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 the philosophical benefit of a detailed description of a particular experimental practice. 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?

This problem arises in the context of 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; Mättig, 2022; O'Malley, 2007; Waters, 2004; Cobb, 2009). A significant 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.

Jutta Schickore argues that, although the discussion of EE has provided HPS with a framework for analyzing historically significant experimental practices and types of instrumentation, it has not 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.

Historiographically, 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? In this question we can see the outlines of a more general challenge to practices-oriented HPS: Granting that a detailed understanding of how scientists confront particular research problems in practice has value simply in increasing the descriptive adequacy of our accounts, in what way does the study of scientific practice contribute to our understanding of normative features of scientific inquiry, whether in terms of the *warranting* of research results, the *rationality* or *reasonableness* of scientific decisions and inferences, or otherwise? Call this the *normativity challenge for practices* or NCP.[1]

Our response to this challenge employs a pragmatist account of inquiry, inspired by but distinct from John Dewey's theory 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.

---

[1] We are not the first to give a name to this kind of a problem. See, for example, (Tal, 2017) for what Tal calls the "challenge of practice" for measurement.

To demonstrate the potential of this framework for responding to the challenge, we take up Schickore's challenge to EE. To understand how existing accounts of EE might respond to the NCP, we need to understand what those accounts aim to achieve. We consider the possibility that they aim to articulate criteria by which to classify experiments with reference to the manner in which they are conducted, so as to separate those experiments that belong to the exploratory class as distinct from other, non-exploratory experiments.[2] The classificatory repsonse to Schickore's challenge proposes that EE is epistemologically significant because it constitutes a type of experiment capable of generating experimental knowledge that is not accounted for in other philosophical approaches focusing on falsification or confirmation of theories. We will argue that the attempt to pursue this path leads to an impasse from which the only escape is a pragmatic approach of the sort we offer. Using the example of Signature Based Model Independent (SBMI) searches for new physics at the LHC, we argue: first, that SBMI experiments constitute an exploratory experimental inquiry; second, that applying classificatory criteria to SBMI either fails to distinguish SBMI experiments from clear examples of theory driven experiments (TDE) at the LHC or fails to offer a path to answering the NCP; and third, that a pragmatic approach allows us to both make sense of the exploratory character of SBMI experiments and answer the NCP.

We also consider the account of experimental exploration offered by Koray Karaca (2017) as a potential source of a response to the NCP.[3] In agreement with our conclusions, Karaca eschews distinguishing experiments into exploratory and non-exploratory classes.[4] He instead provides criteria for classifying *procedures* as

---

[2] We do not assert that such articulation of criteria for classification purposes is the intent of any particular author who has contributed to the EE literature. We propose instead that discussions of EE are better understood as operating within an implicit pragmatism that we are now elucidating and making explict. We propose the classificatory interpretation only as a prima facie plausible non-pragmatic route toward a response to the NCP.

[3] Karaca's 2017 paper was published shortly after Schickore's 2016 paper; it neither cites Schickore nor responds explicitly to the problem Schickore poses.

[4] Karaca does not entirely avoid classification of experiments however. He does distinguish between *theory-driven* and *strategy-driven* experiments.

exploratory, and thus contributing to *experimental exploration*. If there is some distinguishing intrinsic feature of exploratory procedures, we might then ask whether that feature is the source of the epistemological significance of experimental exploration, and hence the basis for responding to the NCP. However, Karaca's criterion for identifying exploratory procedures relies not on any intrinsic feature of a procedure, but on its relationship to the aims of the experiment to which it contributes. Drawing on the example of SBMI searches, we reinforce Karaca's argument for not regarding any procedure as inherently exploratory, but we also show that his criterion as stated renders the classification of a procedure as exploratory either ambiguous or trivial, and hence not well-suited for responding to the NCP.

On our approach, exploration is not a separate kind of experiment carried out with its own distinct procedures, but a *context of inquiry* 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 falsification or confirmation of theoretical claims. We do not propose to reject analyses offered by Steinle, Burian, and other contributors to the EE literature, nor even to propose an alternative account of EE. Instead we wish to situate the aspects of exploration in experimental inquiry that these authors have highlighted within a new, more explicitly epistemological, perspective that is well-suited to elucidate the normative significance of their detailed descriptions. We argue that in the absence of a successful classificatory approach to responding to the NCP regarding exploratory experimentation, recourse to a pragmatic approach to scientific inquiry is required. Our defense of this pragmatic approach to EE, however, does not depend on rejecting the possibility of a scheme for classifying experiments as either EE or TDE. Rather, our view is that if any such scheme for responding to the NCP were feasible, it would need to incorporate pragmatic considerations of how particular aims structure the choices made by investigators and the ways in which they balance epistemic risks, precisely along the lines of our pragmatic model. Whether any philosophical benefit could be derived from such a classification scheme would require a separate inquiry.

By meeting Schickore's challenge, our approach demonstrates its own resources that can meet the more general challenge we have called the NCP. 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 outcomes that are warranted as 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 Robert Boyle's philosophy of experiment (Sargent, 1995a). Boyle himself articulated a notion of exploratory experiments, in contrast to experiments performed to test a theory (ibid., 137). Contemporary discussion has grown out of works by Friedrich Steinle (1997) and Richard Burian (1997), who both used the term "exploratory experimentation" (henceforth, EE) while studying distinct scientific fields (early nineteenth-century experimentation in electromagnetism for Steinle, early twentieth-century molecular biology for Burian) and emphasizing different, but overlapping, features of the experiments they discussed. We will not undertake a comprehensive survey of the literature (see (Mättig, 2022) for a thorough discussion), but highlight features of EE most salient for our analysis.

Steinle characterizes EE in terms of “typical guidelines” such as 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). Subsequent work by Steinle has emphasized EE as means to the formulation and development of concepts (Steinle, 2006, 2016). 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,” thus enabling the comparison of “different, seemingly incompatible, epistemic objects[5] via . . . cross-checking” (ibid., 45).

---

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

Textual evidence is inconclusive as to whether these early discussions of EE were intended to offer a set of criteria for a dichotomous classification between EE and what Steinle refers to as "theory-driven experimentation" (TDE). Steinle's characterization is in terms of "typical guidelines" rather than "criteria" or even "distinguishing characteristics." He acknowledges that the "new conceptual frameworks" that emerge from EE are "necessarily compared with, and measured against, already established concepts and methodological standards" (Steinle, 1997, S72). Burian describes Brachet's experiments as having been "carried out in a *style* of exploratory experimentation" (Burian, 1997, 28, emphasis added) and focuses on how Brachet's research succeeded in localizing epistemic objects in a manner partially independent of his and others' theoretical commitments. On the other hand, as noted by Maureen O'Malley, "the attraction of EE for philosophical accounts of scientific practice" rests precisely on the opposition between EE and TDE "as two mutually opposed types" (O'Malley, 2007, 351). And Burian, looking to advance philosophy of experiment in its "liberation from . . . excessive theory-centrism," refers to the research he describes as having "involved *exploratory experiments*" (Burian, 1997, 41, emphasis in original), which suggests (although by no means demonstrates) a conception of EE as a distinct category of experiments identifiable by reference to a set of distinguishing characteristics.

Early work by Steinle and Burian also left open the question of what roles for theory are compatible with EE. Mättig notes a "general consensus" that EE is not theory free, and several authors have described positive roles for theory in the conduct of EE (Mättig, 2022, 6). Laura Franklin-Hall's (2005) discusses how the use of "wide" or "high-throughput instrumentation" in genomics draws upon background theory to enable 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 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). 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 Collider (LHC). Although theory-laden, including theory about the target phenomena, Karaca argues that the ATLAS experiment 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. 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."

Our strategy requires extracting from our survey of the EE literature a prima facie plausible set of criteria for identifying an experiment as an instance of EE based on its intrinsic features. We propose the following as reflecting themes emphasized in that literature. An experiment is exploratory only if:

1. No prediction from a prior theory is tested by the experiment.
2. Experimental parameters are varied in the performance of the experiment.
3. The experiment helps generate new representational resources (e.g., concepts).
4. The role of theory in performing the experiment is limited to those modes already discussed in the literature on EE.

We assume that the application of such criteria to identify instances of EE would have to treat these criteria conjunctively rather than disjunctively, since the latter would weaken the concept to the point that clear instances of TDE would be classified as instances of EE.

We emphasize again that we are not proposing this classification scheme as providing an appropriate interpretation of the approach of Steinle, Burian, or anyone else writing on EE. Our argument against this classificatory approach to EE is thus not an objection to what has been written by these authors. Our ultimate aim is instead 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 involves a significant philosophical development beyond these previous efforts insofar as we situate exploration within a more systematic pragmatist epistemological approach to experimentation.

3 Sketch of a pragmatic model of experimental inquiry

We propose to view the exploratory character of experimental inquiry through the perspective of 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,"[6] 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. In the pragmatist spirit, we also emphasize that we offer this model as a useful, but not necessarily *uniquely* useful, framework aimed at greater understanding of the normative significance of practices of exploratory experimentation in particular, and of experimental practices more generally.[7]

---

[6] 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).

[7] Hasok Chang provides another pragmatist-oriented framework that may also shed some light on these issues (Chang, 2022). Although it shares some important features with our view, Chang's approach involves different concepts and emphases. A systematic comparison between his pragmatism and ours would require a separate treatment.

We regard knowledge as the product of a successfully executed process of inquiry[8] 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 (more generally, 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. Our aim here is to articulate a framework for modeling this process of inquiry by considering how the tasks of inquiry – manifested as scientific practices – contribute to the production of judgments and other resources 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 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

---

[8] For Dewey, such an outcome takes the form of a "warrantably assertible" judgment. As explained below, we are open to a wider category of knowledge that includes non-propositional products of inquiry.

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 for completing 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. 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. The pursuit of an experimental inquiry may promote multiple aims; an especially important aspect among the aims of an experimental inquiry is the question of to what phenomena the inquirers intend the experiment 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.

Our discussion of tasks has largely been oriented toward the mode of inquiry in which resources are put to use in the execution of tasks to meet the objectives of inquiry – the *use mode*. Our reference to uncertainty estimation invokes a second, *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 distinction between the use mode and the critical mode is therefore *not* a distinction that creates two separates classes of activities or tasks. 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 use resources to produce assertions in the course of any critical mode conduct. Critical mode warranting activities may also modify or contribute content (such as uncertainty statements) to the assertions that emerge as warranted at the end of an inquiry. 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., propositional knowledge, as well as other products that provide stable and suitable resources for subsequent inquiry. This understanding can already provide some insight into important features of the EE literature discussed in the previous section.

3.1 Experimentation as inquiry

Steinle and Burian, commonly cited as having independently and simultaneously initiated the discussion of "exploratory experimentation" differ significantly in their emphases, as Schickore notes. 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). Schickore thus draws attention to important contrasts between Steinle's and Burian's concerns. Schickore concludes, "the two notions of exploratory experimentation they present are really quite different" (ibid., 21).[9]

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.[10] 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 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

---

[9] 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 manipulation 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.

[10] "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).

innovation of new experimental techniques alongside more established ones.[11] 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, as we will illustrate through our analysis of SBMI searches at the LHC, lies in the interplay between the use mode and the critical mode, an interplay that can only be understood through detailed studies of experimental practices.

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,

---

[11] "[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).[12]

Steinle has not been entirely silent on the question posed by Schickore's challenge, insofar as he has considered that one might regard the generative function of EE as subsumed within the "context of discovery" as opposed to the "context of justification." He argues in response that the formulation of theory requires a stable "highly specialized conceptual framework," and such frameworks are achieved by historically and socially complex processes that are, from an "epistemic perspective . . . of utmost interest and significance," and for the study of which the distinction between contexts of discovery and justification are "of little value" (Steinle, 2016, 334).

Although Steinle's rejection of the distinction between the context of discovery and context of justification is apt, it leaves an epistemological gap. If the warranting of outcomes of experimental inquiry is not captured by the philosophical frameworks associated with that traditional distinction (whether in the idiom of confirmation or falsificationism), then what shall take their place? Steinle resorts here to metaphors: "processes of formation, stabilization, and sedimentation" serve to establish linguistic and conceptual tools or frameworks (ibid.).

Our approach is different, and addresses the gap left by Steinle's response. 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

---

[12] 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.

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 exploration plays out in experimental inquiry, which we illustrate by discussing a class of experiments in HEP (SBMI searches for new physics) that have a clear exploratory intent, yet elude classification as such by the characteristics that have been emphasized in the literature on EE. 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 objectives, such as those relating to epistemic risks, guide decisions about tasks performed in the use and critical modes and their relationships.

3.3 An ambiguity resolved

First we will address an ambiguity concerning the term “experiment” that is especially salient to the HEP context, exploiting 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). This idea certainly gets at something important about exploration in experimental inquiry that our account will not only endorse but expand: it is not the procedure in isolation that can be judged as exploratory or not exploratory. For Karaca, the status of a procedure as exploratory is to be judged by considering it in relation to "the range of possible outcomes of an experiment." However, this account does not tell us what counts as an experiment. 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 (tremendously large and complicated) experimental apparatus: the LHC collider in combination with the ATLAS detector (plus, arguably, 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 all seem to qualify: increasing the beam energy at the LHC, searching for a signature associated with a newly introduced theoretical model, slowing the deterioration of performance for a detector component, and postponing the 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. Labeling all these "exploratory" 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? (Aad et al., 2022) What is the mass of the W boson? (Aaboud et al., 2018) Is the CKKW-L prescription for parton shower to matrix element matching (Lonnblad, 2002) adequate to model $Z+$ jets data? (Aad et al., 2019)) 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.

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.

A remedy lies close at hand in the model of inquiry proposed by Dewey that we take as our inspiration: 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 or more 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?" (Altarelli et al., 1989) 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 inquiry in facts about relations amongst aims (and the relations of tasks and resources to those aims) at a variety of levels: the aims of sub-inquiries that contribute to the experimental inquiry in question, the larger aims that that inquiry itself pursues, and yet broader aims of inquiry to which the experimental inquiry seeks to contribute. Whereas the exploratory character of an experimental undertaking for Karaca is to be located simply in the effect that a procedure has on the range of possible outcomes of the experiment, our pragmatic approach focuses instead on thinking of an experimental inquiry as an undertaking with an intentional structure of objectives exhibiting a degree of complexity and interdependence, guiding choices about what to do and how, and it is to these relations that one should look for understanding how an inquiry takes on an exploratory character.

In Sect. 5 we show the value of this pragmatic model as a means of illuminating the exploratory character of SBMI searches in HEP, while illustrating how doing so requires attention to experimental practices. But first 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 new physics in the form of 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 (Langacker, 2009), 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 specific 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 the level of agreement between the observed data and the corresponding SM predictions. This strategy uses theories about BSM 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

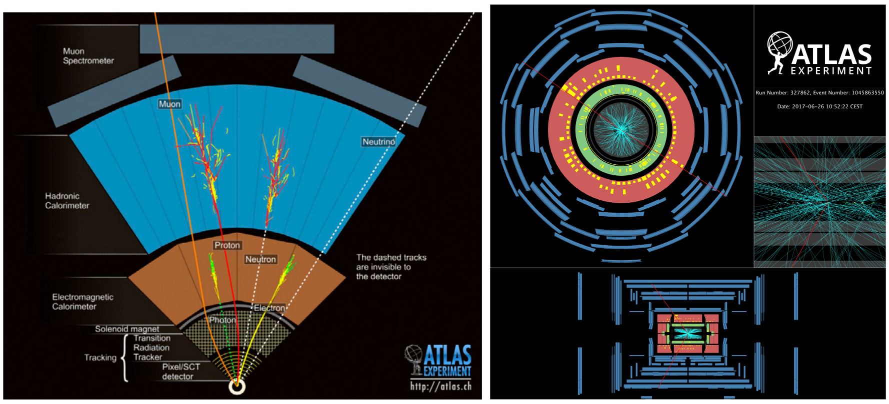

Fig. 1 **Left:** Schematic of how different particles interact with the various ATLAS detector components and the signals they yield. What is identified as a given particle, for example an electron, could be due to another particle mistaken for an electron (a "fake" electron). **Right:** Event display 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).

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 the physics process they are looking for and can control the epistemic risks of an erroneous discovery claim. Securing evidence for or against such physics processes and the theories describing them constitutes an explicit objective of these inquiries.

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. As such, no specific BSM theories constitute explicit objectives of these inquiries, even if the search achieves sensitivity to some specific BSM model that could be targeted by a BSM-oriented search. A search can be sensitive to (even optimally sensitive to) a phenomenon for which it has not been optimized. Instead, the phase spaces to be probed in SBMI searches are selected in light of instrument-based arguments such as the efficiency 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. SBMI searches contrast to BSM-oriented searches in that explicit theories are part of the explicit objectives of the latter and not the former. As we will explain, this does not preclude the use of theoretical resources in both contexts.

5 Can SBMI searches be classified as exploratory?

In this section we first consider SBMI searches in relation to the classificatory criteria introduced in section 2. We reiterate that we do not propose that these criteria should be relied upon as an interpretation of the views of EE advocated by any particular author. Rather, we propose this interpretation as a way in which one might respond to the NCP: Perhaps the philosophical significance of exploratory experimentation or experimental exploration lies in some distinctive normative character of a certain class of experiments or experimental procedures that can be identified by their intrinsic features. We argue that such an approach fails, and that SBMI searches illustrate that failure: exploratory features are not intrinsic either to experiments or to procedures. It is unclear, based on the existing EE literature, what response to the NCP should be proposed in place of the classificatory one, however. In section six, we argue that our pragmatic model of inquiry provides that response in a way that is compatible with the findings of previous authors regarding episodes of EE. Our arguments should not be understood as opposed to any existing discussions of EE. Instead, we propose a preliminary argument *against* the idea that the philosophical significance of EE depends upon identifying a distinct class of exploratory experiments or procedures. We then present our primary argument *for* the idea that a pragmatist account of inquiry enables the philosophical significance of exploratory practices to be explained by appeal to the relations amongst aims and other elements of particular episodes of inquiry. Such relations have been extensively documented in the literature on EE, but not placed explicitly in the philosophical context that enables a response to the NCP.

To illustrate the failure of the classificatory response to the NCP, we first highlight the features of SBMI that reflect its exploratory character in contrast to BSM-oriented searches. This part of the argument relies on features of EE emphasized in the literature without reducing these to the classificatory criteria we propose in section two. We then argue that SBMI searches fail to meet those classificatory criteria for identifying instances of EE proposed earlier, and that efforts to modify those criteria to remedy this will either prevent the classificatory approach from answering the NCP or will lead to the adoption of the very same pragmatic model that we defend.

5.1 SBMI searches appear exploratory in comparison with BSM-oriented 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 those emphasized by Steinle's discussion of exploratory experiments, while BSM-oriented experiments better fit his characterization 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.

SBMI and BSM-oriented searches can also be contrasted from a methodological point of view in a manner that reflects both Steinle's contrasts in the EE literature between EE and TDE and Karaca's description of exploratory procedures. 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 from BSM-oriented searches.

5.2 Classificatory criteria fail to distinguish SBMI as exploratory

However, applying the classification criteria previously introduced fails to distinguish SBMI searches as exploratory in relationship to BSM-oriented searches. The problem concerns the role of theory in SBMI searches: SBMI searches depend on theories about specific BSM phenomena in a way that conflicts with a classificatory approach to exploratory experiments or procedures.

The arguments for the dependence of the SBMI searches on theory about BSM 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:

- *Constraining*: Designing an SBMI search commonly uses theory about BSM phenomena[13] when determining the event signature of interest. Such usage confers

---

[13] Theories about BSM phenomena constitute a spectrum from *fully worked-out models* to *"toy models"* that leave many details unspecified to *parametrizations* that simply expand upon the content of the SM by adding terms to an effective Lagrangian, such as in the SM-EFT approach. Toy models are commonly used in both SBMI and BSM-oriented searches. SM-EFT Lagrangians constitute theoretical representations of potential BSM signatures because of the specific terms in the Lagrangian that they include.

multiple advantages, including avoiding suffering from a large multiple-trial factor which would diminish the discovery potential of the search.[14]

- *Reporting*: The results of SBMI searches are commonly used to constrain some particular BSM theories (e.g., Aad et al. (2013); Chatrchyan et al. (2013)). This is not a mere interpretation of the result, but requires a modification of the observed limits to account for the theory-dependent acceptance and uncertainties on the new physics signal employed.[15] These limits are reported in part as benchmark results to help guide theory development and evaluate gains in sensitivity.
- *Complementing*: 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 phase space covered by the BSM-oriented searches. The two different approaches to search for new physics are therefore not independent of each other. [16]

All three of these modes go beyond what has been countenanced in previous discussions identifying experiments as exploratory, meaning that in spite of their apparently exploratory character explained previously, SBMI searches fail criterion 4. It is also worth noting that "reporting" dependence constitutes a failure of SBMI to satisfy criterion 1. Data from from SBMI searches are commonly used to test predictions of BSM models.

Here one might object that excluding dependence on theory is not a characteristic that emerges clearly from the literature on EE. Some involvement of theory in EE

---

[14] 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. In SBMI searches, this problem can be thought of in terms of the many possible signatures that one might look for, and BSM models help limit the searched phase space not by providing a target for optimization but by guiding the choice of signature (Beauchemin, 2020, 1241).

[15] To our knowledge, previous discussions of exploratory inquiry in HEP, including (Karaca, 2017) and (Mättig, 2022), have not acknowledged this mode of theory-dependence.

[16] This strategy is not dictated by any command from the top of the ATLAS organization but is a consequence of separate decisions of investigators in light of a broadly shared understanding of the general objectives of the experiment. The interdependence of the two approaches is not made less real by the fact that they emerge from a collective solution to a coordination problem.

has been acknowledge by e.g., Elliott (2007), Franklin (2005), and Waters (2004). Our fourth criterion was constructed so as to acknowledge this very point. These discussions have not, however, included the modes of theory-dependence shown here, and so cannot be invoked to classify SBMI searches as satisfying the criteria for EE. Of course, one might then insist that the fourth criterion should be updated to allow for these new modes of theory dependence. Such a strategy threatens to make the scheme unstable, however, as scientists continue to find new ways to deploy theory as a resource for inquiry (something that our model in fact predicts).

What about Karaca (2017)? Karaca allows that BSM models play a role in experimental exploration at ATLAS. However Karaca's notion of experimental exploration does not classify experiments as exploratory, only procedures. He concludes that the ATLAS search for new physics constitutes experimental exploration because it includes exploratory procedures that increase the range of possible outcomes of the ATLAS experiment. We have already argued that this classification of strategies cannot serve to answer Schickore's challenge, and hence cannot address the more general NCP. Thus the only discussion of EE that *might* accommodate some of the modes of theory dependence we have discussed is unable to provide an account of philosophical importance of exploratory inquiry. It also cannot be invoked to account for the exploratory character of SBMI searches in relation to BSM-oriented searches, for reasons that we explain next.

The prospects for learning about what makes for an exploratory inquiry from the classification criteria look even worse when we consider the criteria that SBMI searches do satisfy, particularly criterion 2, regarding the variation of experimental parameters.

Criterion 2 is inadequate for distinguishing SBMI and BSM-oriented searches because of a broader issue: these two types of searches cannot be distinguished on the basis of the procedures involved in producing results. The exact same process of taking reconstructed events, from the same pool of trigger chains, and using a 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 significant difference between the way SBMI searches and BSM-oriented searches are designed and performed resides in considerations that enter into determining such details of the search as 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. 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. These searches are performed practically identically.

The systematic failure of criterion 2 extends far beyond the context of comparing SBMI and BSM-oriented searches. This criterion fails to discriminate among any experiments that involve quantitative determinations depending on assumptions subject to uncertainty. The reason for this failure is that the procedures employed in estimating systematic uncertainty involve the varying of experimental conditions to observe empirical regularities without depending on theories about the phenomenon under investigation. When estimating the systematic uncertainties for results of experiments, such variations enable determining how robust observed patterns are to various experimental conditions, and quantify the impact these have on 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 that emphasizes the cross-checking of a 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 analyses. Even the pioneers of electromagnetism used such systematic variations, agnostic about the phenomena under investigation, 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 cancellation (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, criterion 2 fails to enable any distinction to be drawn with regard to experiments involving evaluation of systematic uncertainty.

SBMI searches exemplify a broader class of exploratory efforts in experimental HEP, yet applying characterizations of EE as classificatory criteria does not result in a clear distinction between SBMI and BSM-oriented searches. The features of HEP experiment we have discussed are 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 experimental procedures and strategies, and feature a similar level of interplay between theory about some phenomenon of interest and experiment (see for example (Pössel, 2020)). Absent a detailed study of these other fields of science, it is at least plausible that such failure of classificatory EE applies similarly to some, if not many, other fields of contemporary science.

The arguments given thus far strongly suggest that while previous accounts of EE provide 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, the characteristics emphasized in those accounts fail when deployed as criteria for dichotomizing experiments or procedures into complementary exploratory and non-exploratory classes.

One might object that the discussion of EE in, for example, Steinle’s treatment of the experiments of Ampére, succeeds precisely because he is relying on classificatory criteria that are more nuanced than those we have employed. According to this objection, it is not enough for EE simply to vary a single experimental parameter. Ampére’s experimentation was exploratory in virtue of the way he varied of a number of different parameters in different ways in order to achieve a representation of phenomenon not anticipated by any theory. By stripping away the context and aims of the experimenter, we render Steinle’s guidelines incapable of functioning in the way that he deploys them.

Our response is that the omission of context and aims from these classificatory criteria is precisely what makes them a non-pragmatic alternative to our approach. As indicated when we introduced them, they are not intended to capture the approach of any particular author writing about EE. They represent rather what one might extract from this literature if one were to attempt to classify experiments as EE or otherwise, based on features of those experiments that *exclude* context and aims.

Steinle and Burian succeed in elucidating the exploratory character of inquiries undertaken by Ampére and Brachet, respectively, not by indicating the presence of intrinsic features of their experiments (particular procedures, theory independence of a particular sort) but by relating the details of their experimental inquiries to broader contexts of inquiry and showing how the aims adopted in those contexts (stabilizing and representing novel electromagnetic phenomena or developing novel techniques needed to identify previously unknown biological agents) serve to underwrite the choices of tasks, resources, and immediate aims in the conduct of particular experimental inquiries. In this section we have shown that attempting to classify SBMI experiments as exploratory by applying classificatory criteria relying on intrinsic features of experiments fails. Had Steinle, Burian, or other authors discussing episodes of EE relied only on such a classificatory approach, we conclude they also would have failed to elucidate any philosophically significant aspect of these inquiries. We take our pragmatic approach to provide a clearer and more explicit account of how these discussions *succeed* in bringing forth an important epistemological feature of experimental inquiry. One might propose that the incorporation of aims and context would allow the formulation of a more nuanced set of classificatory criteria that would succeed better than those we have considered. We have our doubts, but our primary claim would stand because even in such a case we maintain that it would be the pragmatic elements of the account that provide the philosophical significance of exploratory experimental inquiry rather than the classification as such.

To vindicate this last claim, in the next section we provide a more detailed account of how this pragmatic framework applies to the particular context of LHC physics, providing an understanding of the exploratory aspects of experimental inquiry in the case of SBMI searches that also reveals such exploration to have an epistemological value irreducible to theory test, answering the challenge posed by Schickore's critique.

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, i.e., to optimize the sensitivity of the search to the BSM signal of interest. Such a sensitivity is estimated from a quantification of the expected discrepancy between the prediction of the targeted BSM hypothesis and that of the SM-only hypothesis, and from its comparison with the uncertainty estimated for the reported result. The larger the ratio of the former to the latter, the greater the expected sensitivity of the search as designed with respect to the BSM model of interest. 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 (with high probability) 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 optimized for a specific BSM model may still be capable of yielding evidence of other possibilities besides those predicted by that model, but one could not guarantee that the sensitivity to another model would be optimized absent a careful sensitivity a study for that model. Data collected through such an optimized search will be generally less useful for setting limits on models other than that for which it has been optimized.

The role of SBMI searches in the experimental program complements that of BSM-oriented searches: they aim to reveal evidence for new physics phenomena that might be missed by searches undertaken in the experimental program that have been 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. As noted in Sect. 4, they are chosen based on instrument-based considerations that promote sensitivity with respect to phenomena not restricted to specific predictions of a particular BSM model. The point of these considerations is to increase the sensitivity to some range of phenomena by reducing the uncertainty of the SM prediction that is to be compared to a model of the data. The pursuit of SBMI searches 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, the data from that search will not support limits on a specific BSM model 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 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.

BSM-oriented and SBMI searches relate in contrasting ways to distinct 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 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 suggests a complementary contrast between the uses by theorists of BSM-oriented and SBMI searches: BSM-oriented searches are well-suited to support the efforts of top-down model development while SBMI searches fuel the efforts of bottom-up approaches. (Of course, either type of search is capable of yielding results a theorist might find helpful, whatever their model-development strategy.)

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 different, because of the different approaches to developing cuts in the two contexts just described. The arguments for deploying a particular set of selection criteria for an SBMI search will not rely principally on a specific BSM model as in the case of a BSM-oriented search, 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. Although SBMI and BSM-oriented searches employ experimental procedures that are of the same type, as noted in Sect. 5, the details of how those procedures are carried out differ in ways related to their distinct and complementary aims.

The pragmatic account allows us to understand the modes of theory-dependence of SBMI searches that proved problematic for the classificatory approach to EE discussed in Sect. 5. From the pragmatic perspective, theories are simply resources among others (instrumental, representational, etc.) available to an experimental inquiry. The decision whether or not to rely on resource – theoretical or otherwise – is not based on a purely epistemic criterion such as the complete avoidance of theoretical assumptions. 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 in a manner that survives critical assessment. 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 types of BSM dependence in SBMI searches discussed in Sect. 5 (constraining, reporting, complementing) indicate how such resources are exploited in the process of designing, conducting, interpreting, and publishing an SBMI search. Using the results of an SBMI search to constrain a BSM theory is not done to satisfy a higher-level epistemological principle governing SBMI searches, but promotes the objective of learning useful things from data collected in an exploratory manner. The resource and its use are chosen to meet an immediate objective of the inquiry. Here one can see how, rather than forbidding or severely limiting the opportunity to use theory about what is sometimes called the "target phenomenon" in the EE literature, our pragmatic framework makes sense of how an exploratory inquiry like an SBMI search might depend on a theory describing a phenomenon potentially supported by that inquiry: Such theories are useful for promoting specific objectives of an inquiry, and their use is subject to critical assessment to prevent bias or other epistemically objectionable outcomes.

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, although sensitivity to a positive signal in BSM-oriented searches is an important criterion, they are subjected to 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 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. In the case of a negative result, limits will be reported for one or more BSM models, but typically these will be less strict than one would obtain from negative results of a search optimized with respect to a particular BSM model because of the looser constraints on the space of possible discoveries. The difference is a matter of how the risks of an erroneous discovery claim and a missed discovery opportunity are balanced. 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 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, 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. These accounts thus do not provide a sufficient epistemological framework for understanding the theory-directed experiments that have served as contrasts to EE. At the same time, remedying these defects in confirmation-oriented and falsificationist accounts will not yield a reduction of the exploratory character of experimentation to instances of theory-testing or -confirmation precisely because it will require attending to details of the dynamics of inquiry – the warranting that comes about through the interplay of use- and critical-mode tasks, that we have shown to differ in the two cases.[17]

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

---

[17] 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.

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. In the case of BSM-oriented searches, work done in the critical mode serves to argue that the tasks done and the resources used achieved the best sensitivity to or constraints on the targeted BSM model, in light of the resources available. In the case of SBMI searches, work done in the critical mode serves to argue that if there is any new physics signal in the final state probed, the inquiry conducted would have found it. These arguments are very different, and yet they both serve to warrant the conclusions of their respective inquiries.

Moreover, an experimental inquiry will include among the tasks that contribute to its completion, smaller inquiries with their own outcomes. It is 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.[18] 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

---

[18] Karaca has acknowledged this phenomenon in the context of experimental exploration (Karaca, 2017). Here we provide the rationale behind its philosophical importance.

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. 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 (Steinle, 2016). 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.

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.[19] On this view, the use of exploratory modes of experimentation to develop new concepts is one among numerous kinds of knowledge-making that experimental practices generate, not in virtue of special classes of experiments defined in terms of 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.

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 NPC amounts to making clear why the fact that "experimentation has a life of its own" has a philosophical significance 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 the normative dimensions of any given inquiry are its own. 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, 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

---

[19] 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).

(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). To choose well in the setting of cuts, for example, it is not sufficient to competently execute a series of relevant tasks in the use mode alone. One must subject that execution and its results to a critical mode assessment that is apt for the specific experiment and its objectives. 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? and 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, and an important reason for these differences can be attributed to the way that the former seek to fulfill exploratory aims that do not play a similar role in the design and execution of the latter. 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.

If our argument seems to suggest that the understanding of knowledge production requires attention to rather mundane facts regarding the obvious fact that scientists have to engage in the same kind of adaptation of means to ends and adjustment of ends to one another and to the available means that can be found in any other practical undertaking, we happily grant that this is precisely what we mean to claim. If one were to further conclude that this makes the point we are making trivial, our response is that it is far from trivial to claim (as we do) that accounting for the power of scientific inquiry consists in nothing more than attending to precisely these "mundane" considerations.

In other words, what makes for the significance of the ways that methods and tools are adjusted to the aims of an inquiry is the significance of the aims, tools, and methods themselves and differences amongst them. The importance of EE as a distinction from other modes of inquiry is no more and no less than the importance of the difference between the aims of exploration and the aims of non-exploration and the fruitfulness of those aims in fostering important new tools (resources) and methods (tasks).

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 criteria for 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 new physics in $e\nu + jj$ events in a search that has been optimized for a specific BSM model will appeal to criteria specific to that model, invoked through a process of 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 do not aim at being specifically sensitive to any BSM model. The pragmatic conception of knowledge adopted here transforms this quite ordinary coordination of means and ends into the very substance of epistemological analysis.

No inquiry is an island. A pragmatic approach is also the natural home for a point that motivates many analyses undertaken in HPS studies of experiment, including the contributions to the EE literature considered here: 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 our discussion of SBMI searches and their implementation 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. 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 with respect to a theoretical model?) 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" (Steinle, 1997, S68)) and as an 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” (ibid., S67)).

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 procedures, 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

Aaboud, M. et al. 2018. Measurement of the W-boson mass in pp collisions at $\sqrt{s} = 7$ TeV with the ATLAS detector. *Eur. Phys. J. C 78*(2): 110. https://doi.org/10.1140/epjc/s10052-017-5475-4. arXiv:1701.07240 [hep-ex].

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].

Aad, G. et al. 2019. Measurement of the inclusive cross-section for the production of jets in association with a Z boson in proton-proton collisions at 8 TeV using the ATLAS detector. *Eur. Phys. J. C 79*(10): 847. https://doi.org/10.1140/epjc/s10052-019-7321-3. arXiv:1907.06728 [hep-ex].

Aad, G. et al. 2022. A detailed map of higgs boson interactions by the ATLAS experiment ten years after the discovery. *Nature 607*(7917): 52–59. https://doi.org/10.1038/s41586-022-04893-w .

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].

Altarelli, G., B. Mele, and M. Ruiz-Altaba. 1989. Searching for new heavy vector bosons in $p\bar{p}$ colliders. *Zeitschrift für Physik C Particles and Fields 45*(1): 109–121. https://doi.org/10.1007/BF01556677 .

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.

Chang, H. 2022. *Realism for Realistic People: A New Pragmatist Philosophy of Science.* Cambridge: Cambridge University Press.

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].

Cobb, A. 2009. Michael Faraday's "historical sketch of electro-magnetism" and the theory-dependence of experimentation. *Philosophy of Science* 76: 624–636 .

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 .

Langacker, P. 2009, Aug. The physics of heavy $Z^{'}$ gauge bosons. *Rev. Mod. Phys.* 81: 1199–1228. https://doi.org/10.1103/RevModPhys.81.1199 .

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 .

Lonnblad, L. 2002. Correcting the color dipole cascade model with fixed order matrix elements. *JHEP* 05: 046. https://doi.org/10.1088/1126-6708/2002/05/046. arXiv:hep-ph/0112284 .

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.

O'Malley, M.A. 2007. Exploratory experimentation and scientific practice: Metagenomics and the proteorhodopsin case. *History and Philosophy of the Life Sciences 29*(3): 337–360 .

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.

Pössel, M. 2020. A beginner's guide to working with astronomical data. *The Open Journal of Astrophysics 3*(1). https://doi.org/10.21105/astro.1905.13189 .

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.

Tal, E. 2017. A model-based epistemology of experiment, 233–253. New York: Routledge.

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 .