Asking the Right Questions About the Psychology of Human Inquiry: Nine Open Challenges

Anna Coenen¹, Jonathan D. Nelson^{2,3}, Todd M. Gureckis¹

New York University¹, Max Planck Institute for Human Development², University of Surrey³

The ability to act on the world with the goal of gaining information is core to human adaptability and intelligence. Perhaps the most successful and influential account of such abilities is the Optimal Experiment Design (OED) hypothesis, which argues that humans intuitively perform experiments on the world similar to the way an effective scientist plans an experiment. The widespread application of this theory within many areas of psychology calls for a critical evaluation of the theory's core claims. Despite many successes, we argue that the OED hypothesis remains lacking as a theory of human inquiry and that research in the area often fails to confront some of the most interesting and important questions. In this critical review, we raise and discuss nine open questions about the psychology of human inquiry.

Keywords: inquiry; information search; information gain; optimal experiment design; active learning; question asking

Introduction

The ability to ask questions, collect information, and actively explore one's environment is a powerful tool for learning about the world. How do people decide which information to collect in any given situation? One influential idea is that information-seeking or inquiry behaviors are akin to *scientific experiments*. According to this metaphor, a child shaking a new toy, a student asking a question, or a person trying out their first smartphone, can all be compared to a scientist conducting a carefully designed experiment to test their hypotheses (Gopnik, 1996; Montessori, 1912/1964; Siegel, Magin, Tenenbaum, & Schulz, 2014). The core assumption in this work is that people optimize their queries to achieve their learning goals in the most efficient manner possible.

To model everyday inquiry as scientific experimentation, psychologists have been inspired by

the concept of “optimal experiment design” (OED) from the statistics literature (Fedorov, 1972; Good, 1950; Lindley, 1956). OED is a general statistical framework that quantifies the value of a possible experiment with respect to the experimenter's beliefs and learning goals and can help researchers plan informative experiments. The psychological claim is that humans perform “intuitive experiments” that optimize the information gained from their action in a similar way. Within psychology, the OED hypothesis has been applied in many different areas including question asking, exploratory behavior, causal learning, hypothesis testing, and active perception (for overviews, see Gureckis & Markant, 2012; J. D. Nelson, 2005; Schulz, 2012b).

It is easy to see why this metaphor is attractive to psychologists. Not only does the OED hypothesis offer an elegant mathematical framework to understand and predict information-seeking behaviors, it also offers a flattering perspective on human abili-

ties by suggesting that everyone is, on some level, an intuitive scientist. However, the status of OED as the dominant formal framework for studying human inquiry calls for a critical evaluation of its explanatory merits.

This paper addresses two overarching issues concerning the current use of OED as a psychological theory. First, existing OED models rely on a wealth of non-trivial assumptions concerning a learner's prior knowledge, beliefs, cognitive capacities, and goals. Our analysis critically examines these assumptions and lays out future research directions for how to better constrain these choices. Second, some forms of human inquiry cannot be easily expressed in terms of the OED formalism. For example, inquiry does not always start with an explicit hypothesis space, and it is not always possible to compute the expected value of a question. To that end, we highlight research questions that lie outside the realm of the OED framework and that are currently neglected by the focus on inquiry as scientific hypothesis testing.

Our hope is that this paper will serve both as a critical comment on the limits of the OED hypothesis within psychology and a roadmap of some of the hardest but most interesting psychological questions about human inquiry. The main part of the paper takes the form of laying out nine questions about inquiry. For each question, we review the current literature on the topic, examine how past work has dealt with particular challenges, and suggest promising future directions for work within and outside the OED framework. Before turning to these nine key questions, we review the origin and core principles of the OED hypothesis, and its history within psychology. We then consider how best to evaluate the past successes of the framework.

Human inquiry as optimal experiments

The metaphor of intuitive human inquiry as scientific experimentation dates to the 1960s. This

early work compared people's hypothesis testing to philosophical norms of scientific experimentation, and most prominently to principles of falsification (Popper, 1968). Falsification turns out to be a relatively poor description of human behavior, however, and is now widely rejected as an explanatory model (Klayman & Ha, 1989; Klayman, 1995; Nickerson, 1998; Wason, 1960). In contrast, the OED framework, which was inspired by Bayesian norms of experiment design from statistics (Horwich, 1982), has a number of successes as a predictive theory and is gaining in popularity among psychologists.

The origins and use of OED models. OED methods were originally developed as statistical tools to help researchers plan more informative scientific experiments (Good, 1950; Fedorov, 1972; Lindley, 1956). The idea is to create some formal measure of the "goodness" of a particular experiment with respect to the possible hypotheses that the experimenter has in mind. Using this normative measure, researchers can then choose an experiment that is most conducive to discriminating among the possible hypotheses. This is an alternative to experiments that are intuitively designed by researchers themselves but that might not be optimally informative. For example, a cognitive scientist studying human memory might choose different delay intervals for a recall test following study. Parameters like these are typically set using intuition (e.g., to cover a broad range of values). An OED method might instead output specific time intervals that have the best chance to differentiate competing theories (e.g., a power law or exponential forgetting function, see Myung & Pitt, 2009). The advantage of the OED method is that seemingly arbitrary design choices are made based on principled analyses of the researcher's current knowledge about possible hypotheses (or models).

Starting from a problem or situation that the experimenter (or human learner) is attempting to understand, most OED models are based on the following components (see below for more mathematical detail):

Figure 1. An overview of human inquiry from the perspective of OED theories. Such theories begin with an ambiguous situation that provokes an inquiry goal. For example, the learner might wonder why the cat is in a bag. In thinking about how to best obtain the answer the learner is assumed to consider alternative hypotheses about the explanation. Next the learner evaluates possible actions or questions they could ask. Such questions are evaluated largely on how informative the answers to the questions would be. Finally, a question is chosen and the learner updates their belief based on the answer. OED theories capture the information processing elements within the thought bubbles.

1. A set of hypotheses (e.g., statistical models or range of parameter values) the experimenter or learner wants to discriminate among;
2. A set of experiments or questions that the experimenter or learner can choose from (e.g., parameters of a design or types of conditions);
3. A model of the data that each experiment or question could produce, given the experimenter's current knowledge;
4. A measure of the value of these outcomes with respect to the hypotheses (e.g., the difference in model likelihood, or confidence about parameter values).

Together, these components enable a researcher to compute an “expected value” of every possible

experiment, and choose the experiment that maximizes this value. This involves a *preposterior analysis* (Raiffa & Schlaifer, 1961), during which experimenters have to simulate the potential outcomes of every experiment and compute how helpful each of these outcomes would be for achieving their goal.

OED methods have been used by experimenters to improve parameter estimation and model comparison. For example, psychologists have used them to discriminate different memory models (Cavagnaro, Myung, Pitt, & Kujala, 2010; Myung & Pitt, 2009), to compare models of temporal discounting (Cavagnaro, Aranovich, McClure, Pitt, & Myung, 2014), to improve teaching tools for concept learning (Rafferty, Zaharia, & Griffiths, 2014), to fit psychophysical functions (Kim, Pitt, Lu, Steyvers, & Myung, 2014), and even to discriminate between different

models of human inquiry (J. D. Nelson, McKenzie, Cottrell, & Sejnowski, 2010).

Aside from scientific applications, OED concepts are also widely used in machine learning to develop algorithms that rely on *active learning*. Such algorithms have the capacity to self-select their training data in order to learn more efficiently (Mackay, 1992; K. P. Murphy, 2001; Settles, 2010). For example, they can decide when to ask a human to provide a label of an unclassified training instance (e.g., a document). Active learning is especially useful when it is costly or time-consuming to obtain such corrective feedback.

OED modeling in psychology. Somewhat separately from these applied domains, researchers in psychology have used the OED formalism as a theory or hypothesis about human inquiry behavior. OED models have been used to explain how young children ask questions or play with an unfamiliar toy (Bonawitz et al., 2010; Cook, Goodman, & Schulz, 2011; Gopnik, 2009; McCormack, Bramley, Frosch, Patrick, & Lagnado, 2016; J. D. Nelson, Divjak, Gudmundsdottir, Martignon, & Meder, 2014; Ruggeri & Lombrozo, 2015; Schulz, Gopnik, & Glymour, 2007), how people ask about object names in order to help them classify future objects (Markant & Gureckis, 2014; J. D. Nelson et al., 2010; J. D. Nelson, Tenenbaum, & Movellan, 2001), and how people plan interventions on causal systems to understand how variables are causally related to one another (Bramley, Lagnado, & Speekenbrink, 2015; Steyvers, Tenenbaum, Wagenmakers, & Blum, 2003). They can also model how learners would search an environment to discover the position of objects in space (Gureckis & Markant, 2009; Markant & Gureckis, 2012), and where they would move their eyes to maximize the information learned about a visual scene (Najemnik & Geisler, 2005, 2009). Figure 1 illustrates how these basic components of the OED framework might map onto an everyday scenario facing a human learner.

What is common to all these approaches is the claim that the mind operates, at least indirectly, to

optimize the amount of information available from the environment just as OED methods optimize the information value of experiments. It is the broad application and success of this theory that makes it both interesting and worthy of critical evaluation. We will start our discussion of the OED framework by laying out its principles in more mathematical detail.

Formal specification of OED models

An OED model is a mathematical way to quantify the *expected value of a question*, query, or experiment for serving a learner’s goals. The basic approach is related to expected utility models of economic decision making, but uses utilities that are informational in nature, rather than costs and benefits of correct or incorrect decisions. Importantly, OED models are designed to not depend on which hypotheses a researcher personally favors or dislikes. OED models define the expected utility of a question as the average utility of that question’s possible answers. Formally, a question $Q = \{a_1, a_2, \dots, a_m\}$ is a random variable with possible answers $a_1, a_2, \dots, a_m$. The expected utility of that question, $EU(Q)$, is defined as the average utility that will be obtained once its answer is known, i.e.: $EU(Q) = \sum_{a_j \in Q} P(Q = a_j)U(Q = a_j)$.

Utility can be any function that measures a learner’s progress towards achieving their goal of inquiry, which could be pure information gathering, planning a choice, or making a judgment. The learner’s goal is often to identify the correct hypothesis. The possible hypotheses (or states of the world) are defined by a random variable $H = h_1, h_2, \dots, h_n$. Many OED utility functions are based on the prior and possible posterior probabilities of each hypothesis $h \in H$, and on how the distribution of probabilities would change according to each possible answer that could be obtained.

For concise notation in this paper, rather than writing out both the random variable and the value that it takes, we will specify the value that the ran-

dom variable takes. For instance, suppose we wish to denote the probability that a specific question Q , if asked, would result the specific answer a . Rather than writing $P(Q = a)$, we will write $P(a)$. It is important to emphasize that a specific answer a is associated with a specific question Q . Or suppose we wish to denote the probability of a specific hypothesis h , given that question Q has been asked and that answer a has been obtained. Rather than writing $P(H = h|Q = a)$, we will simply write $P(h|a)$. Thus, the expected utility (usefulness) of a question Q can be concisely written as

$$EU(Q) = \sum_{a \in Q} P(a)U(a) \quad (1)$$

A learner is typically faced with a set of possible questions $\{Q\}$. (The curly braces denote that we are referring to a set of questions, $Q_1, Q_2, Q_3, \dots$, each of which is a random variable, rather than to a specific single question.) To determine the optimal question, a learner has to calculate the expected utility of each possible individual question by simulating the possible answers of each question, calculating the usefulness of each answer, and weighting each possible answer's usefulness as in Equation 1.

One of the most prominent OED usefulness functions is the expected value of a learner's *gain in information* or *reduction in uncertainty* (Austerweil & Griffiths, 2011; Cavagnaro et al., 2010; Lindley, 1956; Najemnik & Geisler, 2005; J. D. Nelson et al., 2014; Oaksford & Chater, 1994). A common metric of uncertainty is Shannon entropy, although alternative ways of measuring the value of an outcome will also be discussed below. The information gain of a particular answer, a , to question Q , is the difference between the prior and the posterior entropy:

$$U_{IG}(a) = ent(H) - ent(H|a) \quad (2)$$

The prior Shannon entropy is

$$ent(H) = \sum_{h \in H} P(h) \log \frac{1}{P(h)} = - \sum_{h \in H} P(h) \log P(h), \quad (3)$$

and the posterior entropy is

$$ent(H|a) = \sum_{h \in H} P(h|a) \log \frac{1}{P(h|a)} = - \sum_{h \in H} P(h|a) \log P(h|a) \quad (4)$$

where the posterior probability of each particular hypothesis h is derived using Bayes' (1763) rule:

$$P(h|a) = P(h)P(a|h)/P(a). \quad (5)$$

The combination of Equations (1) and (2) yields the *Expected Information Gain* (EIG) of a query, $EU_{IG}(Q)$.

How Psychologists Develop OED Models of Inquiry

To illustrate how the key components of the OED framework can be mapped onto different experiment paradigms in psychology, consider the list of examples in Table 1. What is impressive about this list is the broad range of human behaviors that have been modeled by way of the OED hypothesis. While this table gives a cursory impression, in the following section we review in detail three example studies that use OED to model hypothesis testing, causal learning, and children's exploratory play. We particularly aim to highlight what types of conclusions theorists have drawn from their models and behavioral findings.

Example 1: Logical Hypothesis Testing. In the most well-known psychological application of OED, Oaksford and Chater (1994) revisit the classic Wason card selection experiment (Wason, 1966). The experiment tests whether people are able to logically test hypotheses by falsifying them, that is, by checking that there are no counter examples to a logical rule. Participants are asked to test for a simple conditional rule involving a

Table 1

How the main OED components map onto particular research topics in psychology.

Examples	Inquiry Goal	Hypotheses (<i>H</i>)	Queries (<i>Q</i>)	Answers (<i>a</i>)
Causal Learning Bramley, Lagnado, & Speekenbrink, 2015; Steyvers et al., 2003; Coenen, Rehder, & Gureckis, 2015	What is the structural pattern of how different variables in a system influence one another?	Causal models that relate the variables of the system in different ways (e.g., causal Bayes nets).	Setting the value of one or more variables through an intervention.	The effects of making an intervention on the causal system.
Categorization Markant & Gureckis, 2014; Markant, Settles, & Gureckis, 2015; J. D. Nelson et al., 2010	What types of shapes are called “daxes”?	Rules about category membership (e.g., “all red things are daxes”).	Different category members or objects to test.	The category membership of the queried item.
Spatial Search Gureckis & Markant, 2009; Markant & Gureckis, 2012; Rothe, Lake, & Gureckis, 2016	What is the hidden spatial pattern?	Different hypotheses about the identity of hidden patterns.	Tiles to turn over in a game or verbal questions that can be asked about the pattern.	The color of a hidden tile or information about the pattern.
Eye movements Najemnik & Geisler, 2005	Where should I move my eyes next to search for a target?	Possible target locations within the visual field.	Locations to move the eyes.	What is visible as a particular location.
Rule learning Wason, 1960; Klayman & Ha, 1989; J. D. Nelson et al., 2001	What rule determines if a sequence numbers is valid or invalid?	A rule about number sequences (e.g., “odd, increasing numbers”).	Different sequences of numbers (e.g., “2-4-6”).	Valid or invalid according to the rule.
Twenty Questions J. D. Nelson et al., 2014; Ruggeri & Lombrozo, 2015; Mosher & Hornsby, 1966	What object does the answerer have in mind?	All possible objects, animals, or people.	Features of the object (e.g., “is it blue?”).	Yes or No.
Logical Reasoning Wason, 1966; Oaksford & Chater, 1994	Find out if cards with vowels have even numbers on the other side.	Vowels always predict even numbers or the two are independent.	Turn over cards with a vowel, consonant, odd, or even number.	Symbol on the other side of the card.

set of four cards. The four cards are labeled “A”, “K”, “2”, or “7” and participants are asked to test if “cards with a vowel on one side have an even number on the other side” (a rule of the form, *if p, then q*). The dependent measure is which of the four cards participants turn over. (Participants are allowed to select all, none, or any subset of the four cards.) An often-replicated pattern of results is that most people select the “A” (*p*) card, many choose the “2” (*q*) card and few choose the “7” (*not-q*) card. This pattern of preferences violates the logical norms, which dictate that one needs to test “A” (*p*) and “7” (*not-q*), but not “2” (*q*). The “7” card (*not-q*) card is crucial, because it could potentially be a counterexample if it had a vowel on the other side.

To explain the discrepancy between people’s choices and reasoning norms, Oaksford and Chater (1994) interpret the task as a problem of inductive inference (how does a learner anticipate to change their beliefs based on data), rather than as checking for violation of a logical rule. Oaksford and Chater propose that people choose queries to reduce their uncertainty about two hypotheses: The dependence hypothesis specifies that the logical rule holds perfectly. The independence hypothesis specifies that the letters (A vs K) are assigned independently of the numbers (2 and 3) on the other side of the cards. Oaksford and Chater compute the *expected information gain* (see Equation 2 above) for each query (card). The model assigns values to different queries (each

of the four cards that can be turned over) and considers possible outcomes from these queries (observing a vowel, consonant, even number, or odd number). In the model, Oaksford and Chater further assume that the “A” and the “2” are *rare*, that is, learners do not expect many cards to have either vowels or even numbers on them. Given these assumptions, it turns out that the expected information gain from testing “2” is actually greater than that of testing “7”, which matches the pattern of behavior often found in this task.

In this article, Oaksford and Chater (1994) apply OED methods as part of a *rational analysis* of the card selection task (Anderson, 1990) that uses an optimal model to capture people’s behavior given some additional assumptions, but without any commitment to a particular set of cognitive processes that underlie this behavior. Regarding the actual implementation of the computation, the authors note that “The reason that people conform to our analysis of the selection task might be due to innate constraints or learning, rather than sophisticated probabilistic calculation.” (Oaksford & Chater, 1994, p. 628). This is an example of a successful OED analysis that does not involve any algorithmic or implementational claim. Oaksford and Chater’s model also illustrates how researchers in the rational analysis framework have often adopted assumptions that make human behavior seem reasonable, rather than looking for deviations from a particular set of logic- or other norm-based assumptions of how people should behave.

Example 2: Causal Learning. Another type of inquiry that has been modeled with OED norms is causal intervention learning. Steyvers et al. (2003) used expected information gain to predict which variables participants would manipulate to figure out the underlying causal relationships. In their experiment, participants first spent some time passively observing the behavior of a causal network of mind-reading aliens. The aliens’ thoughts were depicted as strings of letters appearing over their heads. Participants were told to figure out

which aliens could read which other aliens’ minds (which resulted in them thinking the same thought as the one they were reading from). After the observation phase, participants gave an initial guess about how they thought the aliens were causally connected. Then, they were asked to make an intervention by inserting a thought into one of the aliens’ heads and observing the thoughts of the other aliens.

Again, the authors modeled these choices using an OED model based on expected information gain, which aims to reduce the uncertainty about possible causal structure hypotheses. Here, queries corresponded to the different aliens that could have a thought inserted, and outcomes corresponded to all possible ways in which the other aliens could change their thoughts as a consequence. The hypothesis space contained possible causal structures that described how the aliens were connected (i.e., who could read whose mind). The authors considered a number of implementations of the OED model that differed with respect to the space of hypotheses. An unconstrained version of the model, which considered all possible causal structures connecting the aliens, was not a good fit to people’s choices. However, a more constrained version, which assumed that people were only comparing their top hypothesis to its own subgraphs (i.e., graphs containing a subset of the edges of the most likely graph) and an unconnected graph, fit the human data well.

The authors concluded that “people bring to bear inferential techniques not so different from those common in scientific practice... they choose targets that can be expected to provide the most diagnostic test of the hypotheses they initially formed through passive observation.” (Steyvers et al., 2003, p. 486). Unlike the previous example, this conclusion suggests that people actually implement the underlying computations associated with the OED model. This interpretation is also common in work on OED models of inquiry.

Example 3: Exploratory Play. Finally, consider an example from the developmental litera-

ture on children's capacities for inquiry. Cook et al. (2011) gave preschoolers different information about the causal properties of toys (beads) and examined their subsequent behavior during exploratory play. Children were either shown that *all* beads were causally effective (they could turn on a machine and make it play music) or that only *some* beads are effective (some could not turn on the machine). Subsequently, children were given a new set of two beads that were attached to each other. Children who had learned that *only some* beads are effective proceeded to take apart the two new beads and test them with the machine individually. By contrast, children who had previously learned that *all* beads worked rarely bothered to check the new beads individually.

This behavior can also be modeled with expected information gain, by assuming that learners are choosing between three possible queries (testing both beads, testing the first bead, and testing the second bead) and anticipating one of two outcomes (the machine turning on or not). The experimenter's demonstration is designed to set children's hypotheses about the new pair of connected beads. Children in the all-beads-work condition only have a single hypothesis (both beads work), while those in the some-beads-work condition have four (both work, one works, the other works, neither works). To reduce their uncertainty about these hypotheses, the model predicts that the beads must be tested in isolation, which matches the behavioral data.

This example illustrates a trend in the developmental literature to draw analogies between children and scientists. Without making concrete algorithmic claims, Cook et al. (2011) interpret their findings as evidence that even young children optimize information gain during inquiry in a scientific manner, and conclude that "these results tighten the analogy to science that has motivated contemporary theories of cognitive development" (Cook et al., 2011, p. 348).

These three examples illustrate not only different psychological applications of OED models

but also the different types of explanatory claims that OED analyses have supported, ranging from the computational-level observation that people behave *as if* they optimize informational value (in Oaksford & Chater, 1994) to the more ambitious idea that people, like scientists, actually implement OED computational principles to some degree (in Steyvers et al., 2003). Although the actual explanatory claims may vary significantly from study to study, a common thread remains the tight analogy between empirical science and human information-seeking.

It should be noted that the history of psychology also offers examples of researchers using the OED framework to support the opposite claim that human information-seeking does *not* follow rational and scientific principles. For example, some studies in the heuristics and biases tradition (Kahneman, Slovic, & Tversky, 1982) highlighted ways in which human judgments deviate from OED norms (Baron, Beattie, & Hershey, 1988; Skov & Sherman, 1986; Slowiaczek, Klayman, Sherman, & Skov, 1992). Similarly, prior to the Bayesian approach used by Oaksford and Chater (1996), research on logical rule learning showed many discrepancies between OED principles and human behavior (Klayman & Ha, 1987; Klayman, 1995; Wason, 1960). Despite this history, the people-as-scientists metaphor has by far outweighed these account in recent years.

Merits of the OED hypothesis

The OED approach has greatly contributed to the study of human inquiry. Perhaps most saliently, it has provided a computationally precise approach to some very open-ended aspects of human behavior. In addition, the OED hypothesis provides a theoretical account of diverse information-seeking behaviors, ranging from visual search to question asking. In doing so, it also builds a theoretical bridge to models of a wide array of other cognitive processes, which, on the surface, bear little or no resemblance to information search. For example, Information Gain and related principles

have been used in models of receptive properties of visual neurons (Ruderman, 1994; Ullman, Vidal-Naquet, & Sali, 2002) and auditory neurons (Lewicki, 2002). They are also key components of recent models of visual saliency, which predict human eye movements as a function of image properties (Borji & Itti, 2013; Itti & Baldi, 2005, 2006; Zhang, Tong, Marks, Shan, & Cottrell, 2008). They also connect to Friston and colleagues' (e.g., 2009; 2017) free energy principles, which posit that *all* neuronal activity is aimed at minimizing uncertainty (or maximizing information).

Finally, the close connections between OED models in psychology and formal methods in mathematics, physics, statistics, and epistemology make it straightforward for psychological theory to benefit from advancements in those areas. For example, research in computer science on computationally efficient active learning machines has inspired new theoretical approaches to inquiry behavior in humans (Markant et al., 2015; Rothe et al., 2016).

Limitations of the OED hypothesis

Despite its successes, this article critically examines some of the basic elements of the OED research approach. Our critique springs from two main points, which, at first glance, may seem contradictory. On the one hand, applications of the OED framework in psychology often rely on a wealth of non-trivial assumptions about a learner's cognitive capacities and goals. There is a risk that this makes the models too flexible to generate testable predictions. On the other hand, we will argue that the framework is in some cases not rich enough to capture the broad types of inquiry behavior exhibited by humans. These latter cases are particularly important because, as OED gains in popularity as a theoretical framework, there is a risk that important aspects of behavior are being overlooked.

Elaborating on the first point, the three example studies reviewed above demonstrate a frequent re-

search approach that is shared by many applications of the OED hypothesis within psychology. First, it is assumed that people inquire about the world around them in order to maximize gain in knowledge. Second, this assumption is instantiated as a specific OED model which assigns values to different questions, queries, or actions in a particular task. Finally, additional assumptions about cognitive processes (hypotheses, priors, etc.) may be added to the model to improve its fit.

Importantly, this research strategy does not set out to directly test the core claims of the OED hypothesis. For some researchers the framework provides a set of starting assumptions and novel psychological insights are more likely to emerge from modifications of a model's peripheral components that get adjusted in the light of behavioral data. For instance, in Oaksford and Chater's (1994) analysis of the card selection task the model fits behavior under the assumption that events (p and q) occur rarely. Similarly, Steyvers et al.'s (2003) best-fitting rational test model relies on a very restricted space of causal graph hypotheses. It is common for OED models to rely on very specific assumptions, but less common for researchers to treat these assumptions as discoveries in their own right. The rarity prior in Oaksford and Chater (1994) is an exception in this respect and provides a good example of integration between OED models and their assumptions. The rarity assumption is implicated in other hypothesis testing research, has normative support from the Bayesian literature, and it has generated a number of follow-up studies that systematically manipulate it and find that behavior changes accordingly (McKenzie, Ferreira, Mikkelsen, McDermott, & Skrabble, 2001; Oaksford & Chater, 1996; Oaksford, Chater, Grainger, & Larkin, 1997). In general, however, it is rare for OED applications to examine and justify their auxiliary assumptions in such detail.

This general lack of integration between a formal framework and its assumptions about requisite cognitive components is a common criticism leveled against other classes of models, particu-

larly Bayesian approaches for modeling higher-level cognition (Jones & Love, 2011; Marcus & Davis, 2013). Importantly, as both critics and defenders of Bayesian models have pointed out (see peer commentary on Jones & Love, 2011), this kind of criticism does not require rejecting the entire framework, but can be addressed by promoting greater efforts towards theory integration at different levels of explanation (e.g., computational, algorithmic, and ecological). The same holds for OED models of inquiry. Many of the current limitations of the framework could be overcome by moving beyond the mere metaphor of people as intuitive scientists and beginning to take the role of auxiliary assumptions seriously. This is the approach we advocate in some parts of this paper.

On the second point, there are also ways in which using the OED hypothesis as a starting assumption limits the kinds of behavior studied in the field. Recall that to make an inquiry problem amenable to an OED analysis, a researcher must quantify the set of hypotheses a learner considers, their prior beliefs over these hypotheses, the set of possible queries available to them, and their probability model for the outcome of each query. As we will note throughout this paper, there are many kinds of inquiry behaviors that would be difficult or impossible to express in those model terms, either because they are not based on the same components (e.g., inquiry in the absence of well-defined hypotheses), because of the computational complexity of applying OED, or because we do not yet know how to specify them computationally as part of a model (e.g., query types with computationally complex outcomes). Of course, no psychological theory is able to capture every single interesting cognitive phenomenon in a broad area like inquiry. However, we believe that it is important to pay close attention to the kinds of limits a theory imposes and make sure they do not lead to an overly narrow focus on a small set of questions that happen to be amenable to a particular analysis. Our review highlights the challenges of capturing important inquiry behaviors with OED and aims to

encourage future research in these directions. We also highlight a number of questions that fall entirely outside of the purview of OED analyses, but that we believe deserve more attention in the study of human inquiry.

Nine Questions about Questioning

In the following sections we address what we think are some of the most interesting unresolved psychological questions about human inquiry. The criticism is built around the following nine questions about human inquiry:

1. How do people construct a set of hypotheses?
2. How do people generate a set of candidate queries?
3. What makes a “good” answer?
4. How do people generate and weight possible answers to their queries?
5. How does learning from answers affect query selection?
6. How do cognitive constraints influence inquiry strategies?
7. What triggers inquiry behaviors in the first place?
8. How does inquiry-driven learning influence what we learn?
9. What is the developmental trajectory of inquiry abilities?

Each section is designed to operate somewhat independently so readers are encouraged to read this article in a nonlinear fashion. In addition, at the beginning of certain sections that deal with variables or terms in the standard OED equations (i.e., Equations 1-5), we reprint the relevant equation and highlight the particular component of the OED framework that is discussed.

Question 1: How do people construct the space of hypotheses?

A crucial foundation for being able to use an OED model is the set of hypotheses or hypothesis space, H , that a learner considers. One reason is that the most common measure of information quality (Information Gain, Equation 4) depends on changes in the entropy over the space of possible hypotheses:

$$ent(H|a) = - \sum_{h \in H} P(h|a) \log P(h|a)$$

What
is in
 H ?

The genesis of hypothesis spaces and priors in models of cognition is an issue that has been raised with respect to Bayesian models of cognition (Goodman, Frank, Griffiths, & Tenenbaum, 2015; Griffiths, Chater, Kemp, Perfors, & Tenenbaum, 2010; Jones & Love, 2011; Marcus & Davis, 2013), but plays out in particularly interesting ways in the OED framework.

What is a hypothesis or hypothesis space? Hypotheses often are thought of as reflecting different possibilities about the true state of the world (related to possible world semantics, Ginsberg & Smith, 1988). Hypothesis sets may contain discrete objects (like causal structures, category partitions, or even dynamic physics models, Battaglia, Hamrick, & Tenenbaum, 2013). Alternatively, a hypothesis space might reflect a distribution over continuous quantities (e.g., locations in space), or model parameters. The examples in this article often focus on discrete cases, since they tend to be more commonly used in OED models of higher-level cognition. However, the issues we raise also apply to some continuous hypothesis spaces.

How do current psychological applications of OED models define this hypothesis space? If the domain of inquiry is sufficiently well-defined,

modelers often assume that learners consider an exhaustive set of hypotheses. For example, in categorization tasks the full set includes every possible partition of the space of objects into categories (Anderson, 1991; Markant & Gureckis, 2014; Meder & Nelson, 2012; J. D. Nelson, 2005). In causal learning scenarios, the hypotheses might be all possible (direct and acyclical) graphs (or possible parameterizations of graphs) that might explain the causal relationships between a number of variables (Bramley, Lagnado, & Speekenbrink, 2015; K. P. Murphy, 2001; Steyvers et al., 2003). In a spatial search task, the hypothesis set could consist of all possible locations and orientations of an object (Markant & Gureckis, 2012; Najemnik & Geisler, 2005). This exhaustive approach can lead to the following three problems.

First, fully enumerated hypothesis spaces can be very large and complex, even in relatively simple tasks with a well-defined structure. For example, the number of possible partitions of objects into categories grows exponentially with the number of objects (Anderson, 1991; Berge, 1971). Similarly, the number of possible causal graph hypotheses increases rapidly with each additional variable (2 variables yield 3, 3 variables 25, 4 variables 543, and 5 variables 29281 possibilities). In real world situations, the number of candidate category members and potential causal variables often far exceeds the situations used in psychological experiments, exacerbating the issue.

Given limited cognitive capacities, it seems unlikely that people can consider hundreds or thousands of discrete hypotheses and update their relative plausibility with every new piece of data. In fact, empirical studies often find that people appear to consider only a limited number of hypotheses in probabilistic reasoning tasks (Dougherty & Hunter, 2003b). Hypothesis set size in some tasks also scales with working memory capacity (Dougherty & Hunter, 2003a), which suggests that cognitive load could influence hypothesis set size. Some studies even argue that people consider only one hypothesis at a time in various

learning and decision-making tasks (Bramley, Dayan, & Lagnado, 2015; Courville & Daw, 2007; Sanborn, Griffiths, & Navarro, 2010; Vul, Goodman, Griffiths, & Tenenbaum, 2014).

Another conceptual problem is that hypothesis sets are not always easy to define from the perspective of the modeler. Although it is sometimes obvious what belongs in a hypothesis set for a particular task, there are many cases in which it is much less clear. For example, imagine a child shaking a new toy for the first time. What should we assume about her hypotheses, given that she has never seen a toy like this before? And how should she reduce her uncertainty about these hypotheses as efficiently as possible? Of course, it is possible that, based on prior experience with other toys, she is testing some high-level possibilities, for example whether or not the toy makes any noise when shaken. However, it is also possible that she chooses actions in line with more low-level principles of reducing prediction error about the outcome of her own motor actions. In that case, her hypothesis space might consist of a generative model that links actions, world states and percepts, and that can be used to quantify the expected surprise associated with self-generated actions (for such a general formulation of action as active inference, see Friston, 2009). Alternatively, the best model of this kind of behavior might not involve *any* hypotheses. Instead, the child's behavior might be the outcome of some internal drive to explore and act on the world that is independent of particular beliefs or goals (Hoch, Rachwani, & Adolph, in review).

Confronting these conceptual and practical challenges is critical for models of inquiry. Here we address three possible approaches that have been used in recent research and discuss the merits of each. They include restricting hypothesis spaces, focusing on single hypotheses, and forming queries with no hypotheses whatsoever.

Curtailed hypothesis spaces. One solution to the combinatorial explosion of hypotheses is to select only a few hypotheses at a time and try to be-

have optimally given this subset. This is viable when it is possible to enumerate all hypotheses *in principle*, but the complexity of the full space is large and cognitive limitations forbid considering the whole set.

There is some evidence that people consider such pared down sets of hypotheses when seeking information. For example, Steyvers and colleagues' (2003) causal intervention study, the best-fitting OED model was one that restricted the hypothesis set to a single working hypothesis (causal graph), as well as its "subgraphs" and a null model in which all variables were independent. Oaksford and Chater (1994) made a similar modeling assumption by considering only two possibilities about the world, one in which the conditional *if p then q* holds, and one in which *p* and *q* are entirely independent. However, there are many other logical relationships that could exist between them (e.g., the inverse conditional or a bi-conditional).

If some reduction of a hypothesis space provides a better account of human inquiry, an interesting question for the field becomes how to develop theories of this process. One approach is to model more directly the processes that might be used to construct a hypothesis set. Currently there are few such algorithmic theories, with the exception of a model called *HyGene* (Dougherty, Thomas, & Lange, 2010; Thomas, Dougherty, Sprenger, & Harbison, 2008). When encountering new data, *HyGene* generates hypotheses that have served as explanation for similar types of data in the past. Past data is retrieved from memory based on its similarity to the current data, and working memory capacity places an upper bound on the number of retrieved items. This subset of hypotheses is then evaluated with respect to the current data, and inconsistent hypotheses are ruled out. Since hypothesis generation in *HyGene* is based on memory retrieval processes, this approach would be particularly useful for modeling inquiry in domains where learners have a certain degree of prior knowledge (e.g., clinicians diagnosing diseases).

Alternatively, hypothesis spaces may be constructed on the basis of other processes. For example, comparison has been shown to promote relational abstraction, which in some cases might help bootstrap new types of hypotheses (Christie & Genter, 2010). According to this idea, comparison between two objects invokes a process of structural alignment where different features and relations of the objects are brought into correspondence with one another. In doing so, comparison has been shown to help focus people on shared relational structure, making these commonalities more salient for subsequent processing (e.g., similarity judgments). Thus, comparison might also help alter the hypothesis space considered for inquiry behaviors, by highlighting relational features.

One approach to formalize curtailed hypothesis generation comes from *rational process models* (also often simply referred to as *sampling algorithms*, see Bonawitz, Denison, Griffiths, & Gopnik, 2014; Courville & Daw, 2007; Denison, Bonawitz, Gopnik, & Griffiths, 2013; Gershman, Vul, & Tenenbaum, 2012; Sanborn et al., 2010; Vul et al., 2014). These models explain how a computationally limited organism can approximate Bayesian inference by making simplifying assumptions about how hypotheses are maintained and updated. Instead of representing the complete posterior probability distribution over possible hypotheses, the idea is that learners *sample* from this distribution, and thus only maintain a subset of hypotheses at any point in time. One feature of these models is that they can account for sequential dependencies during learning. For example, under certain parameterizations particle filter models yield hypotheses that are “sticky”, that is, that once considered will be maintained and only dropped when a learner encounters strong conflicting evidence (related to win-stay-lose-shift models of belief updating, Bonawitz, Denison, Gopnik, & Griffiths, 2014). This stickiness property matches human learning data in some tasks and is therefore considered an

advantage of rational process models over “purely rational” models of hypothesis generation and belief updating (Bonawitz, Denison, Gopnik, & Griffiths, 2014; Bramley, Dayan, & Lagnado, 2015; Brown & Steyvers, 2009).

However, current sampling models lack a robust coupling of model terms and psychological processes. For example, it is unclear how the (re-)sampling of new hypotheses from the current posterior might be implemented. A promising direction is to integrate ideas from algorithmic models like HyGene that ground similar computations in mechanistic accounts of memory retrieval (Gershman & Daw, 2017; Shi, Griffiths, Feldman, & Sanborn, 2010).

Rational process models face another challenge. Much of their appeal is based on the fact that, under certain limiting conditions, they converge toward the true Bayesian posterior. Consequently, many have argued that they might bridge between optimal analyses and mechanistic accounts of behavior (Bonawitz, Denison, Griffiths, & Gopnik, 2014; Brown & Steyvers, 2009; Jones & Love, 2011; Sanborn et al., 2010). In reality, however, many of these algorithms require hundreds or thousands of samples in order to converge. Cognitive psychologists, on the other hand, often find that humans use considerably fewer samples, even as few as one (Vul et al., 2014), possibly because sampling incurs cognitive or metabolic costs. One skeptical interpretation of this work is that it implies that Bayesian inference is too costly for the brain. Also, if people sample stochastically, it should be rare that any single person acts optimally during inquiry (S. Y. Chen, Ross, & Murphy, 2014). Instead, these theories predict that people will be optimal or unbiased *on average* (across people or situations). This property of sampling models, if correct, would suggest significant changes to the way OED models are evaluated. For instance, researchers would need to start quantifying optimality at a group level rather than for individuals (e.g., Mozer, Pashler, & Homaei, 2008) or based on data from repeatedly testing a participant

on the same task. This may require larger experimental populations and new experiment designs.

Single-hypothesis queries. One common finding is that learners seem to seek information for a *single* hypothesis at a time. Although this can be seen as just a special (most extreme) case of curtailing hypothesis sets, single-hypothesis queries have rather unique characteristics and have motivated countless psychological experiments and models. Since OED models so fundamentally rely on a process of *discrimination* between competing hypotheses (see Figure 1), single-hypothesis queries have been particularly difficult to explain.

For example, in Wason’s (1960) “2-4-6” task, participants are asked to find out which numeric rule the experimenter is using, knowing only that the sequence “2-4-6” satisfies this rule. In this task, many participants immediately generate the working hypothesis that the rule is “even numbers increasing by 2” and proceed to test this rule with more positive examples, like “4-6-8” (Klayman & Ha, 1989; Wason, 1960). This has been called a *positive testing strategy* (PTS). Because it can yield suboptimal behaviors, it is also cited as an example of confirmation bias, that is, the tendency to verify one’s current beliefs instead of seeking and considering conflicting evidence (Klayman & Ha, 1987, 1989; Nickerson, 1998).

Single hypothesis use and the failure to consider alternatives have been observed in many areas of cognition besides information search. For example, during sequential learning people often only maintain a single hypothesis, which gets adapted with new evidence over time (Bramley, Dayan, & Lagnado, 2015; Gregg & Simon, 1967; Markant & Gureckis, 2014; Nosofsky & Palmeri, 1998; Trueswell, Medina, Hafri, & Gleitman, 2013). When dealing with objects that have uncertain category membership, people often base their inference on the most likely category, ignoring its alternative(s) (Malt, Ross, & Murphy, 1995; G. L. Murphy, Chen, & Ross, 2012; Ross & Murphy, 1996). During causal

reasoning, people frequently make predictions based on single causes and neglect the possibility of alternatives (Fernbach, Darlow, & Sloman, 2010, 2011; Hayes, Hawkins, & Newell, 2015).

The ubiquity of single hypothesis reasoning is not easily reconciled with the metaphor that people act like intuitive scientists, even after conceding that they are subject to cognitive limitations. Since model discrimination lies at the heart of the metaphor, it seems difficult to argue that single-hypothesis queries are the output of an optimal learner in the OED sense. However, it turns out that the PTS maximizes information under certain assumptions. For example, the PTS is optimal (in the OED sense) when hypotheses only have few positive instances, when instances only occur under a single hypothesis (during rule learning or categorization, see Navarro & Perfors, 2011; Oaksford & Chater, 1994; Thomas et al., 2008), or when hypotheses are deterministic (when predicting sequences; see Austerweil & Griffiths, 2011). Although this explanation cannot explain away *all* cases of single-hypothesis inquiry, it does raise the intriguing question of whether these factors actually influence whether people generate alternative hypotheses. For example, Hendrickson, Navarro, and Perfors (2016) manipulated the number of positive instances of a hypothesis and found that participants behaved in a less confirmatory fashion as hypothesis size increased. Similarly, Oaksford et al. (1997) manipulated people’s beliefs about the frequency of features associated with a hypothesis in the Wason card selection task (for example, participants were told that p and q both occurred often). People were more likely to try to falsify the rule when both features were common. These findings highlight how a learner’s prior beliefs about the structure of their environment impacts the hypotheses they generate, and the kinds of evidence they seek to test them (see also Coenen, Bramley, Ruggeri, & Gureckis, 2017).

Zero-hypotheses queries. The assumption that people make queries to test specific hypotheses is central to OED models of cognition. Yet

many reasonable questions do not require any hypotheses at all. For example, upon visiting a city for the first time, you may ask your local friend “Where’s a good place to eat?”. This is an incredibly common kind of query that does not require considering any hypotheses beforehand. Another example of zero-hypothesis information gathering occurs in early childhood, when children exhibit unstructured, exploratory play (e.g., Hoch, Rachwani, & Adolph, in review). Although uncertainty about many aspects of their environment is presumably high, it is difficult to imagine that young children always represent hypotheses about what might happen as a consequence of their information seeking behaviors. These examples raise the question of how it is possible for a learner to quantify their uncertainty or notice a knowledge gap without hypotheses. We provide an in-depth discussion of constraints on zero-hypothesis queries in the next section that addresses how people generate questions in the first place.

Summary. A critical challenge for OED models is to explain the set of hypotheses that the learner considers. Although there is some recent work exploring how people reason with subsets of hypotheses, core psychological principles guiding this process have remained elusive and choices are sometimes made after experimental data have been collected. In addition, the OED framework does not easily apply to situations where learners 1.) consider the wrong hypotheses for the task, 2.) consider only one hypothesis, or 3.) do not consider hypotheses at all. These are not insurmountable challenges to the OED research program, especially in light of recent ideas about adaptive hypothesis sampling or online hypothesis space construction (Christie & Genter, 2010). However, these issues are critical to establishing the broader utility of the OED approach, outside of simple experimental tasks.

Question 2: How do people generate a set of candidate queries?

In standard use, an OED modeler computes the utility or informativeness of each possible query available in the task and then asks if people select the best option. For example, this could be which cards to turn over in the Wason selection task (see above) or where to fixate one’s eyes in a visual search task. However, what comprises the set of possible queries, $\{Q\} = Q_1, Q_2, \dots$, that are available in any situation?

$$\begin{array}{c} \text{What} \\ \text{is in} \\ \{Q\}? \\ \downarrow \\ EU(\boxed{Q}) = \sum_{a \in Q} P(a)U(a) \end{array}$$

Consider, for instance, a young child asking a parent, “Can ducks fly?” Perhaps this is an informative question for the current situation, but there seems no limit to the number of questions that could be asked (e.g., “Do ducks sleep?”, “How many babies do ducks have?”, “Is the weight of a duck in kilograms less than ten times the square root of seven?”), even though only a subset might be relevant for any particular inferential goal. In order for OED principles to apply to this fairly typical situation, *every* possible question or query would need to be evaluated and compared to others. OED models currently provide no guidance on this process, ignoring almost completely how the set of questions or queries is constructed.

For OED to be applied to more general types of inquiry (such as asking questions using language), the framework must be able to deal with the wide range of human questions. As we will argue below, the existing OED literature has tended to focus on relatively simple inquiry behaviors (e.g., turning over cards in a game, asking the category label of an object), which are more amenable to mathematical analysis. However, once one considers modeling the rich and sophisticated set of questions people can ask using natural language, com-

putational issues become a significant challenge. Although this section focuses on question asking in natural language, the concern is not limited to the language domain. For example, interacting with a complex system (like the physical world) often requires us to construct novel actions or interventions (Bramley, Gerstenberg, & Tenenbaum, 2016) from a potentially unbounded space. When playing a new video game, for instance, a person might initially perform a wide range of complex actions to understand the game dynamics and physics. Each action sequence reveals information about the underlying system’s rules but is selected from a potentially large space of possible action sequences.

Searching for the right question. Many researchers have had the experience of sitting through the question portion of a talk and hearing a very clever question asked by an attendee. Often, we would not think to ask that question ourselves, but we immediately recognize it as informative and insightful. While in some cases we might attribute this to differences in knowledge (e.g., perhaps a colleague thinks about an analysis in a slightly different way) it also seems clear that coming up with a question is often a significant intellectual puzzle (Miyake & Norman, 1979).

Consider a recent study by Rothe et al. (2016). In Experiment 1 of the paper, participants played a game where they had to discover the shape and configuration of a set of hidden ships on a gameboard (similar to the children’s game Battleship). Rather than playing an entire game, participants were presented with partially uncovered gameboards (i.e., some of the tiles were uncovered, see Figure 2) and then were given the opportunity to ask questions in natural language, which would be helpful for learning the true configuration of this gameboard (the only limitations were that questions had to be answerable with a single word and that multiple questions could not be combined). Example questions are “Where is the upper right corner of the blue object?”, “How

long is the yellow object?”, or “How many tiles are not occupied by ships?”. Interestingly, while participants generated a wide variety of different questions, they rarely came up with questions that came even close to the highest expected information gain (EIG). This is somewhat surprising, because one assumption of the OED framework is that people will ask the most informative question in a given context. Given the simple setup of the task, this should be the same question for each participant in this game. Yet few subjects asked really clever and revealing questions. The modal participant asked much more mundane and only moderately informative questions.

Interestingly, although participants were not good at devising useful questions, they were highly skilled at recognizing good questions. In a follow-up experiment, participants were presented with the same set of ambiguous game situations and a list of potential questions derived from the questions asked in the previous study. Here people’s selections closely mirrored the predictions of OED models with people preferring the objectively most informative questions.

The Rothe et al. (2016) study highlights how the demands of generating questions “from scratch” may limit optimal information-seeking behavior. In general, this work helps to clarify the distinction between question *generation* and question *evaluation* (the latter being the primary emphasis of contemporary OED approaches). One future research topic raised by this work is how people formulate questions in a given context and how they search the large space of possible questions. While presently underexplored, these topics have natural solutions in computational or algorithmic approaches. For example, Rothe et al. (in prep) develop a question generating model that creates the semantic equivalents to human questions using a context-free grammar. This approach defines the basic semantic primitives of questions and rules for the composition of these primitives, then uses OED models as an objective function of a search process that explores the space of expres-

Figure 2. Top: Example of the Battleship game. Hidden gameboards are created by randomly selecting ships of different sizes and orientations and placing them in a grid at random, non-overlapping locations. A context is defined as a partially unveiled gameboard (center). The goal of the learner is to identify the true gameboard by asking questions. Bottom: Task sequence from Rothe, Lake, & Gureckis (2016). Subjects first turned over individual tiles one by one following instructions of experimenter (clicking on the ?). Next they indicate the possible ship locations. Finally people asked whatever question they wanted in natural language in order to best discover the underlying gameboard.

sions within the grammar to find the optimal question.

An alternative approach would be to construct questions “bottom-up” from the current situation. For example, questions could be constructed around hypothetical propositions or facts that might hold in a given situation (e.g., the size of the red ship *could* be six) but that are currently unknown. In any event, increased emphasis on question-generation is likely to open up new avenues for empirical research and process-level models. In some cases, it might also help expand the range of situations that are addressable within

the OED framework. For example, question asking behavior has long been of interest to educators (Graesser, Langston, & Bagget, 1993), and models that apply to more complex and realistic types of inquiry behaviors might have greater impact.

A mosaic of question types. The question of how to apply OED principles to more open-ended natural language question asking exposes more than just the issue of how this large space can be searched. Once one allows for broader sets of questions additional computational complexities are often encountered. Our intention here is not

to provide an exhaustive taxonomy of different question types (placing questions or queries into categories may not be particularly meaningful), but to compare and contrast a few different types of queries to illustrate the computational issues at stake.

Label queries. As noted above, most information search studies give people the option of choosing from a set of narrowly defined query types. In categorization experiments, for instance, participants can typically only inquire about the category membership of unlabeled items (MacDonald & Frank, 2016; Markant & Gureckis, 2014; Markant et al., 2015). During spatial search the task is usually to query specific locations and learn what they contain (Gureckis & Markant, 2009; Markant & Gureckis, 2012; Najemnik & Geisler, 2005).

In machine learning, these types of queries are called “label queries” and, similar to psychological experiments, they constitute a large part of active machine learning research (Settles, 2010). During a label query, an oracle (knowledgeable human) is asked to produce the label or class of an unlabeled instance, which helps a classification algorithm learn over time (roughly, “What is the name of this?”). An appealing feature of label queries is that they intuitively match some real world question asking scenarios. For example, children often learn by pointing out objects in the environment and having an adult label them. Vocabulary learning of a foreign language has a similar property.

The computational evaluation of label queries in an OED framework is relatively simple, assuming the learner has a well-defined hypothesis space (see Question 1 for why that might not be the case). For example, when encountering an animal that is either a cat or a dog, a child might point at it and ask “what is that?” Knowing that there are two possible answers (“cat” or “dog”), it is relatively easy to compute the sum in Equation 1 (see Question 4).

Feature queries. Instead of requesting labels or examples, learners can also ask about the importance of entire features or dimensions of cat-

egories. For example, a naive learner might ask whether the ability to sing is an important feature for telling whether something is a bird. Unlike label queries, this type of question does not request the class membership of a single exemplar, but instead asks more generic information about the class. Such *feature queries* have proven to be successful in machine learning, in particular when human oracles are experts in a domain and can quickly help improve a classifier’s feature weights to accelerate learning (Raghavan, Madani, & Jones, 2006).

The distinction between item and feature queries holds psychological significance as well. For example, a growing literature in developmental psychology (see Question 9) explores information-gathering strategies in simple games such as “Guess Who?” or “20-questions”. When used as an experiment, the subject tries to identify a hidden object by asking a series of yes/no questions. There are two broad strategies commonly used by human participants in the game: *hypothesis-scanning* questions target a specific instance (e.g., “Is it Bill?”), whereas *constraint-seeking* questions ask about features that are present or absent across multiple objects (e.g., “Is the person wearing a hat?”). A classic finding in this literature is that younger children (aged 6) tend to ask more hypothesis-scanning questions, while older children (aged 11) and adults use more constraint-seeking questions (Mosher & Hornsby, 1966).

From a computational perspective, the informational value of some feature queries is easy to compute (e.g., the constraint-seeking feature questions in the “Guess Who?” game) and researchers have used OED models as a yardstick for human performance (Kachergis, Rhodes, & Gureckis, 2016; J. D. Nelson et al., 2014; Ruggeri & Lombrozo, 2015). A more difficult problem arises when questioners do not yet know what the relevant features might be. For example, I might ask my friend who works in the tech industry what features are relevant for predicting the survival of a new startup.

This question would help me narrow down the set of features that I might then proceed to ask more targeted questions about.

This issue is widely recognized in applied Machine Learning as the problem of “Feature Engineering” (Blum & Langley, 1997). When building a model in a new domain, the modeler first needs to figure out which features to use (or to build from other features or from raw data). This process often relies on human input from experts with domain knowledge, and it precedes the actual learning phase of the model. It is thus difficult to compute the informational value of this kind of feature query in a way that makes it comparable to other types of queries, even though it undoubtedly serves an important purpose when starting inquiry in a new domain.

Demonstration queries. Consider learning a complex skill like how to play “Chopsticks” on the piano. A skill is essentially a category under which some actions count as performing the skill and others do not. Taking a label query approach, the learner would play a random sequence of notes and then ask the teacher or oracle “Is that ‘Chopsticks’?”, eventually learning how to perform the piece. An alternative strategy would be to request an example of a category (e.g., “What is an example performance of ‘Chopsticks’?”). This type of *active class selection* or *demonstration query* provides a positive example of the category, which can be highly informative, especially early in learning (Lomasky, Brodley, Aernecke, Walt, & Friedl, 2007). For example, one might want to ask to see a good or typical example of a category member (“What does a typical bird look like?”) before making queries about new exemplars or specific features. Similarly during causal structure discovery, one can often learn a lot about a system by seeing a demonstration of how it works before making a targeted intervention. The idea of demonstration queries has been considered for teaching motor skills to social robots, who can ask a human to demonstrate a full movement trajectory rather than providing feedback about the robot’s own at-

tempts at a task (Cakmak & Thomaz, 2012). In humans, demonstration queries are particularly useful for learning new skills. Importantly, the usefulness of demonstration queries depends on the level of knowledge or expertise of the answerer, which means that they should be chosen more or less often based on the learner’s beliefs about the answerer. This is a topic we discuss in more detail in Question 6.

Demonstration queries are computationally complex. As noted above, OED models average across all potential answers to a question, but a question like “What does a cat look like?” could be answered by providing any reasonable example of the category (a cat photo, pointing at a cat, a drawing of a cat). For complex hypotheses or categories it does not seem possible for the naive question asker to simulate this set of potential answers via explicit pre-posterior analysis. It is thus hard to imagine how the OED framework could provide a satisfactory explanation of how people assess the usefulness of demonstration queries (“What does a cat look like?”), compared to, for example, label queries (“Is this a cat?”). Explaining how people choose demonstration queries, and when people deem a demonstration query to be more helpful than other queries, will likely require an understanding of people’s metareasoning about query-type selection.

The role of prior knowledge in question generation

So far, this section has highlighted the problem of modeling question types, and generating questions to serve particular goals of a learner. However, there exists a more fundamental puzzle about the way certain questions are generated. Consider the following examples.

- What’s the English translation of the German term “Treppenwitz”?
- Where do raccoons sleep?
- What makes that object float in mid-air?

- Why is that person wearing only purple?
- What do you do for a living?

What these examples have in common is not that they expect particular types of answers (they could ask for features, events, labels, mechanisms, etc.), but that can be asked in the absence of any concrete hypotheses and may be triggered by context and prior knowledge alone. For a non German speaker coming across the term “Treppenwitz” it not necessary to actually consider particular English translations. Simply knowing that most words or phrases can be translated between German and English is sufficient to know that there is information to be gained. Instead of concrete hypotheses, such questions can be generated if the questioner realizes that there exists some currently unknown fact that is knowable in principle. Since the number of unknown facts is infinite, there must be some way of constraining the questions to those that address specific “knowledge gaps” that can realistically be closed. To frame this puzzle in another way, consider how an artificial agent would have to be programmed to generate these questions in an appropriate situation. Perhaps asking for a translation of an unknown phrase would be the easiest to implement if the agent’s goal is to parse and translate sentences. But we are currently still very far away from developing artificial intelligence that spontaneously asks about raccoons’ sleeping places or questions people’s odd clothing choices in the same way a human might do on a walk through the forest or a stroll through the city.

We propose that the structure and content of current knowledge alone can act as a strong constraint on query generation in the absence of hypotheses. Abstract knowledge in the form of categories, schemata, or scripts, can play an important role in highlighting knowledge gaps (e.g., Bartlett & Burt, 1933; Mandler, 2014; Minsky, 1974). Knowing that raccoons are mammals, and that a broadly shared feature of members of the mammal category is the need to sleep, can help us identify a gap in our knowledge about raccoons. (In fact,

this seems to be a common question. When the authors typed “where do raccoons” into a well-known search engine, “sleep?” was among the top suggested completions of the query.) Conversely, most people would be much less likely to spontaneously generate the question “where do raccoons get their nails done?”, because we have no prior knowledge to suggest that there even exists an answer to this question. Asking about the motivation behind a person’s odd clothing choices similarly requires prior knowledge. At the very least, one has to know that people generally act based on goals and intentions, and that an all-purple wardrobe is an unusual choice. Conventions or conversational scripts are another source of queries. For example, we learn that it is typical to ask for someone’s name, place of residence, or profession upon first meeting them. It is much less common to ask about a person’s preferred sleeping position, which might be similarly unknown, but is not part of the conventions that apply to small talk. Conventional sets of questions exists in many domains, which makes the task of generating questions much easier.

What types of knowledge constrain these types of queries? While some of them, for example social conventions, undoubtedly have to be learned, others may be more fundamental. Foundational knowledge, sometimes referred to as core knowledge (Carey & Spelke, 1996; Spelke & Kinzler, 2007), may constrain query generation already in early childhood, when specific world knowledge is still sparse. For example, we know that infants are endowed with a system of object representation involving spatio-temporal principles (such as cohesion, continuity, and support, Spelke & Kinzler, 2007). Furthermore, children as young as 2 years old make relatively sophisticated assumptions about causal relationships between objects (Gopnik et al., 2004). Such early knowledge can be leveraged to help find opportunities for inquiry. For example, it has been shown that children engage in more information seeking behaviors when their prior expectations about causal re-

lationships or spatio-temporal principles are violated than when they are confirmed (Legare, 2012; Stahl & Feigenson, 2015). Upon seeing an object suspended in mid-air, children might therefore proceed to seek further information to explain the now apparent knowledge gap about how the object is supported (Stahl & Feigenson, 2015). Another kind of core knowledge that emerges early in life is the ability to represent animate beings as intentional agents (Spelke & Kinzler, 2007). Young children expect people, but not objects, to execute actions based on goals and plans (Meltzoff, 1995; Woodward, 1998). This means that, similar to adults, young children observing a person behave in an intentional but strange manner might become aware of a knowledge gap and try to find out what goals or intentions could explain this behavior.

Summary. Current models of inquiry assume that questions are generated to satisfy a particular set of inferential goals, by testing specific hypotheses about the world. However, the examples above illustrate the wide variety of questions that arise from knowledge gaps identified and are formulated in other ways. Given how common such questions are, future work on inquiry should devote more attention to query generation that goes beyond the hypothesis-testing framework. Knowledge-based queries also raise an entirely new set of computational challenges. Accounting for these questions will often require models of domain knowledge, structured representations, and fundamental beliefs about causality and intentionality. Most interesting is that these types of queries seem to fall outside the domain of OED models, as formulated to date, in that no alternative hypotheses need be considered, and the set of answers may not be explicitly enumerated.

Question 3: What makes a “good” answer?

In the OED framework, a question’s expected value is a weighted average of the value of each of its possible answers (Equation 1). In this sense, the value of answers is a more basic concept than the expected value of a question. However, what

makes an answer to a query “good”? More formally:

$$EU(Q) = \sum_{a \in Q} P(a) \begin{array}{c} \text{What} \\ \text{is a} \\ \text{good} \\ \text{an-} \\ \text{swer?} \\ \downarrow \\ U(a) \end{array} .$$

The importance of this issue is reflected in a variety of scientific literatures. For example, psychologists and philosophers have discussed what counts as a good “explanation” of a phenomenon. Although there are differences in people’s preference for certain explanation types (e.g., the teleological or ontological distinction, Kelemen & Rosset, 2009; Lombrozo & Carey, 2006), this work does not usually involve computationally precise ways of evaluating the quality of answers (or explanations). Despite its foundational nature, in the OED framework there is very little research on how people evaluate an answer’s usefulness (but see, Rusconi, Marelli, D’Addario, Russo, & Cherubini, 2014).

To develop an initial intuition for answer quality, consider the following example dialogs. If a learner asks someone “Where exactly do you live?”, an answer including exact Global Positioning System (GPS) coordinates completely answers the question and removes any lingering doubt. In contrast, a more imprecise answer like “New York City” might leave uncertainty absent any other information. The point is that our intuition is that some answers (like the GPS) are better than others because they are more informative. The quality of an answer also depends on what the question asker already knows. Consider the following exchange:

Q: What city were you born in?

A: New York City

Q: Do you live in the same city you were born?

A: Yes

Q: Which city do you live in?

A: New York City

Here, the final question-answer pair is identical to the one above but now the answer contains no new information.

These examples highlight a few key points about answers. A good answer is relevant to the given query and adds information above and beyond what is already known by the learner. Answers differ in quality based on the amount of information they provide, but it is possible for two answers to be equally good if they offer the same query-specific information (that is, it does not matter if one answer provides additional information that was not called for by the query). A major topic of research within the OED framework is determining a general-purpose, mathematically rigorous way of defining the quality of an answer to a question. The most common approach is to assume there is a type of utility associated with answers. In the remainder of this section, we will give a more detailed account of specific utility measures that OED models have used to quantify the quality of answers.

Determining the utility of answers. In the broadest sense, it is useful to distinguish between *informational* (or *disinterested*) and *situation-specific* (or *interested*) utility functions (Chater, Crocker, & Pickering, 1998; Markant & Gureckis, 2012; Meder & Nelson, 2012). Pure information utility functions are based solely on probabilities and on how answers change probabilities. Situation-specific functions take into account that learners collect information for a specific purpose beyond pure knowledge gain (e.g., saving time, money, or cognitive resources¹). Both approaches reflect hypotheses about the overall goals and purpose of human inquiry, although the difference between them is not always clearly acknowledged in psychological literature.

Informational utility functions. Most OED models evaluate answers according to *how they change a learner's beliefs about possible hypotheses*. These metrics are a thus a function of the learner's prior belief before receiving an answer,

Table 2

Distributions showing a hypothetical learner's subjective belief about a parameter, θ , before and after learning a new piece of information. The columns show the predictions of Information Gain (IG), Kullback-Leibler divergence (KL), and probability gain (PG). A (+) indicates a positive utility, (o) indicates zero utility and (-) indicates a negative utility. The different approaches frequently disagree.

Situation		IG	KL	PG
1		+	+	+
	+	+	+
	o	o	o
	o	+	o
	-	+	-
	+	+	o
θ				

¹Note that we use the term *situation-specific* to include both external and internal costs

$P(H)$, and their posterior belief having received that answer to a their question $P(H|Q = a)$, which in our shorthand notation can be written as $P(H|a)$. (Recall that a denotes a particular individual answer to a particular question Q .) Information Gain, from Equations (1) and (2), is one of the most popular functions used within psychology, but there exist a number of interesting alternatives, including impact (expected absolute belief change), diagnosticity, KL divergence, and probability gain (J. D. Nelson, 2005).

The differences between these measures may sometimes seem subtle, but comparing them more carefully raises interesting and fundamental questions. Consider the six scenarios depicted in Table 2. Each scenario shows how the distribution of a learner's belief about some parameter θ changes as a result of an answer to their query. The three rightmost columns show how three different utility measures evaluate the usefulness of this change in belief. (To keep things simple, we just focus on the *sign* of the model outputs. For in-depth comparison, see J. D. Nelson, 2005). The models are Information Gain (IG), Probability Gain (PG), and Kullback-Leibler divergence (KL). An answer's probability gain is the reduction in probability of incorrect guess that the answer provides. Interestingly, it can be obtained by replacing Shannon Entropy, $ent(H)$, in Equations (1) and (2) with $pFalse(H)$, also known as Bayes's error:

$$pFalse(H) = 1 - \max_{h \in H} P(h). \quad (6)$$

Kullback-Leibler (KL) divergence is an alternative information-theoretic measure to Shannon entropy, which is useful in comparing two distributions (in this case, a posterior and prior). When evaluating the expected usefulness of a question, that is, $EU(Q)$, KL divergence and Expected Information Gain (EIG) give exactly the same value, for every possible question, in every possible scenario (Oaksford & Chater, 1996). However, KL divergence and IG can make contradictory predictions when the usefulness of different specific answers, i.e., $U(a)$, is evaluated, as the examples in

Table 2 demonstrate. The KL-divergence resulting from an answer a given question Q is

$$U_{KL}(a) = \sum_{h \in H} P(h|a) \log \frac{P(h|a)}{P(h)}. \quad (7)$$

For the first two situations in Table 2, the three models agree that the answers' values are positive. In both cases, the variance of the posterior has narrowed, implying that the learner is now more confident in the estimate of θ . Likewise in the third example, all models assign zero value to an answer that has not changed a learner's beliefs at all. This scenario captures any situation in which a learner is told something they already know or that is irrelevant for answering their question.

Examples four to six show more divergent cases. In scenario four, a learner changes their belief about the value of θ but does not narrow their posterior. For example, imagine learning that your friend's car is in fact a Toyota, not a Chevrolet as you previously assumed. This would change your estimate of the car's costs without necessarily affecting your uncertainty around the precise value. The IG in this example is zero, since uncertainty does not change. The same holds for the probability of making a correct guess (PG), since the probability of the most likely hypothesis has stayed the same. This assessment runs counter to some intuitive definitions of what constitutes a good answer, since the learner has in fact *changed* their belief quite substantially. A measure like KL-divergence, which assigns positive value to this scenario, thus may be more in line with these intuitions.

Scenario 5 is even more puzzling. Here, the learner receives an answer that *increases* their uncertainty. This leads to negative IG and PG, although KL divergence is positive. (In fact, KL divergence is always positive unless prior and posterior are exactly the same.) Returning to the car example, this could happen upon learning that a friend's car is either a Toyota or a Ford, having previously assumed that it was probably a Chevrolet. Now, you might end up being more uncertain about its cost than before. Again, this con-

clusion is somewhat at odds with the intuition that *something* was learned in this scenario even if the learner ended up more uncertain as a consequence.

Finally, scenario 6 shows that sometimes IG and PG make diverging predictions. Here, the learner has narrowed their posterior around the smaller peak, and has therefore reduced their overall uncertainty. However, the probability of the most likely hypothesis has stayed the same; thus, the answer has no value in terms of PG. As an example, imagine that you are trying to guess the breed of your friend's dog and you are pretty sure it is a German shepherd. Finding out that it is not a chihuahua because your friend is allergic to chihuahuas might slightly change your hypothesis space about much less likely possibilities (and therefore lead to positive IG), but not at affect your high confidence regarding your top hypothesis (hence zero PG).

These examples demonstrate that assigning values to answers, even from a completely disinterested perspective (i.e. when one is only concerned with quantifying belief change), is not at all trivial. These examples raise some interesting psychological questions, such as how people treat answers with negative IG, or how they balance information and the probability of making a correct choice. An important area for future research will be to consider information gain based on other types of entropy metrics, and not only based on Shannon entropy. For instance, Crupi and Tentori (2014) discuss information gain based on quadratic (rather than Shannon) entropy. In fact, in mathematics, physics, and other domains, there are many different entropy models, several of which could be important in a descriptive theory of human behavior (Crupi, Nelson, Meder, Cevolani, & Tentori, in press). We will briefly return to these questions below, after discussing situation-specific utility functions.

Situation-Specific utility functions. According to situation-specific (“interested”) theories of information search, the utility of an answer (and therefore a query) depends on concrete goals of the learner, irrespective of or in addition to the goal

of increasing information. Question-asking strategies that are based on situation-specific goals can yield strongly different predictions than disinterested models (Meder & Nelson, 2012). For example, consider a categorization task in which payoffs are asymmetric, such that correctly or incorrectly classifying items into different categories yields different costs or penalties. This could be the case during medical diagnosis, where there might be greater costs associated with misclassifying a potentially fatal condition than a benign one, which leads to asymmetrical decision thresholds for treatment (lower for the fatal condition). This asymmetry should also affect the medical tests that are administered. Tests that have the potential to change the treatment decision are more valuable than those that do not, irrespective of their pure informational value (Pauker & Kassirer, 1980). Cost-sensitive components also matter when learners have some pure information goals (e.g., to minimize Shannon entropy across possible hypotheses) but wish to simultaneously minimize time spent, cognitive demands, or number of queries made.

Interestingly, people are not always sensitive to costs of questions (Baron & Hershey, 1988) and tend to make queries in line with pure information strategies, like probability gain or information gain on some tasks (Markant et al., 2015; Meder & Nelson, 2012). An interesting question for future work is to understand when and why this might be the case. A preference for disinterested search may be adaptive, for instance, if people expect to re-use information later on in a different task. This could be investigated by manipulating people's beliefs about future re-usability to see how the use of disinterested versus interested question asking strategies changes. It is also possible that it is computationally intractable in some cases to assess situation-specific utilities. For example, Gureckis and Markant (2009) explored how even for a simple task this can require not only computing the utility of each individual answer, but also how information from that answer might influence a future decision-making policy. Computing this can

be a significant computational burden. Finally, sometimes people only realize what the value of an answer is when they actually see it and process it. This would suggest that people might have to learn to adjust their inquiry strategy as they learn more about a given situation-specific utility function. This possibility calls for experiments that have people assess the value of both questions and of answers, in tandem, to test how the latter influences the former (also Question 5, on learning from answers).

Summary. Determining the value of an answer is no easy feat. Even when learners have a good probabilistic model of the task at hand, there are many different approaches to measure the utility of answers, many of which have some degree of plausibility. The lack of consensus on the ‘right’ kind of answer utility poses an interesting challenge for OED models, all of which define a question’s expected usefulness as the probability-weighted average of its possible answers’ individual usefulness values. To address this challenge, we see several possible strategies.

First, there are a number of efforts to try to isolate domain-general principles of assigning values to answers. Using carefully designed experiments, this approach might ultimately reveal that some functions are simply a better match for human intuitions about answer utilities than others. One example is work by J. D. Nelson et al. (2010) that found that expected probability gain was the best-fitting psychological model among several candidates, including EIG, for information search behavior in a probabilistic classification task. Future studies will be required to explore more systematically if this finding holds in other domains as well.

Second, if no domain-general information metric can be found, then modeling inquiry in a new domain or task will require an understanding of how people assign value to received answers in that domain. Since this is such a fundamental building block of any OED model, it might be sensible to study the value of answers in isolation, before trying to build models of the expected useful-

ness of questions.

Question 4: How do people generate and weight possible answers to their queries?

OED models define the expected usefulness of a question as a probability-weighted average of the usefulness of each possible answer (Equation 1):

$$EU(Q) = \sum_{a \in Q} P(a) U(a)$$

How do
people
weight
answers?

We have just discussed the problem of evaluating the utility of individual answers. An entirely different question is which answers people anticipate to begin with and what probabilities are assigned to them. For example, if you ask someone “Which city do you live in?”, an OED model requires you to consider each possible answer (“New York”, “Boston”, “Austin”, “Denver”, etc...) and to weight the utility of that answer by the probability of receiving it. If you know nothing else about an individual, the probabilities might be the base rates from the general population. However, if you meet a new colleague whom you know is a professor, cities with universities or colleges might be more probable. Importantly, the above equation assumes that the learner knows the possible answers a question might receive, and the probability of each of those answers. In real-world tasks, as well as in more complex experimental tasks, such as models of eye movements in visual search or of causal learning, models based on the OED framework must make a number of usually implicit assumptions about these variables.

What is a possible answer?. The OED framework treats question asking as following from the goal of obtaining information about something. As a psychological model, OED presumes that people know the possible answers and their probabilities. Returning to an earlier example, if someone asks

“Where do raccoons sleep?” it seems nonsensical that the answer would be “blue,” improbable that the answer is “underwater,” and likely that the answer is “in a den”.

Surprisingly little research in psychology has attempted to understand what people expect as answers to different types of questions. Given the tight coupling between answers and questions implied by the OED framework, this could be a fertile research topic. For example, how do differences in how readily people consider different answers affect information seeking-behaviors? Some questions have rather obvious or simple answer spaces (e.g., a true/false question returns either of two answers). In addition, in some cases the possible answers to a question are basically the same as the hypothesis space. For example, for the question “What city do you live in?”, the possible hypotheses are cities, as are the answers. This suggests that issues about hypothesis generation discussed in Question 1 might hold relevance. The space of answers that people consider possible might strongly influence the value they assign to a question. Furthermore, the type of learning that happens after receiving an unexpected versus expected answer might be somewhat different (see Question 5). Despite the theoretical importance of these issues to the OED hypothesis, little research has addressed them.

Dealing with intractable answer spaces

As noted throughout this article, theories of inquiry based on OED principles share much in common with theories of decision making. This is particularly clear given that the value of a question depends on a “tree” of possible future outcomes similar to how in sequential choice theories the value of an action depends on a “tree” of later actions (see Figure 3). However, as many authors in the decision-making literature have noted, it is computationally intractable to consider all possible future outcomes or scenarios (e.g., Huys et al., 2012; Sutton & Barto, 1988). A variety of methods have been proposed to approximate this vast

Figure 3. Top: A typical decision tree. The value of the current choice is often assumed to depend on the outcomes and available choices at later points in the tree. Bottom: Structure of OED models showing how the value of a question similarly depends future states (i.e., answers to the question).

search space of outcomes, two of which we briefly summarize here.

Integration by Monte-Carlo sampling. The key to Monte-Carlo approximation (e.g., Guez, Silver, & Dayan, 2012) is the fact that the quality of a question is basically a weighted sum or integral (i.e., Equation 1). One way to approximate this integral is to sum over a set of samples:

$$EU(q) = \sum_{a \in A} P(a)U(a) \approx \frac{1}{m} \sum_{\ell=1}^m U(a^{(\ell)}) \quad (8)$$

where $a^{(1)}, \dots, a^{(m)}$ are a set of m samples from the $P(a)$ distribution. In the limit as $m \rightarrow \infty$, the approximation based on samples converges to the true value of EU found by weighting the value of each answer by its appropriate probability. Under the Monte-Carlo approach, people might repeatedly mentally simulate different answers they could receive and evaluate the utility of each.

Highly probable answers would be generated often whereas less probable answers might rarely be simulated. In the case where the number of answers is large, or where some answers are very unlikely this approximate sum may be more computationally efficient. In addition, when m is small certain biases might be introduced (e.g., rare answers are less likely to be sampled and thus less likely to enter into the evaluation of a question).

Integration by tree pruning. An alternative approach assumes explicit “tree pruning” where certain future paths of the decision tree are selectively ignored. For example, Huys et al. (2012) consider tree pruning in a sequential decision-making task. The basic idea is that rather than considering all possible paths of a decision tree unfolding from a particular choice (e.g., Figure 3, top), an agent might selectively drop certain paths. In the Huys et al. setting this including pruning sequential paths that likely lead to particular types of outcomes (e.g., punishment). An analogous strategy in the OED setting might mean removing from consideration answers for which $P(a)$ falls below some threshold. While such ideas have yet to be tested in the inquiry literature, certain heuristic strategies should bias choices in specific ways. For example, it may be possible to experimentally detect a tendency to discard low probability answers with high information utility.

Integration by generalized means. A final approach considers alternative ways of computing $P(a)$, and the possibility of averaging some function of answer utility values, rather than the raw answer utility values themselves. The General Theory of Means (Muliere & Parmigiani, 1993) provides a general mathematical framework. One extension of Equation 1 is to use answer weights that are nonnegative and sum to 1, but which do not necessarily correspond to answer probabilities:

$$EU(q) = \sum_{a \in A} w(a)U(a) \quad (9)$$

Defining expected utility in terms of answer weights, rather than answer probabilities, high-

lights that in the normative theoretical sense, there is a decision to make about what kind of weights to use (e.g., maximum entropy consistent with known constraints, or a minimax strategy, etc).

The basic constraint in the General Theory of Means framework is that the weights should be nonnegative and should sum to 1. For example, if the probability of some answers is well-understood, but the probability of other answers is not known, people might assign higher weight to answers with less-well-understood probabilities, other things being equal. The important points, theoretically, are: (1) from a normative standpoint, we seldom really know the answer probabilities and (2) from a descriptive standpoint, although answer weighting is central to OED models, we still lack a good understanding how people actually evaluate answer utilities.

Summary. The OED framework defines the value of a question as the probability-weighted average of the value of its individual answers. We have reason to suspect that this is not the full story, given that the probability of individual answers is not always knowable, that it is combinatorially difficult or impossible to integrate all possible answers in some circumstances, and that various heuristic strategies might be simpler. Proposals from the decision-making literature suggest some computationally feasible strategies to handle the combinatorics of evaluating all possible answers’ usefulness. Assessing how people weight individual answers is ripe for future research, as alternate proposals can be well-specified, and there has been virtually no research in this area to date.

Question 5: How does learning from answers affect query selection?

Like a scientist who considers what they could learn from possible outcomes of their experiments, an optimal question asker anticipates how an answer would *change their current beliefs*. For example, computing the expected new Shannon entropy in the EIG model entirely relies on the degree of belief change:

How
do
people
change
be-
liefs?

$$ent(H|a) = - \sum_{h \in H} P(h|a) \log P(h|a)$$

This aspect of question evaluation is a key idea behind the concept of preposterior analysis (Raiffa & Schlaifer, 1961) and lies at the heart of the OED approach.

Leaving aside the computational challenges of simulating all possible answers (see previous section), how people update their beliefs based on new data is one of the most fundamental (and contentious) questions in many areas of higher-level cognition, including language acquisition, categorization, stochastic learning, and judgments under uncertainty (e.g., Tenenbaum, Griffiths, & Kemp, 2006). Findings from this longstanding line of work can inform the study of inquiry in a number of ways, two of which will be discussed below. First, we will discuss how deviations from OED norms during inquiry can emerge from particular violations of inference norms. Second, we will show that inductive inference strategies are often heavily influenced by the current context and the identity and intentions of the person providing the information. Since a vast number of inquiry scenarios are embedded in some form of social or educational context, understanding this pragmatic aspect of inference is pivotal for a complete account of question-asking.

Inductive norm violations. There are many ways in which people deviate from (Bayesian) inductive principles when integrating new evidence with prior knowledge. Consider the following well-known examples.

- It has been shown that in some situations people exhibit what is often called *base-rate neglect* (Doherty, Mynatt, Tweney, & Schiavo, 1979; Kahneman & Tversky, 1973).

Base rate neglect is the tendency to evaluate the posterior probability of a hypothesis, $P(h|e)$, mostly based on its ability to account for the new evidence, $P(e|h)$, while largely ignoring its prior probability, $P(h)$.

- When evidence is presented sequentially, people often reveal the opposite phenomenon. That is, they assign too much weight on their initial beliefs and behave conservatively when updating these beliefs in light of new evidence (Edwards, 1968; Phillips & Edwards, 1966).
- In other tasks, it has been shown that people exhibit a positivity bias. That is, they assign more weight to positive evidence (e.g., learning that something is true) compared to negative evidence (learning that something is false), even when both types of evidence are equally diagnostic (Hodgins & Zuckerman, 1993; Klayman, 1995).

There is ongoing debate on whether these phenomena count as biases and whether they can be explained based on people's task-specific beliefs or preferences (Griffiths & Tenenbaum, 2006; Kahneman et al., 1982; Krynski & Tenenbaum, 2007). What's important for the present discussion is that they can have significant impact on the expected information value of possible questions. For example, base-rate neglect could lead people to ask questions about hypotheses that can be tested easily, even if the hypothesis in question is unlikely a priori. Among other things, this could lead to an unwarranted preference for medical tests with a high hit-rate, even if they produce many false positives (some authors would argue that frequent mammograms are an example for the tendency to seek such tests; see Elmore et al., 1998; Gigerenzer, Mata, & Frank, 2009). Conservatism during question asking could lead to a type of "question-asking myopia" whereby askers make a greater effort to test their initial hypotheses, instead of considering alternatives that appeared less likely

in the beginning but are supported by incoming data. This could explain the finding that people who are asked to state their hypotheses *early* during a mock police investigation were subsequently more biased in their information-seeking strategies than those who were not asked to do so (O'Brien & Ellsworth, 2006). (The former group not only showed higher confidence in their initial hypothesis, but also sought more evidence for it, irrespective of the alternatives.) Overweighting positive evidence could lead to a preference for questions that people expect to yield “yes” answers. This possibility in particular could provide another explanation for people’s use of a *positive testing strategy* (discussed above, see also, Klayman & Ha, 1989; Wason, 1960).

These examples show that deviations from optimal induction principles and violations of inquiry norms can be intimately intertwined. However, even though the relationship between them has been pointed out before (Klayman & Ha, 1989; Nickerson, 1998), it is rare for psychologists to consider the two in tandem (but see Coenen & Gureckis, 2015).

Pragmatic & Pedagogical Reasoning. Human inquiry does not take place in a vacuum, nor are people’s questions typically directed at an anonymous oracle with unknown properties. Instead, many question-asking scenarios involve a social context that is shared between the questioner and the answerer. Furthermore, questioners usually have at least some expectations about the knowledge, beliefs, and intentions of answerers. This means that evaluating the usefulness of potential answers crucially depends on pragmatic and (in a teaching context) pedagogical considerations.

Shared context. Imagine that at the end of a meal your friend asks “Are you going to finish that?” Your interpretation and potential answer will be completely different if your friend is currently the host of a dinner party (they want to clear the table) or simply sharing a meal with you at a restaurant (they want to eat your food). It’s of course not a new insight that interpretations of lan-

guage depend on our understanding of the shared context between speaker and listener (Grice, 1975; Lewis, 1969). However, recent advances in probabilistic pragmatics have made it possible to formalize them as part of a Bayesian inference framework (Frank & Goodman, 2012; Goodman & Stuhlmüller, 2013; Goodman & Frank, 2016), which can be integrated with other probabilistic models, including OED models. To illustrate the main idea behind a probabilistic model of pragmatic interpretation, consider the example in Figure 4 from Goodman and Frank’s (2016) *Rational Speech Act* (RSA) model. Here, a speaker is referring to one of three friends and the listener has to infer which one. The listener does so by recursively simulating the speaker’s beliefs about their own beliefs, starting from a simplistic, “literal” (Lit) version of the listener who updates their beliefs about the world based on Bayes’ rule and a flat prior over referents. Based on this literal listener, the simulated speaker infers that the most informative way of pointing out the hat-wearing friend would have been to refer to the hat directly. Thus, the mention of glasses must refer to the hat-less friend with glasses.

A good demonstration of how this probabilistic pragmatic framework can be combined with OED comes from a study by Hawkins, Stuhlmüller, Degen, and Goodman (2015). They used the RSA model together with EIG to model people’s behavior in a guessing game. In this task, participants were assigned the roles of questioners and answerers. Questioners had the task of finding out the location of hidden objects (e.g., “find the poodle”) by directing questions at the answerers, who could see all of the objects (e.g., a poodle, a Dalmatian, a cat, etc.). Questioners were placed under a set of restrictions on the types of questions they could ask (e.g., must not ask about poodles, but may ask about Dalmatians or dogs) and answerers were equally aware of those restrictions. The study showed that questioners could come up with clever indirect questions (e.g., “where’s the

Figure 4. The RSA (rational speech act) framework models pragmatic reasoning as a recursive process. Figure adapted from Goodman & Frank (2016).

dog?”) that were correctly interpreted by the answerers who then gave helpful answers (revealing the location of the poodle, not the Dalmatian). The authors found that both questioners and answerers were better captured by the combination of an RSA model and EIG than by a “pure” EIG model that just used the literal meaning of both questions and answers. This finding demonstrates that when learners try to anticipate the likelihood of different answers, they also take into account the context or state of the world that is shared with their counterpart.

Features of the teacher. Another important factor that affects what we learn from our questions is the intention and expertise of the person providing the answer. For example, we would expect to receive different answers from a knowledgeable and helpful teacher (Tenenbaum, 1999) than from someone who is uninformed or ill-intentioned. This difference between learning in pedagogical and non-pedagogical situations has recently been explored computationally and experimentally (Shafto, Goodman, & Griffiths, 2014; Shafto, Goodman, & Frank, 2012), showing

that learners and teachers can make sophisticated inferences about each others’ minds in order to improve learners’ success. A good demonstration is learning from examples. In a teaching context, learners can usually expect examples to carry more information than just labels, since they expect teachers to choose particular examples that will help the learner generalize (Gweon, Tenenbaum, & Schulz, 2010; Tenenbaum & Griffiths, 2001; Xu & Tenenbaum, 2007). For example, teachers might provide prototypical examples of a category to allow the learner to pick up on the relevant features needed for future classification.

An important question for future research is how askers and answerers simulate the mental states of their counterpart and how many levels of recursive inference (“I think that they think that I think that they think, etc. ...”) are involved in this process. Recent work in probabilistic pragmatics has demonstrated individual variability in terms of levels of recursion (Franke & Degen, 2016). Given the evidence that even young children make pedagogical assumptions about teaching adults (Bonawitz et al., 2011; Kushnir, Wellman, & Gelman, 2008), another question concerns the developmental trajectory of these abilities and how world knowledge (what do people generally assume about one another in question asking scenarios?) and social reasoning (what are the intentions of this particular individual?) contribute and interact to shape the extremely sophisticated inferences that adults make about each other during inquiry.

Summary. Many OED models assume that learners anticipate how the answers to their queries will change their current beliefs. Here, we pointed out two important factors that may constrain this process and consequently affect how queries are chosen. First, given what we know about the plethora of inductive inference biases that people exhibit in other tasks, there is little reason to believe that anticipating future belief-change during inquiry should follow normative principles (Bayes’s Rule) in every respect. Thus, when there

is reason to believe that people *are* anticipating future belief change (like OED models suggest), one has to take into account how biases in this process would affect potential biases during query selection. Second, when answers are provided by other people, as it is often the case during inquiry, learners' inferences will be constrained by pragmatic and pedagogical considerations. Thus, to build realistic inquiry models, we need a better understanding of the psychological underpinnings of inferences in social contexts.

Question 6: How do cognitive constraints influence inquiry strategies?

Previous sections of this paper have pointed out that the OED framework, if interpreted in a mechanistic way, makes very ambitious computational demands that would indubitably exceed learners' memory and processing limitations. In earlier sections we discussed the idea that learners may sometimes restrict their hypothesis space, sample from their posterior beliefs, or approximate the aggregation of answer utilities into a question utility. These ideas fall largely within the OED framework in the sense that they represent cognitively plausible but statistically principled approximations. However, another possibility is that people use an entirely different set of strategies that are not curtailed versions of OED models to balance the trade-off between computation, accuracy and ease of processing (Simon, 1976).

One inquiry strategy that has received a lot of attention in educational psychology is the principle of *controlling variables* (CV). A CV strategy says that learners design experiments by changing one experimental variable at a time and holding everything else constant. Besides the benefit of yielding unconfounded evidence, this strategy is considered desirable because it is relatively easy to use and teach (Case, 1974; Z. Chen & Klahr, 1999), even though children do not often generate it spontaneously (Kuhn et al., 1995; Kuhn, Black, Keselman, & Kaplan, 2000). By focusing on only one variable at a time, it reduces the number of

items to be held in working memory and also creates easily interpretable evidence (Klahr, Fay, & Dunbar, 1993; Tschirgi, 1980). Although CV is often treated as a normative strategy (Inhelder & Piaget, 1958), its effectiveness in an OED sense actually depends on very specific features of the system of variables at hand. For example, when there are many variables but very few of them have any effect on the outcome, it can be much more efficient to manipulate multiple variables at once, assuming that testing for the occurrence of the outcome is costly. However, adults often still test variables in isolation, even when testing multiple variables is more efficient (Coenen et al., 2017). Empirically, these results may reflect the prominence of controlling variables in educational settings, or because people experience the CV strategy to effectively balance effectiveness and ease of use. The key point for the present purposes is that the CV strategy is not entirely equivalent to OED.

There are other ways in which people might trade off informativeness and computational tractability. Klayman and Ha (1987, 1989) found that participants often engage in a strategy they called *limit testing*. According to this approach, people restrict their hypothesis set to one focal hypothesis and seek confirmatory evidence for it. However, within that focal hypothesis people still test regions of higher uncertainty. For example, if a learner's focal hypothesis in a rule testing task was that "countries in South America" satisfy the rule, they might test this hypothesis by asking about South American countries at geographical extremes (e.g., Venezuela and Uruguay), to make sure that the true hypothesis is not in fact smaller than the current one (e.g., "countries in South America that are South of the Equator"). This strategy allows learners to refine their beliefs while still engaging in positive testing, which violates OED norms in many circumstances (see Introduction). Like a controlling variables strategy, limit testing thus does not count as an "optimal" strategy without significant additional assumptions (J. D. Nelson et al., 2001). However,

it might be a very reasonable approach given constraints of a learner’s ability to represent the full set of hypotheses.

Other examples include the idea that people can simply mentally compare two alternative hypotheses, look for places where they diverge, and then ask queries specifically about such diverging points. This process does not require enumerating all possible queries or answers, but it may be a reasonable heuristic in many cases. For example, when deciding between two hypotheses about the structure of a causal system, it is possible to choose which variables to manipulate by comparing the two structures and finding points where they differ (e.g., links that go in opposite directions). In fact, such a “link comparison” heuristic can sometimes closely mimic predictions from an EIG model (Coenen et al., 2015).

Finally, some inquiry behaviors might be selected via a reinforcement learning strategy where questions or actions that lead to positive outcomes are repeated (Sutton & Barto, 1988). For example, you might ask a speaker in a psychology talk “Did you consider individual differences in your study?” because in the past this has been a useful question to ask no matter the speaker or content. While this might lead to highly stereotyped and context-inappropriate questions, it is in fact possible to train sophisticated reinforcement learning agents to adapt question asking to particular circumstances based on intrinsic and extrinsic reward signals (Bachman, Sordoni, & Trischler, 2017). Importantly, the reinforcement learning approach arrives at the value of an action in an entirely different way than an OED model. Instead of prospectively evaluating possible answers and their impact on current beliefs, it relies on a history of past reinforcement. Depending on the specific assumptions, this approach may be discriminable from OED models, particularly during early learning of inquiry strategies.

Adaptive strategy selection. These alternative information-gathering strategies deserve consideration alongside OED not only as alter-

native theoretical frameworks (as for instance the reinforcement learning approach might represent) but also because they might represent cognitive strategies that trade off against more OED-consistent approaches in different situations. Following on this latter idea, what determines whether people follow an optimal OED norm or a heuristic that is easier to use, like controlling variables or limit testing? While determinants of strategy selection have been studied extensively in other domains, like decision making (Lieder et al., 2014; Marewski & Schooler, 2011; Otto, Raio, Chiang, Phelps, & Daw, 2013; Rieskamp & Otto, 2006), relatively little work addresses this question in the inquiry literature. One exception is a recent study by Coenen et al. (2015), who investigated the use of an OED norm (EIG) and a simpler heuristic (positive testing strategy) in a causal inquiry task. Across multiple experiments, participants were asked to intervene on three-variable causal systems to determine which of two possible causal structures governed the behavior of each system (similar to Figure 5, top). Figure 5 (bottom) shows posterior inferences over the hyperparameter μ from a hierarchical Bayesian model of people’s intervention choices. This parameter measures the degree to which participants, on average, relied on an EIG strategy ($\mu = 1$), compared to a positive testing heuristic ($\mu = 0$), which cannot be explained as an approximation of EIG (see paper for full argument). The different distributions are posterior distributions of this parameter for different between-subject experiments that varied a number of task parameters. In the “Baseline” experiment, participants’ behavior was best described by a mixture of the two strategies. In subsequent experiments, however, behavior spanned a wide spectrum of strategy profiles. In the experiment corresponding to the rightmost distribution, labeled “EIG superior”, participants received an additional set of problems before completing the baseline task. These problems were specifically designed to penalize non-OED strategies (i.e. positive testing would yield com-

pletely uninformative outcomes most of the time, costing participants money). Having worked on these problems, participants were more likely to use EIG in the baseline part of the experiment, which indicates that, in principle, most people are able to implement the normative solutions if they learn that their performance would suffer severely otherwise. In contrast, in three experiments that added time pressure to the baseline task (see three leftmost distributions), participants' behavior was much more in line with the positive testing heuristic. This indicates that the availability of cognitive resources can determine how people trade off the use of more complex OED norms and simpler inquiry heuristics.

In a related example, Gureckis and Markant (2009) explored how people searched for information in a simple spatial game based on Battleship (see Figure 2). They identified two distinct search “modes” as the task unfolded. At the beginning of the task when the hypothesis space was relatively unconstrained, people's choices were less in accordance with specific OED predictions and instead appeared more random and exploratory. These decisions were also made relatively quickly. However, at later points in the game, people seemed to behave more in line with OED predictions and their reaction times slowed significantly. This particularly happened in parts of the task where a small number of highly similar hypotheses became viable (i.e., situations where OED might be more computationally tractable). This suggests that even within the context of a single learning problem, people might shift between strategies that are more exploratory (i.e., less directed by specific hypothesis) and more focused on the disambiguation of specific alternative hypotheses.

Summary. There are many factors that have yet to be explored with respect to their impact on strategy selection during inquiry, including task difficulty, working memory capacity, fatigue, and stress. Research into these topics will allow the

Figure 5. Top: Examples of two possible causal graphs relating three nodes (variables). The nodes can take on one of two values (on or off). In the experiment, participants had to intervene on a similar system by setting the value of the nodes in order to determine which of two possible causal graphs actually described the operation of a unknown system. Bottom: Inferred posterior probability over hyperparameter μ in different experiments reported in Coenen, Rehder, and Gureckis (2015). μ captures the average strategy weight of participants in a causal intervention task. When $\mu = 1$, behavior is completely captured by the OED norm Expected Information Gain (EIG), when $\mu = 0$, it is best fit with a heuristic positive testing strategy (PTS). Values in-between correspond to mixed strategies.

field to move beyond simple demonstrations of the OED principle and help explain and predict inquiry behavior in different environments and given particular circumstances of the learner. This topic is of practical importance because inquiry plays a crucial role in a number of high-stakes situations that happen under both external (e.g., time) and internal (e.g., stress) constraints, like emergency medical diagnosis or criminal investigations. Finally, this

line of research also dovetails with a growing interest in cognitive science for models that take into account the cost of computation, and could contribute to the empirical basis for the development of these models (Hamrick, Smith, Griffiths, & Vul, 2015; Lieder et al., 2014; Vul et al., 2014).

Question 7: What triggers inquiry behaviors in the first place?

OED models describe how people query their environment to achieve some particular learning goal. The importance of such goals is made clear by the fact that in experiments designed to evaluate OED principles, participants are usually instructed on the goal of a task and are often incentivized by some monetary reward tied to achieving that goal. Similarly, in developmental studies, children are often explicitly asked to answer certain questions, solve a particular problem, or choose between a set of actions (e.g., play with toy A or toy B, see Bonawitz et al., 2010). However, many real world information-seeking behaviors are generated in the absence of any explicit instruction, learning goal, or monetary incentive. What then inspires people to inquire about the world in the first place?

This is an extremely broad question and there are many possible answers. According to one approach, the well-specified goals that are typically used in OED experiments are representative of a more general information-maximizing “over-goal” that always accompanies people while navigating the world (e.g., Friston et al., 2015). This view is particularly well represented by research on children’s exploratory play, where the claim is often made that this behavior represents sophisticated forms of self-guided inquiry that arise spontaneously during unsupervised exploration (Schulz, 2012b). For example, Cook et al. (2011), whose study is described in more detail above, argue that OED computations form an integral part of preschoolers’ self-guided behavior even in the absence of concrete goals:

“...many factors affect the optimal

actions: prior knowledge and recent experience enter through the term $P(H)$, while knowledge about possible actions and likely affordances enters through the term $P(D|A, H)$ Our results suggest that children are sensitive to all of these factors and integrate them to guide exploratory play” (p. 348).

Under this view, the constraints of popular experimental paradigms simply help control for and standardize the behavior across participants, while still capturing the key aspects of self-motivated inquiry.

One objection to this view is that at any given moment there are many possible inquiry tasks a learner might decide to pursue. While reading this paper you might be tempted to take a break and read about the latest world news, track down the source of a strange sound you hear from the kitchen, or start learning a new instrument. All of these actions might reduce your uncertainty about the world in various ways, but it seems difficult to imagine how OED principles would help explain which task you choose to focus on.

An alternative view acknowledges these limitations of the OED framework and instead argues that OED applies specifically to inquiry devoted to some particular task “frame” (i.e., a setting in which certain hypotheses and actions become relevant). For example, a task frame might be a person in a foreign country trying to determine if the local custom involves tipping for service. The set of hypotheses relevant to this task deal specifically with the circumstances where people might be expected to tip (never, only for bar service, only for exceptional service, etc...), and do not include completely irrelevant hypotheses (e.g., how far away the moon is in kilometers). In psychology experiments, such tasks are made clear by the instructions, but in everyday settings a learner must choose a task before they engage in OED-like reasoning or learning strategies. This latter view seems somewhat more likely because absent a task frame the hypothesis generation issue (see Question 1) be-

comes even more insidious (imagine simultaneously enumerating hypotheses about events in the news, possible sources of noise in the kitchen, and strategies for improving your piano play). However, this leaves open the question of how people define these tasks or goals in the first place. Here we consider two elements to this selection: subgoal identification and intrinsic curiosity.

Subgoal construction. When learning it often makes sense to divide problems into individual components that can be tackled on their own. For example, if a learner's broader goal is to find out which ships are hidden in which location during a game of Battleship (see Figure 2), they might break down the problem into first approximating all the ships' locations, and then determining their shapes. For example, Markant, Settles, and Gureckis (2015) describe an empirical task that led people to decompose a 3-way categorization task into a series of 2-way classification problems while learning via self-guided inquiry. This happened despite the fact that the overall goal was to learn all three categories.

Many learning problems have a hierarchical structure of over-goals and subgoals. Whereas OED norms make predictions about how to address each individual subgoal, they do not naturally capture the process of dividing a problem space into different subsets of goals.

Understanding how people identify subgoals while approaching a complex learning problem is difficult (although there exists efforts in the reinforcement learning literature to formalize this process, see e.g., Botvinick, Niv, & Barto, 2009). A full account of subgoal development would probably require knowing a person's representation of the features of a problem and their preferences for the order of specific types of information.

However, there also exist cases in which goal partitions emerge not from an informational analysis of every individual problem, but via a learning process across many problems that yields a kind of "template" for asking questions in some

domain. A college admissions interviewer might learn, for example, that in order to estimate the quality of a prospective student, it is a useful subgoal to find out what types of books they've read in high school. This may not be the most efficient subgoal to learn about each individual student (it is probably more useful for potential English majors than Physics applicants), but may lead to good enough outcomes on average. In many domains such templates do not even need to be learned, because they have been developed by others and can be taught easily. Consider for example the "Five Ws" (Who? What? When? Where? Why?) that serve as a template for question-asking subgoals in many different areas of inquiry and for many different types of over-goals (solving a crime, following a storyline, understanding the causal structure of an event, etc.). It would be interesting to study how such conventional templates influence people's preferences for establishing hierarchies of goals, and how learned and conventional partitions trade-off or compete with the expected value of information, in particular tasks.

The subgoal/over-goal framework might provide a useful way for thinking about how OED principles might be selected in the first place. A learner might have a generic over-goal to "be an informed citizen" and this then leads to a variety of smaller inquiry tasks that help learn about the impact of proposed changes to tax policy or in the political maneuvering of various parties. Behaviors within these subgoals may look more like OED inquiry where alternative hypotheses are considered; by contrast, the over-goal is more nebulous and is not associated with enumerable hypotheses.

In sum, at least one piece of the puzzle for what triggers inquiry behavior is to consider how people select task frames. The subgoal idea may be an additional fruitful direction, because it makes clear how self-defined objectives might be constructed during learning.

Curiosity and Intrinsic Motivation. Of course, aside from specific goals, we might decide

to spend more time learning about a topic or task frame simply because we are curious. While disinterested OED models (i.e., those with a value function that does not include internal or external costs) are agnostic about *why* learners seek out information, there is a longstanding parallel research tradition in psychology that studies the cognitive and neural bases of curiosity and intrinsic motivation. For example, it is well-known that children spontaneously explore objects with some level of complexity or uncertainty without any instruction to do so (Cook et al., 2011; Kidd, Piantadosi, & Aslin, 2012; Schulz & Bonawitz, 2007; Stahl & Feigenson, 2015). Meanwhile, adults care about the answers to otherwise useless trivia questions (Kang et al., 2009). Experiments have also shown that humans and other primates are even willing to sacrifice primary rewards (like water, money, and time) in exchange for information without obvious use (Blanchard, Hayden, & Bromberg-Martin, 2015; Kang et al., 2009; Marvin & Shohamy, 2016).

To exhaustively review this literature is beyond the scope of this paper, and would be largely redundant in light of recent review articles on the subject (Gottlieb, 2012; Gottlieb, Oudeyer, Lopes, & Baranes, 2013; Kidd & Hayden, 2015; Loewenstein, 1994; Oudeyer, Gottlieb, & Lopes, 2016). However, there are some particularly intriguing findings and theoretical developments in the curiosity literature that we think deserve attention by psychologists studying inquiry with OED models. In particular, they point out factors and mechanisms that add value to certain sources of information over others. These sources of value could potentially be integrated with OED models to yield more accurate predictions about how people choose subjects of inquiry.

To explain curiosity, researchers have traditionally suggested that it is a primary drive (perhaps as a consequence of some evolutionary process that favors information seekers), or an expression of some innate tendency for sense-making (Berlyne, 1966; Chater & Loewenstein, 2015; Loewenstein,

1994). Similarly, recent work has proposed that people seek information because it generates a type of intrinsic reward, similar to “classic” extrinsic rewards, like food or money (Blanchard et al., 2015; Marvin & Shohamy, 2016). In support of this claim, some studies have found activation in primates’ neural reward circuitry during information search that is similar to activation during other types of value-based choice (specifically, the primate data was collected in dopaminergic midbrain neurons Bromberg-Martin & Hikosaka, 2011; Redgrave & Gurney, 2006). Furthermore, a set of recent fMRI studies with humans has found correlations between people’s self-reported curiosity about trivia-questions and activation in areas of the brain involved in processing other rewards (Gruber, Gelman, & Ranganath, 2014; Kang et al., 2009).

What types of information can trigger such intrinsic reward signals? A key component of many theories of curiosity-driven learning is an inverse U-shaped relationship between a learner’s current knowledge and their expressed curiosity about some fact, domain, or stimulus (Kang et al., 2009; Kidd, Piantadosi, & Aslin, 2014; Kidd & Hayden, 2015; Loewenstein, 1994). This means that curiosity is often highest for domains or tasks in which people’s knowledge is at an *intermediate* level. This finding tallies with learning and memory research showing that items with intermediate difficulty are often learned most efficiently (Atkinson, 1972; Metcalfe & Kornell, 2003). Thus, asking questions about facts or relationships that are “just slightly beyond the individual’s current grasp” (Metcalfe & Kornell, 2003) might be an adaptive strategy that helps direct people’s attention to worthwhile opportunities for learning (Vygotsky, 1962). This suggests that intrinsic reward from information can stem from a learner’s expected learning progress, an idea which is already used to build algorithms in self-motivated robots (Oudeyer, Kaplan, & Hafner, 2007).

Another source of informational value is the

anticipation of extrinsic rewards in the *future* that might come from obtaining information in the present. This was demonstrated empirically by Rich and Gureckis (2014, 2017), who showed that people's willingness to explore risky prospects increased with their expectation to encounter them again in the future. This instrumental motivation to explore may actually underlie many kinds of seemingly intrinsically motivated information-seeking behaviors, or at least play some part in shaping people's motivation to seek information. For example, one might be intrinsically curious about the answer to a trivia question, but this effect could be enhanced if the question also pertains to one's future goals (a trivia question about the capital cities in South America would have more appeal for someone who has to take a geography test next week).

While work on curiosity does not specifically focus on how people choose particular task frames and subgoals, it does identify factors that affect what kinds information people seek and offer some hints about why they do so in the first place. Future work is needed to disentangle when people seek information for instrumental (future extrinsic reward) or epistemic (knowledge progress) purposes, and what types of information have evolved to yield particularly strong intrinsic rewards.

Summary. Identifying the source of people's thirst for information lies outside the realm of the OED framework. However, it also lies at the very core of what makes inquiry such a fundamental and fascinating human activity, and thus deserves further study. To arrive at a unified set of computational principles that underlie curiosity, motivation, and informational value will likely require overlapping efforts by cognitive psychologists, neuroscientists and developmental researchers (Kidd & Hayden, 2015). Furthermore, recent advances in reinforcement learning models of intrinsic motivation (Oudeyer et al., 2007, 2016; Singh, Barto, & Chentanez, 2004) may serve as an important inspiration for computationally informed theories.

Question 8: How does inquiry-driven learning influence what we learn?

The OED framework emphasizes effective or even optimal information gathering which, in turn, implies more effective learning. One of the major reasons that inquiry behavior is a topic of study is because it has implications for how best to structure learning experiences in classrooms and other learning environments. For example, in the machine learning literature, active information selection can be proven to reduce the number of training examples needed to reach a particular level of performance (Settles, 2010). However, a key question is if active inquiry reliably conveys the same advantages for human learners. In the following section we review existing work on this topic, first considering the relevant benefits of active over passive learning, and next considering the effect of the decision to stop gathering information on learning. Our core question here concerns when active learning improves learning outcomes and how knowledge acquired during active learning can deviate from underlying patterns in the environment.

Active versus passive learning. A number of studies have attempted to compare active and passive learning in simple, well controlled environments (Castro et al., 2008; Markant & Gureckis, 2014; Sim, Tanner, Alpert, & Xu, 2015). For example, Markant and Gureckis (2014) had participants learn to classify simple shapes into two categories. The experiment contrasted standard, passive learning against a self-directed learning condition. In the passive learning condition, an exemplar was presented on the screen and after the delay the category label of the item was provided. Across trials participants attempted to learn how to best classify new exemplars. In the active learning condition, participants could design the exemplars themselves on each trial, and received the category label of the designed item. The critical difference between these conditions is if the learner themselves controls which exemplar is presented (active learning) or if it is selected by the experimenter or some external process (pas-

sive learning). The study found that active learning led to faster acquisition of the category than passive learning. Furthermore, a third condition, in which yoked participants viewed the designed examples of the active group in a passive setting, showed significantly worse performance than the other groups. The results showed that allowing participants to control the selection process during learning improved outcomes (a process the authors referred to as the hypothesis-dependent sampling bias). This bias is the tendency of active learners to select information to evaluate the hypothesis they currently have in mind, which often does not transfer to other learners (such as the yoked condition) who have alternative hypotheses in mind (Markant & Gureckis, 2014). One study has since tested this idea with children (Sim et al., 2015), while another has explored the boundaries of the effect (MacDonald & Frank, 2016; Markant, 2016).

A potential downside of selecting data based on one's current hypotheses and goals is highlighted by an idea that Fiedler (2008) calls the "ultimate sampling dilemma" According to this idea, there are two main ways people obtain information from the world around them. The first is *natural sampling*, a learning process in which the ambient statistical patterns of the environment are experienced through mostly passive observation. Natural sampling is related to unsupervised learning (Gureckis & Love, 2003; Pothos & Chater, 2005) but focuses more on the data generating process (i.e., how examples are encountered) rather than the lack of supervision and corrective feedback. For example, by walking around a new city one might get a sense of the typical size of a car in the region. Artificial sampling refers to situations where learners intervene on the world to influence what they learn about (e.g., asking about the size of one particular brand of car), thereby interrupting the natural way that data is generated. The ultimate sampling dilemma points out that these two forms of learning can sometimes trade off with each other because they expose different aspects of the en-

vironment. As Fiedler (2008) points out, natural sampling is less likely to bias learners, because they are exposed to the true patterns in the world without altering them through their own behavior. This allows them to learn about natural variation in the world and enables them to gather information about typicality or frequency, for example. On the other hand, artificial sampling, for instance based on an OED model, can have the benefit of being much more efficient for answering a particular question or for seeking out exemplars that occur rarely but are highly informative. In those cases, learning only via natural sampling can require waiting a long time for these particularly informative or infrequent patterns to occur. Of course in some domains, such as causal reasoning, artificial sampling or active intervention is actually *necessary* for uncovering certain relationships in the environment (Pearl, 2009; Schulz, Kushnir, & Gopnik, 2007). As a result, some combination of natural and artificial sampling may be best for promoting more robust learning (MacDonald & Frank, 2016). The best way to do this is still up for debate, however, and there remain key questions about how other elements, such as the learners' conception of a problem, influence learning.

By highlighting the benefits and potential pitfalls of artificial and natural sampling, the ultimate sampling dilemma quite naturally suggests how the tension between the two might be resolved. Since natural sampling helps build an accurate representation of the statistical properties of the world, it might be particularly beneficial for sampling environments that are novel and about which a learner lacks knowledge regarding the most important features or variables. Thus, natural sampling through passive observation forms a natural first step during inquiry in a novel domain, before being followed by more targeted hypothesis-driven inquiry. Of course, an important empirical question is if people are able to determine the best point at which to switch from one mode of questioning to the other (Tversky & Edwards, 1966).

Stopping rules. Besides choosing *what* to learn, people often face the question *when to stop* searching for information. This question is particularly relevant when reaching absolute certainty is unlikely or impossible and learners have to decide when their current level of information is “enough”. OED models (see Question 3) make predictions about when stopping becomes more desirable than making additional queries, but the problem can be formalized more broadly for scenarios in which the content of samples is trivial (i.e., problems that do not require an OED model to select particular queries). A common approach is to use principles of optimal control theory and dynamic programming (some of the work presented in this section takes this approach, see also Edwards, 1965; Ferguson, 1989, 2012). In psychology, the stopping question has been approached in different ways. While some researchers have studied whether people collect too much or too little information given the cost structure of a task, others have looked at the impact of stopping decisions on subsequent beliefs and behaviors. Here, we will investigate the second question, as it reveals some subtle ways in which control during learning can affect our beliefs.

How stopping decisions shape our experience.

A separate line of research investigated what effect (optimal) stopping strategies have on people’s experiences and beliefs about the world.

Stopping decisions can be the source of distorted views of people’s environment. For example, stopping rules can lead to asymmetric knowledge about different options if these options have valence of some sort (i.e. they can be rewarding or not). A good example is the so-called *hot stove effect* (Denrell & March, 2001). Loosely speaking, it is the tendency to underestimate the quality of novel or risky prospects that happen to yield low rewards early on and are subsequently avoided. Having a single bad experience at a new restaurant might deter customers from re-visiting and potentially correcting that bad impression in the future.

Since some bad experiences happen to be exceptions, some restaurants end up being undervalued as a consequence. On the other hand, a coincidental positive experience will not lead to a corresponding *overvaluation* because decision makers will likely revisit these restaurants to reap more benefits and eventually find out their “true” value via regression to the mean. If it turns out the initial good experience was an exception, customers will eventually realize this and correct their good impression. Similar effects have been observed in other tasks that involve choice-contingent information, like approach-avoid decisions (Rich & Gureckis, 2014), or situations in which access to feedback is asymmetric across prospects (Le Mens & Denrell, 2011).

This work demonstrates the potentially large impact that the seemingly innocuous decision of stopping search can have on how we perceive the world. By choosing to learn more about options that we think are rewarding and ignoring those that we suspect to be bad, we can end up with widely asymmetric beliefs. Such asymmetries can be the source of misconceptions with potentially problematic effects. On a social level, for example, they can produce and solidify stereotypes about people or whole social groups, or increase social conformity (Denrell & Le Mens, 2007). They can also lead to unnecessary risk-aversion and resistance to change (because good but variable prospects are more likely to yield low initial rewards), which can be harmful for both individuals and organizations in the long run. Future work should further investigate the impact of stopping decisions on people’s beliefs and judgments as well as determining methods of mitigating stopping-induced biases.

Summary. The results reviewed in this section highlight that optimal inquiry is not just a function of selecting the right queries to answer one’s questions in an OED sense. It also involves knowing the right time to switch between active and passive learning, realizing the right moment to terminate search, and being aware of the conditions under

which our self-selected data was generated. While existing results primarily stem from the judgment and decision-making literature, these issues hold relevance for educators because they help to lay out expectations about when active inquiry will succeed or fail as a pedagogical strategy. Similarly, the problem of deciding when to stop searching for information is crucial for many inquiry tasks, like dividing up study time for different material that will appear on a quiz, asking questions in emergency situations when time is of the essence, or deciding when one has collected enough data to finally start writing a manuscript. Making the wrong stopping decisions in any of these scenarios can have unintended negative consequences and undo any benefits that carefully executed OED methods have accrued.

Question 9: What is the developmental trajectory of inquiry abilities?

The OED hypothesis has been particularly influential in work on children's exploration and learning (Gopnik, 2012). To provide an extensive review of the large literature on children's inquiry behavior is beyond the scope of this paper (see Gopnik & Wellman, 2012; Schulz, 2012b, for excellent summaries). However, it is important to consider a few of the developmental issues involved in inquiry skills, particularly when these touch on core concepts related to OED.

The child as optimal scientist. A growing number of studies suggest that even young children are surprisingly sophisticated at detecting opportunities to obtain useful information. For instance, Stahl and Feigenson (2015) showed infants objects, some of which moved in ways that violate physical laws (e.g., solid objects that appear to move through solid walls). Subsequently, infants were found to explore these objects preferentially, even going so far as to perform actions like banging them on the floor to test their solidity (see also Bonawitz, van Schijndel, Friel, & Schulz, 2012). Similarly, preschool aged children have been shown to devote more exploratory play to

a toy after being shown confounded evidence for how it works (Schulz & Bonawitz, 2007; Cook et al., 2011; van Schijndel, Visser, van Bers, & Raijmakers, 2015). Children also seem to integrate subtle social cues to help guide their exploration, such as exploring more when an adult teacher provides uninformative instruction (Gweon, Palton, Konopka, & Schulz, 2014). Still further, some evidence suggests that children can effectively test simple causal hypotheses through interventions that maximize information (Kushnir & Gopnik, 2005; McCormack et al., 2016; Schulz, Gopnik, & Glymour, 2007).

Based on such findings, researchers have argued that children act in ways analogous to scientific experimentation (Schulz, 2012b). Gopnik (2009) writes "When they play, children actively experiment on the world and they use the results of these experiments to change what they think." (p. 244). In some areas of cognitive development these abilities are viewed as directly supporting the idea of the "child as an optimal scientist". The core of this idea, and what brings it into alignment with OED, is that children's prior knowledge, beliefs, and goals help to structure their information gathering behaviors.

While it is important and intriguing that young children show so many early signs of successful information gathering, not all of these behaviors need be thought of as following exclusively from OED principles. For example, a child might selectively play with a toy after being shown confounded evidence about how it works without considering alternative hypotheses about the causal structure (Schulz, 2012a). Likewise, if exploration always follows violations in expectations, eventually learning will cease because most of the time (e.g., outside the lab) the world works in reliable and predictable ways (Schulz, 2015). As a result, it is important to keep in mind alternative views on children's exploration. For example, Hoch, Rachwani, and Adolph (in review) describe how infants have seemingly haphazard exploration tendencies. Using head-mounted eyetrackers, these

studies show that infants rarely move directly toward focal locations (toys placed in different areas of a room) while walking or crawling that they had previously fixated while stationary, as might be expected with goal directed exploration. In addition, they move around empty rooms just as much as ones filled with interesting and novel toys (Hoch, O’Grady, & Adolph, in review). The opposing perspective offered by this work is that infants are not identifying possibilities for new information and strategically orienting toward them but instead engage in “high-variance” motor plans that discover information almost by serendipity.

The difference between these viewpoints reflects one of the key issues in evaluating the OED framework that we have raised throughout this article. It is possible that there is *some* set of goals and beliefs that makes the apparently haphazard behavior of infants in Hoch et al. (in review) make sense as an optimal information seeking strategy. However, it is also useful not to lose sight of questions about what might change and develop across childhood. As we have reviewed through this article, actually implementing OED-style computations is a complex cognitive ability requiring the coordination of hypotheses, actions, evidence, and learning. It is clear that precisely adhering to the OED framework (especially in messier, real world environments) requires more than what young children have so far been shown to do. For example, after kids identify opportunities for knowledge gain, they also have to figure out the best way to get that knowledge, and (as reviewed below) that has proven difficult, especially in complicated situations.

In the following sections we review three key developmental issues related to OED. First, we consider how the issue raised in Question 1 (hypothesis generation) bears on developmental changes in inquiry behavior. Next we review evidence about inquiry in more formal classroom situations. Finally, we discuss children’s question asking, an important type of inquiry behavior available after acquiring language. Throughout we attempt to focus

our review on how existing evidence bears on the core computations assumed by OED models and how components of this model might change over the course of development.

Explaining children’s variability via hypothesis sampling. One attempt to reconcile the view that children are optimal scientists but also seemingly random in their exploration is to acknowledge that children do not apply OED principles consistently or as well as adults, and instead exhibit more variable behavior (Cook et al., 2011; Schulz, 2012b). Children might simply enter the world with a broader hypothesis space, weaker priors, and/or fewer cognitive resources, which translates to seemingly noisy or more exploratory behavior (Bonawitz, Denison, Griffiths, & Gopnik, 2014; Bonawitz, Denison, Gopnik, & Griffiths, 2014; Denison et al., 2013; Gopnik & Wellman, 2012). Computationally, this might be consistent with approximate Bayesian inference by sampling hypotheses from the posterior, similar to the rational process models described under Question 1. In fact, Gopnik and colleagues have recently argued that the development of both internal (hypothesis sampling) and external (i.e., exploratory play, information generation actions) search may be akin to simulated annealing (Kirkpatrick, Gelatt, & Vecchi, 1983) where increasingly random and undirected search strategies in infancy slowly transition to more stable and structured patterns through to adulthood (Buchsbaum, Bridgers, Skolnick-Weisberg, & Gopnik, 2012). In some cases this can even lead younger learners to find solutions that evade adults by avoiding local minima (e.g., Gopnik, Griffiths, & Lucas, 2015).

Hypothesis sampling models can capture more or less variable behavior given different parameters (e.g., the number of samples taken) and thus provide one computational mechanism that naturally accommodates developmental change. Such an account might also accommodate the undirected exploration of Hoch et al. with the idea that the variability in the behavior is slowly “turned down”

over the course of development². Overall this approach seems promising as long as one keeps in mind the caveats raised about this approach under Question 1 above. For instance, sampling models so far tend to ignore deep integration with other cognitive processes (e.g., memory retrieval), and they also raise the question whether extremely variable behavior generated by such models can even be properly described as optimal in a way that captures the spirit of the child-as-scientist metaphor. In addition, applying such a model to explain the behavior of very young children can be very difficult because it is hard to identify what hypothesis space should be sampled (e.g. in Stahl & Feigenson, 2015, what hypothesis spaces about the physical world do children consider?).

Nevertheless, this approach remains a fertile area for exploration. One obvious empirical prediction of this theory is that the major change in inquiry behavior across development does not necessarily manifest in absolute performance but in inter-subject variability. This suggests a slightly different focus for developmental research which is often framed in terms of when children achieve adult performance in a task. If children are optimal but noisy, the key issue should be characterizing changes in variability.

Inquiry in the science classroom. The concept of inquiry as a cognitive activity has been hugely influential in science education (Inhelder & Piaget, 1958; Z. Chen & Klahr, 1999; Kuhn et al., 2000). A key focus has been to teach general strategies for learning about causal structure (e.g., how variables such as water, fertilizer, or sunlight might influence the growth of a plant in a science lab, see Klahr et al., 1993). However, compared to the developmental literature reviewed above, the conclusion of much of this research is that children often struggle into early adolescence with learning and implementing such strategies. For instance, as reviewed in Question 6, children famously have trouble learning the principle of *controlling variables* (i.e., changing one thing at a time) and applying it spontaneously to new problems (Z. Chen

& Klahr, 1999; Klahr & Nigam, 2004). That is, without the right kinds of instruction, young children tend to want to change many things at once, rather than testing individual factors or variables of a causal system in isolation. One reason for this preference, identified by Kuhn et al. (2000), is that children often have not developed a metastrategic understanding of *why* controlling variables works and what different inferences are warranted by conducting a confounded versus a controlled experiment. Interestingly, recent analyses show that the control of variables is an effective, even optimal (in the OED sense) strategy, only given particular assumptions about the causal structure of the environment (Coenen et al., 2017).

A related example stems from work on children's causal reasoning. Bonawitz et al. (2010) presented young toddlers, preschoolers, and adults with a sequence of two events (first a block contacted a base of an object and then a toy connected to the base started spinning). The question was whether participants subsequently generated an intervention (moving the block to the base) to test if this event was causally related to the second one (i.e., the spinning toy). Unlike preschoolers and adults, toddlers did not perform this hypothesis test spontaneously, although they did come to anticipate the second event from the first. To successfully generate actions they required additional cues, like causal language used by the experimenter, or seeing direct contact between two objects (here, the two toys not separated by the base). The authors hypothesize that this failure to generate spontaneous interventions might be due to young children's inability to recognize the relationship between prediction and causality, unless they are explicitly told or shown.

²However, Hoch (personal communication) reports no evidence of developmental trends in goal directed walking with age or experience within the age ranges her research has considered. Cole, Robinson, and Adolph (2016) found similar rates of goal-directed exploration in 13-month-old (novice walkers) and 19-month old (experienced walker) infants.

In sum, while the previous section provided evidence that seems to support the idea that children identify and explore in systematic ways, claims about the “child as intuitive scientist” remain complicated by the evidence that children struggle learning generalizable strategies of acquiring information that are most akin to the actual practice of science (e.g., control of variables). The question of how kids get from the well documented motivations and abilities that emerge in early childhood to the more complex abilities of older children, adults, and scientists remains an important contradiction in the field and is an important area for future work. This is particularly challenging because it is not clear what, in terms of computational components, actually fails when children do not show mastery of the control of variables. One hint is that it is sometimes easier for children to acquire the control of variables strategy for a particular domain or task than it is to identify how to properly transfer that strategy to a new domain or task. This suggests that aspects of problem solving and transfer may be relevant parts of the answer (Gick & Holyoak, 1983; Catrambone & Holyoak, 1989). In addition, the two examples described here point out that some type of metaknowledge about the value and purpose of “fair tests” is an important precursor for being able to reliably implement such strategies. As mentioned above, these issues currently fall somewhat outside the OED framework which deals primarily with information action selection within a defined goal and framework.

Children’s questions. The view of the child as an optimal (i.e., OED) scientist is further complicated by the literature on children’s question asking. As described in detail in Question 2, asking interesting and information questions using language is important over the course of cognitive development. Children are notorious question askers, and even young children seem to acquire question-like utterances within the first few years in their vocabulary (e.g., “Eh?” or “Doh?” to mean “What is that?”) (K. Nelson, 1973). It has been hy-

pothesized that these pseudo-words aid in the development of language acquisition by coordinating information requests between the child and caregiver.

However, it is unclear if children’s questions reflect any particular sense of optimality (Rothe, Lake, & Gureckis, in review). Part of the reason is that most of the education research on question asking in classrooms has focused on qualitative distinctions between good and bad questions (see Graesser et al., 1993; Graesser & Person, 1994; Chin & Brown, 2002). For example, studies might observe the types of questions students ask in a lecture or while reading some text as “deep” (e.g., why, why not, how, what-if, what-if-not) in contrast with “shallow” questions (e.g., who, what, when, where). Interestingly, the proportion of deep questions asked in a class correlate with students’ exam scores (see Graesser & Person, 1994). Such classification schemes are useful starting places but do not yet allow us to assess if this behavior is reflective of specific OED principles.

To that end, more controlled experimental tasks have shown robust developmental changes in children’s question asking behavior. One classic finding is that younger children (e.g., 7 to 8 years old) often use less efficient question-asking strategies while playing a “20 questions”/“Guess who?” game compared to older children (e.g., 9 to 11 years old) and adults (Mosher & Hornsby, 1966; Ruggeri & Lombrozo, 2015). Younger children have been shown to use very specific question asking strategies (e.g., “Is it Bill?” “Is it Sally?”) that potentially rule out one particular hypothesis at a time (sometimes called hypothesis-scanning questions). In contrast, older children and adults ask constraint-seeking questions that can more effectively narrow down the hypothesis space (e.g., “Is your person wearing glasses?”, “Is your person a man?”). This suggests that designing sophisticated testing strategies that pertain to multiple hypotheses is another skill that develops over time. Whether this is due to limitations or changes in working memory capacity, informational goals, or

beliefs about the value of evidence is still an open question.

At least some recent work has attempted to better understand these patterns. For example, Ruggeri, et al. (2015) found a developmental trend in the degree to which children's questions matched the predictions of an OED models based on EIG. However, they explain that the apparent inefficiency in young children's questions stems from younger children adopting inappropriate stopping rules (asking questions after they have already determined the correct answer, see Question 8 above). In addition, recent work that has attempted to unpack the contribution to various component cognitive processes to this ability (e.g., isolating the ability to ask questions from the need to update beliefs on the basis of the answers) has found a complex relationship between these issues even among young learners. For example, forcing children to explicitly update their hypothesis space with each new piece of evidence actually led them to ask more optimal questions than a condition where a computer interface tracked the evidence for them (which the authors interpret as a type of "desirable difficulty" during learning, Kachergis, Rhodes, & Gureckis, 2017).

Summary. The studies reviewed in this section give a nuanced view of the development of inquiry skills. In some cases, even very young children seem remarkably successful at identifying opportunities for learning. However, it is also clear that children show difficulty in many places where they might be expected to succeed (e.g., in learning scientific inquiry skill directly, or in formulating informative questions). To better understand children's inquiry behavior, more work is needed to unpack the individual components that contribute to it, such as children's theories about the world, their cognitive capacities, and their understanding of their own actions as means to test theories.

Final Thoughts

In this paper we examine how people ask questions and inquire about the world. Our review

poses the following question: Are we asking the right questions about human inquiry? Our synthesis offers two summary insights. First, the OED hypothesis has contributed a great deal of theoretical structure to the topic of human inquiry. Qualitative theories have been superseded by principled (and often parameter-free) models that often explain human behavior in some detail. OED models have been successful at providing explanations at different levels of processing including neural, perceptual and higher-level. Human inquiry is a very rich and open ended type of behavior, so the success of the theoretical framework across so many tasks or situations is remarkable.

However, at the same time OED has rather severely limited the focus of research on human inquiry. Of course, constraining research questions and methods is a necessary (and desirable) function of a(ny) cognitive model or scientific paradigm, so we do not claim that finding limitations of OED constitutes a groundbreaking contribution in and of itself. However, being aware of *how* current theories constrain our thinking and critically reflecting on their merits is invaluable for the long-term progress of our field. In this respect we hopefully convinced at least some readers that the OED theories suffer from a number of particularly troubling blind spots. Some of the hardest questions about human inquiry, including the motivational impetus to acquire information about particular phenomena, are hard to accommodate within OED formalisms. Furthermore, the richer set of situations in which inquiry proceeds (e.g., natural language question asking) remain important gaps in our current understanding. These gaps matter because there is not currently a plausible way to account for these behaviors within the bounds of the OED framework and in many cases it is doubtful that there ever will be. In addition, these topics are exactly the situations that are most interesting to other aligned fields where inquiry is a basic concern. Perhaps the best illustration of the continued disconnect is the fact that the OED hypothesis has

become a widely adopted and popular approach to study learning, but has had little or no impact on current thinking in education. Papers about OED models are almost never published in education journals. Certainly some elements of this can be chalked up to the different roles that formal models play in different fields. However, it also must be acknowledged that it still is difficult to apply OED models outside of the carefully constructed experimental contexts studied by psychologists.

The nine questions laid out in this paper hopefully offer a way forward. We have attempted to highlight what we see as the most exciting, difficult, and under-explored questions about human inquiry. As we suggest throughout this paper, answering these questions will likely require a number of different experimental paradigms and modeling approaches, many of which do not follow the classic structure of OED studies. Our hope is that there are enough ideas presented here that a graduate student might build a thesis around any one of the topics raised. Before concluding, we believe it is worthwhile to consider how *answers* to our questions could lead to progress in a number of domains beyond basic cognitive science.

Education. One contribution of this article is to elucidate the set of constraints and prerequisites that surround people's ability to effectively learn from self-directed inquiry and exploration. We argued that a solid understanding of these constraints and their developmental trajectory, as well as, ultimately, the development of computational models that incorporate these constraints will help apply cognitive science within educational contexts. What are some insights of future work that could benefit educational practices?

Take as an example the first question we raise in this paper, which challenges the assumption that people can represent *all* possible hypotheses about some learning domain. We suggested that future work should develop models of hypothesis generation that take into account constraints of the learner, for instance in terms of their memory processes or cognitive resources. Progress in this area

could be directly applicable to the development of adaptive learning systems, which are growing in popularity both in schools (e.g., U.S. Department of Education, 2017) and as part of online learning tools that are used by the broader population. The success of adaptive learning systems crucially relies on being able to predict what information would be most useful to a user (e.g., what materials to train, re-train, and test them on). This, in turn, requires an accurate representation of their current hypothesis space. Integrating process-level theories of memory and resource constraints into models of hypothesis generation could thus lead to significant improvement of these technologies.

Another important line of research we highlight in this paper concerns the relationship between active learning and passive learning (e.g., in the discussion of the "ultimate sampling dilemma" in Question 8). We point out that the two modes of learning yield different benefits and thus work best in different situations, depending on the context and current knowledge of the learner. We hope that future work will develop models that can determine mixtures of those two modes of learning that optimize learning success in particular subject areas. Insights from these models could be used, for example, to design educational interventions in subjects that rely on combinations of teaching and experimentation (like many physical and life sciences).

Machine intelligence. OED computations already play an important role in the field of machine learning, where they are often used to design so-called *active learning algorithms* (Settles, 2010). Active learning is used to select particular unlabeled items (e.g., images, text, or speech) and have them labeled by a (human) oracle with the goal of improving classification of future items. We have discussed how human active learning can far exceed this particular situation, for example with a breadth of different query types and strategies that preserve computational tractability of even complex queries. Some types of queries aim to test concrete hypotheses (these

are the “classic” OED questions), some seek out relevant meta-knowledge (feature queries), some address a particular knowledge gap, and some merely follow shared conventions (asking “How are you?”). Building a computational “repertoire” of these different query types could be especially valuable for the development of conversational machine intelligence, like chatbots and digital assistants, that can ask and answer questions in conversation with a human. Currently, these technologies tend to be limited to fairly narrow domains, beyond which they are unable to respond adequately to users’ questions. Over the past few years, Machine Learning researchers have started to develop models that generate or answer a broader array of questions, specifically about images (e.g., Jain, Zhang, & Schwing, 2017; Ren, Kiros, & Zemel, 2015). However, these algorithms work by training on large datasets of question-image pairs, and have no way of taking into account the context of any given conversation or features of the user (e.g., their current goal). Psychologically-inspired models that can adapt to changes in the subject, context, and goals of the conversation partner would thus be enormously helpful in making these tools more flexible and realistic.

Another point we raised above is the importance of pragmatics in question asking. To be helpful to one’s human counterpart and to draw the right conclusions from their answers requires at least a basic model of their knowledge and expectations, as well as the context of the conversation (e.g., is this polite small talk, or does this person want me to teach them something?). Recent work on human-robot interaction has demonstrated just how important it is that people perceive robots, with whom they collaborate on a joint task, as adapting to their actions (Fisac et al., 2016). It showed, for example, that a robot that acts optimally with respect to the task can be immensely frustrating to their human “partner” if the partner’s strategy happens to be suboptimal. Computational models of how humans interpret each other’s questions in a

given context could be used to improve artificial agents in their ability to account for their conversation partner’s goals and constraints when answering and generating questions.

Experimental methods. What are the ramifications of our discussion for experimental methods used within psychology? We hope that the work we reviewed provides yet another set of examples for why it is informative to study people actively seeking information. Although this view is not new and has gained momentum over the past years, the vast majority of learning experiments still rely on paradigms in which subjects passively observe a sequence of stimuli preselected by the experimenter (Gureckis & Markant, 2012). This approach is desirable because it gives us experimental control over the information presented to participants, but it lacks one of the most crucial components of real-world learning, that is, the ability to change our environment, ask questions, and explore the world around us. Through the development of sophisticated modeling techniques, many of which are highlighted in this paper, researchers are now developing research methodologies that exploit the lack of experimental control, instead of sacrificing validity because of it. Beyond the OED framework, we have for example pointed to models of pragmatic reasoning, sequential sampling, tree search, or optimal stopping. All of these can provide windows into different aspects of inquiry and, taken together, we believe they make giving up *some* experimental control worthwhile.

Acknowledgments

We thank Neil Bramley, Justine Hoch, Doug Markant, Greg Murphy, and Marjorie Rhodes for many helpful comments on a draft of this paper. We also thank Kylan Larson for assistance with illustrations. This work was supported by BCS-1255538 from the National Science Foundation, the John S. McDonnell Foundation Scholar Award, and a UNSW Sydney Visiting Scholar Fellow, to TMG; and by NE 1713/1-2 from the Deutsche Forschungsgemeinschaft (DFG) as part of the

"New Frameworks of Rationality" (SPP 1516) priority program, to JDN.

References

- Anderson, J. R. (1990). *The adaptive character of thought*. Hillsdale, NJ: Erlbaum.
- Anderson, J. R. (1991). The adaptive nature of human categorization. *Psychological Review*, 98(3), 409-429.
- Atkinson, R. C. (1972). Optimizing the learning of a second-language vocabulary. *Journal of Experimental Psychology*, 96(1), 124.
- Austerweil, J., & Griffiths, T. (2011). Seeking confirmation is rational for deterministic hypotheses. *Cognitive Science*, 35, 499-526.
- Bachman, P., Sordoni, A., & Trischler, A. (2017). Towards information-seeking agents. In *Iclr*.
- Baron, J., Beattie, J., & Hershey, J. C. (1988). Heuristics and biases in diagnostic reasoning: II. congruence, information, and certainty. *Organizational Behavior and Human Decision Processes*, 42(1), 88-110.
- Baron, J., & Hershey, J. C. (1988). Outcome bias in decision evaluation. *Journal of personality and social psychology*, 54(4), 569.
- Bartlett, F. C., & Burt, C. (1933). Remembering: A study in experimental and social psychology. *British Journal of Educational Psychology*, 3(2), 187-192.
- Battaglia, P., Hamrick, J., & Tenenbaum, J. (2013). Simulation as an engine of physical scene understanding. *Proceedings of the National Academy of Sciences*, 110(45), 18327-18332.
- Bayes, T. (1763). An essay towards solving a problem in the doctrine of chance. *Philosophical Transactions of the Royal Society of London*, 53(370-418).
- Berge, C. (1971). *Principles of combinatorics*. San Diego, CA: Academic Press.
- Berlyne, D. E. (1966). Curiosity and exploration. *Science*, 153(3731), 25-33.
- Blanchard, T. C., Hayden, B. Y., & Bromberg-Martin, E. S. (2015). Orbitofrontal cortex uses distinct codes for different choice attributes in decisions motivated by curiosity. *Neuron*, 85(3), 602-614.
- Blum, A. L., & Langley, P. (1997). Selection of relevant features and examples in machine learning. *Artificial intelligence*, 97(1), 245-271.
- Bonawitz, E. B., Denison, S., Gopnik, A., & Griffiths, T. L. (2014). Win-stay, lose-sample: A simple sequential algorithm for approximating bayesian inference. *Cognitive psychology*, 74, 35-65.
- Bonawitz, E. B., Denison, S., Griffiths, T. L., & Gopnik, A. (2014). Probabilistic models, learning algorithms, and response variability: sampling in cognitive development. *Trends in cognitive sciences*, 18(10), 497-500.
- Bonawitz, E. B., Ferranti, D., Saxe, R., Gopnik, A., Meltzoff, A. N., Woodward, J., & Schulz, L. (2010). Just do it? investigating the gap between prediction and action in toddlers causal inferences. *Cognition*, 115(1), 104-117.
- Bonawitz, E. B., Shafto, P., Gweon, H., Goodman, N. D., Spelke, E., & Schulz, L. (2011). The double-edged sword of pedagogy: Instruction limits spontaneous exploration and discovery. *Cognition*, 120(3), 322-330.
- Bonawitz, E. B., van Schijndel, T. J., Friel, D., & Schulz, L. (2012). Children balance theories and evidence in exploration, explanation, and learning. *Cognitive psychology*, 64(4), 215-234.
- Borji, A., & Itti, L. (2013). State-of-the-art in visual attention modeling. *IEEE transactions on pattern analysis and machine intelligence*, 35(1), 185-207.
- Botvinick, M. M., Niv, Y., & Barto, A. C. (2009). Hierarchically organized behavior and its neural foundations: A reinforcement learning perspective. *Cognition*, 113(3), 262-

280.

- Bramley, N. R., Dayan, P., & Lagnado, D. A. (2015). Staying afloat on neuraths boat—heuristics for sequential causal learning. In *Proceedings of the 36th annual conference of the cognitive science society* (pp. 262–267).
- Bramley, N. R., Gerstenberg, T., & Tenenbaum, J. B. (2016). Natural science: Active learning in dynamic physical microworlds. In A. Papafragou, D. Grodner, D. Mirman, & J. C. Trueswell (Eds.), *Proceedings of the 38th annual meeting of the cognitive science society*. Austin, TX.
- Bramley, N. R., Lagnado, D., & Speekenbrink, M. (2015). Conservative forgetful scholars - how people learn causal structure through sequences of interventions. *Journal of Experimental Psychology: Learning, Memory, and Cognition*.
- Bromberg-Martin, E. S., & Hikosaka, O. (2011). Lateral habenula neurons signal errors in the prediction of reward information. *Nature neuroscience*, 14(9), 1209–1216.
- Brown, S. D., & Steyvers, M. (2009). Detecting and predicting changes. *Cognitive Psychology*, 58(1), 49–67.
- Buchsbaum, D., Bridgers, S., Skolnick-Weisberg, D., & Gopnik, A. (2012). The power of possibility: causal learning, counterfactual reasoning, and pretend play. *Philosophical Transactions of the Royal Society of London*, 367(1599), 2202–2212.
- Cakmak, M., & Thomaz, A. L. (2012). Designing robot learners that ask good questions. In *Proceedings of the seventh annual acm/ieee international conference on human-robot interaction* (pp. 17–24).
- Carey, S., & Spelke, E. (1996). Science and core knowledge. *Philosophy of science*, 63(4), 515–533.
- Case, R. (1974). Structures and strictures: Some functional limitations on the course of cognitive growth. *Cognitive Psychology*, 6(4), 544–574.
- Castro, R., Kalish, C., Nowak, R., Qian, R., Rogers, T., & Zhu, X. (2008). Human active learning. In *Advances in neural information processing systems* (Vol. 21). Cambridge, MA: MIT Press.
- Catrambone, R., & Holyoak, K. (1989). Overcoming contextual limitations on problem-solving transfer. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 15(6), 1147.
- Cavagnaro, D. R., Aranovich, G. J., McClure, S. M., Pitt, M. A., & Myung, J. I. (2014). On the Functional Form of Temporal Discounting: An Optimized Adaptive Test. *Journal of Risk and Uncertainty*, 52(3), 233–254.
- Cavagnaro, D. R., Myung, J. I., Pitt, M. A., & Kujala, J. V. (2010). Adaptive design optimization: A mutual information-based approach to model discrimination in cognitive science. *Neural computation*, 22(4), 887–905.
- Chater, N., Crocker, M., & Pickering, M. (1998). The rational analysis of inquiry: The case of parsing. In M. Oasford & N. Chater (Eds.), *Rational models of cognition* (p. 441–468). Oxford University Press.
- Chater, N., & Loewenstein, G. (2015). The underappreciated drive for sense-making. *Journal of Economic Behavior & Organization*.
- Chen, S. Y., Ross, B. H., & Murphy, G. L. (2014). Implicit and explicit processes in category-based induction: Is induction best when we don't think? *Journal of Experimental Psychology: General*, 143(1), 227.
- Chen, Z., & Klahr, D. (1999). All other things being equal: Acquisition and transfer of the control of variables strategy. *Child development*, 70(5), 1098–1120.
- Chin, C., & Brown, D. E. (2002). Student-generated questions: A meaningful aspect of learning in science. *International Journal of Science Education*, 24(5), 521–549.
- Christie, S., & Genter, D. (2010). Where hypotheses come from: Learning new relations by

- structural alignment. *Journal of Cognition and Development*, 11(3).
- Coenen, A., Bramley, N. R., Ruggeri, A., & Gureckis, T. M. (2017). Beliefs about sparsity affect causal experimentation. In *Proceedings of the 39th Annual Conference of the Cognitive Science Society*. Austin, TX.
- Coenen, A., & Gureckis, T. M. (2015). Are biases when making causal interventions related to biases in belief updating? In R. D. Noelle, A. S. Warlaumont, J. Yoshimi, T. Matlock, C. D. Jennings, & P. P. Maglio (Eds.), *Proceedings of the 37th annual conference of the cognitive science society*. Austin, TX: Cognitive Science Society.
- Coenen, A., Rehder, B., & Gureckis, T. M. (2015). Strategies to intervene on causal systems are adaptively selected. *Cognitive Psychology*, 79, 102–133.
- Cole, W., Robinson, S., & Adolph, K. (2016). Bouts of steps: The organization of infant exploration. *Developmental Psychobiology*, 58, 341–354.
- Cook, C., Goodman, N. D., & Schulz, L. (2011). Where science starts: Spontaneous experiments in preschoolers exploratory play. *Cognition*, 120(3), 341–349.
- Courville, A. C., & Daw, N. D. (2007). The rat as particle filter. In *Advances in neural information processing systems* (pp. 369–376).
- Crupi, V., Nelson, J., Meder, B., Cevolani, G., & Tentori, K. (in press). Generalized information theory meets human cognition: Introducing a unified framework to model uncertainty and information search. *Cognitive Science*.
- Crupi, V., & Tentori, K. (2014). State of the field: Measuring information and confirmation. *Studies in History and Philosophy of Science*, 47, 81–90.
- Denison, S., Bonawitz, E., Gopnik, A., & Griffiths, T. L. (2013). Rational variability in children's causal inferences: The sampling hypothesis. *Cognition*, 126(2), 285–300.
- Denrell, J., & Le Mens, G. (2007). Interdependent sampling and social influence. *Psychological review*, 114(2), 398.
- Denrell, J., & March, J. G. (2001). Adaptation as information restriction: The hot stove effect. *Organization Science*, 12(5), 523–538.
- Doherty, M. E., Mynatt, C. R., Tweney, R. D., & Schiavo, M. D. (1979). Pseudodiagnosticity. *Acta Psychologica*, 43(2), 111–121.
- Dougherty, M. R. P., & Hunter, J. (2003b). Probability judgment and subadditivity: the role of working memory capacity and constraining retrieval. *Memory & cognition*, 31(6), 968–982. doi: 10.3758/BF03196449
- Dougherty, M. R. P., & Hunter, J. E. (2003a). Hypothesis generation, probability judgment, and individual differences in working memory capacity. *Acta Psychologica*, 113(3), 263–282. doi: 10.1016/S0001-6918(03)00033-7
- Dougherty, M. R. P., Thomas, R., & Lange, N. (2010). Toward an integrative theory of hypothesis generation, probability judgment, and hypothesis testing. *Psychology of learning and motivation*, 52, 299–342.
- Edwards, W. (1965). Optimal strategies for seeking information: Models for statistics, choice reaction times, and human information processing. *Journal of Mathematical Psychology*, 2(2), 312–329.
- Edwards, W. (1968). Conservatism in human information processing. *Formal representation of human judgment*, 17, 51.
- Elmore, J. G., Barton, M. B., Mocer, V. M., Polk, S., Arena, P. J., & Fletcher, S. W. (1998). Ten-year risk of false positive screening mammograms and clinical breast examinations. *New England Journal of Medicine*, 338(16), 1089–1096.
- Fedorov, V. V. (1972). *Theory of optimal experiments*. Access Online via Elsevier.
- Ferguson, T. S. (1989). Who solved the secretary problem? *Statistical science*, 282–289.
- Ferguson, T. S. (2012). *Optimal stopping and ap-*

- lications. Electronic Text. Retrieved from <https://www.math.ucla.edu/~tom/Stopping/Contents.html>
- Fernbach, P. M., Darlow, A., & Sloman, S. A. (2010). Neglect of alternative causes in predictive but not diagnostic reasoning. *Psychological Science*, 21(3), 329–336.
- Fernbach, P. M., Darlow, A., & Sloman, S. A. (2011). When good evidence goes bad: The weak evidence effect in judgment and decision-making. *Cognition*, 119(3), 459–467.
- Fisac, J. F., Liu, C., Hamrick, J. B., Sastry, S., Hedrick, J. K., Griffiths, T. L., & Dragan, A. D. (2016). Generating plans that predict themselves. In *Proceedings of wafr*.
- Flammer, A. (1981). Towards a theory of question asking. *Psychological Research*, 43(4), 407–420.
- Frank, M. C., & Goodman, N. D. (2012). Predicting pragmatic reasoning in language games. *Science*, 336(6084), 998–998.
- Franke, M., & Degen, J. (2016). Reasoning in reference games: Individual-vs. population-level probabilistic modeling. *PloS one*, 11(5), e0154854.
- Friston, K. (2009). The free-energy principle: a rough guide to the brain? *Trends in cognitive sciences*, 13(7), 293–301.
- Friston, K., FitzGerald, T., Rigoli, F., Schwartenbeck, P., & Pezzulo, G. (2017). Active inference: A process theory. *Neural computation*, 29(1), 1–49.
- Friston, K., Rigoli, F., Ognibene, D., Mathys, C., Fitzgerald, T., & Pezzulo, G. (2015). Active inference and epistemic value. *Cognitive neuroscience*, 6(4), 187–214.
- Gershman, S., & Daw, N. (2017). Reinforcement learning and episodic memory in humans and animals: An integrative framework. *Annual Review of Psychology*, 68, 1–28.
- Gershman, S., Vul, E., & Tenenbaum, J. B. (2012). Multistability and perceptual inference. *Neural computation*, 24(1), 1–24.
- Gick, M., & Holyoak, K. (1983). Schema induction and analogical transfer. *Cognitive Psychology*, 15(1), 1–38.
- Gigerenzer, G., Mata, J., & Frank, R. (2009). Public knowledge of benefits of breast and prostate cancer screening in europe. *Journal of the National Cancer Institute*, 101(17), 1216–1220.
- Ginsberg, M., & Smith, D. (1988). Reasoning about action i: A possible worlds approach. *Artificial Intelligence*, 35(2), 165–195.
- Good, I. J. (1950). *Probability and the weighting of evidence*. New York: Charles Griffin.
- Goodman, N. D., Frank, M., Griffiths, T., & Tenenbaum, J. (2015). Relevant and robust. a response to marcus and davis. *Psychological Science*, 26(4), 539–541.
- Goodman, N. D., & Frank, M. C. (2016). Pragmatic language interpretation as probabilistic inference. *Trends in Cognitive Sciences*, 20(11), 818–829.
- Goodman, N. D., & Stuhlmüller, A. (2013). Knowledge and implicature: Modeling language understanding as social cognition. *Topics in cognitive science*, 5(1), 173–184.
- Gopnik, A. (1996). The scientist as child. *Philosophy of science*, 63(4), 485–514.
- Gopnik, A. (2009). *The philosophical baby: What children's minds tell us about truth, love & the meaning of life*. Random House.
- Gopnik, A. (2012). Scientific thinking in young children: Theoretical advances, empirical research, and policy implications. *Science*, 337(6102), 1623–1627.
- Gopnik, A., Glymour, C., Sobel, D. M., Schulz, L., Kushnir, T., & Danks, D. (2004). A theory of causal learning in children: causal maps and bayes nets. *Psychological review*, 111(1), 3.
- Gopnik, A., Griffiths, T., & Lucas, C. (2015). When younger learners can be better (or at least more open-minded) than older ones. *Current Directions in Psychological Science*, 24(2), 87–92.
- Gopnik, A., & Wellman, H. M. (2012). Re-

- constructing constructivism: Causal models, bayesian learning mechanisms, and the theory theory. *Psychological Bulletin*, 138(6), 1085.
- Gottlieb, J. (2012). Attention, learning, and the value of information. *Neuron*, 76(2), 281–295.
- Gottlieb, J., Oudeyer, P.-Y., Lopes, M., & Baranes, A. (2013). Information-seeking, curiosity, and attention: computational and neural mechanisms. *Trends in Cognitive Sciences*, 17(11), 585–593.
- Graesser, A., Langston, M., & Bagget, W. (1993). Exploring information about concepts by asking questions. In G. Nakamura, R. Taraban, & D. Medin (Eds.), *The psychology of learning and motivation: Vol. 29: Categorization by humans and machines* (p. 411–436). Academic Press.
- Graesser, A., & Person, N. K. (1994). Question asking during tutoring. *American Educational Research Journal*, 31(1), 104–137.
- Gregg, L. W., & Simon, H. A. (1967). Process models and stochastic theories of simple concept formation. *Journal of Mathematical Psychology*, 4(2), 246–276.
- Grice, H. P. (1975). Logic and conversation. 1975, 41–58.
- Griffiths, T. L., Chater, N., Kemp, C., Perfors, A., & Tenenbaum, J. B. (2010). Probabilistic models of cognition: Exploring representations and inductive biases. *Trends in Cognitive Sciences*, 14, 357–364.
- Griffiths, T. L., & Tenenbaum, J. B. (2006). Optimal predictions in everyday cognition. *Psychological science*, 17(9), 767–773.
- Gruber, M. J., Gelman, B. D., & Ranganath, C. (2014). States of curiosity modulate hippocampus-dependent learning via the dopaminergic circuit. *Neuron*, 84(2), 486–496.
- Guez, A., Silver, D., & Dayan, P. (2012). Efficient bayes-adaptive reinforcement learning using sample-based search. In *Advances in neural information processing systems* (pp. 1025–1033).
- Gureckis, T. M., & Love, B. C. (2003). Human unsupervised and supervised learning as a quantitative distinction. *International Journal of Pattern Recognition and Artificial Intelligence*, 17, 885–901.
- Gureckis, T. M., & Markant, D. B. (2009). Active learning strategies in a spatial concept learning game. In *Proceedings of the 31st annual conference of the cognitive science society* (pp. 3145–3150). Austin, TX.
- Gureckis, T. M., & Markant, D. B. (2012). Self-directed learning a cognitive and computational perspective. *Perspectives on Psychological Science*, 7(5), 464–481.
- Gweon, H., Palton, H., Konopka, J., & Schulz, L. (2014). Sins of omission: Children selectively explore when teachers are underinformative. *Cognition*, 132, 335–341.
- Gweon, H., Tenenbaum, J. B., & Schulz, L. (2010). Infants consider both the sample and the sampling process in inductive generalization. *Proceedings of the National Academy of Sciences*, 107(20), 9066–9071.
- Hamrick, J., Smith, K., Griffiths, T., & Vul, E. (2015). Think again? the amount of mental simulation tracks uncertainty in the outcome. In R. D. Noelle, A. S. Warlaumont, J. Yoshimi, T. Matlock, C. D. Jennings, & P. P. Maglio (Eds.), *Proceedings of the 37th annual conference of the cognitive science society*. Austin, TX: Cognitive Science Society.
- Hawkins, R. X., Stuhlmüller, A., Degen, J., & Goodman, N. D. (2015). Why do you ask? good questions provoke informative answers. In R. D. Noelle, A. S. Warlaumont, J. Yoshimi, T. Matlock, C. D. Jennings, & P. P. Maglio (Eds.), *Proceedings of the 37th annual conference of the cognitive science society*. austin, tx: Cognitive science society. Austin, TX: Cognitive Science Society.
- Hayes, B. K., Hawkins, G. E., & Newell, B. R.

- (2015). Consider the alternative: The effects of causal knowledge on representing and using alternative hypotheses in judgments under uncertainty. *Journal of experimental psychology. Learning, memory, and cognition*.
- Hendrickson, A. T., Navarro, D. J., & Perfors, A. (2016). Sensitivity to hypothesis size during information search. *Decision*, 3(1), 62.
- Hoch, J., O'Grady, S., & Adolph, K. (in review). It's the journey, not the destination: Locomotor exploration in infants.
- Hoch, J., Rachwani, J., & Adolph, K. E. (in review). Why do infants move? locomotor exploration in crawling and walking infants.
- Hodgins, H. S., & Zuckerman, M. (1993). Beyond selecting information: Biases in spontaneous questions and resultant conclusions. *Journal of Experimental Social Psychology*, 29(5), 387–407.
- Horwich, P. (1982). *Probability and evidence*. CUP Archive.
- Huys, Q. J. M., Eshel, N., O'Nions, E., Sheridan, L., Dayan, P., & Roiser, J. P. (2012). Bonsai trees in your head: How the pavlovian system sculpts goal-directed choices by pruning decision trees. *PLoS Computational Biology*, 8(3), e1002410.
- Inhelder, B., & Piaget, J. (1958). *The growth of logical thinking*. New York: Basic Books.
- Itti, L., & Baldi, P. (2005). A principled approach to detecting surprising events in video. In *Computer vision and pattern recognition, 2005. cvpr 2005. ieee computer society conference on* (Vol. 1, pp. 631–637).
- Itti, L., & Baldi, P. (2006). Bayesian surprise attracts human attention. In B. Weiss, J. Schoelkopf, & Platt (Eds.), *Advances in neural information processing systems* (Vol. 18, pp. 547–554).
- Jain, U., Zhang, Z., & Schwing, A. (2017). Creativity: Generating diverse questions using variational autoencoders. *arXiv preprint arXiv:1704.03493*.
- Jones, M., & Love, B. C. (2011). Bayesian fundamentalism or enlightenment? on the explanatory status and theoretical contributions of bayesian models of cognition. *Behavioral and Brain Sciences*, 34(04), 169–188.
- Kachergis, G., Rhodes, M., & Gureckis, T. (2017). Desirable difficulties in the development of active inquiry skills. *Cognition*, 166, 407–417.
- Kachergis, G., Rhodes, M., & Gureckis, T. M. (2016). Desirable difficulties in the development of active inquiry skills. In *Proceedings of the 38th Annual Conference of the Cognitive Science Society*.
- Kahneman, D., Slovic, P., & Tversky, A. (1982). *Judgment under uncertainty: Heuristics and biases*. Cambridge: Cambridge University Press.
- Kahneman, D., & Tversky, A. (1973). On the psychology of prediction. *Psychological review*, 80(4), 237.
- Kang, M. J., Hsu, M., Krajbich, I. M., Loewenstein, G., McClure, S. M., Wang, J. T.-y., & Camerer, C. F. (2009). The wick in the candle of learning epistemic curiosity activates reward circuitry and enhances memory. *Psychological Science*, 20(8), 963–973.
- Kelemen, D., & Rosset, E. (2009). The human function compunction: Teleological explanation in adults. *Cognition*, 111(1), 138–142.
- Kidd, C., & Hayden, B. Y. (2015). The psychology and neuroscience of curiosity. *Neuron*, 88(3), 449–460.
- Kidd, C., Piantadosi, S. T., & Aslin, R. N. (2012). The goldilocks effect: Human infants allocate attention to visual sequences that are neither too simple nor too complex. *PloS one*, 7(5), e36399.
- Kidd, C., Piantadosi, S. T., & Aslin, R. N. (2014). The goldilocks effect in infant auditory attention. *Child development*, 85(5), 1795–1804.
- Kim, W., Pitt, M. A., Lu, Z.-L., Steyvers, M., &

- Myung, J. I. (2014). A hierarchical adaptive approach to optimal experimental design. *Neural computation*.
- Kirkpatrick, S., Gelatt, C., & Vecchi, M. (1983). Optimization by simulated annealing. *Science*, 220, 671–680.
- Klahr, D., Fay, A. L., & Dunbar, K. (1993). Heuristics for scientific experimentation: A developmental study. *Cognitive psychology*, 25(1), 111–146.
- Klahr, D., & Nigam, M. (2004). The equivalence of learning paths in early science instruction effects of direct instruction and discovery learning. *Psychological Science*, 15(10), 661–667.
- Klayman, J. (1995). Varieties of confirmation bias. *Psychology of learning and motivation*, 32, 385–418.
- Klayman, J., & Ha, Y.-W. (1987). Confirmation, disconfirmation, and information in hypothesis testing. *Psychological review*, 94(2), 211.
- Klayman, J., & Ha, Y.-w. (1989). Hypothesis testing in rule discovery: Strategy, structure, and content. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 15(4), 596.
- Krynski, T. R., & Tenenbaum, J. B. (2007). The role of causality in judgment under uncertainty. *Journal of Experimental Psychology: General*, 136(3), 430.
- Kuhn, D., Black, J., Keselman, A., & Kaplan, D. (2000). The development of cognitive skills to support inquiry learning. *Cognition and Instruction*, 18(4), 495–523.
- Kuhn, D., Garcia-Mila, M., Zohar, A., Andersen, C., White, S. H., Klahr, D., & Carver, S. M. (1995). Strategies of knowledge acquisition. *Monographs of the society for research in child development*, i–157.
- Kushnir, T., & Gopnik, A. (2005). Young children infer causal strength from probabilities and interventions. *Psychological Science*, 16(9), 678–683.
- Kushnir, T., Wellman, H. M., & Gelman, S. A. (2008). The role of preschoolers’s social understanding in evaluating the informativeness of causal interventions. *Cognition*, 107(3), 1084–1092.
- Lagnado, D. A., & Sloman, S. (2004). The advantage of timely intervention. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 30(4), 856.
- Legare, C. H. (2012). Exploring explanation: Explaining inconsistent evidence informs exploratory, hypothesis-testing behavior in young children. *Child Development*, 83(1), 173–185.
- Le Mens, G., & Denrell, J. (2011). Rational learning and information sampling: On the ‘naivety’ assumption in sampling explanations of judgment biases. *Psychological review*, 118(2), 379.
- Lewicki, M. S. (2002). Efficient coding of natural sounds. *Nature neuroscience*, 5(4), 356.
- Lewis, D. (1969). *Convention: A philosophical study*. Cambridge: Harvard University Press.
- Lieder, F., Plunkett, D., Hamrick, J. B., Russell, S. J., Hay, N., & Griffiths, T. (2014). Algorithm selection by rational metareasoning as a model of human strategy selection. In *Advances in neural information processing systems* (pp. 2870–2878).
- Lindley, D. V. (1956). On a measure of the information provided by an experiment. *The Annals of Mathematical Statistics*, 986–1005.
- Loewenstein, G. (1994). The psychology of curiosity: A review and reinterpretation. *Psychological bulletin*, 116(1), 75.
- Lomasky, R., Brodley, C. E., Aernecke, M., Walt, D., & Friedl, M. (2007). Active class selection. In *Machine learning: Ecml 2007* (pp. 640–647). Springer.
- Lombrozo, T., & Carey, S. (2006). Functional explanation and the function of explanation. *Cognition*, 99(2), 167–204.
- MacDonald, K., & Frank, M. C. (2016). When

- does passive learning improve the effectiveness of active learning? In A. Papafragou, D. Grodner, D. Mirman, & J. Trueswell (Eds.), *Proceedings of the 38th annual conference of the cognitive science society*. Austin, TX.
- Mackay, D. (1992). Information-based objective functions for active data selection. *Neural Computation*, 4, 590-604.
- Malt, B. C., Ross, B. H., & Murphy, G. L. (1995). Predicting features for members of natural categories when categorization is uncertain. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 21(3), 646.
- Mandler, J. M. (2014). *Stories, scripts, and scenes: Aspects of schema theory*. Psychology Press.
- Marcus, G. F., & Davis, E. (2013). How robust are probabilistic models of higher-level cognition? *Psychological science*, 24(12), 2351-2360.
- Marewski, J. N., & Schooler, L. J. (2011). Cognitive niches: an ecological model of strategy selection. *Psychological review*, 118(3), 393.
- Markant, D. B. (2016). The impact of biased hypothesis generation on self-directed learning. In A. Papafragou, D. Grodner, D. Mirman, & J. Trueswell (Eds.), *Proceedings of the 38th annual conference of the cognitive science society*. Austin, TX: Cognitive Science Society.
- Markant, D. B., & Gureckis, T. M. (2012). Does the utility of information influence sampling behavior? In N. Miyake, D. Peebles, & R. P. Cooper (Eds.), *Proceedings of the 34th annual conference of the cognitive science society* (pp. 719-724). Austin, TX.
- Markant, D. B., & Gureckis, T. M. (2014). Is it better to select or to receive? learning via active and passive hypothesis testing. *Journal of Experimental Psychology-General*, 143(1), 94-122.
- Markant, D. B., Settles, B., & Gureckis, T. M. (2015). Self-directed learning favors local, rather than global, uncertainty. *Cognitive Science*.
- Marvin, C., & Shohamy, D. (2016). Curiosity and reward: Valence predicts choice and information prediction errors enhance learning. *Journal of experimental psychology. General*.
- McCormack, T., Bramley, N. R., Frosch, C., Patrick, F., & Lagnado, D. (2016). Children's use of interventions to learn causal structure. *Journal of experimental child psychology*, 141, 1-22.
- McKenzie, C. R., Ferreira, V. S., Mikkelsen, L. A., McDermott, K. J., & Skrabbe, R. P. (2001). Do conditional hypotheses target rare events? *Organizational Behavior and Human Decision Processes*, 85(2), 291-309.
- Meder, B., & Nelson, J. D. (2012). Information search with situation-specific reward functions. *Judgment and Decision Making*, 7(2), 119-148.
- Meltzoff, A. N. (1995). Understanding the intentions of others: re-enactment of intended acts by 18-month-old children. *Developmental psychology*, 31(5), 838.
- Metcalf, J., & Kornell, N. (2003). The dynamics of learning and allocation of study time to a region of proximal learning. *Journal of Experimental Psychology: General*, 132(4), 530.
- Minsky, M. (1974). A framework for representing knowledge. *MIT-AI Laboratory Memo 306*.
- Miyake, N., & Norman, D. (1979). To ask a question one must know enough to know what is not known. *Journal of Verbal Learning and Verbal Behavior*, 18, 357-364.
- Montessori, M. (1912/1964). *The montessori method*. New York, NY: Schocken.
- Mosher, F. A., & Hornsby, J. R. (1966). Studies in cognitive growth. In (chap. On asking questions). New York, NY: Wiley.

- Mozer, M., Pashler, H., & Homaei, H. (2008). Optimal predictions in everyday cognition: The wisdom of individuals or crowds? *Cognitive Science*, 32(7), 1133-1147.
- Muliere, P., & Parmigiani, G. (1993, November). Utility and means in the 1930s. *Statistical Science*, 8(4), 421-432.
- Murphy, G. L., Chen, S. Y., & Ross, B. H. (2012). Reasoning with uncertain categories. *Thinking & Reasoning*, 18(1), 81-117.
- Murphy, K. P. (2001). Active learning of causal bayes net structure. *Technical Report. Department of Computer Science, U.C. Berkeley.*
- Myung, J. I., & Pitt, M. A. (2009). Optimal experimental design for model discrimination. *Psychological review*, 116(3), 499.
- Najemnik, J., & Geisler, W. S. (2005). Optimal eye movement strategies in visual search. *Nature*, 434(7031), 387-391.
- Najemnik, J., & Geisler, W. S. (2009). Simple summation rule for optimal fixation selection in visual search. *Vision Research*, 49, 1286-1294.
- Navarro, D. J., & Perfors, A. F. (2011). Hypothesis generation, sparse categories, and the positive test strategy. *Psychological Review*, 118(1), 120.
- Nelson, J. D. (2005). Finding useful questions: on bayesian diagnosticity, probability, impact, and information gain. *Psychological Review*, 112(4), 979.
- Nelson, J. D., Divjak, B., Gudmundsdottir, G., Martignon, L. F., & Meder, B. (2014). Children's sequential information search is sensitive to environmental probabilities. *Cognition*, 130(1), 74-80.
- Nelson, J. D., McKenzie, C. R., Cottrell, G. W., & Sejnowski, T. J. (2010). Experience matters: information acquisition optimizes probability gain. *Psychological Science*, 21(7), 960-969.
- Nelson, J. D., Tenenbaum, J. B., & Movellan, J. R. (2001). Active inference in concept learning. In J. D. Moore & K. Stenning (Eds.), *Proceedings of the 23rd conference of the cognitive science society* (p. 692-697). Austin, TX.
- Nelson, K. (1973). Structure and strategy in learning to talk. *Monographs of the Society for Research in Child Development*, 38(1-2, Serial No. 149), 1-135.
- Nickerson, R. S. (1998). Confirmation bias: A ubiquitous phenomenon in many guises. *Review of General Psychology*, 2(2), 175.
- Nosofsky, R. M., & Palmeri, T. J. (1998). A rule-plus-exception model for classifying objects in continuous-dimension spaces. *Psychonomic Bulletin & Review*, 5(3), 345-369.
- Oaksford, M., & Chater, N. (1994). A rational analysis of the selection task as optimal data selection. *Psychological Review*, 101(4), 608.
- Oaksford, M., & Chater, N. (1996). Rational explanation of the selection task. *Psychological Review*, 103(2), 381-391.
- Oaksford, M., Chater, N., Grainger, B., & Larkin, J. (1997). Optimal data selection in the reduced array selection task (rast). *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 23(2), 441.
- O'Brien, B., & Ellsworth, P. C. (2006). Confirmation bias in criminal investigations. In *1st annual conference on empirical legal studies paper*.
- Otto, A. R., Raio, C. M., Chiang, A., Phelps, E. A., & Daw, N. D. (2013). Working-memory capacity protects model-based learning from stress. *Proceedings of the National Academy of Sciences*, 110(52), 20941-20946.
- Oudeyer, P.-Y., Gottlieb, J., & Lopes, M. (2016). Intrinsic motivation, curiosity, and learning: Theory and applications in educational technologies. *Progress in brain research*, 229, 257-284.
- Oudeyer, P.-Y., Kaplan, F., & Hafner, V. V. (2007). Intrinsic motivation systems for autonomous mental development. *IEEE transactions on*

- evolutionary computation*, 11(2), 265–286.
- Pauker, S. G., & Kassirer, J. P. (1980). The threshold approach to clinical decision making. *New England Journal of Medicine*, 302, 1109–1117.
- Pearl, J. (2009). *Causality*. Cambridge university press.
- Phillips, L. D., & Edwards, W. (1966). Conservatism in a simple probability inference task. *Journal of experimental psychology*, 72(3), 346.
- Popper, K. R. (1968). *Logik der forschung: zur erkenntnistheorie der moderner naturwissenschaft*. Mohr Siebeck.
- Pothos, E., & Chater, N. (2005). Unsupervised categorization and category learning. *The Quarterly Journal of Experimental Psychology*, 58A(4), 733–752.
- Rafferty, A. N., Zaharia, M., & Griffiths, T. L. (2014). Optimally designing games for behavioural research. *Proceedings of the Royal Society A*, 470(2167), 20130828.
- Raghavan, H., Madani, O., & Jones, R. (2006). Active learning with feedback on features and instances. *The Journal of Machine Learning Research*, 7, 1655–1686.
- Raiffa, H., & Schlaifer, R. O. (1961). *Applied statistical decision theory*. New York: Wiley.
- Redgrave, P., & Gurney, K. (2006). The short-latency dopamine signal: a role in discovering novel actions? *Nature reviews neuroscience*, 7(12), 967–975.
- Ren, M., Kiros, R., & Zemel, R. (2015). Exploring models and data for image question answering. In *Advances in neural information processing systems* (pp. 2953–2961).
- Rich, A. S., & Gureckis, T. M. (2014). The value of approaching bad things. In *Proceedings of the 36rd annual conference of the cognitive science society. austin, tx: Cognitive science society*.
- Rich, A. S., & Gureckis, T. M. (2017). Exploratory choice reflects the future value of information. *Decision*.
- Rieskamp, J., & Otto, P. E. (2006). Ssl: a theory of how people learn to select strategies. *Journal of Experimental Psychology: General*, 135(2), 207.
- Ross, B. H., & Murphy, G. L. (1996). Category-based predictions: influence of uncertainty and feature associations. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 22(3), 736.
- Rothe, A., Lake, B., & Gureckis, T. (in review). Do people ask good questions? *Computational Brain and Behavior*.
- Rothe, A., Lake, B. M., & Gureckis, T. M. (2016). Asking and evaluating natural language questions. In *Proceedings of the 38th annual conference of the cognitive science society*.
- Ruderman, D. L. (1994). Designing receptive fields for highest fidelity. *Network: Computation in neural systems*, 5(2), 147–155.
- Ruggeri, A., & Lombrozo, T. (2015). Children adapt their questions to achieve efficient search. *Cognition*, 143, 203–216.
- Ruggeri, A., Lombrozo, T., Griffiths, T., & Xu, F. (2015). Children search for information as efficiently as adults, but seek additional confirmatory evidence. In D. C. Noelle et al. (Eds.), *Proceedings of cogsci 37*.
- Rusconi, P., Marelli, M., D’Addario, M., Russo, S., & Cherubini, P. (2014). Evidence evaluation: Measure z corresponds to human utility judgments better than measure l and optimal-experimental-design models. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 40(3), 703–723.
- Sanborn, A. N., Griffiths, T. L., & Navarro, D. J. (2010). Rational approximations to rational models: alternative algorithms for category learning. *Psychological Review*, 117(4), 1144.
- Schulz, L. (2012a). Finding new facts; thinking new thoughts. *Advances in Child Development and Behavior*, 43, 269–294.
- Schulz, L. (2012b). The origins of inquiry: Induc-

- tive inference and exploration in early childhood. *Trends in Cognitive Sciences*, 16(7), 382–389.
- Schulz, L. (2015). Infants explore the unexpected. *Science*, 348, 42–43.
- Schulz, L., & Bonawitz, E. B. (2007). Serious fun: preschoolers engage in more exploratory play when evidence is confounded. *Developmental psychology*, 43(4), 1045.
- Schulz, L., Gopnik, A., & Glymour, C. (2007). Preschool children learn about causal structure from conditional interventions. *Developmental Science*, 10(3), 322–332.
- Schulz, L., Kushnir, T., & Gopnik, A. (2007). Learning from doing: Interventions and causal inference. In A. Gopnik & L. Schulz (Eds.), *Causal learning: Psychology, philosophy, and computation*. Oxford University Press.
- Settles, B. (2010). Active learning literature survey. *University of Wisconsin, Madison*.
- Shafto, P., Goodman, N. D., & Frank, M. C. (2012). Learning from others the consequences of psychological reasoning for human learning. *Perspectives on Psychological Science*, 7(4), 341–351.
- Shafto, P., Goodman, N. D., & Griffiths, T. L. (2014). A rational account of pedagogical reasoning: Teaching by, and learning from, examples. *Cognitive Psychology*, 71, 55–89.
- Shi, L., Griffiths, T., Feldman, N., & Sanborn, A. (2010). Exemplar models as a mechanism for performing bayesian inference. *Psychonomic Bulletin & Review*, 17(4), 443–464.
- Siegel, M., Magin, R., Tenenbaum, J., & Schulz, L. (2014). Black boxes: Hypothesis testing via indirect perceptual evidence. In A. Papafragou, D. Grodner, D. Mirman, & J. Trueswell (Eds.), *Proceedings of the 36th Annual Conference of the Cognitive Science Society*. Austin, TX: Cognitive Science Society.
- Sim, Z., Tanner, M., Alpert, N., & Xu, F. (2015). Children learn better when they select their own data. In *Proceedings of the 34th annual conference of the cognitive science society*. Austin, TX.
- Simon, H. A. (1976). From substantive to procedural rationality. In *25 years of economic theory* (pp. 65–86). Springer.
- Singh, S. P., Barto, A. G., & Chentanez, N. (2004). Intrinsically motivated reinforcement learning. In *Nips* (Vol. 17, pp. 1281–1288).
- Skov, R. B., & Sherman, S. J. (1986). Information-gathering processes: Diagnosticity, hypothesis-confirmatory strategies, and perceived hypothesis confirmation. *Journal of Experimental Social Psychology*, 22(2), 93–121.
- Slowiaczek, L. M., Klayman, J., Sherman, S. J., & Skov, R. B. (1992). Information selection and use in hypothesis testing: What is a good question, and what is a good answer? *Memory & Cognition*, 20(4), 392–405.
- Spelke, E. S., & Kinzler, K. D. (2007). Core knowledge. *Developmental science*, 10(1), 89–96.
- Stahl, A. E., & Feigenson, L. (2015). Observing the unexpected enhances infants learning and exploration. *Science*, 348(6230), 91–94.
- Steyvers, M., Tenenbaum, J. B., Wagenmakers, E.-J., & Blum, B. (2003). Inferring causal networks from observations and interventions. *Cognitive Science*, 27(3), 453–489.
- Sutton, R. S., & Barto, A. G. (1988). *Reinforcement learning: An introduction*. MIT Press.
- Tenenbaum, J. B. (1999). *A bayesian framework for concept learning* (Unpublished doctoral dissertation). MIT.
- Tenenbaum, J. B., & Griffiths, T. L. (2001). Generalization, similarity, and bayesian inference. *Behavioral and brain sciences*, 24(04), 629–640.
- Tenenbaum, J. B., Griffiths, T. L., & Kemp, C. (2006). Theory-based bayesian models of inductive learning and reasoning. *Trends in*

- cognitive sciences*, 10(7), 309–318.
- Thomas, R. P., Dougherty, M. R. P., Sprenger, A. M., & Harbison, J. I. (2008). Diagnostic hypothesis generation and human judgment. *Psychological Review*, 115, 155–185.
- Trueswell, J. C., Medina, T. N., Hafri, A., & Gleitman, L. R. (2013). Propose but verify: Fast mapping meets cross-situational word learning. *Cognitive psychology*, 66(1), 126–156.
- Tschirgi, J. E. (1980). Sensible reasoning: A hypothesis about hypotheses. *Child development*, 1–10.
- Tversky, A., & Edwards, W. (1966). Information versus reward in binary choices. *Journal of Experimental Psychology*, 71(5), 680.
- Ullman, S., Vidal-Naquet, M., & Sali, E. (2002). Visual features of intermediate complexity and their use in classification. *Nature neuroscience*, 5(7), 682.
- U.S. Department of Education. (2017). *Reimagining the Role of Technology in Education: 2017 National Education Technology Plan Update* (Tech. Rep.). Office of Educational Technology.
- van Schijndel, T., Visser, I., van Bers, B., & Raijmakers, M. (2015). Preschoolers perform more informative experiments after observing theory-violating evidence. *Journal of Experimental Child Psychology*, 131, 104–119.
- Vul, E., Goodman, N. D., Griffiths, T. L., & Tenenbaum, J. B. (2014). One and done? Optimal decisions from very few samples. *Cognitive Science*, 38(4), 599–637. doi: 10.1111/cogs.12101
- Vygotsky, L. (1962). *Thought and language*. Cambridge, MA: MIT Press.
- Waldmann, M. R., & Hagmayer, Y. (2005). Seeing versus doing: two modes of accessing causal knowledge. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 31(2), 216.
- Wason, P. C. (1960). On the failure to eliminate hypotheses in a conceptual task. *Quarterly journal of experimental psychology*, 12(3), 129–140.
- Wason, P. C. (1966). Reasoning. In B. Foss (Ed.), *New horizons in psychology*. Harmondsworth: Penguin Books.
- Woodward, A. L. (1998). Infants selectively encode the goal object of an actor's reach. *Cognition*, 69(1), 1–34.
- Xu, F., & Tenenbaum, J. B. (2007). Word learning as bayesian inference. *Psychological review*, 114(2), 245.
- Zhang, L., Tong, M. H., Marks, T. K., Shan, H., & Cottrell, G. W. (2008). Sun: A bayesian framework for saliency using natural statistics. *Journal of vision*, 8(7), 32–32.