註：本科考題共為三大題，請三題選二題作答，若三題均回答，則以最低分之二題計算。

1. Please read the following article thoroughly and answer the questions either English or in Chinese.

(1) Please derive a research framework of this article. You may draw a figure to express the causal relationships among the variables shown in this article. (20%)

(2) If you are asked to propose a research plan, how would you "empirically" develop the following sections for the research? Please describe the: (a) sample; (b) data; (c) methodology; and (d) expected contributions which you would like to employ. (30%)

Animosity on the Home Front: The Intifada in Israel and Its Impact on Consumer Behavior

Aviv Shoham,
Moshe Davidow,
Jill G. Klein, and
Ayalla Ruvio

The research we present herein uses animosity theory (Klein, Ettenson, and Morris 1998) and extends it to an intracountry context. For this purpose, we conducted the study in Israel and extended animosity theory to account for the buying behavior of Jewish Israelis from Arab Israelis following the advent of the second Intifada (the Palestinians' uprising against Israel's occupation of the West Bank and Gaza).

Thousands of Israelis and Palestinians have been killed since the second Intifada began in late 2000. This article deals specifically with the Arab Israelis' uprising (not Palestinians, a term that describes citizens of the Palestinian Authority in the West Bank and Gaza). Although this article focuses on anger held by Jewish Israelis in reaction to violent acts, we do not imply that Jewish Israelis were not involved in violence against Arab Israelis, and we do not argue the political justifications of either side in this conflict, which is beyond the scope of an article about marketing issues. We examine the consumer choice implications of this conflict, concentrating particularly on the relationship between the Israeli Jews and the Israeli Arabs who joined the Intifada in its early stages. Under normal (pre-Intifada) circumstances, both groups freely purchased each other's products (e.g., Dayan

Submitted August 2005
Accepted January 2006

Journal of International Marketing
© 2006, American Marketing Association
Vol. 14, No. 3, 2006, pp. 92–114
ISSN 1069-031X (print)
1547-7215 (electronic)

2001). In this article, we examine the impact of the conflict between Arab and Jewish Israelis on the consumer behavior of Jewish Israelis and its managerial implications.

There is ample anecdotal evidence in the popular Israeli press to support the view that the uprising has affected Jewish Israelis' consumer behavior profoundly. A leading daily reported (Katz 2002, p. 4) that "the last [postuprising] year has seen a 50% decrease in sales in the [Israeli] Arab sector. The industries hurt most include … tourism, furniture, and food…. The ones that were hurt most were those businesses in the Arab sector, which depended on a Jewish clientele." The report documented persistent sales decreases of 60%–70% in numerous Arab-owned stores, restaurants, and production plants, much larger than Intifada-related sales decreases in the Jewish sector. The industry that has been hurt the most is Jewish Israelis' domestic tourism to Arab cities: "Tourism in Nazareth has decreased 90% in 2001 over 2000," and "occupancy in hotels has been lower than 5%" (Nir 2002, p. B4). Both tourism and hotel occupancy far exceeded those in the parallel non-Arab sectors (due to the general downturn caused by the Intifada), reflecting a decision by many Jewish Israelis to avoid Arab Israeli products.

In searching for a theoretical lens to guide our research on these attitudinal and behavioral changes, we used the animosity model of foreign product purchase (Klein 2002; Klein and Ettenson 1999; Klein, Ettenson, and Morris 1998). Klein and colleagues argue and document that consumer animosity—defined as the remnants of antipathy related to previous or ongoing military, political, or economic events toward current or former enemies—affects consumers' willingness to buy that country's products. Although the animosity model was originally developed to explain intercountry animosity, we test its applicability to intracountry animosity between two ethnic groups.[1]

A key tenet of the animosity model is that people who are angry do not denigrate the quality of a country's products, and previous research has found that animosity does not predict product judgments (Klein 2002; Klein, Ettenson, and Morris 1998). We expect a departure from this aspect of the original model and predict a link between animosity and product judgments. Klein and colleagues' studies examine anger toward Japan (due to World War II and more recent economic issues) among consumers in Nanjing, China, and the United States. Japanese products in these markets tended to be electronics, appliances, and cars. Nanjing consumers kept their animosity toward Japan separate from their assessment of product quality; that is, products were judged positively regardless of animosity levels.[2]

P93

However, imagine a context, such as that of the current study, in which products of an ethnic group are integral to that group's culture and personality. Anger toward this group might translate into a dislike for its customs, and products that represent the culture of that group will be denigrated. Although anger toward Japan might not lead consumers to believe that Japanese cars or stereos are inferior, anger toward Arabs might lead a consumer to denigrate the quality of Arab products and services, if they are perceived as personal and culture specific. Cognitive dissonance theory supports this contention (Festinger 1957). Animosity would produce dissonance and cause people to have a negative attitude toward Arabs but positive beliefs about the way they interact in a service encounter, their clothes, or their food. This dissonance could be reduced by denigrating the quality of Arab products and services that are laden with cultural symbolism. In addition, because this is within-country animosity, there is more contact between the groups, making animosity more personal. In contrast, animosity toward Japan among people in the United States or people in China would likely involve anger toward a more abstract national entity than toward individuals and their habits. We revisit this issue subsequently.

In summary, this article makes the following contributions to the understanding of the role of animosity in consumer judgments: (1) We test the animosity model in an intracountry context; (2) we test a relationship between animosity and product-quality judgments not found in previous research; and (3) we examine whether dogmatism, nationalism, and internationalism serve as antecedents of animosity. In addition, we rule out the possibility that fear of entering Arab neighborhoods explains our findings about the effects of animosity, and we also rule out the possibility that Israeli Jews refrain only from buying from Arab Israelis but do not necessarily refrain from purchasing products produced by Israeli Arabs.

We begin our discussion by identifying potential antecedents of animosity. We then describe animosity's outcomes. After describing the study and its findings, we discuss theoretical and managerial implications.

Literature Review

Antecedents of Animosity

The original animosity model did not include any potential antecedents to consumer animosity, and thus there is little understanding of the personality traits or belief systems that underlie consumer animosity. This is important to marketing managers who face animosity toward their country or religious/ethnic group and want to examine ways to counter the negative impacts of such animosity. We examine three potential antecedents to animosity—dogmatism, nationalism, and internationalism—and justify their selection.

p.94

"Dogmatism" is defined as "the extent to which a person asserts his/her opinion in an unyielding manner" (Bruner and Hensel 1992, p. 194).[3] It attests to the degree of openness/closeness in people's belief systems (Rokeach 1960). Previous research established that dogmatic people tend to be conservative. For example, dogmatism is related to conservative Christian beliefs (Wilson 1985) and to low levels of empathy, defined as affective sensitivity (Carozzi et al. 1995). Furthermore, highly dogmatic people are more orthodox doctrinally (McNeel and Thorsen 1985), more self-righteous (Falbo and Belk 1985), less prone to process new information readily (Flake 1991; Temkin 1987), and more resistant to changing their beliefs (Davies 1993) than are less dogmatic people.

Of particular relevance to this study are findings that people that are higher in dogmatism are less tolerant of other groups and minorities (Mangis 1995). For example, dogmatism predicted Israeli attitudes about the rights of Arab Israelis; specifically, Israelis who were higher in dogmatism were against increased rights for Arabs (Kedem, Bilu, and Cohen 1987). Thus, we expect that dogmatism will predict animosity toward Israeli Arabs.

Balabanis and colleagues (2001, p. 160) assert that "nationalism" refers to a perspective that "encompasses views that one's country is superior and should be dominant (and thus implies a denigration of other nations)" (see also Kosterman and Feshbach 1989). We prefer this definition of nationalism to that of patriotism because nationalism corresponds to an intergroup comparison with other nations, whereas patriotism corresponds with temporal or standard comparisons (Mummendey, Klink, and Brown 2001). Our interest is in intergroup comparisons within Israel, thus making nationalism the focal construct.

Ample evidence exists that nationalism predicts attitudes and behaviors. Nationalistic people are more aggressive and prejudiced toward other nations and ethnic groups than are less nationalistic people (Druckman 1994). Nationalists tend to have stereotypical images of out-groups, have exaggerated national self-images, hold isolationist viewpoints, and support trade protectionism (Sidanius et al. 1997). Furthermore, a history of war and conflict has been found to raise ethnic identification (Hong, Wong, and Liu 2001), which in turn results in negative out-group attitudes (Brown et al. 2001). Thus, we expect that nationalism will predict animosity toward Israelis Arabs.

According to Kosterman and Feshbach (1989, p. 260), "*Internationalism* should not be seen as one pole of a single dimension for which nationalism is at the other pole," because such a view could obscure what might be separate dimensions. Thus, they defined (p. 260) nationalism as "the view that [a country] is superior and should be dominant" and internationalism as "attitudes toward other nations." Balabanis and colleagues (2001, p. 158) assert that high internationalism reflects "positive feelings for other nations and their people and an open-mindedness and acceptance concerning other nations and cultures." Thus, we predict that internationalism will be inversely related to animosity.

Animosity Consequences: Product-Quality Judgments, Willingness to Buy, and Changes in Purchase Behavior

Klein, Ettenson, and Morris (1998) hypothesize that animosity reduces consumers' willingness to buy products from a country toward which animosity is held, and their data support this expectation. However, they provide little in the way of theory-based arguments for this relationship.

Here, we draw on the principle of cognitive consistency, which suggests that consumers value harmony among their thoughts, feelings, and actions and strive to reduce dissonance when such harmony does not exist (Festinger 1957). Moreover, their attitudes should determine actual behavioral responses (Fishbein 1963; Fishbein and Ajzen 1975). According to balance theory (Heider 1958), harmony is achieved in the triad if all three relationships among the elements are positive or if two relationships are negative and one is positive. When all three relationships are negative or when there are two positive relationships and a negative relationship, the triad is imbalanced. This imbalance produces unpleasant tension, and people are motivated to reduce it by changing one of the relationships. The underlying implication in balance theory is that it is easier to change conflicting attitudes than to experience the tension.

Applying these ideas to our context, we posit that consumers who feel heightened animosity due to the current situation will be motivated to adapt their behavioral intentions (willingness to buy) downward, which will then change their purchase behavior accordingly (downward), thus restoring the balance. Similarly, striving to balance attitudes of willingness to buy and product-quality judgments with animosity should lead to a lower product-quality assessment. Thus, an increase in animosity levels will bring about a decrease in the willingness to buy, which will then bring about a decrease in product-quality judgment to maintain the harmony. Thus, an increase in Jewish Israeli animosity will probably affect their cognitions (product quality) and behavior (willingness to buy and changes in purchase behavior).

In summary, cognitive consistency theory provides an argument for a negative link between animosity and willingness to buy, a negative link to product judgments, and a positive link between willingness to buy and changes in actual pur-

P96

chases. In addition, the theory implies a positive relationship between product-quality perceptions and willingness to buy (Figure 1 shows our model; all estimates are in the hypothesized direction). As we discussed previously, the prediction about the negative link between animosity and product judgments is not consistent with previous animosity studies. In previous studies that focused on Japan (Klein 2002; Klein, Ettenson, and Morris 1998), the most visible products on the Chinese and U.S. markets (i.e., electronics and cars) were universal and were distinct from the customs or culture of Japan. Thus, for example, a consumer in Nanjing could hold that Japanese products are of high quality yet not want that product in his or her home because of animosity toward Japan. However, products such as food and services, which are culturally imbedded, are much more difficult to disentangle from the people who produce them. In this study, we focus on such products and services, and therefore we predict that animosity will lead to quality denigration.

P.97

2 請評述以下之文獻:

(1) 文獻內容敘述(20%)

(2) 文獻內容及架構之分析及評述(30%)

Knowledge acquisition for expert systems in accounting and financial problem domains

Abstract

Since the mid-1980s, expert systems have been developed for a variety of problems in accounting and finance. The most commonly cited problems in developing these systems are the unavailability of the experts and knowledge engineers and difficulties with the rule extraction process. Within the field of artificial intelligence, this has been called the 'knowledge acquisition' (KA) problem and has been identified as a major bottleneck in the expert system development process. Recent empirical research reveals that certain KA techniques are significantly more efficient than others in helping to extract certain types of knowledge within specific problem domains. This paper presents a mapping between these empirical studies and a generic taxonomy of expert system problem domains. To accomplish this, we first examine the range of problem domains and suggest a mapping of accounting and finance tasks to a generic problem domain taxonomy. We then identify and describe the most prominent KA techniques employed in developing expert systems in accounting and finance. After examining and summarizing the existing empirical KA work, we conclude by showing how the empirical KA research in the various problem domains can be used to provide guidance to developers of expert systems in the fields of accounting and finance.

Keywords: Expert systems; Accounting expert systems; Finance expert systems; Knowledge acquisition; Problem domain

1. Introduction

Since the mid 1980s Expert Systems have been developed for a number of different accounting and finance applications. These systems have either been developed completely in house, purchased as proprietary software, or developed using an expert system shell. The most commonly cited problems in developing these systems are the unavailability of both the experts and knowledge engineers and difficulties with the rule extraction process [1,2]. This has been called the 'knowledge acquisition' (KA) problem and has been identified as a major bottleneck in the expert system development process [3,4]. Simply stated, the problem is how to efficiently acquire the specific knowledge for a well-defined problem domain from one or more experts and represent it in the appropriate computer format.

Given the 'paradox of expertise', the experts have often proceduralized their knowledge to the point that they have difficulty in explaining exactly what they know and how they know it. However, new empirical research in the field of expert systems reveals that certain KA techniques are significantly more efficient than others in different KA domains and scenarios. Adelman [5] was one of the first to design experiments to objectively compare the effectiveness of different KA techniques. He identified five determinants of the quality of the resulting knowledge base. These are:

1. Domain experts
2. Knowledge engineers
3. Knowledge representation schemes
4. Knowledge elicitation methods
5. Problem domains.

This paper presents a mapping between the body of KA empirical studies and the different problem domains within accounting and finance to guide developers of accounting and finance expert systems in their choice of KA techniques.

2. A generic problem domain taxonomy

Research in the field of knowledge acquisition has

* Corresponding author. Tel.: +1-610-519-6449.
E-mail address: william.wagner@villanova.edu (W.P. Wagner).

0950-7051/02/$ - see front matter
PII: S0950-7051(02)00026-6

Table 1
Generic problem domain taxonomy

Analysis problems
- Classification—categorizing based on observables.
- Diagnosis—inferring system malfunctions from observables.
- Debugging—prescribing remedies for malfunctions.
- Interpretation—inferring situation descriptions from sensor data.

Synthesis problems
- Configuration—configuring collections of objects under constraints in relatively small search spaces.
- Design—configuring collections of objects under constraints in relatively large search spaces.
- Planning—designing actions.
- Scheduling—planning with strong time and/or space constraints.

Problems combining analysis and synthesis
- Command and control—ordering and governing overall system control.
- Instruction—diagnosing, debugging, and repairing student behavior.
- Monitoring—comparing observations to expected outcomes.
- Prediction—inferring likely consequences of given situations.
- Repair—executing plans to administer prescribed remedies.

focused on several dimensions of the problem as determining factors. As previously mentioned, one primary determinant of the KA technique used to develop an expert system is the problem domain. To enhance research in the KA field, generic problem domain taxonomies have been developed that cut across functional areas. The most commonly used taxonomy breaks problems into general categories of analysis, synthesis, and those that combine analysis and synthesis [4,6,7]. This taxonomy is shown as Table 1.

2.1. Mapping of accounting and finance tasks to the generic problem domains

The generic task domains defined in the taxonomy can be mapped into accounting and finance task domains. Such a mapping is presented as Table 2, along with selected examples of accounting and finance expert systems. To do the mapping, we conducted an analysis and survey of accounting and finance expert system ease studies.

A wide variety of problems have been addressed with expert systems with varying levels of success. The most common applications were in the domains of classification and planning with no clear examples of systems in the repair or scheduling domains. These tasks were then placed in the generic taxonomy based upon the generic task descriptions. Due to the complex nature of accounting and finance business functions, some of them may also fall under multiple task domains and the categories are not completely exclusive. For example, depending on the scope of the inputs and the resulting advice from the expert system, a classification expert system might also fall in the categories of debugging or diagnosis.

In addition, the process of mapping specific functions to the higher level categories of analysis, synthesis, and the combination reveals some interesting characteristics of accounting and financial problems. Looking at the specific tasks that fall within the analytic category shows that all of these tasks involve taking a set of data inputs and identifying patterns in them. In contrast, the synthetic problems require that solutions be generated based upon the

Table 2
Accounting and finance task domains and expert system examples

Task domains: Categories	Generic	Accounting and finance	Accounting and finance expert systems
Analysis	Classification	Classify firms into writedown or non-writedown categories	WDXPERT [36]
	Debugging	Evaluating loan losses and proposing solution strategies	ES for evaluating loan losses. [41]
	Diagnosis	Consumer relations analysis	ESCFE [37]
	Interpretation	1. Analyze accounting variances	1. Model variance analysis expert system [39]
		2. Interpret tax consequence of stock redemptions	2. CORPTEX [42]
Synthesis	Configuration	Managing a business loan portfolio.	MARBLE [43]
	Design	1. Design a personal financial plan	1. Planman [44]
		2. Design a compensation system	2. IMIS [45]
	Planning	1. Audit planning	1. Expertest [46]
		2. Investment planning	2. Vanguard online planner [40]; capital investment system [47]
	Scheduling		None found
Combination	Command and control	Estimate control risk	C and L control risk assessor [46]
	Instruction	Training insolvency counselors	PISCES: ([48])
	Monitoring	Fraud detection	Fraud detection system [49]
	Prediction	Risk estimation	APX [50]
	Repair		None found

Table 3
Excerpt from an unstructured interview (adopted from Ref. [2])

I:	What information are you given about airports?
E:	Now, just some rudimentary information comes with it. Common name, latitude and longitude, ah…no information comes with it about the ah…maximum number of airplanes on the ground or the port capability that is at the field. None of that comes along with it
I:	What's the difference between MOG and airport capability?
E:	Ah…MOG maximum on the ground is parking spots… on the ramp. Airport capability is how many passengers and tons of cargo per day it can handle at the facilities
I:	Throughout…ah…throughout as a function of …
E:	It all sorta goes together as throughput. If you've only got…if you can only have ah…if you've only got one parking ramp with the ability to handle 10,000 tons a day, then your…your throughput is gonna be limited by your parking ramp. Or, the problem could be vice versa
I:	Yeah…
E:	So it's a (unintelligible phrase).

more general goals of the system and involve the search of a much larger set of potential solutions. Combinations of the two are typically the most ambitious types of expert systems in that they must perform in-depth analysis of large amounts of diverse input data, identify the problems and causes and design a possible solution. The inherent difficulties of developing systems for these types of problems may be the underlying reason so that few of these types of expert systems have been attempted in the area of accounting and finance [52].

When the task categories presented in Table 2 are combined with the comparative empirical KA studies in Table 4 (which is to be discussed later in detail), these mappings may serve as a guide for which KA techniques might be the most appropriate for the different problem domains within accounting and finance. A general discussion of the most prominent KA techniques is provided in the Section 3.

3. Knowledge acquisition techniques

In the strict sense, knowledge elicitation should be viewed as one phase of the knowledge acquisition process. However, in much of the research knowledge elicitation and knowledge acquisition are used interchangeably. The role that the human knowledge engineer will play in the knowledge acquisition process will vary considerably depending on the particular elicitation technique or method used. In some cases, it may be appropriate for the knowledge engineer to become an apprentice to the expert or participate somehow in the actual problem-solving process. Other times it may be better for the knowledge engineer to conduct an unstructured interview or simply observe the expert perform a given task.

Many different techniques have been developed especially for knowledge engineers in these different situations or have been drawn from existing research in fields such as psychology, and several researchers have attempted to summarize these [2,4,9–11]. Of these techniques, it should not be surprising that a survey [12] found that the most commonly used knowledge elicitation technique was the 'unstructured interview', where the knowledge engineer asks general questions and just hopes for the best. However, each technique requires different abilities from the knowledge engineer, the knowledge source, and allows a different set of knowledge representations to be used.

Although the human-based knowledge acquisition techniques described in the following sections are certainly the most common used today, they are certainly not without their problems. Not only do they require an enormous amount of time and labor on the part of both the knowledge engineer and the domain expert but they also require the knowledge engineer to have an unusually wide variety of interviewing and knowledge representation skills in order for them to be successful. Research on the most common elicitation and representation techniques used by human agents are presented in the following sections.

3.1. Unstructured interviewing techniques

Undoubtedly the most common technique currently used by knowledge engineers [2,12], it is difficult to describe this as a true technique, since as its name implies it is just a wandering conversation between the expert and the knowledge engineer. Table 3 is a transcript of a typical unstructured interview given as an example.

Even though this may be the case, unstructured interviewing still has a valuable place in the knowledge engineer's toolkit since it allows the greatest possible freedom for knowledge engineer and expert alike to explore the topic. Many researchers have documented and described its usage, although researchers in the field usually downplay its value as a real tool. In fact, an anthropological study of knowledge engineering cited the use of the unstructured interview as one of the biggest failings of knowledge engineers who were attempting to develop expert systems [13].

While it is certain that no interview is completely unstructured, different types of unstructured interviews have been suggested by authors of general surveys of knowledge acquisition techniques. In the most extreme case, the knowledge engineer does not have a prepared set of detailed

Table 4
Task types [2]

Familiar task activities	• Unobtrusive observation • Simulated familiar tasks
Interview tasks	• Unstructured • Structured
Constrained professing tasks	• Case-based reasoning • Event recall • Scaling and Sorting tasks • Creative problem solving • Decision analysis

questions and the expert does not have ready replies or information at his/her fingertips. The result is that the interview takes the form of a wandering dialog in which open-ended questions are asked by the knowledge engineer. However, this can be helpful both in acclimating the expert to the ideas of artificial intelligence and in helping the knowledge engineer learn general ideas about the problem domain. It serves the additional purpose of building an essential rapport between the system developer and the human source.

Unfortunately, this has been shown to be a time-consuming and inefficient process [2,14,15] and can offend the expert as being a 'waste of time' [13]. The difficulties of the unstructured interview become apparent when one views a sample from an actual interview and sees how inefficient it can be. (See Table 3).

3.2. Structured interviewing techniques

Recognizing that unstructured interviews are inefficient, researchers in the area of psychotherapy have been developing structured interviewing techniques for many years [16]. Basically, they provided structure by developing a carefully pre-planned series of ordered questions. From this work, psychologists developed other interviewing techniques and tools which were designed to structure the interview process. This work has been generally applied to the knowledge elicitation problem. These techniques can often be applied to situations where the expert is being interviewed while actually performing a task or where the task is simulated or reconstructed by case studies or scenarios or simply from the expert's own past experience.

Elicitation techniques most commonly discussed in the literature include protocol analysis [12,17,18], repertory grids [4], prototyping [19,20], multidimensional scaling [4, 21], cluster analysis [14], event recall [2], discourse analysis [51], and card sorting [15].

Some rudimentary structuring can be given to the interview process by having the expert perform a particular task while the knowledge engineer asks questions freely. The task may be typical of the problem-solving situation which the knowledge engineer wishes to explore or it may be a special case, identified in earlier sessions which the knowledge engineer wishes to use to have the expert refine previously elicited knowledge. The simplest task the knowledge engineer could give the expert could be to prepare a brief lecture designed to lay out the main themes and ideas associated with the particular problem domain. Obviously, this type of task would be more appropriate for early knowledge acquisition sessions whereas the special task would be better for when the knowledge engineer was more familiar with the particular domain [2].

It should be noted that the tasks which are used as the basis for structuring the interview can be either actual tasks or simulated tasks. This method of structuring the interview process by using specific tasks has been termed 'constrained-processing' [2] and the different tasks were grouped as shown in Table 4.

3.3. Protocol analysis

Protocol analysis is one of the most frequently mentioned elicitation techniques in the knowledge acquisition literature. [12] found it to be second only to unstructured interviews in actual usage. Suggested by Newell and Simon [18], subjects are asked to 'think aloud' while solving a problem or making a decision. These verbalizations are usually taped and then transcribed and the transcription is analyzed using a particular coding scheme. The transcript itself is termed a 'protocol' and may be used to refer to a word-for-word record or a summary of the major points. Whatever the form of the protocol, it should enable the knowledge engineer to easily access, index, and code specific pieces of information.

Depending on the problem domain it may be desirable to also generate 'motor' protocols or even 'eye-movement' protocols to more clearly understand an expert's performance of a task [10]. Motor protocols require that the expert's physical movements be closely observed and noted by the knowledge engineer, which may appropriate for acquiring certain types of expertise. At an even more subtle level, noting the movement and visual focus of the eyes of the expert as a task is being performed may reveal something of the sensory experience of the expert as he/she performs the task [10].

However, all protocols can be classified as being either 'concurrent' or 'retrospective' [19]. Concurrent protocols are records of the expert's thought processes at the same time he/she is solving a problem while retrospective protocols are records of the expert's review of his/her verbalizations after the task is completed. These are often used when it is felt that the task of verbalization has interfered with the expert's performance of the actual task [19]. For example, an expert in quality control for autos may be asked to think aloud as he/she goes through the auto inspections process and the resulting recorded protocol might form the basis for the next phase of knowledge acquisition. The transcript may in turn be translated by the knowledge engineer into a more formal protocol which attempts to summarize the major points in a format designed for easy access (e.g. using indexing, notation, or special coding systems).

Once the protocol has been worked into the desired format, the actual analysis of such a protocol by a knowledge engineer can begin. The usual method knowledge engineers use is something called 'process tracing' drawn from cognitive psychology [9]. Analysis involves breaking down the decision rules used by the expert into typical, naturally recurring decision rules. These can then be refined further by either the domain expert or another, external expert before and after they are implemented in the final system.

Protocol analysis has become popular as an elicitation tool because it forces the expert to focus on a specific task or problem without interruptions from the knowledge engineer. It forces the expert to consciously consider the problem-solving process and so may be a source of new self-understanding. It is also very flexible in that many different types of tasks (simulations, special cases, etc.) may serve as a basis for the protocol. Having a record encourages the knowledge engineer to identify specific topics and also missing steps in the process. It has been successfully applied to developing expert systems [2] and early results of comparative experiments show that it is more efficient than unstructured interviewing [15], although the same set of experiments shows clearly that it is less efficient than other non-traditional knowledge acquisition methods such as card-sorting and goal decomposition. Also, on a practical level, protocol analysis requires little equipment or special training for the knowledge engineer.

The main disadvantage to protocol analysis is the very necessity of forcing the expert to verbalize his/her actions. It is often the case that expertise has become so proceduralized that the expert is either unable to express it or is completely unaware of it. This phenomenon is more commonly referred to as the 'paradox of expertise' [2], and is one of the major motivations for research in the field of knowledge acquisition. Not only they may be unaware of their problem-solving methods, but they may actually verbalize them incorrectly and thus introduce error or bias into the resulting system. Especially when special or difficult test cases are used as cues the expert may experience considerable discomfort in trying to verbalize the problem-solving process. Thus the appropriateness of protocol analysis may depend heavily on the type of task being studied and the personality and ability of the expert to be introspective and verbalize thought processes. Protocols can also be very time-consuming to generate and may result in more data than the knowledge engineer can efficiently handle; this is especially true of larger knowledge acquisition tasks.

While protocol analysis involves little interaction between the expert and the knowledge engineer, several elicitation techniques have been suggested which require the knowledge engineer to actively participate in the problem-solving process. These techniques capitalize on the idea that the knowledge engineer must become somewhat of an expert in order to successfully translate the expert's knowledge into a machine representation. Thus the interview may be treated as a tutorial where the expert delivers a lecture which the knowledge engineer may paraphrase or use to solve similar problems [22,23]. The knowledge engineer may become even more actively involved by playing the role of an apprentice or otherwise participating in the expert's problem-solving process [22]. Making the knowledge engineer become like the expert is certainly the most time-consuming approach to knowledge elicitation but ensures the highest quality resulting system. The inherent difficulties of requiring the knowledge engineer to learn all the expertise in order to translate it into a suitable machine representation are what have motivated much of the work at designing expert system shells with more naturalistic interfaces which enable the expert to enter his/her expertise directly into the system.

3.4. Psychological scaling

Other interviewing techniques that have been proposed in the literature have been drawn directly from psychology. These include multidimensional/psychological scaling, network scaling, discourse analysis, cluster analysis, and card sorting. Many of these techniques combine elicitation and structuring aspects and thus are difficult to consider as simply 'elicitation' or 'structuring' techniques. Originally researchers thought that such techniques such as these would be more objective than more traditional interviewing methods [14]. One empirical comparison of elicitation techniques supports this contention somewhat in that it found that non-traditional techniques such as card sorting and goal decomposition performed better than protocol analysis and interviewing [15].

A number of different techniques fall under the heading of what may be called 'psychological scaling' techniques. These include multidimensional scaling (MDS), network scaling, and hierarchical cluster analysis. Generally speaking, experts are asked to rate the similarity of different objects (usually chosen prior) and this rating is portrayed as a distance on a seven-point scale ranging from no similarity to completely similar. The purpose of this is to discover the expert's rank ordering of objects within a problem domain.

Probably the most complicated method in this group is MDS. It is based on the use of the least squares method of fit the elicited data and has been in use for quite a long time [24]. Experts evaluate the similarity of a set of key factors or objects using a numbered scale along different dimensions, generating a grid of numbers. This is supposed to give an overall picture of the problem space. The location of the objects in the different dimensions is then inferred using the least squares method. This requires that both the objects and dimensions be identified beforehand and that they should be representative of the larger problem domain without contradictions. Thus a considerable amount of knowledge is already embedded in the selection of factors or objects and dimensions on which they are to be evaluated, so this technique may more appropriate for structuring previously acquired knowledge.

In this sense it can be considered to be an elicitation tool much as the Analytic Hierarchy Process [25] elicits data about objects on a hierarchy of dimensions. Even though Cullen and Bryman's survey [12] found this to be the least used knowledge elicitation technique, MDS has been used extensively by researchers [14,26] and is well-suited to use in automated knowledge acquisition programs such as AQUINAS, KITTEN, KSSO, PATHFINDER, and PLANET [4].

3.5. Card sorting

Card or concept sorting techniques are also used to help structure an expert's knowledge. As its name implies, the knowledge engineer write the names of previously identified objects, experiences and rules on cards which the expert is asked to sort into groups. The expert describes for the knowledge engineer what each group has in common and the groups can then be organized to form a hierarchy. Like MDS, some empirical research [15] suggests that card sorting may be a more efficient elicitation technique than some of the more traditional techniques such as protocol analysis or interviewing. It has been used with some success to develop applications described in the literature [27,28]. It has also been suggested that it is a tool which could be easily implemented on a computer as an automated knowledge acquisition tool [9].

3.6. Empirical research on knowledge acquisition techniques

Work on the knowledge acquisition problem currently follows along three major, interlocking lines. We describe these as technique-oriented, empirical studies, and conceptual research. As has been noted in the literature [4,29], the primary emphasis to date has been on developing new knowledge acquisition tools and methods. This article focuses on examining the impact of recent empirical work.

A review of the KA literature shows that both conceptual and empirical research has lagged behind technique-oriented research. Experiments and case studies have focused on comparing and evaluating knowledge acquisition techniques. However, the empirical work has suffered from a general lack of control and precision [29].

There have been a few recent efforts to empirically test the usability of different knowledge acquisition tools and techniques. [15] tested the ability of various knowledge elicitation methods to elicit knowledge about classifying different rocks. And [30] compared the relative efficiency of several automated knowledge acquisition tools. Such research is important because it serves to break new ground, but it needs to be conducted in a more systematic and rigorous manner [31].

Previous researchers have recognized the need for sound empirical research to compare the effectiveness and efficiency of KA tools and methods. Fellers [32] concluded that more research was needed to answer the following:

1. Is there one best elicitation technique for knowledge acquisition?
2. If not, what is the best combination of techniques?
3. Which techniques are most suitable under which circumstances?
4. What skills are required in order to utilize each of these techniques?

One KA researcher [19] designed an experiment to test the ability of three different KA methodologies to elicit different types of heuristics. The three methods tested were scenarios, simulated different tasks, and actual familiar tasks. Heuristics were divided into two categories: those that-all subjects identified regardless of knowledge acquisition method and those that only individual subjects identified. These are further broken down as conceptual, operational, and logistical heuristics. Overall, she found a 30% overlap in the heuristics generated by each of the knowledge acquisition methods she tested. Of the heuristics that did not overlap, she identified conceptual, logistical, and operational heuristics that were distinct to each method. But given that the task studied (piloting a boat in a harbor) was operational in nature her results were not surprising.

Adelman [5] varied the domain experts, the elicitation methods, and the knowledge engineers in an attempt to see which if any had the greatest effect on the quality of the knowledge base. Five of the six knowledge engineers had PhD, degrees and one was ABD, but all had extensive training in both top-down and bottom-up elicitation techniques. The relative accuracy of each was compared to a 'golden mean' rule set derived prior to the elicitation sessions. Although a long line of psychological research has been devoted to describing interviewer effects which are analogous to the potential effects of a knowledge engineer [33], no significant effects were observed in this set of experiments. Interestingly, the only significant source of variation came from the domain experts themselves.

The best-known experimental research on KA methods is that of Burton et al. [15]. By varying the different knowledge acquisition techniques among different groups of experts, each of whom was tested for cognitive style, they discovered several specific things. Among their findings was that protocol analysis took the most time and elicited less knowledge than the other three techniques they tested (interviewing, card sorting, and goal decomposition). Not surprisingly, they also found that introverts needed longer interview sessions but generated more knowledge than extroverts. Interestingly, the rarely used techniques of goal decomposition and card sorting proved to be as efficient as

Table 5
Summary of empirical KA research (Adapted from Ref. [29])

Studies	KA techniques	Mod. Vars.	P. Domain	Dep. Vars.	Results
[38]	Interviewing Inductive learning	Not considered	Diagnosis	Percentage of correct diagnosis generated	Inductive learning performed better than interviewing
[17]	ID3 Induction and interviews	Not considered	Diagnosis	Comparison to known cases	Induction performed much better than interviewing
[53]	Interviews protocol analysis expert walk-throughs	Human vs. reconstructed knowledge sources	Interpretation	KE's opinion of the quality of K acquired	Protocol analysis had limited usefulness for certain types of K
[15]	Interviews protocol analysis goal decomposition card-sorting multi-dimensional scaling	Introvert vs. extrovert; cognitive styles	Classification	Time taken for capturing K; time for coding into rules; number of rule clauses; completeness of rule set	Protocol an. takes longer and yields less K; Introverts need longer interviews but generate more K than extroverts
[35]	Interviewing and protocol analysis	Domain complexity	Planning	Efficiency and quality of K as measured by number of nodes and arcs and their accuracy	Interviewing is more efficient and accurate for simple cases but protocol is more efficient for complex cases
[8]	Structured interviews protocol analysis card sorting laddered grids	Expert vs. non-expert; two classification domains	Classification	Efficiency of process	Protocol analysis performed poorly in classification domain; card sorting and grids performed better than interviewing; external validation of experts important
[8,19]	Scenarios, simulated different tasks, and actual familiar tasks	Not considered	Command and control	Overlap of heuristics	Found 30% overlap between heuristics elicited by the 3 different methods
[5]	top-down vs. bottom-up interviewing	Knowledge engineer, and domain expert	command and control	Accuracy of elicited rules as compared to 'golden mean' set	Found no sig. Variation except for that due to domain expert

the more common interviewing technique and more efficient than the commonly used protocol analysis.

This experiment was criticized somewhat for its lack of rigor [29,31]. One measure of technique efficiency was the time it took to code the transcripts into pseudo-rules while the number of rules or clauses was taken as a measure of acquired knowledge. Coding time does not fully account for temporal differences among KA methods and there are also serious drawbacks to using the number of coded rules as a measure [29,34]. The results may also have been confounded by unmeasured differences among the experts and the knowledge engineers.

These various experimental studies are symptomatic of a recognized need empirical investigation of KA phenomena. The small number of such studies is at least in part, indicative of the difficulty in conducting them. The few pioneering studies are typified by confusing terminology, conflicting operationalizations, and the proliferation of ad hoc taxonomies. In addition, results are conflicting and no clear pattern has emerged. There are problems controlling for effects of moderator variables and in operationalizing the measurement of dependent variables. In light of these problems, [29] concluded that empirical KA work should concentrate on case studies rather than experiments, at least in the near term. A strategy for addressing some of these experimental obstacles has also been proposed [31]. A fairly comprehensive summary of empirical KA research (Adapted from Ref. [29]) is presented as Table 5.

4. Conclusions

The application of empirical KA research to the problem of choosing an appropriate KA technique for developing an expert system application in the fields of accounting and finance suggests several directions. First, if one is going to be working in an analytic problem domain such as classifying asset writedowns. KA techniques that provide a high degree of structure to the interviewing process seem to work best. Protocol analysis, though fairly commonly used, is relatively inefficient for analytic problems while the most popular technique of using an unstructured interviewing is one of the least efficient and least satisfying from the standpoint of the expert. So it may be worth exploring some of the non-traditional KA techniques when working on these type applications.

For the more difficult synthetic and combination problem domains the evidence is not as clear. However, the Holsapple and Raj study [35] seems to indicate that problem complexity may be one determinant of the appropriate KA technique to choose. So if one were to develop an expert system for estimating control risk then we might suppose that protocol analysis might be more efficient than interviewing. The fact that interviewing is more efficient for simple domains may imply that it is best used for initial KA sessions, when the problem complexity is not yet developed clearly.

For those studies that did consider the effect of moderator variables, it seems clear that no matter what type of problem domain, developers of expert systems in the fields of accounting and finance should consider their potential impact. The impact of the cognitive style of the expert, domain complexity, along with other attributes of the domain expert all seem to be important factors in the quality of an expert system regardless of the problem domain. It is hoped that further research will clarify some of these issues with respect to the effect of moderator variables and problem domains.

References

[1] O. Owrang, F. Grupe, Database tools to acquire knowledge for rule-based expert systems, Information and Software Technology 39 (1997) 607–616.

[2] R. Hoffman, The problem of extracting the knowledge of experts from the perspective of experimental psychology, AI Magazine 8 (2) (1987) 53–67.

[3] B. Le Roux, Knowledge acquisition as a constructive process: a methodological issue, Decision Support Systems 18 (1) (1996) 33–43.

[4] J. Boose, A survey of knowledge acquisition techniques and tools, Knowledge Acquisition 1 (1) (1989) 3–38.

[5] L. Adelman, Measurement issues in knowledge engineering, IEEE Transactions on Systems, Man, Cybernetics 19 (1989) 483–488.

[6] T. Shaft, I. Vessey, The relevance of application domain knowledge: the case of computer program comprehension, Information Systems Research 6 (3) (1995) 286–299.

[7] W. Clancey, Heuristic classification, Artificial Intelligence 27 (1985) 298–350.

[8] A. Burton, R. Schweickert, N. Taylor, E. Corlet, N. Shadbolt, A. Hedgecock, Comparing knowledge elicitation techniques; a case study, Artificial Intelligence 1 (14) (1990) 245–254.

[9] K. McGraw, K. Harbison-Briggs, Knowledge Acquisition: Principles and Guidelines, Prentice-Hall, Englewood Cliffs, NJ, 1989.

[10] G. Tuthill, Knowledge Engineering: Concepts and Practices for Knowledge-Based Systems, Blue Ridge, PA, TAB Books, 1990.

[11] J. Kim, J. Courtney, A survey of knowledge acquisition techniques and their relevance to managerial problem domains, Decision Support Systems 4 (3) (1988) 269–285.

[12] J. Cullen, A. Bryman, The knowledge acquisition bottleneck: time for a reassessment? Expert Systems 5 (3) (1988) 216–224.

[13] D. Forsythe, J. Buchanon, Knowledge engineer as anthropologist, IEEE Transactions on Systems, Man and Cybernetics (1989) 3.

[14] N. Cooke, J. McDonald, The applications of psychological scaling techniques to knowledge elicitation for knowledge-based systems, International Journal of Man Machine Studies 26 (1987) 81–92.

[15] A. Burton, N. Shadbolt, A. Hedgecock, G. Rugg, A formal evaluation of knowledge elicitation techniques for expert systems: domain 1, Proceedings of the First European Workshop on Knowledge Acquisition for Knowledge-Based Systems, September, Reading University, 1987.

[16] R. Merton, M. Fiske, P. Kendall, The Focused Interview, The Free Press, Glencoe, IL, 1956.

[17] A. Hart, Experience in the use of an inductive system in knowledge engineering, in: M.A. Bramer (Ed.), Research and Development in Expert Systems, Cambridge University Press, Cambridge, 1985, pp. 117–126.

[18] A. Newell, H. Simon, Human Problem Solving, Prentice-Hall, Englewood Cliffs, NJ, 1972.
[19] M. Grabowski, Knowledge Acquisition Methodologies: Survey and Empirical Assessment, Proceedings of ICIS, 1988.
[20] D. Waterman, A Guide to Expert Systems, Addison-Wesley, Reading, MA, 1986.
[21] L. Elliot, Analogical problem solving and expert systems, IEEE Expert 1 (2) (1986) 17–28.
[22] M. Welbank, Knowledge Acquisition: A Survey and British Telecom Experience, Proceedings of the First European Workshop on Knowledge Acquisition for Knowledge-Based Systems, Reading University, UK, 1987, pp. C6.1–C6.9.
[23] L. Johnson, N. Johnson, Knowledge elicitation involving teachback interviewing, in: A. Kidd (Ed.), Knowledge Elicitation for Expert Systems: A Practical Handbook, 1987.
[24] J.B. Kruskal, Multidimensional scaling and other methods for discovering structure, in: Enslein, Ralston, Wilf (Eds.), Statistical Methods for Digital Computers, Wiley, New York, 1977.
[25] T. Saaty, The Analytical Hierarchy Process, McGraw-Hill, New York, 1981.
[26] K. Butler, J. Carter, The use of psychologic tools for knowledge acquisition: a case study, in: W.A. Gale (Ed.), Artificial Intelligence and Statistics, Addison-Wesley, Menlo Park, 1986, pp. 295–320.
[27] M. Chi, P. Feltovich, R. Glaser, Categorization and representation of physics problems by experts and novices, Cognitive Science 5 (1981) 121–152.
[28] J. Gammack, R. Young, Psychological techniques for eliciting expert knowledge, in: M. Brammer (Ed.), Research and Development in Expert Systems, Cambridge University Press, London, 1985, pp. 105–112.
[29] J. Dhaliwal, I. Benbasat, A framework for the comparative evaluation of knowledge acquisition tools and techniques, Knowledge Acquisition 2 (2) (1990) 145–166.
[30] C.M. Kitto, Progress in automated acquisition tools: how close are we to replacing the knowledge engineer? in: Banff (Ed.), Proceedings of the Third Knowledge Acquisition for Knowledge-Based Systems Workshop, 1988, pp. 14.1–14.13.
[31] C. Holsapple, V. Raj, W. Wagner, Knowledge acquisition: recent theoretic and empirical developments, in: C. Holsapple, A. Whinston (Eds.), Recent Developments in Decision Support Systems, Springer, Berlin, 1993.
[32] J. Fellers, Key factors in knowledge acquisition, Computer Personnel 11 (1) (1987) 10–24.
[33] K. Hammond, Subject and sampling—a note, Psychological Bulletin 45 (1948) 530–533.
[34] C. Holsapple, A. Whinston, Manager's Guide to Expert Systems, Dow Jones-Irwin, Homewood, IL, 1986.
[35] C. Holsapple, V. Raj, An exploratory study of two KA methods, Expert Systems 11 (2) (1994) 77–87.
[36] S. Ragothaman, B. Naik, Using Rule Induction for Expert System Development: The Case of Asset Writedowns, International Journal of Intelligent Systems in Accounting, Finance and Management 3 (3) (1994) 187–203.
[37] A. Houghton, M. Gardner, Expert Systems at the Customer Interface, International Journal of Intelligent Systems in Accounting, Finance and Management 1 (4) (1992) 289–296.
[38] R. Michalski, R. Chilausky, Knowledge acquisition by encoding expert rules versus computer induction from examples—a case study involving soybean pathology, International Journal of Man-Machine Studies 12 (1980) 63–87.
[39] A.S. Hollander, Using Rules to Model Variance Analysis, Expert Systems Review 2 (3) (1990) 46–52.
[40] Vanguard Web Site (http://majestic.vanguard.com/guide/da), May 2000.
[41] G. Ribar, Expert systems technology at peat Marwick Main, Expert Systems Review for Business and Accounting 1 (1) (1987) 1–5.
[42] H. Schwartz, Expert systems in accounting, CPA Journal Online April (1989) .http://www.nysscpa.org/cpajournal.
[43] M. Shaw, J. Gentry, Using an expert system with inductive learning to evaluate business loans, Financial Management (1988).
[44] C. Brown, N. Nielson, M. Phillips, Expert systems for personal financial planning, Journal of Financial Planning 3 (3) (1990) 137–143.
[45] B. David, N. U, Automates payroll with AI system, Datamation 42 (17) (1996) 129–132.
[46] D. Murphy, C. Brown, The uses of advanced information technology in audit planning, International Journal of Intelligent Systems in Accounting, Finance and Management 1 (3) (1992) 187–193.
[47] C. Brown, M. Phillips, Expert systems for management accountants, Management Accounting 71 (7) (1990) 18–23.
[48] N. Doherty, K. Pond, An expert system solution to mortgage arrears problems, The Services Industry Journal 15 (2) (1995) 267–288.
[49] K. Lecot, Using expert systems in banking: the case of fraud detection and prevention, Expert Systems Review for Business and Accounting 1 (3) (1988) 17–20.
[50] A. Bharadwaj, V. Karan, R. Mahapatra, U. Murthy, A. Vinze, APX: an integrated knowledge-based system to support audit planning, International Journal of Intelligent Systems in Accounting, Finance, Management 3 (3) (1994) 149–164.
[51] N.J. Belkin, H.M. Brooks, P.J. Daniels, Knowledge Elicitation Using Discourse Analysis, International Journal of Man–Machine Studies 27(2) (1987) 127–144.
[52] S. Eom, A survey of operational expert systems in business (1980–1993), Interfaces 26 (5) (1996) 50–70.
[53] W.F. Messier, J.V. Hansen, Expert Systems in Auditing: The state of the art, Auditing: A Journal of Practice and Theory 7 (1) (1987) 94–105.

科目：管理文獻評述

編號：321

考生注意：
1.依次序作答，只要標明題號，不必抄題。
2.答案必須寫在答案卷上，否則不予計分。
3.限用藍、黑色筆作答；試題須隨卷繳回。

3. Based on the ideas revealed from the below article, if you are asked to propose a research plan, what will you do? Please develop a research proposal, including, but not limited to, motivations, objectives, framework, research design and methods, and expected contributions. (50%)

SUSTAINING BREAKTHROUGH INNOVATION

More than 100 practitioners of innovation assembled in Cambridge, Massachusetts in November 2008 to discuss the importance of innovation as a management discipline and share leading-edge practices critical to its future (*1*). In this rich environment, I had the opportunity to deliver a keynote address, in which I positioned innovation as an emerging business function, just as marketing emerged as such in the mid-1950s.

Innovation as a distinct and separate management discipline is moving from an academic concept into real-world practice at major corporations. The focus of our research program at Rensselaer Polytechnic Institute has been to understand how firms can build a sustainable breakthrough innovation (BI) capability, knowing that back in 2001 the average life expectancy of these initiatives was only about four years. I am pleased to report that the life expectancy is increasing as companies embark upon a more systematic approach to innovation, become more strategic about it and learn how to orchestrate their innovation capacity.

Research Insights

Our insights are based on the cumulative learning of 28 companies directly participating in this research and over 100 companies involved in validating the research and tools through the Industrial Research Institute (IRI). We learned, first, that breakthrough innovation can be characterized by high uncertainty, which is no surprise. But our initial insights led us to identify four distinct categories of uncertainty (market, technical, resource, and organizational uncertainty), which fuel a host of management challenges. In addition to technical and market issues that project teams typically face, resource (accessing talent, partners and money) and organization (establishing and maintaining organizational legitimacy) uncertainties must also be managed. In fact, the latter two are the ones most likely to impede project success.

Along with the importance of reducing uncertainty across multiple dimensions, the management challenges we observed include: 1) capturing breakthroughs (clarifying and articulating breakthrough concepts); 2) managing projects that experience the chaos of multidimensional uncertainty and ambiguity; 3) learning about markets that do not yet exist; 4) developing new business models; 5) acquiring resources (both funding as well as appropriate competencies); 6) transitioning projects successfully from R&D to operating units, particularly when the fit is not clear; and 7) harnessing the contributions of individuals. We noted in our first book (*2*) that an innovation focal point, or group, is needed to screen and evaluate new business opportunities, create the right environment for these projects to be successful, and manage an open innovation model with multiple interfaces.

We ultimately defined a mature breakthrough innovation capability as an embedded system for initiating, supporting and sustaining BI activities. The problem is, most firms are designed to meet the objectives of operational excellence: customer intimacy, responsiveness to market needs, and strict cost containment. Such a system may tend to snuff out breakthroughs. When leadership clarifies that BI is important to the firm, we observe that one or two design elements of the management system may be put in place, but rarely, if ever, are all of them installed simultaneously. In order to have a sustained innovation management function, all the elements of a management system must be attuned to that objective, *and* they must work in parallel with the management

Gina Colarelli O'Connor is associate professor of marketing and academic director of the Radical Innovation Research Program at Rensselaer Polytechnic Institute's Lally School of Management and Technology in Troy, New York. She conducts research on how firms build management systems for innovation, and teaches corporate entrepreneurship, business implications of emerging technologies, and marketing of high technology products to M.B.A. and executive M.B.A. students. She is academic director of RPI's Professional Development Program on Innovation and Corporate Entrepreneurship. She and the research team at RPI published their second book based on the Radical Innovation Research Program in 2008, titled Grabbing Lightning: Building a Capability for Breakthrough Innovation *(Jossey-Bass). She received her Ph.D. in marketing and corporate strategy from New York University.* **oconng@rpi.edu.**

system for operational excellence as well. We identify the elements of the BI management system as follows:

- Defining mandate and scope.
- Ensuring a supportive leadership/culture.
- Establishing an appropriate organizational structure and managing interfaces.
- Developing governance and decision-making mechanisms.
- Introducing learning-based processes and tools.
- Identifying skill sets and developing talent.
- Developing innovation metrics and rewards.

Indeed, these management system elements are the same for both an innovation as well as an operational excellence agenda. But they certainly play out differently for each.

Beyond the elements of a management system, we observed in our research that three distinct competencies are necessary for breakthrough innovation to occur with any degree of frequency in a firm. These three distinct BI competencies we call *Discovery*, *Incubation* and *Acceleration* (or DIA for short). The DIA Model is not linear due to the iterative learning that takes place when faced with high levels of uncertainty. Most important, each of these three competencies has its own set of leadership/culture, structure, governance, processes/tools, skills, and metrics to consider.

1. The DIA Competencies and System

Discovery is the creation and identification of opportunities that may have major impact in the marketplace, either through the delivery of new performance benefits or greatly improved performance. Discovery does not equal invention or R&D. It incorporates scientific work, for sure; but in addition to scientific work, there must be a focus on opportunity generation, elaboration and development, as well as the skills to articulate those opportunities in a manner that is of strategic importance to the company and can be understood by senior non-technical leaders.

Incubation is a competency of experimentation. We call it the long and winding road (agree that we lack originality!). It requires experimenting with the technology and business concepts/models simultaneously to arrive at a demonstrated model of a new business that brings breakthrough value to the market and consequently the company. The focus is on learning and redirecting based on the creation and pursuit of options, with allowances for failures. Enriching and extending internal and external networks is critical to enlarging the scope of the company's knowledge base and commercial potential. The *Incubation* staff coaches, brokers and nurtures these fledgling opportunities, as well as thins and enriches the portfolio. In Incubation, many avenues are explored initially, yet few enter *Acceleration*.

> **Amazing progress has been made among the companies we studied.**

Acceleration is a competency of developing critical mass. The objective is to scale these nascent businesses so they can compete with mature businesses in their ultimate home (in an existing business unit or as a new strategic business unit or division) for resources and attention. The objectives are to build critical mass of sales and operational infrastructure, establish market presence, develop the management team, and prepare the new business to blend into the fabric of the rest of the organization. The goal is to develop predictable sales forecasts, achieve acceptable yields and demonstrate a path to profitability.

The *DIA System* is defined as the set of activities that manage the links and interfaces within DIA and oversee its health in terms of the BI mandate, its perceived role in the company and its portfolio of businesses. System imbalances can occur when D, I and/or A are not managed well, when one of the three competencies is not as healthy as the others, or when the three do not interact. For example, we observed several companies that had an excellent *Discovery* capability, but lacked *Incubation*. Nascent opportunities moved straight to *Acceleration*, and were typically under-leveraged, meaning that one product may have been introduced, but the larger prize of an entire business platform was never executed. Thus D and A are healthy, but I(*ncubation*) is not. We call this situation one of "Big ideas...incrementally executed."

A second example of a system imbalance occurs when *Incubation* work that was done by one team is ignored by the handoff team receiving the opportunity for *Acceleration*. This failure to leverage learning signifies a weakness in the interfaces of the DIA competencies. There are many other sources of DIA system imbalance. You can probably think of those that plague your own organization.

All this is to say that someone needs to be concerned with the health of the DIA system. This involves monitoring and managing system imbalances in conjunction with organizational capacity, attending to portfolio health and diversity, assembling and reassembling project teams, providing an enabling project infrastructure, removing barriers, brokering external and internal liaisons, guiding strategic alignment activities, providing help for project resource acquisition, educating others about the role of BI compared with the rest of the company, pacing projects, and overseeing transitions from D to I to A to landing zones.

We observed these roles emerging in some of the companies we studied, although none were formalized as such. Firms still seem shy about formalizing such responsibilities, largely because they're not yet aware of the need to do so. Still, the fact that there is now a senior VP of Strategic Growth at Corning, or a vice president of Emerging Business Opportunities at IBM, or a portfolio overseer at DuPont are all signs of things moving in the right direction.

2. BI Portfolio Considerations and System Metrics

In managing the BI portfolio, degree and nature of diversification, portfolio churn rate, portfolio size, portfolio pacing objectives across DIA, and cross-portfolio management are key considerations. Diversification is interesting: typically one diversifies one's portfolio to hedge against risk. But in the world of BI, all projects are risky. Thus, we observed companies diversifying along other dimensions: technology domains/competencies, business/market domains, time horizon, and organizational fit.

Churn rate should be monitored within D, I and A. If it is too high, this could be attributable to poor ideas, transitioning too early, execution problems, inappropriate team composition, inadequate resources, or lack of fit with intended strategy. Portfolio size objectives should be set based on how many projects/platforms the firm envisions as strategically key for its future health. For portfolio pacing, it is easy to get caught up with I and A, and forget to replenish the pipeline. Finally, in cross-portfolio management, the focus is on checking for synergistic effects, convergence, redundancies, and spillover between the innovation efforts.

Measuring the BI system's success cannot be done with a financial returns measure alone. There are too many unknowns, lengthy timeframes involved, and too many benefits that accrue to the company that are not directly traceable to dollars. Metrics require a focus on the health/activity of the portfolio, interface management, impact of the BI portfolio on the market, and, finally, impact on the company.

Portfolio health relates to number of new ideas, projects started and projects transitioned between the stages, as well as synergies, diversity and pacing of projects compared with objectives. Interface management is about the smoothness of handoffs and communication flows within the system. Market impact is about game-changing attributes and gaining external recognition for technology expertise, as well as enrolling others such as partners in the opportunity. Impact on the company takes many forms, such as number of projects transitioned out into businesses, dollar impact of those projects, learning spillover to other areas of high uncertainty, development of talent (serial entrepreneurs), and increased robustness of new ideas.

Questions for R&D Leaders

Amazing progress has been made among the companies we studied in the time that we have been observing them. Other firms as well have contacted us and requested talks, workshops and interaction as they invest in building this capability. Innovation is becoming more embedded throughout the organization, with BI's role better understood within the innovation system. Along with this comes increased confidence, a focus on portfolios of BI opportunities and a willingness to stay the course, even when times are not great. Although new roles are emerging in leadership, exploratory marketing and new-business-creation coaching capacities, career paths for new business creation remain a concern.

As a management discipline, innovation should not be viewed as a program; it is a constant, as is any other function in the company, such as engineering, marketing, finance, and manufacturing. Despite this progress, our companies admit, this discipline is still very new. Most feel they are on the right track, yet wish they had more direction.

Consequently, companies continue to struggle with developing the right management system, understanding relative priorities and finding the talent to manage innovation and growth in complex corporate environments. We believe that the elements of the system are becoming clearer and new roles are emerging that require training and skills beyond those of traditional marketing, strategy, technology management, or business development.

As R&D leaders, I ask you the following questions:

- What are you doing to implement a management system for innovation in your company? Which elements do you need to address now? In the future?
- How well established is *Discovery* in your company? What about *Incubation*?? *Acceleration*??
- Where will you find people with the skills to manage the complexity of innovation? Has your company developed a career path for innovation experts?
- Do you have the right climate drivers to enable strategic innovation? How will you measure success?

Best wishes on your journey! ⊛

References

1. The forum was the 5th annual Entrepreneurship, Innovation and Growth (EIG) conference, titled "Innovation...Cultivating a Discipline," co-sponsored by Product Genesis, the Radical Innovation Group and the National Venture Capital Association's Corporate Venture Group. See *http://www.nvca.org/EIG08.html* for conference details.
2. Leifer, R. et al. 2000. *Radical Innovation: How Mature Firms Can Outsmart Upstarts*. Boston: Harvard Business School Press.