A Complete Method for Assessing the Effectiveness of Eyewitness Identification

Procedures: Expected Information Gain

Jeffrey J. Starns, Andrew L. Cohen, and Caren M. Rotello

Department of Psychological and Brain Sciences

University of Massachusetts, Amherst

Author Note:

Preparation of this article was supported by Grant No. 1454868 from the National Science Foundation.

Correspondence concerning this article should be addressed to Jeffrey Starns, Department of Psychological and Brain Sciences, University of Massachusetts, Amherst, Tobin Hall, 135 Hicks Way, Amherst, MA 01003, United States. Email: jstarns@umass.edu.

Abstract

We present a method for measuring the efficacy of eyewitness identification procedures by applying fundamental principles of information theory. The resulting measure evaluates the Expected Information Gain (EIG) for an identification attempt, a single value that summarizes an identification procedure's overall potential for reducing uncertainty about guilt or innocence across all possible witness responses. In a series of demonstrations, we show that EIG often disagrees with existing measures (e.g., diagnosticity ratios or area under the ROC) about the relative effectiveness of different identification procedures. Each demonstration is designed to highlight key distinctions between existing measures and EIG. An overarching theme is that EIG provides a complete measure of evidentiary value, in the sense that it factors in all aspects of identification performance. Collectively, these demonstrations show that EIG has substantial potential to inspire new discoveries in eyewitness research and provide a new perspective on policy recommendations for the use of identifications in real investigations.

Keywords: Eyewitness identification, lineup, showup, Shannon information, information theory

A Complete Method for Assessing the Effectiveness of Eyewitness Identification

Procedures: Expected Information Gain

Many critical functions of the criminal justice system hinge on the ability to perform a fundamental decision task: Judging which suspects are guilty and which are innocent. A perfect system would be capable of determining a suspect's guilt or innocence with 100% accuracy. In reality, of course, decision makers in the criminal justice system must cope with many sources of uncertainty and make the best judgments possible given the available evidence. This challenge applies in obvious cases, such as a jury deciding whether or not to convict a defendant, but it is also manifest at many other decision points, such as a judge deciding whether the current evidence against a suspect justifies issuing a search warrant or a detective deciding whether the evidence against one suspect is sufficient to terminate investigation of other suspects. In scenarios like these, decision makers can achieve better outcomes if they have access to better evidence, that is, evidence that supports a higher level of certainty about the guilt or innocence of the suspect. Thus, developing principled methods for assessing the effectiveness of an evidence-gathering procedure is a critical goal for effective research and practice. This goal goes beyond the question of how to make optimal use of evidence once it becomes available and necessitates a measure that can be used to compare different policies for *obtaining new evidence*.

One well-studied source of evidence in criminal investigations is an eyewitness identification of the suspect (Wells & Olson, 2003). Identification attempts come in a wide variety of formats (National Research Council, 2015; Police Executive Research Forum, 2013). In the simplest form, the suspect is presented to the witness – either in person or in a photograph – and the witness indicates whether or not the suspect is the person they saw commit the crime. This procedure is known as a *showup*, a name inspired by the practice of apprehending a suspect

soon after a crime and “showing up” at the crime scene to present the suspect to a witness who is still being interviewed. Another common procedure is to present a suspect along with *fillers* who are known to be innocent, creating a *lineup*. For the showup procedure, a witness can either affirm that the suspect is the culprit (a *suspect identification* or suspect ID) or fail to do so (a *rejection*). Lineups have an additional possible outcome, as the witness might claim that one of the fillers was the culprit (a *filler identification* or filler ID). These basic witness responses are sometimes supplemented with ancillary information, such as the witness’ level of confidence in their choice (Dallenbach, 1913; Deffenbacher, 1980) or the response time (RT) associated with the decision (Smith, Lindsay, & Pryke, 2000; Sporer, 1993).

The ultimate goal of an identification attempt is to use the witness response to judge whether or not the suspect is guilty (Clark, 2012; Wells & Lindsay, 1980; Wells, Yang, & Smalarz, 2015). Decades of eyewitness memory research demonstrates that the success of this enterprise varies widely across situations; for example, positive identifications isolate guilty suspects with extreme reliability in some circumstances (Seale-Carlisle, Wetmore, Flowe, & Mickes, 2019), but are very likely to promote false convictions of innocent suspects in other circumstances (Gronlund, Carlson, Dailey, & Goodsell, 2009).

The factors affecting the effectiveness of identification procedures include *system* variables that are under the control of the identification administrator and *estimator* variables that cannot be controlled, but can potentially be assessed to aid in the interpretation of witness responses (Cutler, Penrod, & Martens, 1987; Fulero, 2009; Wells, 1978; Wells & Olson, 2003). Eyewitness researchers have explored a wide range of system variables, such as whether the suspect is presented to the witness alone (i.e., in a showup) or as a member of a lineup (e.g., Wetmore et al., 2015), whether lineup members are presented simultaneously or sequentially

(e.g., Mickes, Flowe, & Wixted, 2012), whether the witness is encouraged to make an identification (e.g., Steblay, 1997), and many others. Researchers have also thoroughly tested many estimator variables such as the viewing conditions of the witnessed crime (e.g., Lindsay, Semmler, Weber, Brewer, & Lindsay, 2008; Loftus, 1985; Shapiro & Penrod, 1986), whether the race of the witness is the same as that of the suspect (Malpass & Kravitz, 1969), and the level of confidence reported by the witness (Dallenbach, 1913).

Given the huge range of effectiveness across different identification procedures, researchers need to be able to assess a given procedure's potential to inform judgments about guilt or innocence. Researchers have compared the effectiveness of identification procedures using a range of quantitative measures such as the change in the probability of guilt from before to after an identification (Wells & Lindsay, 1980), diagnosticity ratios (Clark, 2012), area under the receiver operating characteristic (ROC; Mickes et al., 2012), deviation from perfect performance (Smith, Lampinen, Wells, Smalarz, & Mackovichova, 2019), and expected utility (Ceci & Freidman, 2000). The relative merits of these measures are a matter of continuing debate (Lampinen, 2016; Lampinen, Smith, & Wells, 2019; Rotello & Chen, 2016; Smith, Wells, Smalarz, & Lampinen, 2018; Wixted & Mickes, 2012; Wixted, Mickes, Wetmore, Gronlund, & Neuschatz, 2017).

This issue is not merely academic; eyewitness research influences policies adopted by individual investigative units or recommended by government agencies (National Institute of Justice, 1999; National Research Council, 2014), and sometimes different measures support different policy recommendations (e.g., Wixted & Mickes, 2012). Thus, developing and improving measures of the evidentiary value of an identification procedure is critical for driving future progress in eyewitness research and ensuring that policy makers have the guidance they

need to select identification procedures that provide the best possible information about guilt or innocence.

The Current Approach

The goal of this work is to introduce a new measure of evidentiary value that combines all aspects of identification performance to characterize the full evidentiary potential of an identification procedure. Details of the approach are provided below; in this section, we provide a high-level overview.

The approach we outline here is closely related to past work using Bayesian inference to evaluate identification procedures (Wells & Lindsay, 1980; Wells & Olson, 2002; Wells & Turtle, 1986; Wells, Yang, & Smalarz, 2015). The key insight of the Bayesian approach is that the evidentiary value of a witness response can be measured by contrasting the *prior* probability of guilt (before the identification attempt) with the *posterior* probability of guilt (after the attempt). Bayes' theorem defines the shift from prior to posterior based on the likelihood of obtaining the observed witness response for guilty suspects and for innocent suspects (Wells & Lindsay, 1980).

Although the Bayesian approach has supported important discoveries, it has a critical limitation: it is used to assess the evidentiary value of a single witness response at a time, e.g., a suspect ID. The focus on single responses offers no way to resolve commonly observed tradeoffs that highlight offsetting strengths and weaknesses of different lineup procedures, e.g., whereas a suspect ID may support higher certainty of guilt for Procedure A, a rejection may support higher certainty of innocence for Procedure B (Wells, Yang, & Smalarz, 2015; also see Clark, 2012). Moreover, considering responses in isolation is not an effective way to compare procedures that have non-corresponding response options, such as different methods for obtaining confidence

ratings or different ways to use variables that predict witness accuracy (Brewer, Weber, & Guerin, 2020). Finally, a consideration of tradeoffs is necessary for a complete accounting of the overall pattern of witness confidence (National Research Council, 2014, p. 108) and both positive and negative responses, e.g., suspect IDs and rejections (National Research Council, 2014, p. 90).

The single-response limitation can be overcome by measuring the *Expected Information Gain* (EIG) of an identification procedure.¹ In short, EIG is the average amount of information about guilt or innocence that is gained by performing an identification attempt, and it is calculated by applying information-theory principles to eyewitness identification data. Rather than assessing an individual witness response within a given procedure, the resulting value evaluates the identification procedure as a whole. Like the Bayesian approach, EIG is based on the change in certainty of guilt or innocence produced by observing a witness response. A critical difference, however, is that certainty is defined by entropy instead of probability, providing a way to summarize the total evidentiary value across all responses.

Entropy is a measure of *uncertainty* given that higher values indicate less certainty. More specifically, entropy is the average Shannon surprisal for a discrete probability distribution (Shannon, 1948). Shannon surprisal is an information-theory measure of judgment error, and measuring certainty/uncertainty with a surprisal-based measure like entropy, rather than by “raw” probability, confers a number of advantages that will be discussed below (Gneiting & Raftery, 2007; McElreath, 2016). To preview, however, the primary value of entropy is linked to one critical property: entropy measures uncertainty on an equal-interval scale (Shannon, 1948),

¹ The expected information gain for random variable A produced by learning the outcome of random variable B is also commonly called the mutual information of A and B (Cover & Thomas, 2007).

meaning that a given change in uncertainty is associated with the same change in entropy throughout the entire range of the scale.

In contrast, probability is *not* an equal-interval measure of certainty. That is, the same change in certainty can have different effects on probability in different situations. For example, imagine that your friend hands you two fair coins and asks you to try to guess what will happen if you flip them both. Before flipping any coins, you can be 25% certain of the outcome of both flips (there are 4 equally likely combinations of heads, H, and tails, T,: HH, HT, TH, and TT). You then flip the first coin and it lands heads. Now you can be 50% certain of the outcome of both flips (the remaining possibilities are HH and HT), for a certainty gain of .25 on the probability scale [= .50-.25]. You then flip the second coin and it lands tails. Now you are 100% certain of the outcome of both flips, for a certainty gain of .50 on the probability scale [=1.00-.50]. In both cases, your level of certainty changed by the same amount because you learned the outcome of exactly the same event – a flip of a fair coin. The probability scale, however, represents this change in certainty differently for different starting points (high versus low starting uncertainty for the first and second coin, respectively). This example demonstrates that probability does not have an equal-interval relationship with change in certainty.

We will show that failing to use an equal-interval measure of certainty introduces averaging artifacts, and that entropy provides a useful way to avoid these artifacts. Thus, using entropy provides a way to summarize the average level of certainty across all possible witness responses, allowing researchers and policy makers to identify procedures that will produce the most information about guilt or innocence across multiple identification attempts.

The application of EIG to eyewitness lineups is consistent with the NAS recommendation to explore “other powerful tools for evaluating a ‘binary classifier’” in eyewitness research

(National Research Council, 2014, p. 4). Indeed, entropy has been an extremely influential measure across many fields of study, such as computer science (Chaitin, 1975), medicine (Benish, 1999), and engineering (Cover & Thomas, 2006). Techniques based on entropy are also beginning to play a critical role in forensic applications (Ramos & Gonzalez-Rodriguez, 2008; Ramos, Gonzalez-Rodriguez, Zadora, Zieba-Palus, & Aitken, 2007), such as automated recognition systems for faces (Gong, Li, Tao, Liu, & Li, 2015), speakers (Campbell, Reynolds, Campbell, & Brady, 2005), and fingerprints (Osterburg, Parthasarathy, Raghavan, & Sclove, 1977).

EIG indicates the extent to which a procedure tends to support high levels of certainty about guilt or innocence when all possible witness responses are considered. Thus, EIG can be used to compare procedures with non-corresponding response options. This property allows researchers to compare different methods for obtaining witness responses and explore new research questions, such as the optimal way to use variables that predict witness accuracy (Brewer et al., 2020). Indeed, we will show that EIG can be easily calculated even for continuous variables like response time (RT). Critically, the consideration of all possible witness responses also allows EIG to directly address tradeoffs across responses, meaning that it can be used to answer questions like “Does the advantage of Procedure A for suspect IDs outweigh the advantage of Procedure B for lineup rejections?” (e.g., Clark, 2012; Wells, Yang, & Smalarz, 2015).

EIG resolves tradeoffs by evaluating the average amount of information about guilt or innocence generated by identification attempts, where information is defined as a reduction in uncertainty as measured by Shannon (1948) entropy. The strategy of resolving tradeoffs by averaging across different witness responses is also adopted by the Expected Utility approach

(EU; Ceci & Freidman, 2000; Clark, 2012; Lampinen, Smith, & Wells, 2019; Malpass, 2006; Smith et al., 2019). Indeed, the EIG and EU approaches differ only in how they assign a value to an identification outcome, such as correctly identifying a guilty suspect or correctly rejecting an innocent suspect. In the EU approach, the researcher assigns a value representing the utility of each identification outcome. The utilities are meant to represent the personal and societal costs and benefits of identification outcomes. Procedures are then assessed in terms of the average utility obtained across identification attempts, i.e., the Expected Utility. Researchers who have adopted the EU approach have been very open about one of its critical limitations: there is no way to ensure that the utilities selected by the researcher accurately represent the real consequences of identification outcomes. Furthermore, different circumstances may require different utilities (e.g., deciding to issue a search warrant versus deciding to sentence a defendant to a life sentence) and changing the utilities can change conclusions about the relative effectiveness of identification procedures (Yang, Smalarz, Moody, Cabell, & Copp, 2019).

EIG builds off of the EU approach but, critically, provides a way to replace the arbitrarily assigned utilities with a principled measure of the *intrinsic* value of an identification outcome, that is, the value of an outcome *as a source of evidence about guilt or innocence*. This approach has the added benefit of assessing evidence without having to assume that identification outcomes always lead to the same consequences. The information-theory metrics used by EIG assign higher values to stronger evidence, i.e., evidence that supports higher certainty. The precise values are determined by the statistical properties of the distribution of uncertainty, not asserted by the researcher (Shannon, 1948). In short, whereas EU provides a way for researchers and policy makers to choose the identification procedures that produce the highest utility across a large number of identification attempts, EIG provides a way to choose identification procedures

that produce *the most information about guilt or innocence* across a large number of identification attempts. We argue that informational value is important to consider, especially when there is a questionable link between a chosen set of utilities and the realities of actual investigations. Throughout, we will limit our focus to measures of informational value, as we are interested in evaluating identification outcomes as a source of evidence.

EIG and Policy

In the current work, we mainly present EIG as tool for researchers and policy makers who are considering different identification procedures. Clark (2012) recommends that policy makers consider a variety of measures to evaluate alternative identification procedures. We agree with the wisdom of this approach, as each measure is best suited to a particular policy goal. EIG is tailored to the goal of obtaining the most possible information about guilt or innocence across a set of identification attempts in which witness responses vary from one attempt to the next, and the *EIG in the Justice System* section below outlines a rationale for why policy makers should value this goal. The crux of our argument is that EIG provides a *complete* measure of identification effectiveness; that is, a measure that integrates every aspect of identification performance (see the *Complete versus Incomplete Measures* section). Many existing measures are not complete in this sense, as changing certain aspects of identification performance has no effect on the measure. We acknowledge that isolating individual factors can be illuminating, but it is also useful to have a way to put all the pieces together and make a comprehensive judgment. We will show that considering all aspects of identification performance can change conclusions about the relative effectiveness of different procedures, and we present targeted demonstrations to illustrate this point.

Overview

In what follows, we explain the information-theory principles that motivate the EIG measure and present a series of demonstrations designed to help readers understand the properties of EIG and distinguish it from extant measures of evidentiary value. Our demonstrations focus on ways that EIG can both redefine interpretations of existing results and inspire important new directions in eyewitness memory research. We discuss a number of concepts, some of which are fairly complex, so this section previews the general structure and some of the major results.

The *Information Theory and Eyewitness Identifications* section explains the theoretical basis for the EIG measure, which combines a Bayesian approach to interpreting witness responses (Wells & Lindsay, 1980) with foundational measures from information theory (Shannon, 1948). Using data from Colloff and Wixted (2020), we (1) explain that all measures of evidentiary value are based on the statistics of a reference population; (2) demonstrate that EIG is a *complete* measure in the sense that it uses all information from the reference population; (3) review the process of using Bayesian inference to guide probabilistic judgments of guilt or innocence; (4) explain how Shannon surprisal provides a measure of the degree of error in probabilistic judgments; (5) introduce the concept of *entropy*, the average surprisal across a probability distribution, and explain how it quantifies uncertainty about guilt or innocence on an equal-interval scale; and (6) show how the change in entropy produced by witness responses can be used to calculate EIG.

This section also presents examples that explore some of the core theoretical properties of EIG, including (1) an example that demonstrates the importance of having an equal-interval scale of uncertainty by showing that averaging probabilities produces misleading results, (2) an

example that explores the effect of uneven base rates, and (3) an example that explores the effect of miscalibration in the interpretation of identification evidence.

The *Demonstrations Exploring EIG* section presents a series of example applications designed to illustrate key properties of EIG. A common theme throughout the demonstrations is that EIG considers all factors relevant to the evidentiary value of an identification procedure, and we highlight situations in which EIG's comprehensive assessment supports conclusions that differ from existing measures. We also provide examples of how EIG can inspire new lines of inquiry in eyewitness memory research. The goal of the demonstrations is to help readers understand the implications of using EIG as a measure of the effectiveness of an identification procedure, not to make empirical conclusions about particular issues. Accordingly, we generally use well-validated decision models to generate example data sets for the demonstrations, which allows us to ensure that we are working with realistic data patterns while also isolating the key processes that distinguish EIG from existing measures. We also consider some targeted data sets from eyewitness experiments, but these were selected to highlight key differences between EIG and the measures used in the original papers, not to provide a comprehensive assessment of the associated research questions. Because the task of using EIG to reevaluate the eyewitness memory literature cannot be achieved in a single paper, we focus instead on more limited goals; namely, demonstrating that EIG provides a new perspective on many central issues and exploring the theoretical properties of this measure that distinguish it from existing approaches.

We report six demonstrations in total. The first demonstration contrasts EIG and diagnosticity ratios to explore the impact of *liberal* or *conservative* biases in witness responses, where liberal responding means that witnesses are willing to make an identification even if they are not too confident that they are selecting the culprit and conservative responding means that

witnesses are reluctant to make an identification unless they are very confident that they are selecting the culprit. Whereas diagnosticity ratios increase with increasingly conservative responding (e.g., Rotello & Chen, 2016; Wixted & Mickes, 2012), EIG reveals that any bias in responding, whether liberal or conservative, limits the evidentiary value of witness responses. This demonstration also explores a scenario in which the procedure that optimizes EIG is affected by the prior probability of guilt. Reassuringly, simulation results show that optimizing EIG under the assumption of a 50% prior probability of guilt identifies near-optimal procedures for *any* actual prior probability. Thus, researchers can identify useful identification procedures even when the prior probability of guilt varies widely across different identification attempts, as in real investigations.

The second demonstration contrasts EIG with area under the Receiver Operating Characteristic (ROC; Macmillan & Creelman, 2005; Mickes et al., 2012). The key conclusion from this section is that these two measures are designed to measure subtly different constructs: Whereas area under the ROC measures how well witnesses can discriminate guilty and innocent suspects based on their memory for the culprit, EIG measures how well the witness *responses* can be used by an outside observer to improve assessments of a suspect's guilt or innocence (Wells, Yang, & Smalarz, 2015, also noted this distinction in comparing the ROC and Bayesian approaches). *Demonstration 2* also presents an example evaluating variation in lineup size to show that EIG and area under the ROC lead to different conclusions in some situations. This example shows how EIG can lead to new theoretical discoveries, as combining EIG with a signal-detection model of identification performance revealed a novel “filler curtain” effect whereby fillers prevent direct responding to the suspect and degrade the evidentiary value of an identification procedure.

The third demonstration contrasts EIG and Deviation from Perfect Performance (DPP; Smith et al., 2019) by exploring the potential evidentiary value of obtaining confidence ratings from witnesses. For example, witnesses might rate their confidence that the suspect is the culprit on a scale of 0-100 (e.g., Colloff & Wixted, 2020). The DPP approach collapses the confidence rating scale at different cutoff values to find the one that produces the lowest overall error rate, with all confidence ratings below the cutoff value collapsed into the “rejection” category and all confidence ratings above the cutoff value collapsed into the “identification” category. EIG reveals that collapsing a confidence scale in this way reduces the evidentiary value of witness responses compared to considering all confidence levels separately. Thus, DPP and related measures cannot capture the full evidentiary potential of an identification procedure.

The fourth demonstration shows how EIG can be used to evaluate continuous predictors of witness accuracy, such as the response time (RT) of identification responses. Considering variables that predict witness accuracy increases the evidentiary value of witness responses, and EIG can be used to quantify this advantage to compare the relative value of different predictor variables or different prediction models. This example outlines a potential research application of EIG that cannot be supported by any of the existing measures of evidentiary value.

The fifth demonstration uses the simultaneous and sequential lineup data from Clark’s (2012) meta-analysis to explore the different policy implications of EIG and alternative measures of evidentiary value. Existing measures suggest that sequential lineups provide better evidence of guilt when the witness makes a suspect ID, but simultaneous lineups provide evidence of innocence when the witness fails to make a suspect ID. The existing measures, however, have no way to combine results across outcomes, and they take no account of how the different procedures affect the overall probability of observing the different identification outcomes.

Because EIG combines across outcomes, it reveals that the simultaneous and sequential data sets are nearly identical in terms of overall evidentiary value. That is, the two procedures produced very similar amounts of total information about guilt or innocence over multiple identification attempts.

The sixth, and final, demonstration continues the Clark (2012) simultaneous versus sequential example to explore the effect of ignoring identification attempts that do not result in a suspect ID (non-identifications). This demonstration highlights a key advantage for EIG over existing measures, as EIG can directly quantify the cost of mis-using evidence; instantiated here as a failure to acknowledge that non-identifications are evidence that the suspect is innocent. EIG decreased sharply when non-identifications were ignored for both the simultaneous and sequential data sets. The two procedures were nearly equivalent in overall evidentiary value, again highlighting a key difference between EIG and existing measures. That is, existing measures show that the sequential procedure produces stronger evidence of guilt when an identification is obtained, but EIG also considers the fact that identifications are obtained less often in the sequential procedure. The overall prevalence of identification outcomes is an important factor for policy makers to consider, as it has a big impact on the total information about guilt or innocence that a procedure will produce when it is used for multiple identification attempts.

We have included a number of *Conceptual Summary* statements, typically at the end of major sections, that review important concepts without quantitative detail. Readers who are not interested in mathematical implementation can skip from the section introductions to the conceptual summaries to obtain a complete picture of our major claims and the qualitative arguments that support them.

Information Theory and Eyewitness Identifications

This section introduces the mathematics for analyzing eyewitness identifications that are applied in the demonstrations reported below. Our goal is to have a theory-driven measure representing the overall potential of a given identification procedure to provide information about guilt or innocence. Such a measure will allow researchers to directly compare different procedures and inform policy recommendations for investigative practices when all aspects of identification performance are considered.

There are three critical steps in pursuing this goal: (1) determine an appropriate reference population for interpreting witness responses; (2) use Bayes' theorem to assess the probability that a given suspect is guilty given the identification response (Wells & Lindsay, 1980); (3) compute Expected Information Gain (EIG), a score that characterizes the evidentiary value of the identification procedure as a whole. Within this framework, a procedure with high evidentiary value is one that tends to dramatically reduce uncertainty about guilt or innocence across multiple identification attempts with varying outcomes.

The third step, computing EIG, is itself comprised of four sub-steps: (3A) use Shannon *surprisal* to measure the degree of error in probabilistic judgments (Shannon, 1948); (3B) evaluate the average judgment error, or *entropy*, across a full probability distribution representing uncertainty about the guilt or innocence of the suspects; (3C) assess *information gain* by comparing the certainty of suspect guilt before and after the witness response, as measured by entropy; and (3D) combine information gain across all responses to yield an *EIG* value.

This section also discusses the important measurement considerations involved in selecting EIG as an index of overall evidentiary value. We explain why defining overall

evidentiary value requires a measure that puts uncertainty on an additive scale, one of the key properties that has made Shannon surprisal an extremely useful measure of uncertainty (Cover & Thomas, 2006; McElreath, 2016; Shannon, 1948; Quinlan, 1986). We also introduce PG_{EIG} , a translation of EIG onto the probability scale, which may aid the interpretation of evidentiary value metrics.

Readers who are not interested in quantitative details can safely skip to the *Conceptual Summary of Information Theory and Identifications* sub-section at the conclusion of this section.

Step 1: The Reference Population

Considering a relevant reference population facilitates the ability to effectively interpret a source of evidence. Take the familiar example of a doctor interpreting a diagnostic test to judge whether or not one of her patients has a certain disease. The doctor would ideally consider the characteristics of a larger population of similar patients; namely, the overall prevalence of the disease, the proportion of people with the disease who get the same test result as her patient, and the proportion of people without the disease who get the same test result as her patient (Gigerenzer, Gaissmaier, Kurz-Milcke, Schwartz, & Woloshin, 2008).

Witness responses play the same role in investigations as diagnostic tests in medicine. Just as positive test results are more likely when the disease is present, suspect IDs are more likely when the suspect is guilty. Just as negative test results are more likely if the disease is absent, rejections and filler IDs are more likely when the suspect is innocent. For lineups, the reference population defines the overall proportion of guilty and innocent suspects and the probability of obtaining specific witness responses for guilty and innocent suspects (Wells, Yang, & Smalarz, 2015).

As an example, Table 1 provides this information for the standard showup condition in Colloff & Wixted (2020; Exp. 3). To simplify the example, we collapsed confidence ratings into two levels. The shaded section provides response proportions for guilty and innocent suspects. For example, 31% of the responses were high-confidence identifications of guilty suspects. A similar reference population could be defined for any eyewitness memory experiment regardless of the procedure used and the available responses.

Reference populations are easy to define for eyewitness experiments because the guilt or innocence of each suspect is known. One might be tempted to conclude that it is not possible to consult a reference population in actual investigations, as the guilt or innocence of each suspect is unknown. It is, of course, possible to determine the proportion of suspect IDs, filler IDs, and rejections. The difficulty lies in separating those proportions by guilty and innocent suspects. Recently developed modeling techniques, however, allow researchers to recover the properties of reference populations, even for data that are collapsed across guilty and innocent suspects (Cohen, Starns, & Rotello, 2020; Wixted, Mickes, Dunn, Clark, & Wells, 2016). We will return to this issue in the *General Discussion*. For now, we note that readers should not dismiss the importance of reference populations by assuming that they are irrelevant in actual cases.

Throughout the paper, we will assume that each data set we consider defines a reference population and explore how different measures of identification effectiveness would characterize that reference population. Of course, sometimes researchers will want to treat a data set as a sample from a larger population and characterize uncertainty about the possible population statistics; that is, they will want to perform statistical inference. We will avoid this additional complication to an already quite complex paper, and readers should keep in mind that our

analyses are meant to explain the EIG measure and show how it is different from alternative measures, not to make general conclusions about empirical phenomena in eyewitness memory.

Complete versus Incomplete Measures. All measures for assessing identification procedures are based on the reference population, albeit in different ways. Indeed, differences across these measures are generated by the differing relationships between the reference population and the measures. In particular, many measures are “blind” to changes in some aspects of the reference population, that is, changes in these aspects do not result in changes in the measure. We refer to measures that do not incorporate all aspects of the reference population as *incomplete*. In contrast, EIG is a *complete* measure, in that it integrates all aspects of the reference population. We will return to the implications of this distinction in more depth in the demonstrations below; in particular, *Demonstration 5* offers a detailed exploration of the differences between conclusions based on EIG and the measures discussed in this section. For now, we briefly note a number of the more common measures of identification effectiveness and identify the reference population information they do and do not rely on. In particular, we consider the measures discussed in Clark (2012) and Wells, Yang, & Smalarz (2015).

Using data derived from previous identification experiments, Clark (2012) explored a number of measures of probative value, which he defined as “the usefulness of the evidence in discriminating between the guilty and the innocent” (p. 244). It is important to emphasize that each of the measures considered by Clark (2012) was designed to ask a particular question of the data. As noted by Clark (2012), “There is more than one way to calculate probative value, and the conclusions that one draws will depend on precisely how probative value is calculated” (p. 244). That is, each of these measures was designed to assess a different aspect of line-up performance and each has led to important insights about eyewitness lineup procedures. Thus,

although we claim that these measure are incomplete, we are not asserting that these measures are necessarily flawed or that EIG is always the appropriate measure to use. Instead, our purpose is to highlight the fact that EIG is a complete measure, and, as such, provides a perspective on identification performance that is different from any of the extant measures.

Table 2 summarizes the relationship between various measures and properties of the reference population. Each column of Table 2 represents a different probative measure considered in Clark (2012) or Wells, Yang, and Smalarz (2015), plus EIG. The diagnosticity ratio (diag.) for a given witness response indicates the relative support for guilt over innocence provided by a witness response (e.g., Price & Fitzgerald, 2016; Wells & Lindsay, 1980; Wells & Turtle, 1986; Wells, Yang, & Smalarz, 2015), and is calculated as the probability of observing the response for a guilty suspect divided by the probability of observing the response for an innocent suspect.² Innocence risk (IR) indicates the probability that an identified suspect is innocent under the assumption of a 50% prior probability of guilt (Clark, 2012), and is given by the false ID rate divided by the sum of the true and false ID rates. Clark also considered two signal-detection theory (SDT) measures (Macmillan & Creelman, 2005), d' , which measures how well witnesses can discriminate guilty from innocent suspects, and c , which measures the overall tendency to make an identification. Probability gain (PG_{Resp}) indicates the change in the probability of guilt after learning a particular witness response, calculated as the difference between the posterior and prior probability.

Each row of Table 2 is a different aspect of the reference population. We acknowledge that there could be other useful ways to divide the information provided by the reference population. These entries were selected to highlight critical differences across the probative

² What we described in specifically the diagnosticity for guilt. The inverse gives the diagnosticity for innocence.

measures. Note that, with the exception of EIG, the measures are applied to a single response class, x_1 . For example, we can compute the diagnosticity of the high confidence suspect IDs in Table 1, in which case x_1 would be the class of high-confidence suspect IDs. In Table 2, $x_{2..n}$ are the other response categories. Continuing the example from Table 1, $x_{2..n}$ are low-confidence suspect IDs and high- and low-confidence rejections.

We are now in a position to consider the relationship between the measures and the reference population. First, by necessity, all of these measures are sensitive to changes in conditional response proportions for guilty, g , and innocent, i , suspects, $p(x_1 | g)$ and $p(x_1 | i)$. Second, with the exception of EIG, none of the measures, as typically applied, consider changes in performance or distributions for responses other than x_1 , i.e., $p(x_{2..n} | g)$ and $p(x_{2..n} | i)$. Third, because diagnosticity ratios, innocence risk, and the signal-detection measures are based on the conditional probability of a response for guilty or innocent suspects, changing the prior probability of guilt, $p(g)$, has no effect. Thus, only EIG and PG_{Resp} are sensitive to changes in base rates. Finally, with the exception of EIG, none of the measures consider the overall prevalence of all of the various witness responses, that is, both $p(x_1)$ and $p(x_{2..n})$. For example, a procedure with true and false identification rates of 10% and 2.5% has the same diagnosticity ratio as a procedure with corresponding values of 80% and 20% [$.10/.025 = .80/.20 = 4$], even though identifications are much rarer overall in the first procedure.

The overall prevalence of x_1 , $p(x_1)$, does affect the position of the signal-detection response criterion. The signal-detection measures, however, are not affected by the distribution across $p(x_{2..n})$. For example, if the target response, x_1 , is high-confidence suspect IDs, the signal-detection measures are not affected by changes in the proportion of high- and low-confidence rejections. As discussed in detail below, although a Receiver Operating Characteristic (ROC) can

be formed by calculating hit and false alarm rates for suspect IDs across multiple confidence ranges (Mickes, Flowe, & Wixted, 2012), the ROC points are still not affected by the distribution of rejections across confidence categories (or filler IDs for lineups; Wells, Smalarz, & Smith, 2015; Wells, Smith, and Smalarz, 2015; Wells, Yang, & Smalarz, 2015). For this reason, we do not include $p(x_{2..n})$ under SDT in Table 2.

Step 2: Bayesian Updating After a Witness Response

In the Bayesian approach to interpreting eyewitness identifications, uncertainty about the guilt or innocence of a particular suspect is represented by a probability distribution, and this distribution is updated when new evidence is encountered (Wells & Lindsay, 1980). Here, there are two hypotheses that might be true (guilty or innocent), so uncertainty is represented by a Bernoulli distribution. The Bernoulli distribution is characterized by a single parameter, which, in this case, can be defined as the probability that the suspect is guilty, $p(g)$, where the probability that the suspect is innocent is the complement of this value, i.e., $p(i) = 1 - p(g)$. For the Table 1 data, there are an equal number of showup trials with a guilty suspect and an innocent suspect (see the right column of Table 1), so initial uncertainty about guilt is represented by a Bernoulli distribution with a 50% chance of guilt and a 50% chance of innocence. In other words, with no information about the witness response, one is resigned to being completely uncertain about whether the suspect in a given identification attempt is guilty or innocent.

When the witness response is considered, the posterior probability of guilt can be determined by applying Bayes Theorem (Wells & Lindsay, 1980). Taking the showup data in Table 1 as an example, there were .31 high-confidence identifications of guilty suspects out of .42 overall high-confidence identifications, yielding a 74% posterior probability of guilt

$[.31/.42=.74]^3$. In other words, the posterior distribution is a Bernoulli distribution with $p(g | hid) = .74$ and $p(i | hid) = .26$, where *hid* indicates a high-confidence suspect ID.

The posterior distributions for all witness responses are provided in the unshaded region of Table 1. Note that high-confidence identifications provided stronger evidence of guilt than low-confidence identifications (74% versus 67% certainty of guilt). Rejections provided evidence of innocence, which was again stronger when witnesses expressed high versus low confidence (88% and 65% certainty of innocence, respectively).

Characterizing the level of certainty in guilt or innocence plays a critical role in making effective decisions across situations with varying potential costs and benefits (Robert, 2007). For example, a 74% level of certainty that the suspect is guilty would reasonably justify issuing a search warrant by a “probable cause” standard, but would not justify a conviction by a “beyond a reasonable doubt” standard (Kagehiro & Stanton, 1985). Of course, identification evidence can be combined with other sources of evidence to achieve higher levels of certainty about a suspect’s guilt or innocence (Clark, Benjamin, Wixted, Mickes, & Gronlund, 2015).

The Bayesian approach has provided valuable insight into effective policies for identification procedures (Wells & Lindsay, 1980; Wells & Olson, 2002; Wells & Turtle, 1986; Wells, Yang, & Smalarz, 2015). For example, Wells, Yang, and Smalarz (2015) explored a number of variables in terms of the shift from the prior to the posterior probability of guilt for individual witness responses. Although they called this measure information gain, we call it *probability gain* (PG) to avoid confusion with our use of Shannon (1948) information below. More specifically, we labeled this measure PG_{Resp} in Table 2, with the subscript used to

³ Using Bayes rule: $p(g | hid) = p(hid | g) p(g) / (p(hid | g) p(g) + p(hid | i) p(i)) = .62 \times .50 / (.62 \times .50 + .22 \times .50) = .31/.42 = .74$ and $p(i | hid) = p(hid | i) p(i) / (p(hid | i) p(i) + p(hid | g) p(g)) = .22 \times .50 / (.22 \times .50 + .62 \times .50) = .11/.42 = .26$.

emphasize that this measure applies to a single witness response class. For example, high-confidence suspect IDs bring the probability of guilt from .50 to .74, representing a PG_{Resp} of .24 $[=.74-.50]$. Higher values are preferred, as they indicate that observing the witness response produced a larger increase in certainty about the guilt or innocence of the suspect.⁴

Unfortunately, variations in lineup procedures can, and often do, introduce trade-offs in a Bayesian analysis by improving the evidentiary value associated with some witness responses and decreasing evidentiary value for others (e.g., Clark, 2012; Wells, Yang, & Smalarz, 2015). Figure 1 illustrates such a trade-off when comparing the showup data from Table 1 to data from the standard simultaneous lineup condition in the same experiment (Colloff & Wixted, 2020, Exp. 3). Figure 1A shows the overall probability of obtaining each possible witness response regardless of suspect guilt or innocence. Note that high- and low-confidence filler IDs are included as possible responses for the lineup condition, but these responses cannot occur in the showup procedure. Figure 1B shows the gain in certainty of guilt or innocence associated with each response measured on the probability scale, i.e., PG_{Resp} . For the showup, the bar heights for suspect IDs can be determined by subtracting the prior probability of guilt (.50) from the posterior probabilities of guilt reported in Table 1. The same goes for the rejection responses, substituting the prior and posterior probability of innocence.

The results in Figure 1B show that some responses are more informative in the showup data, and some are more informative in the lineup data. For example, high-confidence suspect

⁴ Wells, Yang, and Smalarz (2015) defined information gain as the absolute difference between the prior and posterior probability of guilt. Using the absolute difference introduces the possibility of obtaining a positive value even if certainty is lower after a witness response than before; for example, if the prior probability of guilt was 80% and observing a rejection response brought the posterior probability of guilt to 60%. Wells, Yang, and Smalarz would describe that as an information gain of 20%. In our analysis, that situation would signify a *loss* of information based on the witness response. We only apply the Wells, Yang, and Smalarz measure (probability gain, in our terminology) to cases with a 50% prior probability of guilt, so this difference in interpretation does not affect any of the probability gain results that we report (i.e., the original authors would apply the measure in the same way that we did in these cases).

IDs produce a larger increase in certainty of guilt in the lineup, but high-confidence rejections produce a larger increase in certainty of innocence in the showup. As will be discussed more fully below, evidentiary value in the lineup is hampered by filler IDs that provide poor evidence, i.e., low gain in certainty. These sorts of tradeoffs are pervasive in empirical data sets (Clark, 2012).

Introducing Information Theory Measures

Before showing how to compute EIG, in this section we provide a high-level introduction to the information-theory measures and an example of how they can be used to resolve tradeoffs in evidentiary value, as in the showup and lineup data in Figure 1.

As outlined in the previous section, Bayesian analysis has been used to determine the evidentiary value of a *particular* witness response, e.g., a high confidence suspect response. The main goal of this work is to extend the Bayesian approach to measure the *overall* evidentiary value of an identification procedure by analyzing a full set of identification attempts, e.g., *all* potential witness responses. This goal has great practical importance. Any identification procedure adopted by real investigators will be used across numerous cases with differing witness responses. Accordingly, we seek a measure that can be used to select the identification procedure that will provide the most information about guilt or innocence when it is used across multiple investigations.

One straightforward way to combine evidentiary value across the complete set of possible witness responses is to evaluate how much a procedure increases certainty of guilt or innocence *on average*. It is, of course, possible to average probability gain across all response categories, and this is also true for other measures like diagnosticity ratios, innocence risk, or logarithms of these measures (Clark, 2012). However, in order for the average value of the measure to

correspond the average level of certainty, the measure must be on an equal-interval scale with certainty. That is, the same change in certainty must produce the same change in the measure over the full range of the scale. As will be discussed below, probability gain does not have an equal-interval relationship with certainty, nor do any of the other alternative measures.

Fortunately, a measure with this property exists: Shannon entropy.

Following up on work by Hartley (1928), Shannon (1948) introduced an equal-interval measure of certainty that he called *entropy*, based on its formal relationship to concepts in physics. Shannon focused on transmission of information by communication systems, but the same math that allowed him to characterize uncertainty in a signal per unit time allows eyewitness researchers to characterize uncertainty in guilt or innocence per identification attempt. Entropy expresses uncertainty, and so it is highest when the probability of guilt is 50% and decreases as the probability of guilt approaches 100% (certain guilt) or 0% (certain innocence; see the right panel of Figure 2). Thus, a reduction in entropy represents a gain in information, as rigorously defined by information theory (Shannon, 1948).

The next section formally describes how entropy is calculated, but first we illustrate the use of entropy for the lineup versus showup example from Colloff and Wixted (2020). Figure 1C shows the gain in certainty about guilt or innocence, this time on the entropy scale, given each possible response. That is, we calculated the change in entropy from before to after a witness response for each response category. Positive values indicate a *decrease* in entropy, i.e., a loss of uncertainty or, equivalently, a gain of information. For example, in the showup, whereas a high-confidence rejection greatly increased certainty of innocence, a high-confidence suspect ID generated a more modest increase in the certainty of guilt.

Once this certainty gain is computed for each response, these values enter into a weighted average to produce the Expected Information Gain (EIG). That is, EIG is computed by multiplying the probability of obtaining each witness response (Figure 1A) by the change in certainty for the response (Figure 1C), and summing across responses. Put simply, EIG is the average change in certainty across all identification attempts. Note that EIG satisfies our goal of using all of the information from the reference population. In particular, EIG is a combination of the certainty of guilt or innocence before an identification attempt, the certainty after an identification attempt for each class of witness response, and the overall probability of observing each class of witness response.

EIG values for the Colloff & Wixted data are shown in the far right of Figure 1C, and they indicate that the showup acted as a somewhat better evidence source than the lineup. That is, on average, witness responses from the showup supported higher certainty of guilt or innocence than the lineup. Another way to say this is that the showup produced more information about suspect guilt or innocence than the lineup. Eyewitness researchers would likely be surprised to learn that the showup has higher evidentiary value than the lineup for the data in Figure 1. Indeed, the authors of the original paper concluded that the lineup was superior based on another standard of identification effectiveness, i.e., area under the ROC (Colloff & Wixted, 2020), and that conclusion echoes a clear consensus in the literature (e.g., Gronlund et al., 2012; Wells, Yang, & Smalarz, 2015). Thus, this example provides an excellent opportunity to consider cases in which different measures produce different conclusions for the same data (Clark, 2012). We return to this point in *Demonstration 2* below, which includes a more in-depth analysis of the Colloff and Wixted (2020) data.

Step 3: Calculating Expected Information Gain

In this section, we formally introduce how to compute EIG and associated measures.

Step 3A: Surprisal. Entropy expresses uncertainty as the average Shannon surprisal for a probability distribution. Shannon (1948) derived the equation for the amount of information gained by a single outcome, or *surprisal*, to satisfy a number of reasonable and intuitive properties, among them: (1) a definite outcome produces no information; (2) less probable outcomes produce more information; and (3) the information produced from multiple events is additive. There is a unique function that satisfies these axioms and has many other desirable properties. Namely, the surprisal or information content I of an outcome y is defined as

$$I(y) = \log_2 \left(\frac{1}{p(y)} \right), \quad (1)$$

where $p(y)$ is the judged probability of outcome y before the outcome was revealed. The logarithm base is arbitrary, but, when base 2 is used, as we adopt here, information is measured in *bits* (Shannon, 1948).

Surprisal is a natural measure for the degree of error in a probabilistic judgment of an unknown outcome, in this case, whether a suspect is guilty or innocent. For example, consider a suspect who is in fact guilty. Assigning a low probability of guilt, say 1%, would result in high judgment error. In this case, learning that the suspect is in fact guilty would be a big surprise, and, equivalently, a large gain in information:

$$\log_2(1/.01) = 6.64. \quad (2)$$

Assigning a high probability of guilt, say 99%, would result in low judgment error. In this case, learning that the suspect is indeed guilty would not come as much of a surprise, and would therefore represent a small gain in information:

$$\log_2(1/.99) = 0.01. \quad (3)$$

Conversely, if the suspect is in fact innocent, then assigning a high probability of innocence (i.e., a low probability of guilt) represents low judgment error.

For another example, imagine a set of jurors who decided to convict a defendant because they were convinced of his guilt with 99% certainty; that is, they judged only a 1% chance that the defendant was innocent. These jurors should be very surprised to learn that the defendant was later exonerated by DNA evidence, and the large surprisal signifies a high degree of error in their judgment (in this scenario, an error with disastrous consequences).

The relationship between probability and surprisal is provided in the left panel of Figure 2. In brief, good judgments assign a high probability to the true state of affairs, and bad judgments assign a low probability to the true state of affairs.

Step 3B: Entropy. Defining Shannon surprisal sets the stage for calculating entropy, traditionally denoted by H (Shannon, 1948). Let Y be a discrete random variable with outcomes y_j , each with probability $p(y_j)$, then, the entropy of Y is defined as

$$H(Y) = \sum_j p(y_j) \log_2 \left(\frac{1}{p(y_j)} \right). \quad (4)$$

That is, entropy is a sum over all possible outcomes in which the surprisal of each outcome is weighted by the probability of that outcome. In other words, entropy is the average surprisal. Because events that never happen cannot affect average surprisal, outcomes with probability zero are dropped. In short, entropy measures the uncertainty associated with a probability distribution – the less predictable the outcome, the higher the entropy.

In the current lineup context, Y has two possible outcomes: guilty, g , associated with probability $p(g)$, and innocent, i , associated with probability $p(i)$. The right panel of Figure 2 shows how entropy for judgments about guilt or innocence changes across different distributions

of uncertainty, with each distribution characterized by the probability that a suspect from the distribution is guilty, $p(g)$. Notice that entropy is highest at .50 and decreases as the probability of guilt moves closer to 0 or 1.

In everyday terms, entropy embodies the intuition that it is easier to judge guilt for an individual taken from a set of predominantly guilty or predominantly innocent suspects than for a set that is an even mix of guilty and innocent suspects. For current purposes, the set of suspects could include all of the suspects who received the same witness response.⁵ When judging the guilt of suspects sampled from a set of 50% guilty and 50% innocent suspects, the optimal strategy is to judge that every suspect in the set has a 50% chance of being guilty or innocent. This situation leads to high judgment error, i.e., high surprisal at learning the true outcome, for both guilty suspects,

$$I(g) = \log_2(1 / p(g)) = \log_2(1 / 0.50) = 1.00 \text{ bit}, \quad (5)$$

and innocent suspects,

$$I(i) = \log_2(1 / p(i)) = \log_2(1 / 0.50) = 1.00 \text{ bit}, \quad (6)$$

and a high average judgment error (entropy) across the entire set of suspects,

$$H(Y) = p(g) I(g) + p(i) I(i) = 0.50 \times 1.00 + 0.50 \times 1.00 = 1.00 \text{ bit}. \quad (7)$$

The value of 1 here represents the 1 bit of information that would be gained, on average, if the suspect's true guilt or innocence was revealed.

An unbalanced set, say 80% guilty suspects and 20% innocent suspects, supports a higher level of certainty. Entropy measures *uncertainty*, so this higher certainty corresponds to lower

⁵ Eyewitness experiments often use the same suspect across many identification attempts, so it would be more precise to refer to "a set of identification attempts" than "a set of suspects" for experimental data. For convenience, we adopt terminology that is appropriate for the situation that experiments are meant to simulate, where identification attempts come from different cases with different suspects. Certainly, a real investigation would not have hundreds of identification attempts for the same suspect, as in many experiments.

entropy. Specifically, with this distribution, the most reasonable judgment is that every suspect has an 80% chance of being guilty, and this judgment is associated with low surprisal for most members of the set. Thus judgment error decreases relative to the 50-50 set. Specifically,

$$I(g) = \log_2(1 / p(g)) = \log_2(1 / 0.80) = 0.32 \text{ bits}, \quad (8)$$

$$I(i) = \log_2(1 / p(i)) = \log_2(1 / 0.20) = 2.32 \text{ bits}, \quad (9)$$

$$H(Y) = p(g) I(g) + p(i) I(i) = 0.80 \times 0.32 + 0.20 \times 2.32 = 0.72 \text{ bits}. \quad (10)$$

Overall, judgments of guilt and innocence are associated with just 0.72 bits of uncertainty in the 80-20 set compared to 1 bit of uncertainty in the 50-50 set. Notice that judgment error has actually *increased* for innocent suspects relative to the 50-50 set of suspects (from 1 to 2.32 bits). This increase occurred because the innocent suspects in the 80-20 set had the bad luck of being mixed in with mostly guilty suspects, creating the misleading impression that they were probably guilty. Because innocent suspects are rare in this set, however, the *average* judgment error has decreased relative to the 50-50 set.

The same reasoning follows for a set 20% guilty suspects and 80% innocent suspects. The most reasonable judgment is that each suspect in the set has a 20% chance of being guilty, and the average judgment error is

$$I(g) = \log_2(1 / p(g)) = \log_2(1 / 0.20) = 2.32 \text{ bits}, \quad (11)$$

$$I(i) = \log_2(1 / p(i)) = \log_2(1 / 0.80) = 0.32 \text{ bits}, \quad (12)$$

$$H(Y) = p(g) I(g) + p(i) I(i) = 0.20 \times 2.32 + 0.80 \times 0.32 = 0.72 \text{ bits}. \quad (13)$$

In colloquial terms, overall uncertainty has decreased relative to the 50-50 set because it is not surprising to learn that a suspect from the 20-80 set is innocent, so judgment error is low for the majority of suspects (i.e., the 80% who are actually innocent).

In the most extreme case, in which a set of suspects are either all guilty or all innocent, one can make perfect judgments, and learning whether the suspect is actually guilty or innocent would provide no new information. For example, in a set of only guilty suspects, the most reasonable judgment is a 100% chance of guilt for each suspect, and this judgment is always confirmed, yielding zero judgment error,

$$I(g) = \log_2(1 / p(g)) = \log_2(1 / 1) = 0 \text{ bits}, \quad (14)$$

$$H(Y) = p(g) I(g) = 1 \times 0 = 0 \text{ bits}. \quad (15)$$

Note that innocent suspects are eliminated from the entropy calculation because there are none in the set. Likewise, for a set of only innocent suspects, the most reasonable judgment is a 100% chance of innocence for each suspect, and this expectation is always confirmed, again yielding zero judgement error,

$$I(i) = \log_2(1 / p(i)) = \log_2(1 / 1) = 0 \text{ bits}, \quad (16)$$

$$H(Y) = p(i) I(i) = 1 \times 0 = 0 \text{ bits}. \quad (17)$$

In both cases, an entropy of zero corresponds to 100% certainty in our ability to correctly specify guilt or innocence.

We have emphasized the fact that entropy provides an equal-interval measure of uncertainty, and we will reinforce this point by returning to an earlier example. In the *Introduction*, we used the level of certainty in the outcome of two coin flips to show that probability is not an equal-interval measure of certainty.. On the probability scale, certainty in the outcome started at 25%, went to 50% after a single flip was observed, and went to 100% after both flips were observed. Thus, 1 coin flip worth of information produced different changes in probability depending on the starting point. Switching to a Shannon surprisal scale, uncertainty in the outcome started at 2 bits [=log₂(1/.25)], went to 1 bit after a single flip was observed

$[\log_2(1/.50)]$, and went to 0 bits after both flips were observed $[\log_2(1/1.00)]$. Thus, one coin flip provides 1 bit of information, regardless of the starting level of certainty. Appendix A further explores this equal-interval property for interested readers.

Step 3C: Information Gain. The information gain associated with a particular witness response is the difference in entropy from the prior to posterior distribution.⁶ There are three steps in defining this quantity: computing the prior entropy, computing the posterior entropy, and then taking the difference. We will again consider the showup data from Table 1. First, compute the prior entropy. This was done previously. Recall that the prior probability of guilt for the showup is .50, corresponding to a prior entropy of 1.00 bit (see Equation 7).

Second, compute the posterior entropy. For the showup data, observing a high-confidence identification brought the probability of guilt to .74. In order to compute the posterior entropy, we need to condition the calculations on the witness response. Let $I(y | X=x)$ or, more simply, $I(y | x)$, be the surprisal for outcome y given response x , and let $H(Y | X=x)$ or $H(Y | x)$ be the entropy for distribution Y given response x . The posterior entropy can then be found by computing the surprisal for guilty suspects (g) given a high-confidence identification (hid),

$$I(g | hid) = \log_2(1 / p(g | hid)) = \log_2(1 / 0.74) = 0.43 \text{ bits}, \quad (18)$$

the surprisal for innocent suspects (i) given a high-confidence identification,

$$I(i | hid) = \log_2(1 / p(i | hid)) = \log_2(1 / 0.26) = 1.94 \text{ bits}, \quad (19)$$

and, the average surprisal, or entropy⁷, given a high-confidence identification,

$$H(Y | hid) = p(g | hid) I(g | hid) + p(i | hid) I(i | hid) = \quad (20)$$

⁶ The term “information gain” is sometimes used to denote the mutual information between two random variables, a quantity that we call “expected information gain” (e.g., Quinlan, 1986). For continuity with previous eyewitness studies (e.g., Wells & Lindsay, 1980), we will use “information gain” when referring to a particular witness response and “expected information gain” when referring to the average information gain across a set of identification attempts.

⁷ Entropy conditioned on a particular value, i.e., $H(Y | X = x)$, is called the specific conditional entropy.

$$0.74 \times 0.43 + 0.26 \times 1.94 = 0.82 \text{ bits}.$$

Finally, the information gain, IG , for response x and distribution Y is given by the difference:

$$IG(x) = H(Y) - H(Y | x). \quad (21)$$

For the showup data, the information gain after a high confidence identification is given by

$$IG(hid) = H(Y) - H(Y | hid) = 1.00 - 0.82 = 0.18, \quad (22)$$

which is the value depicted in Figure 1C for high-confidence showup identifications. This process is repeated for each response category.

Step 3D: Expected Information Gain (EIG). Once the information gain for each response has been calculated, the *Expected Information Gain (EIG)* is calculated by taking the weighted average of information gain across responses:⁸

$$EIG = \sum_i p(x_i)IG(x_i). \quad (23)$$

This process produces EIG values of 0.22 and 0.18 for the showup and lineup data from Colloff and Wixted (2020; Experiment 3), respectively, as shown in Figure 1C.

Probability Gain from EIG (PG_{EIG}). Many readers will have little or no experience working with Shannon surprisal as a unit of measurement. To facilitate the interpretation of overall evidentiary value, we will often translate the change in entropy back to the corresponding change in the probability parameter of a Bernoulli distribution representing uncertainty about guilt, yielding a value that we will call the *Probability Gain from EIG (PG_{EIG})*. The subscript EIG in PG_{EIG} differentiates the probability gain computed from EIG from the probability gain for a specific response, PG_{Resp} , discussed previously. Where clear in context, we sometimes simply

⁸ In some circumstances, it might be useful to consider other distributional characteristics of the information gain values, such as the variance. We focus on the average, because a procedure that produces higher average information gain will also produce higher total information across multiple investigations.

refer to probability gain. Because the process is complex and not particularly informative, we do not show how to compute PG_{EIG} here. The equation to estimate PG_{EIG} can be found in Appendix B.

PG_{EIG} simplifies interpretation in two ways. It uses the more familiar probability scale and it translates the same amount of information represented by the EIG value to the most basic identification scenario. Specifically, PG_{EIG} is based on a procedure that has a 50% prior probability of guilt and has equal evidentiary value across all witness responses. In this scenario, PG_{EIG} is the PG_{Resp} for each of the response categories. To be clear, EIG does not have to assume these simplifications, so for example, a researcher could calculate EIG for a showup procedure with a 20% prior probability of guilt and with responses that vary in evidentiary value, e.g., perhaps the evidence of guilt provided by identifications is stronger than the evidence of innocence provided by rejections. To provide perspective on the amount of information represented by this EIG value, PG_{EIG} translates the same amount of information to the simplified scenario. For example, if the showup has a PG_{EIG} of 20%, this result corresponds to a scenario in which identifications bring the probability of guilt from 50% to 70% and rejections bring the probability of innocence from 50% to 70%. This scenario produces the same amount of information as the actual showup procedure. Readers should note that using PG_{EIG} for interpretation is *not* meant to imply that a procedure would produce the same amount of information with any prior probability of guilt, and in fact we demonstrate below that EIG (and thus PG_{EIG}) is affected by prior probability. Again, PG_{EIG} only provides a way to think about an amount of information in more familiar terms.

Figure 3 shows the relationship between EIG and PG_{EIG} . A probability gain of zero means that the identification procedure had no evidentiary value, corresponding to a situation in

which certainty about guilt or innocence is 50% in both the prior and the posterior distribution. A probability gain of 50% means that the identification procedure provided perfect evidence, corresponding to a situation in which the 50% certainty for the prior distribution increased to 100% certainty for the posterior distribution. For the example shown in Figure 1, PG_{EIG} was .27 for the showup procedure and .24 for the simultaneous lineup procedure. Thus, on average, an identification attempt in the showup procedure produced enough information to increase certainty of guilt or innocence from 50% to 77% and an identification attempt in the lineup procedure produced enough information to increase certainty from 50% to 74%.

It is important to recognize that PG_{EIG} values do not indicate the change in the probability of guilt or the probability of innocence; instead, they indicate the change in the level of certainty about guilt or innocence. For example, both a 90% probability of guilt and a 10% probability of guilt correspond to 90% certainty, because, in the former case, judgments of guilt have a 90% chance of being correct and, in the latter case, judgments of innocence have a 90% chance of being correct.

Further Properties of Information Theory Measures

In this section, we further discuss properties of information theory measures and consider some extensions on their use that go beyond the initial example presented above. In particular, we discuss (1) why entropy is the appropriate measure when averaging; (2) the relationship between confidence and information gain; (3) the effect of base rates on EIG; and (4) the effect of miscalibrated interpretation of witness responses on EIG.

Averaging Entropy. As shown previously, EIG resolves tradeoffs in evidentiary value by computing the average gain in certainty across all witness responses. This section further examines why the average must be taken on the entropy scale, rather than, say, the probability

scale. The explanation uses a hypothetical example designed to highlight key concepts, but, as shown below, the same concepts can also be applied to real data.

Figure 4 provides hypothetical identification data from two showup procedures: Procedures A and B. Both procedures include 80 identification attempts, 40 with guilty suspects and 40 with innocent suspects. In both procedures, witnesses identified 30 of the 40 guilty suspects and 10 of the 40 innocent suspects.

The only difference between the two procedures is that witnesses in Procedure B were also asked to indicate their level of confidence as either high or low. In Procedure B, for the 30 witnesses who either correctly identified a guilty suspect or correctly rejected an innocent suspect, 19 indicated high confidence (HC) and 11 indicated low confidence (LC). For the 10 witnesses who either incorrectly identified an innocent suspect or incorrectly rejected a guilty suspect, 1 indicated high confidence and 9 indicated low confidence.

Hopefully, the immediate impression is that Procedure B has higher evidentiary value than Procedure A. All witness responses in Procedure A support 75% certainty in guilt or innocence [30 out of 40]. The same would be true of Procedure B if witness confidence was ignored, but high and low confidence responses differ greatly in the extent to which they distinguish guilty and innocent suspects. High confidence responses support 95% certainty of guilt or innocence [19 out of 20]. Low confidence responses support 55% certainty of guilt or innocence [11 out of 20]. Thus, high and low confidence responses dramatically differ in the extent to which they reliably determine guilt, and knowing this information certainly improves the ability to make effective decisions guided by the witness responses as a source of evidence.

Figure 4 shows the computations involved in assessing the evidentiary value of the witness responses for the two procedures. For both procedures, we compute the average across

probability gain (PG_{Resp}) and information gain. Our goal is to demonstrate that averaging across information gain appropriately characterizes the full evidentiary value of each procedure, whereas averaging across probability gain does not.

First consider Procedure A. Each of the identified suspects is best judged as having a 75% chance of being guilty and each of the rejected suspects is best judged as having a 75% chance of being innocent. Thus, relative to the 50% prior probability, both identifications and rejections achieve a PG_{Resp} of .25 [.75 - .50], resulting in an average probability gain across identification and rejections of .25.

To compute EIG, we first compute the prior entropy, which in this case is 1 given that half of the suspects in the reference population are guilty. We next compute entropy conditioned on the response, which, because of the symmetry in the data, is the same for both identifications and rejections:

$$H(Y|id) = H(Y|rej) = \frac{30}{40} \times \log_2 \left(\frac{1}{.75} \right) + \frac{10}{40} \times \log_2 \left(\frac{1}{.25} \right) = 0.81. \quad (24)$$

For both identifications and rejections suspects, information gain is 0.19 [1.00 – 0.81], resulting in an average information gain, or EIG, of 0.19.

Next, consider Procedure B. Each suspect identified with high confidence is best judged as having a 95% chance of being guilty, and each suspect rejected with high confidence is best judged as having a 95% chance of being innocent. Thus, high-confidence responses have a PG_{Resp} of .45 [= .95 - .50]. Each suspect identified (rejected) with low confidence is best judged as having a 55% chance of being guilty or (innocent), resulting in a PG_{Resp} of .05 [= .55 - .50]. Because each response category has the same number of data points, we can simply average to determine an average probability gain of .25.

The entropy for high confidence identifications and rejections is

$$H(Y|id) = H(Y|rej) = \frac{19}{20} \times \log_2 \left(\frac{1}{.95} \right) + \frac{1}{20} \times \log_2 \left(\frac{1}{.05} \right) = 0.29, \quad (25)$$

producing an information gain of 0.71 [1.00 – 0.29]. Similarity, entropy for low confidence identifications and rejections is

$$H(Y|id) = H(Y|rej) = \frac{11}{20} \times \log_2 \left(\frac{1}{.55} \right) + \frac{9}{20} \times \log_2 \left(\frac{1}{.45} \right) = 0.99, \quad (26)$$

producing an information gain of 0.01 [1.00 – 0.99]. Averaging information gain produces an EIG of 0.36.

For both procedures, the average PG_{Resp} is 0.25, suggesting that Procedure A and Procedure B have equivalent evidentiary value. This result violates the (accurate) intuition that witness confidence provides valuable information by distinguishing responses that are highly accurate from responses that are often errors. In contrast, the average information gain increases from 0.19 for Procedure A to 0.36 for Procedure B, correctly indicating that Procedure B produces more overall information about guilt and innocence than Procedure A.

The key result is that averaging on the different scales produces different conclusions about the relative value of the procedures. Averaging on the entropy scale produces the correct conclusion. This is not a coincidence; instead, it is a natural consequence of the fact that entropy values have an equal-interval relationship with certainty, and thus can be validly averaged to characterize average certainty (Shannon, 1948). Probability gain does not have an equal-interval relationship with certainty, so, as mentioned previously, averaging PG_{Resp} values produces misleading results. Appendix A shows that none of the other measures mentioned above satisfy an equal-interval relationship with certainty.

In summary, it is critical that the averages are taken on the entropy scale even if the probability scale is used for interpretation. For applications that evaluate a single response class with a 50% prior probability of guilt, the two scales will always agree on the relative ranking of

procedures (as they do in Figure 1B and C). However, averaging on the probability scale does not accurately characterize average uncertainty, potentially leading to incorrect conclusions about the relative evidentiary value of different procedures. Appendix A evaluates whether a wide range of alternative measures have an equal-interval relationship with certainty. In short, the other measures of identification performance do not have this property, so none of these measures are appropriate for characterizing average certainty across identification attempts. This fact does not in any way invalidate the usefulness of these measures when they are applied to a single response category, but it does mean that they are not well suited to characterizing overall evidentiary value across all response categories.

Confidence and Information Gain. The previous example clearly illustrates that including confidence, as in Procedure B, improved evidentiary value. This result holds whenever confidence is appropriately related to accuracy. Although early research often failed to find a strong relationship between confidence and accuracy (e.g., Deffenbacher, 1980), this relationship is clearly supported by more recent research (e.g., Wixted, Mickes, Clark, Gronlund, & Roediger, 2015). To illustrate this relationship, this time with experimental, rather than hypothetical data, we again consider the Colloff and Wixted (2020) showup experiment. However, whereas the earlier example collapsed confidence into two levels, here, we consider the full confidence scale used by Colloff and Wixted, with the exception of collapsing across rare low-confidence responses (the scale ranged from 0 to 100 in steps of 10, and we collapsed levels 0-40).

These data showed a typical confidence-accuracy relationship, i.e., higher confidence responses were more accurate (Wixted et al., 2015). Identifications made at the highest and lowest confidence levels were 77% and 60% accurate, respectively, and the corresponding

accuracy rates for rejections were 95% and 63%. An EIG analysis shows a modest increase in information when considering witness confidence; namely, EIG increased from 0.19 when confidence was ignored to 0.23 when confidence was included. This boost provides further support for recording witness confidence (National Research Council, 2014), especially in light of other benefits of confidence ratings, such as their role in model-based recovery of guilty-suspect base rates and witness accuracy (Cohen et al., 2020; Wixted et al., 2016).

The Effect of Base Rates. As discussed previously, EIG is sensitive to the proportion of guilty suspects in the reference population, i.e., the guilty base rate. We now consider the relationship between EIG and guilty base rate using the showup and simultaneous lineup data of Colloff and Wixted (2020; also see Wells, Yang, & Smalarz, 2015, and Cohen et al, 2020, for further discussion of the importance of base rates).

In experimental contexts, the base rate of guilty suspects is determined by the experimenter and tends to be around 50% (e.g., Cohen et al., 2020). In investigative contexts, however, base rates can be influenced by a variety of factors, such as police standards for which suspects are subjected to a lineup (Wells, Yang, & Smalarz, 2015). Thus, researchers may be interested in comparing procedures across a range of potential base rates. This comparison can be achieved by defining alternative reference populations that cover a range of base rates while preserving the observed conditional response rates within each response category. EIG can then be calculated for each of these constructed reference populations using the same method described previously.

For the showup and simultaneous lineup data in from Colloff and Wixted (2020; Experiment 3), Figure 5 shows how EIG varies across the full range of possible base rates of

guilty suspects. The values for $p(g) = .5$ are the same as those calculated above, .22 and .18 for the showup and lineup, respectively, indicating an information advantage for the showup.

The curves in Figure 5 are similar to the probability-gain curves reported in Wells, Yang, and Smalarz (2015). As discussed previously, however, whereas those probability-gain curves apply to a single response category, these EIG curves represent the informational value across all response categories. EIG is highest near a 50% prior probability of guilt and approaches zero as the prior probability of guilt approaches 0 or 100%. That is, if the prior distribution already conveys a lot of information about guilt or innocence, i.e., prior entropy is low, then there is less information that can be gained by consulting a source of evidence. For example, if all suspects are innocent or all suspects are guilty (0 or 100% prior probability of guilt), then the prior entropy is 0 and observing a witness response cannot further decrease uncertainty in guilt or innocence.

Figure 5 shows that the showup procedure produces more information about guilt and innocence for all base rates above .26, with very little difference between the two procedures below this value. As shown in Figure 1, the primary advantage of the showup was that rejections were better evidence of innocence than in the lineup. With low base rates, the prior probability already indicates a high chance of innocence, so obtaining a further indication of innocence is less informative. That is, a rejection mostly just confirms what was already strongly suspected.

Demonstration 1 below thoroughly explores how base rates interact with information-based measures and *Demonstration 2* further explores the showup advantage. For now, we will just note that EIG can be evaluated for any desired base rate, and that researchers can construct base rate curves to compare alternative procedures in a comprehensive fashion.

The Effect of Miscalibration. All of our previous examples have assumed that identification evidence was properly interpreted, that is, the judged posterior probability of guilt associated with a response matched the true proportion of guilty suspects in the reference population for that response. For example, if 74% of high-confidence IDs involve a guilty suspect, as in Table 1, then we assumed that each suspect identified with high-confidence was judged to have a 74% chance of being guilty. However, eyewitness researchers cannot ensure that criminal justice officials will make exactly correct interpretations of every identification response, that is, judgments may be miscalibrated.

In real investigations, the exact statistics of the reference population are unknown, and must be approximated as accurately as possible based on analogy to experimental tasks, analysis of outside data sources like conviction rates or expert strength-of-evidence ratings (Amendola & Wixted, 2015), and modeling approaches that can be applied to data sets in which the guilt or innocence of each suspect is unknown (Cohen et al., 2020; Wixted et al., 2016). For example, an eyewitness researcher might conduct a study that attempts to mimic the lineup procedures, viewing conditions, and witness characteristics relevant for an investigative unit, and officials could use the resulting data to guide the interpretations of identification outcomes. It is then hoped that the witness accuracy in the study will be comparable to accuracy levels for the actual witnesses in the investigations, but there will inevitably be some “wobble.” Therefore, it is important to be able to factor in the cost of miscalibration and consider the extent to which miscalibration could change procedural recommendations.

Fortunately, the assumption of accurate judgment is not necessary for computing EIG – indeed, information theory provides an elegant means to compute the information loss incurred by such miscalibration. For example, if the true proportion of guilty suspects is 74%, not only

can we compute EIG under the assumption that suspects with high-confidence identifications are judged to have, say, an 80% chance of being guilty (over-certainty) or a 65% chance of being guilty (under-certainty), but we can also determine how much additional judgement error is introduced by this miscalibration. In the remainder of this section, we present the equations needed to compute EIG when evidence is miscalibrated. We also use the opportunity to explore the extent to which miscalibration degrades the informational value of an evidence source and the potential for miscalibration to change the relative performance of alternative identification procedures.

Calculating EIG when evidence is miscalibrated involves all of the same steps outlined above: the information gain associated with each response is the change in entropy from before to after the response was observed, and EIG is found by taking the weighted average of information gain across responses. Computing entropy, however, requires a modification. In particular, we now compute the *cross*-entropy between the actual probability distribution of guilty and innocent suspects and the *judged* probability of guilty and innocent suspects (Cover & Thomas, 2006).

Formally, let Y be a discrete random variable with outcomes y_j , each with probability $p(y_j)$, and let Y' be a discrete random variable representing the *judged* probability, $p'(y_j)$, of each y_j in Y . The cross-entropy of Y and Y' is defined as

$$H(Y, Y') = \sum_j p(y_j) \log_2(1/p'(y_j)), \quad (27)$$

still measured in bits. Note the similarity to the entropy calculation, Equation 4. The key difference is that, in cross-entropy, the information gain associated with each outcome, i.e., the surprisal or $\log_2(1/p'(y_j))$, is calculated based on the judged probabilities, i.e., the p' , with the surprisals weighted by the actual outcome probabilities, i.e., the p . When $p = p'$, entropy and

cross-entropy are identical. Miscalibration will always increase entropy, that is, $H(Y) \leq H(Y, Y')$.

The amount of information lost by a miscalibration is given by the Kullback-Leibler divergence

$$D_{KL}(Y, Y') = H(Y, Y') - H(Y), \quad (28)$$

again measured in bits.

As an example, consider again the Colloff and Wixted showup data from Table 1.

Previously, in Equation 20, we computed the entropy associated with a high-confidence suspect ID, $H(Y | hid) = 0.82 \text{ bits}$. Now, assume that these high-confidence suspect IDs are mistakenly interpreted as indicating an 80% probability of guilt, rather than the correct 74% probability of guilt. Then the cross-entropy is given by

$$H(Y, Y' | hid) = p(g | hid) \log_2 \left(\frac{1}{p'(g | hid)} \right) + p(i | hid) \log_2 \left(\frac{1}{p'(i | hid)} \right) = \quad (29)$$

$$0.74 \times \log_2 \left(\frac{1}{0.80} \right) + 0.26 \times \log_2 \left(\frac{1}{0.20} \right) = 0.84 \text{ bits}.$$

The amount of information loss associated with this miscalibration is given by

$$D_{KL}(Y, Y' | hid) = H(Y, Y' | hid) - H(Y | hid) = 0.84 - 0.82 = 0.02 \text{ bits}. \quad (30)$$

Notice that, in this example, the cost of these misjudgments is quite small, as the cross-entropies are nearly as low as the entropy. This result is reassuring, and suggests that it is possible to make very effective judgments of guilt and innocence even if one does not know exact values for the relevant statistics of the reference population. Of course, judgement error increases if the judged distribution of uncertainty diverges from the true distribution to a greater degree, for example, the cross-entropy for 85% judged probability of guilt would be

$$H(Y, Y' | hid) = .74 \times \log_2 \left(\frac{1}{.85} \right) + .26 \times \log_2 \left(\frac{1}{.15} \right) = 0.89 \text{ bits},^9 \quad (31)$$

⁹ As a point of reference, this result is similar to the judgment error achieved if 69% of identified suspects are guilty and decision makers appropriately adopt a 69% level of certainty in guilt for identified suspects, $H(Y) = .69 \times \log_2 \left(\frac{1}{.69} \right) + .31 \times \log_2 \left(\frac{1}{.31} \right) = 0.89$.

and so divergence rises to,

$$D_{KL}(Y, Y' | hid) = H(Y, Y' | hid) - H(Y | hid) = 0.89 - 0.82 = 0.07 \text{ bits.} \quad (32)$$

The divergence also increases as the judgment gets close to 0% or 100%. In short, misinterpretation of evidence imposes a cost to informational value, and EIG provides a measure that appropriately quantifies this cost with information theory metrics.

In the previous example, we computed cross-entropy for a single response category. To determine EIG, the cross-entropy is computed for all response categories and then combined as previously, i.e., as a weighted average. We now use this technique, in conjunction with the Colloff and Wixted (2020) data, to explore the impact of miscalibration on EIG, with a particular focus on the effect of miscalibration on the relative performance of alternative procedures.

Figure 6 shows EIG for the showup and lineup data across various levels of miscalibration. The x -axis is the average absolute deviation between the actual and judged posterior probabilities across all response categories. Positive and negative values represent over-certainty (e.g., actual=.74, judged=.80) and under-certainty (actual=.74, judged=.65) of guilt or innocence.¹⁰

EIG is highest when evidence is properly interpreted (an offset of 0), and decreases with larger deviations between actual and judged posterior probabilities. Over-certainty decreases

¹⁰ The curves in Figure 6 were created in three steps. First, the proportions for all four response categories (e.g., high-confidence suspect ID) were made either more or less extreme by $X\%$. For over-certainty, which we express as positive miscalibration values, the probabilities were made more extreme. That is, if $X=23$, then the judged probability of guilt for high-confidence suspect IDs moved from .74 to .80 [$=.74+.23 \times (1.00-.74)$]. Proportions below .50 move towards 0. For under-certainty, which we express as negative miscalibration values, the probabilities were made less extreme. That is, if $X=37$, then the judged probability of guilt for high-confidence suspect IDs moved from .74 to .65 [$=.74-.37 \times (.74-.50)$]. Proportions both above and below .50 move towards .50. Second, for a given X , the miscalibration offset is defined as the absolute deviation between the true and judged probability of guilt averaged across the four response categories. For example, for over-certainty with $X=23$, the miscalibrated judgments are .80, .75, .09, and .27 producing a mean absolute deviation of .06 [$=(|.80-.74|+|.75-.67|+|.09-.12|+|.27-.35|)/4$]. Third, to create the curves, X was allowed to vary from 0 (calibrated) to 95, for both over- and under-certainty, and the miscalibration offset (x -axis) and EIG (y -axis) was determined for each X .

informational value more severely than under-certainty. Critically, the relationship between EIG values for the showup and lineup is very consistent over different levels of miscalibration.¹¹ In other words, the conclusion based on applying EIG under the assumption of properly interpreted evidence holds even if evidence is misinterpreted, as long as the misinterpretation is not egregious.

Note that EIG can go below zero with extreme miscalibration. That is, if evidence is severely misinterpreted, it can become a source of *misinformation*, in that judgments of guilt and innocence are more effective without knowledge of the witness response. Within the range of positive EIGs, however, the EIG difference between showups and lineups is remarkably consistent. These results suggest that conclusions drawn under the assumption of proper interpretation will be robust to low or moderate levels of miscalibration.

Based on these results, and because our focus is on best-practices in which judgments are well-calibrated, nearly all of the demonstrations below assume that evidence is appropriately interpreted. The one exception is *Demonstration 6*, which explores a situation in which decision makers fail to recognize that non-identifications provide evidence for innocence (Wells & Lindsay, 1980) and, instead, assume that they are not informative.

It is important to note the obvious: the effect of miscalibration is not unique to EIG. Miscalibration can occur for any measure of probative value. Miscalibration occurs because the person interpreting the identification results has an inaccurate impression of the statistics of the reference population or, perhaps more commonly for many real-world investigations, fails to consider the reference population in any sense and adopts a level of certainty based on intuition.

¹¹ An analysis in which the miscalibration of IDs and non-IDs went in opposite directions, i.e., IDs are over-certain and non-IDs are under-certain and vice versa, produced qualitatively similar results.

As explained above, all measures of identification performance are based on the reference population, and all are subject to the deleterious consequences of making unjustified conclusions. For example, although the showup data in Table 1 indicate that low-confidence identifications are twice as likely for guilty (.10) than innocent (.05) suspects, i.e., a diagnosticity ratio of 2, such a response could be taken to be a much stronger indicator of guilt. In other words, a criminal justice official might be under the impression that low-confidence identifications are 4 times more likely for guilty than innocent suspects, or at least behave as if this is the case (i.e., make decisions based on an unjustifiably high degree of certainty). What *is* unique to EIG, at least to our knowledge, is the ability to quantify the effects of miscalibration. The theory behind EIG provides a principled measure of the cost of misalignment between judgments and the true statistics of the reference population. This issue deserves further research, which we mostly leave for future work.

Conceptual Summary of Information Theory and Identifications

The previous sections explained our proposed method for defining the overall evidentiary value of an eyewitness identification. This conceptual summary provides a synopsis of those results and is written in sufficient detail for readers who may have skipped the previous sections. We used data sets from Colloff and Wixted (2020) to explain the concepts that underlie the Expected Information Gain (EIG) measure and to demonstrate how EIG can be calculated. To aid in the interpretation of EIG, we also introduced Probability Gain from EIG (PG_{EIG}), which translates EIG from the entropy to the probability scale and provides context by projecting the amount of information represented by the EIG value onto a simplified identification scenario with a 50% prior probability of guilt and with equal evidentiary value across all responses.

All measures of evidentiary value, including EIG, are based on the statistical properties of a reference population. The reference population is defined by a set of identification attempts that share similar properties, e.g., identification attempts using the same experimental paradigm or from suspects accused of similar crimes in the same area (Wells, Campbell, Li, & Swindle, 2016). All existing measures of evidentiary value are incomplete in the sense that they are unaffected by changes in some subset of the reference population statistics. Our goal was to develop a complete measure that incorporates every aspect of the reference population - EIG.

EIG builds off the Bayesian approach. Bayesian inference defines how gaining knowledge of a particular witness response changes the probability that a suspect is guilty (Wells & Lindsay, 1980). In the Bayesian approach, the *prior* and *posterior* distributions represent the state of uncertainty before and after the witness response is observed, respectively. For the Colloff and Wixted (2020) data, the Bayesian approach revealed that high-confidence suspect IDs increased the probability of guilt more than low-confidence suspect IDs, and that high-confidence rejections increased the probability of innocence more than low-confidence rejections. The Bayesian approach also revealed tradeoffs between the showup and simultaneous lineup data. Whereas rejections were a stronger indication of innocence in the showup procedure, suspect identifications were a stronger indication of guilt in the lineup procedure.

The Bayesian approach provides a meaningful and useful measure of the evidentiary value of a particular identification response, e.g., suspect IDs (Wells and Lindsay, 1980). Identification procedures, however, are used across many identification attempts with varying responses, e.g., suspect IDs, filler IDs, and rejections. Thus, measuring the full evidentiary value of a procedure necessitates a method for combining across possible responses. This goal is facilitated by switching from measuring uncertainty on the probability scale, as past application of the Bayesian

approach, to measuring uncertainty with information-theory measures, specifically Shannon surprisal and entropy (Shannon, 1948).

Shannon surprisal is a measure of the degree of error in a probabilistic judgment. A person who assigns a high probability to the true state of the world would not be very surprised to learn the truth, signaling good judgement or low judgment error. In contrast, a person who assigns a low probability to the true state of the world would be very surprised to learn the truth, signaling poor judgment or high judgment error. High judgment error indicates that a decision maker is unlikely to choose options that lead to good outcomes, as their choices will be guided by an inaccurate assessment of reality.

Combining these key innovations from the Bayesian approach and information theory produces a theoretically-grounded method for measuring evidentiary value (e.g., Ramos & Ganzalez-Rodriguez, 2008). The change in average Shannon surprisal, or *entropy*, from the prior to the posterior distribution defines the *information gain*, or the reduction in uncertainty about guilt or innocence achieved by the witness response. Importantly, because entropy measures uncertainty on an equal-interval scale (Shannon, 1948), entropy can be averaged across all identification attempts to produce the *Expected Information Gain (EIG)*, i.e., the overall evidentiary value of a lineup procedure. This property is not shared by raw probability or any other existing measure of lineup performance.

For the Colloff and Wixted (2020) data sets, calculating EIG revealed that the showup procedure produced more information about guilt or innocence than the simultaneous lineup procedure. Although the lineup produced stronger evidence of guilt when the witness identified the suspect, suspect IDs were fairly rare for the simultaneous lineup, and many identification attempts produced filler IDs that were quite low in evidentiary value. Notably, using area under

the ROC, Colloff and Wixted (2020) advocated for the lineup over the showup. These contrasting results illustrate that EIG provides a perspective on identification effectiveness that is different than other existing measures, and we explore this different perspective more thoroughly in the demonstrations below.

This section also included an example to highlight the importance of using an equal-interval measure of certainty when characterizing overall evidentiary value. The example compared hypothetical data sets without and with confidence ratings in the presence of a strong confidence-accuracy relationship. Averaging across confidence levels on an entropy scale appropriately characterized the obvious evidentiary value of considering witness confidence, producing a higher EIG value for the procedure with confidence ratings. In contrast, averaging on a probability scale obscured the evidentiary value of confidence ratings; that is, when averaging on a probability scale, there was no advantage to including confidence ratings – a clearly misleading result. This example is a specific demonstration of a general principle. Because entropy is an equal-interval measure of certainty, entropy values can appropriately express average certainty. Because probability is not an equal-interval measure of certainty, averaging probability values produces artifacts.

As a complete measure, EIG is sensitive to and can be calculated for any prior probability of guilt. Although not unique to EIG, this property is important because the proportion of guilty suspects in the reference population is likely to vary widely in different real-world scenarios (e.g., different police departments or suspects accused of different crimes; Wells, Yang, & Smalarz, 2015). Extending the Colloff and Wixted (2020) example, EIG revealed that the showup consistently produced more overall information about guilt or innocence than the simultaneous lineup across a broad range of potential prior probabilities of guilt.

Judgments of lineup evidence can be misinterpreted or miscalibrated. For example, an investigator may be 80% certain of guilt after observing a witness response that actually indicates guilt only 70% of the time. We show that EIG can be used to quantify the effects of miscalibration, as EIG values progressively decrease with higher levels of miscalibration. Indeed, EIG values can be negative with extreme miscalibration, showing that identification outcomes can become a source of *misinformation* if they are egregiously misinterpreted. This point is critical, as *any* measure of evidentiary value can be subject to miscalibration, but, to our knowledge, only EIG offers an elegant, theory-based method for assessing the effects of miscalibration.

In summary, EIG is a complete measure that considers all relevant aspects of the reference population – namely, the information gain associated with all possible witness responses and how often these responses are observed – to characterize the full evidentiary value of an identification procedure. That is, EIG indicates how much the identification procedure reduces uncertainty about guilt or innocence, on average, across multiple identification attempts. To aid in interpretation, the equivalent probability gain for the EIG value (PG_{EIG}) translates EIG from the entropy scale back to the probability scale to show how much certainty would need to change from a 50% baseline to produce the same amount of information represented by the EIG value. Thus, in direct contrast to all existing measures, EIG and PG_{EIG} allow eyewitness researchers to directly compare different identification procedures with a metric that considers every factor relevant to evidentiary value. Higher EIG and PG_{EIG} values indicate procedures that produce more information about guilt or innocence across multiple identification attempts.

Demonstrations Exploring EIG

The previous sections developed EIG as a measure for comparing eyewitness identification procedures. In the current section, we present six demonstrations that are designed

to highlight key properties of EIG and illustrate how it differs from existing measures. In *Demonstrations 1-3*, we contrast EIG with diagnosticity ratios (*Demonstration 1*, Wells & Lindsay, 1980), area under the ROC (*Demonstration 2*, Mickes et al., 2012), and deviation from perfect performance (*Demonstration 3*, Smith et al., 2019). In each case, we demonstrate that EIG provides a more complete assessment of the effectiveness of an identification procedure when the goal is using the witness responses to improve judgments of guilt or innocence. Indeed, EIG considers *all* factors relevant to a procedure's evidentiary value. Using response time (RT) as an example, *Demonstration 4* shows how EIG can define the evidentiary value of continuous predictors of witness accuracy. This demonstration also shows that EIG can be used to address important research questions that are not accessible to the existing measures. In *Demonstration 5*, we explore the differences between EIG and other measures of probative value using the sequential and simultaneous lineup data from Clark's (2012) meta-analysis. Finally, in *Demonstration 6*, we consider the effect of ignoring non-suspect IDs.

Readers who are not interested in quantitative details can skip to the *Conceptual Summary of EIG Demonstrations* section.

Decision Modeling

Many of the demonstrations below rely on decision models to isolate key processes while also ensuring realistic data patterns. Mathematical models play an important role in eyewitness memory research (Wixted, Vul, Mickes, & Wilson, 2018), as well as research on memory and decision making more generally. Signal-detection models were used to generate example data sets for many of the following demonstrations (Macmillan & Creelman, 2005). Signal-detection theory (SDT) is probably the most broadly applied mathematical model in psychology (Wixted, 2020), and signal-detection models have proven to be quite useful for analyzing identification

data (e.g., Wixted et al., 2018). Using simulated data from a well-validated model allowed us to isolate individual processes that distinguish EIG from the other measures, while, at the same time, ensuring that we are working with realistic data patterns.

The signal-detection model assumes that witnesses base their identification responses on the match between the suspect and their memory for the culprit, with separate distributions of match strength for identification attempts with a guilty suspect and identification attempts with an innocent suspect (for detailed model descriptions, see Cohen et al., 2020; Wixted et al., 2016). Match strength tends to be higher for identification attempts with a guilty suspect, of course, but match strength varies across attempts for both guilty and innocent suspects. This variation is driven by many factors, including factors that affect how well witnesses remember the culprit (e.g., whether or not they got a clear view), factors that affect how much the appearance of a guilty suspect changes from the witnessed crime to the identification attempt (e.g., whether or not the culprit grew a beard), and factors that affect how much an innocent suspect looks like the culprit (e.g., whether or not the culprit and innocent suspect happen to share many facial features). We assumed that variability in match strength was equal for identification attempts with guilty and innocent suspects, an assumption that is often supported by data from eyewitness experiments (Wixted et al., 2016, 2018), in contrast to data from list-learning recognition tasks (e.g., Ratcliff, Sheu, & Gronlund, 1992).

The collective ability of witnesses to discriminate guilty and innocent suspects is represented by d' , the standardized distance between the means of the memory match distributions for innocent and guilty suspects. A d' of 0 means that the witness cannot discriminate between guilty and innocent suspects. As d' increases, the match strength

distributions for guilty and innocent suspects become increasingly distinct, representing an improved ability to distinguish guilty and innocent suspects.

The response criterion is the cutoff for how well the suspect must match the witness' memory of the culprit to justify an identification response. A low, or liberal, criterion means that even a relatively poor match justifies an identification response. In contrast, a high, or conservative, criterion means that a strong match is needed to justify an identification response.

Variations of this basic signal-detection model are described in each demonstration. *Demonstration 4* also uses the diffusion model (Ratcliff, 1978), which implements a dynamic version of the signal-detection process in which the match strength varies over time within a single decision (for more information, see Ratcliff & McKoon, 2008). Although this model has not been applied to lineup procedures (to our knowledge), it has been extensively validated for two-choice decision tasks that match the form of a showup procedure (i.e., see a single stimulus and decide which of two stimulus classes it belongs to; Ratcliff & McKoon, 2008). In fact, if signal-detection theory is the most broadly applied mathematical model in psychology (Wixted, 2020), the diffusion model is a strong candidate for the second-most broadly applied (Wagenmakers, 2009).

We used model-generated performance patterns to explore critical properties of EIG that distinguish it from other measures. These properties are not dependent on the specific assumptions of the models that we applied. One of our examples references a theoretical process that is linked to a specific signal detection model, the filler-curtain effect. For that application, we also evaluate an empirical data set and show that it is consistent with the filler curtain effect.

Demonstration 1: EIG, Diagnosticity Ratios, and Criterion Shifts

The Diagnosticity Ratio One popular approach to evaluating identification procedures is to compare diagnosticity ratios to assess the probative value of witness responses (e.g., Clark, 2012; Price & Fitzgerald, 2016; Wells & Lindsay, 1980; Wells & Turtle, 1986; Wells, Yang, & Smalarz, 2015). The diagnosticity ratio for a given witness response is calculated by dividing the probability of observing the response for a guilty suspect by the probability of observing the response for an innocent suspect. Although diagnosticity ratios can be calculated for any type of response, we will mostly focus on the diagnosticity of suspect IDs, as this is often the outcome of greatest concern to investigators (Wells & Turtle, 1986).

The showup experiment summarized in Table 1 can serve as a simple example of using diagnosticity ratios to evaluate an identification procedure. In this experiment, 62% [= .31/.50] of guilty suspects and 22% [= .11/.50] of innocent suspects were identified by a witness with high confidence, producing a diagnosticity ratio of 2.82 [= .62/.22]. Similarly, the diagnosticity ratio for low-confidence suspect IDs was 2.00 [= (.10/.50)/(.05/.50)]. Higher ratios indicate that identifications are stronger evidence of guilt, so a high diagnosticity ratio is a desirable quality for an identification procedure. In this case, the diagnosticity ratios suggest that high confidence enhances the evidentiary value of witness responses.

Wixted and Mickes (2012) and Rotello and Chen (2016) provided a clear demonstration of the limitations of diagnosticity ratios by exploring how this measure is affected by changes in response bias, that is, a witness' overall willingness to make an identification response. This point is important, so we explored the effect of changes in response bias to illustrate some critical differences between diagnosticity ratios and EIG. For simplicity, we considered a showup procedure and used a signal-detection model to precisely and uniquely manipulate response bias

while ensuring realistic patterns of simulated data. The differences that we demonstrate between EIG and diagnosticity also apply to lineup procedures. We chose a showup procedure because it is easier to define the criterion value that produces unbiased responding in the showup model.

The Effect of Criterion Shifts. Using a signal-detection model for a showup procedure, we varied the response criterion along a full range from very low (liberal) values to very high (conservative) values. Variation in the response criterion could represent the effect of policy choices that influence witness conservativeness via system variables, for example, whether or not the witnesses is warned that the suspect might not be the culprit (Clark, 2012; Wixted & Mickes, 2012). Of course, these policies could never achieve the precise control over witness conservativeness represented by the range of criterion values in our demonstration, but that is exactly why it is instructive to use a “sanitized,” model-based example. We assume a reasonable level of witness accuracy, $d' = 1.5$. For our purposes, it was convenient to make zero the reference point for an unbiased criterion, so we set the position of the match-strength distributions to $d'/2$ for guilty suspects and $-d'/2$ for innocent suspects (Macmillan & Creelman, 2005). For our first simulation, we used a model in which half of the witnesses saw a guilty suspect and half saw an innocent suspect.

Figure 7 summarizes the results of the criterion demonstration. Figure 7A shows how changing the criterion changes identification rates for guilty and innocent suspects (dashed and dotted lines, respectively). Guilty suspects are more likely to be identified than innocent suspects, and identification rates for both guilty and innocent suspects decrease with increasingly conservative responding (higher criterion values).

For each criterion value, Figure 7B shows the associated diagnosticity ratio, which is found by dividing the height of the dashed line by the height of the dotted line at a given x -axis

value in Figure 7A. This panel shows the pattern noted by Wixted and Mickes (2012; see also Rotello & Chen, 2016) - When the criterion is liberal, a high proportion of identifications are made overall and the diagnosticity ratio is low because the numerator cannot exceed 1 and the denominator is close to 1. When the criterion is conservative, a low proportion of identifications are made overall and the diagnosticity ratio becomes arbitrarily large. Thus, using the diagnosticity of identifications to measure evidentiary value helps policy makers understand that conservative responding makes identifications better evidence of guilt; for example, telling witnesses there is a very low chance that the suspect was the actual culprit would strengthen the evidence of guilt provided by an identification (Wixted & Mickes, 2012).

Although this information is useful, policy makers would ideally consider other impacts of making witnesses very conservative, such as changes in the overall probability of obtaining a suspect ID and changes in the evidentiary value of alternative responses. Intuitively, the fact that identifications are strong evidence of guilt is not very useful if identifications almost never happen, so just getting witnesses to respond as conservatively as possible is probably not a good strategy. Moreover, conservative responding makes rejections more ambiguous, i.e., they cannot be interpreted as strong evidence for innocence given that witnesses fail to identify many of the guilty suspects in addition to many of the innocent suspects.

The EIG measure combines all relevant factors; namely, the probability of observing of every possible witness response and the evidentiary value of every possible witness response. The right-hand panels of Figure 7 show how the factors influencing EIG work together for the criterion example.

Figure 7C shows the overall probability of suspect IDs (blue curve) and rejections (red curve) across different possible values for the response criterion, and Figure 7D shows the

associated information gain for identifications and rejections. As before, the information gain values are the difference in entropy from the prior to posterior distribution of uncertainty about the suspect's guilt category. The information gain for suspect IDs increases as responding becomes more conservative. With very liberal responding, witnesses almost always identify the suspect, and information gain for suspect IDs is low because there is almost as much uncertainty about guilt after a suspect ID as before. As the criterion becomes more conservative, suspect IDs become rarer and more concentrated to identification attempts that involved a guilty suspect (innocent suspects rarely have a strong match to the culprit). This criterion shift increases information gain via the same mechanism that increases the diagnosticity ratio, but the information gain measure approaches a value of 1 rather than becoming an arbitrarily high value (recall that 1 bit represents going from complete uncertainty to complete certainty about guilt or innocence; see Equation 5). Moreover, unlike diagnosticity, EIG considers all possible lineup responses, and Figure 7D shows that more conservative responding decreases the information gain for rejections.

Figure 7E shows EIG, which is a product of the information gain for each possible response (Figure 7D) and its corresponding probability (Figure 7C). Generally, comparing Figure 7A and B to Figure 7C and D illustrates that EIG considers more aspects of the reference population than the diagnosticity ratio; indeed, EIG considers *every* aspect of the reference population that is relevant to the informational value of identifications as an evidence source. In this example, EIG is highest for a criterion value of 0, representing unbiased responding.

The patterns in Figure 7C and D provide an explanation for why EIG prioritizes unbiased responding: Response biases simultaneously make one type of response more common overall and degrade the evidentiary value of that response, resulting in a lower EIG value. Thus, whereas

optimizing the diagnosticity of identifications motivates policies that encourage witnesses to respond as conservatively as possible, EIG suggests that the best policy for gaining information about the guilt or innocence of the suspect is to encourage unbiased responding.

A Note on Prior Probability. The criterion-setting example provides an opportunity to highlight an important point: In some situations, changing the prior probability of guilt can change the procedure that produces the highest EIG. Figure 8 shows EIG curves across criterion settings for a 10%, 50%, and 90% prior probability of guilt, using the same signal-detection model as before. Because there is less information to be gained from the witness responses, EIG is lower overall with 10% or 90% prior probabilities than with a 50% prior probability. That is, with a 10% or 90% prior probability of guilt, the prior distribution already conveys a lot of information about guilt or innocence. More critically, compared to a 50% prior probability, the response criterion that maximizes EIG shifts from 0 to 0.31 for a 10% prior probability and to -0.31 for a 90% prior probability. These criterion locations are denoted with dashed and dotted vertical lines in Figure 8. Thus, EIG sometimes prescribes different procedures for different prior probabilities, but note that, even for these relatively extreme 10% and 90% priors, the unbiased criterion value of zero has an EIG value that is very close to the EIG value at the optimal criterion setting, a point we that will elaborate on below.

These changes in the optimal criterion location are a natural product of the asymmetric uncertainty about guilt introduced by prior information (Wells, Yang, & Smalarz, 2015). With a 10% prior probability of guilt, for example, we start with a strong indication that a suspect is innocent. For the 90% of suspects that are in fact innocent, there is not much room to improve judgment success by reducing uncertainty. That is, learning that a suspect is innocent is not very surprising if one already had information indicating a 90% chance of innocence, i.e., $\log_2(1/0.90)$

= 0.15. For the 10% of suspects who are actually guilty, however, there is a lot of room for improvement in judgment success. That is, learning that a suspect is guilty is very surprising, $\log_2(1/0.10) = 3.32$. This inherent asymmetry prioritizes conservative criteria, which make identifications strong evidence of guilt, over liberal criteria, which make rejections strong evidence of innocence. The EIG measure balances the asymmetry in informational value with the frequency of identifications and rejections for each criterion value, revealing the criterion position that optimizes the potential for witness responses to reduce uncertainty about guilt or innocence. Increasing the prior probability of guilt has the opposite effect, and prioritizes procedures with the potential to produce strong evidence for innocence (a liberal criterion, in this example).

The fact that prior information can change the effectiveness of information-gathering procedures is a direct consequence of Bayesian inference and was acknowledged decades ago by eyewitness memory researchers (Wells & Lindsay, 1980). Nevertheless, some researchers might be troubled by the dependency on prior probability, so we now consider the best response to this dependency. The prior probability of guilt for real identification attempts will vary from case to case and is influenced by a range of factors, such as the other available evidence outside of the identification (Clark, 2012) and the policy for including suspects in an identification procedure (Wells, Yang, & Smalarz, 2015). In light of the substantial variability in prior probability of guilt across different real-world identification attempts, the best policy is to assume a prior probability that is maximally robust across different actual prior probabilities, and we will now show that assuming a 50% prior probability of guilt achieves this desired robustness.

Figure 9 shows EIG from a simulated showup procedure for different *actual* prior probabilities of guilt (different plotting symbols) across a range of criterion values that optimize

EIG under different *assumed* prior probabilities of guilt (x -axis). These results are again based on a signal-detection model. Because the effect of prior probability on the optimal criterion value is larger for high d' , we used a high d' value of 3. Thus, the robustness illustrated below would be even greater at lower d' values. Because we are using a higher d' , EIG values will be higher than in Figure 8.

The circles in Figure 9 show the results with a true prior probability of 0.50, and EIG is accordingly highest when the criterion is optimized for an assumed prior probability of 0.50 (which results in an unbiased criterion value; see Figure 8). EIG decreases as the assumed prior probability moves farther below 0.50 (which motivates an increasingly conservative response criterion) or farther above 0.50 (which motivates an increasingly liberal response criterion, see Figure 8).

The open and filled squares show EIG for a true prior probability of 0.25 and 0.75, respectively. As discussed previously (see Figure 8), EIG is lower for these prior probabilities than for a 50% prior probability of guilt because entropy is relatively low even before a witness response is considered, leaving less room for improvement. Although the functions are peaked at the true prior probabilities, the functions are very flat between an assumed prior probability of 0.50 and the true prior probabilities. In other words, using a criterion optimized for a .50 prior probability works well for these true prior probabilities, and using a criterion optimized for the true prior probability would not provide a meaningful improvement. In fact, the difference in PG_{EIG} is only .002.

The open and filled triangles show even more extreme true prior probabilities of guilt of 0.05 and 0.95, respectively. EIG is low across the full range for these functions, but this should not be concerning because entropy is already very low in the prior distributions. Again, the

functions are peaked at the true prior probabilities, but, again, they do not drop much from the true prior probabilities to an assumed prior probability of .50. The difference in PG_{EIG} is still only .01. Overall, Figure 9 shows that there is not much to be gained by using prior probabilities of guilt other than 0.50, even if the true prior probability of guilt is far from this value. Moreover, assuming a prior probability far from .50 can be quite detrimental when this assumption is false; for example, the EIG for true base rates of .25, .5, and .75 drops steeply as the assumed base rate approaches 0 or 1.

An alternative approach would be to try to specify a value – or better yet, a probability distribution – representing the best guess about the prior probability of guilt for real identification attempts, and to seek procedures that optimize information gain under those conditions. For example, we might assume that investigators typically have good reason to believe that a suspect is guilty if they subject that suspect to a lineup, motivating the use of a relatively high prior probability of guilt. To our knowledge, however, there is only one attempt to use data from real investigations to estimate the proportion of guilty suspects. This study was based on a data set for robbery suspects in Houston, and analysis suggested that 35% of the suspects were actually the culprit (Wixted et al., 2016). This estimate could be used as the base rate when calculating EIG; however, the prior probability could be different for different precincts, different crimes, and a whole host of other factors (not to mention different estimation procedures, although the method used by Wixted et al. seems to be accurate across a number of data sets with fair lineups, Cohen, Starns, & Rotello, 2020). Figure 9 shows that calculating EIG with an assumed 50% prior probability of guilt identifies procedures that will be generally effective regardless of the true prior probability of guilt, which eliminates the need to try to anticipate particular prior probabilities.

Summary. In summary, EIG and diagnosticity ratios support different conclusions. Diagnosticity ratios measure the strength of evidence provided by a particular type of response, such as suspect IDs, but take no account of how often witnesses make the response and ignore the evidence provided by alternative responses. In contrast, EIG considers all factors relevant to a procedure's evidentiary value, including the information gain associated with every possible witness response and how often these responses are obtained from witnesses. Although EIG changes for different prior probabilities of guilt, assuming a 50% prior probability of guilt allows researchers to identify procedures that are either optimal or nearly optimal across the full range of prior probability.

Demonstration 2: EIG versus Area Under the ROC

Receiver Operating Characteristics (ROCs), in the eyewitness context, show the relationship between the probability of identifying a guilty suspect and the probability of identifying an innocent suspect across different levels of response bias, which are typically defined based on confidence-rating data (e.g., Mickes et al., 2012). Mickes et al. (2012) demonstrated how ROCs could be used to compare different eyewitness identification procedures, and the ROC methodology has proven to be a valuable tool that has helped to redefine theoretical debates (Lampinen, 2016; Lampinen et al., 2019; Rotello & Chen, 2016; Smith et al., 2018; Wixted & Mickes, 2012; Wixted et al., 2017).

The goal of this demonstration is to show that EIG provides a different perspective on identification performance than the most common ROC-based measure, area under the curve (AUC). We focus on lineup size, a basic choice for any identification procedure. EIG and AUC support different conclusions for this variable, and the basis of this discrepancy hinges on a fundamental difference between the two measures. We first consider model simulations as a

clean demonstration of the difference between the measures, then we evaluate a key model prediction using empirical data from Colloff and Wixted (2020).

Eyewitness researchers have explored lineup sizes from 1 (a showup) to 120 or more (e.g., Laughery, Alexander, & Lane, 1971; Levi, 2006, 2012; Seale-Carlisle et al., 2019). Law enforcement predominantly uses lineups that contain between 4 and 9 photos, as well as mugbook searches that could include many more than 9 faces (Police Executive Research Forum, 2013). Both the empirical data and simulation-based research (Rotello & Chen, 2016) suggest that area under the ROC does not differ meaningfully with lineup size.

We used a signal-detection model to vary lineup size while holding all other factors constant. The model simulates a simultaneous lineup in which the witness sees all lineup members at the same time. Witnesses select the lineup member with the highest match to their memory for the culprit and respond based on this person alone (Clark, 2003; Duncan, 2006; Macmillan and Creelman, 2005, p. 255-259). Alternative decision processes sometime outperform this sort of “take-the-best” model (Wixted, Vul, Wilson, & Mickes, 2018), but we used it because it has been explored more extensively and it provides a good description of empirical data (Cohen, Starns, Rotello, & Cataldo, 2020; Lampinen, 2016; Rotello & Chen, 2016; Smith, Wells, Lindsay, & Penrod, 2017; Wixted et al., 2017).

Two different lineup sizes were simulated using this model, a 15-person lineup and a 1-person lineup (i.e., a showup), both with a d' value of 1.5. We used a large lineup size to accentuate the differences between measures, but the Colloff and Wixted (2020) analyses below will show that the theoretical mechanism identified is still relevant to more typical lineup sizes (6). To generate the confidence ratings needed to generate an ROC, we assumed that witnesses

made confidence ratings from 0-100 in steps of 10 (i.e., 0, 10, 20 ..., 90, 100), which requires 10 response criteria for both identifications and rejections (see the caption of Figure 10 for details).

Figure 10 shows the ROCs for the simulated 15-person and 1-person lineups. The first (left-most) ROC point only includes identification responses with a confidence level of 100 as identifications, with the x - and y -coordinates determined by the proportion of ID attempts with guilty-suspects and innocent-suspects that meet this standard. The second point includes identification responses with both confidence 100 and confidence 90, and each point adds the next confidence level until the last point, which includes suspect IDs at all confidence levels.

The 15-person lineup ROC is shorter than the showup ROC because of a process known as *filler siphoning* whereby identifications that might have been suspect selections are realized as filler selections instead (Wells, 2001; Wells, Smalarz, & Smith, 2015). In fact, if innocent suspects are no more likely to match the guilty suspect than the fillers (i.e., if lineups are fair; e.g., Colloff, Wade, & Strange, 2016), then the maximum identification rate for lineups with innocent suspects is $1/LS$ where LS is the lineup size (Mickes et al., 2012)¹². This value for the 15-person lineup is marked by the dashed vertical line in Figure 10. Note that the model used for this demonstration implements the filler siphoning process: When an innocent suspect passes the criterion level of match strength required to make an identification, there is a chance that one or more of the fillers will nevertheless exceed the suspect's match strength and prevent a false suspect ID.

The effectiveness of identification procedures is usually measured by the area under the ROC curve (Mickes et al., 2012). For lineup ROCs, it is customary to evaluate the partial area

¹² The maximum false identification rate is defined by assuming that witnesses always make a selection; that is, they never make a rejection. The highest *observed* false identification rate will be lower than the maximum if witnesses sometimes reject the lineup. For example, the maximum possible false identification rate for a showup is 1, but the highest observed false identification rate (the right-most ROC point) will usually be much lower than this.

under the ROC up to the maximum false identification rate (Mickes et al., 2012). Figure 10 shows that the ROCs for the 15-person lineup and the showup overlap almost completely across the range of possible false identification rates for the former. Thus, the partial area under the ROC would be very similar for the two procedures. Rotello and Chen (2016) also demonstrated little or no change in partial area under the ROC for lineups with lengths from 1 to 6. In other words, an ROC analysis suggests that changing the lineup size over a wide range has no meaningful impact on evidentiary value when all other factors are held constant.

Lineups vs. Showups. We will now consider how lineup size highlights differences between the EIG and ROC measures.

The procedure for calculating EIG from confidence data for the 15-person lineup is shown in Table 3. For simplicity, this table only shows the steps for a subset of the possible confidence levels. The full version is provided in the Appendix C. The signal-detection model is used to determine the likelihood of each response, x , for guilty, $p(x|g)$, and innocent, $p(x|i)$, suspects. From these likelihoods, we compute the overall probability of each response, $p(x)$ the posterior probability of guilt for each response, $p(g|x)$, the posterior entropy associated with each response, $H(Y|x)$, the information gain for each response, $IG(x)$, and, finally, the Expected Information Gain by weighting the information gain of each response by the probability of obtaining the response from a witness, $p(x) \times IG(x)$, and summing across responses. For the 15-person lineup, the resulting EIG value is 0.16. Following the same process for the showup yields an EIG of 0.32, indicating that the showup procedure has substantially higher evidentiary value compared to the 15-person lineup procedure. The corresponding PG_{EIG} values are 32% for the showup and 23% for the 15-person line-up. Thus, the procedures substantially differ in the

amount of information about guilt or innocence that they provide, a difference that is not evident in the ROCs, which overlap each other nearly completely (see Figure 10).

This example suggests that EIG has the potential to redefine research questions in the eyewitness memory literature. Reviewing previous identification studies reveals that the field has come to a clear consensus on at least one conclusion: Showups are a bad identification procedure. In fact, a recent paper noted that showups are “much maligned” in the literature and jovially called them “the bane of all identification procedures” (Wells, Yang, & Smalarz, 2015, p. 116). By factoring in all relevant aspects of the identification procedure and using a principled measure of uncertainty, however, EIG reveals a fairly shocking possibility: If all other factors are held constant, showups have more evidentiary value than lineups!

We hasten to add that we are not advocating the use of showups as an empirical or practical conclusion. The showup advantage relies on holding all other factors constant between showups and lineups, and past results convincingly show that these two procedures differ in a number of ways that affect evidentiary value. We provide two possible reasons here. First, in some circumstances, showups can promote a liberal bias to identify a suspect (Kassin, Ellsworth, & Smith, 1989), and Figure 7 shows that bias can severely reduce the evidentiary value of witness responses. Second, the opportunity to compare faces in a lineup helps witnesses discriminate guilty and innocent suspects based on their memory for the culprit (Colloff & Wixted, 2020), and higher witness accuracy increases the evidentiary value of witness responses. Thus, although we are not necessarily advocating showups over lineups, the comparison here does suggest the intriguing possibility that lineups have an inherent cost that accompanies their benefits (see Smith & Ayala, submitted, for a similar claim). We will now consider the

theoretical basis of this cost. After that, we provide an example of how the cost could be investigated experimentally using data from Colloff and Wixted (2020).

The Role of Fillers. Why do showups provide better information about guilt than lineups when all other factors are held constant? Showups have one very useful property when it comes to judging the guilt or innocence of a suspect based on a witness response: The witness always responds to the suspect. The suspect is the only person considered, and the witness must say whether or not they think the suspect was the culprit and, in some cases, provide a confidence rating. This advantage is lost when fillers are added to the identification procedure. Adding fillers introduces the possibility that a witness will fail to respond to the suspect and instead respond to one of the fillers. In hindsight, it seems obvious that it is difficult to judge the guilt of the suspect based on the witness response if the witness does not respond to the suspect. Parallel to our work defining this cost with EIG, Smith and Ayala (submitted) recently demonstrated the same cost by simulating ROC functions that include all possible responses.

We will illustrate this principle by exploring the distributions of match strength across a large range of lineup sizes from 2 to 100. As a reminder, the signal-detection model applied here assumes that witnesses choose the one lineup member that best matches their memory for the culprit and then select a response (and a confidence rating) based on the match strength for this lineup member (Clark, 2003; Duncan, 2006; Macmillan and Creelman, 2005, p. 255-259; Wixted et al., 2018). Sometimes, the lineup member with the strongest match will be one of the fillers, and this result is more likely to occur with larger lineups, which introduce more chances for a filler to randomly provide a good match to the witness' memory for the culprit.

As an example of this principal, consider the simulations shown in Figure 11. The top panel shows match-strength distributions across all guilty and innocent suspects (the filler

strength distribution, not shown, is identical to the innocent suspect distribution). Using these distributions, and simulating lineup sizes of 2, 10, and 100, produces a 50%, 90%, and 99% probability that a filler match strength will be stronger than an innocent suspect match strength. The corresponding probabilities that a filler match strength will be stronger than a guilty suspect match strength are 8%, 33%, and 68%. So, with large lineups there is a high probability that the witness will fail to respond to the suspect because the suspect match strength is exceeded by a filler match strength.

This example illustrates that eyewitness researchers should consider an additional metaphor for the role of fillers in a lineup procedure. Fillers function not only as a *siphon* that reduces the number of false suspect IDs (Wells, 2001; Wells, Smalarz, & Smith, 2015) but also as a *curtain* that obscures information about the suspect's guilt or innocence. To appreciate why a larger lineup size degrades the evidentiary value of witness responses, we can look "behind the filler curtain" to evaluate whether the match strength values for suspects who "lose" to fillers contain valuable information about guilt. In a standard eyewitness experiment, the investigator does *not* have access to this information, because, on these trials, the witness responded based on their memory for one of the fillers, not the suspect (in our metaphor, the suspect is behind the filler curtain). If the suspect match distributions convey strong information about guilt on these trials, then the investigator missed out on valuable information.

The bottom panel of Figure 11 shows conditional distributions of match strength for suspects when the strongest match was from a filler (i.e., filler-victory lineups). These distributions are shifted down relative to the distributions across all suspects in the top panel. This shift occurs because the conditionalization selects suspects who tend to have relatively low match strengths (making them susceptible to losing to a filler).

With a lineup size of 2, the selection effect is very strong, and so the distributions for guilty (solid line) and innocent (dashed line) suspects overlap noticeably more than the distributions across all suspects in the top panel. This result means that the information obscured by filler victories is of moderate value for discriminating guilty and innocent suspects. In other words, not much information about guilt has been lost behind the filler curtain, and getting a response based uniquely on the suspect's match to the culprit would not dramatically improve judgments of guilt and innocence.

As lineup size increases, the selection effect decreases because, with many fillers, the suspect match can be relatively high and still have a good chance of being exceeded by a filler match. This abatement of the selection effect has a larger influence on guilty than innocent suspects, a fact that is easiest to see in comparing lineup sizes 10 and 100 (red and green). The distributions for innocent suspects (dashed lines) change very little between these two lineup sizes, but the distributions for guilty suspects (solid lines) show that high match strengths are more common for a lineup size of 100 than a lineup size of 10. Thus, the conditional match-strength distributions for guilty and innocent suspects have less and less overlap as lineup size increases. The dispersed distributions for guilty and innocent suspects for large lineup sizes mean that investigators are missing out on some very valuable information. In summary, Figure 11 shows that using more fillers puts more information about guilt behind the filler curtain, a pattern that is not evident in an ROC analysis.

A Filler Curtain Example with Empirical Data. Colloff and Wixted (2020) included a condition that was closely matched to a standard simultaneous lineup, but in which witnesses only responded directly to the suspect, that is, they could not respond to the fillers. They called this procedure a *simultaneous showup*. Comparing the simultaneous showup, in which filler IDs

are not permitted, to the standard simultaneous lineup, in which filler picks are permitted, provides an excellent opportunity to look for the pattern predicted by the filler curtain hypothesis. The simultaneous showup was designed to test the prediction from diagnostic feature detection theory (Wixted & Mickes, 2014) that including comparison faces in a lineup, i.e., the fillers, would enhance the ability to discriminate guilty and innocent suspects relative to a standard showup. ROC analysis confirmed this prediction, in that AUC increased from a standard showup to the simultaneous showup, with the latter producing an AUC value similar to a standard simultaneous lineup. Our primary interest is whether eliminating the possibility of filler selections increases EIG for the simultaneous showup relative to the standard simultaneous procedure, as predicted by the filler curtain hypothesis.

We calculated EIG for the data from the simultaneous showup and simultaneous lineup conditions of Colloff and Wixted (2020). To promote stable response proportions, however, we collapsed the original confidence scale into three levels that were used with roughly equal frequency (0-60, 70-80, and 90-100).

EIG was 0.18 for the standard simultaneous lineup and 0.25 for the simultaneous showup. Thus, these data demonstrate the filler curtain effect, as EIG was lower in the procedure that allowed filler selections to prevent direct responses to the suspect. These data also show that EIG and AUC can produce different conclusions, as Colloff and Wixted reported no difference between the two procedures. Indeed, partial AUC values were .03 for both procedures. Colloff & Wixted (2020) also ran a standard showup condition. Interestingly, EIG for the standard showup, 0.23, was also higher than for the simultaneous lineup. Relative to the standard simultaneous lineup, Colloff and Wixted report a deficit for the standard showup, with a partial AUC of .01, illustrating that EIG and AUC can produce opposite conclusions for the same data.

This demonstration reveals that EIG and area under the ROC can support different conclusions for a key procedural choice in any identification procedure, i.e., lineup size. In the remainder of the section, we further discuss the differences between EIG and area under the ROC that underlie these sorts of differences in interpretation.

The Fundamental Difference between EIG and ROCs. The preceding demonstration is one manifestation of a general principle: EIG and area under the ROC measure different constructs. Whereas area under the ROC is designed to measure how effectively *witnesses* can discriminate guilty and innocent suspects, EIG measures how effectively the witness *responses* can be used as a source of evidence about guilt and innocence (Wells, Yang, & Smalarz, 2015, highlighted this same distinction in comparing the ROC and Bayesian approaches). Thus, from the perspective of EIG, the ultimate judgments of guilt and innocence are made not by the witnesses themselves, but by criminal justice officials with access to the witness reports, representing people in the role of, say, jurors, detectives, or judges.

Figure 12 offers a simple demonstration of this difference between EIG and ROCs. The curves show the ROCs implied by the most basic signal-detection model for a showup procedure, with equal-variance match-strength distributions for guilty and innocent suspects. d' is the distance from the mean of the innocent-suspect match distribution to the mean of the guilty-suspect match distribution, and c is the distance of the identification criterion (the minimum match strength needed to identify a suspect) from the midpoint of the two distributions. Higher d' values signify a greater difference in match strength between guilty and innocent suspects; in other words, high d' means that witnesses generally remember the culprit clearly. Every point along a given function has the same d' value; in other words, calculating d'

serves the goal of measuring memory discriminability while negating the influence of differences in response bias (Macmillan & Creelman, 2015).

The points in Figure 12 correspond to different possible positions for the response criterion, where a value of zero indicates unbiased responding, negative values indicate liberal responding (i.e., willingness to make an identification even for relatively poor matches to memory for the culprit), and positive values indicate conservative responding (i.e., reserving identifications only for very strong match values). The number beside each point is the EIG value for that true and false ID rate with a 50% prior probability of guilt.

For unbiased responding ($c = 0$), EIG is highest for $d' = 2$ and decreases with lower d' values. This result is intuitive: If witnesses have poor memory for the culprit, then their identification responses cannot provide much information about the guilt or innocence of suspects. For a given level of witness memory, EIG also decreases with increasingly liberal or conservative responding. This result is perhaps less intuitive, but becomes clear when one thinks about limiting cases. The bottom-left edge of the curves represents a state in which no guilty suspects and no innocent suspects are identified, and the top-right edge represents a state in which all guilty and all innocent suspects are identified. If identification attempts always produce the same outcome regardless of suspect guilt, then they have no value as evidence. Thus, EIG is not equivalent to measures like d' or AUC, which are designed to measure the overall ability of witnesses to discriminate guilty and innocent suspects regardless of response bias. EIG is designed to accommodate the effect of both discriminability and response bias on the informational value of witness responses.

Critics of the ROC approach have noted another reason why ROCs may fail to capture the full evidentiary value of an identification procedure; namely, ROCs are not affected by the

proportion of non-suspect IDs that are filler IDs versus rejections, nor are ROCs affected by the distribution of confidence for either filler IDs or rejections (Wells, Smalarz, & Smith, 2015; Wells, Smith, and Smalarz, 2015; Wells, Yang, & Smalarz, 2015). This issue is a minor concern if the different types of non-identification responses have similar distributions for guilty and innocent suspects, but these distributions can be starkly different in some circumstances (Wells & Olson, 2002; Wells, Yang, & Smalarz, 2015). EIG considers all responses, and any differences in response patterns between guilty and innocent suspects will increase EIG according to how much these differences improve judgments of guilt or innocence.

The properties measured by EIG and ROCs are closely related, of course, and one can devise many scenarios in which EIG and area under the ROC support the same conclusions about the relative efficacy of identification procedures. However, these measures can also diverge in important respects. For example, as we have already seen, even if we assume that adding more fillers has no effect on the overall distributions of match-to-culprit strength for guilty and innocent suspects, it can still degrade the evidentiary value of witness responses by making it less likely that the witness responds based on the suspect.

Therefore, from the standpoint of policy recommendations, ROCs should not be the only basis for identifying the most effective identification procedures; indeed, ROCs are specifically designed to compensate for some of the factors affecting evidentiary value (Wells, Yang, & Smalarz, 2015). Although ROCs will certainly continue to play an important role in eyewitness memory, researchers and policy makers should also consider EIG, as it provides a complete measure of evidentiary value and reveals the procedure that produces the most information about guilt or innocence across multiple identification attempts.

Summary. This demonstration shows that EIG and area under the ROC can support different conclusions about the relative effectiveness of different identification procedures. For example, using EIG reveals that filler IDs impair the informational value of lineup procedures, especially those with large lineup sizes. This effect of lineup size is not detected by area under the ROC, and so provides a good example of how EIG can inspire new research questions in the eyewitness literature. For example, a corollary of this property is that lineup paradigms in which witnesses always respond to the suspect may be particularly useful (e.g., Wilson et al, 2019). Even if investigators are reluctant to adopt such procedures, EIG suggests that policy makers might want to avoid large lineup sizes even though they reduce false identifications of innocent suspects with no appreciable decrease in area under the ROC.

Demonstration 3: EIG versus Deviation from Perfect Performance (DPP)

Most eyewitness studies that evaluate lineup success with ROCs have used measures based on area under the curve, but ROCs have also been interpreted in terms of the accuracy achieved at individual ROC points. We will focus on one such measure proposed by Smith et al. (2019), *Deviation from Perfect Performance (DPP)*, and use this measure to further explore the implications of using EIG with confidence-rating data.

In its simplest form, the DPP measure is the conditional error rate for guilty suspects plus the conditional error rate for innocent suspects at a given ROC point. Smith et al. (2019) proposed that an identification procedure can be evaluated in terms of the average or minimum DPP value across all points on the ROC (the minimum was favored by Smith et al.). DPP prefers the procedure producing a lower value, as errors are minimized.

For example, consider Figure 13A. The blue circles show an ROC generated from an SDT model by varying the response criterion (see the figure caption for model details). Here we

assume that the criteria are associated with confidence ratings. The red plus symbol shows lineup results generated from the same SDT model, but collapsed across all identification confidence levels and all rejection confidence levels. That is, these data were generated from a simpler model that only has a single response criterion, corresponding to the identification-rejection boundary. At this collapsed point, witnesses identified 77.3% of the guilty suspects and 22.7% of the innocent suspects

Figure 13B shows the DPP value for each individual ROC point. For example, the DPP for the central ROC point (red plus) is .454 $[(1-.773)+.227]$. The horizontal lines mark the minimum DPP and average DPP measures. The lowest DPP value is produced by the procedure that ignores confidence ratings; that is, the model with a single response criterion. If the minimum DPP is used to judge the efficacy of a lineup procedure, then the inclusion of confidence ratings does not improve the evidentiary value of the identification procedure, as the ROC point corresponding to the ID criterion has the lowest DPP value. The average DPP is higher when confidence ratings are considered, indicating worse performance. Thus, if the efficacy of this lineup procedure is judged by the average DPP, including confidence ratings actually *decreases* the evidentiary value.

EIG provides a starkly different evaluation. The blue circles in Figure 13C shows the EIG associated with each individual ROC point. The red plus shows the EIG calculated on the data collapsed across confidence ratings. Consistent with the DPP evaluation, the highest individual EIG is for the central point, which, again, corresponds to the procedure that ignores confidence ratings. The horizontal line indicates the EIG value produced by using the full confidence distributions. In contrast to DPP, the EIG value for the full distribution of confidence ratings is higher than the EIG for any of the individual ROC points, even though those ROC points were

created from the same confidence data. In fact, using the full confidence scale increased EIG from 0.23 to 0.31 (corresponding to PG_{EIG} values of 27% and 31%). Table 4 outlines the process of calculating EIG for both the full confidence data and the central ROC point.

Thus, unlike the DPP measures, EIG suggests that obtaining confidence ratings is very helpful for distinguishing guilty and innocent suspects. The primary reason for this difference is that, unlike DPP, the EIG measure does not collapse across different levels of the confidence scale; instead, it evaluates the information gain associated with each confidence level and then averages across levels weighted by their probability. High-confidence responses can be associated with substantial information gain (e.g., Wixted et al., 2016), so segregating them from the lower confidence responses reduces overall uncertainty about guilt, as illustrated in Table 4

Summary. In summary, EIG and DPP will often disagree as measures of the effectiveness of an identification procedure: Whereas DPP collapses the confidence scale to determine the error rate for each individual ROC point, EIG considers the evidentiary value of each separate confidence level and then averages across all identification attempts. This example shows that DPP, and indeed *any* measure based on individual ROC points, cannot capture the full evidentiary value of confidence rating data. Furthermore, this demonstration suggests that allowing witnesses to provide confidence ratings can meaningfully enhance the evidentiary value of an identification procedure.

Demonstration 4: EIG and Continuous Variables – A Response Time Example

We have shown that EIG has different properties and can support different research conclusions compared to other common measures of the effectiveness of an identification procedure. This section illustrates how EIG can promote new discoveries by allowing

researchers to address questions that cannot be effectively answered with existing measures. In particular, we present a theoretical demonstration of how EIG can assess the value of using identification response times (RTs) to predict the accuracy of witness responses.

We chose RT as a convenient example that has a well-known theoretical relationship with accuracy (Ratcliff & McKoon, 2008), but we stress that these methods could be used for any variable, or combination of variables, that could be used to predict witness accuracy (Brewer et al., 2020). Such variables include estimator variables (Wells & Olson, 2003) such as viewing distance, viewing time, level of intoxication, and many others. EIG offers a principled method for assessing which estimator variables are important to consider by quantifying the degree to which an estimator variable improves judgments of guilt and innocence. Moreover, EIG does not require that researchers discretize continuous predictors, in contrast to previous applications of the Bayesian approach (e.g., Chapman & Cahil, 2012, Footnote 4).

For our demonstration, we generated responses from a showup task using the diffusion model, which is, essentially, a dynamic version of signal-detection theory (Ratcliff & McKoon, 2008; see the Figure 14 caption for model parameters). Figure 14 shows the resulting RT distributions, which have the positive skew that is universally observed for empirical RT distributions in standard decision tasks (Wagenmakers, 2009). Another standard pattern is that errors (i.e., identifying an innocent suspect or rejecting a guilty suspect) tend to be slower than correct responses (Ratcliff & McKoon, 2008), and this property is important for the current example because it means that RT can be used to predict the accuracy of a witness response, and in turn, the probability that the suspect is guilty or innocent given the witness response (Smith et al., 2000; Sporer, 1993).

We used a standard logistic regression to predict the accuracy of witness responses from RT, where the response variable coded the suspect's guilt (1) or innocence (0) and the predictor variable was the z -score of the log-transformed RT (we took logs to better approximate a normal distribution and z -scores to aid in interpretation of slope coefficients). We ran separate models for identification and rejection responses. The maximum likelihood estimates of the intercept coefficient were 1.87 for identifications and -1.98 for rejections, indicating that the probability of guilt was higher overall for identifications (0.87) than for rejections (0.12). The estimated slope coefficient for identifications was -0.99 , meaning that the predicted probability of guilt was highest for fast identifications and decreased for slower identifications. The slope coefficient for rejections was 0.94, meaning that the predicted probability of guilt was lowest for fast rejections and increased for slower rejections. Overall, these slope coefficients show the desired pattern: Fast responses were more likely to be correct than slow responses, so fast identifications provide particularly strong evidence for guilt and fast rejections provide particularly strong evidence for innocence.

The logistic equations return the predicted probability that the suspect is guilty for every identification attempt in the data set from Witness w based on Response x_w and Response Time t_w . To calculate EIG, we used these predictions as the posterior probability of guilt, $p(g | x_w \& t_w)$, or innocence, $p(i | x_w \& t_w)$, for each identification attempt, calculated the resulting change in Shannon surprisal relative to the appropriate prior (e.g., Equation 21), and averaged these values across attempts (i.e., across witnesses). We assume here that each witness sees either a single guilty or innocent suspect. Thus, the information gain, IG, or change in Shannon surprisal, is given by

$$IG(x_w \& t_w) = \log_2 \left(\frac{1}{p(g)} \right) - \log_2 \left(\frac{1}{p(g | x_w \& t_w)} \right) \quad (33)$$

for a guilty suspect, and

$$IG(x_w \& t_w) = \log_2 \left(\frac{1}{p(i)} \right) - \log_2 \left(\frac{1}{p(i|x_w \& t_w)} \right) \quad (34)$$

for an innocent suspect. The expected information gain is then calculated as

$$EIG = \frac{1}{N} \sum_w IG(x_w \& t_w), \quad (35)$$

where N is the total number of identification attempts. For the simulated data set, the resulting EIG value is 0.47, corresponding to a PG_{EIG} of 38%.

To assess the value of using RT as a predictor of witness accuracy, we compared this value to the EIG value produced by using witness responses without considering RT. This comparison EIG value is 0.37, corresponding to a PG_{EIG} of 34%. Thus, using RT as a predictor of witness accuracy improved the evidentiary value of the witness responses. The EIG difference is clearly influenced by the strength of the relationship between the predictor variable and witness accuracy, but it is not just another way of expressing this relationship strength. Instead, the EIG difference combines all relevant factors to measure how much using the predictor variable can reduce uncertainty about guilt, given the witness responses.

Summary. The outcome of a given identification attempt is influenced by a wide range of continuous estimator variables, such as the viewing distance, the amount of time that the culprit was in view, and the retention interval separating the witnessed crime and the identification attempt. Considering these variables can improve the evidentiary value of identification responses; for example, criminal justice officials who obtain a suspect ID could judge a higher probability of guilt if the witness saw the culprit from a short distance for a long period of time and responded to a lineup the next day than if the witness saw the culprit from far away for a brief moment and responded to a lineup months later.

EIG can be used to quantify the evidentiary value of continuous predictors, and this section provided a simple, but illustrative, example of this functionality by using RT as a predictor of witness accuracy for showup data generated by the diffusion model. We computed two EIG values. One EIG value was computed from posterior probabilities that were based on the overall response rates and ignoring RT, as in previous examples. The other EIG value was computed from posterior probabilities of guilt that were estimated by a logistic regression model using RT as a continuous predictor within each response category. EIG increased for the latter case, quantifying the amount of additional information about guilt or innocence obtained by considering RT. This method can be easily generalized to any prediction model, including those that consider multiple predictors (e.g., Brewer et al., 2020). This quality makes EIG an important tool for addressing novel research questions in the eyewitness memory literature.

Demonstration 5: Probative Value

In this demonstration, we further explore the differences between EIG and other measures of probative value using the sequential and simultaneous lineup data from Clark's (2012) meta-analysis. We chose this data set because it provides a clear example of the unique perspective EIG brings to identification performance and helps illustrate why this perspective should be of interest to policy makers and eyewitness researchers.

Clark (2012) Data and Measures. Table 5 shows two reference populations based on Clark's (2012) estimated performance levels for simultaneous and sequential lineups with a 50% prior probability of guilt. These data only distinguished suspect IDs (*Sus ID*) and non-suspect-IDs (which we shorten to *Non ID*), without considering more detailed response patterns, such as differentiating filler IDs and rejections or responses made at different confidence levels. The *ID*

rate column provides the identification rate for guilty and innocent suspects, e.g., for simultaneous lineups, 54% $[=.27/(.27+.23)]$ of guilty suspects were identified.

The measures of probative value investigated by Clark (2012) help policy makers understand how lineup format affected the evidentiary value of individual identification outcomes (suspect IDs or non-identifications). Here, we will focus on diagnosticity ratios and innocence risk, and we will also discuss change in the probability of guilt (Wells & Lindsay, 1980). For the data in Table 5, the diagnosticity ratios for suspect IDs were 3.6 $[=.54/.15]$ and 4.78 $[=.43/.09]$ for simultaneous and sequential lineups, respectively, indicating that suspect IDs were more diagnostic of guilt in the sequential procedure. Innocence risk tells a similar story, as there was a 22% chance that an identified suspect was innocent for the simultaneous procedure compared to just a 17% chance for the sequential procedure. If policy makers followed a Bayesian approach (Wells & Lindsay, 1980), they could alternatively note that the posterior probability of guilt for identification responses was lower in the simultaneous (78%) than the sequential (83%) lineup.

For non-IDs, the diagnosticity ratios for innocence were 1.85 $[=(1-.15)/(1-.54)]$ and 1.60 $[=(1-.09)/(1-.43)]$ for simultaneous and sequential lineups, respectively, indicating that non-identifications were more diagnostic of innocence in the simultaneous procedure. So the sequential advantage in the diagnosticity of suspect IDs is offset by a disadvantage in the diagnosticity of non-IDs, but the difference between procedures is much smaller for non-IDs. This sequential disadvantage for non-IDs is also apparent in the posterior probability of innocence for identification outcomes without a suspect ID, which was 65% for the simultaneous procedure and 61% for the sequential procedure. For both procedures, the evidence of guilt produced by suspect IDs is stronger than the evidence of innocence produced by non-IDs. That

is, compared to non-identifications, suspect IDs have diagnosticity ratios farther from 1 and posterior probabilities of guilt farther from 50%.

All of these results provide information to policy makers considering the relative merits of simultaneous versus sequential lineups. As discussed previously, these measures are designed to isolate specific aspects of performance, and as such, they do not provide an overall assessment of evidentiary value. We have argued that it is also useful for policy makers to consider all of the implications of choosing a particular procedure. To demonstrate how using a complete measure can impact interpretations, we now consider EIG for this example.

Expected Information Gain. We calculated EIG using the same method outlined above: express prior and posterior uncertainty about guilt or innocence as entropy values; compute the information gain for each witness response as the reduction in entropy; and take the weighted average across responses. For the simultaneous procedure, the prior entropy is 1.00 and the posterior entropy is 0.76 and 0.93 for suspect IDs and non-IDs, respectively, corresponding to information gain values of 0.24 and 0.07. Using the overall response probabilities from Table 5, EIG is 0.13 [$=.345 \times 0.24 + .655 \times 0.07$] and PG_{EIG} is .21; on average, a witness response brought certainty of guilt or innocence from 50% to 71%. For the sequential procedure, the posterior entropy is 0.66 and 0.96 for IDs and non-IDs, corresponding to information gain values of 0.34 and 0.04. EIG is 0.12 [$=.26 \times 0.34 + .74 \times 0.04$], and PG_{EIG} is .20. Thus, on average, a witness response brought certainty of guilt from 50% to 70%.

These results show that the simultaneous and sequential formats had nearly identical informational value for these data. Although suspect IDs are a stronger indication of guilt for the sequential procedure (83% versus 78% certainty of guilt, corresponding to information gain values of 0.34 and 0.24), this advantage is offset by two facts: (1) suspect IDs are rarer in the

sequential procedure (26% versus 34.5% of identification attempts), and (2) non-suspect IDs are a weaker indication of innocence for the sequential procedure (61% versus 65% certainty of innocence, corresponding to information gain values of 0.04 and 0.07).

Summary. In summary, the existing measures of probative value are helpful in determining which procedure acted as a better evidence source *conditioned on a specific type of witness response*. EIG is helpful in determining which procedure acted as a better evidence source *in general*. Policy makers choose procedures that will be used across a large number of investigations, so they should not solely focus on one response category. EIG provides a measure that synthesizes all relevant factors, making it a useful tool for policy decisions.

Demonstration 6: Ignoring Non-identifications.

In this demonstration, we continue to explore the Clark (2012) data used in *Demonstration 5* and consider the effect of ignoring non-suspect IDs (non-IDs). Although some eyewitness researchers have recommended that non-IDs should be recorded and interpreted as evidence for innocence (e.g., Wells & Lindsay, 1980), criminal justice officials have not fully embraced this recommendation. In fact, a survey of 619 law enforcement agencies across the United States found that between 31% and 38% of agencies failed to document non-IDs in showups and lineups (Police Executive Research Forum, 2013, Table 13). Without such documentation, non-IDs cannot be used as evidence for innocence.

This section addresses two questions: (1) How much is evidentiary value affected by ignoring non-IDs? and (2) If criminal justice officials only interpret suspect IDs and not any other identification outcomes, is EIG equivalent to any of the other measures of probative value?

EIG is an invaluable tool for addressing the first question by quantifying the extent to which evidentiary value is affected by ignoring non-IDs. As noted in the *The Effect of*

Miscalibration section above, EIG is able to measure the informational value of an evidence source regardless of how that evidence source is interpreted. Thus, the standard EIG value, which assumes proper interpretation of all witness responses, can be compared to an EIG value that ignores non-IDs as a potential source of evidence. To denote the fact that only identifications are used as evidence, we label the latter measure EIG_{ID}.

Returning to the sequential and simultaneous data from Clark (2012), the bottom two rows of Table 5 report the true distributions of uncertainty for these data following each witness response. If non-IDs are ignored as an evidence source, then the judged posterior distribution of uncertainty associated with non-IDs is equal to the prior distribution. In other words, non-IDs have no influence on the judged probability that a suspect is guilty or innocent. The average judgment error for non-IDs is thus the cross-entropy (see *The Effect of Miscalibration* section) of the actual posterior distribution and the prior distribution of uncertainty about guilt or innocence.

For example, consider non-IDs in the simultaneous procedure. For standard EIG, these responses are correctly interpreted as supporting 65% certainty that the suspect was innocent, yielding a posterior entropy of 0.93 and an information gain of .07. For EIG_{ID}, these responses are ignored, so the judged posterior distribution of uncertainty about guilt is equal to the prior distribution, in this case, a 50% chance of guilt and 50% chance of innocence. The cross-entropy is then given by

$$H(Y, Y' | non-ID) = .35 \times \log_2 \left(\frac{1}{.50} \right) + .65 \times \log_2 \left(\frac{1}{.50} \right) = 1, \quad (36)$$

producing an information gain of 0 [$IG(non-ID) = H(Y) - H(Y, Y' | non-ID) = 1 - 1 = 0$].

We previously reported that EIG was 0.13 and 0.12 for the simultaneous and sequential lineups, respectively, corresponding the PG_{EIG} values of .21 and .20. The corresponding EIG_{ID} values are 0.08 and 0.09 for the two procedures, corresponding to a PG_{EIG} value of .17 for both

procedures. Thus, failing to interpret non-identification data impairs the overall information value for each procedure, and this effect is actually much larger than the difference between procedures.

The cost of ignoring non-identification evidence is potentially higher with more complex response options, e.g., responses that include confidence. For example, high-confidence rejections were a strong signal of innocence for the showup data in Table 1. Thus, these results strongly support the recommendation that decision makers should consider evidence for innocence in addition to evidence for guilt, and EIG provides a way to quantify the cost of failing to do.

We now turn to the second question. Irrespective of whether decision makers *should* interpret responses that provide evidence for innocence, what is the effect if policy makers determine that non-IDs *will not* be considered? In that case, do existing measures give policy makers the same information as EIG? In short, no. The reason is that the existing measures omit a factor that it is important for policy makers to consider: the overall prevalence of identification responses.

For the data from Clark (2012) covered in *Demonstration 5*, the sequential procedure had an advantage in terms of the diagnosticity of suspect IDs (4.78 versus 3.6), innocence risk (.17 versus .22, where small values are better), and posterior probability of guilt following a suspect ID (.83 versus .78, with a 50% prior probability of guilt). Clark (2012) noted, however, that the sequential procedure also led to fewer true suspect IDs, and indeed, Table 5 shows that the *overall* identification rate was only 26% $[=.50 \times .43 + .50 \times .09]$ for sequential lineups compared to 34.5% $[=.50 \times .54 + .50 \times .15]$ for simultaneous lineups.

The EIG_{ID} measure incorporates this difference in identification rates and reveals that the procedures are nearly equivalent in evidentiary value, even if evidence from non-IDs is ignored, with the PG_{EIG} values indicating a change in certainty of guilt or innocence from 50% to about 67% in both cases. In other words, in contrast to the impression given by the other measures, the sequential procedure does not produce more information about guilt or innocence overall even in a situation in which only suspect IDs are considered. The extra information that the sequential procedure provides when a suspect ID occurs is offset by the fact that identifications happen less often in this procedure, and the result is that the sequential procedure provides no extra information about guilt or innocence when both procedures are used across a large number of identification attempts.

Summary. In this demonstration, we used the Clark (2012) simultaneous and sequential lineup data to explore how EIG can quantify the cost of ignoring non-IDs and to demonstrate that EIG sometimes supports different conclusions than existing measures of probative value, even if suspect IDs are the only outcome of interest. EIG decreases dramatically if non-IDs are ignored; that is, if the judged probability of guilt following non-IDs is equal to the prior probability of guilt. Even if criminal justice officials follow the unadvisable practice of only considering suspect IDs, EIG still considers an additional aspect of the reference population ignored by other measures: the overall prevalence of identifications.

To illustrate, imagine a policy maker relying on the Clark (2012) data to assess the effectiveness of these alternative lineup procedures. The measures of probative value considered by Clark (2012) and Wells, Yang, & Smalarz (2015) could help this policy maker understand that the sequential procedure produced better evidence *when a suspect ID was observed*. Considering EIG, however, could help the policy maker understand that the sequential procedure

nevertheless failed to produce more total information about suspect guilt or innocence *when all identification attempts were considered*. EIG provides a different perspective by factoring in important differences in overall response rates, and this perspective is valuable for making effective policy decisions.

To take an extreme example, would police want to use an identification procedure in which a suspect ID yields nearly incontrovertible evidence of guilt, but that only produces suspect IDs in 1 out of 10,000 identification attempts? Clearly, the extreme rarity of suspect IDs would impair the usefulness of this procedure. EIG can help policy makers focus on reforms that will produce the largest potential benefits when all aspects of the reference population are considered.

Conceptual Summary of EIG Demonstrations

We provided six demonstrations of the strengths of using EIG to evaluate the evidentiary value of a lineup procedure. In this section we review the basic results from those demonstrations.

The first demonstration compared EIG to diagnosticity ratios, which are computed by dividing the likelihoods of observing a particular witness response (typically a suspect ID) for guilty and innocent suspects (Wells & Lindsay, 1980). We used a signal-detection model to determine the effect of varying response bias from very liberal to very conservative in a showup procedure, representing situations in which witnesses are very willing or very reluctant to identify the suspect, respectively. The diagnosticity ratio for identifications increases to an arbitrarily high value as responding becomes more conservative (as noted by Wixted & Mickes 2012 and Rotello & Chen, 2016). Conversely, for a 50% prior probability of guilt, EIG is highest when witnesses are unbiased, that is, when they are equally likely to err by identifying an

innocent suspect and by rejecting a guilty suspect. This difference in interpretation arises because diagnosticity ratios and EIG consider different aspects of identification performance. Whereas the diagnosticity ratio ignores many of the factors that are relevant to the evidentiary value of an identification procedure, EIG considers all of these factors.

The first demonstration also explored the effect of differences in prior probability on EIG. Changing the prior probability of guilt can change the relative EIG values of different procedures. With a low prior probability of guilt, EIG is highest when witnesses respond conservatively, and with a high prior probability, EIG is highest when witnesses respond liberally. We noted that the prior probability of guilt is likely to vary across different investigations (Clark, 2012; Wells, Yang, & Smalarz, 2015). Thus, policy recommendations should prioritize procedures that are generally effective across a broad range of prior probabilities. We showed that optimizing EIG under the assumption of a 50% prior probability of guilt achieves this desired robustness. That is, even when the true prior probability of guilt is far from 50%, the optimal procedure identified under the assumption of a 50% prior probability of guilt is nearly as effective as the optimal procedure identified with the true prior probability.

The second demonstration compared EIG and area under the receiver operating characteristic (ROC). ROCs relate the probability of identifying a guilty suspect to the probability of identifying an innocent suspect across different levels of response bias. The area under the ROC is often used to measure the efficacy of a given identification procedure (Wixted & Mickes, 2012). This demonstration showed that EIG and area under the ROC can support different conclusions about the relative effectiveness of different identification procedures. For example, whereas changing response tendencies, such as shifts in overall response bias, can dramatically reduce EIG, area under the ROC is designed to negate the influence of bias

differences (Macmillan & Creelman, 2005). Moreover, any information about guilt or innocence carried by the different types of non-suspect ID responses (e.g., rejections or filler IDs) will be detected by EIG but will have no effect on area under the ROC.

The second demonstration also considered the effect of adding fillers to a lineup. Using a common signal-detection model for lineup tasks (Macmillan and Creelman, 2005, p. 255-259), we showed, surprisingly, that adding fillers to a lineup degrades the evidentiary value of witness responses if all other factors are held constant. In fact, a showup (i.e., a 1-person lineup) produces the highest EIG value. Area under the ROC gave little to no indication of this drop in evidentiary value. Empirical evidence suggests that many factors vary between lineups and showups, and many of these changes are advantageous for the lineup procedure (Kassin et al., 1989; Colloff & Wixted, 2020). Our demonstration suggests that these factors might be hiding an inherent disadvantage of the lineup procedure, that is, that fillers often prevent direct responding to the suspect, obscuring information about the suspect's guilt or innocence. We showed that data from Experiment 3 of Colloff and Wixted (2020) are consistent with this *filler curtain* effect. Compared to a standard simultaneous lineup, EIG was higher in a simultaneous showup condition that prevented filler responses by highlighting the suspect in a 6-person lineup. The filler curtain concept motivates future work on new procedures that maintain the advantages of a lineup while ensuring that the witness always directly responds to the suspect. Although ROCs have proven to be a valuable tool for exploring various research questions in eyewitness research, EIG can add to this toolkit and inspire new research questions.

The third demonstration compared EIG and Deviation from Perfect Performance (DPP), an ROC-based measure that has been proposed as an alternative to area under the ROC (Smith et al., 2019). Each point of a confidence-rating ROC is created by collapsing across a range of

confidence levels. DPP defines the effectiveness of an identification procedure based on the error rate at these individual ROC points. We used a signal-detection model to determine the effect of either including or ignoring confidence ratings in an identification procedure. Whereas DPP found no benefit of adding confidence ratings, EIG revealed that using a full confidence scale meaningfully increased evidentiary value. Indeed, the evidentiary value produced by using a full confidence scale was higher than the evidentiary value produced by collapsing the confidence scale at any confidence level (see Figure 13C). In other words, collapsing a confidence scale to create an individual ROC point erases much of the information about guilt or innocence carried by the confidence data. This demonstration shows that DPP, and indeed *any* measure based on individual ROC points, cannot capture the full evidentiary value of confidence rating data.

The fourth demonstration showed how EIG can be used to assess the value of using continuous predictors of witness accuracy to improve judgments of guilt or innocence. We focused on a particular predictor, response time (RT), i.e., how long it took the witness to make an identification decision. Using the diffusion model (Ratcliff, 1978), we simulated data with a known relationship between RT and accuracy, whereby faster identification decisions were more accurate than slower decisions (Ratcliff & McKoon, 2008). We used logistic regression to predict witness accuracy for suspect IDs and rejections with RT as a continuous predictor. The regression was then used to define the posterior probability of guilt on each identification attempt. The results showed that using RT to predict witness accuracy produced a meaningful increase in EIG. This method could be adapted to any combination of predictors and any prediction model, and indeed, it is a standard metric for evaluating the success of classification algorithms in computer science and engineering (Quinlan, 1986). Because there are a huge range of ecological variables that influence witness accuracy (Wells & Olson, 2003), exploring the

information value of different predictors of witness accuracy is an important direction for eyewitness research. Developing effective methods for predicting witness accuracy could dramatically improve judgments of guilt and innocence, and EIG is a useful tool in the search for better prediction methods.

The fifth demonstration used the simultaneous and sequential lineup data from Clark's (2012) meta-analysis to further demonstrate how EIG provides a different perspective on policy recommendations when compared to existing measures of evidentiary value. Diagnosticity ratios, innocence risk, and probability gain are applied to individual response classes. For suspect IDs, all of these measures suggested that the sequential lineup produced stronger evidence of guilt. When the suspect was not identified (rejections and filler IDs), neither procedure produced strong evidence of innocence, although the simultaneous procedure had a small advantage. Presented with this information, a policy maker might reasonably conclude that the sequential lineup served as a better evidence source. EIG, however, provides a way to combine across identification attempts that both did and did not lead to suspect IDs, and, critically, EIG also accounts for the fact that the sequential procedure is less likely to produce exactly the response that makes it appealing, i.e., a suspect ID. With all aspects of the reference population considered, EIG was very similar for the simultaneous and sequential procedures. In other words, the choice of procedure has little impact on the amount of information about guilt or innocence achieved across multiple identification attempts with varying witness responses.

The sixth, and final, demonstration further analyzed the Clark (2012) simultaneous and sequential data under the assumption that non-IDs are ignored; that is, identification attempts that fail to produce suspect IDs are *not* interpreted as evidence for innocence. Although neither procedure produced strong evidence for innocence, ignoring this evidence still produced a

meaningful drop in EIG. In fact, the difference in EIG within each procedure when incorporating and ignoring non-IDs was much larger than the difference in EIG between the two procedures. As in the fifth demonstration, EIG led to a different conclusion than the other measures of evidentiary value. Although suspect IDs are better evidence for guilt in the sequential procedure, EIG again corrects this advantage for the fact that suspect identifications occur less frequently in the sequential procedure. The result is that the sequential procedure fails to produce more evidence of guilt across all identification attempts. This is an important result for policy makers to consider, as they need to think about how to get the best outcomes across all investigations, not just the investigations that happen to have a suspect ID.

General Discussion

The principles of probabilistic reasoning (e.g., McElreath, 2016) and information theory (Shannon, 1948) provide invaluable quantitative tools for assessing the effectiveness of evidence-gathering practices. We applied these principles to develop a comprehensive method for measuring the evidentiary value of eyewitness identification procedures. We built on previous work that defined the effectiveness of identification procedures based on how much the procedure changed the probability of guilt in a Bayesian analysis (Wells & Lindsay, 1980). We showed how this approach can be extended to define the *Expected Information Gain (EIG)* of an identification procedure by measuring uncertainty with entropy values and averaging the change in uncertainty across a set of identification attempts. This method provides a complete measure of a procedure's overall evidentiary value.

In a series of demonstrations, we showed that EIG often supports different research conclusions than existing measures of the effectiveness of an identification procedure. In every case, the difference arose because the existing measure overlooked one or more aspects of the

reference population, whereas EIG incorporates all relevant aspects of the reference population. EIG also has broader applicability than the existing assessment tools, because it can be used to compare any procedures, even when the procedures give witnesses different response options and even when there are tradeoffs in evidentiary value (e.g., whereas one procedure better reduces uncertainty for identifications, another procedure better reduces uncertainty for rejections). EIG can also be used to evaluate the value of using continuous predictors of witness accuracy, such as response time, allowing researchers to identify important estimator variables using a metric that is targeted to the ultimate goal of identification procedures; namely, reducing uncertainty about guilt or innocence.

EIG in the Justice System

In the criminal justice system, people in a variety of positions are called on to make decisions that depend on the guilt or innocence of a suspect. For example, detectives need to decide if they should continue investigating a suspect or turn their attention to other suspects, judges need to decide if they should grant or deny a search or arrest warrant, and jurors need to decide whether or not a defendant should be convicted. To ensure the best possible outcomes, these decision makers need to consider the level of certainty supported by the evidence available to them and the potential consequences associated with different decision options (Robert, 2007). This task is complicated by the fact that the costs and benefits vary enormously across different decision scenarios. For example, the decision to grant or deny an arrest warrant for a shoplifting suspect likely has very different potential consequences than the decision to convict or release an accused murderer.

What is the role of identifications in this process? Identification attempts provide one source of evidence that informs judgments of guilt or innocence. Decision makers need to

integrate identification evidence with other sources of evidence to judge how certain they can be that each suspect is guilty or innocent. Obviously, eyewitness identification policy is just one component of performing this task effectively. Eyewitness researchers play an important role in helping decision makers understand the level of certainty supported by particular identification outcomes. For example, decision makers need to consider questions like “If a given suspect was identified with medium confidence from a 6-person simultaneous lineup, to what extent is this a strong indication of guilt?” The statistical properties of an appropriate reference population are an invaluable guide for these sorts of judgments, as they support the application of Bayesian inference to determine the posterior probability of guilt given a witness response (Wells & Lindsay, 1980). Eyewitness researchers potentially play a critical role in discovering and communicating information about the evidentiary value of particular witness responses in a variety of situations.

Against this backdrop, what do eyewitness researchers have to offer to policy makers, who need to consider outcomes over many investigations? We assume that good policy means putting decision makers in a position to make the best decisions possible across the full range of scenarios that they might encounter. One way to help decision makers is to give them evidence that supports stronger certainty of guilt or innocence. High certainty sets the stage for effective decision making even when potential costs loom large; for example, if jurors can be 99% certain that a defendant is guilty, then they might be willing to convict the defendant even if they ascribe a very high cost to falsely convicting innocent suspects.

To support high levels of certainty, decision makers need access to sources of evidence that provide a lot of information about guilt or innocence. Different identification responses produce different amounts of information, so policy makers need a way to assess the *overall*

level of certainty supported by a procedure across a large number of identification attempts. For a given reference population, EIG answers this question by factoring in all relevant aspects of performance to quantify how much a procedure increases certainty on average. If policy makers select procedures with the highest EIG, they will ensure that the various decision makers across all the various scenarios collectively have the most information about guilt and innocence possible. In other words, EIG quantifies the extent to which a procedure tends to ease the burden on decision makers by supporting high certainty of guilt or innocence.

Applying EIG to Real Investigations

Is it possible to assess EIG for real investigations? EIG is based entirely on the statistics of the reference population, so this question can be re-phrased as, “Is it possible to accurately estimate the statistics of the reference population for real investigations?” For example, a useful reference population could be a set of suspects accused of similar crimes in similar areas. Before addressing this question, we wish to remind readers that *every* measure of evidentiary value is based in some way on the statistics of the reference population. Moreover, these statistics determine the reasonable interpretation of identification outcomes by revealing the risk of making incorrect judgments based on particular witness responses. Thus, if it is not possible to make useful estimates of the statistics of the reference population for real investigations, then eyewitness researchers and investigators must acknowledge a foundational problem with the practice of using identifications as an evidence source.

Fortunately, researchers have several ways to learn about the statistics of a reference population. The task of defining a reference population can be segmented into two components: defining the likelihood of obtaining different witness responses for guilty and innocent suspects and defining the base rate of guilty suspects, i.e., the overall prior probability of guilt.

Likelihoods can be estimated with experiments designed to mimic the important properties of actual investigations. Base rates can be roughly estimated by evaluating conviction rates or expert ratings of guilt (Amendola & Wixted, 2014). Recent work also shows that decisions models can be used to accurately estimate witness accuracy and guilty-suspect base rates for data sets collapsed across guilty and innocent suspects, as in data from real investigations (Cohen et al., 2020; Wixted et al., 2016). A critical goal for future research is continuing to test and improve these techniques, creating tools that can be used to define reference populations for real investigations.

Our miscalibration analyses show that evidentiary value is fairly robust to misinterpretation of evidence, so even rough estimates of the reference population characteristics would dramatically improve the use of identifications as an evidence source. Moreover, EIG quantifies the cost of miscalibration, so it is poised to be a critical tool in the development of analysis techniques that can better characterize the reference population. Our results also show that researchers do not need detailed knowledge of base rates to effectively compare alternative identification procedures. Sometimes the relative evidentiary value of two procedures remains the same across a wide range of base rates (e.g., Figure 5). Moreover, even when the best procedure is sensitive to base rate, near-optimal procedures can be identified by assuming a base rate of 50% (see *Demonstration 1* and Figure 9).

Eyewitness researchers have focused more on exploring the effects of variables than characterizing reference populations from real investigations. Both are important goals, and EIG provides a new perspective on the effects of system and estimator variables. Real investigators might not be willing or able to implement every practice that improves EIG in eyewitness experiments, but EIG is a useful tool to help researchers and policy advocates “pick their battles”

by focusing on the changes that achieve the biggest improvements in overall information about guilt or innocence.

Information Versus Consequences

In the *Introduction*, we discussed the Expected Utility (EU) approach that assigns a value, or utility, to each possible identification outcome based on its assumed societal and personal consequences. This approach allows policy makers to value certain responses over others, for example, they might decide that evidence of guilt is more valuable than evidence of innocence. Understanding the implications of “over-weighting” certain outcomes is an important goal for policy-focused research. We mostly leave this complex issue for future studies, but here we will just note that EIG can quantify the decrease in evidentiary value that will result if policy makers select procedures that optimize an assumed value structure instead of optimizing information about guilt or innocence. Again, identification responses can be used to inform a wide range of decisions with a wide range of potential consequences, making it difficult (or perhaps impossible) to assign a single set of utilities that captures this full complexity. Being more certain about guilt or innocence is generally helpful across the full range of decision scenarios, so informational value is important to consider, independently of consequences.

EIG and Theory

We have mostly focused on EIG as a tool for policy-focused research, but EIG could play an important role in theoretical research as well. For example, we discovered a “filler curtain” mechanism by combining EIG with a signal-detection model of eyewitness decision making. Alternative decision models likely make different predictions for how much EIG should decrease as additional fillers are added to a lineup, as this result is affected by the probability that a filler

will produce a higher match to memory for the culprit than a guilty suspect. In contrast to a continuous signal-detection model, some models assume that guilty suspects can produce a memory state that can never be achieved by a filler, such as the detect state in the two-high-threshold model (Snodgrass & Corwin, 1988) or the recollection state in the dual-process signal-detection model (Yonelinas & Parks, 2007). Future research can explore whether testing the filler curtain effect on EIG is a useful way to distinguish these models.

We also discussed how EIG is a perfect tool for exploring the ability of decision models to recover the statistics of a reference population based on data that are collapsed across guilty and innocent suspects (Cohen et al., 2020; Wixted et al., 2016). Beyond eyewitness research, EIG could also establish a common performance metric across different decision models, facilitating model comparison and interpretation. For example, regardless of whether a model defines performance as the discriminability of continuous evidence distributions (Wixted, 2020) or discrete distributions of the probability of entering various evidence states (Snodgrass & Corwin, 1988), researchers can express performance in terms of the average information about stimulus class available from the model's evidence distributions.

Conclusion

Overall, the current work shows that EIG is an important measure for eyewitness memory research, and that it provides a complete assessment of a procedure's potential to inform judgments of guilt or innocence.

References

- Amendola, K. L., & Wixted, J. T. (2015). Comparing the diagnostic accuracy of suspect identifications made by actual eyewitnesses from simultaneous and sequential lineups in a randomized field trial. *Journal of Experimental Criminology*, 11, 263-284.
- Benish, W. A. (1999). *Relative entropy as a measure of diagnostic information. Medical Decision Making: an International Journal of the Society for Medical Decision Making*, 19, 202-206.
- Brewer, N., Weber, N., & Guerin, N. (2020). Police lineups of the future? *American Psychologist*, 75, 76-91.
- Campbell, W.M., Reynolds, D. A., Campbell, J. P., & Brady, K. J. (2005). Estimating and evaluating confidence for forensic speaker recognition. in *Proc. of ICASSP, 2005*, pp. 717–720.
- Ceci, S. J., & Friedman, R. D. (2000). The suggestibility of children: Scientific research and legal implications. *Cornell Law Review*, 86, 33-108.
- Chaitin, G. J. (1975). Randomness and mathematical proof. *Scientific American*, 232, 47–52.
- Chapman, S. D., & Cahill, B. S. (2012). Witnesses' memory for lineup fillers postdicts their identification accuracy. *Journal of Applied Research in Memory and Cognition*, 1, 11-17.
- Clark, S. E. (2003). A memory and decision model for eyewitness identification. *Applied Cognitive Psychology*, 17, 629-654. Doi: 10.1002/acp.891
- Clark, S. E. (2012). Costs and benefits of eyewitness identification reform: Psychological science and public policy. *Perspectives on Psychological Science*, 7, 238–259.
doi:10.1177/1745691612439584

- Clark, S. E., Benjamin, A. S., Wixted, J. T., Mickes, L., Gronlund, S. D. (2015). Eyewitness identification and the accuracy of the criminal justice system. *Policy Insights from the Behavioral and Brain Science*, 2, 175-186.
- Cohen, A. L., Starns, J. J., Rotello, C. M., & Cataldo, A. M. (2020). Estimating the proportion of guilty suspects and posterior probability of guilt in lineups using signal-detection models. *Cognitive Research: Principles and Implications*. 5:21 Doi: 10.1186/s41235-020-00219-4
- Colloff, M. F., Wade, K. A., & Strange, D. (2016). Unfair lineups make witnesses more likely to confuse innocent and guilty suspects. *Psychological Science*, 27, 1227-1239.
- Colloff, M. F., & Wixted, J. T. (2020). Why are lineups better than showups? A test of the filler siphoning and enhanced discriminability accounts. *Journal of Experimental Psychology: Applied*, 1, 124-143.
- Cover, T. M. & Thomas, J. A. (2006). *Elements of Information Theory*, 2nd ed., Wiley Interscience, 2006.
- Cutler, B. L., Penrod, S. D., & Martens, T. K. (1987). The reliability of eyewitness identification: The role of system and estimator variables. *Law and Human Behavior*, 11, 233–258.
- Dallenbach, K. M. (1913). The relation of memory error to time-interval. *Psychological Review*, 20, 323-337.
- Deffenbacher, K. (1980). Eyewitness accuracy and confidence: Can we infer anything about their relationship? *Law and Human Behavior*, 44, 243-260.
- Duncan, M. (2006). A signal detection model of compound decision tasks. (Tech Note DRDC TR 2006-256). Toronto: Defence Research and Development Canada.

- Fulero, S. M. (2009). System and estimator variables in eyewitness identification: A review. In D. A. Krauss & J. D. Lieberman (Eds.), *Psychological Expertise in Court: Psychology in the Courtroom* (pp. 57-78). Burlington VT: Ashgate Publishing.
- Gigerenzer, G., Gaissmaier, W., Kurz-Milcke, E., Schwartz, L. M., & Woloshin, S. (2008). Helping doctors and patients make sense of health statistics. *Psychological Science in the Public Interest*, 8, 53-96.
- Gneiting, T., & Raftery, A. E. (2007). Strictly proper scoring rules, prediction, and estimation. *Journal of the American Statistical Association*, 102, 359-378.
- Gong, D., Li, Z., Tao, D., Liu, J., & Li, X. (2015). A maximum entropy feature descriptor for age invariant face recognition. *Proceedings of the IEEE conference on computer vision and pattern recognition*, 5289-5297.
- Gronlund, S. D., Carlson, C. A., Dailey, S. B., & Goodsell, C. A. (2009). Robustness of the sequential lineup advantage. *Journal of Experimental Psychology: Applied*, 15, 140–152. <https://doi.org/10.1037/a0015082>.
- Gronlund, S. D., Carlson, C. A., Neuschatz, J. S., Goodsell, C. A., Wetmore, S. A., Wooten, A., & Graham, M. (2012). Showups versus lineups: An evaluation using ROC analysis. *Journal of Applied Research in Memory and Cognition*, 1, 221–228. doi:10.1016/j.jarmac.2012.09.003
- Hartley, R. V. L. (1928). Transmission of Information. *Bell System Technical Journal*, VII (3), 535–563.
- Kagehiro, D. K. & Stanton, W. C. (1985). Legal versus quantified definitions of standards of proof. *Law and Human Behavior*, 9, 159-178

- Kassin, S. M., Ellsworth, P. C., & Smith, P. L. (1989). The 'general acceptance' of psychological research on eyewitness testing: A survey of the experts. *American Psychologist*, 44, 1089-1098.
- Lampinen, J. M. (2016). ROC analyses in eyewitness identification research. *Journal of Applied Research in Memory and Cognition*, 5, 21-33. Doi: 10.1016/j.jarmac.2015.08.006
- Lampinen, J. M., Smith, A. M., & Wells, G. L. (2019). Four utilities in eyewitness identification practice: Dissociations between receiver operating characteristic (ROC) analysis and expected utility Analysis. *Law and Human Behavior*, 23, 36-44. Doi: 10.1037/lhb0000309
- Laughery, K. R., Alexander, J.F., & Lane, A. B. (1971). Recognition of human faces: Effects of target exposure time, target position, pose position, and type of photograph. *Journal of Applied Psychology*, 55, 477-483.
- Levi, A. M. (2006). An analysis of multiple choices in MSL lineups, and a comparison with simultaneous and sequential ones. *Psychology, Crime & Law*, 12, 273-285. Doi: 10.1080/10683160500238782
- Levi, A. M. (2012) Much better than the sequential lineup: A 120-person lineup. *Psychology, Crime & Law*, 18, 631-640. Doi: 10.1080/1068316X.2010.526120
- Lindsay, R. C. L., Semmler, C., Weber, N., Brewer, N., & Lindsay, M. R. (2008). How variations in distance affect eyewitness reports and identification accuracy. *Law and Human Behavior*, 32, 526-535. DOI: 10.1007/s10979-008-9128-x
- Loftus, G. R. (1985). Picture perception: Effects of luminance on available information and information-extraction rate. *Journal of Experimental Psychology: General*, 114, 342-356.

- Macmillan, N. A., & Creelman, C. D. (2005). *Detection Theory: A User's Guide* (2nd ed.). Mahwah, NJ: Lawrence Erlbaum Associates.
- Malpass, R. S. (2006). A policy evaluation of simultaneous and sequential lineups. *Psychology, Public Policy, and Law*, 12, 394 – 418. Doi: 10.1037/1076-8971.12.4.394
- Malpass, R. S., & Kravitz, J. (1969). Recognition for faces of own and other race. *Journal of Personality and Social Psychology*, 13, 330-334.
- McElreath, R. (2016). *Statistical Rethinking: A Bayesian Course with Examples in R and STAN*. Boca Raton: CRC Press/Taylor & Francis Group.
- Mickes, L., Flowe, H. D., & Wixted, J. T. (2012). Receiver operating characteristic analysis of eyewitness memory: Comparing the diagnostic accuracy of simultaneous and sequential lineups. *Journal of Experimental Psychology: Applied*, 18, 361–376.
- National Institute of Justice (1999). *Eyewitness evidence: A guide for law enforcement*. Washington DC: US Department of Justice
- National Research Council (2014). *Identifying the culprit: Assessing eyewitness identification*. Washington, DC: National Academies Press.
- Osterburg, James & Parthasarathy, T. & Raghavan, T. (1977). Development of a Mathematical Formula for the Calculation of Fingerprint Probabilities Based on Individual Characteristics. *Journal of The American Statistical Association*, 72. 772-778. 10.1080/01621459.1977.10479954.
- Police Executive Research Forum (2013). *A National Survey of Eyewitness Identification Procedures in Law Enforcement Agencies*. Washington, DC: US Department of Justice

- Price, H. L., & Fitzgerald, R. J. (2016). Face-off: A new identification procedure for child eyewitnesses. *Journal of Experimental Psychology: Applied*, 22, 366-380.
- Quinlan, J. R. (1986). Induction of decision trees. *Machine Learning*, 1, 81-106.
- Ramos, D. & J. Gonzalez-Rodriguez, J. (2008). Cross-entropy Analysis of the Information in Forensic Speaker Recognition. *Proceedings of IEEE Odyssey*, Odyssey 2008: The Speaker and Language Recognition Workshop.
- Ramos, D., Gonzalez-Rodriguez, J., Zadora, G., Zieba-Palus, J., & Aitken, C. (2007). Information-theoretical comparison of likelihood ratio methods of forensic evidence evaluation. *Third International Symposium on Information Assurance and Security*, 411-416.
- Ratcliff, R. (1978). A theory of memory retrieval. *Psychological Review*, 85, 59-108.
- Ratcliff, R., & McKoon, G. (2008). The diffusion decision model: Theory and data for two-choice decision tasks. *Neural Computation*, 20, 873-922.
- Ratcliff, R., Sheu, C.-f., & Gronlund, S. D. (1992). Testing global memory models using ROC curves. *Psychological Review*, 99, 518-535.
- Robert, C. P. (2007). *The Bayesian Choice: From Decision-Theoretic Foundations to Computational Implementation*. Springer.
- Rotello, C. M., & Chen, T. (2016). ROC curve analyses of eyewitness identification decisions: An analysis of the recent debate. *Cognitive Research: Principles and Implications*, 1:10. DOI 10.1186/s41235-016-0006-7
- Schooler, J. W., & Engstler-Schooler, T. Y. (1990). Verbal overshadowing of visual memories: Some things are better left unsaid. *Cognitive Psychology*, 22, 36-71. doi:10.1016/0010-0285(90)90003-m.

- Seale-Carlisle, T. M., Wetmore, S. A., Flowe, H. D., & Mickes, L. (2019). Designing police lineups to maximize memory performance. *Journal of Experimental Psychology: Applied*, 25, 410-430. DOI: 10.1037/xap0000222
- Shannon, C.E. (1948) A Mathematical Theory of Communication. *Bell System Technical Journal*, 27, 379-423, 623-656.
- Shapiro, P. N., & Penrod, S. (1986). Meta-analysis of facial identification studies. *Psychological Bulletin*, 100, 139-156.
- Smith, A. M., Lampinen, J. M., Wells, G. L., Smalarz, L., & Mackovichova, S. (2019). Deviation from perfect performance measures the diagnostic utility of eyewitness lineups but partial area under the ROC curve does not. *Journal of Applied Research in Memory and Cognition*, 8, 50-59. DOI: 10.1016/j.jarmac.2018.09.003
- Smith, S. M., Lindsay, R. C. L., & Pryke, S. (2000). Postdictors of eyewitness errors: Can false identifications be diagnosed? *Journal of Experimental Psychology: Applied*, 85, 542-550. Doi: 10.1037//0021-9010.85.4.542
- Smith, A. M. & Ayala, N. T. (submitted). Do traditional lineups undermine the capacity for eyewitness memory to rule out innocent suspects?
- Smith, A. M., Wells, G. L., Lindsay, R. C. L., & Penrod, S. D. (2017). Fair lineups are better than biased lineups and showups, but not because they increase underlying discriminability. *Law and Human Behavior*, 41, 127–145.
- Smith, A. M., Wells, G. L., Smalarz, L., & Lampinen, J. M. (2018). Increasing the similarity of lineup fillers to the suspect improves the applied value of lineups without improving memory performance: Commentary on Colloff, Wade, and Strange (2016). *Psychological Science*, 29, 1548-1551.

- Snodgrass, J. G., & Corwin, J. (1988). Pragmatics of measuring recognition memory: Applications to dementia and amnesia. *Journal of Experimental Psychology: General*, 117, 34-50.
- Sporer, S. L. (1993). Eyewitness identification accuracy, confidence, and decision times in simultaneous and sequential lineups. *Journal of Applied Psychology*, 78, 22-33.
- Steblay, N. M. (1997). Social influence in eyewitness recall: A meta-analytic review of lineup instruction effects. *Law and Human Behavior*, 21, 283-297.
- Wagenmakers, E.-J. (2009). Methodological and empirical developments for the Ratcliff diffusion model of response times and accuracy. *European Journal of Cognitive Psychology*, 21, 641-671.
- Wells, G. L. (1978). Applied eyewitness-testimony research: System variables and estimator variables. *Journal of Personality and Social Psychology*, 36, 1546.
- Wells, G. L. (2001). Police lineups: Data, theory, and policy. *Psychology, Public Policy, and Law*, 7, 791-801.
- Wells, G. L., & Lindsay, R. C. L. (1980). On estimating the diagnosticity of eyewitness nonidentifications. *Psychological Bulletin*, 88, 776-784. doi:10.1037/0033-2909.88.3.776
- Wells, G. L., & Olson, E. A. (2002). Eyewitness identification: Information gain from incriminating and exonerating behaviors. *Journal of Experimental Psychology: Applied*, 8, 155-167. Doi: 10.1037//1076-898X.8.3.155
- Wells, G. L., & Olson, E. A. (2003). Eyewitness testimony. *Annual Review of Psychology*, 54, 277-295.

- Wells, G. L., Smalarz, L., & Smith, A. M. (2015). ROC analysis of lineups does not measure underlying discriminability and has limited value. *Journal of Applied Research in Memory and Cognition*, 4, 313-317. Doi: 10.1016/j.jarmac.2015.08.008
- Wells, G. L., Smith, A. M., & Smalarz, L. (2015). ROC analysis of lineups obscures information that is critical for both theoretical understanding and applied purposes. *Journal of Applied Research in Memory and Cognition*, 4, 324-328. Doi: 10.1016/j.jarmac.2015.08.010
- Wells, G. L., & Turtle, J. W. (1986). Eyewitness identification: The importance of lineup models. *Psychological Bulletin*, 99, 320-329.
- Wells, G. L., Yang, Y., & Smalarz, L. (2015). Eyewitness identification: Bayesian information gain, base-rate effect-equivalency curves, and reasonable suspicion. *Law and Human Behavior*, 39, 99-122. doi:10.1037//lhb0000125
- Wells, W., Campbell, B., Li, Y., and Swindle, S. (2016). The characteristics and results of eyewitness identification procedures conducted during robbery investigations in Houston, TX. *Policing*, 39, 601-619.
- Wetmore, S. A., Neuschatz, J. S., Gronlund, S. D., Wooten, A., Goodsell, C. A., & Carlson, C. A. (2015). Effect of retention interval on showup and lineup performance. *Journal of Applied Research in Memory and Cognition*, 4, 8-14. doi:10.1016/j.jarmac.2014.07.003
- Wilson, B. M., Donnelly, K., Christenfeld, N., & Wixted, J. T. (2019). Making sense of sequential lineups: An experimental and theoretical analysis of position effects. *Journal of Memory and Language*, 104, 108-125.
- Wixted, J. T. (2020). The forgotten history of signal detection theory. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 46, 201-233.

- Wixted, J. T., & Mickes, L. (2012). The field of eyewitness memory should abandon “probative value” and embrace Receiver Operating Characteristic analysis. *Perspectives on Psychological Science*, 7, 275-278.
- Wixted, J. T., & Mickes, L. (2014). A signal-detection-based diagnostic-feature-detection model of eyewitness identification. *Psychological Review*, 121, 262.
- Wixted, J. T., Mickes, L., Clark, S. E., Gronlund, S. D., & Roediger, H. L. III. (2015). Initial eyewitness confidence reliably predicts eyewitness identification accuracy. *American Psychologist*, 70, 515–526.
- Wixted, J. T., Mickes, L., Dunn, J. C., Clark, S. E., & Wells, W. (2016). Estimating the reliability of eyewitness identifications from police lineups. *Proceedings of the National Academy of Science*, 113, 304–309. doi:10.1073/pnas.1516814112
- Wixted, J. T., Mickes, L., Wetmore, S. A., Gronlund, S. D., & Neuschatz, J. S. (2017). ROC Analysis in Theory and Practice. *Journal of Applied Research in Memory and Cognition*, 6, 343-351.
- Wixted, J. T., Vul, E., Mickes, L., & Wilson, B. M. (2018). Models of lineup memory. *Cognitive Psychology*, 105, 81–114. <https://doi.org/10.1016/j.cogpsych.2018.06.001>.
- Yang, Y., Smalarz, L., Moody, S. A., Cabell, S. A., & Copp, C. J. (2019). An expected cost model of eyewitness identification. *Law and Human Behavior*, 43, 205-219.
- Yonelinas, A. P., & Parks, C. M. (2007). Receiver operating characteristics (ROCs) in recognition memory: A review. *Psychological Bulletin*, 133, 800-832.
- Zhou, X., Cheng, M., He, X., & Matsumoto, T. (2014). Exact and approximated outage probability analyses for decode-and-forward relaying system allowing intra-link errors.

IEEE Transactions on Wireless Communications, 137062-7071.

10.1109/TWC.2014.2354337

Tables

Table 1

Table 1. Showup data from Colloff and Wixted (2020, Exp. 3) at high and low confidence levels.

	Suspect ID		Rejection		Total
	High	Low	High	Low	
Guilty	.31	.10	.03	.06	.50
Innocent	.11	.05	.23	.11	.50
Total	.42	.15	.26	.17	1.00
<i>p</i> (Guilty Resp.)	.74	.67	.12	.35	
<i>p</i> (Innocent Resp.)	.26	.33	.88	.65	

Notes: ID is identification; Resp. is response.

Table 2

Table 2. Aspects of the reference population included in probative measures.

Reference population aspect	Probative measure				
	EIG	Diag.	IR	SDT	PG _{Resp}
$p(x_I g)$	✓	✓	✓	✓	✓
$p(x_I i)$	✓	✓	✓	✓	✓
$p(x_{2..n} g)$	✓				
$p(x_{2..n} i)$	✓				
$p(g)$	✓				✓
$p(x_I)$	✓			✓	
$p(x_{2..n})$	✓				

Notes: EIG is expected information gain; Diag. is diagnosticity; IR is innocence risk; SDT is signal-detection theory; PG_{Resp} is probability gain for a single response; g is guilty; i is innocent; x_I is the response of interest for all measures but EIG (e.g., suspect ID); $x_{2..n}$ are responses other than the response of interest for all measures but EIG (e.g., filler ID, rejection, etc).

Table 3

Table 3. Steps in calculating EIG for the 15-person lineup in Demonstration 2

Response (x)	$p(x g)$	$p(x i)$	$p(x)$	$p(g x)$	$H(Y x)$	$IG(x)$	$p(x) \times IG(x)$
Suspect ID 100	0.066	0.001	0.0335	0.9851	0.1119	0.8881	0.02980
⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮
Suspect ID 0	0.010	0.006	0.0080	0.6250	0.9544	0.0456	0.00040
Filler ID 100	0.018	0.019	0.0185	0.4865	0.9995	0.0005	0.00001
⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮
Filler ID 0	0.032	0.079	0.0555	0.2883	0.8665	0.1335	0.0074
Rej. 100	0.000	0.000	0.0000	-----	-----	-----	-----
⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮
Rej. 0	0.019	0.047	0.0330	0.2879	0.8660	0.1340	0.0044
$\Sigma p(x) \times IG(x) = EIG = 0.16$							

Notes. Rows for intermediate confidence levels were removed for simplicity; see Appendix C for the full table with all confidence levels. Model predictions were rounded to the nearest 1000th, resulting in zero probability for some responses. Zero-probability responses are ignored in the *EIG* calculations. $p(x|g)$ is probability of Response x on identification attempts with guilty suspects; $p(x|i)$ is probability of Response x on identification attempts with innocent suspects; $p(x)$ is overall probability of Response x ; $p(g|x)$ is posterior probability of guilt for lineup Response x (posterior probability of innocence is $1 - p(g|x)$); $H(Y|x)$ is entropy for Response x ; $IG(x)$ information gain for Response x ; EIG is expected information gain.

Table 4

Table 4. Steps for calculating EIG for the showup in Demonstration 3 with the full confidence scale and with the scale collapsed to create a single ROC point.

Response (x)	$p(x g)$	$p(x i)$	$p(x)$	$p(g x)$	$H(Y x)$	$IG(x)$	$p(x) \times IG(x)$
Full Confidence Data							
Suspect ID 4	0.382	0.036	0.2090	0.9139	0.4234	0.5766	0.1205
Suspect ID 3	0.138	0.038	0.0880	0.7841	0.7526	0.2474	0.0218
Suspect ID 2	0.135	0.062	0.0985	0.6853	0.8985	0.1015	0.0100
Suspect ID 1	0.118	0.091	0.1045	0.5646	0.9879	0.0121	0.0013
Rej. 4	0.036	0.382	0.2090	0.0861	0.4234	0.5766	0.1205
Rej. 3	0.038	0.138	0.0880	0.2159	0.7526	0.2474	0.0218
Rej. 2	0.062	0.135	0.0985	0.3147	0.8985	0.1015	0.0100
Rej. 1	0.091	0.118	0.1045	0.4354	0.9879	0.0121	0.0013
$\Sigma p(x) \times IG(x) = EIG = 0.31$							
Center ROC Point (Confidence Collapsed for ID and Rej. Responses)							
Suspect ID 1-4	0.773	0.227	0.500	0.773	0.7727	0.2273	0.1136
Rej. 1-4	0.227	0.773	0.500	0.227	0.7727	0.2273	0.1136
$\Sigma p(x) \times IG(x) = EIG = 0.23$							

Notes. Model predictions were rounded to the nearest 1000th. $p(x|g)$ is probability of Response x on identification attempts with guilty suspects; $p(x|i)$ is probability of Response x on identification attempts with innocent suspects; $p(x)$ is overall probability of Response x ; $p(g|x)$ is posterior probability of guilt for lineup Response x (posterior probability of innocence is $1 - p(g|x)$); $H(Y|x)$ is entropy for Response x ; $IG(x)$ information gain for Response x ; EIG is expected information gain; ID is identification; Rej. is rejection.

Table 5

Table 5. Simultaneous and sequential lineup data from Clark (2012).

	Simultaneous				Sequential			
	Sus ID	Non ID	Total	ID rate	Sus ID	Non ID	Total	ID rate
Guilty	.270	.230	.500	.540	.215	.285	.500	.430
Innocent	.075	.425	.500	.150	.045	.455	.500	.090
Total	.345	.655	1.00		.260	.740	1.00	
<i>p(g Resp.)</i>	.780	.350			.830	.390		
<i>p(i Resp.)</i>	.220	.650			.170	.610		

Note: Sus ID is a suspect ID; Non ID is any response that is not a suspect ID; ID rate is the identification rate.

Figures

Figure 1

Figure 1. Showup and lineup data from Colloff and Wixted (2020). H is High confidence; L is Low confidence; Sus is Suspect; Fill is Filler; Rej is Rejection.

Figure 2

Figure 2. Relationship of Shannon surprisal and entropy to probability. In the left panel, $p(y)$ denotes the judged probability of an outcome, and I denotes the degree of judgment error as measured by surprisal, or equivalently, the amount of information gained by observing the outcome. In the right panel, $p(g)$ denotes the probability of guilt for a Bernoulli distribution defining uncertainty about guilt or innocence, and H denotes the entropy of the distribution (i.e., the average Shannon surprisal).

Figure 3 Figure 3. Relationship between EIG and PG_{EIG} .

Figure 4

Figure 4. A demonstration of the effect of averaging probability and information gain. The top and bottom panels provide results and analyses for hypothetical showup experiments that either do not (A, top) or do (B, bottom) include confidence ratings. ID is a suspect identification, Rej is a rejection, HC is high-confidence, LC is low-confidence.

Figure 5

Figure 5. Base rate curves for the standard showup and simultaneous lineup conditions from Colloff and Wixted (2020; Exp. 3).

Figure 6

Figure 6. The effect of miscalibration on EIG for the standard showup and simultaneous lineup conditions from Colloff and Wixted (2020; Exp 3). Miscalibration is the average absolute deviation between the actual and judged posterior probabilities across all response categories. Positive and negative values represent over-certainty and under-certainty of guilt or innocence. See text for details.

Figure 7

Figure 7. Illustration of the differences between EIG and diagnosticity ratios for the effect of changing the response criterion in a signal-detection model to simulate differences in witnesses' overall willingness to identify a suspect. The model used Gaussian distributions with standard deviations of 1 and means of -0.75 and 0.75 for identification attempts with innocent and guilty suspects, respectively. Thus, d' was 1.5. ID is suspect identification; Rej is rejection.

Figure 8

Figure 8. EIG across different criteria in a signal-detection model, simulating differences in witnesses' overall willingness to identify a suspect. $p(g)$ is the prior probability of guilt. The dashed lines mark the criterion position that optimizes the evidentiary value of the witness responses. The model used Gaussian distributions with standard deviations of 1 and means of -0.75 and 0.75 for identification attempts with innocent and guilty suspects, respectively. Thus, d' was 1.5.

Figure 9

Figure 9. EIG across different actual prior probabilities of guilt, $p(g)$, with criterion values optimized using different assumed prior probabilities of guilt. The model used Gaussian distributions with standard deviations of 1 and means of -1.5 and 1.5 for identification attempts with innocent and guilty suspects, respectively. Thus, d' was 3.

Figure 10

Figure 10. ROCs for lineup sizes (LS) of 15 and 1. Identification decisions were based on a signal-detection model in which the match-strength distribution had a mean of 1.5 for guilty suspects and 0 for innocent suspects (and fillers in the 15-person lineup). All distributions had a standard deviation of 1. The identification criterion (minimum match to culprit required for an identification) was set to 1, and 10 confidence criteria were arrayed above and below the identification criterion in equal steps of 0.20 to define confidence bins. For both identifications and rejections, the bin closest to the identification criterion was for confidence level 0, the next bin out was for confidence level 10, and so forth until the most extreme bin for confidence level 100. Responses for the 15-person lineups were determined by selecting the lineup member with the highest match strength and comparing this member's strength values to the identification and confidence criteria. Sus. is suspect; ID is identification; Inn. is innocent.

Figure 11

Figure 11. Match-strength distributions from a lineup signal-detection model across all suspects (top panel) or just suspects from lineups for which one of the fillers provided the strongest match to the witness' memory for the culprit (bottom panel). The filler strength distribution is not shown, but it is identical to the distribution for innocent suspects in the top panel. The model used Gaussian distributions with standard deviations of 1 and means of -1 and 1 for identification attempts with innocent and guilty suspects, respectively. LS is lineup size; Prob. Dens. is probability density.

Figure 12

Figure 12. The relationship between ROCs and EIG. The ROCs were determined from showup data simulated from a signal-detection model assuming different levels of witness discriminability. Each point and associated value correspond to the EIG value for a given level of discriminability and response criterion.

Figure 13

Figure 13. ROC functions, deviation from perfect performance (DPP), and expected information gain (EIG) values. Identification decisions were based on a signal-detection model in which the match-strength distribution had a mean of 1.5 for guilty suspects and 0 for innocent suspects. Both distributions had a standard deviation of 1. The identification criterion (minimum match to culprit required for an identification) was set to 0.75, and 3 confidence criteria were arrayed above and below the identification criterion in equal steps of .35 to define confidence bins. For both identifications and rejections, the bin closest to the identification criterion was for confidence level 1, the next bin out was for confidence level 2, and so forth until the most extreme bin for confidence level 4. Responses were determined by comparing the suspect's match strength to the identification and confidence criteria. Sus. is suspect; ID is identification; Inn. is innocent; min. is minimum; av. is average; Conf. is confidence.

Figure 14

Figure 14. Response time (RT) distributions for showup data generated from the diffusion model. See Ratcliff & McKoon (2008) for a detailed explanation of the model and the meaning of the parameters. The parameter values used to generate the data were s (diffusion coefficient) = 1; a (boundary width) = 5; z (relative starting point) = .5; sz (variability in starting point) = 0; v_G (average drift rate for guilty-suspect trials) = 3; v_I (average drift rate for innocent-suspect trials) = -3; sv (variability in drift rate across trials) = 3; ter (mean non-decision time) = .5; st (variability in non-decision time) = .05. IDs is suspect identifications; Rejs is rejections.

Appendix A: Equal-interval Relationship with Certainty

This appendix further explores the concept of measuring levels of certainty on an equal-interval scale. It includes an additional example of what it means to say that entropy, but not probability, provides an equal-interval scale of uncertainty and an analysis of whether or not any other measures of evidentiary value provide an equal-interval measure of uncertainty, as does information gain.

Equal-interval Example

The goal of this example is to explain what it means to have an equal-interval scale of certainty/uncertainty and to illustrate how using a scale that is not equal-interval can distort conclusions. The key property of an equal-interval scale is that the same change in certainty/uncertainty always produces the same amount of change on the scale, regardless of the starting level of certainty. To illustrate this property, the example will compare the relationship between two states of uncertainty in two scenarios, one with high uncertainty overall and another with low uncertainty overall.

Imagine that two people are playing the board game Clue. Each game begins with the random selection of cards to determine which one of six characters is the murderer, which one of six implements is the murder weapon, and which one of nine rooms is the location of the murder. To win the game, a player must be the first to specify all three details (murderer, murder weapon, and location). We will consider 2 players who have different levels of certainty regarding the correct solution. At this particular point in the game, both players have figured out which character committed the murder (Professor Plum, say). Player 1 has also figured out the location of the murder (the Library, say), but still does not know the exact murder weapon, which Player 1 has narrowed down to any 1 of 4 equally likely possibilities. Player 2 has the same information

as Player 1 in regards to the murder weapon, but Player 2 has only been able to eliminate 1 room, leaving 8 equally likely possibilities for the location. Thus, Player 1 is more certain of the correct solution than Player 2, as Player 2 must deal with an additional source of uncertainty (the unknown room).

The difference in certainty/uncertainty between the two players can be measured on either the probability scale or the surprisal scale. Player 1 is 25% certain of the correct solution; that is, Player 1's only barrier to solving the case is that there are 4 equally likely murder weapons. On the surprisal scale, Player 1 has 2 bits of uncertainty about the correct solution, $\log_2(1/.25) = 2$. Player 2 is 3.125% certain of the correct solution; that is, as far as Player 2 knows, the correct solution could involve any one of 4 equally likely weapons and any one of 8 equally likely rooms, for 32 total possibilities ($1/32 = .03125$). On the surprisal scale, Player 2 has 5 bits of uncertainty about the correct solution, $\log_2(1/.03125) = 5$. On the probability scale, the advantage in certainty for Player 1 over Player 2 is $.25 - .03125 = .21875$. On the surprisal scale, the difference in *uncertainty* between Player 2 and Player 1 is $5 - 2 = 3$.

Now consider another scenario in which the actual difference in certainty between the two players is held constant but the overall level of certainty is increased. Imagine that the murder weapon is revealed (the wrench, say). The only difference between Player 1 and Player 2 is still that Player 1 knows the location whereas Player 2 does not, but both Players are now more certain of the correct solution. How does this change affect the difference in certainty between the two players on the two different scales? Player 1 is now 100% certain of the correct solution, corresponding to 0 bits of uncertainty, $\log_2(1/1) = 0$. Player 2 is still uncertain as to which of 8 different rooms was the location of the murder. Thus, Player 2's certainty is 12.5% ($1/8$) on the probability scale, corresponding to 3 bits of uncertainty, $\log_2(1/.125) = 3$. Player 1 is still more

certain than Player 2, of course, and now the difference on the probability scale is $1 - .125 = 0.875$. Note that this is much larger than the probability difference from before the murder weapon was revealed, .21875. On the surprisal scale, Player 1 is still less uncertain than Player 2, and the difference is $3 - 0 = 3$. Note that this is exactly equal to the difference in surprisal before the murder weapon was revealed.

In summary, this example compares the behavior of the probability and surprisal scales for measuring uncertainty. The key outcome is the difference between Player 1 and Player 2 in two different scenarios, which represent two different ranges of the certainty/uncertainty scale. That is, one scenario has high overall uncertainty (murder weapon unknown) and the other has low overall uncertainty (murder weapon known). The actual uncertainty difference between Player 1 and Player 2 is the same in both scenarios; that is, Player 1 has the same piece of extra information that Player 2 lacks (the location of the murder). On the probability scale, the difference between Player 1 and Player 2 is larger in the low-overall-uncertainty scenario than the high-overall-uncertainty scenario, demonstrating that this scale does not provide an equal-interval measure of certainty.

Thus, probability is not a useful scale for answering questions like “What is the informational value of knowing the exact location versus only knowing that the location is 1 of 8 possibilities?”, as it will give different answers to this question in different scenarios. On the surprisal scale, the difference between Player 1 and Player 2 is constant across scenarios, demonstrating that this scale *does* provide an equal-interval measure of uncertainty. Thus, surprisal *is* useful for defining the informational value of knowing the location of the murder, as it provides the same answer to this question across different scenarios. For the same reason, surprisal is a useful measure for answering questions like “What is the informational value of

knowing that a witness identified the suspect with high confidence?” (The same is true for measures derived from surprisal, like entropy and cross-entropy.)

To expand on the last point, imagine that a researcher is interested in the relative information value of high-confidence identifications with two alternative procedures, e.g., sequential versus simultaneous lineups. Further, imagine that the researcher specifically wants to explore how the difference between procedures interacts with high or low levels of witness memory for the culprit. On the probability scale, this interaction would be distorted by measurement artifacts given that high witness memory corresponds to a low-overall-uncertainty state and low witness memory corresponds to a high-overall-uncertainty state. Measurements on the surprisal scale do not have this problem.

Finally, an equal-interval measure is required to take valid averages, that is, for the average scale value to correspond to the average amount of whatever characteristic is measured by the scale. Thus, researchers must use a measure like surprisal to validly answer the question “What is the average informational value across an entire set of identification attempts when informational value changes from one attempt to the next?”

Checking Other Measures for an Equal-Interval Relationship with Certainty

This section explores the relationship between various measures of evidentiary value and information gain. The goal is to determine if any alternative measures have an equal-interval relationship with uncertainty, as does information gain. Our strategy is to see if any of the measures has a linear (equal-interval) relationship with information gain, which would mean that the measure provides an equal-interval measure of uncertainty. The logic of this method is that information gain is an equal-interval measure of uncertainty (Shannon, 1948), so any other

measure that has an equal-interval relationship with information gain also has an equal-interval relationship with uncertainty.

For simplicity, we focused on the relationship between measures for a single response category, e.g., suspect IDs. To visualize the relationship between measures, we generated a sequence of true and false likelihoods of observing a particular witness response and calculated the measures using these values. False likelihoods ranged from .01 to .98 in steps of .01. For each false likelihood, we generated a sequence of true likelihoods from the given false likelihood plus .01 to .99 in steps of .01. This process ensured that we were exploring the relationship between measures across the full range of potential performance patterns for a response. We assumed a 50% prior probability of guilt for all measures influenced by this factor.

In short, perusing the plots in Figure A1 shows that none of the other measures of identification effectiveness considered here have a linear relationship with information gain, so none of them measure uncertainty on an equal-interval scale. This result means that averaging any of the other measures across responses, conditions, or anything else will distort conclusions about the average level of certainty.

For each measure, we include a plot for the “raw” measure and a plot for the measure on a log scale. The log scales deserve special consideration for the probability-based measures, i.e., probability gain and innocence risk. The plots in Figure A1 show that the log of these measures does not have an equal-interval relationship with information gain. This result may be confusing because entropy – the basis for information gain – is based on log probability values. The key to resolving this seeming contradiction is realizing that the different measures are based on probabilities at a different levels of analysis. For example, say that 80% of identified suspects are guilty. This situation would represent a probability gain of .3 [= .8-.5] on the raw scale and 0.47

$[\log(.8)-\log(.5)]$ on the log scale. Alternatively, this situation has an innocence risk of .2 on the raw scale and $-1.61 [\log(.2)]$ on the log scale.

Thus, these measures evaluate probability in terms of a distributional characteristic; namely, the probability parameter of a Bernoulli distribution representing uncertainty about guilt or innocence. In other words, for these measures, probability is a summary statistic for a distribution as a whole. In contrast, entropy evaluates probability at the level of individual judgments; namely, the judged probability assigned to the true state of affairs for each individual case represented in the distribution.

In this example, the optimal judgment is 80% certainty in guilt. For the 80% of suspects who are in fact guilty, the probability assigned to the true state of affairs is 80% (the person making the judgment is 80% certain that the suspect is guilty, and the suspect is indeed guilty). For the 20% of suspects who are in fact innocent, the probability assigned to the true state of affairs is 20% (the person making the judgment is 80% certain that the suspect is guilty, but the suspect is actually innocent). Entropy is based on the logs of *both* of these probabilities, and this property is why entropy does not have an equal-interval relationship with the log of the other probability-based measures. Specifically, entropy is based on the surprisal for both guilty and innocent suspects, found by taking the log of the inverse of the probability assigned to the true state of affairs for guilty and innocent suspects (or, equivalently, taking the negative of the log of each probability). These two surprisal values are weighted by the overall proportion of guilty and innocence suspects and combined; in other words, entropy is the average surprisal over all the individual judgments across the entire probability distribution.

In summary, even though (1) probability gain and innocence risk are probability-based measures and (2) information gain is based on log probabilities, taking the log of probability gain

or innocence risk does not achieve the nice measurement properties of information gain. This occurs because the first two measures are based on probability as a summary description of a distribution, and entropy is based on the probability assigned to the true state of affairs across all the individual judgments for all the samples in a probability distribution.

Figure A1

Figure A1. Relationship between information gain and other measures, including probability gain for a response, innocence risk, diagnosticity ratios, and their log-scale transformations for a single response. See text for details.

Appendix B: EIG and PG_{EIG} Computations for Eyewitness Researchers

For ease of reference, in this section, we compile all of the equations needed to calculate EIG and PG_{EIG} for an eyewitness memory data set. The equations reported here are designed for general application across a variety of different data sets regardless of the identification format (e.g., lineup or showup) and the potential witness responses (e.g., responses with or without confidence ratings).

To facilitate comparisons across data sets and conditions, the equations allow researchers to compute EIG for any prior probability of guilt regardless of the true base rate of guilty suspects in the data set. Using a prior probability other than 50% reduces the EIG value, with a larger reduction for prior probabilities that are closer to 0 or 100%. For example, if all of the suspects in a data set are guilty (e.g., Schooler & Engstler-Schooler, 1990), then the prior distribution has an entropy of zero and it is impossible to gain any information from the witness responses. We suggest that the standard practice should be calculating EIG for a 50% prior probability of guilt. Under this assumption, EIG represents the level of certainty supported by an identification procedure when it is the only source of information about guilt or innocence, and in this sense it is a “pure” measure of the information provided by the procedure. Moreover, the *Demonstrations Exploring EIG* section below establishes that identification procedures that optimize EIG under a 50% prior probability are near-optimal for *any* actual prior probability.

Computing EIG

Throughout, we use Y to refer to the distribution of uncertainty about guilt, g , and innocence, i . We use X to refer to the distribution of witness responses with x identifying a particular response, e.g., x could be an identification, id , or rejection, rej . For example, $H(Y|X=id)$ is the entropy of the distribution of guilt and innocence given that the response was an

identification. Where clear in context, we simplify as $H(Y | id)$. The possible responses in X can vary across situations, as will be shown below. EIG is a function of both Y and X , but, because these values are always provided by context, we do not include them in the notation.

The equations have terms for the prior probability of guilt, $p(g)$, where the prior probability of innocence is the complement of this value, $p(i) = 1 - p(g)$. For brevity, we use $p(i)$ instead of $1 - p(g)$ in equations, so researchers must ensure that $p(g) + p(i) = 1$ or the equations will give nonsensical results. Observed results enter the equations only in the form of the conditional probability of observing each potential witness response for guilty and innocent suspects (i.e., the response likelihoods), which is what allows researchers to find EIG for any desired prior probability of guilt. Of course, when known, a researcher could choose to set the prior probability of guilt to the actual base rate of guilty suspects, in which case all elements in the EIG equation will be based on the data set at hand. When comparing multiple conditions, researchers will almost always want to use the same prior probability of guilt for all conditions even if the base rate of guilty suspects differs across conditions.

The overall probability that a witness will make Response x , e.g., an identification or rejection (but not restricted to these responses), is given by

$$p(x) = p(g) \times p(x|g) + p(i) \times p(x|i), \quad (B1)$$

where $p(x|g)$ and $p(x|i)$ are the observed likelihoods of Response x for guilty and innocent suspects, respectively. The posterior probability of guilt, $p(g|x)$, for each possible witness response is given by Bayes rule as

$$p(g|x) = \frac{p(g) \times p(x|g)}{p(g) \times p(x|g) + p(i) \times p(x|i)}, \quad (B2)$$

and the posterior probability of innocence is given by

$$p(i|x) = 1 - p(g|x). \quad (B3)$$

The entropy for the prior distribution of uncertainty about guilt is given by

$$H(Y) = p(g) \log_2 \left(\frac{1}{p(g)} \right) + p(i) \log_2 \left(\frac{1}{p(i)} \right). \quad (\text{B4})$$

The entropy for the posterior distribution of uncertainty associated with witness Response x is given by

$$H(Y | x) = p(g | x) \log_2 \left(\frac{1}{p(g|x)} \right) + p(i | x) \log_2 \left(\frac{1}{p(i|x)} \right) \quad (\text{B5})$$

and the information gain, IG , is given by

$$IG(x) = H(Y) - H(Y | x). \quad (\text{B6})$$

Finally, the expected information gain is given by

$$EIG = \sum p(x) IG(x), \quad (\text{B7})$$

where the sum is over all possible responses.

Computing PG_{EIG}

As a useful supplement for interpreting EIG, we introduce PG_{EIG} (probability gain computed from EIG) by converting EIG from an information scale to a probability scale and switching to a simplified identification scenario for which the prior probability of guilt is 50% and all witness responses have equal evidentiary value. Specifically, PG_{EIG} is the probability gain across all responses in this simplified identification scenario that produces an amount of information equal to the EIG value (see the main text for more details). .

Unfortunately, there is no closed form equation for translating the entropy of a dichotomous variable (like guilt or innocence) to the probability parameter of a Bernoulli distribution, which is the operation needed to find PG_{EIG} . PG_{EIG} values can be found either with a

search algorithm or with the following approximation based on the formula reported by Zhou, Cheng, He, & Matsumodo (2014):

$$PG_{EIG} \approx .5 - \left(2^{0.6794 \times (1-EIG)^{0.7244}} - 2^{-0.1357 \times (1-EIG)^{21.8026}} \right)^{1.9920}. \quad (B8)$$

When calculated under the assumptions imposed here, a PG_{EIG} value of 50% means that the witness responses perfectly discriminated the guilty and innocent suspects, such that the average certainty about guilt increased from 50% to 100%. A PG_{EIG} value of zero means that the witness responses had no evidentiary value.

Appendix C: Full table of EIG calculations for the 15-person lineup

Response (x)	$p(x g)$	$p(x i)$	$p(x)$	$p(g x)$	$H(Y x)$	$IG(x)$	$p(x) \times IG(x)$
Suspect ID 100	0.066	0.001	0.0335	0.9851	0.1119	0.8881	0.0298
Suspect ID 90	0.029	0.001	0.0150	0.9667	0.2108	0.7892	0.0118
Suspect ID 80	0.037	0.002	0.0195	0.9487	0.2918	0.7082	0.0138
Suspect ID 70	0.044	0.003	0.0235	0.9362	0.3425	0.6575	0.0155
Suspect ID 60	0.050	0.005	0.0275	0.9091	0.4395	0.5605	0.0154
Suspect ID 50	0.052	0.007	0.0295	0.8814	0.5255	0.4745	0.0140
Suspect ID 40	0.049	0.009	0.0290	0.8448	0.6226	0.3774	0.0109
Suspect ID 30	0.041	0.010	0.0255	0.8039	0.7140	0.2860	0.0073
Suspect ID 20	0.030	0.010	0.0200	0.7500	0.8113	0.1887	0.0038
Suspect ID 10	0.019	0.008	0.0135	0.7037	0.8767	0.1233	0.0017
Suspect ID 0	0.010	0.006	0.0080	0.6250	0.9544	0.0456	0.0004
Filler ID 100	0.018	0.019	0.0185	0.4865	0.9995	0.0005	0.00001
Filler ID 90	0.015	0.016	0.0155	0.4839	0.9992	0.0008	0.00006
Filler ID 80	0.025	0.028	0.0265	0.4717	0.9977	0.0023	0.0001
Filler ID 70	0.038	0.045	0.0415	0.4578	0.9949	0.0051	0.0002
Filler ID 60	0.054	0.068	0.0610	0.4426	0.9905	0.0095	0.0006
Filler ID 50	0.071	0.096	0.0835	0.4251	0.9838	0.0162	0.0014
Filler ID 40	0.082	0.122	0.1020	0.4020	0.9721	0.0279	0.0028
Filler ID 30	0.084	0.138	0.1110	0.3784	0.9569	0.0431	0.0048
Filler ID 20	0.073	0.137	0.1050	0.3476	0.9319	0.0681	0.0071
Filler ID 10	0.054	0.115	0.0845	0.3195	0.9039	0.0961	0.0081
Filler ID 0	0.032	0.079	0.0555	0.2883	0.8665	0.1335	0.0074
Rej. 100	0	0	0	-----	-----	-----	-----
Rej. 90	0	0	0	-----	-----	-----	-----
Rej. 80	0	0	0	-----	-----	-----	-----
Rej. 70	0	0	0	-----	-----	-----	-----
Rej. 60	0	0	0	-----	-----	-----	-----
Rej. 50	0	0	0	-----	-----	-----	-----
Rej. 40	0	0	0	-----	-----	-----	-----
Rej. 30	0	0.001	0.0005	0	0	1	0.0005
Rej. 20	0.002	0.006	0.004	0.25	0.8113	0.1887	0.0008
Rej. 10	0.007	0.020	0.0135	0.2593	0.8256	0.1744	0.0024
Rej. 0	0.019	0.047	0.033	0.2879	0.8660	0.1340	0.0044
$\Sigma p(x) \times IG(x) = EIG = 0.16$							

Notes. Model predictions were rounded to the nearest 1000th, resulting in zero probability for some responses. Zero-probability responses are ignored in the *EIG* calculations. $p(x|g)$ is probability of Response x on identification attempts with guilty suspects; $p(x|i)$ is probability of Response x on identification attempts with innocent suspects; $p(x)$ is overall probability of Response x ; $p(g|x)$ is posterior probability of guilt for Response x (posterior probability of innocence is $1 - p(g|x)$); $H(Y|x)$ is entropy for Response x ; $IG(x)$ is information gain for Response x ; *EIG* is expected information gain; ID is identification; Rej. is rejection.