My work cuts against the current trend in philosophy towards ever narrower specialization. Though I list my central research areas as philosophy of mind, epistemology, and philosophy of science, I am, in truth, a generalist. I approach my research in perception with an eye towards understanding mathematical and scientific reasoning; I assess debates in formal epistemology in light of my views on the metaphysics of mental states; and I address questions concerning scientific methodology through the lens of my account of the nature of conscious experience. It is my view that we can hope to make progress on the central questions of contemporary philosophy only by taking an appropriately broad perspective, combining insights from many different areas of research.

I have three central research projects, each of which exemplifies this approach.

Dissertation-Related Work

Two of my research projects stem from my dissertation, titled *Sensible Concepts: Experience and the A Priori*, in which I develop a new account of spatial experience that situates our experience of space within a broader context of non-sensory cognitive activities. On my account, to perceive an object as square is, in part, to deploy the very same Euclidean concept of squareness that we utilize in *a priori* mathematical reasoning. Such geometrical concepts *feature in*, but are not *derived from*, experience.

My first research project applies this account of perceptual experience to questions about the nature of our mathematical reasoning and the interpretation of scientific theory.

A PRIORI ASPECTS OF MATHEMATICAL COGNITION: In a paper titled "*A Priori* Concepts in Euclidean Geometry" (a version of which has been invited to be considered for publication in the *Proceedings of the Aristotelian Society*), I argue that we must recognize a *sui generis* form of *a priori* cognition at work in the practice of Euclidean proof. Classical geometrical reasoning, unlike the logical deduction systems of later mathematicians such as Hilbert, is not purely formalistic – it does not reduce to pure logic. And yet, at the same time, the concepts involved in such reasoning are genuinely *a priori*: as I show, many of our Euclidean geometrical concepts, such as the concept of a continuous curve, cannot be derived from experience. I consider several recent attempts, from the literature on the philosophy of mathematical practice, to explain Euclidean reasoning as stemming from empirical cognition, and I argue that none can offer a satisfying account of our geometrical thought.

CONTEMPORARY PHYSICS AND THE INTERPRETATION OF SCIENTIFIC THEORY: In a paper titled "Shape Perception in a Relativistic Universe" (a version of which is forthcoming in *MIND*), I consider how we should evaluate the veridicality of our experience of shape in light of Einstein's special theory of relativity (STR). The challenge from STR is a particular instance of a larger question that I explore within the framework developed in my dissertation: How should we understand the relationship between the conceptual resources we deploy in our perceptual experience of the world—resources that are, on my view, largely *a priori* in origin—and the picture of physical reality we get from scientific investigation? According to my account, our very *possession* of a conception of the physical world depends on our deployment of *a priori* concepts in perception. So scientific theories, which themselves depend on perception and aim to reveal truths about the physical world, cannot consistently renounce our *a priori* conceptual scheme in its entirety: doing so would leave those theories with no coherent account of the very evidence on which they are based, or of the world they are intended to describe. This imposes a useful constraint on scientific theorizing: we must interpret our scientific theories in a way that makes room for an intelligible account of our perceptual contact with the empirical world. I argue that some interpretations of central theories in contemporary physics—including not only STR, but also the general theory of relativity and quantum mechanics—run afoul of this constraint.

My second research project further elaborates the account of spatial experience developed in my dissertation and extends the framework to other domains of perception.

A PRIORI CONCEPTS IN SPATIAL EXPERIENCE: According to a widely-held view of spatial experience known as structuralism, perceptual representations of spatial features are merely *structurally isomorphic* to abstract Euclidean geometry; they do not themselves comprise substantive Euclidean concepts. In a paper titled "More Than Mere Structure," I show that this structuralist view fails to explain the way in which we apply our Euclidean concepts to the spatial features we perceive. For, on the structuralist picture, the results of Euclidean geometry would be equally applicable in perception to *any* set of features isomorphic to Euclidean space. Colors are one such set of features: their variations along the dimensions of hue, saturation, and brightness can used to generate a (metaphorical) color "space" that maps onto the structure of Euclidean space. But we do not perceive colors, in spite of their being isomorphic to the features we reason about in Euclidean proof, *as* instances of Euclidean spatial relations – we do not see groups of objects as, say, square in virtue of their color properties. It is only when we perceive the *literal* spatial features of objects—for example, when we see a chessboard *as a square*—that we take our geometrical concepts to be applicable. This shows that, unlike in the case of color, the connection between our *spatial* experience and our geometrical reasoning is more than merely structural.

CORE COGNITION AND SPATIAL PROTO-CONCEPTS: Given the connection I have highlighted between our *a priori* and empirical representations of space, we need an account of spatial experience that can explain *how* the very same set of *a priori* concepts could show up both in Euclidean proof and in perceptual experience. Building on recent work by developmental psychologists Susan Carey and Elizabeth Spelke on "core cognition" systems, I argue (in a paper titled "*A Priori* Concepts in Spatial Cognition") that we have an innate, primitive grasp of basic spatial properties—a set of core "proto-concepts"—that feature in our cognitive lives in two ways. On the one hand, these proto-concepts are deployed in our experience of the world around us, representing, in a qualitatively rich format, the presence of particular spatial features. On the other hand, with the acquisition of more sophisticated cognitive and linguistic capacities, these same proto-concepts are built up into full-blown geometrical concepts, which allows for their explicit articulation in the practice of Euclidean proof – a practice through which we explore in detail the spatial features that our innate proto-concepts represent.

TEMPORAL COGNITION: Going forward, I plan to utilize the framework developed in my analysis of spatial cognition in investigating the nature of our concepts and experience of time. Unlike spatial properties, the temporal properties we perceptually represent are properties our experiences *themselves* can have: both a headache and the transit of Venus can last for six hours; both an itch and a phone's ring can occur at regular five-second intervals. This provides the grounds for an intriguing proposal about the origins of our temporal concepts: our experiences supply us with concepts of time in virtue of *instantiating* temporal properties. On this picture, we can grasp the nature of the temporal properties we represent *worldly* events as having because we have immediate access to *experiences* that instantiate those very properties. Temporal concepts, then, share features with both our *a posteriori* color concepts and our *a priori* spatial concepts. On the one hand, temporal concepts, like color concepts, are derived from experience. On the other hand, temporal concepts, like spatial concepts, give us insight into the *nature* of the properties they represent.

Formal Theories of Rational Decision-Making and Scientific Investigation

My third research project focuses on foundational issues in formal epistemology, decision theory, and philosophy of science. These fields, I argue, suffer from inadequate attention to the relation between their mathematically-articulated theories and our intuitive understanding of the phenomena analyzed. Our commonsense concepts of mental states like belief and preference provide the only possible starting point for the development of a precise formal theory of rationality. So a firm understanding of those mental

states is a prerequisite for any such theory. In my research, I investigate the nature of rational decision-making and scientific discovery through the lens of my background in philosophy of mind. This approach opens up a novel perspective on questions typically addressed through more purely formalistic methods.

BAYESIAN CONFIRMATION THEORY: It is typically assumed that, in applying the standard Bayesian rule for assessing hypotheses in the light of evidence, we ought to reason on the basis of *all* of the available information. In a paper titled "The Fine-Tuning Argument and the Requirement of Total Evidence" (forthcoming in *Philosophy of Science*), I argue that, surprisingly, this "requirement of total evidence" often fails to deliver the correct verdicts. In analyzing the evidential import of a given observation, we must not be overly-specific in our *classification* of that observation. Using a series of simple experimental cases, I show that, by classifying observations in the most-specific way possible, we run the risk of masking the confirmation our experimental procedures in fact supply. This puzzling phenomenon raises important questions about how to model sound reasoning about probabilistic processes. I go on to explore these questions in the context of the so-called "fine-tuning argument" for the existence of multiple universes. I argue that, in assessing questions about the process that gave rise to our own universe, the possible outcomes of that process should be partitioned coarsely, without reference to our own unique perspective as the beings who actually resulted from it. Thinking in this way avoids the kind of *post hoc* assessment of evidence we recognize as problematic in more mundane cases.

UTILITY IN DECISION THEORY: In a paper titled "Against Operationalism," I argue for a realist interpretation of the decision theoretic notion of *utility*. Anti-realism or "operationalism" about utility—which is the orthodox view amongst decision theorists—denies that the notion of utility that features in decision theory represents any real psychological phenomenon. Instead, on the operationalist interpretation, utility is just a convenient way to represent an agent's choice behavior. Such a picture, I argue—by divorcing the notion of utility from our pre-theoretical conception of the *strengths of our preferences*—leaves the axioms at decision theory's foundation without their needed mooring. If decision theory is to serve as a theory of practical rationality, its axioms must be defensible as rational requirements. This defense, in turn, is possible only if we understand the mathematically-precise concept of utility employed in the theory as deriving from our commonsense understanding of the preferences that drive us to act. And those preferences, given the kinds of agents we are, come in degrees of strength. Thus, we need to embrace a *realist* conception of utility as a measure of the strength of our preferences in order to explain how decision theory can provide normative guidance for us as human agents.

SELF-LOCATING BELIEF: Two very different kinds of mental states are often lumped together in formal epistemology under the label "credence." The first is the kind of state an agent is in when she believes that a probabilistic process has some objective chance of resulting in a particular outcome – for instance, when she believes that a coin has a 50% chance of landing heads. The second kind of state occurs when an agent simply lacks relevant information – when she is *ignorant* about some fact. In some instances—particularly those involving self-locating belief—this second kind of state tracks propositions that cannot intelligibly be taken to involve any kind of "probabilistic process" at all. I may not know whether it is Monday or Tuesday; but that is not the same as my having the bizarre belief that some random process, like the tossing of a coin, has somehow "determined what day it is." Thus, classifying these two different kinds of mental states simply as "credences" obscures an important distinction between them. In a paper titled "Sleeping Beauty and the Principal Principle," I show that the standard argument for the "1/3" answer to the Sleeping Beauty problem ignores this distinction, making it guilty of a kind of equivocation in its use of terms like "likelihood." I offer an alternative analysis, arguing that we should not let the mere ignorance embodied in self-locating uncertainty undermine the contentful judgments about probabilistic processes embodied in our assessments of objective chances. As a result, we should accept the ½ answer to the Sleeping Beauty problem.