Forth coming in *Auditing: A Journal of Practice and Theory* (2012)

Causal inference in auditing: A framework

Rajendra P. Srivastava
Ernst & Young Distinguished Professor & Director
School of Business, The University of Kansas
1300 Sunnyside Drive, Lawrence, KS 66045, USA
Phone: (785) 864-7590, Fax: (785) 864-5328
Email: rsrivastava@ku.edu,

Theodore J. Mock
Distinguished Professor of Auditing & Assurance
University of California, Riverside,
900 University Avenue.
Riverside, Ca. 92521 and
Professor of Audit Research
University of Maastricht
Email: tmock@marshall.usc.edu

Karen V. Pincus
Doyle Z. and Maynette Derr Wiliams Chair of Professional Accoutning
University of Arkansas
Sam M. Walton College of Business
Fayetteville, AR 72701
Email: kpincus@walton.uark.edu

Arnold Wright
Joseph M. Golemme Research Professor
College of Business Administration
Northeastern University
360 Huntington Avenue, Boston, MA 02115, USA
Phone no: 617-373-7351; Fax: 617-373-8814
Email: a.wright@neu.edu

February 22, 2012

Acknowledgment: We wish to gratefully acknowledge the comments received from workshop participants at Boston College, the University of Arkansas and the University of Southern California and from Stan Biggs, Jean Bedard, Peter Gillett, Karim Jamal, Ganesh Krishnamoorthy, and Prakash Shenoy. In addition, we would like to express our sincere thanks to the anonymous reviewers and Robert Knechel for their thoughtful and constructive suggestions to improve the paper.

Causal inference in auditing: A framework

ABSTRACT

Causal inference—that is, determining the "root cause(s)" of an observed anomaly—is one of the most fundamental audit tasks. This study develops an analytical framework to formally model conditions present in many audit settings and provides illustrations related to performing substantive analytical procedures. We examine four conditions not fully considered in prior research: multiple hypotheses about what may cause an anomaly; multiple items of evidence with varying diagnosticity; observed effects that may not be certain; and hypotheses sets that may not be exhaustive. The results reveal when the following phenoma should occur: (1) discounting or inflating of posterior probabilities; (2) superadditive probabilities of various causes; and (3) unchanged probability of a potential cause given evidence in support of a different cause. The analytical findings have implications for the design and interpretation of experimental auditing research, for educating novice auditors, and for potentially improving audit practice.

Key Words: Causal Inference, Audit Judgment, Audit Analytical Procedures, Applied Probability, Multiple Hypotheses, Discounting, Uncertain Reasoning, Causal Schema

I. **INTRODUCTION**

I would rather discover one true cause than gain the kingdom of Persia.
−Democritus (460 B.C. – 370 B. C.)

Causal inference—that is, determining the root cause(s) of an observed event—is a pervasive and crucial component of many audit tasks, such as analytical procedures. Auditing standards stress the need to approach audit tasks with professional skepticism, encouraging the consideration of multiple hypotheses or causes, and the careful evaluation of evidence. For example, *International Standard on Auditing 200* (IAASB 2009, Paragraph 13(l)) defines "professional skepticism" as "an attitude that includes a questioning mind, being alert to conditions which may indicate possible misstatement due to error or fraud, and a critical assessment of audit evidence."

As Fornelli and Desmond (2011) note: "Skepticism involves the validation of information through probing questions, critical assessment of evidence, and attention to inconsistencies." Analytical procedures are particularly useful in focusing auditor attention on determining the causes of inconsistencies or anomalies. For U.S. public companies, analytical procedures are required as a risk assessment procedure (PCAOB 2010a) and at the overall review stage of the audit (PCAOB 2010b). They may also be used as substantive tests (AICPA 2009, AU 329). As large audit firms began using more powerful technology and business risk-based audit methodologies, the use of analytical procedures increased and grew more sophisticated (Curtis and Turley 2007; Knechel 2007; Knechel et al. 2010; Trompeter and Wright 2010).

Findings in operations research, management science, psychology, business, and auditing (e.g., MacDuffie 1997; Van Wallendael 1989; Bonner and Pennington1991) have shown that the consideration of potential causes ("hypothesis evaluation") is critical to successful performance of a diagnostic task. Further, research in psychology and auditing indicates that individuals have

difficulty performing hypothesis evaluation (e.g., Van Wallendael and Hastie 1990; Jamal et al. 1995; Asare and Wright 1997a, 1997b, 1995; Anderson and Koonce 1998; Johnson et al.2001; Mock et al. 2008). Such tasks are often quite complex, since the relationship between the hypothesized causes may vary dramatically across different settings.

Causal inference tasks in auditing share four important characteristics:

- First, there are typically multiple potential causes for an observed effect, that is, the decision-maker inherits or generates multiple hypotheses. For example, Wright et al. (2004) found that auditors considered operational changes, a change in economic conditions, unintentional errors, and deliberate fraud as potential causes of an unexpected fluctuation in an account balance. Several studies demonstrate that auditors readily generate and test multiple hypotheses (e.g., Libby 1985; Ismail and Trotman 1995; Kida and Luippold 2009).

- Second, hypotheses may be *non-exclusive,* that is, more than one cause may contribute to an observed effect. For example, Coglitore and Berryman (1988) presented archival evidence that financial statement fluctuations often have multiple causes, and Mock et al. (2008) provided verbal protocol evidence that auditors frame their hypotheses as non-mutually exclusive. Practitioner and educator materials provide many specific examples. For example, Blocher and Willingham (1993) offer over 30 potential causes for a change in gross margin and Knapp (2011) provides examples of recent fraud cases where losses were hidden by simultaneously manipulating multiple accounts.

- Third, the hypotheses set under consideration may be *non-exhaustive* – that is, the auditor acknowledges not all potential causes may have been identified. Bedard et al. (1998), Asare and Wright(2003). and Green and Trotman (2003) reported experimental evidence that auditors do not always include the correct candidate cause in their initial hypotheses set, yet may still identify the correct cause as audit evidence is reviewed. As one auditor in Mock et al. (2008, 133) said, "there is always the possibility of other causes."

- Finally, causal inference tasks usually entail relying on evidence that is not perfectly diagnostic, that is, evidence that may be convincing but does not establish the cause(s) with certainty.

Multiple research approaches can contribute to our understanding of causal inference. Archival studies provide evidence about the existence and association of certain potential causes and effect. However, causal inference is difficult to examine in archival studies since control

over extraneous and confounding factors can be quite difficult to achieve. Experimental studies have the advantage of being able to control such factors and thus focus our attention more clearly on causal paths, but are often limited in complexity (and therefore generalizability) by practical constraints on participant availability or time. For instance, many experimental studies of analytical procedures limit complexity by instructing auditors to assume there is only *one* cause of an observed financial statement fluctuation (e.g., Libby 1985; Bedard and Biggs 1991; Asare and Wright 1997a, 1997b). The purpose of this study is to extend prior research by providing an analytical framework for causal inference in which we address common complexities present in professional audit settings. Such a framework provides an alternative, but complementary, research method to both archival and experimental research.

The genesis of this study was a debate about how to interpret the results of a prior study (Asare and Wright 1997b) where subjects were instructed to assume there was a single cause for a material unexpected fluctuation in the gross margin of an audit client. The participants were presented with five hypotheses (candidate causes) which were described as coming from competent audit staff. The subjects then received 12 pieces of audit evidence (some confirming, some disconfirming, and some neutral) and were asked to update their assessments of the likelihoods (posterior probabilities) of each candidate cause being the reason for the fluctuation. Given the instructions, the authors expected that as auditors adjusted the likelihood for one cause up (down), they would balance this with an equal adjustment in the opposite direction to the other candidate causes. However, instead of summing to 1, the posterior probabilities summed to greater than 1 (an assessment process known as "superadditive").

Asare and Wright (1997b) interpreted these findings as evidence of suboptimal behavior resulting from cognitive strain, as suggested in prior psychology research. On the other hand, an

alternative interpretation is possible: in spite of the instruction to assume a single cause for the fluctuation in gross margin, the participants' audit experience may have led them to consider the possibility of multiple causes and/or to perceive interrelationships among the hypotheses; in this case, the superadditive results could be rational and appropriate because multiple hypotheses can co-occur as we will illustrate with examples later in this paper. The irresolvable debate about which interpretation was correct led to the our decision to develop an analytical framework that could address the complexities present in many audit judgment tasks. This paper describes the resulting framework, including appendices containing the key proofs and derivations related to the debate. The framework provides valuable insights for conducting experimental audit judgment research on causal inference and for developing the skills of novice auditors. Specifically, we extend and generalize past research by Morris and Larrick (1995) and Srivastava et al. (2002) by presenting a framework for causal inference in auditing. We examine four realistic conditions that are present in audit practice but have not been fully considered in prior research: multiple hypotheses about what may cause an anomaly; multiple uncertain items of evidence pertaining to a single hypothesis; observed effects that may not be certain; and hypotheses sets that may not be exhaustive.

The remainder of the paper is divided into five sections. The next section provides an overview of the causal inference task in an audit setting, discusses two important issues relevant to hypotheses evaluation, and introduces the theoretical framework chosen to deal with these issues. Section III develops a framework for analyzing multiple non-exclusive non-exhaustive hypotheses. Section IV analyzes discounting and inflation, and section V illustrates the model for several special cases of interest. The final section discusses implications for designing and

interpreting experimental research and offers suggestions for use of the framework in educating novice auditors and potentially improving audit practice.

II. ISSUES IN HYPOTHESES EVALUATION

Audit judgments are evidence-based. As Mautz and Sharaf noted in *The Philosophy of Auditing* (1961, p. 68), "The hypotheses which we develop. . . become strong enough to justify belief in them only if adequately supported by the evidence. Evidence gives us a rational basis for forming judgments." The auditing tasks we are interested in include cases where judgment must be based on evidence that is not perfectly diagnostic. That is, as Mautz and Sharaf (1961, p. 84) observed "compelling evidence is available to support only a limited number of financial statement propositions and . . . many other propositions, probably the great majority, are such that the mind of the auditor is not compelled but rather only persuaded."

Auditing standards acknowledge that audit evidence is often "persuasive rather than conclusive"(e.g., AICPA 2009, AU 326.13) and prior research (e.g., Spires 1991; Caster and Pincus 1996; Janvrin 2008) has examined various characteristics that determine the persuasiveness of evidence sets, such as source reliability and directness of evidence. Items of evidence may be confirming (supporting a hypothesis), disconfirming (contradicting a hypothesis), or neutral (non-diagnostic). Consequently, a set of evidence may range from strongly positive (all confirming) to strongly negative (all disconfirming) or, as often happens, may be mixed.

Auditing causal inference tasks involve four phases: (1) identification of an effect of interest, such as an unexpected financial statement fluctuation; (2) consideration of potential underlying causes, with each candidate cause being viewed as a hypothesis about what led to the effect; (3) development of a causal schema, that is a causal model of the task setting; and (4)

obtaining and evaluating evidence in order to update assessments of the probabilities of the potential causes. As evidence is accumulated, some causes are judged more likely and others are discounted (judged less likely) until a final judgment is rendered.

The focus of the current study is on hypothesis evaluation, the final two phases. We illustrate this framework in the context of performing substantive analytical procedures, an audit task that has been critiqued as deficient by the Public Company Accounting Oversight Board (PCAOB 2008). Additionally, prior research has shown that one of the most difficult phases of analytical procedures is hypothesis evaluation (Asare and Wright 1995, 1997a, 1997b; Mock and Turner 2005). To develop a framework for hypotheses evaluation, we must first consider the representation of cause-effect relationships (causal schemas) and representation of the diagnostic strength of the evidence pertaining to each hypothesis. For convenience, Table 1 summarizes terms used throughout the discussion of the framework.

----- Table 1 here -----

Morris and Larrick(1995, 347) describe two basic approaches to causal inference, "induction" and "ascription:"

> *Induction is a bottom-up process that draws on co-variation across multiple cases to infer a cause-and-effect connection or causal rule. Ascription is a top-down process that draws on a causal schema (a set of assumptions about causal relationships) to infer the causal relationships present.*

Our model takes the ascription approach. This approach is appropriate when there is a lack of rich data to determine the co-variation across multiple cases to infer a causal schema. Many audit settings lack such data and thus better fit the ascription approach.

Causal schemas may be modeled as the joint probability distribution of all the causes under consideration. To illustrate, several examples of causal schemas are presented in Figure 1 in terms of Venn diagrams and joint probabilities of the causes before and after observing an

effect. Assume that there are only two independent and "sufficient" potential causes, A and B, for effect F. A "sufficient" cause is one that *always* results in the effect (i.e., P(F|A) = 1 and P(F|B) = 1). For example, overstatement of sales revenue and overstatement of ending inventory are independent causes that would each *always* result in overstatement of gross margin. This is the most simple case and is depicted in Panels A and B where P(A)=.20, P(B)=.10 and P(AB)=.02.

While some sufficient causes exist in audit settings, non-sufficient causes are also relevant. Non-sufficient causes *sometimes*, but not always, result in the effect. For example, assume an auditor finds an unexpected increase in a client's gross margin. A change in the sales mix will *sometimes*, but not always, cause an unexpected increase in gross margin. Panels C and D in Figure 1 depict the case of "non-sufficient" causes, A and B, of the effect F, where P(F|A) < 1, and P(F|B) < 1. Specifically, P(A) is still equal to .20 but now there are conditions where A can occur (or A and B together) without F occurring. Similarly, P(B) is still equal to .10 but there are conditions where B can occur (or B and A together) without F occurring. The resulting joint probabilities for A and B for this more complex case are shown in Panels C and D.

----- Figure 1 here -----

For our more general analysis, we use likelihood ratios (λ) to express the nature and strength of evidence (e.g., see Edwards 1984; Pearl 1990; Dutta and Srivastava 1993, 1996; Srivastava et al. 2002) about the potential causes of an effect F[1]. For an individual item of positive evidence supporting that cause A is present, the likelihood ratio takes values between 1 and infinity ($1<\lambda<\infty$). A value of $\lambda\rightarrow\infty$ is equivalent to saying the evidence conclusively establishes A as the cause, the situation discussed by Morris and Larrick (1995) but of less

[1] Our model also considers that the observed effect may itself be uncertain. For example, Schultz et al. (2010, p. 240) note that fluctuations which at first appear to be inconsistencies are sometimes in fact consistent with subtle changes in business conditions.

interest in an audit context because evidence is not often conclusive. A value of 1 implies the evidence has no diagnostic value. Values between 1 and ∞ indicate increasingly strong evidence in support of the cause examined (A). For negative evidence, the likelihood ratio takes values between zero and one ($0<\lambda<1$) where a value of zero implies that cause A is certainly not present. In general, if we have q independent items of evidence pertaining to a single hypothesis, say B, then the combined strength of the evidence set is determined by multiplying the individual likelihood ratios as: $\lambda_B = \prod_{i=1}^{q} \lambda_{Bi}$ This result may be derived using Bayes rule (Dutta and Srivastava 1993).

III. ANALYTICAL FORMULA FOR MULTIPLE HYPOTHESES EVALUATION

In this section we develop our analytical framework by deriving a formula for the posterior probability of a potential cause, H_i, after observing an anomalous effect F, and gathering evidence pertaining to multiple potential causes. We begin by considering the general case where there are n potential causes (hypotheses) of F as depicted in Figure 2. For this model, we assume there is one piece of evidence pertaining to each hypothesis and the presence of the effect F is not known with certainty. In other words, there is evidence E_F about the presence of the effect F, but not absolute assurance. The likelihood ratio corresponding to this evidence, $\lambda_F = P(E_F|F)/P(E_F|\sim F)$, is assumed to be between one and infinity.

----- Figure 2 here -----

This model may be extended to the situation of multiple items of evidence for each hypothesis by first combining all the items of evidence pertaining to a hypothesis using the multiplicative model described in the previous section and then treating this combined evidence set similarly to a single piece of evidence. Let E_i represent the evidence pertaining to hypothesis

H_i and the corresponding strength of evidence E_i pertaining to H_i will be expressed in terms of the likelihood ratio $\lambda_i = P(E_i|H_i)/P(E_i|\sim H_i)$. The posterior probability that a particular potential cause is present given that we have evidence about the other potential causes and the effect F can now be stated in terms of Theorem 1 (see Appendix A for the proof):

Theorem 1: The posterior probability that hypothesis H_k is present given that we have knowledge about the presence or absence of various hypotheses and the effect F through the evidence pertaining to each hypothesis and the effect F is:

$$\mathbf{P(H_k|E_1E_2\ldots E_nE_F) = N_k/(N_k + D_k)} \quad (1)$$

where N_k and D_k are defined in Appendix A (see equations A9b and A9c).

Theorem 1 is the direct result of Bayes' rule of updating probabilities, which at first yields a rather daunting expression that simplifies nicely to Equation (1). Intuitively, the N_k in Equation (1) represents partly the impact of all the items of evidence on the assessed likelihood of the k^{th} candidate cause *being responsible for the observed effect* (H_k is true) and partly the impact of the *relationships among the hypotheses* (positively correlated, negatively correlated, or independent) on H_k. In contrast, D_k represents partly the impact of all the items of evidence on the assessed likelihood of the k^{th} candidate cause *not* being responsible for the observed effect ($\sim H_k$ is true) and partly the impact of the relationships among the hypotheses on $\sim H_k$.

This framework represents the most general relationship among the potential causes. One can analyze special cases using Equation (1) by considering the appropriate number of hypotheses, relevant values of the strength of evidence, and by choosing appropriate values for the joint probabilities to represent the relationships among the hypotheses. The next two sections elaborate on this theorem for a number of cases relevant to auditing.

IV. DISCOUNTING OR INFLATING POSTERIOR PROBABILITIES

To illustrate the issue of discounting or inflating posterior probabilities, consider Asare and Wright's (1997b) experiment in which participants were provided with 12 pieces of evidence and asked to determine posterior probabilities of five hypotheses (potential causes) for an unexpected increase in the gross margin. The participants were instructed to assume there was a *single* cause for the fluctuation. Asare and Wright thus expected to observe discounting because when there is a single cause, evidence to *support* one candidate cause is indirect evidence *refuting* the other potential causes. However, discounting did not occur. Ignoring indirect evidence by failing to discount the likelihoods of other potential causes could be inefficient – that is, it could prompt unnecessary additional audit effort to investigate the other potential causes.

On the other hand, assuming a single cause when multiple causes are present can lead to an ineffective audit. Consider, for instance Mautz and Sharaf's (1961, p. 125) list of ways to manipulate accounts to perpetrate fraud: omission of entry, false entry, false footing or other calculations, false posting or other bookkeeping procedures, destruction of documents, preparation of false documents, and alteration of legitimate documents. Suppose for a moment a fraudster had used all of these techniques (candidate causes) to misstate inventory. The audit would likely be ineffective if an auditor assumed a single cause and therefore interpreted evidence supporting one cause as implying the other potential causes were less likely.

Using Theorem 1, we can derive a general formula for the amount of discounting (or its opposite, *inflation*) of posterior probabilities that will occur under different conditions that may be present in audit practice. In particular, hypotheses may be independent (i.e., the presence of one cause does not imply another is also present), positively correlated (i.e., if one cause is

present, another is also likely to be present), or negatively correlated (i.e., if one cause is present, another is less likely to be present). The following sections contain illustrations in audit practice of each of these relationships between potential causes. The general formula (Equation 2 below) is developed in Appendix B:

$$\textbf{Discounting of A} = P(A|F) - P(A|FE_B) = (\lambda_B - 1)[P(A|F)P(B|F) - P(AB|F)]/[1+(\lambda_B - 1)P(B|F)] \quad (2)$$

This formula generalizes the expression for the magnitude of discounting for: 1) any relationship between two candidate causes A and B, whether they are independent, positively correlated or negatively correlated, 2) whether B is known with certainty or not, 3) whether A and B are sufficient or non-sufficient causes of effect F, and 4) whether A and B are the only possible causes of F or not. Consider the following special cases.

A and B are the Only Possible Causes of the Effect F

While auditors are not often in a situation to consider an exhaustive set of hypotheses due to the large number of potential causes for many anomalies such as an increase in gross margin, it is helpful to understand that discounting will *always* occur in this situation, as analytically demonstrated in Appendix B. Later, we will contrast this result with the more common situation in audit practice, where the set of causes under consideration is not exhaustive.

An example of such a situation in an audit context is a retailer with an unexpectedly high gross margin where we assume there are only two potential reasons, overstated sales revenue (Cause A) and understated cost of sales (Cause B). Figure 3 shows how discounting of A will increase as the strength of evidence in support of B increases, using P(A) = 0.2 and P(B) = 0.1 for illustration as we did in the examples in Figure 1. Case 1 describes the situation where A and B are negatively correlated and sufficient causes of F. Case 2 describes the situation where A and

B are independent and sufficient. Case 3 describes the situation where A and B are positively correlated and sufficient.

Although it is somewhat counterintuitive, discounting occurs even when A and B are positively correlated under the situation when A and B are the only possible causes. However, the magnitude of discounting is lower than when they are either negatively correlated or independent.[2] Recall the example of the retailer with unexpectedly high gross margin and two potential causes (sales revenue overstated and cost of sales understated). These causes could be positively correlated if the cost of sales includes material purchase discounts because of a higher than usual volume of purchases to fill the higher sales demand. If the auditor finds evidence supporting the overstated revenue hypothesis, then discounting means the chance of cost of sales being understated is reduced.

----- Figure 3 here -----

A and B are Not the Only Possible Causes of F

Next, consider a common situation in audit practice: there are more than two potential causes, A and B, of effect F which may be related in different ways. Discounting will occur in some situations, for example when A and B are positively correlated before observing F, but conditionally negatively correlated given F.

Discounting's opposite, inflation, will occur in other situations, for example when A and B are positively correlated both before observing F and given F. Yet in other situations, there will be neither discounting nor inflation – for example, when A and B are positively correlated

[2]Morris and Larrick(1995) obtained similar results, however only for the special case where B was known with certainty. One can obtain their model for discounting by considering E_B in (Equation B3) to be infinitely strong and positive (i.e., $\lambda_B \rightarrow \infty$) implying that B is known with certainty (i.e. $E_B \equiv B$).

before observing F, but independent given F. Figure 4 provides numerical illustrations of these three possibilities and we illustrate the third possibility with an auditing example.

----- Figure 4 here -----

In Case 1 of Figure 4, the discounting of A increases as the strength of evidence in support of B increases, since in this case, A and B are negatively correlated given that the auditor has observed the effect F ($\rho_{AB|F} = -0.258$). However, in Case 2, discounting of A is zero for all values of the strength of evidence in support of B. In this case, A and B are independent given that the auditor has observed the effect F ($\rho_{AB|F} = 0$).

In Case 3 of Figure 4, the discounting of A decreases as the strength of evidence in support of B increases. This result implies that when there is evidence in support of an alternative cause B which is conditionally positively correlated to A, then the probability that A is present increases. This *inflating* phenomenon is the inverse of *discounting*. In this case, the correlation between A and B given F is positive and equal to 0.258.

To illustrate Case 3 with an example, suppose the auditor is considering two potential causes for an unexpected increase in gross margin: higher sales and lower cost of goods sold. The introduction of a marketing program may increase sales, leading to a higher gross margin and may also result in greater raw material discounts due to larger purchases, which simultaneously results in an increase in the gross margin.

V. ALTERNATIVE SCENARIOS IN MULTIPLE HYPOTHESES EVALUATION

In this section, the goal is to finish with an illustration of a complex realistic scenario, a non-exhaustive and non-exclusive case with mixed correlations and increasing amounts of evidence. As before, we first consider the simpler case of an exhaustive set of two potential causes, A and B, of an observed effect F and then add more realistic complexity. In general, we

assume the causes are non-sufficient, that is A and B would not always cause the effect. We will now investigate the changes in the probability of the presence of each cause as new items of evidence are gathered.

Starting Case: Two Possible Causes with One Item of Evidence per Cause

Mutually Exclusive and Exhaustive Schema: While mutually exclusive and exhaustive schema may not be typical in practice, in experimental research participants sometimes are given a set of potential causes for an unexpected analytical procedure result and are told the single cause is among this set. The mutually exclusive and exhaustive set of causes, A and B, of the effect F, implies that Cause A cannot co-exist with Cause B, i.e., $A \cap B = \emptyset$, or $P(A \cap B) = 0$. Under this condition, the revision of probabilities of the two causes would always be complementary and their sum will be always one. In other words, $P(A|FE_AE_B) + P(B|FE_AE_B) = 1$ when $P(A \cap B) = 0$, and A and B are the only two candidate causes. This is the result that is assumed in Asare and Wright (1997b) and similar experimental studies.

Non-Exclusive and Exhaustive Schema: What kind of results would have been expected in Asare and Wright (1997b) if the auditor subjects treated the hypotheses as non-exclusive, i.e. the two causes can co-occur? Under this scenario, we still have only two causes (exhaustive), A and B, but they are not mutually exclusive, which means $A \cap B \neq \emptyset$, or $1 > P(AB) > 0$. Figure 1 represents such a schema where Panel A assumes A and B to be "sufficient" causes and Panel C assumes A and B to be non-sufficient causes.

An auditing example of non-exclusive and exhaustive schema would be a situation where the auditor thinks the sales revenue for the current period is unusually high compared to prior periods (representing the effect F) and has reason to suspect deliberate overstatement. The auditors considering two hypotheses: prematurely recording the sales (Cause A), and creating

fictitious sales (Cause B). The causes are clearly not mutually exclusive in that management can do both. Figure 5 represents such a case, where we observe that the revised probabilities of A and B *increase* as the strength of evidence, pertaining to cause A (i.e., λ_A) and the strength of evidence pertaining to cause B (i.e., λ_B) increase. The sum of the two revised probabilities is more than one, referred to as "superadditive," and increases as the values of λ_A and λ_B increase. In the extreme case when the items of evidence are extremely strong, i.e., $\lambda_A \rightarrow \infty$ and $\lambda_B \rightarrow \infty$, the sum of the two revised probabilities would be two as seen in Figure 5. This is an important result as it demonstrates superadditivity is consistent with a Bayesian approach under these circumstances, whereas experimental research, including Asare and Wright (1997b) that started our debate, has sometimes viewed superadditivity as inconsistent with a Bayesian model.

----- Figure 5 here -----

Final Case: Non-Exclusive and Non-Exhaustive Hypotheses with Mixed Correlations

Finally, we show how the model presented above can be extended to deal with even more complex situations that occur in auditing. In particular, we consider cases where there are several interrelated potential causes of an effect. Some of these hypotheses may be positively correlated, some negatively correlated, and some independent. Further, the auditor may recognize that not all potential causes are known.

Consider again an auditor investigating F, an unexpected increase in gross margin, and assume the auditor identified the following initial hypotheses:

- H_1 = greater sales discounts yielded increased sales volume;
- H_2 = price increases yielded a better profit margin;
- H_3 = greater raw material discounts due to higher volume purchases;
- H_4 = improper sales cutoff; and
- H_O = other non-identified reasons.

Note that H_1 and H_3 are expected to be positively correlated; that is a marketing strategy offering greater sales discounts is likely to result in greater raw material price discounts due to volume purchases. In contrast, H_1 and H_2 would be expected to be negatively correlated; greater sales discounts would be expected to increase sales quantity while a price increase would be expected to yield a higher profit margin on the existing volume and perhaps even a decreased volume. Moreover, since H_1 and H_3 are positively correlated, H_2 and H_3 would be negatively correlated. H_4 is expected to be independent of the set of hypotheses $\{H_1, H_2, H_3\}$ given F (the observed increase in gross margin). This independence means finding evidence that there are greater sales discounts (H_1), price increases (H_2), or greater raw materials price discounts (H_3), does not provide any information about whether there has been an improper sales cutoff (H_4). Also, for simplicity of modeling, we assume H_O will not co-occur with $\{H_1, H_2, H_3, H_4\}$. A schema of the above relationships with a numerical illustration is represented through a Venn diagram in Figure 6.

----- Figure 6 here -----

Figure 7 shows how the revised probabilities of these various hypotheses change as additional evidence is collected, assuming for illustration that each piece of evidence is moderately strong. The lower five lines show the changes in the probability of each hypothesis. As the graph indicates, when evidence E_1 is observed, the probability of H_1 and H_3 goes *up*, the probability of H_2 and H_O goes *down*, and the probability of H_4 remains *unchanged*. This result is what one intuitively expects based on the correlation among the hypotheses.

Similarly, from Figure 7, note that when E_2 is observed, the probability of H_1 and H_3 *decrease*, the probability of H_2 *increases*, and the probability of H_4 remains *unchanged*. Importantly, not only does the model provide results consistent with directional expectations, it

also determines the *magnitude* of probability revision, which entails a very complex cognitive task that individuals are likely to have great difficulty performing.

When evidence E_3 is obtained, we have similar results to that of evidence E_1, the probabilities of H_1 and H_3 go up, the probability of H_2 goes down, but the probability of H_4 remains unchanged. However, when evidence E_4 is observed, the probability of H_4 goes up, but the probabilities of H_1, H_2, and H_3 remain unchanged. This result is because the hypotheses in the set $\{H_1, H_2, H_3\}$ are independent of H_4. That is, evidence that there is an improper sales cutoff does not tell us anything about whether H_1, H_2, or H_3 are likely causes.

The top solid line in Figure 7 represents the sum of all the five probabilities. It is important to note that under the assumed scenario, the sum of all the five probabilities is greater than one at each stage, once again illustrating superadditivity is possible under the Bayesian framework. Superadditivity will occur if some of the hypotheses are positively correlated and some hypotheses are independent of others.

----- Figure 7 here -----

VI. CONCLUSIONS

In this paper we develop a general framework for auditing causal inference and explore the implications for various cases an auditor may encounter. Our key analytical findings include:

- The assessment of the potential causes of an unexpected fluctuation (anomaly) and the related *discounting* or *inflating* of causes is contingent on the number of potential causes and their interrelationships.
- The type (positive, negative or mixed), quantity (single or multiple items), and strength of audit evidence have a dramatic impact on the magnitude of probability revisions.

- In a common audit setting with non-exhaustive and non-exclusive hypotheses, *inflating* should occur in the posterior probability of a cause, say A, if there is positive evidence that another cause, say B, is present as long as A and B are positively correlated. In some situations, this will lead to *superadditive probabilities*.
- In the audit setting with non-exhaustive and non-exclusive hypotheses, the posterior probabilities of all the hypotheses may logically add to more than one (*superadditive*) when there are strong positive items of evidence in support of several of the positively correlated hypotheses or when some of the hypotheses are independent. Knapp (2011) includes examples of financial statement frauds that fit this situation such as Crazy Eddie.

These findings have implications for both research and practice. Most importantly, given that both practitioners and experimental researchers often consider a fairly small initial set of hypotheses (e.g. Libby 1985, Bedard and Biggs 1991), it is critically important to know whether the set is considered by the auditor or by the experimental participants to be exhaustive or non-exhaustive and to know the expected interrelationships (positive, negative or independent) among the hypotheses. Differences in these characteristics imply different interpretations of audit evidence or research results.

The framework and our related analysis show that either side in our original debate over Asare and Wright (1997b) could be correct. That is, interpretation of the appropriateness of participants' hypothesis evaluation behavior is contingent upon the manner in which auditors encoded the nature of the conditions of the hypothesis set. The same problem also applies to other prior studies that asked participants to assume a mutually exclusive and exhaustive set of hypotheses (e.g., Asare and Wright 1997a; Bedard and Biggs 1991; Libby 1985). Thus, our

findings highlight the importance of explicitly measuring or controlling interdependencies among hypotheses evaluated in future experiments.

While our focus has been on causal inference in auditing research, similar ideas apply to causal inference in other areas of accounting. For example, Koonce et al. (2011) suggest a program of future research directed toward studying causal reasoning in financial reporting and voluntary disclosure. Researchers pursuing their questions involving diagnostic reasoning would also benefit from consideration of our framework.

Researchers should consider the framework when designing experimental task materials and use pilot tests and debriefing questions to establish how participants perceive the interrelationships among the hypotheses provided in the experiment. Such debriefing questions would be much like manipulation checks commonly employed in experimental research. For example, if a researcher has tried to create an exhaustive set of hypotheses, a pilot test question could ask "Can you think of any other potential causes for (the effect)?" If new hypotheses are identified, the researcher can modify the experimental materials either by adding new hypotheses or instead recalculating the expected results assuming the set is non-exhaustive. Another option would be to add one additional hypothesis in the experimental materials that covers "other potential causes" and treat the set as exhaustive.

Similar questions could be asked in debriefings to be sure experimental subjects viewed the interrelationships among the hypotheses as intended. If hypotheses are self-generated, participants could be asked to describe any expected relationships. Knowing these interdependencies will allow clearer prediction and interpretation of the results.

In the classroom we have also found the framework useful in instructing novice auditors. For example, one coauthor uses a five hypothesis case with students in both master's and

undergraduate auditing courses. The students review audit evidence about the potential causes of a fluctuation and adjust their likelihoods after each piece of evidence. The framework helps students understand how relationships among potential causes should affect probability assessments and see for themselves what audit judgment challenges exist in realistic cases. The framework also is useful when discussing audit judgment research design with doctoral students.

From a practice perspective, consider the challenges faced by auditors in interpreting audit evidence and balancing the efficiency and effectiveness of procedures. For many fluctuations in account balances, the number of potential causes is quite large (e.g. Blocher and Willingham 1993). Yet, auditors typically consider a much smaller initial set of hypotheses (e.g., Ismail and Trotman 1995: Cianci and Bierstaker 2009), raising the question of whether auditors fully consider the possibility of other unidentified potential causes. Further, do auditors incorporate knowledge about the interrelationships between potential causes? Could assessments be improved by focusing auditors' attention to these matters? One approach to address these questions could be combining the implications of this framework with the use of systems-thinking tools (e.g., Brewster 2011; O'Donnell and Perkins 2011).

Based on classroom experience using the framework to teach novice auditors, it is possible a decision support tool based on our framework could be beneficial in auditor training or in the field when investigating potential causes for a material fluctuation, including consideration of fraud hypotheses. As the Center for Audit Quality (2010, 31) recently noted in its call for improvements in deterring and detecting fraud, the development of tools to enhance "the review and analysis of a company's financial results and related complex information" has the potential to strengthen professional skepticism.

Finally, while this paper has focused on finding the root cause or causes of client anomalies, the same type of analysis might be useful when analyzing audit firm deficiencies. As James Doty, Chairman of the PCAOB, predicted in an April 2011 address to the Council of Institutional Investors, the future of audit oversight includes an enhanced PCAOB focus on the root causes of audit deficiencies. In particular, PCAOB inspections will "press firms to identify root causes of deficiencies and address them" (Doty 2011).

REFERENCES

AICPA (American Institute of Certified Public Accountants). 2009. *AICPA Professional Standards*. American Institute of Certified Public Accountants – AICPA. New York, N.Y.

Anderson, U., and L. Koonce. 1998. Evaluating the sufficiency of causes in audit analytical procedures. *Auditing: A Journal of Practice & Theory* 17 (1): 1-12.

Asare, S., and A. Wright. 1995. Normative and substantive expertise in multiple hypotheses evaluation. *Organizational Behavior and Human Decision Processes* 66: 171-184.

Asare, S., and A. Wright. 1997a. Evaluation of competing hypotheses in auditing. *Auditing: A Journal of Practice & Theory* 16 (1): 1-13.

Asare, S., and A. Wright. 1997b. Hypothesis revision strategies in conducting analytical procedures. *Accounting, Organizations and Society* 22 (8): 737-755.

Asare, S. and A. Wright. 2003. A note on the interdependence between hypothesis generation and information search in conducting analytical procedures. *Contemporary Accounting Research* 20 (2): 235-251.

Bedard, J. and S. Biggs. 1991. Pattern recognition, hypothesis generation, and auditor performance in an analytical task. *The Accounting Review* 66 (July): 622-643.

Bedard, J., S. Biggs, and J. DiPietro. 1998. The effects of hypothesis quality, client management explanations, and industry experience on audit planning decisions. *Advances in Accounting* 16: 49-73.

Blocher, E. and J. Willingham. 1993. *Analytical Review: A Guide to Performing Analytical Procedure.* New York: McGraw-Hill.

Bonner, S., and N. Pennington. 1991. Cognitive processes and knowledge as determinants of auditor expertise. *Journal of Accounting Literature*,1-50.

Brewster, B. 2011. How a systems perspective improves knowledge acquisition and performance in analytical procedures. *The Accounting Review* 86 (3): 915-943.

Caster, P. and K. Pincus. 1996. An empirical test of Bentham's theory of the persuasiveness of evidence. *Auditing: A Journal of Practice & Theory* 15 (Supplement): 15, -22.

Center for Audit Quality. 2010. *Deterring and Detecting Financial Reporting Fraud: A Platform for Action.* Available at: http://www.thecaq.org/Anti-FraudInitiative/CAQAnti-FraudReport.pdf.

Cianci, A. and J. Bierstaker. 2009. The impact of positive and negative mood on the hypothesis generation and ethical judgments of auditors. *Auditing: A Journal of Practice & Theory* 28 (2): 119-144.

Coglitore, F. and R. Berryman. 1988. Analytical procedures: a defensive necessity. *Auditing: A Journal of Practice & Theory* 7 (2): 150-163.

Curtis, E. and S. Turley. 2007. The business risk audit: a longitudinal case study of an audit engagement. *Accounting, Organizations and Society* 32: 439-461.

Doty, J. 2011. Looking Ahead: Auditor Oversight. April 4, 2011 speech to the Council of Institutional Investors 2011 Spring Meeting, Washington, D.C.

Dutta, S. and R. Srivastava.1993. Aggregation of evidence in auditing: A likelihood perspective. *Auditing: A Journal of Practice & Theory*12 (Supplement):137-160.

Dutta, S. and R. Srivastava. 1996. A Bayesian perspective on the strength of evidence in auditing. *Indian Journal of Accounting* XXVII (June): 4-18.

Edwards, A.1984.*Likelihood: An Account of the Statistical Concept of Likelihood and its Application to Scientific Inferences.* Cambridge: Cambridge University Press.

Fornelli, C. and A. Desmond. 2011. The art of skepticism. *Directorship:* February 16, 2011.

Green, W. and K. Trotman. 2003. An examination of different performance outcomes in an analytical procedures task. *Auditing: A Journal of Practice & Theory* 22 (2): 219-235.

International Audit and Assurance Standards Board (IAASB). 2009. *International Standard on Auditing (ISA) 200: Overall Objectives of the Independent Auditor and the Conduct of an Audit in Accordance with International Standards on Auditing*

Ismail, Z. and K. Trotman. 1995. The impact of the review process in hypothesis generation tasks. *Accounting, Organizations and Society* 20 (5): 345-357.

Jamal, K., P. Johnson, and R. Berryman. 1995. Detecting framing effects in financial statements. *Contemporary Accounting Research*12 (1):85-105.

Janvrin, D. 2008. To what extent does internal control effectiveness increase the value of internal evidence? *Managerial Auditing Journal* 23 (3): 262-282.

Johnson, P., S. Grazioli, K. Jamal, and R. Berryman. 2001. Detecting deception: adversarial problem solving in a low base rate world. *Cognitive Science* 25 (3): 355-392.

Knapp, M. 2011. *Contemporary Auditing: Issues & Cases, 8th Edition.* South-Western Cengage Learning, Mason, OH, USA.

Kida, T. and B. Luippold. 2009. The impact of initial information ambiguity on the accuracy of analytical review judgments. http://ssrn.com/abstract=1113496.

Knechel, W. R. 2007. The business risk audit: origins, obstacles and opportunities. *Accounting, Organizations and Society* 32: 383-408.

Knechel, W. R., S. Salterio, and N. Kochetova-Kozloski. 2010. The effect of benchmarked performance measures and strategic analysis on auditors' risk assessments and mental models. *Accounting, Organizations and Society* 35: 316-333.

Koonce, L., N. Seybert, and J. Smith. 2011. Causal reasoning in financial reporting and voluntary disclosure. *Accounting, Organizations and Society* 36: 209-225.

Libby, R. 1985. Availability and the generation of hypotheses in analytical review. *Journal of Accounting Research* 23 (2): 648-667.

MacDuffie, J. 1997. The road to "root cause": Shop-floor problem-solving at three auto assembly plants. *Management Science* 43 (4): 479-502.

Mautz, R. and H. Sharaf. 1961. *The Philosophy of Auditing*. American Accounting Association.

Mock, T. and J. Turner. 2005. Auditor identification of fraud risk factors and their impact on audit programs. *The International Journal of Auditing* 9 (1): 59-77.

Mock, T., A. Wright, R. Srivastava, and H.Lu. 2008. The Framing and Evaluation of Multiple Hypotheses. *Asia Pacific Journal of Accounting and Economics* 15: 123–140

Morris, M. and R. Larrick. 1995. When one cause casts doubt on another: A normative analysis of discounting in causal attribution. *Psychological Review* 102 (2): 331-355.

O'Donnell, E. and J. Perkins. 2011. Assessing Risk with Analytical Procedures: Do Systems-Thinking Tools Help Auditors Focus on Diagnostic Patterns? *Auditing: A Journal of Practice & Theory* 30 (4): 273-283.

PCAOB (Public Company Accounting Oversight Board). 2008. PCAOB Release No. 2008-008: *Report on the PCAOB's 2004, 2005, 2006 and 2007 Inspections of Domestic Annually Inspected Firms*.

PCAOB. 2010a. Auditing Standard #12: *Identifying and Assessing Risks of Material Misstatement.*

PCAOB. 2010b. Auditing Standard #14: *Evaluating Audit Results*.

Pearl, J. 1990. Bayesian decision methods. In *Uncertain Reasoning* edited by G. Shafer and J. Pearl, 345-352. San Mateo, California: Morgan Kaufmann Publishers, Inc.

Schultz, J. Jr., J. Bierstaker, and E. O'Donnell. 2010. Integrating business risk into auditor judgment about the risk of material misstatement: The influence of a strategic-systems-audit approach. *Accounting, Organizations and Society* 35: 238-251.

Shenoy, P. and G. Shafer. 1990. Axioms for probability and belief-function computation. In *Uncertainty in Artificial Intelligence* edited by R. Shachter, T. Levitt, J. Lemmer, and L. Kanal, Vol. 4: 169-198. North-Holland.

Spires, E. 1991. Auditors' evaluation of test-of-control strength. *The Accounting Review* 66 (April): 259-276.

Srivastava, R., A.Wright, and T. Mock. 2002. Multiple hypothesis evaluation in auditing. *Journal of Accounting and Finance* (Australia) 42 (3): 251-277.

Trompeter, G. and A. Wright. 2010. The world has changed – have analytical procedure practices? *Contemporary Accounting Research* 27 (2): 669-700.

Van Wallendael, L. 1989. The quest for limits on noncomplementarity in opinion revision. *Organizational Behavior and Human Decision Processes*43 (3):385-405.

Van Wallendael, L., and R. Hastie. 1990. Tracing the footsteps of Sherlock Holmes: Cognitive representations of hypothesis testing. *Memory and Cognition*18: 240-250.Wright, W., N. Jindanuwat, and J. Todd. 2004. Computational models as a knowledge management tool: a process model of the critical judgments made during audit planning. *Journal of Information Systems* 18 (1): 67-94.

APPENDIX A: PROOF OF THEOREM 1

We use the induction approach to prove Theorem 1. First we develop analytical formulas for two potential causes (hypotheses), H_1 and H_2, of an effect F and then by induction we derive the formula for n causes (hypotheses) of an effect F. In general, we assume the presence of effect F to be not known with certainty. That is, there is uncertainty in knowing whether the effect F is present. The presence of effect F is observed through a piece of evidence E_F.

Two Hypotheses Case

Consider a situation where there are two hypotheses, H_1 and H_2. We want to determine the posterior probabilities of the two hypotheses after observing the evidence, E_F, pertaining to an observed effect F, and two other items of evidence, E_1, pertaining to H_1, and E_2, pertaining to H_2. In the present discussion we assume the three variables, H_1, H_2, and F, to be binary, i.e., either they are present represented by their names or not present represented by the symbol '~' in front of the name. For example, H_1 represents the state that hypothesis H_1 is present or true and $\sim H_1$ represent the state that the hypothesis H_1 is not present or not true.

To derive the formulas for the posterior probabilities of H_1 and H_2 based on Bayes' Rule, we use Shenoy and Shafer (1990) approach of combining probability information. Under this approach we first need to identify all the probability information relevant in our problem. Probability information on a variable (variables being H_1, H_2, and F) is expressed in terms of what Bayesian literature refers to as probability potentials (Shenoy and Shafer, 1990). The probability potentials at a variable essentially are probabilities or conditional probabilities associated with the variable, but are not necessarily normalized, i.e., they do not necessarily add to one. For example, the conditional probabilities associated with variable F due to the evidence

E_F could be expressed as potentials at F with two values, one for 'F' that it is present and the other for '~F' that it is not present, and represented as $P(E_F|F)$ and $P(E_F|\sim F)$. In general, we use the symbol $\phi(.)$ to express the potential for the argument given in the parenthesis. For example, the potentials on the state space {F, ~F} of the binary variable 'F' based on the conditional probabilities can be expressed as: $\phi(F) = P(E_F|F)$, and $\phi(\sim F) = P(E_F|\sim F)$. The following discussion provides the details of combining all the potentials and finally determining the overall potentials at H_1, and H_2, the variables of interest.

Step 1: Identify all the Probability Potentials for two hypotheses case in Figure 2

Probability Potentials at 'F 'due to Evidence E_F: $\begin{bmatrix} \Phi(F) \\ \Phi(\sim F) \end{bmatrix} = \begin{bmatrix} P(E_F|F) \\ P(E_F|\sim F) \end{bmatrix}$

Probability Potentials at variable H_1 *due to Evidence* E_1: $\begin{bmatrix} \Phi(H_1) \\ \Phi(\sim H_1) \end{bmatrix} = \begin{bmatrix} P(E_1|H_1) \\ P(E_1|\sim H_1) \end{bmatrix}$

Probability Potentials at Variable H_2 *due to evidence* E_2: $\begin{bmatrix} \Phi(H_2) \\ \Phi(\sim H_2) \end{bmatrix} = \begin{bmatrix} P(E_2|H_2) \\ P(E_2|\sim H_2) \end{bmatrix}$

Probability Potentials related to Causal Schemata:

The most general form of the relationship among three variables, H_1, H_2, and F, can be written in terms of the joint probability distribution over the entire joint space $\{H_1H_2F, H_1H_2\sim F, H_1\sim H_2F, H_1\sim H_2\sim F, \sim H_1H_2F, \sim H_1H_2\sim F, \sim H_1\sim H_2F, \sim H_1\sim H_2\sim F\}$ as:

$$\begin{bmatrix} \Phi(H_1H_2F) \\ \Phi(H_1H_2{\sim}F) \\ \Phi(H_1{\sim}H_2F) \\ \Phi(H_1{\sim}H_2{\sim}F) \\ \Phi({\sim}H_1H_2F) \\ \Phi({\sim}H_1H_2{\sim}F) \\ \Phi({\sim}H_1{\sim}H_2F) \\ \Phi({\sim}H_1{\sim}H_2{\sim}F) \end{bmatrix} = \begin{bmatrix} P(H_1H_2F) \\ P(H_1H_2{\sim}F) \\ P(H_1{\sim}H_2F) \\ P(H_1{\sim}H_2{\sim}F) \\ P({\sim}H_1H_2F) \\ P({\sim}H_1H_2{\sim}F) \\ P({\sim}H_1{\sim}H_2F) \\ P({\sim}H_1{\sim}H_2{\sim}F) \end{bmatrix}$$

Step 2: Combination of Potentials

We use Shenoy and Shafer (1990) approach to combine various potentials defined on different state spaces. First the potentials need to be vacuously extended to a common state space of the joint space and then they need to be point-wise multiplied. Under point-wise multiplication, each element of a potential is multiplied with the same element of another potential. Given below is an example of the vacuous extension of the potential at F.

$$\begin{bmatrix} \Phi(H_1H_2F) \\ \Phi(H_1H_2{\sim}F) \\ \Phi(H_1{\sim}H_2F) \\ \Phi(H_1{\sim}H_2{\sim}F) \\ \Phi({\sim}H_1H_2F) \\ \Phi({\sim}H_1H_2{\sim}F) \\ \Phi({\sim}H_1{\sim}H_2F) \\ \Phi({\sim}H_1{\sim}H_2{\sim}F) \end{bmatrix} = \begin{bmatrix} P(E_F|F) \\ P(E_F|{\sim}F) \\ P(E_F|F) \\ P(E_F|{\sim}F) \\ P(E_F|F) \\ P(E_F|{\sim}F) \\ P(E_F|F) \\ P(E_F|{\sim}F) \end{bmatrix} \quad \text{(A1)}$$

Similar results are obtained for the potentials at $H_1$and H_2.

Next, we point-wise multiply the four sets of potentials (three sets from the three variables H_1, H_2, and F, and one set from the Causal Schemata). In other words, we point-wise multiply the potentials defined earlier at F, H_1, H_2, and the potential for causal schema after

vacuously extending the potentials onto the joint space of the three variables. This multiplication yields the following potentials:

$$\begin{bmatrix} \Phi(H_1H_2F) \\ \Phi(H_1H_2{\sim}F) \\ \Phi(H_1{\sim}H_2F) \\ \Phi(H_1{\sim}H_2{\sim}F) \\ \Phi({\sim}H_1H_2F) \\ \Phi({\sim}H_1H_2{\sim}F) \\ \Phi({\sim}H_1{\sim}H_2F) \\ \Phi({\sim}H_1{\sim}H_2{\sim}F) \end{bmatrix} = \begin{bmatrix} P(E_1|H_1)P(E_2|H_2)P(E_F|F)P(H_1H_2F) \\ P(E_1|H_1)P(E_2|H_2)P(E_F|{\sim}F)P(H_1H_2{\sim}F) \\ P(E_1|H_1)P(E_2|{\sim}H_2)P(E_F|F)P(H_1{\sim}H_2F) \\ P(E_1|H_1)P(E_2|{\sim}H_2)P(E_F|{\sim}F)P(H_1{\sim}H_2{\sim}F) \\ P(E_1|{\sim}H_1)P(E_2|H_2)P(E_F|F)P({\sim}H_1H_2F) \\ P(E_1|{\sim}H_1)P(E_2|H_2)P(E_F|{\sim}F)P({\sim}H_1H_2{\sim}F) \\ P(E_1|{\sim}H_1)P(E_2|{\sim}H_2)P(E_F|F)P({\sim}H_1{\sim}H_2F) \\ P(E_1|{\sim}H_1)P(E_2|{\sim}H_2)P(E_F|{\sim}F)P({\sim}H_1{\sim}H_2{\sim}F) \end{bmatrix}. \quad \text{(A2)}$$

To determine the combined potentials at variable H_1, we marginalize the above potential in (A2) onto the state space of H_1. The marginalization process yields the following potentials at H_1.

$$\begin{bmatrix} \Phi(H_1) \\ \Phi({\sim}H_1) \end{bmatrix} = \begin{bmatrix} T1+T2+T3+T4 \\ T5+T6+T7+T8 \end{bmatrix}, \quad \text{(A3)}$$

where T1, T2, … T8 are defined below:

$$T1 = P(E_1|H_1)P(E_2|H_2)P(E_F|F)P(H_1H_2F);\ \ T2 = P(E_1|H_1)P(E_2|H_2)P(E_F|{\sim}F)P(H_1H_2{\sim}F),$$

$$T3 = P(E_1|H_1)P(E_2|{\sim}H_2)P(E_F|F)P(H_1{\sim}H_2F);\ T4 = P(E_1|H_1)P(E_2|{\sim}H_2)P(E_F|{\sim}F)P(H_1{\sim}H_2{\sim}F),$$

$$T5 = P(E_1|{\sim}H_1)P(E_2|H_2)P(E_F|F)P({\sim}H_1H_2F);\ T6 = P(E_1|{\sim}H_1)P(E_2|H_2)P(E_F|{\sim}F)P({\sim}H_1H_2{\sim}F),$$

$$T7 = P(E_1|{\sim}H_1)P(E_2|{\sim}H_2)P(E_F|F)P({\sim}H_1{\sim}H_2F),\ T8 = P(E_1|{\sim}H_1)P(E_2|{\sim}H_2)P(E_F|{\sim}F)P({\sim}H_1{\sim}H_2{\sim}F). \quad \text{(A4)}$$

Equation (A3) provides the overall potentials at node H_1 after combining all the items of evidence including the priors. The potentials in (A3) when normalized yield the posterior probability that H_1 is present given all the evidence about F, H_1, and H_2. To simplify the final result, we divide both the numerator and denominator of the normalized potentials by $P(E_1|{\sim}H_1)P(E_2|{\sim}H_2)P(E_F|{\sim}F)$ and define the following likelihood ratios, λ's:

$$\lambda_1 = P(E_1|H_1)/P(E_1|{\sim}H_1);\ \ \lambda_2 = P(E_2|H_2)/P(E_2|{\sim}H_2);\ \ \lambda_F = P(E_F|F)/P(E_F|{\sim}F). \quad \text{(A5)}$$

We obtain the following posterior probability of H_1: $P(H_1|E_1E_2E_F) = N_1/(N_1+D_1)$, where

$$N_1 = P(H_1H_2F) + P(H_1H_2{\sim}F)/\lambda_F + P(H_1{\sim}H_2F)/\lambda_2 + P(H_1{\sim}H_2{\sim}F)/(\lambda_2\lambda_F), \quad \text{(A6a)}$$

$$D_1 = P({\sim}H_1H_2F)/\lambda_1 + P({\sim}H_1H_2{\sim}F)/(\lambda_1\lambda_F) + P({\sim}H_1{\sim}H_2F)/(\lambda_1\lambda_2) + P({\sim}H_1{\sim}H_2{\sim}F)/(\lambda_1\lambda_2\lambda_F). \quad \text{(A6b)}$$

Similarly the posterior probability $P(H_2|E_1E_2E_F)$ can be written as $P(H_2|E_1E_2E_F) = N_2/(N_2+D_2)$, where

$$N_2 = P(H_1H_2F) + P(H_1H_2{\sim}F)/\lambda_F + P({\sim}H_1H_2F)/\lambda_1 + P({\sim}H_1H_2{\sim}F)/(\lambda_1\lambda_F), \quad \text{(A7a)}$$

$$D_2 = P(H_1{\sim}H_2F)/\lambda_2 + P(H_1{\sim}H_2{\sim}F)/(\lambda_2\lambda_F) + P({\sim}H_1{\sim}H_2F)/(\lambda_1\lambda_2) + P({\sim}H_1{\sim}H_2{\sim}F)/(\lambda_1\lambda_2\lambda_F). \quad \text{(A7b)}$$

We can rewrite the above posterior probabilities as a combined equation as follows:

$$P(H_i|E_1E_2E_F) = N_i/(N_i+D_i), i = 1, \text{or } 2. \quad \text{(A8a)}$$

$$N_i = [P(H_iH_jF) + P(H_iH_j{\sim}F)/\lambda_F] + [P(H_i{\sim}H_jF) + P(H_i{\sim}H_j{\sim}F)/\lambda_F]/(\lambda_j), \quad \text{(A8b)}$$

$$D_i = [P({\sim}H_iH_jF) + P({\sim}H_iH_j{\sim}F)/\lambda_F]/(\lambda_i) + [P({\sim}H_i{\sim}H_jF) + P({\sim}H_i{\sim}H_j{\sim}F)/\lambda_F]/(\lambda_i\lambda_j), \quad \text{(A8c)}$$

where i = 1, 2, j = 1, 2, and $i \neq j$. Equations (A8a) - (A8c) represent the desired posterior probabilities for the two hypotheses in terms of the likelihood ratios, λ's, given the evidence E_1, E_2, and E_F. Equations (A8a-A8c) are general expressions for the revised probabilities of H_1 and H_2. The joint probabilities, $P(H_1H_2E_F)$, $P(H_1H_2{\sim}E_F)$, $P(H_1{\sim}H_2E_F)$, $P(H_1{\sim}H_2{\sim}E_F)$, $P({\sim}H_1H_2E_F)$, $P({\sim}H_1H_2{\sim}E_F)$, $P({\sim}H_1{\sim}H_2E_F)$ and $P({\sim}H_1{\sim}H_2{\sim}E_F)$, in (A8b) and (A8c) fully determine the causal schema for the two causes, H_1 and H_2. Next we extend the above for the situation where there are 'n' potential causes for the effect F.

Extension to 'n' Causes or Hypotheses of an Effect F

As we can see from (A8a)–(A8c), the two-cause case, that there is a definite pattern in the formula. For example, whenever an element of the joint space contains a hypothesis in its negation form, the probability of that element is divided by the corresponding likelihood ratio. For example, consider the second and fourth terms in (A8b). The probability of the second

element (H_iH_j~F) is divided by λ_F, the likelihood ratio pertaining to H_F, and the probability of the fourth element (H_i~H_j~F) is divided by $\lambda_j\lambda_F$, the product of the likelihood ratios for H_j and F. This pattern continues in all the terms in (A8b), and (A8c). Thus, by induction we can generalize the results for the case where we have n causes as follows:

$$P(H_k|E_1E_2\ldots E_nE_F) = N_k/(N_k+D_k), \text{ where,} \tag{A9a}$$

$$N_k = [P(\bigcap_{i=1}^{n} H_i\, F) + (1/\lambda_F)\, P(\bigcap_{i=1}^{n} H_i \sim F)] + \sum_{\substack{j=1 \\ j\neq i\neq k}}^{n} \left[P(H_k \bigcap_{\substack{i=1 \\ i\neq k\neq j}}^{n} H_i \sim H_j F) + (1/\lambda_F) P(H_k \bigcap_{\substack{i=1 \\ i\neq k\neq j}}^{n} H_i \sim H_j \sim F)\right]/(\lambda_j)$$

$$+ \sum_{\substack{j=1,\, m=1 \\ j\neq i\neq k\neq m}}^{n} [P(H_k \bigcap_{\substack{i=1 \\ i\neq k\neq j\neq m}}^{n} H_i \sim H_j \sim H_m F) + (1/\lambda_F) P(H_k \bigcap_{\substack{i=1 \\ i\neq k\neq j\neq m}}^{n} H_i \sim H_j \sim H_m \sim F)]/(\lambda_j\lambda_m)$$

$$+ \ldots\ldots + [P(H_k \bigcap_{\substack{i=1 \\ i\neq k}}^{n} \sim H_i\, F) + (1/\lambda_F)\, P(H_k \bigcap_{\substack{i=1 \\ i\neq k}}^{n} \sim H_i \sim F)]/\prod_{\substack{i=1 \\ i\neq k}}^{n} \lambda_i\,, \tag{A9b}$$

$$D_k = \left[P(\sim H_k \bigcap_{\substack{i=1 \\ i\neq k}}^{n} H_i F) + (1/\lambda_F) P(\sim H_k \bigcap_{\substack{i=1 \\ i\neq k}}^{n} H_i \sim F)\right]/\lambda_k + \sum_{\substack{j=1 \\ j\neq i\neq k}}^{n} \left[P(\sim H_k \bigcap_{\substack{i=1 \\ i\neq k\neq j}}^{n} H_i \sim H_j F) + (1/\lambda_F) P(\sim H_k \bigcap_{\substack{i=1 \\ i\neq k\neq j}}^{n} H_i \sim H_j \sim F)\right]/(\lambda_j\lambda_k)$$

$$+ \sum_{\substack{j=1,\, m=1 \\ j\neq i\neq k\neq m}}^{n} [P(\sim H_k \bigcap_{\substack{i=1 \\ i\neq k\neq j\neq m}}^{n} H_i \sim H_j \sim H_m F) + (1/\lambda_F) P(\sim H_k \bigcap_{\substack{i=1 \\ i\neq k\neq j\neq m}}^{n} H_i \sim H_j \sim H_m \sim F)]/(\lambda_j\lambda_m\lambda_k)$$

$$+ \ldots\ldots + [P(\bigcap_{i=1}^{n} \sim H_i\, F) + (1/\lambda_F)\, P(\bigcap_{i=1}^{n} \sim H_i \sim F)]/\prod_{i=1}^{n} \lambda_i\,. \tag{A9c}$$

QED

APPENDIX B: DERIVATION OF DISCOUNTING OR INFLATING OF POSTERIOR PROBABILITIES

In this section we derive a general formula for the *discounting* or *inflating* of the posterior probability of a cause, say A, of an effect F, when we have evidence about another cause, say B. Mathematically, one can express discounting of cause A in relation to cause B as:

$$\text{Discounting of A} = P(A|F) - P(A|FE_B) \tag{B1}$$

Morris and Larrick (1995) analyzed discounting under the condition where B is known with certainty. Here, we extend their model to the more likely audit situation where B is not known with certainty using Theorem 1 and Equation B1. This is achieved by setting the number of hypotheses n = 2; $H_1 = A$; $H_2 = B$; $E_2 = E_B$; $\lambda_1 = 1$ (no direct evidence that A is present); $\lambda_2 = \lambda_B$; and $\lambda_F \rightarrow \infty$ (we assume here the effect F is observed with certainty) in (1)-(3). This process yields the following expression for $P(A|FE_B)$:

$$P(A|FE_B) = [P(AB|F) + (1/\lambda_B)P(A\sim B|F)]/[(1/\lambda_B) + (1 - 1/\lambda_B)P(B|F)] \tag{B2}$$

Combining Equations (B1) and (B2), one obtains the following expression for discounting:

$$\textbf{Discounting of A} = P(A|F)-P(A|FE_B) = (\lambda_B-1)[P(A|F)P(B|F)-P(AB|F)]/[1+(\lambda_B-1)P(B|F)] \tag{B3}$$

Under the condition where A and B are the only causes of F, one can show[3] that discounting will always occur whether A and B are sufficient or non-sufficient causes of F, as long as we have a positive piece of evidence pertaining to cause B (i.e., $\lambda_B>1$). One can also

[3]For $\lambda_B>1$, in general, the factor $(\lambda_B-1)[P(A|F)P(B|F)-P(AB|F)]$ in the numerator of (B3) is always positive irrespective of the relationship between A and B and irrespective of whether A and B are sufficient or non-sufficient causes of F. First, consider the case of A and B being the sufficient causes of F, i.e., $F=A\cup B$. Using Bayes' rule, the above condition yields: $[P(A|F)P(B|F)-P(AB|F)]=[P(A)P(B)-P(AB)P(F)]/P(F)^2$. Since $P(F)=P(A)+P(B)-P(AB)$, the above expression reduces to $[P(A)-P(AB)][P(B)- P(AB)]/P(F)^2$, which is always positive except when $A\subseteq B$ or $B\subseteq A$, it is zero. Next, consider the second case where A and B are non-sufficient causes of F (i.e., $P(F|A)<1$, and $P(F|B)<1$). Assume that $A'=A\cap F$, $B'=B\cap F$. Since A' and B' are the only causes, $F=A'\cup B'$. For this case, one can write the discounting of A from (B3) as: Discounting of $A = (\lambda_B-1)[P(A')P(B') -P(A'B')P(F)]/[P(F)2 + (\lambda_B-1)P(B')P(F)]$. Again, using the above approach one can show that the numerator is always positive.

show that in this case, discounting will always occur whether they are assumed to be negatively correlated, independent, or positively correlated prior to observing F.

Morris and Larrick (1995) define the correlation coefficient between two causes A and B as: $\rho_{AB}=[P(AB) - P(A)P(B)]/[P(A)P{\sim}A)P(B)P(\sim B)]^{1/2}$ where ~ represents the negation of the cause. Similarly, the correlation coefficient between A and B given F can be written as: $\rho_{AB|F}=[P(AB|F) - P(A|F)P(B|F)]/[P(A|F)P{\sim}A|F)P(B|F)P(\sim B|F)]^{1/2}$. One can show that irrespective of the correlation between the two exhaustive causes of an effect before the effect is observed, the two causes will always be negatively correlated after the effect is observed regardless of whether the two causes are sufficient or non-sufficient. This is the reason that discounting will always occur in such a situation. This finding is logical because, if we observe an effect and know that there are only two possible causes and one of them is already found to be present, then it is less likely that the other cause is present and vice versa.

Figure 1: Causal Schemas for Two Causes, A and B, of an effect F with the following probabilities: P(A) = 0.2, P(B) = 0.1. A and B are assumed to be the only causes for F.

Panel A: Schema before observing F for sufficient and independent causes, A and B. F = A∪B

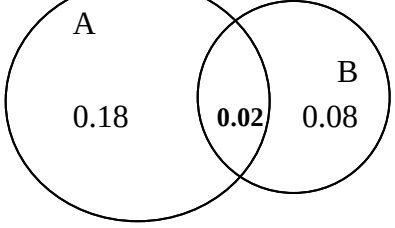

P(F) = 0.28,
Joint prior probabilities representing the causal schema:
P(ABF) = 0.02, P(A~BF) = 0.18,
P(~ABF) = 0.08, P(~A~BF) = 0
Correlation coefficient: $\rho_{AB} = 0$.

Panel B: Schema after observing F for sufficient causes, A and B. A and B are independent before observing F (see Panel A).

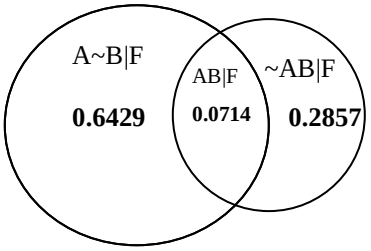

Joint conditional probabilities representing the causal schema:
P(AB|F) = 0.0714, P(A~B|F) = 0.6429,
P(~AB|F) = 0.2857, P(~A~B|F) = 0,
Correlation coefficient: $\rho_{AB|F} = -0.848$

Panel C: Schema before observing F for non-sufficient and independent causes, A and B.

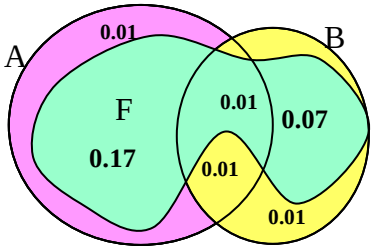

P(F) = 0.25
Joint prior probabilities representing the causal schema:
P(AB) = 0.02, P(A~B) = 0.18,
P(~AB) = 0.08, P(~A~B) = 0.72,
Correlation coefficient: $\rho_{AB} = 0$.

Panel D: Schema after observing F for non-sufficient causes, A and B. A, and B are independent before observing F (see Panel C). The shaded area represents F.

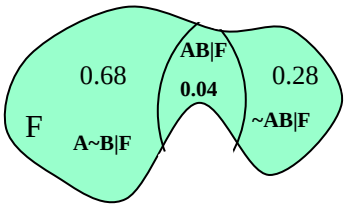

Conditional joint probabilities representing the causal schema:
P(AB|F) = 0.04, P(A~B|F) = 0.68,
P(~AB|F) = 0.28, P(~A~B|F) = 0
Correlation coefficient: $\rho_{AB|F} = -0.909$

Figure 2: A General Model for Multiple Hypotheses Evaluation

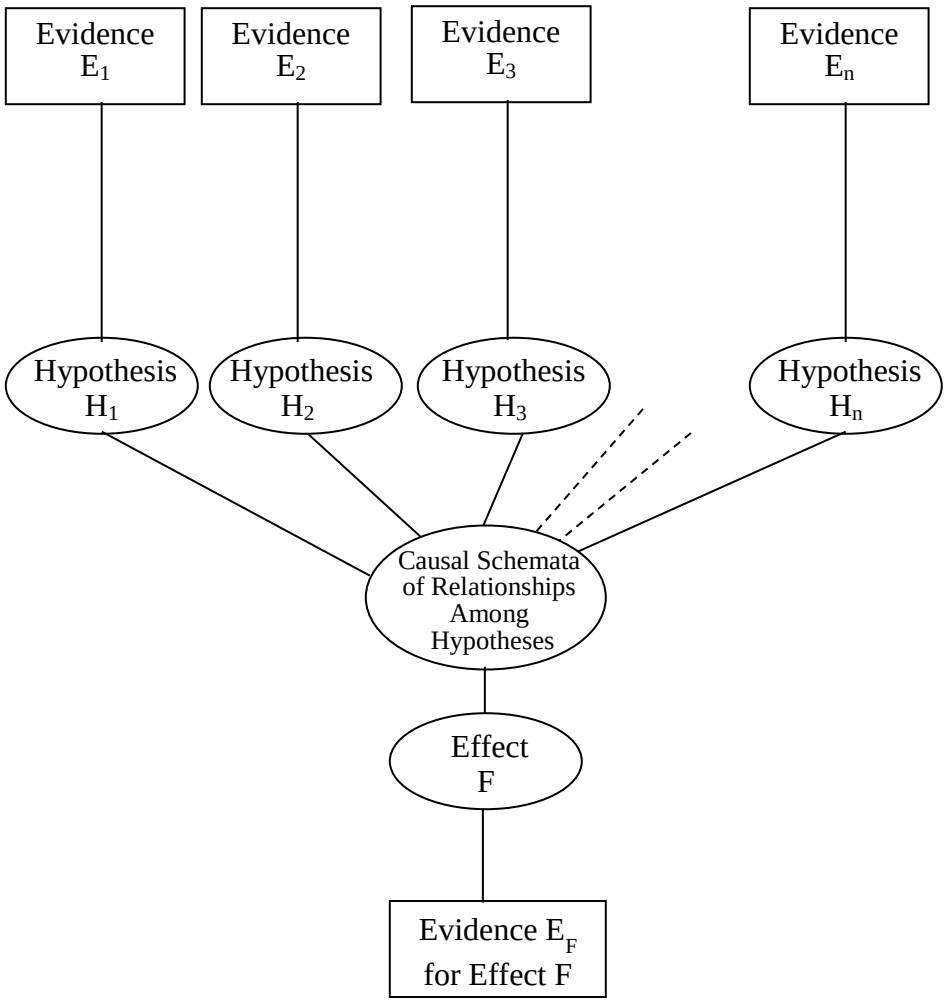

Figure 3: Discounting of Cause A versus the strength of evidence in support of Cause B when A and B are sufficient and the only possible causes of F with P(A)=0.2, and P(B)=0.1.

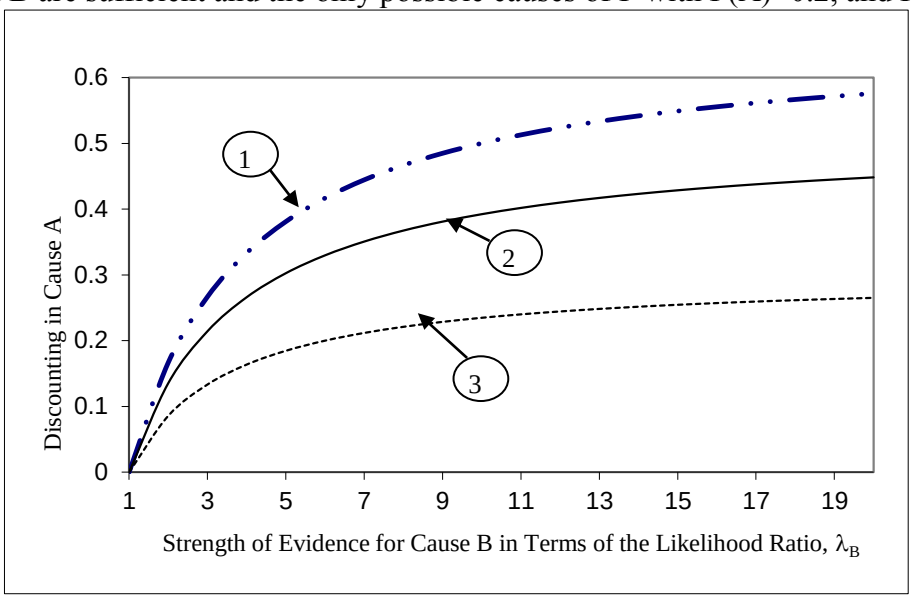

(1) A and B are negatively correlated with $\rho_{AB} = -0.667$, $\rho_{AB|F} = -1.0$, P(AB) = 0.0, P(F) = 0.3

(2) A and B are independent with $\rho_{AB} = 0$, $\rho_{AB|F} = -0.849$, P(AB) = 0.02, P(F) = 0.28

(3) A and B are positively correlated with $\rho_{AB} = 0.25$, $\rho_{AB|F} = -0.612$, P(AB) = 0.05, P(F) = 0.25

Figure 4: Discounting of Cause A versus the strength of evidence in support of Cause B when A and B are non-sufficient and not the only possible causes of F with P(A)=0.2, P(B)=0.1, P(AF)=0.18, P(BF)=0.07, and P(F)=0.28.

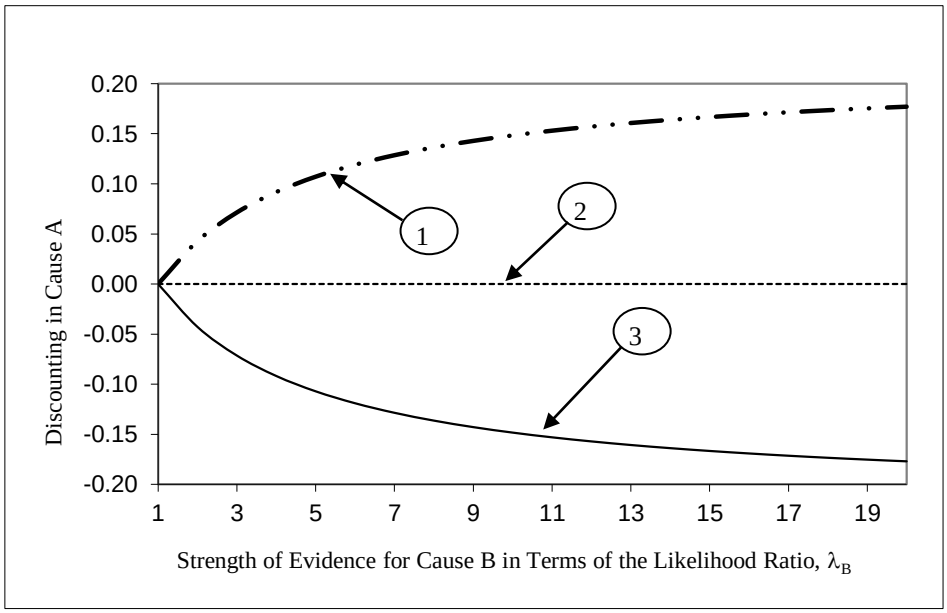

(1) Correlation coefficients, $\rho_{AB}=0.083$, $\rho_{AB|F} = -0.258$, P(AB)=0.03, P(ABF)=0.03;

(2) $\rho_{AB}=0.208$, $\rho_{AB|F} = 0$, P(AB)=0.045 and P(ABF)=0.045 and

(3) $\rho_{AB}=0.33$, $\rho_{AB|F} = 0.258$, P(AB)=0.06 and P(ABF)=0.06.

Figure 5**: Posterior Probabilities of two Causes, A and B, of an Effect F that are *non-exclusive* and *non-exhaustive*, as a function of the strength of evidence in support of both causes assuming A and B are positively correlated.**

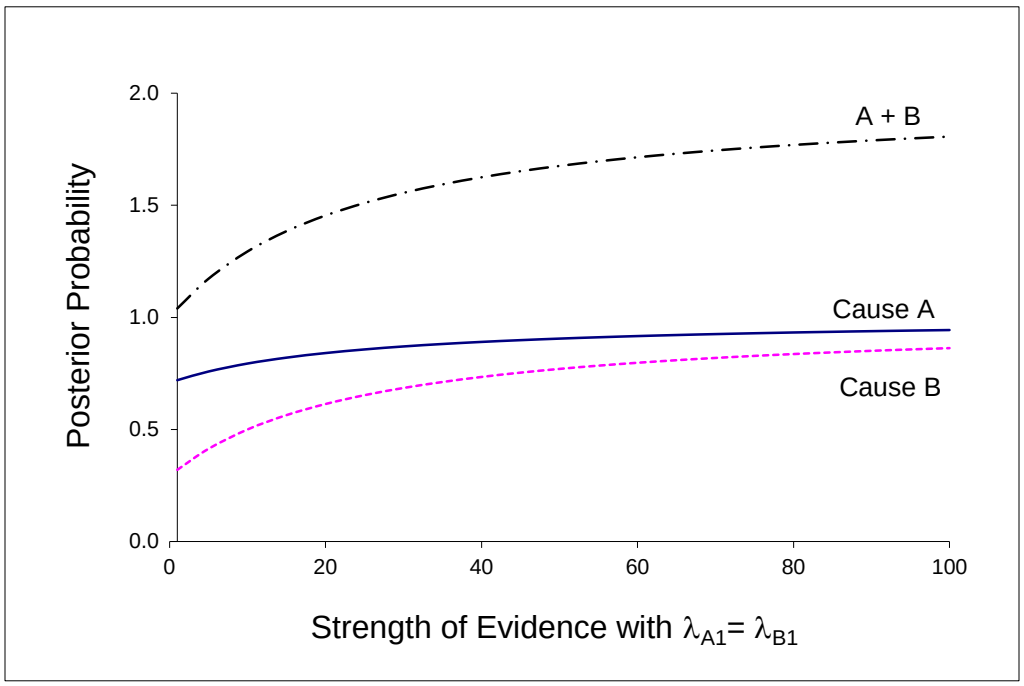

Figure 6: Illustration of Causal Schema after observing the effect F for a case where* $P(H_1|F) = 0.4$, $P(H_2|F) = 0.1$, $P(H_3|F) = 0.2$, and $P(H_4|F) = 0.4$, and positive correlation of 0.612 between H_1 & H_3, negative correlation of -0.272 between H_1 & H_2, and -0.167 between H_2 & H_3, and an independent relationship among H_4 and the set $\{H_1, H_2, H_3\}$.

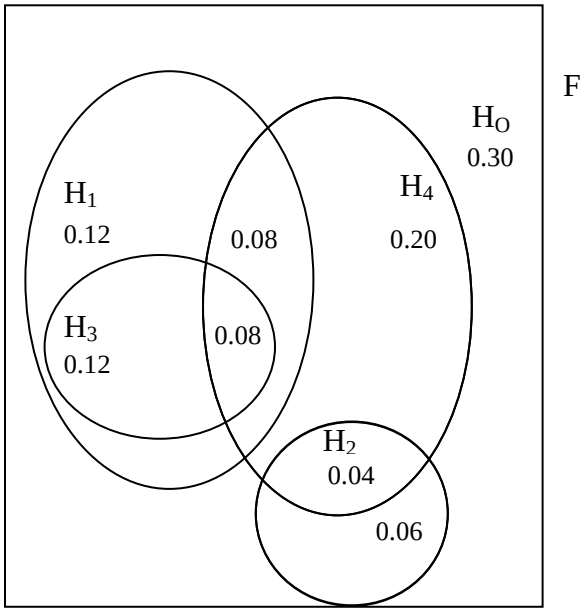

*****These values were chosen to fit the example used by Asare and Wright(1997a). The joint probability distribution is obtained in two steps. In the first step, we determine the joint probability distribution for H_1, H_2, and H_3, based on their marginal probabilities, and their correlation coefficients. In the second step we cross-multiply the distribution obtained in the first step with the probability distributions for H_4 for them to be independent of each other.

Figure 7: Graph of Revised Probabilities for Non-Exhaustive and Non-Exclusive Case with Mixed Correlations as a function of successive piece of evidence

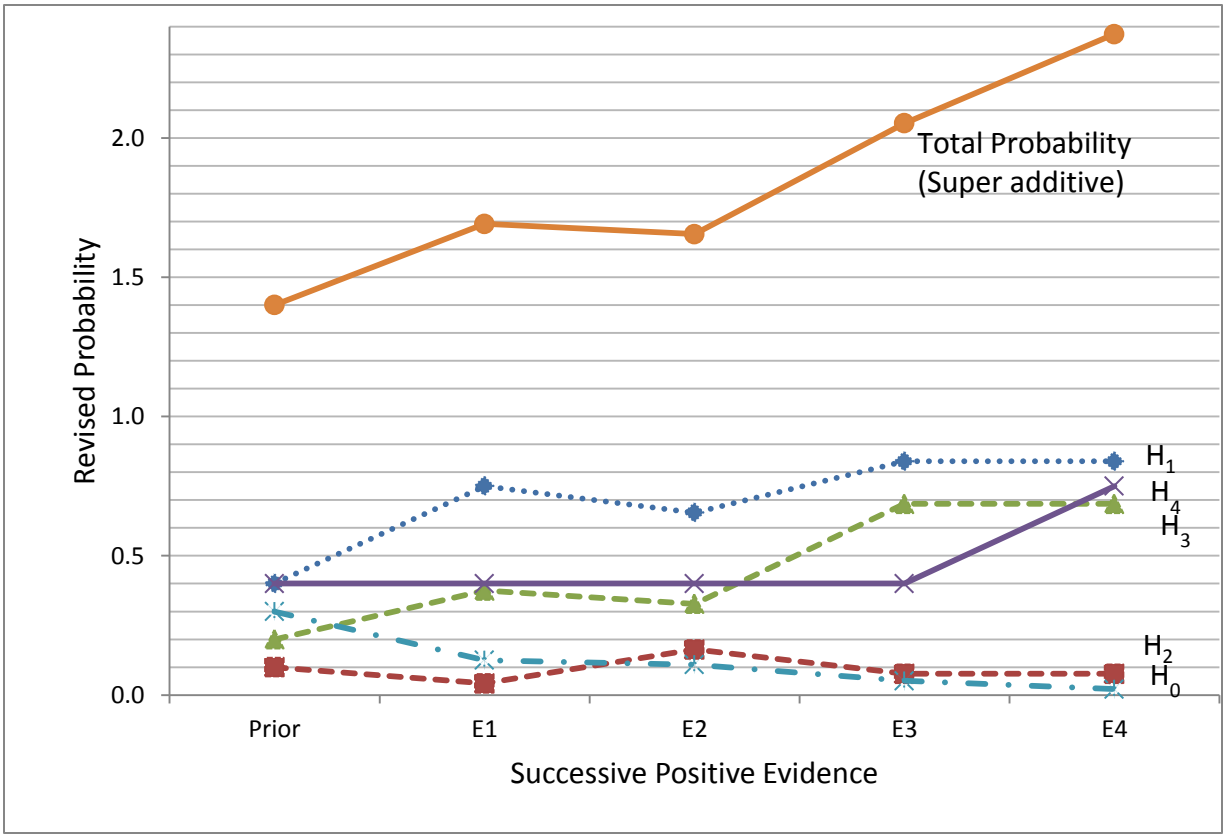

Table 1: Terms used in the Discussion of the Framework

Terms describing hypothesis sets

Mutually exclusive: If one hypothesis (candidate cause) is supported, then others will not be supported. There is just one cause responsible for the observed effect.	**Non-exclusive**: More than one hypothesis may be supported. Multiple causes may be responsible for the observed effect.
Exhaustive: All potential causes are identified in the initial hypothesis set.	**Non-exhaustive**: All potential causes may not be identified in the initial hypothesis set.

Terms describing evidence

Likelihood ratio: For any single item of evidence E pertaining to candidate cause A, E_A, the likelihood ratio λ_A is defined as: $\lambda_A = P(E_A|A)/P(E_A|\sim A)$.	
Perfectly diagnostic: evidence that conclusively supports or confirms that a particular cause is present ($\lambda \rightarrow \infty$) or is not present ($\lambda=0$).	**Non-diagnostic**: evidence that is neutral concerning a particular cause being present. Evidence with no diagnostic value has $\lambda=1$.
Positive evidence: evidence that supports or confirms that a particular cause is present. For an individual item of positive evidence (support that cause A is present), the likelihood ratio takes values between 1 and infinity ($1<\lambda<\infty$).	**Negative evidence**: evidence that counters or disconfirms that a particular cause is present. For negative evidence, the likelihood ratio takes values between zero and one ($0<\lambda<1$).

Terms describing relationships between hypotheses A and B for an effect F

Sufficient causes: The sufficiency condition implies that cause A or cause B *always* results in the effect F implying that P(F|A)=1 and P(F|B)=1, which yields F = A∪B when only two causes are present.	**Non-sufficient causes**: When candidate causes are non-sufficient, the assumed potential causes may not always cause the effect (i.e., P(F|A) < 1, and P(F|B) < 1).