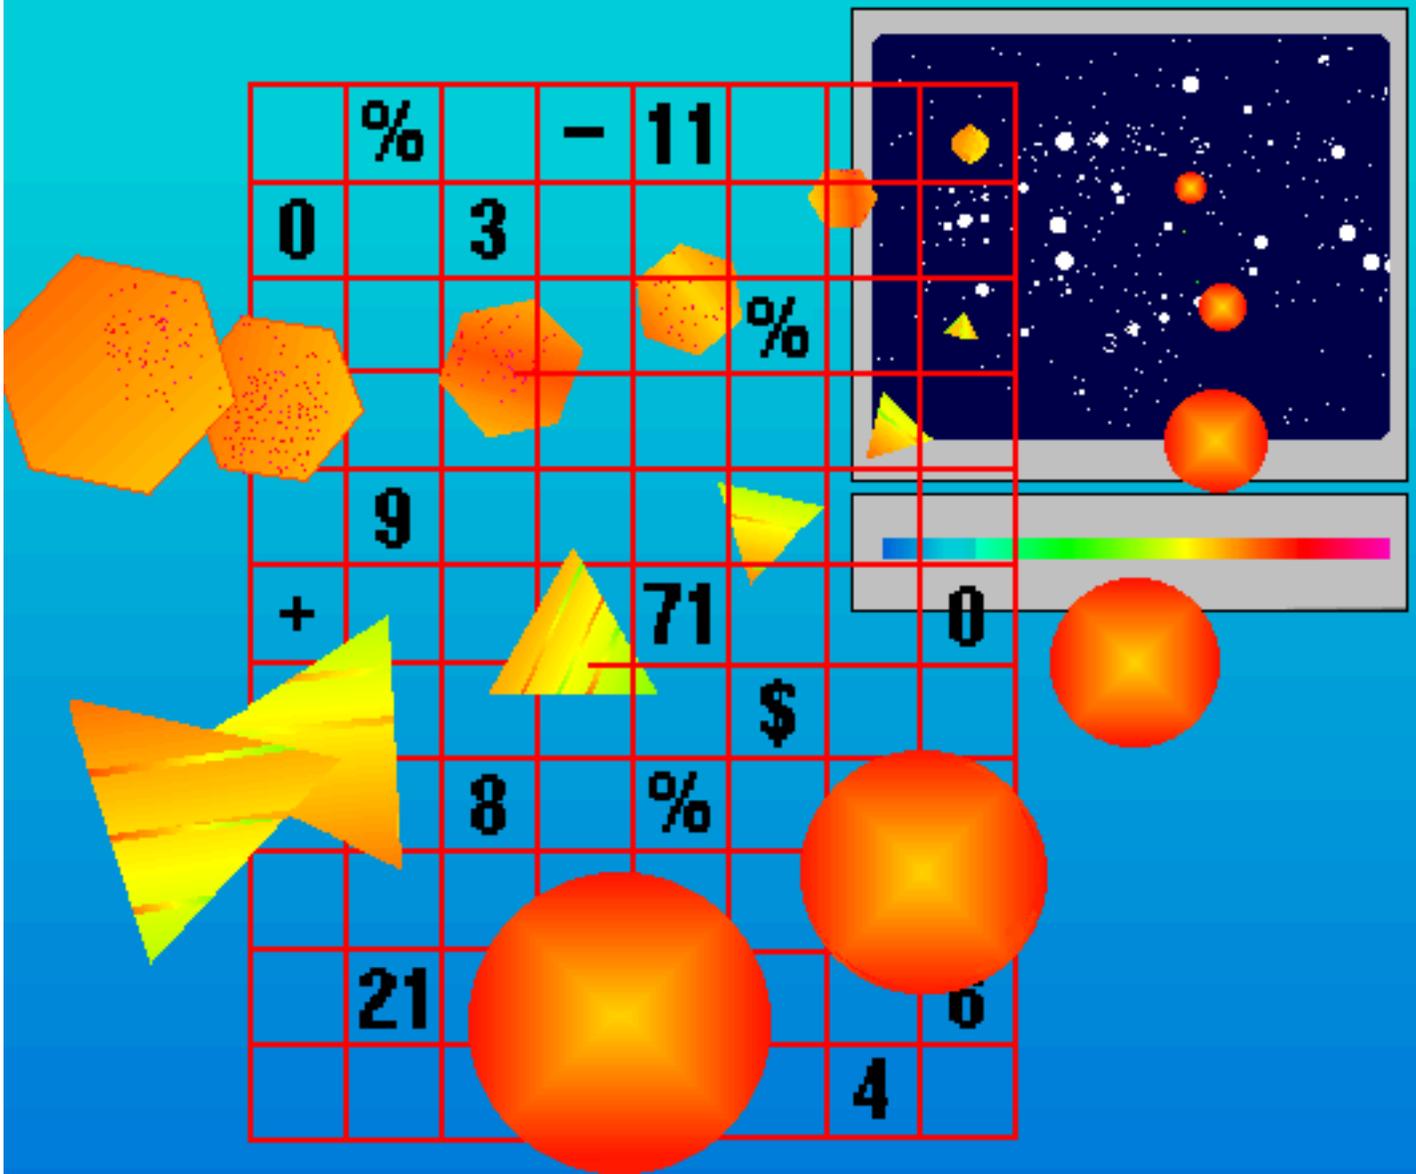


A Knowledge-Based

Vasant Dhar, Barry Lewis, & James Peters



Model of Audit Risk

Within the academic and professional auditing communities, there has been growing concern about how to accurately assess the various risks associated with performing an audit. These risks are difficult to conceptualize in terms of numeric estimates. This article discusses the development of a prototype computational model (computer program) that assesses one of the major audit risks—inherent risk. This program bases most of its inferencing activities on a qualitative model of a typical business enterprise.

Many decision-making problems involve making choices with incomplete information. Models of decision making under conditions of risk are well established in decision-theory literature. In these models, risk and return (payoffs) are specified in terms of numeric estimates, and the goal is to make a decision that maximizes some expected value. In addition, new information can be combined using a decision rule (such as Bayes' rule) for deriving revised estimates of risk.

There are several problems, however, for which quantifiable risk estimates are difficult to obtain. In these problems, the decision maker might have to perform a significant amount of problem solving if asked to provide a numeric estimate of risk associated with a proposition or a state of nature, and even then, attach qualifying comments to the numeric estimate (Dhar and Pople 1987). In this article, we describe one such problem in which the concept of risk is difficult to conceptualize in terms of quantitative estimates and hence inappropriate to manipulate in terms of standard belief calculi. Specifically, we focus on the problem of an auditor faced with the task of assessing the likelihood that a client's financial statements will contain material errors. Especially for large accounts, it is important that the auditor be confident that the risk of making an erroneous decision is low. Given the increased competitiveness of the auditing profession, however, the auditor is limited from a pragmatic standpoint in terms of time and other resources that can be assigned to each client. Good planning is therefore essential for an efficient audit; before beginning the actual process of information gathering and substantive test-

ing, the auditor must have expectations about specific general ledger accounts that might be particularly risky in a specific case and plan tests accordingly.

We have studied audit planning in two large, international accounting firms over the last two years. We have come to recognize the process of risk assessment as a complex one that involves understanding the effects of a variety of economic and organization-specific factors on accounts. We characterize the problem of risk assessment as knowledge-based because knowledge about the client's history, about recent events specific to a firm or industry, and about the specific internal characteristics of a firm are crucial in shaping the auditor's judgment about the risks associated with the accounts and hence crucial to the audit plan.

There is currently much interest in developing knowledge-based support systems to support audit planning. First, accounting firms are interested in developing systematic methods for risk assessment and audit planning. They feel that knowledge-based systems can help in this respect. Second, accounting firms are concerned with losing information about a client when an auditor leaves the audit team. By preserving the results of past experiences with auditing a client, knowledge-based systems might also help in training new members of an audit team. In this research project, we concentrated on the development of systematic methods of risk assessment by trying to understand and model the risk assessment process. Once the details of this process are understood, the model should provide a sound basis for designing intelligent support systems.

In the next section, we define more precisely the different types of risks associated with auditing and which of these we are interested in modeling. We also make certain assertions about how auditors assess risk. These assertions are based on empirical observations of the audit planning process (tape recordings of actual audit planning meetings as well as interviews). In the third section, we describe the knowledge-based model of audit risk, focusing on the knowledge representation employed in a system designed to assess a certain component of audit risk called inherent risk. We conclude with auditors' reactions about the inadequacies of this system and the types of information that must be gathered and modeled in order to overcome these limitations.

Risks in Auditing

The problem confronting an auditor is referred to in the accounting literature as one of audit risk assessment. Cushing and Loebbecke (1983) provide a detailed discussion of the historical development of a model of audit risk, a review of the literature dealing with the appropriateness of the model, and a discussion of the problems that auditors might encounter in trying to implement it. Most of this literature has focused on the American Institute of Certified Public Accountants' (AICPA) audit risk model as presented in Statements on Auditing Standards (SAS) 47:

audit risk = inherent risk * control risk * detection risk .

Inherent risk is the susceptibility of an account balance or class of account balances to error that could be material, assuming that there are no related internal accounting controls. *Control risk* is the risk that this error in the account balance will not be caught by the client's internal control system. *Detection risk* is the risk that any error not detected by the control system will not be detected by the audit procedures, thereby affecting final financial statement balances. The AICPA emphasizes that this model of audit risk is a purely conceptual one and says little about whether or how its various components can be measured.

According to the AICPA, the risk-assessment process should occur during audit planning. Auditors should determine an acceptable level for audit risk, assess the levels of inherent risk and control risk, and then determine the level of detection risk based on these assessments. Auditors should assess the inherent risk of specific accounts by reviewing a variety of factors that are specific to the client, to the client's industry, or to the economy in general and by determining the impact of these factors on individual accounts. This risk assessment can help determine the nature, timing, and extent of tests of the client's internal control systems (AICPA 1985) and account balances.

The professional literature also provides lists of factors that should be considered in assessing inherent risk. Peat Marwick's audit manual (Peat Marwick International 1985), for example, presents a representative list:

1. Monetary amount associated with the account
2. Susceptibility of asset to theft
3. Complexity required to determine amounts to be entered in the account
4. Degree of management judgment involved in valuing the account
5. Degree to which external events affect values in the account
6. Past history of error
7. Degree to which client's financial condition motivates management to misstate the amount in the account
8. Experience of the personnel performing accounting functions involving the account

Although this model of audit risk and the list of factors provide a conceptual foundation for understanding the role of inherent risk, they do not indicate how experienced auditors actually use these factors in conjunction with industry- and client-specific knowledge to form judgments about inherent risk. Our central objective in this research has been to explicate this model and use it to build a system that can be used to assess inher-

ent risk. In this model, inherent risk is not computed as a numeric estimate for each account. Rather, risk at the account level is expressed in the form of potential reasons or hypotheses that are a by-product of a more general reasoning process involving analysis of financial statements and client and industry factors as a whole. In other words, financial statements are analyzed in the context of industry- and client-specific data; this analysis results in potential reasons for errors in specific accounts, without quantified estimates of error. These potential reasons or hypotheses help to target the auditor's evidence gathering in subsequent phases of the audit. The following quote from an experienced auditor provides a succinct description of the information-gathering process:

I think the process you go through to obtain that knowledge really is to gain an understanding of the client's business, an understanding of the client, an understanding of how the fluctuations in the economy might affect a client's business. You compare the client's business to other businesses in the same industry to see if they are having consistent operating results and if not, if there are logical reasons for it; if they are having consistent operating results, is that what you expected? I mean, you develop expectations in your mind of what you expect to see and, to the extent results don't conform to that yet, you begin asking questions to obtain the necessary knowledge.

In the rest of this section, we provide a general description of the process of inherent risk assessment before describing the details of the model. For descriptive empirical studies of the various types of risks, the reader is referred to Gibbins and Wolf (1982); Jiambalvo and Waller (1984); Libby, Artman, and Willingham (1985); and Boritz, Graber, and Lemon (1986).

Inherent Risk Assessment

Based on the results of our two-year field research project (more completely described in Peters, Lewis, and

Dhar 1987), we characterize the approach to assessing the inherent risk of material error in a given general ledger account balance as one involving differential analysis. The analysis is essentially change-driven, in that specific external or internal changes generate expectations of change in accounts.

In our model, the inherent risk evaluation process begins by generating expectations for account balances. Specifically, the auditor identifies changes that have occurred in the firm or its environment and determines how those changes should interact with historic trends to produce an expected balance in the account. In order to do this, the auditor uses an understanding of the relationships between environmental factors and general ledger accounts. By making use of these relationships, the auditor develops expectations about how the changes observed should affect the balances in a given account. For accounts in which actual balances are outside the expected range, the auditor first reviews factors that might create or affect management incentives to misstate the account balances (for example, the existence of a compensation plan keyed to reported earnings). At the same time, the auditor considers factors that might influence the likelihood that management could or would misstate that particular account balance, because of the degree of judgment allowed in the determination of account balances. The auditor also considers the complexity of the transactions or accounting for a particular account, because such factors might also be responsible for deviations from expected balances. Based on this analysis, the auditor decides if additional evidence will be needed to determine whether the difference between the expected balance and the actual balance was caused by an error in the expectation generating process, a legitimate response by management to a change in the environment, a questionable response by management, or an unintentional error.

Assertions of the Inherent Risk Evaluation Model

Here we list five assertions that char-

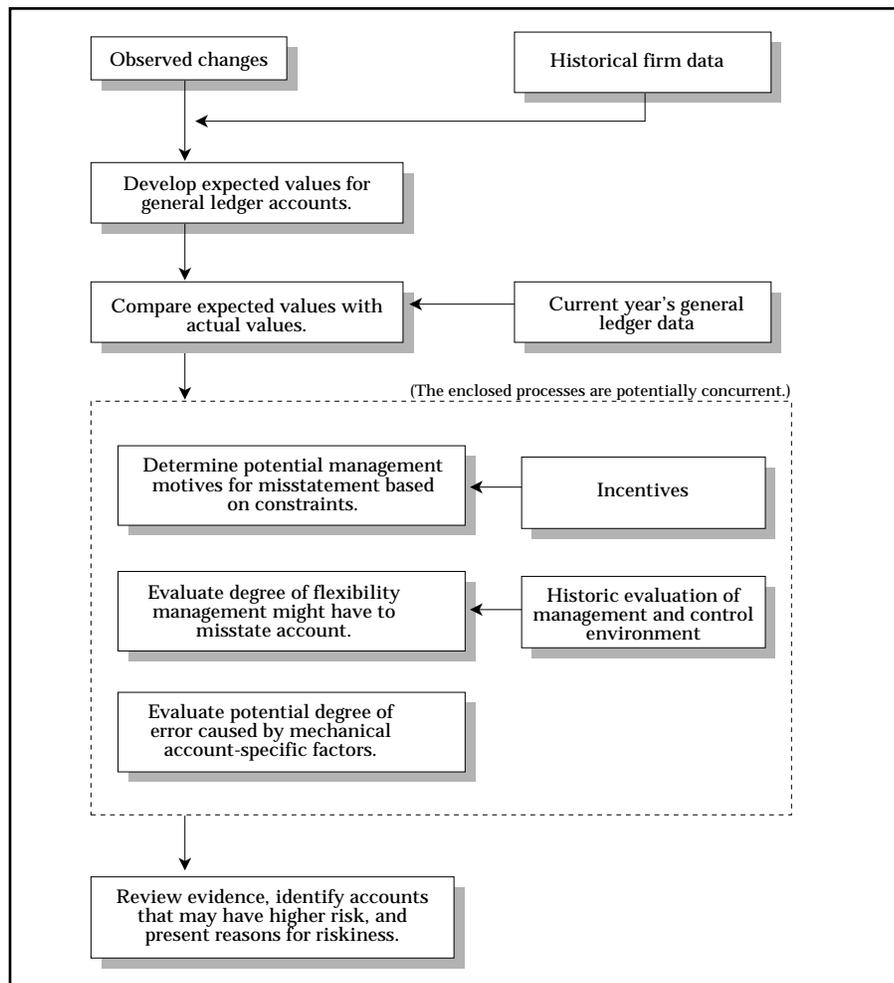


Figure 1. The Process of Inherent Risk Assessment.

acterize our model of inherent risk determination:

Assertion 1: Auditors generate expectations concerning account balances and investigate balances that differ from these expectations.

Assertion 2: Auditors generate expectations about accounts based on changes in events or circumstances relative to prior years.

Assertion 3: Management's incentives and ability to manipulate account balances affect the assessment of inherent risk.

Assertion 4: Inherent risk assessments are generated on an account-by-account basis.

Assertion 5: To be useful, inherent risk assessments should explain why a given account is risky rather than merely provide a quantitative estimate of risk.

In the next section, we describe in detail how we model the various knowledge components referred to in the assertions.

Knowledge Representation

Figure 1 shows the general process of inherent risk assessment. First, expected values of general ledger accounts are generated. The data input to this process are historical data and changes in external and internal conditions from the previous year. The output of the first stage, expected account balances, is then compared with the actual general ledger data for the period being audited. A materiality judgment is incorporated in this stage. The output of the second stage is a list of accounts for which expectations are not in accord with reality. The third stage involves

determining whether some of management's incentives or constraints might have caused managers to take actions that account for the differences between observed and expected values of accounts. The fourth stage involves making a judgment about management's track record (that is, its ability and inclination) in manipulating accounts. Finally, auditors indicated that the complexity associated with computing account balances (for example, "Last in first out (LIFO) is more complex than first in first out (FIFO) for valuing inventory") also affects inherent risk. We refer to these as mechanical account-specific factors.

In summary, the five main processes are expectation generation, evidence checking, incentive checking, the judging of management's track record, and the evaluation of the mechanical complexity associated with each account. These processes compose the overall control structure of the system, as shown in figure 1. In our current system, we have not focused on modeling the last two processes. We comment on the consequences of these missing components in the next section. In the rest of this section, we discuss the first three processes, including the knowledge used in each process and how it is represented.

Expectation Generation

Auditors generally have what we term base expectations for a client's accounts. These base expectations are formed from observations about the client. For example, the auditor might expect that the client's market share this year should have fallen by 5 percent because a major customer was acquired by a competitor in a vertical merger. For our model, the relevant input is a quantitative one, accompanied by a descriptive explanation of the reason for the quantitative input. Specifically, an input is a trituple of the form

<firm factor, percent change, comment> ,

where the firm factor is a general ledger account or an economic factor (for example, market demand) that influences a general ledger account, followed by its percent change and a

comment that denotes the reason for the change. The comment, which is the reason for the base expectation, is for the user's benefit only.

To assess the impacts of these changes on specific accounts, the system makes use of what Bouwman (1978) and we have termed an *internal model of the firm*. This model consists of three types of factors associated with the firm, namely, general ledger accounts, financial statement totals or subtotals, and exogenous factors such as market demand. These factors are related through algebraic equations. In effect, the internal model of the firm has a node-link structure; *nodes* are the arithmetic operators involved in the equations, and *links* represent the three types of firm factors involved in these algebraic equations. Figure 2 is a graphic representation of a segment of the internal model of the firm.

Nodes and links are represented as structured objects. The operator object has three slots, namely, operator, inputs, and output. This object type also has two types of functions associated with it that perform computation (in object-oriented language, these functions are called methods). The first type is for computing an output value (in absolute or percentage terms) from input. The second type computes qualitative directional change in the output given a qualitative change in only one input. The usefulness of both types of methods is apparent in the discussion that follows.

The firm factor object contains slots for the following information: the type of node (whether it represents a general ledger account), historic value, current value, change value, pointers to operator objects that compute the current value of the firm factor, and expectation values (generated as a consequence of changes of other firm nodes' values) that are posted for further analysis. There are several ways of computing the value of a firm factor.

A base expectation is indexed to the internal model by its first component, which is a pointer to a firm-factor object. When a base expectation is entered as a change, it specifies a

change in the firm-factor object to which it points, which in turn causes the operator objects that compute its value to recompute their output values. This procedure is carried out recursively until all changes have been propagated through the network. The end result of the propagation is that each general-ledger-account object's expectation slot consists of a list of expectations for that account—each of these generated using a particular formula.¹

Evidence Checking

The second stage is a straightforward one. The percentage-value change in the expectation is checked against the actual change computed from the general ledger data. A material difference exists if the difference exceeds a pre-set range. Expectations that are not materially different from actual values are excluded from further consideration, while the remainder are put into a "flagged accounts list." The latter set also includes accounts for which no expectations were generated by the system but for which material changes have occurred.

Incentive Checking

Determining discrepancies is only part of the overall auditing process. It is equally important to understand the reasons for the discrepancies. This problem is one of determining whether and how management discretion might have been exercised with respect to these accounts. To make this determination, it is necessary to know management's motivations and how these motivations affect specific accounts.

We refer to motivational factors as incentives, which are represented as constraints. An incentive can represent either a restriction placed on account balances as part of a contract (for example, bond covenants), contractual arrangements that tie management's compensation to levels in account balances, or expectations by outsiders concerning levels of account balances (for example, public expectations of steadily growing earnings). Examples of incentives might include a management bonus plan that provides additional compensation if gross profit exceeds 110 percent of last

year's profit, or a bond covenant that requires the current ratio to be greater than 2.

Incentives are represented as constraints consisting of any relational operator (except equality) with two arguments. The arguments can be constants or functions that return values of firm-factor nodes. At least one of the arguments must be of the latter type. In effect, incentives are indexed to the internal model of the firm using these constraints. To illustrate, the second incentive listed earlier would be represented in prefix Lisp-like notation as

```
(>= cvalue(current-ratio) 2.0),
```

where *cvalue* is a function that returns the current value of the current-ratio object.

The default assumption used by the system is that account balances have been influenced by the incentives. The problem is to determine how the manipulations might have occurred. Given the structure of the internal model of the firm, there are an infinite number of ways of manipulating an account because of the infinite number of combinations of input values that can produce a certain output. The problem is therefore one of limiting this search so that only relevant combinations are considered.

The incentive-checking process proceeds as follows. For each constraint, a materiality computation is applied the same way as it is for general ledger accounts. If we denote the actual value of the object referred to in the constraint as *A*, the expected value as *E*, and the relational operator in the constraint as *R*, then for a constraint of the form (*R arg1 arg2*), the following decision rules are applied.

Incentive Rule 1 states that

If $A = E$

then if *R* is of type ">" or ">="

then management has been motivated to increase *arg1*

(or decrease *arg2*)

else management has been motivated to decrease *arg1*

(or increase *arg2*).

The rationale for this rule is that *A* has been manipulated to be equal to *E*.

Incentive Rule 2 states that

If $A < E$

then if *R* is of type "<" or "<="

then management has been motivated to decrease *arg1*

(or decrease *arg2*)

else management has been motivated to increase *arg1*

(or increase *arg2*).

Thus, when there is a significant difference between actual and expected values, the difference stems from one of two factors: (1) management gave up the constraint as not achievable and hence shifted account values so that they may be favorable for the following period or (2) management satisfied the constraint by a large enough margin and hence shifted account values (in a direction opposite to the operator) to position the firm more favorably for the following period without compromising this year's objective. In effect, the rule incorporates a "shifting logic" that might be

used by management to plan for subsequent periods.

Applying these decision rules results in a hypothesis of the following form:

Account *i* has been increased.

To investigate this hypothesis, the system sets up the following goal:

Determine how account *i* has been increased.

This investigation involves a search beginning at object *i* in the firm model, determining ways in which *i* can be increased. As we pointed out, this result can be achieved by numerous combinations. For example, the current ratio could be increased by increasing current assets or decreasing current liabilities in various combinations, each of which could in turn be manipulated by other firm factors (for example, current assets might be

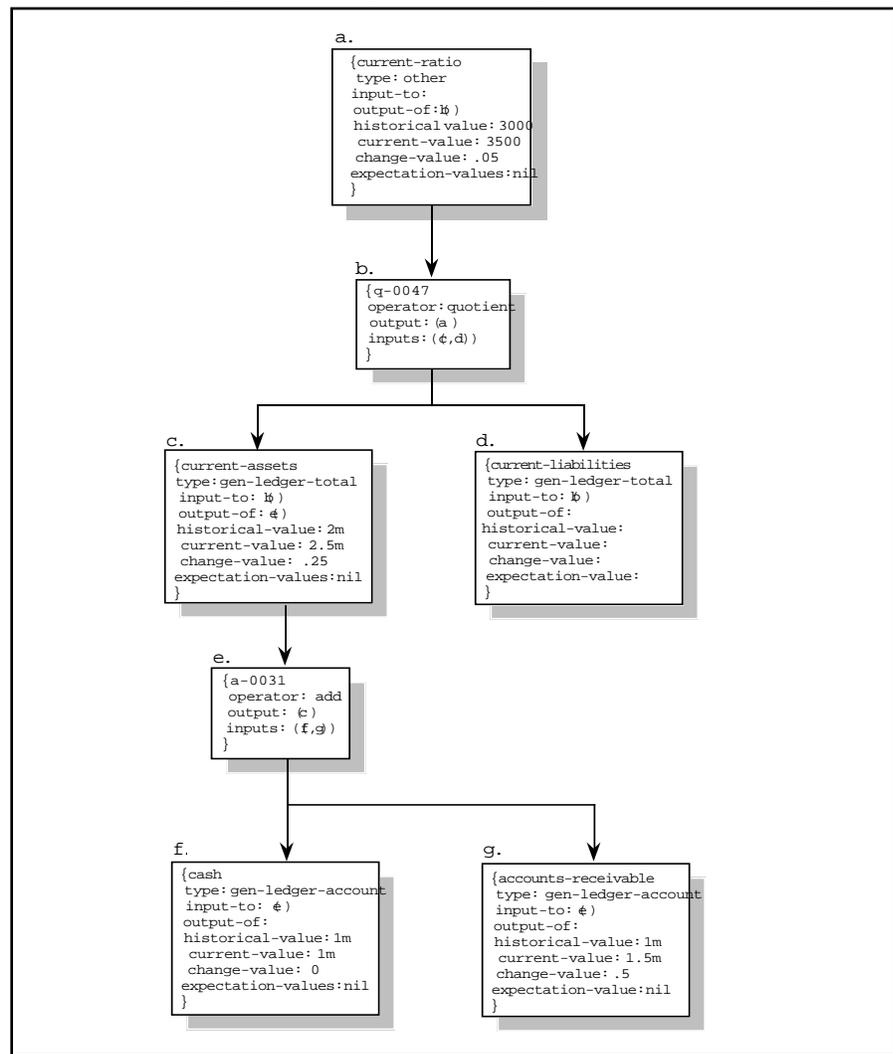


Figure 2. A Segment of the Internal Model of the Firm.

increased by overstating accounts receivable). The heuristic used in focusing the search comes from the deviant expectations generated in the previous stage. Specifically, given the hypothesis “account *i* has been increased,” the system computes which of the input values in its equation must be increased, holding all other components of the equation constant to make the hypothesis true. The analysis makes use of knowledge about the type of operator being applied. Evidence confirming the hypothesis is deemed to be found if there is an account that deviates from expectations in a direction that matches the hypothesis. For example, if there was a bond covenant that the firm maintain a current ratio of greater than 2, then an unexpected increase in accounts receivable (determined in the expectation-generation phase) would suggest that the bond covenant constraint had in fact played a role in inflating accounts receivable. In effect, this reasoning would constitute the explanation for the deviant expectation.

Evaluation and Summary

The model we described has been demonstrated to several auditors who helped us design it. Their comments and critiques have been useful in helping us understand the limitations of the system and what must be done to alleviate them.

Overly Data-Driven Expectation Generation

Our system’s control of attention is driven by the base expectations specified by the user. However, we now realize that auditors’ analyses are to a large extent driven by the structure of the general ledger, beginning with current assets (accounts such as cash) and ending with extraordinary items. More important, as this “downward scanning” proceeds, expectations about values of accounts yet to be reviewed become increasingly constrained. If the observed value does not fit with expectations, it is necessary to reinterpret the accounts reviewed thus far. This type of progressive constraint posting does not take place in our model.

For example, suppose that the auditor observes that inventory is low, in accordance with the base expectation for this account. This observation might cause the auditor to expect that payables will also be low (because of decreasing purchases). If payables are not low, the original interpretation must be altered. It could be that there has been a sharp increase in sales (causing a depletion in inventory) and reorders have been lagging (and hence not yet reflected in payables). A different explanation might be that a technological change in the client’s industry has rendered inventory obsolete, so its value must be marked down. Similarly, other explanations can be generated, some more plausible than others. Fundamentally, the more plausible of these explanations make use of *compiled knowledge*, that is, typical patterns in relationships among firm factors observed by experienced auditors.

Conceptually, the type of analysis we described makes use of knowledge about qualitative (directional) relationships between accounts and other firm factors. This knowledge is in fact available in our model of the firm. However, our system makes use only of the model of the firm in trying to determine whether any incentives could be responsible for a deviant expectation. It is now clear that the compiled knowledge about relationships among firm factors in this model must be used in determining whether an account is deviant in the first place. To implement this functionality, this knowledge will have to be gathered from further observations of experienced auditors and represented as compiled-knowledge links between nodes in the model of the firm. Such links have been used to great advantage in the CADUCEUS system (Pople 1982) as a way of focusing search.

Overemphasis on Incentives

A second criticism, related to the first, was that our model emphasized management motives too strongly. This effect is more noticeable because the system ignores relationships among firm factors, as we noted. Although incentives are an important determinant of inherent risk, we have

found that auditors usually do not consider them unless the value of the constrained item in the incentive formula is close to the boundary. That is, our default assumption that account balances have been influenced by incentives is not accurate. In fact, the opposite is often the case. Also, the second decision—rule that shifting may occur when the value of the constrained item in the incentive formula is some distance from the bound—is often applied inappropriately. This decision rule must be specialized. Auditors also did not perceive all incentives as having an equal potential impact on management’s actions. For example, violating a bond covenant would be considered more serious than not achieving a budgeted goal. In summary, a hierarchy of incentives and accounts is necessary to distinguish the important determinants of risk from the incidental ones.

Extrapolation of Data to Year-End Values

Audit planning typically occurs before the client’s year end. Therefore, auditors do not have actual, unaudited year-end balances with which to test their expectations. The current system makes use of a simple trend extrapolation to compute year-end balances. It turns out that the extrapolation is more complex than we had envisioned and can involve various factors. For example, seasonal fluctuations are an important component in estimating year-end balances, and vary according to the particular industry. Our model does not make distinctions among industries. Furthermore, some accounts, such as extraordinary items, cannot be extrapolated.

Unintentional Errors

An important addition to the current model that auditors deemed necessary was some mechanism for dealing with unintentional errors. *Unintentional errors* are those that occur in a given account balance without any intention on the part of management or employees to create the error. The lack of intention differentiates these errors from the potential errors dealt with by the incentive-checking portion of the model. Examples include

miscalculating inventory values or accidentally failing to include all outstanding invoices in the accounts payable balance. Our system has no way to deal with such errors. Our conceptual model also does not deal with the characteristics of the firm's management that could affect all, or at least a broad range, of accounts (for example, management's concern for internal control or high employee turnover in the accounting department). The auditors believed that these two classes of factors that influence the likelihood of an unintentional account error were extremely important in determining appropriate inherent risk for a given account. Further interactions with auditors will be needed to determine how they use these factors to assess inherent risk and how their assessment of these factors is combined with the expectation and incentive data that the model currently processes.

Summary and Conclusions

The model we described is a summary of a two-year research effort, during which we attempted to observe, in as much detail as possible, the audit planning process in real cases. From our findings, we came to recognize the difference between descriptions in the literature and what actually occurs in practice. It is clear that auditors do not consider it appropriate to generate numeric estimates of risk on an account-by-account basis; they prefer to analyze a client's financial statements using knowledge about changes in the industry or in the client, management's motivations, prior track record, and so on. It has become increasingly apparent to us that if computer-based systems are to support auditors in this task, they must be capable of modeling this range of knowledge. Although our model is still inadequate in the ways we described, it is proving to be very useful in sharpening our understanding of the process of inherent risk assessment. In particular, giving auditors a real system to use (however limited) helped elicit data that would have been virtually impossible to obtain

using interviews and other data-gathering techniques. We are continuing to work with several of the auditors who helped us design the first version of the model. We hope to address the limitations of this model and report on our progress with the new model in the near future.

Acknowledgments

Funding for this project was provided by the Peat Marwick Foundation through their Research Opportunities in Auditing Program and by a Faculty Research Grant of the Graduate School of Business of the University of Pittsburgh. We are especially grateful for the many hours contributed by the nine professional accountants who formed the basis for our conceptual model of risk analysis. We are also grateful to Paul Frishkoff, Jenny Gaver, Bob Clemen, and Barry Floyd for commenting on an earlier draft of this article.

References

- AICPA. 1985. *AICPA Professional Standards, Volume A: U.S. Auditing Standards*. American Institute of Certified Public Accountants.
- Boritz, J. E.; Graber, B. G.; and Lemon, W. M. 1986. An Experimental Study of Review of Preliminary Audit Strategy by External Auditors. In *Proceedings of the Contemporary Accounting Research Conference*, 1-154. Canadian Academic Accounting Association.
- Bouwman, M. J. 1978. *Financial Diagnosis: A Cognitive Model of the Process Involved*. Ph.D. diss., Graduate School of Industrial Administration, Carnegie-Mellon Univ.
- Cushing, B. E., and Loebbecke, J. K. 1983. Analytical Approaches to Audit Risk: A Survey and Analysis. *Auditing: Journal of Practice and Theory* 3(1): 23-41.
- Dhar, V., and Pople, H. E. 1987. Rule-Based Versus Model-Based Representations for Explaining and Generating Expert Behavior. *Communications of the ACM* 30(6): 542-555.
- Gibbins, M., and Wolf, F. M. 1982. Auditor's Subjective Decision Environment—The Case of a Normal External Audit. *Accounting Review* LVII(1): 105-124.
- Jiambalvo, J. J., and Waller, W. 1984. Decomposition and Assessments of Audit Risk. *Auditing: Journal of Practice and Theory* 3(2): 80-88.
- Libby, R.; Artman, J. T.; and Willingham, J. J. 1985. Process Susceptibility, Control

Risk and Audit Planning. *Accounting Review* LX(2): 212-230.

Peat Marwick International. 1985. *PMI Audit Manual—U.S.* Montvale, N.J.: Peat Marwick International.

Peters, J.; Lewis, B.; and Dhar, V. 1987. Assessing Inherent Risk during Audit Planning: A Computational Model. Unpublished.

Pople, H. E. 1982. Heuristic Methods for Imposing Structure on Ill-Structured Problems: The Structuring of Medical Diagnostics. In *Artificial Intelligence in Medicine*, ed. P. Szolovits, 119-184. Boulder, Colo.: Westview.

Notes

1. Because an account can be computed in several ways, the values could be in opposite directions and hence cancel out. We preserve all values because they contain information that would be lost in adding them.