universe “obey” one particular definition of causality?) Or is causality rather a tool which can be utilized (by humans, animals, etc. or by information processing systems in general) to describe some parts of reality well to some extent?⁴⁰

Chalmers’ grounding is strongly dependent on which answer one gives to this question. E.g., it determines which sort of “influence” of the phenomenal domain on the physical domain is compatible with the definitions of the grounding, or which type of condition the definition of phenomenal aspects of consciousness constitutes. One may ignore this problem as long as one applies the grounding to theories of the physical domain which incorporate some notion of causality, such as, arguably, abstract neural networks. However, if one wishes to apply the grounding to fundamental physical theories, whose laws do not refer to, or come equipped with, any notion of causality, this question cannot be ignored.

These issues can be avoided completely if one takes the various uses of the term “causality” in Chalmers’ grounding to jointly mean that the physical domain is not changed in any way by phenomenal aspects of consciousness, i.e., that the various uses of causality simply amount to ensuring that the “physical laws already form a closed system” [30] (p. 127). This seems to be the actual

⁴⁰ E.g., [34] holds that “[i]f you wish to include the entire universe in the model, causality disappears because interventions disappear—the manipulator and the manipulated lose their distinction. However, scientists rarely consider the entirety of the universe as an object of investigation. In most cases the scientist carves a piece from the universe and proclaims that piece *in*—namely, the *focus* of investigation. The rest of the universe is then considered *out* or *background* and is summarized by what we call *boundary conditions*. This choice of *ins* and *outs* creates asymmetry in the way we look at things, and it is this asymmetry that permits us to talk about ‘outside intervention’ and hence about causality and cause-effect directionality” [34] (p. 419f.). “What we conclude (...) is that physicists talk, write, and think one way and formulate physics in another” [34] (p. 407).

intention of the author in [30,31], which is why we have fixed this interpretation in the beginning of this appendix.

References

1. Faw, B. Consciousness, modern scientific study of. In *The Oxford Companion to Consciousness*; Bayne, T., Cleeremans, A., Wilken, P., Eds.; Oxford University Press: Oxford, UK, 2014.
2. Seth, A. Models of consciousness. *Scholarpedia* **2007**, *2*, 1328. [[CrossRef](#)]
3. Dehaene, S.; Kerszberg, M.; Changeux, J.P. A neuronal model of a global workspace in effortful cognitive tasks. *Proc. Natl. Acad. Sci. USA* **1998**, *95*, 14529–14534. [[CrossRef](#)] [[PubMed](#)]
4. Baars, B.J. Global workspace theory of consciousness: Toward a cognitive neuroscience of human experience. *Prog. Brain Res.* **2005**, *150*, 45–53. [[PubMed](#)]
5. Dennett, D.C. *Consciousness Explained*; Penguin: London, UK, 1993.
6. Carruthers, P. Higher-Order Theories of Consciousness. In *The Stanford Encyclopedia of Philosophy*; Zalta, E.N., Ed.; Metaphysics Research Lab, Stanford University: Stanford, CA, USA, 2016.
7. Oizumi, M.; Albantakis, L.; Tononi, G. From the phenomenology to the mechanisms of consciousness: Integrated information theory 3.0. *PLoS Comput. Biol.* **2014**, *10*, e1003588. [[CrossRef](#)]
8. Mayner, W.G.P.; Marshall, W.; Albantakis, L.; Findlay, G.; Marchman, R.; Tononi, G. PyPhi: A toolbox for integrated information theory. *PLoS Comput. Biol.* **2018**, *14*, e1006343. [[CrossRef](#)]
9. Haun, A.; Tononi, G. Why does space feel the way it does? Towards a principled account of spatial experience. *Entropy* **2019**, *21*, 1160. [[CrossRef](#)]
10. Kleiner, J.; Tull, S. The mathematical structure of integrated information theory. *arXiv* **2020**, arXiv:2002.07655.
11. Tull, S.; Kleiner, J. Integrated information in process theories. *arXiv* **2020**, arXiv:2002.07654.
12. Metzinger, T.; Wiese, W. *Philosophy and Predictive Processing*; Number 978-3-95857-138-9; MIND Group: Frankfurt am Main, Germany, 2017.
13. Dolkega, K.; Dewhurst, J.E. Fame in the predictive brain: A deflationary approach to explaining consciousness in the prediction error minimization framework. *Synthese* **2020**. [[CrossRef](#)]
14. Chalmers, D.; McQueen, K. Consciousness and the collapse of the wave function. In *Quantum Mechanics and Consciousness*; Oxford University Press: New York, NY, USA, forthcoming.
15. Kent, A. Quanta and qualia. *Found. Phys.* **2018**, *48*, 1021–1037. [[CrossRef](#)]
16. Kent, A. Toy Models of Top Down Causation. *arXiv* **2019**, arXiv:1909.12739.
17. Penrose, R. *Shadows of the Mind*; Oxford University Press: Oxford, UK, 1994; Volume 4.
18. Kremnizer, K.; Ranchin, A. Integrated Information-induced quantum collapse. *Found. Phys.* **2015**, *45*, 889–899. [[CrossRef](#)]
19. Mason, J.W.D. Quasi-conscious multivariate systems. *Complexity* **2016**, *21*, 125–147. [[CrossRef](#)]
20. Hoffman, D.D.; Prakash, C. Objects of consciousness. *Front. Psychol.* **2014**, *5*, 577. [[CrossRef](#)] [[PubMed](#)]
21. Metzinger, T. The problem of consciousness. In *Conscious Experience*; Metzinger, T., Ed.; Imprint Academic: Exeter, UK, 1995; pp. 3–37.
22. Nagel, T. What is it like to be a bat? *Philos. Rev.* **1974**, *83*, 435. [[CrossRef](#)]
23. Lewis, C. *Mind and the World Order*; C. Scribner's Sons: New York, NY, USA, 1929.
24. Metzinger, T. *Grundkurs Philosophie des Geistes Band 1–3*; Mentis: Paderborn, Germany, 2007.
25. Metzinger, T. *Conscious Experience*; Imprint Academic: Exeter, UK, 1995.
26. Tsuchiya, N.; Taguchi, S.; Saigo, H. Using category theory to assess the relationship between consciousness and integrated information theory. *Neurosci. Res.* **2016**, *107*, 1–7. [[CrossRef](#)]
27. Tononi, G. Consciousness as integrated information: A provisional manifesto. *Biol. Bull.* **2008**, *215*, 216–242. [[CrossRef](#)]
28. Atmanspacher, H. On macrostates in complex multi-scale systems. *Entropy* **2016**, *18*, 426. [[CrossRef](#)]
29. Resende, P. Quanta and Qualia. In Proceedings of the Workshop on Combining Viewpoints in Quantum Theory, ICMS, Edinburgh, UK, 19–22 March 2018; Talk provided by the author.

30. Chalmers, D. *The Conscious Mind: In Search of a Fundamental Theory*; Oxford Univ. Press: New York, NY, USA, 1996.
31. Chalmers, D. *The Character of Consciousness*; Philosophy of Mind, Oxford University Press: New York, NY, USA; Oxford, UK, 2010.:oso/9780195311105.001.0001. [CrossRef]
32. Schlick, M., Form and content: An introduction to philosophical thinking. In *Gesammelte Aufsätze 1926–1936*; Gerold: Vienna, Austria, 1938.
33. Kleiner, J.; Hoel, E. Falsification and consciousness. *arXiv* **2020**, arXiv:2004.03541.
34. Pearl, J. *Causality: Models, Reasoning, and Inference*, 9th ed.; Univ. Press: Cambridge, UK, 2009.
35. Churchland, P.M. Eliminative materialism and propositional attitudes. *J. Philos.* **1981**, *78*, 67–90.
36. Kuehni, R. Color spaces. *Scholarpedia* **2010**, *5*, 9606. [CrossRef]
37. Resnikoff, H.L. Differential geometry and color perception. *J. Math. Biol.* **1974**, *1*, 97–131. [CrossRef]
38. Provenzi, E. Principal Fiber Bundles and Geometry of Color Spaces. In Proceedings of the Second International Conference on Advances in Signal, Image and Video Processing, Barcelona, Spain, 21–25 May 2017.
39. Sharma, G.; Wu, W.; Dalal, E.N. The CIEDE2000 color-difference formula: Implementation notes, supplementary test data, and mathematical observations. *COLOR Res. Appl.* **2004**, *30*, 21–30. [CrossRef]
40. Prentner, R. Consciousness and topologically structured phenomenal spaces. *Conscious. Cogn.* **2019**, *70*, 25–38. [CrossRef]
41. Pervin, W.J. *Foundations of General Topology*; Academic Press: Cambridge, MA, USA, 1964.
42. nLab authors. Pretopological space. Revision 10. Available online: <https://ncatlab.org/nlab/show/pretopological+space> (accessed on 1 May 2020).
43. Levine, J. Materialism and qualia: The explanatory gap. *Pac. Philos. Q.* **1983**, *64*, 354–361. [CrossRef]
44. Woodward, J. Scientific Explanation. In *The Stanford Encyclopedia of Philosophy*; Zalta, E.N., Ed.; Metaphysics Research Lab, Stanford University: Stanford, CA, USA, 2017.
45. Winther, R.G. The structure of scientific theories. In *The Stanford Encyclopedia of Philosophy*; Zalta, E.N., Ed.; Metaphysics Research Lab, Stanford University: Stanford, CA, USA, 2016.
46. Craver, C.F. Structures of Scientific Theories. In *The Blackwell Guide to the Philosophy of Science*; Machamer, P., Silberstein, M., Eds.; Blackwell Publishers: Hoboken, NJ, USA, 2002; Chapter 4, pp. 55–79. [CrossRef]
47. Giulini, D. Concepts of symmetry in the work of wolfgang pauli. In *Recasting Reality: Wolfgang Pauli's Philosophical Ideas and Contemporary Science*; Atmanspacher, H., Primas, H., Eds.; Springer: Berlin/Heidelberg, Germany, 2009; pp. 33–82. 3. [CrossRef]
48. Chalmers, D.J. Absent qualia, fading qualia, dancing qualia. In *Conscious Experience*; Metzinger, T., Ed.; Imprint Academic: Exeter, UK, 1995.
49. Bishop, R.C. The Hidden Premise in the Causal Argument for Physicalism. *Analysis* **2005**, *66*, 44–52. [CrossRef]
50. Dehaene, S.; Changeux, J.P.; Naccache, L. The global neuronal workspace model of conscious access: From neuronal architectures to clinical applications. In *Characterizing Consciousness: From Cognition to the Clinic?* Dehaene, S., Christen, Y., Eds.; Springer: Berlin/Heidelberg, Germany, 2011; pp. 55–84. 4. [CrossRef]
51. Grindrod, P. On human consciousness: A mathematical perspective. *Netw. Neurosci.* **2018**, *2*, 23–40. [CrossRef]
52. Wallace, R. *Consciousness: A Mathematical Treatment of the Global Neuronal Workspace Model*; Springer: Berlin/Heidelberg, Germany, 2005.
53. Dehaene, S.; Naccache, L. Towards a cognitive neuroscience of consciousness: Basic evidence and a workspace framework. *Cognition* **2001**, *79*. [CrossRef]
54. Elitzur, A.C. Consciousness makes a difference: A reluctant dualist's confession. In *Irreducibly Conscious: Selected Papers on Consciousness*; Batthyany, A., Elitzur, A.C., Eds.; Winter: Heidelberg, Germany, 2009.
55. Floridi, L. Semantic conceptions of information. In *The Stanford Encyclopedia of Philosophy*; Zalta, E.N., Ed.; Metaphysics Research Lab, Stanford University: Stanford, CA, USA, 2017.

56. Shannon, C.E. A Mathematical Theory of Communication. *Bell Syst. Tech. J.* **1948**, *27*, 379–423. [[CrossRef](#)]
57. Schaffer, J. The metaphysics of causation. In *The Stanford Encyclopedia of Philosophy*; Zalta, E.N., Ed.; Metaphysics Research Lab, Stanford University: Stanford, CA, USA, 2016.



© 2020 by the author. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<http://creativecommons.org/licenses/by/4.0/>).