Judgmental Errors, Interactive Norms, and the Difficulty of Detecting Strategic Surprises

Joseph Lampel • Zur Shapira
University of Nottingham, Business School, Nottingham, United Kingdom, NG8 1AE
Stern School of Business, New York University, 44 W. 4th Street, New York, New York 10012
joseph.lampel@nottingham.ac.uk • zshapira@stern.nyu.edu

Abstract

The concept of strategic surprise has been used by writers on military strategy to explore the successful amplification of resources during battle. This paper applies the same concept to subcontracting relationships and develops a cognitive framework to explain the phenomenon of strategic surprises, using buyer-supplier relations as an example. We first examine the factors that produce vulnerability to strategic surprise in cooperative situations. Then, we explore the reasons why firms are caught by surprise in spite of their vigilance. We present a model representing false alarms and strategic surprises as judgmental errors. We argue that judgmental errors cause misinterpretation of evidence and a consequent sense of false security. Interactive norms, which exist in some industries and may be taken as proxies for enduring relationships, may increase the likelihood of misjudgment and strategic surprises.

(*Strategic Surprises*; *Noncontractual Relationships*; *Heuristics*; *Norms*; *Trust*)

Introduction

Can a rational actor ever be surprised by the decisions of his or her counterpart? Rational, vigilant actors should not, in principle, be surprised by actions that are foreseeable in light of what the actors know, even when such actions are highly improbable. History nevertheless records numerous instances of strategic surprises: The Japanese surprised the Americans in Pearl Harbor, Hitler surprised Stalin by invading the Soviet Union, and the oil-consuming nations were surprised by OPEC's decision to double oil prices in October 1973.

Less dramatic, but perhaps not less significant, are cases of commercial firms in contractor or subcontractor relationship who are caught by surprise when their partners switch from cooperative to predatory behavior (Provan and Skinner 1989, Sako 1992). Here, actors find themselves in a position in which their partner is taking advantage of the very process that perpetuates the relationship. A situation in which both sides share control now becomes asymmetric. The possibility of such a sharp turn is always present when firms rely on relationships that are noncontractual. The virtue of such relationships is that they combine efficiency and flexibility in a way that is harder to attain through contracts. However, the very informality of noncontractual relationships creates more risk.

Trying to tap the advantages of noncontractual relationships while simultaneously managing the risk of strategic surprises is difficult. The first line of defense is reliance on mechanisms that reduce the risk of opportunism by deterring it. The use of reputation as a way of increasing the costs of opportunism represents such a mechanism. Deterrence, however, is not always reliable. Opportunists may calculate (or miscalculate) the costs of deterrence and find the advantages of predatory behavior to be greater than potential costs. Thus, a rational actor will also practice vigilance, monitoring partners' behavior for signs of opportunistic intent and acting to prevent strategic surprises before they happen. Vigilance can reduce vulnerability to strategic surprises, but for vigilance to be effective actors must correctly estimate the risk of strategic surprise in the light of new information. Such information can come from various sources such as the behavior of the actors themselves.

The problem is that strategic surprises are rare events. In any given relationship they are likely to happen only once—after all, few actors would continue a noncontractual relationship with an overt opportunist. To estimate the risk of a strategic surprise it is therefore necessary to focus on behavior that is likely to signal opportunistic intent. In this paper, we focus on the process of information gathering as well as on the interpretation of this information. We argue that the process may undermine vigilance and increase susceptibility to strategic surprise due to the combined effect of basing judgment on a small

1047-7039/01/1205/0599/$05.00
1526-5455 electronic ISSN

sample of events and on cooperative norms, to infer intent.

First, actors are likely to focus on behavior that deviates sharply from established patterns of interaction. Such departures are perceived as anomalous, attracting attention and raising questions about the opportunistic intent behind such behavior. They push the firm to make a decision: Should it terminate the relationship, or should it regard the unexpected behavior as random or accidental? The problem is that firms have a limited amount of information at their disposal. Deviations from an established pattern of cooperative behavior do not occur frequently, so firms must rely on the first few instances of unexpected behavior. They cannot afford to take a wait-and-see attitude, because the consequences of inaction may be very serious. Thus, they tend to draw inferences from a small sample, which often leads to estimation errors (Tversky and Kahneman 1971). The estimation error may lead firms to terminate the relationship when this is not warranted—a mistake which eliminates the risk of strategic surprise, but which also means the loss of a valuable relationship—or it may lead to underestimation of the risk of opportunism and thus make the firm more rather than less vulnerable to strategic surprise. Those mistakes are defined, respectively, as Type I and Type II errors. In the next section, we present a model that highlights their role in strategic surprises.

Second, in most industries firms do not rely solely on the past behavior of their partners when making inferences about the threat of strategic surprise. They also tend to rely on behavioral norms that constrain opportunism (Swedberg 1994). Norms shape generalized expectations about what is likely to happen in buyer-supplier relations; these expectations are then modified in light of what is known about specific partners. A group of firms may espouse adherence to norms, but this does not guarantee that they embrace and feel bound by norms to the same degree. Thus, a key problem facing firms attempting to evaluate the risk of strategic surprise is deciding how strongly their partners are bound by norms. Because there is also only a small amount of data available on how strongly partners adhere to norms, inference is likely to amplify errors, leading firms to misjudge the trustworthiness of their partners.

There are diverging views on the effect of trust on supplier-buyer relations, ranging from the highly calculative (Williamson 1993) to trust-based relations (Sabel 1993). One way to sort out these competing, though somewhat overlapping, views is to consider the impact of surprise on these relationships. Surprise is a low-probability event, and as such it challenges the parties involved as to whether they should switch from the "business as normal" mode to an adversarial mode. Our paper addresses this question and contrasts the economics perspective with the sociological view. While both approaches have merit, we embark on a new road that combines judgmental issues with a sociological perspective on norms in an attempt to analyze the strategic surprise question.

Our paper is structured as follows. In the second section, we define strategic surprises and present a model to describe the judgmental process firms go through. We present the work of Lorenz (1988) on contracting-subcontracting relationships to illustrate many of the issues that are discussed. In the third section we analyze the role of false alarms in strategic surprises. In the fourth section, we examine the role of norms and their joint effect with judgmental errors on strategic surprises. We conclude with a discussion of a risk-taking perspective on firm behavior in the buyer-supplier context, and of cooperative relationships in general.

Judgment and Strategic Surprises: A Model

In discussing strategic surprises, we must distinguish between surprises that come from outside a relationship and those that result from the interaction itself. Incidents such as Three Mile Island, the Exxon Valdez oil spill, and Bhopal (Bowman and Kunreuther 1988) constitute a major strategic surprise for the companies involved, but they are not considered strategic surprises in the sense discussed here. In this paper, we look at strategic surprises that take place in interactions between interdependent actors. Interactive situations, whether in war or commerce, are jointly developed and perpetuated by the actors. Intent plays a key role in such situations. In military confrontation, the intent driving the presumed interaction is negative, and the process is therefore one of escalating hostility. In business relationships, intent is assumed to be positive, and the process expected to be one of escalating commitment to a joint course of action. In both instances, assumptions about intent are crucial to decision making, but it is also these assumptions that put actors at risk.

Strategic surprises comprise an overt action coupled with a surprise that increases the action's effectiveness. A strategic surprise occurs when an actor switches from behavior that reinforces cooperation to one that expresses the intent of extracting concessions. The surprise itself results from two elements: speed, which allows little time for warning and even less for defensive measures, and contrast between assumptions held by the victim before the action and the intent revealed by the action. The greater this contrast, the greater the surprise, and the more likely it is to achieve its goal.

For victims of strategic surprises, a shock derives not from the switch from cooperative to opportunistic behavior—this possibility has always been there—but from the failure to perceive this switch when it comes. Information alone is not sufficient; the process of interpretation must also be sound because it is the interpretation of a certain piece of information that leads decision makers to classify it as a signal rather than noise. In so doing they may commit either a Type I or a Type II error. For example, rejecting a candidate or failing to proceed with a project may mean missing valuable opportunities, but hiring or proceeding with a project may turn out to be a costly mistake (Einhorn and Hogarth 1978). Decision makers usually collect information to reduce the error of making the wrong choice, but at some point information gathering must come to an end and the decision must be made. At this point, decision makers must decide which error they want to minimize: the error of proceeding with a potentially costly and unrewarding opportunity, or the error of missing a valuable opportunity. The dilemma facing the decision maker is that at the time of decision it is impossible to reduce the probability of one error without increasing the probability of the other. This is the same dilemma which confronts decision makers who are concerned about strategic surprise. To expand on this point further, we turn to a model that has its roots in signal detection theory and has been used more recently by Shapira (1995) to analyze risk taking in organizational settings. The model is described in Figure 1.

In a selection decision (such as project or employee selection), a critical value x_c is determined such that if the ex ante evaluation of the project yields a value x whereby $x > x_c$ the project is accepted. The project is rejected if $x < x_c$. After the decision is made the project's performance is measured against a critical value y_c. If the realized value is greater than y_c the project is considered a success. The degree to which the ex post performance of the project is predicted by the ex ante evaluation and is described by the correlation coefficient r_{xy}. The higher the correlation, the thinner is the ellipsoid in Figure 1; the extreme case (when $r_{xy} = 1$) is represented by a straight line.

Figure 1 **The Relation Between Severity of Signal, Alarm Threshold and Strategic Surprise**

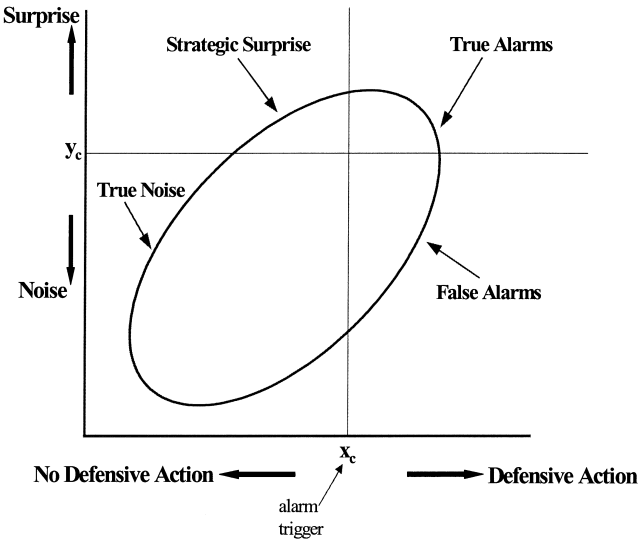

In most cases ($0 < r_{xy} < 1$), we encounter four possible action-outcome combinations as a result of an accept/reject decision. If a project is accepted (i.e., $x \geq x_c$), it can eventually succeed ($y \geq y_c$) or fail ($y < y_c$). The former action-outcome combination is called a "true positive" while the latter is called a "false positive." If, on the other hand, a project is rejected (i.e., $x \leq x_c$), there are also two possibilities: If the project eventually becomes a success (for example, in another firm), it is labeled a "false negative," and if it eventually fails in another trial it is called a "true negative."

The problem facing a decision maker in a situation where a partner may switch from cooperative to predatory behavior is structurally identical to the model described above. Applying the model to strategic surprises, we provide the following description of the four quadrants: (1) True alarms (true positive)—Alarms which lead to precautions that are justified by subsequent events; (2) False alarms (false positive)—Alarms which lead to precautions that are not justified by subsequent events; (3) Strategic surprise (false negative)—No alarms followed by events which show that precautions should have been taken; and finally (4) True noise (true negative)—No alarms and no events to suggest that precautions should have been taken. The horizontal axis in Figure 1 depicts the severity of the signal, namely whether one, two, or more delays occurred and whether those involved a potentially low or high cost. The value x_c can be defined as the trigger value, that is, if the severity of the signal reaches this level defensive action is warranted. If the signal severity is below x_c, it is not sufficient to trigger defensive action. The vertical axis depicts the degree of surprise varying with the degree of speed and the contrast between previously held assumptions and revealed action. Events above y_C are defined as strategic surprises and below it as noise, that is, the other party did not intend to surprise the first party even though some of its activities might have been interpreted as such. Note also that both x_C and y_C can vary when considered from different vantage points (for discussion see Shapira 1995, pp. 106–119).

Looking at Figure 1, it is clear that one can make two errors in a strategic surprise decision, either by responding to a false alarm or by failing to interpret the signal and falling victim to a strategic surprise. These two errors correspond to Type I and Type II errors in statistical inference, respectively, and as stated earlier, they are interrelated. For a given r_{xy} one cannot reduce the probability of one error without increasing the other. The only way to reduce the two errors simultaneously is to increase r_{xy}. To do that one needs to collect more information to validate the model and increase its predictive power. Collecting more data allows us to learn more about the other parties and can potentially reduce strategic surprises. However, r_{xy} can never reach the ultimate value of 1.0 in such situations. Thus, it is a necessary but not sufficient condition for reducing strategic surprises. As will become clear in the examples to be presented next, in examining the possibility of a strategic surprise there exists, in addition to the information layer, a layer of interpretation of the other party's intentions, and it is this layer that makes reasoning about strategic surprises a difficult task.

The two errors correspond to two distinct but related dilemmas facing decision makers when it comes to dealing with strategic surprises. The first is the cost of responding to false alarms, and the second the problem of detecting strategic surprises before they occur. Two historical examples, more than thirty years apart, illustrate these dilemmas.

In her study of the Japanese attack on Pearl Harbor in 1941, Wohlstetter (1962, p. 387) concludes that: "Never before have we had so complete an intelligence picture of the enemy." Japanese codes had been broken, British intelligence provided information from its own monitoring activity, and on-the-spot reporting from diplomats and journalists was highly accurate. In spite of the constant flow of information which indicated that Japanese attack was likely if not imminent, American forces in Pearl Harbor were caught entirely by surprise. The surprise is consistent with our model in two respects. First, information indicating a Japanese attack led to false alarms which ultimately undermined vigilance, and second, assumptions about the low likelihood of a surprise attack led American intelligence to discount information which pointed in that direction.

On three separate occasions information about Japanese intentions led to a state of alert being declared in Pearl Harbor. The first alert was declared on June 17, 1940, but it was directed mainly against the possibility of sabotage by Japanese agents and sympathizers. The second alert was declared on July 25, 1941. This alert was total. It assumed the possibility of all-out attack by Japanese air and naval forces. The third alert was on October 16, 1941. It not only included a warning about the possibility of Japanese attack, it also directed American commanders to deploy their forces in readiness to repel such an attack (Wohlstetter 1962, Prange 1982). These alerts, however, were seen as costly and disruptive false alarms when an attack did not materialize. By the time the third alert was issued, the tendency to discount warnings and to relax vigilance after a short duration was strong.

Pearl Harbor is also a failure of interpretation in the face of powerful evidence. The possibility of a Japanese attack on Pearl Harbor had long been part of American strategic thinking. American war games in and around Hawaii had been based on just such a contingency since 1933. In April and July of 1941 two separate reports were submitted which forecast the Japanese attack in some detail. The second, by Colonel William E. Farthing, was so prescient as to be declared by the investigative report convened after the Japanese attack to be "prophetic in its accuracy and uncanny in its analysis of the enemy's intention" (Prange 1982, p. 188). Notwithstanding all this, the American military discounted information that the Japanese would launch such an attack because they believed the risks involved (for the Japanese) to be too great. As Prange puts it "the Japanese plan appeared fantastic, an inadmissible risk, almost suicidal" (p. 188). Even when U.S. ambassador to Japan, Joseph C. Grew, informed Washington in January 1941 that the Peruvian minister had learned "from many sources, including a Japanese source, that in the event of war breaking out between the United States and Japan, the Japanese intended to make a surprise attack on Pearl Harbor," the response was largely dismissive. Naval Intelligence passed the information to Admiral Kimmel, the commander of the Pacific Fleet at Pearl Harbor, with a note that read, "Naval Intelligence places no credence in these rumors" (Prange 1982, p. 33).

In 1973 Israel confronted a situation that bears considerable similarity to the one that faced the United States 32 years earlier. In April 1973, about six months before the Yom Kippur War, the Egyptians engaged in military movements that were interpreted by Israel's Defense Minister Moshe Dayan, and Chief of Staff of the Israeli Defense Forces (IDF) General David Elazar, as a preparation for war. They decided to respond by a partial mobilization of Israel's reserve forces. They did so despite an estimate of a low likelihood of war by the Chief of Intelligence of the IDF, Major General Eli Zeira. Consequently, the Egyptians (who indeed planned to attack Israel) called off their attack plans. When the Egyptians and Syrians repeated their maneuvers again in October 1973, Israeli decision makers, motivated in part by the high cost of April's mobilization, responded in a conser-

vative manner (Zeira 1993). The events of April 1973 were eventually classified by the Israeli government as a false alarm. Not unlike Pearl Harbor, Israel had ample information to suggest an imminent Egyptian and Syrian attack. For instance, on October 1 a lower-echelon intelligence officer at the IDF's Southern Command wrote a memo to the chief of intelligence alerting him to the possibility of renewal of war, but this statement was deleted from the Chief Intelligence report. On the same day, Syrian troops were redeployed from the Jordanian border to the Golan front, but this move was interpreted as a gesture of goodwill toward Jordan and a defensive measure (Handel 1976). Further, on October 4 the evacuation of the Soviet advisers and technicians started, but this didn't prevent the chief of intelligence of the IDF from putting his estimate of a war at a very low probability on October 5, just 25 hours before the war started. It appears that the high cost of the April mobilisation coupled with the (ex-post) correct estimate of the chief of intelligence led to the Israeli reluctance to mobilize reserve forces in October. It has proven to be a false negative, or a strategic surprise. The Egyptians and Syrians surprised Israel and attacked it when the IDF was not ready to respond. This surprise emanated primarily from the failure of the IDF intelligence branch to interpret the intentions of their enemies.

When dealing with strategic surprises such as Pearl Harbor or the Yom Kippur war, a temptation exists to argue the case of information gap. Such an approach implies that if American or Israeli intelligence had been truly excellent, it would have surmised the date of the attack and possibly even the target. Failure to predict strategic surprise is a failure of information gathering which can be prevented by improving information collection. The difficulty with this prescription is that there is always a gap between the information available and the information needed for completely accurate prediction. This gap persists as along as the evidence is incomplete and is susceptible to contradictory interpretation. Closing it completely calls for information gathering that is difficult, if not impossible, to attain. As Kahn (1965) points out,

> Intelligence officers could perhaps have foreseen the attack if the United States, years before had insinuated spies into high-level Japanese military and naval circles, flown regular aerial reconnaissance of the Japanese navy, put intercept units aboard ships sailing close to Japan to pick up naval messages that a greatly expanded codebreaking unit might have cracked, or recruited a network or marine observers to report on ship movements (p. 144).

To do this would not only call for extraordinary foresight—in effect fixing on an event as probable long before it is even contemplated seriously—but it also calls for major commitment of resources. This points to another dilemma posed by strategic surprises. Reducing the risk of strategic surprise, either by taking precautionary defensive measures or by investing in a comprehensive intelligence system, can be very costly, and the costs escalate as one tries to eliminate the risk entirely. As the leaders of the Soviet Union discovered before the regime collapsed, insuring against the probability of a surprise nuclear attack can impose crippling costs and, paradoxically, damage beyond rescue the very system one is trying to protect (Ansoff 1975, Kylen 1985). This holds true to an even greater extent for firms that have limited resources at their disposal. No matter how much such a firm may wish to insure against the consequences of a subcontractor defaulting on its obligations, there are clear limits to what a firm can do without imposing unacceptable costs on itself in the process.

Balancing Dependence Against the Risk of Surprise: An Example

The limits that firms face when it comes to reducing the risks of strategic surprises are especially evident when they enter into economic exchange, which is regulated by noncontractual relationships rather than by binding contracts. Johnston and Lawrence (1988) cite a list of industries—automobiles, textiles, publishing, movies—in which relationships are displacing contracts as the main conduit for transactions.

Relationships in these industries are not explicitly defined when initiated, but rather are built over time through a stream of orders. With orders come both professional and social contacts which produce mutual understanding and expectations about what the relationship entails (Macaulay 1963). These understandings and expectations, in turn, have important implications for the decision making of the firms involved. Firms are more likely to invest in specialized machinery and production processes if they can rely on future orders from certain customers, and they tend to avoid creating asset specificity if they are uncertain of the level of demand they can expect. Similarly, buyers who require special inputs to launch a new innovation may be reluctant to develop dependence on suppliers who are in exclusive possession of this special capability. For example, Polaroid relied on Kodak for film during its early days. However, when it launched its revolutionary SX-70 camera in 1973, it considered this dependence risky, and moved to build its own internal film-making capability (Berg and Merry 1984).

Such a process may lead to asymmetry in power as the

firm which undertakes specialized investments may be decreasing its relative power. What makes this dilemma pernicious is the fact that it does not come down to a single decision, but emerges gradually over time. More dependence increases both benefits and risks. Because it does so gradually, firms may underestimate the risks, and thereby open themselves to strategic surprises.

This is in effect what happened to TCI Manufacturing Ltd. The Canadian firm had a long-term relationship with Power Computing Corporation of Texas, who produced Mackintosh clones. TCI was its sole supplier of computer cases and power-supply systems. The relationship was governed by a standard purchasing agreement with a thirty-day termination clause. In September 1997, Apple purchased Power's core assets and its license to manufacture clones for $100 million dollars. Shortly thereafter Power announced that it was no longer going to make clones and cancelled all standing orders with TCI. The decision forced TCI to close most of its operations, and lay off most of its staff. Patrick Jabal, TCI's president, accused Power Corporation of "destroying" his company (Saunders 1997). Dave Allen, TCI's director of engineering and a founding partner, expressed anger about the lack of advance warning from Power: "If they'd given us a heads-up earlier in the year that things were a little shaky with Apple, we wouldn't have put ourselves in this situation," he said. "We're not really a part of the Mac community per se, and we had no idea there was even the possibility of them being shut down like this (Schick 1998). Mark Dietz, TCI's attorney, went further, accusing Power Computing of intentionally committing fraud. "We believe they knew in the Spring of 1996 that their licensing talks with Apple were going poorly," he says. "Based upon that knowledge, and the fact that TCI and Power were so closely aligned, they had a duty to tell us what they knew about the prospect of losing the Apple license. And by instead placing orders that took all of our capacity, and that we had to turn down other business. . . that was a misrepresentation to TCI regarding the future of the company" (Schick 1998).

TCI's experience is the specter that haunts all firms in contracting-subcontracting relationships. Lucrative relationships often engender dependence, and dependence increases exposure to opportunistic behavior. Lorenz (1988) studied how small and medium-size French firms dealt with this dilemma. In particular, he looked at 10 firms, located around Lyons, that produce machine tools, textile machinery, packaging machinery, mining equipment, and industrial filters. The products of these firms are large, complex, and customized, and have a high unit price. Production is to order, and the basic outlines of demand are difficult to predict beyond a horizon of about six months to a year.

In 1982, the machinery sector in this region experienced a sharp recession. The recession came at a time when computer numerical control machine tools offered dramatic improvements in productivity. Previously producers would amortize investment in new capacity against expected orders, but demand uncertainty made such forward investment too risky. Suppliers were better placed to bear such risk because they could aggregate demand from multiple clients, but they were less likely to reap the benefits of the investment which accrued to the final producer. Entering into a relationship in which producers outsourced the production of key components while suppliers invested in the new technology made sense, but it also came with certain risks. The producer ran the risk of late delivery or poor-quality parts, and the supplier ran the risk that after it made the costly investment the producer would fail to place sufficient orders to justify the added capacity.

The obvious way of reducing such risks is to stipulate them contractually. The problem with this remedy is not only that it is costly, as Williamson (1985) points out, but that it may also foster risk aversion. An excessive preoccupation with the possibility of opportunistic behavior inevitably leads to a proliferation of hypothetical contingencies under which opportunistic behavior could be advantageous (Ghoshal and Moran 1996). This produces a heightened sense of risk and a defensive posture that may undermine the commitment to outsource. It is for this reason that it is essential in noncontractual relationships for the two parties to deal with risks informally, often by making commitments that are intended to build confidence. For example, in the case that Lorenz examined producers agreed to buy at least 10% of the subcontractor's output but no more than 15%. Less than 10% would undermine the rationale for specialized investment. More than 15% would make the subcontractor dangerously dependent on the market fortunes of the producer. The subcontractor in turn undertook a different set of commitments. In addition to investment in new technologies, the subcontractor undertook to be price competitive relative to other suppliers and to deliver quality components on time. In return for this commitment, the producers informally gave subcontractors the following crucial guarantee. As Lorenz describes it:

> It [was] understood that subcontractors will not be instantaneously dropped by the clients if a differential in terms of price, quality, or delivery time emerges with respect to competitor subcontractors. Rather, a system of advanced warnings is used. A reasonable amount of time is allowed the subcontractors to match the competition (1993, p. 316).

Note the ambiguity of the terms used. No explicit criteria of price, quality, or "lateness" of delivery are stipulated. Consider what is a "reasonable" amount of time allowed before a subcontractor is dropped—a hard-pressed subcontractor and a producer with demanding clients are likely to have diverging interpretations of the term. Nevertheless, to preserve the flexibility of the relationship it is necessary to have a measure of constructive ambiguity. But this constructive ambiguity also opens the way for strategic surprises. Confronted with a request for postponement of delivery, firms will therefore be unsure of how this ought to be interpreted. As Lorenz observes:

> In such cases, how is the firm to interpret a request from the subcontractor for an extension on delivery? Is it a legitimate request due to an unanticipated machine breakdown or some other unfortunate difficulty? Or is it opportunistic, the work having been taken on in the full knowledge that the capacity to complete the order on time was lacking? (1988, p. 204).

Given the specialized nature of production, switching to another subcontractor is a very costly option, and the firm feels that an extension has to be granted. Warnings may be issued, and these warnings may have their intended effect. But even when performance returns to normal, larger questions still remain. Does the lapse represent a clear signal that the subcontractor is behaving with opportunistic intent?

One may argue that actors should focus mainly on the costs and benefits of opportunism at the time that the decision is made, rather than on the past behavior of partners. This is both rational and prudent. Deception can always masquerade as cooperation, taking the past into account is therefore risky. Experimental evidence shows, however, that actors use past behavior as evidence of their partners' character and willingness to cooperate even though it may be risky (Camerer 1997, Dawes and Thaler 1988).

Fischhoff (1982) notes that the interpretation of a sequence of events is subject to the hindsight bias since some events are scrutinized more closely than others. Lorenz's respondents singled out a fall in orders or late delivery as significant when it came to extrapolating the possibility of major adverse discontinuities or, in our language, "strategic surprises." The question is, how late does a late delivery have to be before it deserves close attention? The problem exists when an event has been classified outside what is acceptable, but is not classified as a trigger event that requires action. That is, an event may still fall into an area between the region where deliveries are normal, and the one where they are sufficiently late to require action. This region is ambiguous because it is not stipulated in advance either explicitly or informally, and because partners set the rules and then deal with the apparent violations as they occur. Such a process makes judgment more difficult and may make the partners more vulnerable to strategic surprises.

False Alarms and Strategic Surprises

During the Cold War, each of the superpowers faced the possibility that her rival would launch, without warning, a full-scale nuclear missile attack. Immense early warning systems were built to detect the launching of such missiles. Initially, more information reduced superpower vulnerability to surprise attacks. However, as the systems became more sophisticated, each side faced a new dilemma: Various innocent activities could be interpreted by an overly sensitive system as the initial stages of an all-out attack. To reduce the risk of launching a false preemptive attack, an alert system was put in place which decreased the sensitivity of the system to incoming information (Kahn 1965, Issacs and Downing 1998).

Suppose that firms develop dedicated early warning systems, and that they have three states of alert. The first state, denoted *green,* represents business as usual. The second state, *yellow,* is characterized by increased vigilance by certain managers directly involved in sales or purchasing. The third state of alert, *red*, denotes a high-level attention on the part of managers, including meetings to decide what actions are called for. These alerts are similar to the areas marked in Figure 1, where green equals true noise, red equals true alarm, and yellow marks both false alarms and strategic surprise. Thus, it is the yellow-alert stage that is marked by ambiguity and the decision maker has to resolve it by acting one way or the other. Turning back to the firms in Lorenz's study, how should a subcontractor interpret a drop in orders from an established customer? Should it adopt a green, yellow, or red-alert status as a result of this fluctuation in orders?

The answer depends on where one draws the lines between the three alert states. Let us assume that a subcontractor in Lorenz's study emulates a Cold War superpower. She sets the yellow and red alerts at relatively low levels during their initial dealings with an unfamiliar client. If the client's behavior remains below the yellow-alert level, there is an inevitable tendency to increase trust and raise the threshold. If a client violates the threshold, this may result in less trust, the supplier becomes more vigilant, and spends more time interpreting cues which previously would have been ignored.

It is also possible, however, for the supplier to think that it is the alert which is unreliable rather than the client, and therefore decide to raise the threshold level. First,

lowering the alert threshold raises costs to managers and organizations. A yellow alert forces key managers to rearrange schedules and priorities, while a red alert may force organizations into costly preemptive actions. Further, if the decision makers are confident that the level at which they set their alert is an accurate representation of client reliability, they will decrease their reliance on their partners or rupture the relationship entirely. Second, if they believe that the violation of the yellow-alert threshold is accidental and does not reflect the trustworthiness and reliability of their partners, they will adjust the threshold upward, thereby increasing their vulnerability to strategic surprise.

False alarms therefore play a major role in the way partners estimate the probability of strategic surprise. The model in Figure 1 depicts the ideal case, in which data are available on all four quadrants. The case of strategic surprise is not as clear. Because strategic surprises are rare events and false alarms constitute a small sample, it is more difficult to use the model in the technical sense than in the case of employee selection or project selection (Shapira 1995). This difficulty is related to the interpretation of the *intentions* of the other party. The ambiguity about these intentions is not usually present in the case of employee selection or project selection. Further, since the costs of false alarms may be prohibitive, actors put a lot of weight on the probability of false alarms when setting the alarm threshold. As noted earlier, decision makers also overweigh recent false alarms in estimating the probability of strategic surprise. Therefore, a costly false alarm may sway some firms to raise the trigger threshold. In terms of the model in Figure 1, this represents moving the alarm criterion to x_{c1} on the right, reducing the probability of false alarm but increasing the probability of strategic surprise (see Figure 2).

The example of the Yom Kippur War described earlier in this paper fits the idea of an increase in the level of the alarm trigger. The change in the level of alarm trigger is an updating procedure. Whatever the prior estimate of a strategic surprise had been, the fact that the Israeli response ended in what was classified as a false alarm made the decision makers focus on it in updating the probability of a future strategic surprise. In so doing, they might have exhibited the base-rate bias (Tversky and Kahneman 1982) rather than applying Bayes' formula properly. Of course, without knowing what the prior was, it is difficult to evaluate the updating process. However, given that there were very few relevant data points to begin with and that the cost of the defensive action (that was taken in April) was very high, it appears that in updating, excessive weight was put on the prior false alarm.

PROPOSITION 1. *In noncontractual relationships, the estimation of the probability of a strategic surprise is highly sensitive to false alarms. This sensitivity arises from the necessity of inferring the probability of strategic surprise from a small sample of available data on false alarms. In addition, when recent false alarms are associated with high costs, there is a tendency to raise the alarm threshold, thereby increasing the probability of strategic surprise.*

Figure 2 Raising the Threshold for the Alarm Trigger

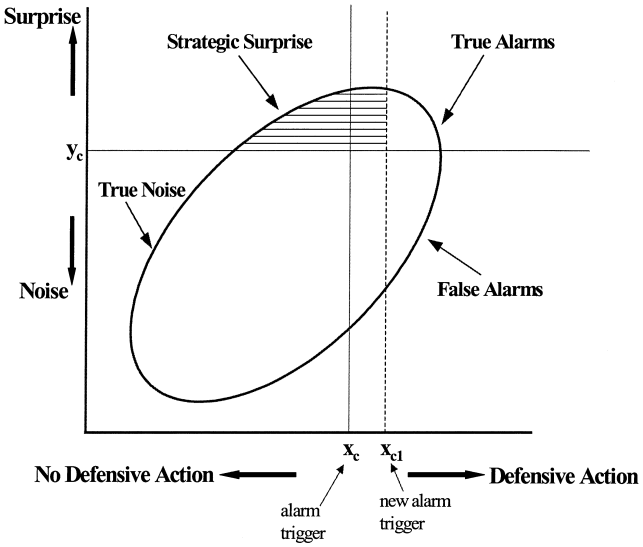

Norms and Strategic Surprises

Arrow (1992) criticizes the definition of rationality in economic theory, arguing that the rationality of actors cannot be divorced from the context in which it is exercised. As he puts it: ". . .rationality is not a property of the individual alone, . . . rather, it gathers not only its force but also its very meaning from the social context in which it is embedded" (p. 63). In other words, inferences about individuals interact with inferences about what the situation entails. Actors observe what others do, and filter these observations through assumptions about the basic character of the situation itself.

These assumptions are often explicit: Competition and cooperation among firms in an industry are not possible without a variety of explicit norms and sanctions that are rooted in the legal and regulatory structure of each economy. The assumptions, however, are often also informal and tacit. Industries are social systems where repeated interactions give rise to informal norms (White 1986). These informal norms may cover areas such as avoidance of price rivalry, mutual forbearance of entry into each other's market, turning a blind eye to a certain amount of

product imitation (Hay and Morris 1991, p. 149–180), or signaling credible commitments (Frank 1987). Such informal norms rely on interfirm relations, and may not get translated into formal norms. This may be the result of the absence of a central legal regime with the authority to enact laws to regulate relationships, a situation that prevails in global markets. On the other hand, informal norms may not become formal because a central legal regime regards such a move as detrimental to public interest—a view that is behind legislating the prohibition of price collusion.

In dependency-producing interactions, repeated interaction can change the basic character of the relationship from a causal arm's-length relationship to one of partnership. (Stinchcombe 1986). Lorenz speaks of the development of "moral contracts" in some of the cases that he observed. Such contracts, as he puts it, involve the following:

> . . .in exchange for improved performance by the subcontractor on quality and delivery, the client firm will make every effort to guarantee a level of work; furthermore, any adaptations to price, quantity and delivery are to be made in a non-opportunistic way by both sides, with full disclosure of the relevant information. (1988, p. 206).

The two sides do not attempt to insert the understanding that underpins the moral contract precisely because this would render the contract almost useless. As one of Lorenz's interviewees put it:

> You can put all the rules you want in place, all the most sophisticated ways of doing things in terms of information processing, follow-up procedures, and quality control. If, at a point in time, men who come into contact and make a commitment to each other act like machines and do this in an administrative manner, you will always end up with failure (1993, p. 317).

Norms and Conventions

For the most part, the understandings contained in Lorenz's moral contracts depend for their effectiveness on relatively ambiguous norms. These norms are different from overt norms or conventions which are relatively unambiguous. For example, in his review of the literature on norms and conventions, Young (1996) discusses the emergence of left-hand and right-hand driving. Before such conventions came into being, there was uncertainty as to whether horse-drawn carriages would drive on the right or the left, but there was no ambiguity about the structure of the choices. Once the conventions were established, drivers expected others to drive on the right or on the left, depending on the country. In short, there may be uncertainty about which conventions will prevail, but once they do there is no ambiguity about how to follow them.

This is not necessarily the case when it comes to norms which govern cooperative relationships. These are structurally more complex then the norms discussed above, and hence their violations are susceptible to multiple interpretations. This interpretative flexibility is both an advantage and a danger. A delivery that is late by a day, or a batch that has several more defective parts than usual, can be explained away. As the French manager quoted above (originally in Lorenz 1993) suggests, one needs to have a very mechanical notion of relationships to regard this as a serious breach of understanding. The problem arises when we enter a "gray" (or yellow in our earlier parlance) zone where delays and defect rates give cause for concern. It is at this point that managers often begin to suspect that their partners are only "paying lip service" to the agreement and may not be truly bound by norms.

The main difference between conventions of the kind that govern driving, dressing, or dining, and norms that govern cooperative behavior in business, is that the former do not depend on motivation, whereas the latter do. It is not important why people abide by conventions when it comes to driving on the left or right, but it is important why they make deliveries on time. Mere conformity does not suffice here precisely because departure from norms is to some extent permitted. Because there is an inevitable gap between following norms and embracing them, it is not safe to assume that a partner will adhere to norms without also evaluating the extent to which he or she embrace them. The former may be a matter of convenience or expediency, unless it is strongly supported by the latter.

Following Norms vs. Embracing Norms

Evaluating the degree to which partners embrace norms as opposed to simply following them is not always easy. In some instances, there are ample opportunities to observe norm-relevant behavior and then use these data to estimate how strongly bound by norms the partners are. In other cases, such as those of strategic surprises, making an accurate estimate is difficult if not impossible. What decision makers tend to do instead is to rely on the tendency of norms governing one kind of behavior to be related to norms governing other types of behaviors. Opportunists by nature are generally less likely to embrace norms, so their pattern of norm following in one context is likely to be correlated with such patterns in other domains. This means that norms can act as a mechanism that signals linkages between one type of behavior and others. For example, a subcontractor who is willing to make last-minute changes to an order without extra charge is likely to embrace norms that constrain taking advantage in other situations.

Norms can therefore be regarded as social and cognitive constructs which link different populations of events

(Coleman 1990). What makes norms interesting in the present context is that they can link common events with those which are uncommon. By observing the frequency of actions in one population of events, a subcontractor can make inferences about the frequency of other events. The stronger the correlation between the two populations, the stronger is the inference that a partner embraces norms as opposed to merely conforming. It is worth noting however that this inference is valid only if the correlation between the populations is relatively high.

The Role of Norms in Estimating the Probability of Strategic Surprises

In the absence of norms there is a higher tendency to lower alarm thresholds, thereby committing Type I errors. Arguably, this is one reason why cooperative norms come into being. Unfortunately, however, we claim that once norms have emerged, they are likely to increase Type II error and thus to increase vulnerability to strategic surprises. The reason, as noted earlier, is that conformity to norms does not automatically mean embracing norms. The relationship is probabilistic and hence subject to judgmental error. The error is due to the difficulties of estimating both how strongly a particular partner adheres to norms, and how powerful the norms are in general.

In any industry where norms are in place, norm espousal usually precedes norm adherence. It is obviously cheaper to talk about norms than to put them into action (Farrell 1987). Such "cheap talk" is often viewed with distrust by economists, but it may play an indispensable function in the evolution of norms. Talking about norms makes it easier to behave in accordance with them. In a sense, norms describe cumulative learning, and as such they provide a substitute for judgment.

The idea that norms represent common wisdom acquired through history is akin to Bowman's (1963) analysis of managerial judgment. He pointed out that ". . .decision rules derived from management's own average behavior might yield better results than the aggregate behavior itself " (1963, p. 316). In our terms, norms can be conceived as average behavior resulting form long-term interactions that at times can provide a better guide for behavior than judgment based on a few recent events. This argument is also echoed in Sunstein's (1996) analysis, in which he claims that economic activity is better described by norms than by preferences and judgment. As he puts it ". . .norms often outstrip the beliefs and receive a kind of moral grounding that is not simply reducible to an instrumental judgment about likely risks" (1996, p. 930). The moral grounding is reinforced by the social nature of the interaction (see also Durkheim 1947). Lorenz (1993) found that the emergence of norms in the population of firms he observed was accompanied by the use of a "language of friendship." Such language, he suggests, became important in situations which raised questions about the trustworthiness of partners. As he puts it: "By using [the language of friendship], clients were conveying to their subcontractors that, when in doubt, they should act as if their actions were guided by the norms of friendship" (1993, p. 320).

Acting "as if" norms are in force is easier if there is sufficient history to suggest that norm espousal is strongly correlated with norm adherence. In industries where norms have only been in existence for a short time, belief in norms receives less support from historical evidence. It follows that decision makers are more likely to see behavior that violates norms as evidence of opportunistic intent than would be the case in industries where there is a longer history of norm espousal and norm conformity.

PROPOSITION 2. *In industries where norms have developed over a* short *period of time, firms are more likely to regard behavior that is contrary to norms and expectations as valid indicators of opportunistic intentions and thus are more prone to lower the alarm thresholds than are firms in industries where norms have developed over a long period of time.*

In terms of the model described in Figure 1, with minimal reliance on norms, firms lower the threshold of setting the alarm trigger, that is, moving the alarm criterion to x_{c2} on the left, thereby increasing the probability of false alarms but decreasing the probability of strategic surprises (see Figure 3).

In contrast, in industries with longer histories of norm espousal and norm conformity the opposite problem exists. Because there is more evidence to support the belief that norms are in force, there is also a greater tendency

Figure 3 Lowering the Threshold for the Alarm Trigger

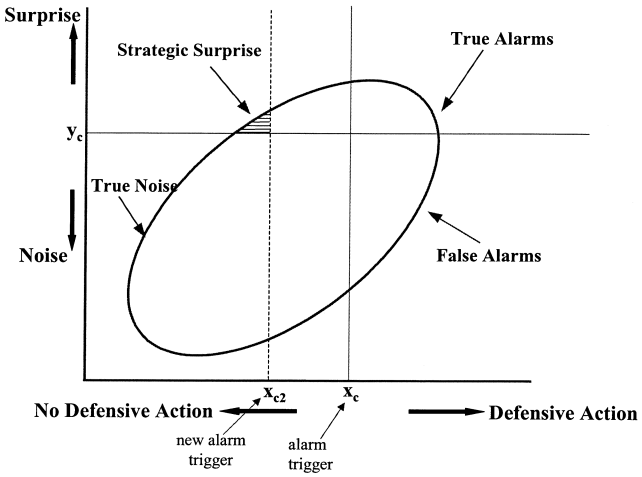

to conclude that norm following is a strong index of norm embracing. Firms are therefore likely to give norms greater weight when confronted with behavior that is unexpected. This means that they are more likely to discount evidence of behavior that violates norms and maintain the relationship.

PROPOSITION 3. *In industries with a long history of interactive norms, firms place more weight on norms when interpreting their partners' intentions and less weight on recent evidence of behavior. This tendency leads to downplaying the importance of recent evidence on false alarms. If in addition the cost of recent false alarms has been high, firms may raise the level of the alarm trigger, thereby reducing the probability of false alarms, but also making themselves more vulnerable to strategic surprises.*

The scenario discussed in Proposition 3 leads to a situation similar to the one described in Figure 4.

As noted earlier, shifting the level of the alarm trigger can be considered as a process of updating the prior probability of a strategic surprise. In the case described in Proposition 2 excessive weight is put on recent evidence, while in the scenario described in Proposition 3 most weight is put on the prior without updating. Due to small samples it is difficult to consider the case as a proper application of a normative model such as Bayes' formula. Norms, however, affect the updating procedure differentially in the two cases.

Figure 4 **Raising the Threshold for the Alarm Trigger—Joint Effect of Errors and Norms**

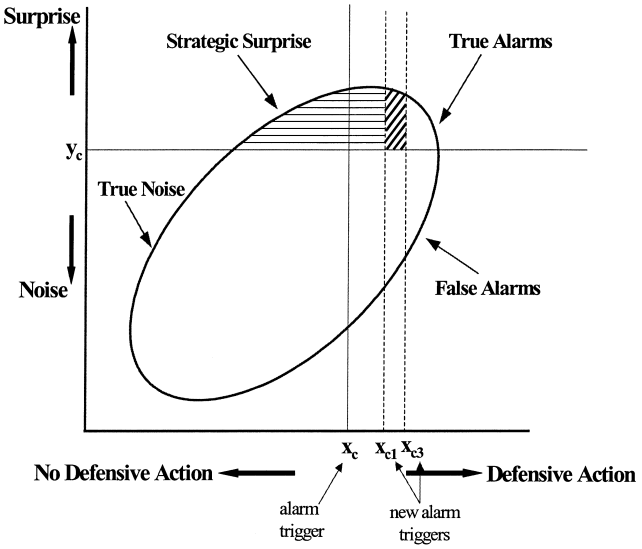

Discussion

Many researchers have looked at buyer-seller relations from a game-theoretic perspective (see, for example, Heide and Miner 1992). Others, such as Levinthal and Fichman (1988) analyzed the relationships between auditors and their clients using the event-history technique to measure the duration of such relationships. Our analysis is based on a somewhat different perspective. We propose a cognitive framework suggesting that managers tend to interpret variations in their environment in a nonoptimal manner from a statistical point of view by focusing on small samples and being prone to judgmental errors such as the misinterpretation of the other party's intentions.

The paradox of having to live with events which are potentially damaging, and yet too costly to prevent, exacerbates this tendency. Managers are forced to pay attention to cues that signal the onset of a strategic surprise. Behavioral decision theory suggests that judgmental biases are amplified when there is pressure to arrive quickly at conclusions without sufficient data (Payne et al. 1993). This increases the susceptibility of firms to strategic surprises. In the case of relationships that call for cooperation, this susceptibility is further exacerbated by norms that emerge during interaction.

We highlighted the issue of norms in this paper because it has been dealt with differently by economists and sociologists. Norms in cooperative relationships are partly situational, emerging as a result of repeated interaction between actors (cf. Heide and Miner 1992). The issue of trust, which has attracted so much discussion and debate, is ultimately concerned with the degree to which norms influence behavior when there are no incentives or sanctions to act as constraints. Economists have argued that no rational actor should trust without clear evidence of incentives or sanctions to back promises, thus promoting a calculative view of trust (Dasgupta 1988, Dixit and Nalebuff 1991, Williamson 1993). Sociologists, on the other hand, argue that normative structures are binding because they are linked to actors' definition of self and values (Lorenz 1993, Sabel 1993). One way to deal with this debate is by treating trust as essentially an evaluation of risk (see Chiles and McMackin 1996, Williamson 1993).

We have made a similar argument by noting that context influences judgment. One of the key judgments that firms have to make is whether their partners have cooperative or opportunistic intent. Judgment of opportunistic intent is influenced by norms that contain information about the behavior that may be expected of a partner. The reliability of this information, however, depends on how rich and detailed the norms are. In most industries, norms that emerge over a longer period of time are richer and

more detailed, providing a more reliable source for estimating the probability of opportunism. This is due to the fact that the longer the span of norm emergence, the more learning by each side about its partners becomes feasible. This learning, however, may produce a willingness to take risks which in retrospect may not be warranted. Clearly, more data and more information can lead to learning more about situations prone to strategic surprise. Yet as noted by Einhorn (1982) and Fischhoff (1982), learning form the past is not an easy task. The learning task is even more difficult when dealing with strategic surprises that are, by their nature, rare events.

The emergence of norms takes place against the background of negotiations, discussions, and other social encounters. Such behaviors are of relatively low cost to the parties involved. They serve to create normative links to actions whose consequences are more costly. For an opportunistic firm, the ability to engage in low-cost actions that link up with norms which predict other more costly actions opens the way for a "normative free-ridership" strategy (Goldstone 1994). Such strategy relies on the fact that firms build their evaluation of the power of norms on the basis of populations rather than individual behavior. Thus, paradoxically, the emergence of a normative environment increases the danger of strategic surprise by encouraging firms to reduce ambiguity through the use of estimated correlations with other populations of events in their inference process. A subcontractor assumes that a client abides by a set of common norms, and using these norms, derives an estimate of the probability that certain actions signal opportunistic intent. In the absence of norms, the subcontractor would probably have made a more conservative estimate of such a distribution with the consequent burden on managerial vigilance. If norms exist, and the subcontractor concludes that the client abides by them, he/she will eventually tend to be less vigilant. Lower vigilance reduces the costs of routine transactions, but at the same time it also increases the damage that may result from a potential strategic surprise.

Finally, while it is not possible to eliminate strategic surprises, firms should attempt to reduce their probability. There are potential measures firms can take regarding the judgmental procedure on the one hand, and the economic/legal perspective on the other. From the judgmental perspective firm are encouraged to do the following: First, as depicted in the model in Figure 1, the higher the predictability (i.e., r_{xy}) the lower the probability of strategic surprise. Thus, firms should collect more information to validate the judgmental model they use. Second, firms should not rely too much on history in predicting their partners' intentions. Even if historical evidence is systematically collected it is not a guarantee against potential changes in partners' intentions. Thus, firms should scrutinize their environments continuously in an attempt to detect changes that may affect their partners' goals and intentions. Third, firms should develop a sound process of interpreting the information that is being gathered. The reliance on a single source of intelligence may prove to be wrong, indeed, following the Yom Kippur War the Israeli government established a few bodies who interpret the information that is being gathered independently rather than relying on one source of intelligence. This may help in getting a better estimate, and it can also serve as a safeguard against the incentives different sources may have in proving that they were right. From the economic/legal perspective, firms should attempt to diversify their dependence on buyers and suppliers so as not to reach an extreme level of dependence, as shown in the case of TCI. Furthermore, firms should include enforcement costs in assessing the probability of opportunistic behavior by partners. Many relationships are built on legal agreements that can deter predatory opportunistic behavior by partners if they are enforceable, and especially if they are enforceable at relatively low cost.

Future research should examine the ways firms interpret what they perceive to be opportunistic behavior of their partners. What information do they collect, how do they change the alarm triggers, and what actions they take to anticipate such behavior? Of particular interest would be research that would examine and compare the above for firms whose relationships with their buyers (suppliers) have developed over a long (short) period of time.

Acknowledgments

An earlier version of this paper was presented at the Conference on Risk, Managers, and Options in honor of Ned Bowman at the Wharton School. Comments of the participants as well as suggestions made by Sari Carp, Huw Davies, Bruce Kogut, Howard Kunreuther, Jonathan Leland, Dan Levinthal, Sridhar Seshadri, Yaacov Vertzberger and three anonymous referees are acknowledged.

References

Ansoff, H. I. 1975. Managing strategic surprise by response to weak signals. *California Management Rev.* **18** 21–33.

Arrow, K. J. 1992. Rationality of self and others in an economic system. Mary Zey, ed. *Decision Making: Alternatives to Rational Choice Models.* Sage, Newbury Park, CA.

Berg, N. A., G. W. Merry. 1984. Polaroid-Kodak. Harvard Business School case, 376–266. Harvard University, Boston, MA.

Bowman, E. 1963. Consistency and optimality in managerial decision making. *Management Sci.* **9** 310–321.

——, H. Kunreuther. 1988. Post-Bhopal behaviour at a chemical company. *J. Management Stud.* **25** 387–401.

Camerer, C. F. 1997. Progress in behavioral game theory. *J. Econom. Perspectives* **11** 167–188.

Chiles, T., J. McMackin. 1996. Integrating variable risk preferences,

trust, and transaction cost theory. *Acad. Management Rev.* **21** 73–99.

Coleman, J. S. 1990. Norm-generating structures. K. S. Cook, M. Levi, eds. *The Limits of Rationality*. The University of Chicago of Press, Chicago, IL.

Dasgupta, P. 1988. Trust as a commodity. D. Gambetta, ed. *Trust*. Basil Blackwell, Oxford, U.K.

Dawes, R. M., R. H. Thaler. 1988. Cooperation. *J. Econom. Perspectives* **2** 187–197.

Dixit, A., B. Nalebuff. 1991. *Thinking Strategically*. W. W. Norton, New York.

Durkheim, E. 1947. *The Division of Labor in Society*. Free Press, Glencoe, IL.

Einhorn, H. 1982. Learning from experience and suboptimal rules in decision making. D. Kahneman, P. Slovic, A. Tversky, eds. *Judgment Under Uncertainty: Heuristics and Biases.* Cambridge University Press, Cambridge, U.K.

——, R. Hogarth. 1978. Confidence in judgment: Persistence of the illusion of validity. *Psych. Rev.* **85** 395–416.

Farrell, J. 1987. Cheap talk, coordination and entry. *Rand J. Econom.* **18** 34–39.

Fischhoff, B. 1982. For those condemned to study the past: Heuristics and biases in hindsight. D. Kahneman, P. Slovic, A. Tversky, eds. *Judgment Under Uncertainty: Heuristics and Biases*. Cambridge University Press, Cambridge, U.K.

Frank, R. 1987. If homo economicus could choose his own utility function, would he want one with a conscience? *Amer. Econom. Rev.* **77** 593–604.

Ghoshal, S., P. Moran. 1996. Bad for practice: A critique of the transaction cost theory. *Acad. Management Rev.* **21** 13–72.

Goldstone, J. A. 1994. Is revolution individually rational—groups and individuals in revolutionary collective action. *Rationality and Society* **6** 139–166.

Handel, M. 1976. Perception, deception and surprise: The case of the Yom Kippur war. Working paper No. 19. Leonard Davis Institute for Peace Studies, Hebrew University, Jerusalem.

Hay, D. A., D. J. Morris. 1991. *Industrial Economics and Organization*. Oxford University Press, Oxford, U.K.

Heide, J., A. Miner. 1992. The shadow of the future: Effects of anticipated interaction and frequency of contact on buyer-seller cooperation. *Acad. Management J.* **35** 265–291.

Issacs, J., T. Downing. 1998. *Cold War*. Little Brown and Company, London, U.K.

Johnston, R., P. R. Lawrence. 1988. Beyond vertical integration. *Harvard Bus. Rev.* **66** (July-August) 94–101.

Kahn, H. 1965. *On Escalation: Metaphors and Scenarios*. Praeger, New York.

Kylen, B. J. 1985. What business leaders do—Before they are surprised. *Adv. Strategic Management* **3** 181–222.

Levinthal, D. A., M. Fichman. 1988. Dynamics of interorganizational attachments—Auditor-client relationships. *Admin. Sci. Quart.* **33** 345–369.

Lorenz, E. H. 1988. Neither friends nor strangers: Informal networks of subcontracting in French industry. D. Gambetta, ed. *Trust: Making and Breaking Cooperative Relations*. Basil Blackwell, London, U.K. 194–210.

——. 1993. Flexible production systems and the social construction of trust. *Politics and Society* **21** 307–324.

Macaulay S. 1963. Non-contractual relations in business: A preliminary study. *Amer. Soc. Rev.* **28** 55–69.

Payne, J. W., J. R. Bettman, E. J. Johnson. 1993. *The Adaptive Decision Maker*. Cambridge University Press, New York.

Prange, Gordon W. 1982. *At Dawn We Slept.* Michael Jospeh, London, U.K..

Provan, K. G., S. J. Skinner. 1989. Interorganizational dependence and control as predictors of opportunism in dealer-supplier relations. *Acad. Management J.* **32** 202–212.

Sabel, C. 1993. Studied trust-building new forms of cooperation in volatile economy. *Human Relations* **46** 1133–1170.

Sako, M. 1992. *Prices, Quality and Trust*. Cambridge University Press, Cambridge, U.K.

Saunders, J. 1997. Power, TCI taking their differences to the courts. *Computing Canada* **23** 22.

Schick, S. 1998. TCI suing Power Computing for US$42.75 million in damages. *Computer Dealer News* **13** 6.

Shapira, Z. 1995. *Risk Taking: A Managerial Perspective*. Russell Sage Foundation, New York.

Stinchcombe, A. L. 1986. *Norms of Exchange, in Stratification and Organization: Selected Papers*. Cambridge University Press, Cambridge, U.K.

Sunstein, C. 1996. Social roles and social norms. *Columbia Law Rev.* **96**(4) 903–968.

Swedberg, Richard. 1994. Markets as social structures. Neil J. Smelser, Richard Swedberg, eds. *The Handbook of Economic Sociology.* Princeton University Press, Princeton, NJ 255–282.

Tversky, A., D. Kahneman. 1971. Belief in the law of small numbers. *Psych. Bull.* **76** 105–110.

——, ——. 1982. Evidential impact of base rates. D. Kahneman, P. Slovic, A. Tversky, eds. *Judgment Under Uncertainty: Heuristics and Biases*. Cambridge University Press, Cambridge, U.K.

Watts, B. D. 1996. Clausewitzian friction and future war. McNair Paper Number 52, *Institute for National Strategic Studies*, Washington, DC.

White, H. C. 1986. Varieties of markets. B. Wellman, S. D. Berkowitz, eds. *Social Structures: A Network Approach.* Cambridge University Press, Cambridge, U.K. 226–260.

Wholstetter, R. 1962. *Pearl Harbor: Warning and Decision.* Stanford University Press, Stanford, CA.

Williamson, O. E. 1985. *The Economic Institutions of Capitalism*. Free Press, New York.

——. 1993. Calculativeness, trust, and economic organization: The analysis of discrete structural alternatives. *J. Law and Econom.* **36** 269–296.

Young, P. 1996. The economics of convention. *J. Econom. Perspectives* **10** 105–122.

Zeira, E. 1993. *The Yom Kippur War*. Yediot Aharonot, Tel Aviv.

Accepted by Arie Lewin.