The epistemological significance of exploratory experimentation: Why practices matter, philosophically

Pierre-Hugues Beauchemin* Kent W. Staley†

We employ a pragmatic model of inquiry to distinguish the epistemological character of exploratory experimentation. Exploratory experimentation is not constituted by any intrinsic characteristics of an episode of experimentation but depends on the context and aims of the experiment and the ways in which these shape decisions about how the experimental inquiry is to be conducted: its tasks, resources, and aims, as well as the critical assessment of all of these. To demonstrate the usefulness of our approach, we apply it to the contrast between two kinds of searches for new physics at the Large Hadron Collider. Some searches are exploratory while others target specific Beyond Standard Model hypotheses, but this contrast can be understood only by considering the relations between these searches, their aims, and the way that these aims shape their respective experimental parameters and procedures. Our approach provides a model for establishing the epistemological significance of details of experimental practice.

*Physics Department, Tufts University
†Department of Philosophy, Saint Louis University

1 Introduction

It has become a commonplace of History and Philosophy of Science (HPS) that careful study of experimental practice opens a valuable perspective onto the scientific process that can correct for the artificiality and misleading abstractness of approaches concerned only with confirmation relations regarding theory. A persistent challenge of such studies, however, is to articulate what is gained philosophically from a detailed description of particular experimental practices. How does the practice contribute epistemically to the production of scientific knowledge? Does it figure into the justification, validation, or support of scientific claims, and if so, how does it do so?

A context in which this problem arises is the discussion of *exploratory experimentation* (EE), a term introduced independently in print by Friedrich Steinle (1997), Richard Burian (1997), and Rose-Mary Sargent (1995b; 1995a) and subsequently taken up by numerous others (Elliott, 2007; Franklin, 2005; Karaca, 2017; Waters, 2004). The motivation for such discussions has been to draw a contrast with "theory-driven" experiments, such as those seeking to test hypotheses articulated prior to experimentation.

According to a critique by Jutta Schickore, the discussion of exploratory experimentation has been successful from a historiographical point of view in providing HPS with a framework for analyzing historically significant experimental practices and types of instrumentation. The discussion of exploratory experimentation has not, however, provided an account of knowledge production and validation that offers a clear, compelling contrast even to such theory-dominated accounts as falsificationism. The discussion of EE has principally revealed the need for better "conceptual tools for the study of experimental practice" while remaining "too unspecific to fill the desideratum whose existence it demonstrated" (Schickore, 2016, 25). Steinle's work allows us to identify distinctive features of a *style* of experimentation employed by, say, Ampère in his experiments with the astatic magnetic needle, just as Burian does for Brachet's investigations of synthesis and degradation of nucleic acids. We can see what these scientists did in the course of their experimentation leading to their conclusions, but what, in comparison to accounts focused on falsification or confirmation of theoretical claims, is distinctive about how those practices contribute to the conclusions *constituting knowledge*, or about the *kind of knowledge* these types of experiments produce?

Our response to this challenge employs a pragmatist account of inquiry (cf. (Dewey, 1938)). We treat knowledge as the product of a successfully executed process of inquiry carried out within a community of inquirers. Inquiry is carried out by performing *tasks*, using a variety of resources (instruments, theories, computations, simulations, etc.) directed at both proximate and distal aims, ultimately directed toward the production of judgements that are sufficiently stable and informative to serve in turn as resources for future inquiry. Crucial to our account is that inquiry is conducted in two entangled modes of *use* and *criticism*. Inquirers use resources in the execution of aim-directed tasks. They also engage in criticism to assess adequacy – of resources for their intended use, of the performance of tasks using those resources, and of attainment of the aim for which the task was performed. By conducting an inquiry in both the use and the critical modes, directed at and articulated with the objectives of the inquiry, scientists are able to warrant the claims they make on the basis of the results of their experiment.

We use this framework to examine a class of experimental inquiries in contemporary High Energy Physics (HEP): the use of Signature Based Model Independent (SBMI) searches for new physics. We argue that these searches constitute contemporary, big-data examples of exploratory experimentation that cannot be assimilated to any existing conception of EE. We then offer a new perspective on EE that allows us to elucidate what is epistemically distinctive about such experiments. This enables us to offer a positive account of the way in which SBMI searches contrast with other

HEP experiments that are paradigmatically theory-driven. We situate the exploratory character of these experiments in the way that investigators adopt and adapt various tasks and resources in the service of aims that promote, in the context of a broader experimental program, an openness to unanticipated phenomena. The components of an SBMI inquiry are chosen and executed to prioritize increasing the potential to discover something unanticipated over maximizing the prospects for discovering some specific phenomenon predicted by a particular theoretical model.

Our approach generalizes beyond HEP: Exploration is not a separate kind of experiment carried out with its own distinct procedures, but a *context* that calls for a distinct evaluation of epistemic risks, and an adjustment of tasks, resources, and aims reflecting that evaluation in a manner sustainable under critical examination. This achievement of epistemic aims through the interplay of use and criticism is not adequately captured by accounts oriented toward testing or confirmation of theoretical claims. We do not propose to reject analyses offered by Steinle, Burian, and other contributors to the EE literature, but to recast them in a new, more explicitly epistemological, perspective.

By meeting Schickore's challenge, our approach demonstrates its own resources that can meet the more general challenge of demonstrating the philosophical significance of attending to experimental practices. We explicate the epistemic significance of exploration by showing how considerations of the success conditions for inquiry, such as epistemic risk, enter the conception of the objectives of an inquiry, the formulation of a strategy to meet those objectives, and the selection and execution of tasks to fulfill that strategy. In this way, a pragmatic analysis reveals how careful study of practices is essential for understanding the variety of ways in which inquiry yields knowledge.

2 Exploratory experimentation and experimental exploration

The term 'exploration' suggests, not aimlessness, but openness. The explorer of a city lets themself be drawn into the eccentric art gallery or disreputable tavern; they do not follow a fixed itinerary of sights to be seen. In science, exploration similarly connotes an approach to inquiry that lends itself to surprise or perhaps revelation.

Discussion in HPS of exploration in the context of experimentation emerged in the mid-1990s, first in Rose-Mary Sargent's book on the philosophy of experiment of Robert Boyle (Sargent, 1995a). But it is Boyle himself who articulated a notion of exploratory experiments, in contrast to experiments performed to test a theory (ibid., 137). The contemporary discussion has been shaped to a large extent by the independent efforts of Friedrich Steinle (1997) and Richard Burian (1997), who both used the term "exploratory experimentation" (henceforth, EE) while studying quite different scientific fields (early nineteenth-century experimentation in electromagnetism for Steinle, early twentieth-century molecular biology for Burian) and emphasizing different, though overlapping, features of the experiments they discussed. We will not undertake a comprehensive survey of the literature here (see (Mättig, 2022) for a thorough discussion), but highlight features of EE that have particular salience to our analysis.

Steinle's approach to EE eschewed conceptual analysis in favor of listing the most important "typical guidelines" for such experimental activity. Included in Steinle's guidelines for EE are the systematic variation of "a large number of different experimental parameters," "looking for stable empirical rules," and "finding appropriate representations by means of which those rules can be formulated" (Steinle, 1997, S70). Burian's study of Jean Brachet's exploratory experiments on the "localization of thymonucleic acid" (DNA) emphasizes the way Brachet used a combination of new and older techniques to "uncover the unknown places in which, and the unknown sequence in which, nucleic acids ... and other substances are present" in the cell, as well as the effects of interventions (Burian, 1997, 41). Brachet and colleagues were seeking to "find correlations" between nucleic acids and "biochemical, physiological, and morphological changes" in order to understand how these substances contribute to the "entire ontogenetic process." To achieve that aim, Brachet and colleagues employed numerous techniques of cross-checking of the "biochemical constitution and spatio-temporal localization" of nucleic acids. Burian's discussion of exploratory experimentation emphasizes this point: "At the heart of the matter is the need for a battery of technically adequate means for cross-checking different techniques, one against another, for reidentifying a 'thing' or process" (ibid., 43-44). By deploying these means, Brachet and other researchers were able to localize entities "in space and time without depending on the theoretical or functional identities assigned to them, however provisionally," the enabling the comparison of "different, seemingly incompatible, epistemic objects[1] via ... cross-checking" (ibid., 45).

A consistent theme of the early EE literature is to distinguish experiments performed in pursuit of exploration from those performed in pursuit of the testing of prior theory, or what Steinle refers to as "research determined by theory in the strict sense" (Steinle, 2016, 314). That leaves unaddressed, however, just what roles for theory are compatible with EE. As noted by Mättig (2022, 6), there is a "general consensus" that EE is not theory free, and several authors have described positive roles for theory in the conduct of EE. Laura Franklin-Hall's (2005) discusses how the use of "wide" or "high-throughput instrumentation" in genomics draws upon background theory to enable the efficient exploration of an experimental system. Kenneth Waters shows how classical genetics constituted a "broad system of scientific knowledge ... organized for exploratory research" (Waters, 2004, 786) that in turn played a

[1] Burian here alludes to the context of his analysis in relation to Rheinberger's concepts of experimental systems and epistemic objects (Rheinberger, 1997).

theory-like role by enabling the classification and organization of experimental knowledge. Kevin Elliott, drawing upon these and other early contributions to the literature on EE, attempted a taxonomy of EE along three dimensions, including "role of theory in the [experimental] activity" (Elliott, 2007, 324). Those roles include "providing background information," "serving as a starting point or foil," and "being constituted by exploratory projects or strategies" (ibid.). As Karaca points out, however, these allowances for theoretical dependence in EE do not include dependence on theories about the *target* or "phenomena under investigation" (Karaca, 2017, 334, 338), and Mättig insists that "the absence of a target theory is a defining moment for EE and separates it from experiments geared to theory testing" (Mättig, 2022, 7).

There is, however, an ambiguity around the notion of a "target," and hence around the question of a "target theory." Koray Karaca describes a program of experimentation that is simultaneously exploratory and relies upon theoretical input regarding the target phenomenon (Karaca, 2017).

Historical motivations for introducing EE first arose in domains seemingly remote from the context of such theory-intensive experimental programs as HEP. Nonetheless, Koray Karaca makes a strong case for regarding some work in experimental HEP as involving exploration in a manner not captured by earlier accounts of EE. Karaca introduces the idea of an *exploratory procedure* as a procedure that "serves to extend the range of possible outcomes of an experiment and thereby the scope of the experimental inquiry to the investigation of a wider range of phenomena" (Karaca, 2017, 340).[2] Experiments that incorporate such procedures may also involve non-exploratory procedures, and procedures of both kinds may be theory-laden, as illustrated by Karaca's example of a data selection employed by the ATLAS collaboration at the Large Hadron

[2]Karaca's use of the term 'experiment' seems to reflect a common usage in HEP that applies the term to the entire undertaking of a group like ATLAS (e.g., "the ATLAS experiment"). His discussion of the ATLAS trigger as an exploratory procedure that expands the potential for discovery by the ATLAS group through its various inquiries relies on this usage. Although well-established within HEP, this usage is not common in other disciplines. Attending to *inquiry* focuses on specific efforts to achieve experimental knowledge within such larger experimental research programs.

Collider (LHC). Although theory-laden, including theory about the target phenomena, Karaca argues that the procedure is not *theory-directed*, but instead is driven by a *strategy* of data selection that extends the range of possible experimental outcomes. Yet the data selected may be used to test specific theoretical hypotheses drawn from specific BSM models. Consequently, the data selection procedure employed by ATLAS, though not an instance of EE (exploratory experimentation) in the sense of Steinle and Burian, constitutes an example of "experimental exploration."

We seek to do justice to the motivations that gave rise to Steinle's and Burian's work on EE, while drawing inspiration from Karaca's shift of focus toward specific procedures and their implications for the aims of a given experiment. Yet our approach differs fundamentally insofar as we situate our understanding of exploration within a more systematic pragmatist epistemological approach to experimentation.

3 Sketch of a pragmatic model of experimental inquiry

Our account of the exploratory character of experimentation uses concepts from a pragmatic model of experimental inquiry in physics we are currently developing. Here we sketch its broad philosophical framework, explicate its aims, and catalogue the most general concepts we employ in its construction, which suffice for the present argument. Our sketch highlights a distinction between a *use mode* and *critical mode* of epistemic activity, crucial for appreciating the philosophical significance of exploration as a mode of experimentation.

Our model is *pragmatic* insofar as it treats knowledge gained in experimental physics as the successful outcome of a certain kind of *doing*. The most salient precedent for our approach is Dewey's account of inquiry, in which the inquirer proceeds from an initial state, described as an "indeterminate situation,"[3] to a final judgment through a process of inquiry involving five non-sequential iterative phases: observation, institution of a problem, suggestion, reasoning, experimentation, resolution of indeterminacy ((Dewey, 1938, 101–119); (Brown, 2012, 270–276)). Although we find Dewey's model to be a fruitful inspiration, our model is based on a distinct set of concepts and aims at a more fine-grained treatment of inquiry applicable to the context of experimental physics.

We regard knowledge as the product of a successfully executed process of inquiry (understood by Dewey in terms of arriving at a "warrantably assertible" judgment) by a community of inquirers. As do both Peirce and Dewey, we take the pragmatic significance of such judgments to rest on their forward-looking stability and suitability for use as premises (resources) in future episodes of inquiry Peirce (1878); (Dewey, 1938, 138–140). They should be sufficiently informative to enable epistemic ends not previously achievable, or to open up means of achieving such ends not previously available. Modeling the process of inquiry that results in knowledge thus understood must proceed by consideration of how the tasks of inquiry contribute to the production of judgments with these features of stability and suitability as resources for future inquiry.

A *task*, in our usage, is a type of action undertaken to serve some aims, which may be epistemic or non-epistemic, that contributes in a specifiable way to the production of knowledge. To understand the tasks of experimentation is to understand what scientists are doing when they are engaged in experimentation (compare Chang's *epistemic activity* 2012, 15.)

The definition of a particular task need not be unique — the actions of a scientist at a particular moment may map onto more than one task — but an adequate characterization of an experimental task will facilitate sense making and analysis of scientific

[3] Dewey's term refers to a relationship between agent and environment giving rise to a "discoordination or disequilibrium" in the conduct of the agent's practices and a corresponding feeling of "doubtfulness or hesitancy" (Brown, 2012, 276).

inquiry. Once defined, one can ask how tasks relate to one another as an inquiry is conducted in order to reach its objectives. Tasks as we conceive them may immediately serve non-epistemic aims, but contribute to an overall experimental activity aimed at producing knowledge. We accommodate thus the entanglement of epistemic and non-epistemic pursuits within experimentation. Tasks may be described and distinguished at higher or lower grades of resolution or generality. 'Establishing a statistical discrepancy from a prediction' is a coarse-grained task included in many experimental undertakings. 'Determining the contribution of the choice of generator level Monte Carlo to the systematic uncertainty of the null hypothesis prediction' is a more fine-grained task required of many experimental efforts in particle physics and related fields, often necessary to complete the coarse-grained task just mentioned. Ultimately, one hopes to provide a comprehensive narrative of what experimentalists do at any desired level of granularity.

To carry out a task requires using *resources*, which may be physical (a calorimeter), computational (a computer simulation), representational (a plot), inferential (a statistical testing method), or epistemological (a theory), just to name a few possibilities. In this picture, data, collected by the instruments central to the experiment or by ancillary apparatuses, are also resources, as are data models Suppes (1962); Leonelli (2019); Bokulich (2020); Antoniou (2021). Many tasks produce new resources that may in turn be used to carry out a subsequent task, as when running a computer simulation produces a set of simulated data that become a resource for estimating a source of background.

Tasks may be performed for both proximate and distal *aims*. These aims may be formulated prior to, and provide the original motivation for, the conduct of inquiry, or they may arise while performing the experiment. Crucial to determining the aims of an experimental inquiry are those phenomena to which the experiment is meant to be sensitive: for what physics phenomena and evidence thereof is the experiment being performed? Although decisions about sensitivity to target phenomena may be revised if necessary, they tend to exert significant influence over decisions regarding experimental design, and hence over the specific tasks and resources that performing the experiment demands. The centrality of sensitivity to experimentation entails a corresponding importance to estimating uncertainty, which bears directly on sensitivity Beauchemin (2017); Staley (2020). This is true whether the experiment is performed to test a specific theory or is meant to be capable of registering a range of phenomena in a more exploratory manner.

The discussion of tasks thus far has been in reference to the *use mode*. A central feature of our pragmatic model is to contrast the mode in which one uses resources to execute tasks with the *critical mode*, in which one assesses the adequacy of those resources for achievement of the aims of their associated tasks, as well as the degree of success in achieving those aims. Such critical mode assessments may also result in changing the resource or its use. At the same time, our model emphasizes the ways in which these two modes are entangled: every critical mode assessment is in turn carried out via tasks that involve their own use of resources. The entanglement of "use" and "critical" modes of inquiry arises from the fact that the inquirer may (in the use mode) make an assertion, but the warranting of that assertion requires actions conducted in the critical mode; conversely, inquirers must always perform tasks that produce assertions before completing any critical mode conduct. The point of inquiry is not the production of warrant or of assertion, but achieving warranted assertability.

To summarize: we understand experimentation to exemplify the pursuit of inquiry through the execution of tasks, that use resources and are subject to criticism, in the pursuit of proximate and distal aims, where the latter includes the achievement of warranted assertability, i.e., knowledge. This understanding can already provide some insight into important features of the EE literature discussed in the previous section.

3.1 Experimentation as inquiry

Although it is commonplace to cite Steinle and Burian as having independently and simultaneously initiated the discussion of "exploratory experimentation" in their respective 1997 publications, Schickore emphasizes some important differences between the two discussions. Steinle characterizes exploratory experimentation in terms of methodological strategies, including parameter variation and purifying of an experimental arrangement, that also produces a "*specific* form of knowledge, namely new concepts." Schickore contrasts this with Burian's approach, which she understands as depicting exploratory experimentation as "the triangulation on a phenomenon with the help of a rich and diverse set of experimental approaches," reflecting a methodological strategy of seeking "independent confirmations in situations where the available techniques and instruments do not appear completely secure" (Schickore, 2016, 22). In this way, Schickore draws our attention to important contrasts between Steinle's and Burian's concerns that can be lost sight of if we simply regarding them as both having introduced the same idea using different examples. Schickore concludes, "the two notions of exploratory experimentation they present are really quite different" (ibid., 21).[4]

Our framework allows us to see a systematic connection between the methodological aspects emphasized in Steinle's and Burian's discussions. Steinle's references to experiments that employ apparatuses and arrangements that allow for many variations in arrangement or outcome consistently show us how the experimenters *use* the material and conceptual resources at hand to explore a phenomenon.[5] Burian's description of Brachet's experimentation also, of course, describes him as using the resources needed for his experimentation, but Burian's emphasis returns continually

[4] Steinle himself notes that Burian uses the term "in a slightly broader sense" that he considers to include his own, and that Sargent uses the term "in the rather different sense of exploring new methods of chemical manipuation or new experimental techniques" (Steinle, 1997, S71). Schickore's claim that Steinle's and Burian's notions of EE are entirely distinct seems to rest on treating triangulation as exhaustive of Burian's conception. We agree with Steinle that Burian's notion of EE is broader than this.

[5] "Faraday systematically varied a lot of parameters of the arrangement such as the direction of motion (relative to the magnetic dip), the mode of motion (e.g., various parts of the circuit or the circuit in its entirety), the form of the circuit, and so on" (Steinle, 1997, S68).

to what we regard as *critical mode* experimental practices of cross-checking, an aspect made especially important by the way in which Brachet and colleagues relied on the innovation of new experimental techniques alongside more established ones.[6] What sets apart Steinle and Burian, from this perspective, may not be so much that they are attempting to characterize two different ways of going about exploration, but rather that they are focusing on complementary modes of experimental inquiry conducted in an exploratory manner.

Once we recognize the importance of the use/criticism distinction for appreciating the relationship between Steinle's and Burian's approaches to EE, we are also equipped to respond to Schickore's critique of the EE literature in a way that reveals the philosophical significance of exploration in the epistemology of experimental science. In particular, we can recognize a contrast between how warranting is accomplished when experimentation is pursued in an exploratory manner as opposed to the theory-driven alternative. That contrast lies in the interplay between the use mode and the critical mode.

3.2 Is EE relevant to the philosophy of science? Meeting the challenge

Schickore notes how Popper's falsificationism provides a foil for Steinle's discussion, and she extends this use of falsificationism to pose a challenge regarding the philosophical relevance of EE as articulated by Steinle. She argues that a Popperian falsificationist would insist that the testing of theories is the only thing accomplished by experiments that "is relevant for the justification of these theories" and the kind of exploratory work discussed by Steinle simply belongs to the process of arriving at concepts to be used in theoretical hypotheses to be tested. The latter process is relevant to the epistemology of science only in a very weak sense. The challenge, for Schickore,

[6] "[T]he key to a number of significant advances was the Rouge-Cloître group's great emphasis on the use of a variety of histochemical and other means of localizing the molecules and on sufficient cross-checking by multiple techniques to be reasonably confident that artefacts were eliminated" (Burian, 1997, 43).

is to show that EE leads to knowledge in a different manner or of a different kind from the falsificationist testing of hypotheses (Schickore, 2016, 22–23).[7]

Our approach is different. To the extent that the philosophical relevance of EE has been elusive, we attribute this to an underappreciation of something shared by experiments performed in exploratory and theory-driven ways: the interplay between tasks executed in the use mode and the critical mode as a means of achieving the epistemic goals of an inquiry. Such goals may center on the testing of a "target" theoretical hypothesis of prior interest, or on the development of new concepts for the characterization of phenomena in a domain under exploration, or on a particular way of balancing different epistemic risks attending the inquiry undertaken. Paying attention to how experimenters engage in carefully adjusted use- and critical-mode activity so as to satisfy the criteria of success for their experimental inquiries, we gain a perspective on experimental knowledge production that is absent from approaches based entirely on falsification or confirmation as central concepts in the epistemology of science. Once this perspective is adopted, the strong epistemological value of exploratory experimentation, and not just its historiographic interest, will come clearly into view.

To vindicate this claim, we need to locate how specifically exploration plays out in experimental inquiry, and for this purpose we examine a class of experiments in HEP (SBMI searches for new physics) that have a clear exploratory intent, yet elude easy characterization using ideas in the extant EE literature. Our discussion will show how exploration in the context of SBMI searches involves an interplay between tasks carried out in the use and critical modes, with an aim towards achieving a particular prioritization of competing epistemic risks. Our analysis will show how these same tasks can be found in both theory-driven and exploratory modes of experimentation. What distinguishes the theory-driven from the exploratory mode is the *way* in which

[7]Schickore in one place expresses the challenge in terms of the need to show that EE "has a justificatory function" (Schickore, 2016, 21). Popper himself states outright that "[s]cientific theories can never be 'justified', or verified" (Popper, 1959, 315). We assume that Schickore's use of the term 'justification' reflects her deployment of falsificationism merely as a convenient analytic device.

epistemic risks guide decisions about the relationship between tasks performed in the use and critical modes.

3.3 An ambiguity resolved

First, however, we need to address an ambiguity concerning the term "experiment" that is especially salient to the HEP context. It turns out that addressing this ambiguity involves harvesting yet another fruit made available by the Deweyan roots of our pragmatic approach.

According to Karaca's account, an *exploratory procedure* extends "the range of possible outcomes of an experiment." By doing this, the procedure extends "the scope of the experimental inquiry to the investigation of a wider range of phenomena" (Karaca, 2017, 340). What this account does not tell us is what counts as an experiment. Moreover, if a given procedure can be thought of as contributing to more than one experiment, how should we think about its exploratory character in relation to different experiments of which it is a part? The case to which Karaca applies this idea is the specification of data acquisition procedures for ATLAS, and hence is closely related to our own study. ATLAS is often referred to, in the jargon of HEP, as a single experiment (i.e., "ATLAS Experiment," the title that appears on the homepage of the ATLAS website, https://atlas.cern). In one sense, this is reasonable, given that ATLAS employs a single (though tremendously large and complicated) experimental apparatus: the LHC collider in combination with the ATLAS detector (plus, arguably, all of the associated computer facilities that extend the processing of ATLAS detector signal outputs into usable data). Following Karaca, we can thus think of the exploratory character of data acquisition procedures of the ATLAS experiment in this global sense: a data acquisition procedure at ATLAS is exploratory to the extent that it extends the scope of possible outcomes of the entire ATLAS experimental enterprise.

This interpretation, however, makes it rather easy for something to count as exploratory for ATLAS. The following would all seem to qualify: increasing the beam energy at the LHC, pursuing a search optimized for a signature associated with a newly introduced BSM model, finding a way to slow the process of deterioration of performance for a detector component, and postponing the planned termination of a data-taking run. It is not obvious that any of these should *not* count as exploratory procedures, but the list does tend to reinforce the concern, raised by Schickore, that the HPS literature on EE (here in the sense of experimental exploration) has not articulated clearly the philosophical significance of its subject matter. There are lots of decisions that will affect the range of possible outcomes of an experiment. Having a single name for all of these does not on its own yield any obvious benefit to our understanding of how scientific experimentation yields knowledge.

Perhaps, though, the problem just raised is an artifact of HEP's peculiar use of language. The "ATLAS experiment" publishes papers at an astonishing rate on a wide range of physics phenomena (the number of papers submitted for publication based on LHC collision data reached 1000 on June 18, 2021). The diversity of physics questions addressed in these papers (Is there a Higgs boson? What is the mass of the W boson? Is the KW-L prescription for parton shower to matrix element matching adequate to model $Z+$ jets data?) suggests that we should regard these papers as reports on the outcomes of many different experiments. This *might* cut down on the variety of things that would count as exploratory, or at least allow for closer connections between the procedures that count as exploratory and specific experimental outcomes. This would, however, be at odds with Karaca's own usage and his argument, which aims to show that the data selection strategy of the ATLAS experiment as a whole constitutes an exploratory procedure insofar as it "enables the widest possible range of interesting events that have the potential to serve the entire range of objectives to be pursued in *the ATLAS experiment*" (Karaca, 2017, 345; emphasis added). It also leaves us with an unsolved problem regarding how to individuate experiments for the purpose of deciding whether a given procedure ought to count as exploratory.

A remedy lies close at hand in the model of inquiry proposed by Dewey that we take as our inspiration. The HEP usage of the term "experiment" is not well suited to serve an analytic purpose in the epistemology of experiment. Yet the shared context of all the many ATLAS physics results (one instrument, a single data taking enterprise, etc.) presents challenges to any effort to break "the ATLAS experiment" into some number of smaller experiments.

We propose a shift from experiment as a unit of analysis to *inquiry*. Identifying an inquiry, in Dewey's sense, is both less difficult and has more evident epistemological salience. ATLAS physicists begin with a wide range of open physics questions that they have some reason to think can be answered using the resources made available in the ATLAS experimental enterprise. (Dewey would call these "unresolved situations" in order to emphasize the way in which these questions, because they are open, correspond to obstacles that prevent physicists from completing some of their projects.) An inquiry consists in an organized and systematic effort to answer one of these questions (to "resolve a situation," in Dewey's own favored language). Conducting such an inquiry might involve its own sub-inquiries, which will have their own structure. (Answering a physics question like "do our data contain a signal for a heavy W' boson decaying to an electron plus neutrino?" will require first answering a different question: "what data selection criteria for separating signal from background will give us the greatest chance of finding a signal of W' if it is present?") This poses no problem for us, insofar as our approach locates the exploratory character of an experimental undertaking precisely in the relationship between proximate aims of sub-inquiries with broader aims of a larger inquiry to which they contribute.

Before we can show the value of this pragmatic model as a means of illuminating the exploratory character of SBMI searches in HEP, we must explain how, and why, SBMI searches are specified and performed, and how HEP physicists use their results.

4 Signature-Based Model-Independent Searches

Searches for New Physics constitute a large fraction of the scientific program of various HEP experiments, including those at the Large Hadron Collider (LHC). These searches generically consist in looking for experimental evidence for a statistically significant discrepancy between the observed data and the predictions obtained with the Standard Model of particle physics (SM). Very frequently, this will take the form of a counting experiment where the number of events of a given "signature" collected by the detector is compared to expectations from the SM. Theoretical extensions of the SM are called Beyond the Standard Model (BSM) theories. Before elaborating on our study case, it is useful to better define some of these terms.

An event operationally consists in everything recorded by a detector like ATLAS (ATLAS, 1999) within a 25 ns time-window centered on each LHC proton bunch crossing at the interaction point of the detector. ATLAS physicists distinguish particle and event signatures as distinct levels.

A *particle signature* consists in the reconstruction, from the digital output of the instruments composing a multipurpose detector like ATLAS, of the passage of a fundamental particle (electrons, muons, hadrons, photons, neutrinos) through the material of the detector, as can be seen on the left panel of Fig. 1. A particle signature consists in not only a type of particle, specified by the components of the detector yielding a signal, but also in its kinematic state, i.e. its energy and momentum obtained from the magnitude and shape of the detector signal. Physicists use the term "signature" because there is an unresolvable non-negligible probability that the reconstructed particle is a "fake" or is miscategorized (e.g. a 63 GeV pion can be confused with a 55

GeV electron). These probabilities, and the kinematic properties of the particle reconstructed, are fundamentally associated with an experimental uncertainty that itself depends on the values estimated for these quantities. The uncertainty associated with statements like "we detected particle X with energy E" must be taken into account in any subsequent reasoning relying on such statements about particle signatures.

An *event signature* consists in a set of reconstructed particles with their kinematics and their correlations (e.g. the invariant mass of two of the particles in the final state) attributed to a physics process. An event signature is therefore the experimental outcome of a phenomenon of interest, but also of all the other physics or instrumental processes that could give the same outcome, called background. For example, the right panel of Fig. 1 features the signature of a low-momentum Z' event (predicted by some BSM models), which, after careful analysis, has been evaluated as more likely coming from a SM background process. The properties of the reconstructed particle signatures and their associated uncertainties are propagated to the event signature and their interpretation in terms of a targeted physics process. A search for new physics therefore consists in selecting all events of a given signature of interest, estimating how many SM background events are expected in the selected dataset, including an uncertainty estimate on the SM predictions, and comparing the observed number of events with this expectation. Evidential claims are obtained from this comparison using statistical methods.

The standard strategy for optimizing searches for new physics consists in choosing a set of BSM models and, with the help of Monte Carlo simulation, scanning over the free parameters of the models to determine, for each parameter-space point, which set of event selections would maximize a given figure-of-merit such as the statistical significance of an excess over SM expectations. A data analysis would then be performed in each of the phase space regions selected with the help of the models to quantify

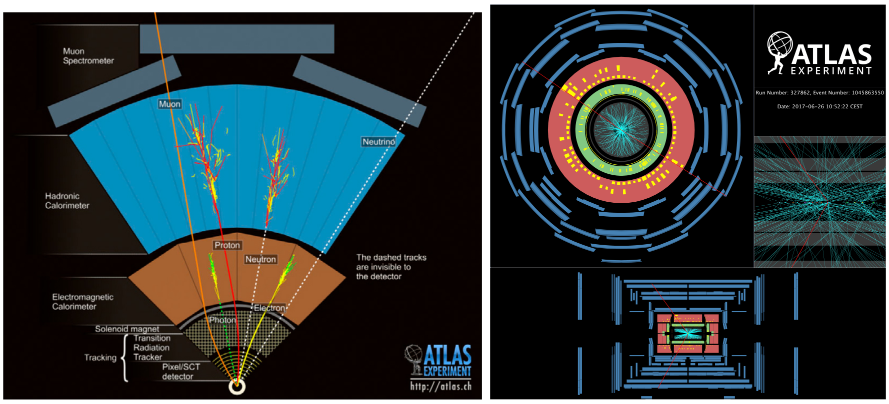

Fig. 1 **Left:** Schematic of how different particles interact with the various ATLAS detector components and the signals they yield. What is identified as a given particle, for example an electron, could be due to another particle mistaken for an electron (a "fake" electron). **Right:** Event display of a dimuon candidate event with the highest invariant mass ($m_{\mu\mu} = 2.75$ TeV) observed in the 2015-18 data taking period. The analysis searched for new heavy particles such as Z' bosons decaying into dilepton (e.g., two muons) final states. Statistical analyses led ATLAS to conclude that the observed signal in the entire dataset is consistent with SM expectations (ATLAS, 2019).

the level of agreement between the observed data and the corresponding SM predictions. This strategy uses theories about BSM target phenomena explicitly to optimize the search for new physics and to account for the impact of BSM model theoretical uncertainties on the measurement results (e.g. acceptance and efficiency corrections, which depend on the kinematics of the assumed underlying target processes). We refer to strategies developed through optimization to BSM models as "BSM-oriented searches." They allow for optimal sensitivity to the theory model of interest, given the uncertainty estimated on the measurement. With such an approach, experimenters have clear expectations about outcomes, about what a discovery would look like if the theory is correct, and about how this theory, if adequate, would impact the uncertainty on the experimental results. Inquirers have a clear idea of what they are looking for and can control the epistemic risks of an erroneous discovery claim.

There is an incredibly rich and broad spectrum of BSM possibilities and theorists have rarely developed models considered complete and realistic. Most BSM models are simplified, approximate, or minimal "toy models," developed to capture gross features of a possible new physics scenario. Failure to observe a significant deviation with respect to Standard Model predictions leads theorists to continue to work on a variety of such models, while not being strongly committed as a group to the ultimate viability of any single model (Mättig and Stöltzner, 2019). These toy models are nevertheless exploited by experimentalists to develop BSM-oriented searches, leaving a possibility of missing an important discovery by looking for new physics in the wrong phase space region. In response to this danger, collider experiments pursue a complementary strategy called signature-based model-independent (SBMI) searches.

SBMI searches look for new phenomena by applying event selections not optimized for a particular BSM theory. Instead, they select the phase spaces to be probed in light of instrument-based arguments such as the inclusivity of a trigger used to collect the dataset or the high resolution expected for the kinematic states of the particles forming the signature (see (Aaltonen et al., 2010; Aad et al., 2013; Chatrchyan et al., 2013), for example). The goal of SBMI searches is to avoid biasing the discovery potential of the overall experimental program of a collaboration like ATLAS toward any particular BSM theory assumptions, and in this way to increase the potential for a discovery of something unexpected. The concern of these searches is therefore less to avoid false discovery claims than to avoid missing a discovery. Decisions about experimental parameters of the search reflect this priority.

5 The Exploratory Character of SBMI Searches

SBMI searches contrast with BSM-oriented searches in a few fundamental aspects. First, the objective of SBMI searches is not to test any specific theory about New Physics, but simply to find some new phenomenon, for example an event count for some signature that the SM does not explain. Such new phenomena need not be predicted by any known BSM theory. To the contrary, the aim is to identify empirical novelties that may guide the development of new theories. Discovering a new phenomenon with SBMI searches would steer BSM theory development by constraining what new models should predict. As such, the epistemic objectives of SBMI searches are consistent with the account Steinle gave of exploratory experiments, while BSM-oriented experiments better fit the account he provides of theory-driven experiments. Second, there is no expectation about the specific new particles that could be found in an SBMI search, while BSM-oriented searches start from the assumption that a specific type of new particle is targeted by the search[8]. Last, SBMI and BSM-oriented searches can be contrasted from a methodological point of view.

While the experimental process put in place in BSM-oriented searches consists in selecting events in the phase space regions optimizing the sensitivity to the BSM theory being tested, SBMI searches typically involve a systematic scan of the kinematic region in which an event signature of interest is reconstructed in order to discover a significant excess compared to SM expectations. The ATLAS search for a resonance, or "bump hunting," in dijet events is one such example ATLAS (2019). In addition, the analysis process of SBMI searches is developed so that it can demonstrate the novelty of a newly discovered phenomenon (via inconsistencies with SM expectations) regardless of what it could be. It would even provide a representation of the target to be explained by a new theoretical extension of the SM. SBMI searches therefore function like exploratory experiments, in contrast to BSM-oriented searches. Note finally that by allowing for the "unexpected" discovery of new particle signals, SBMI searches extend the range of possible outcomes of an experiment like ATLAS. Their experimental procedures are designed for finding just such unexpected signals. They thus incorporate procedures to qualify as experimental exploration in Karaca's sense Karaca (2017). From their objectives, ontological commitments, and experimental processes, SBMI searches for new physics therefore clearly display an exploratory character absent in BSM-oriented searches.

[8]In the following, we use "target phenomenon" to refer to the physical process beyond the Standard Model that is predicted by a specific theoretical hypothesis

However, a more careful evaluation of how SBMI searches are conducted reveals that the exploratory character of these experimental inquiries cannot be mapped to the account of EE outlined in the literature by philosophers like Steinle, Burian, or Karaca. There are two fundamental reasons why existing characterizations of EE fail to apply to SBMI searches, despite their exploratory character outlined above:

1. SBMI searches depend on theory about the target phenomena in a way that has been critically left out of other EE characterizations, including Karaca's;
2. SBMI and BSM-oriented searches cannot be distinguished on the basis of the procedures involved in producing claims about the data and target.

The arguments for the dependence of the SBMI searches on theory about the target phenomena have been presented in Beauchemin (2020). That paper focuses on the ubiquity of theory-ladeness in complex experimental inquiries, but the argument presented applies here, too. Without repeating the argument, we can summarize its three key elements:

- The design of an SBMI search uses theory about some target phenomena when determining the event signature of interest in order to avoid suffering from a large multiple-trial factor which would diminish the discovery potential of the search.[9]
- The results of SBMI searches are always used to constrain some specific BSM theories, and this requires a modification of the observed limits to account for acceptance and uncertainties on the new physics signal targeted in that context that contributes to the selected data.
- SBMI searches are deployed synchronously with BSM-oriented searches as part of a global strategy adopted by large collaboration like ATLAS to "fill the gaps" in the

[9]The multiple-trial factor refers to a consideration in the interpretation of statistical significance when a search is conducted in a manner that could result in claims of a statistical excess across a range of different possible locations in phase space (also called a "look elsewhere effect"). Such an effect is not unknown in more targeted searches (the look-elsewhere effect in the Higgs boson search has been discussed in Dawid (2015) and Staley (2017)), but optimization of such searches limits the effect's potential impact.

phase space covered by the BSM-oriented searches. The two different approaches to search for new physics are therefore not independent of each other.

From the pragmatic perspective presented in Sect. 3, BSM theories are simply resources among others (theoretical, experimental, or practical) available to an experimental inquiry. The decision whether to use an available resource is not to be determined by an epistemological criterion of keeping an identifiable "empirical" core on which experimental knowledge could rest. Available resources are used if they help solve a problem, i.e. if they are useful to reach the epistemic and pragmatic objectives of the inquiry. BSM theory resources may not be central to the development of SBMI searches, but they help these analyses meet their objective of enhancing the possibilities of a discovery or, in the absence of such a discovery, of constraining the possibilities that must be considered in future searches. The three cases of BSM dependence in SBMI searches just discussed indicate how such resources are exploited in the process of designing, conducting, interpreting, and publishing an SBMI search. Rather than forbidding or severely limiting the opportunity to use theory about a target phenomenon to conduct experiment in an exploratory manner, our pragmatic framework makes sense of the way in which exploratory searches do depend on theories of the target phenomena.

The other main reason why EE as characterized in literature fails to apply to SBMI, alluded to in Beauchemin (2020), can be stressed more strongly here: there is no fundamental difference in the way SBMI and BSM-oriented searches are conducted. The exact same process of taking reconstructed events, from the same pool of trigger chains, and of using the same Monte Carlo simulated sample to define some event selections, estimate backgrounds, and apply corrections to the data are used in both types of searches. The same procedures are used to estimate statistical and systematic uncertainties, and the same statistical tools can be used to make inferences from the data. For example, the "bump hunting procedure" used in model-independent dijet resonance searches ATLAS (2019) is very similar to that of the theory-driven search for the SM Higgs boson in the diphoton channel ATLAS (2012). The only difference between the way SBMI searches and BSM-oriented searches are designed and performed resides in the considerations that enter into determining the details of the search, such as the selection criteria for events. In both cases, theoretical considerations about the phase space occupied by potential new physics, as well as experimental considerations of how to select a pure and robust signal, are used to define event selection criteria, but the BSM-oriented searches are more heavily dependent on these BSM theoretical considerations, while SBMI searches focus more on instrument-related criteria to define selections. However, both searches include a combination of elements from both strategies, and once the selections have been defined, the same selected data could equally be used to discover something not expected, or to establish constraints on well-known theories. Hence, even if there were no dependence on theory about the target in the way SBMI searches are conducted, one could not claim, on the basis of procedures, that they belong to a type of experiment distinct from those that are theory-driven. These searches are performed practically identically.

The HEP experiments, understood in the sense defined in Sect. 3, that are similar to SBMI searches have a strong exploratory character and yet these inquiries fail to correspond to the characterization of EE discussed in the relevant literature. It therefore seems that the traditional characterization of EE (or SD) is incapable of properly accounting for the exploratory character of inquiries like SBMI searches. The way experiments and searches for discovery are conducted in HEP is not exclusive to this field of science. A large spectrum of contemporary sciences, from astrophysics to genomics and climate change studies, appear to follow similar general experimental procedures and strategies, and feature a similar level of intertwinement between theory about a target phenomenon and the experiment (see for example **?**). Absent a detailed study of these other fields of science, it is nonetheless plausible that the same restriction in applicability of EE to HEP applies to some, if not many, other fields of contemporary science.

However, the limitations of the traditional characterization of EE are even more severe than this: it fails to discriminate among any experiments that involve quantitative determinations involving assumptions subject to uncertainty. The reason for this failure is that there is no clear demarcation between EE as presented in the literature and systematic uncertainty estimates in this very broad class of experimentation. As briefly summarized in Sect. 2, exploratory experimentation as described by Steinle proceeds by varying experimental conditions to observe empirical regularities without depending on theories about a target phenomenon. When estimating the systematic uncertainties for results of experiments, these same procedures, similarly agnostic about the target, are employed. These variations tell how robust observed experimental patterns are to various experimental conditions, and quantify the impact some of these conditions can have on the observed regularities. This similarity between EE and systematic uncertainty estimation (or other kinds of robustness analysis) is even more striking in Burian's account of EE where the name of the game is to cross-check a large variety of different techniques to robustly establish their findings. This is central to any experimental process. One can make no discovery nor constrain any theoretical model without systematic uncertainty estimates and other robustness analysis. Even the pioneers of electromagnetism, a class of study cases on which Steinle built his arguments, used such systematic variations, agnostic about the target phenomena, to assess the robustness of their observed regularities against the conditions of the experiment. They would for example try to cancel the effect of the Earth's magnetic field and vary the orientation of the apparatus to evaluate how successful they were with such cancelation (see for example Langlois (2005)). Systematic variations to discover patterns of regularities, and systematic variations to establish the robustness of these regularities cannot be disentangled from one another. Any exploratory experimentation would therefore involve procedures that deeply entangle their objectives of pattern discovery and of establishing the robustness of observed regularities. From these considerations and the fact that systematic uncertainty estimates permeate science, the reliance on procedure to distinguish a category of experimentation as exploratory leads to the unwelcome result that either all experimentation that involves estimation of systematic uncertainty (nearly all experimentation in the physical sciences) counts as exploratory, or no such experimentation counts as exploratory.

The arguments given thus far strongly suggest that while the account of EE given in the literature provides excellent "guidelines" (to use Steinle's own term) to identify many episodes of the history of science in which experiments are conducted in an exploratory manner, they fail to demarcate the fundamental aspects of these experimentations that make them exploratory. As we will argue in the next section, our analysis of SBMI searches from the lens of our pragmatic model allows us to identify what makes the exploratory character of an experiment. We for example elucidate this entanglement between the exploratory character of an experimentation and systematic uncertainty estimate and other robustness analysis with our understanding of any experiment as being conducted in both a use and a critical mode of inquiry. We will show that our revised account of exploratory experiment confers on them an epistemological value irreducible to theory test in a way that resolves the challenge posed by Schickore's critique of EE.

6 The epistemological significance of the exploratory character of SBMI searches, and of details of experimental practice

The pragmatic modeling approach directs us toward understanding the epistemic character of any experimental undertaking through careful study of tasks that contribute to its experimental aims: What tasks are performed? What resources are used in those tasks? What aims are pursued through the execution of those tasks? How are the tasks and their associated resources critically assessed with respect to their adequacy for purpose and their successful execution? How are the results of such assessment used in the modification (correction or improvement) of tasks and resources? If the context of exploration has any philosophical significance for our understanding of experimental knowledge, we should expect to find it by answering these kinds of questions.

That SBMI searches and BSM-oriented searches differ in their aims is already evident. Although both probe for evidence of new physics processes, they differ in their priorities.

A BSM-oriented search will be optimized toward its twin objectives: to reveal evidence for a specific phenomenon predicted by a particular BSM model, if that phenomenon exists, and to yield the strongest limits on that model, if it does not. The specification of the search parameters will be chosen so as to maximize the chance of discovery for that specific phenomenon. To carry out such an optimization, the ability to determine the *sensitivity* of the search is crucial. Sensitivity corresponds to the ability of a test to discriminate amongst physical possibilities of interest. In this case, sensitivity requires a discrepancy between the prediction of the targeted BSM hypothesis and that of the SM-only hypothesis that is significantly larger than the uncertainty estimated for the reported result. A search that is very sensitive to the phenomena predicted by a specific BSM model will also be capable of supporting strong limits on that model in the case of a null result. There will be a larger range of theoretical possibilities for that model for which one can say, "if this possibility were the case, we would have seen it in our data, but we did not, so we can exclude it." High sensitivity thus requires, among other things, maintaining the ability to produce results with small uncertainties, since large uncertainties render results compatible with more possibilities in the space of theoretical possibilities. A search that is optimized for a specific BSM model may still be capable of yielding evidence of other possibilities besides those predicted by that model, but will be, in general, less sensitive to those. Data collected through such an optimized search will also be generally less useful for setting limits on models other than that for which it has been optimized.

SBMI searches, in contrast, play a role in the experimental program that complements that of BSM-oriented searches: they aim to reveal evidence for new physics phenomena that may not be targeted by searches optimized for specific BSM models. Consequently, the parameters of SBMI searches are not chosen through an optimization procedure that allows for the determination of sensitivity to the possibilities of a specific BSM model. The pursuit of SBMI searches collectively allows the experimental program to be open to possibilities of new physics that have not been included in the program of optimized BSM searches. In the case that an SBMI search returns a null result, this means that the data from that search will not support limits on a specific BSM model that are as strong as would have been obtained with a search optimized to that model with comparable data. But SBMI results are used to set limits on BSM models, and an SBMI search often yields the best limits set on some model, simply because no optimized search with comparable data has been conducted for that model, or because the event selections used in an SBMI search happen to correspond to those that an optimization for that model would yield.

A further contrast between BSM-oriented and SBMI searches concerns their relations to distinct approaches to theoretical model-building strategies. *Top-down* approaches draw from general theoretical principles (symmetries, capacity to solve long-standing problems such as energy scale hierarchies, theoretical virtues such as naturalness, etc.). *Bottom-up* approaches begin with specific empirical findings, with special attention to any discrepancies or anomalies with respect to SM predictions. The contrast as drawn in this way is quite rough, and deserves a more nuanced discussion than the present context affords. Nonetheless, the contrast is at least suggestive of a complementary contrast between the uses by theorists of BSM-oriented and SBMI searches, with BSM-oriented searches being well-suited to support the efforts of top-down model development and SBMI searches fueling the efforts of bottom-up approaches. (Of course, either type of search is capable of yielding results that a theorist might find helpful whatever their strategy for model development.)

As a consequence of the different aims of BSM-oriented and SBMI searches, and their correspondingly different roles in the experimental programs of research groups at the LHC, these two kinds of searches differ in quite deep levels of detail. A BSM-oriented search and an SBMI search might both involve selecting events with a certain final state, such as an electron, a neutrino, and two hadronic jets ($e\nu + jj$), but the criteria imposed on the data to select an event as an instance of $e\nu + jj$ in these two searches will in general be quite different, because of the differences in how those data selection procedures are developed. Those used in the BSM-oriented search will have been developed through an optimization procedure based on a specific BSM model, while those used in the SBMI search would be developed on the basis of generic considerations of a signature and a detailed understanding of how that signature will be manifested in the instrument outputs of the detector. The arguments for deploying a particular set of selection criteria for an SBMI search will thus not rely principally on a specific BSM model, but instead will be concerned with generic signature attributes and the quantitative determinations that will support acceptable levels of background rejection, efficiency, precision of the efficiency estimate, and so on.

6.1 Answering the challenge: the epistemological significance of practices of inquiry

The complementary contributions in ATLAS's experimental program of BSM-oriented and SBMI searches can be understood in terms of the interplay between use- and critical-mode tasks in pursuit of the aim of discovering new physics. To put it roughly, the apparatus of a BSM-oriented search is used subject to a very rigorous critical scrutiny to ensure that if the targeted physical phenomenon is not manifest, the search will be warranted to not report that it is, but instead to indicate the extent to which it has been ruled out as a possibility. The apparatus of an SBMI search is used subject to a looser critical standard: if some one of a range of – possibly unanticipated – physical phenomena is manifest, the search will result in a warranted (though fallible) claim of discovery, and is not likely to fail to report positively. A generalized version of this idea can serve to characterize the distinctive ways in which experiments conducted in an exploratory manner produce experimental knowledge. That is to say, it can address the challenge posed at the beginning of this essay: what is the philosophical (epistemological) significance of practices of scientific experimentation, such as those documented in the literature on EE?

The key to answering this question lies in the basic architecture of inquiry as a means of producing knowledge. Achieving warrant for assertions about the results of an inquiry requires the use of resources, but it also requires that those resources and the manner in which they are used be capable of withstanding critical scrutiny as to their adequacy for the tasks of inquiry at hand, understood as aimed at that inquiry's objectives. It is insufficient as an account, not only of exploratory experimentation, but of scientific inquiry as a whole, to locate the epistemological significance of experimentation entirely in its role of testing hypotheses for potential falsification or in the terms of any other "theory-dominated" approach to scientific knowledge. The falsificationist seeks to locate the epistemological import of an experiment in the experimenter's logical comparison between a prediction drawn from a hypothesis and a description of some model of the data. The confirmation theorist will locate it in probabilistic calculations based on comparison of the same features. Such comparisons between statements on the "theory" side and the "experiment" side, however, cannot warrant any conclusion at all in the absence of substantial and extensive deployment and scrutiny that warrants both kinds of statements (or models) that experimenters are comparing.

Here we have focused our discussion on the warranting of statements that purport to report the result of an experiment, but in many contexts (certainly in HEP) significant resources and tasks are implicated in arriving at a warranted statement about what a given theoretical model predicts. Such warranting practices are both prior to and independent of judgements about the acceptability of theory based on reported experimental results. How were these predictions and results obtained? How does that matter to the comparison being drawn? An epistemology of experimental science that operates independently of answers to these questions will be at best incomplete, omitting crucial parts of the warranting of scientific knowledge that begin well before a comparison is made between experimental results and theoretical predictions.

The potential epistemological contribution of any experiment must be secured through the execution of specific tasks that are chosen, supplied, and evaluated on the basis of the specific objectives of that experiment. The difference between the exploratory aims of an SBMI search and the theory-driven aims of a BSM-oriented search entails further differences of epistemological significance. Because these aims differ, so will the tasks and the manner in which they are performed. The warranting of assertions issuing from inquiries directed at these aims will likewise differ, insofar as such warranting depends on critical assessment in reference to standards of success for achieving the aims sought. Because different techniques might be alternate routes to a single end, and a single technique might be adapted to different ends, the epistemological difference between SBMI and BSM-oriented searches rests not only on what is done in conducting these inquiries, but on the differing relations between the tasks performed and the aims promoted by those tasks.

Moreover, most experimental inquiries will include among the tasks that contribute to its completion, smaller inquiries with their own outcomes. It was for this reason that procedures of a sort invoked to characterize EE (such as parameter variation) can be found in an apparently theory-driven experiment like a BSM-oriented search. Likewise a procedure that appears to characterize a theory-driven experimentation (comparison of prediction to outcome) can be found in an experiment performed in an exploratory manner. Consequently, any philosophy of experimental science that locates knowledge production only in potential falsfications or only in the confirmation of theory will fail in multiple ways: It will fail, first, to register a great quantity of knowledge production from these sub-inquiries within experiments. As a consequence it will fail to account for any knowledge whatsoever that experimental inquiry yields, because it will not account for the conditions that produce warrant for conclusions of the sort that it does register (whether these concern corroboration or confirmation of theoretical claims). And because it cannot account for such context-specific differences in the manner of warranting outcomes of inquiry, such an approach will epistemologically conflate outcomes of inquiry that are distinct because they constitute knowledge in relation to differing aims, with correspondingly different means of warranting, calling for different critical-mode assessments.

Our account also provides resources to respond to another aspect of Schickore's challenge to Steinle regarding the philosophical significance of EE. Steinle has emphasized the role of EE as a means of developing new concepts in science (Steinle, 2006, 2016). Part of Schickore's challenge asks for a response to the denial that such concept development is epistemologically significant in a strong sense. It is, as a falsificationist might say, a contribution to the psychological stage of conjecture, or, to use the logical positivist jargon, the context of discovery rather than justification. Such contributions are epistemically significant in a weak sense of being part of the process of producing knowledge, by coming up with “epistemically fruitful” ideas (Schickore, 2016, 22), but not in the strong sense of warranting knowledge of a kind different from that achieved in theory-testing experiments. This challenge constitutes another way of asking what do we add to our understanding of the ways in which experimental work warrants knowledge by attending to exploratory approaches to experimentation.

The products of inquiry constitute *resources* to be used in future inquiries. These resources can take many forms, including theories, data sets, models, propositions, and concepts, among others. Schickore, in arguing against the idea that EE is distinctive in producing “know-how” rather than propositional knowledge, writes that “Steinle himself claims that the significant and lasting outcome of exploratory experimentation is conceptual and classificatory, and surely, conceptual and classificatory knowledge must be regarded as propositional knowledge” (Schickore, 2016, 23). Although we do not deny that discoveries involving new concepts can be articulated in the form of propositional knowledge, inquiry that results in warranted knowledge produces resources that need not be propositional, and warranting the usefulness of a concept in particular need not be understood as the production of propositional knowledge.[10] On this view, the use of exploratory modes of experimentation to develop new concepts is just one of various kinds of knowledge-making that experimental practices generate, not in virtue of there being some special class of experiments that qualify for this distinction because they involve distinctive procedures, but in virtue of the way that particular experimental objectives such as the development of representational resources inform the specification and execution of particular experimental tasks and their critical assessment.

[10] Although a full defense of this idea would involve much longer discourse, our view thus embraces a broad range of possible ontologies for knowledge that is friendly to views such as that defended by Davis Baird, who has argued that knowledge can be expressed non-linguistically, in the form of scientific instruments (Baird, 2004).

Our approach provides a comprehensive epistemological framework for a point that has been argued extensively in the practices-oriented literature of the "New Experimentalism," as exemplified in the work of Ian Hacking, Deborah Mayo, and many others. Answering the challenge put forth by Schickore in its most general form amounts to making clear why the fact that "experimentation has a life of its own" has a philosophical significance that goes beyond the fact that producing knowledge by performing experiments requires a great deal of "know-how" and attention to sources of potential error, and beyond pointing out how experimentation involves material engagement with the objects of inquiry through engineered interventions. As significant as these points are, we propose our pragmatic epistemology as a framework for understanding that experimentation has a life of its own also in the sense that epistemological warranting activities pervade the work that experimenters do. For any given experimental task, the epistemological significance can only be understood clearly in relation to both its immediate and overall experimental objectives.

We regard the selection of a specific cut to be applied to the data, for example, as a knowledge-producing task of its own, but only in relation to the objective of that specific choice (to reduce background of specific kinds, by specific amounts, subject to specific efficiency considerations, etc.) and to the objective of the experimental inquiry (to achieve sensitivity to a specific range of BSM phenomena not targeted by optimized searches, or to maximize sensitivity to a phenomenon predicted by a specific BSM model). Whether the choice of a cut has been made well is not just a question of whether someone is good at the craft of particle physics. To make that choice requires using resources (knowledge of the relevant parts of the detector and their performance, actual and simulated data regarding the relevant signals and backgrounds, etc.) and the choice is made to achieve certain objectives (a specific level of discrimination between signal and background, subject to constraints ensuring a certain level of efficiency, in the service of the broader objectives of the experiment). Asking how well a particular choice has been made is a *critical mode* question that involves asking to what degree the resources relied on in making that choice are adequate (how well is the detector and its performance understood? how reliable are the actual and simulated data used as a basis for the decision at hand?), how well has the task of choosing been executed (i.e., were these resources put to use in a manner that is well suited to the objective of this task?), how does the choice impact the results? (as reflected in the systematic uncertainty, for example), and whether it is directed at an appropriate objective (is the choice aiming at the right level of discrimination between signal and background? are the efficiencies aimed at sufficient? do these local criteria sufficiently make sense in light of the overall objectives of the experiment?).

The answers to these questions are, for example, largely different between SBMI and BSM-oriented searches. To explain thoroughly how physicists performing such experiments warrant the claims that result from them, and how the contrasts between them are significant to knowledge of particle physics, requires a detailed study of how these searches are done. We have here indicated some of the details on which such a close study would focus.

In her 1996 *Error and the Growth Experimental Knowledge*, Deborah Mayo argued memorably and compellingly that philosophy of science cannot achieve its aim of understanding how scientific knowledge is warranted by taking a "white glove" approach to the analysis of experimental data. She put detailed studies of experimental practice to work in showing how scientists can and do take steps to ensure that they accept hypotheses on the condition that they have been subjected to sufficiently severe tests on the basis of the data in hand. We endorse this "gloves off" approach here in the pursuit of a broad epistemological account of experimental science in which detailed studies of experimental practice can contribute to our understanding of knowledge production in experimental science by scrutinizing local experimental tasks in the context of both their immediate objectives and broader scientific aims.

7 Conclusion

By deploying a pragmatic epistemological framework to explore the manner in which SBMI searches at the LHC can be understood as having an exploratory character, we have demonstrated the potential for this framework to provide a philosophical account of the warranting function of experimental practices more broadly. By way of advertising the approach, and to encourage others to pursue its potential fruits, we propose two slogans: (1) All epistemology is local. (2) No inquiry is an island.

All epistemology is local. We posit that knowledge can be understood as the outcome of a successfully executed process of inquiry, but criteria of success are not generally available for bulk purchase. Executing a process of inquiry successfully involves carrying out multiple tasks aimed at specific outcomes, and the successful achievement of the outcome cannot be specified independently of the outcome itself, nor independently of the broader context in which that outcome is sought. For example, judging the suitability of a choice of cuts to be applied to jets selected in a search for $e\nu + jj$ events (what is the minimal transverse momentum of the most energetic jet? what is the average angular separation between the two highest energy jets?) in a search that has been optimized for a specific BSM model will appeal to criteria specific that model, at least in the sense that the choice will be made as part of the process of such optimization. If the cuts are being chosen for the purpose of an SBMI search, the relevant criteria will be drawn from more generic or instrument-based considerations that are not so dependent upon any specific BSM model.

No inquiry is an island. Understanding how an experimental inquiry contributes to scientific knowledge involves considering the smaller inquiries that are carried out in the course of that inquiry. As the example just cited shows, those smaller inquiries must also be understood in the context of the larger inquiries of which they are a part. Those larger inquiries also have important relations to one another that, once understood, can help us to see their epistemological import. Dewey's account of inquiry already incorporated these points very clearly (Dewey, 1938). At ATLAS, BSM-oriented and SBMI searches play complementary roles, and are further supported by yet more types of inquiries (measurements of particle properties and other theoretical parameters, for example), to collectively, through a complex network of relations, constitute ATLAS's comprehensive inquiry into fundamental physics. We can see, for example, that a specific BSM-oriented search employs cuts determined in a manner different from that employed in an SBMI search, but to understand that difference (why not optimize all searches?) requires consideration of how the two types of searches complement one another within the experimental program of the search for new physics at the LHC.

We close by returning to Steinle's and Burian's seminal work on exploratory experimentation. Our perspective on these works now reveals them as accomplishing significant philosophical insights that are not dependent on identifying a distinctive class of experiments that can be designated as exploratory in virtue of some particular procedure or aim that they involve. Steinle's discussion of Ampère and of Faraday draws contrasts between exploratory and theory-driven efforts that both undertook, and emphasizes those contrasts that relate to the importance or unimportance of theory in those experimental efforts. The differences concern theory both as a resource (Faraday's discovery of the induction of electrical currents by a changing magnetic field is "driven" by "theoretical speculation") and as aim (Ampère sought in his "attraction experiment" to test a prediction of a theory about the nature of magnetism, while in his "astatic needle" experiment he sought only "a rule giving the behavior of the needle").

It is natural, in reaction to the prevalence of theory-dominated philosophy of science, to see this theory-related distinction as having special significance, but to understand how the inquiries of Faraday or of Ampère succeed in producing knowledge, it is not only the presence or absence of theory in some specific respects that is important but the whole collective of aims, tasks, and resources that they bring to bear. An interest in highlighting ways in which theory might take a backseat to other aims and resources leads Steinle to emphasize certain details of experimental practice and what they are aimed at achieving. Even if the resulting analysis provides only part of the picture, it nonetheless illuminates the very local manner in which knowledge is warranted that lies at the heart of our account.

Burian's discussion is less explicit in attempting to distinguish Brachet's work *from* theory-driven experimentation, choosing instead to place his discussion in the context of Rheinberger's epistemology of experimentation. Nonetheless, his focus on triangulation, on the cross-checking of multiple methods (some novel) for determining spatiotemporal location and biochemical constitution of nucleic acids provides a rich example of the ways in which individual steps of the inquiry process, each with its own aim and using its own resources, knit together a fabric of discovery, warranting judgments about the behavior and characteristics of RNA that, although ultimately fruitful for the development of theory, do "not depend wholly on the specific disciplinary or theoretical background of the experimenters who initiated the work on those objects" (Burian, 1997, 45).

These close considerations of the ways in which experimentalists tailor their practices to their aims, and the wide variety of aims that arise in the context of even a single experiment, let alone an entire domain of inquiry, reveal exploration as one aspect of the production of experimental knowledge. We have argued here that this aspect is philosophically significant not because it denotes a distinct class of experiments or a different kind of knowledge, but because understanding how experimenters pursue their interest in exploring can serve as an example of how to understand the local, connected, and pragmatic nature of experimental activities in general. We have only sketched a framework to pursue this kind of understanding, with much detailed work left to do. The illumination gained here from that sketch leads us to hope for greater insights to come from the more complete picture we are working to produce.

References

Aad, G. et al. 2013. Search for dark matter candidates and large extra dimensions in events with a jet and missing transverse momentum with the ATLAS detector. *JHEP* 04: 075. https://doi.org/10.1007/JHEP04(2013)075. arXiv:1210.4491 [hep-ex].

Aaltonen, T. et al. 2010. Search for new physics with a dijet plus missing transverse energy signature in $p\bar{p}$ collisions at $\sqrt{s} = 1.96$ TeV. *Phys. Rev. Lett.* 105: 131801. https://doi.org/10.1103/PhysRevLett.105.131801. arXiv:0912.4691 [hep-ex].

Antoniou, A. 2021. What is a data model?: An anatomy of data analysis in high energy physics. *European Journal for Philosophy of Science 11*(4): 1–33. https://doi.org/10.1007/s13194-021-00412-2 .

ATLAS. 1999. Atlas: Detector and physics performance technical design report. volume 1 .

ATLAS. 2012. Search for the Standard Model Higgs boson in the diphoton decay channel with 4.9 fb^{-1} of pp collisions at $\sqrt{s} = 7$ TeV with ATLAS. *Phys. Rev. Lett.* 108: 111803. arXiv:1202.1414 [hep-ex].

ATLAS. 2019. Search for high-mass dilepton resonances using 139 fb^{-1} of pp collision data collected at $\sqrt{s}$ =13 TeV with the ATLAS detector. *Phys. Lett. B* 796: 68–87. arXiv:1903.06248 [hep-ex].

Baird, D. 2004. *Thing Knowledge*. Berkeley: University of California Press.

Beauchemin, P.H. 2017. Autopsy of measurements with the ATLAS detector at the LHC. *Synthese* 194: 275–312. https://doi.org/10.1007/s11229-015-0944-5 .

Beauchemin, P.H. 2020. Signature-based model-independent searches at the Large Hadron Collider: An experimental strategy aiming at safeness in a theory-dependent way. *Philosophy of Science 87*(5): 1234–1245. https://doi.org/10.1086/710515 .

Bokulich, A. 2020. Towards a taxonomy of the model-ladenness of data. *Philosophy of Science* 87: 793–806 .

Brown, M.J. 2012. John Dewey's logic of science. *HOPOS: The Journal of the International Society for the History of Philosophy of Science 2*(2): 258–306 .

Burian, R.M. 1997. Exploratory experimentation and the role of histochemical techniques in the work of Jean Brachet, 1938-1952. *History and Philosophy of the Life Sciences 19*(1): 27–45 .

Chang, H. 2012. *Is Water H_2O? Evidence, Realism and Pluralism.* Boston Studies in the Philosophy of Science. Dordrecht: Springer.

Chatrchyan, S. et al. 2013. Searches for long-lived charged particles in *pp* collisions at $\sqrt{s}$=7 and 8 TeV. *JHEP* 07: 122. https://doi.org/10.1007/JHEP07(2013)122. arXiv:1305.0491 [hep-ex].

Dawid, R. 2015. Higgs discovery and the look elsewhere effect. *Philosophy of Science 82*(1): 76–96 .

Dewey, J. 1938. *Logic: The Theory of Inquiry.* New York: Henry Holt and Company.

Elliott, K.C. 2007. Varieties of exploratory experimentation in nanotoxicology. *History and Philosophy of the Life Sciences 29*(3): 313–336 .

Franklin, L.R. 2005. Exploratory experiments. *Philosophy of Science 72*(5): 888–899. https://doi.org/10.1086/508117 .

Karaca, K. 2017. A case study in experimental exploration: Exploratory data selection at the Large Hadron Collider. *Synthese 194*(2): 333–354. https://doi.org/10.1007/s11229-016-1206-x .

Langlois, P. 2005. *Sur la route de l'electricite*. Editions Multimondes.

Leonelli, S. 2019. What distinguishes data from models? *European Journal for Philosophy of Science 9*(2): 22. https://doi.org/10.1007/s13194-018-0246-0 .

Mättig, P. 2022. Classifying exploratory experimentation – Three case studies of exploratory experimentation at the LHC. *European Journal for Philosophy of Science* 12: 66 .

Mättig, P. and M. Stöltzner. 2019. Model choice and crucial tests. On the empirical epistemology of the Higgs discovery. *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics* 65: 73–96. https://doi.org/https://doi.org/10.1016/j.shpsb.2018.09.001 .

Mayo, D.G. 1996. *Error and the Growth of Experimental Knowledge*. Chicago: University of Chicago Press.

Peirce, C.S. 1992[1878]. How to make our ideas clear, In *The Essential Peirce*, eds. Houser, N. and C. Kloesel, Volume 1, 124–141. Bloomington, IN: Indiana University Press.

Popper, K. 1959. *The Logic of Scientific Discovery*. New York: Routledge.

Rheinberger, H.J. 1997. *Toward a History of Epistemic Things: Synthesizing Proteins in the Test Tube*. Stanford: Stanford University Press.

Sargent, R.M. 1995a. *The Diffident Naturalist: Robert Boyle and the Philosophy of Experiment*. Chicago: University of Chicago Press.

Sargent, R.M. 1995b. Exploratory experiments: Scientists at play. unpublished manuscript of a History of Science Society lecture.

Schickore, J. 2016. "Exploratory experimentation" as a probe into the relation between historiography and philosophy of science. *Studies in History and Philosophy of Science Part A* 55: 20–26. https://doi.org/https://doi.org/10.1016/j.shpsa.2015.08.007 .

Staley, K.W. 2017. Pragmatic warrant for frequentist statistical practice: the case of high energy physics. *Synthese* 194: 355–376. https://doi.org/10.1007/s11229-016-1111-3 .

Staley, K.W. 2020. Securing the empirical value of measurement results. *British Journal for the Philosophy of Science 71* (1): 87–113. https://doi.org/10.1093/bjps/axx036 .

Steinle, F. 1997. Entering new fields: Exploratory uses of experimentation. *Philosophy of Science* 64: S65–S74 .

Steinle, F. 2006. Concept formation and the limits of justification. "Discovering" the two electricities, In *Revisiting discovery and justification. Historical and philosophical perspectivces on the context distinction*, eds. Schickore, J. and F. Steinle, 183–196. Dordrecht: Springer.

Steinle, F. 2016. *Exploratory Experiments: Ampère, Faraday, and the Origins of Electrodynamics*. Pittsburgh, PA: University of Pittsburgh Press.

Suppes, P. 1962. Models of data, In *Logic, Methodology and Philosophy of Science Proceedings of the 1960 International Congress*, eds. Nagel, E., P. Suppes, and A. Tarski.

Waters, C.K. 2004. What was classical genetics? *Studies in History and Philosophy of Science Part A 35*(4): 783–809. https://doi.org/https://doi.org/10.1016/j.shpsa.2004.03.018 .