Improving Critical Thinking

Marvin S. Cohen and Jared T. Freeman
Cognitive Technologies, Inc.
Arlington, VA

Is Naturalistic Training Possible?

The naturalistic approach to decision research takes as its starting point the way people actually make decisions in real-world environments, as revealed in interviews, observation, and contextually realistic experimentation . It does not start with a mathematical or logical model of how decisions ought to be made, nor does it typically compare behavior in artificial laboratory tasks to such models (Cohen, 1993). However, there is more agreement about the starting point of naturalistic research than about its destination. What functions does this research ultimately serve? In particular, can it generate *prescriptions* about how to think better or make better decisions? Will it eventually arrive where other approaches begin, and lead to training that mitigates the shortcomings of ordinary thinking?

For a variety of reasons, we might suppose the answer is *no*. First, there is the logical prohibition against deriving an *ought* from an *is*, a mistake which is called by philosophers, appropriately enough, the *naturalistic fallacy*. We cannot conclude that a particular decision process is the best one available simply because real decision makers use it. Second, naturalistic researchers allegedly view real-world decision making through rose-tinted glasses (Doherty, 1993). Indeed, some naturalistic decision researchers have criticized the idea, promoted by Kahneman, Slovic, and Tversky (1982) and others, that ordinary decision making is riddled by systematic errors or biases (e.g., Cohen, 1993). Third, there is an emphasis in naturalistic research on pattern recognition rather than on more explicit processes of reflective reasoning (Klein, 1993). We can summarize these pessimistic points as follows: In naturalistic research, prescription is *impossible* because it would confuse what is and what ought to be, *unnecessary* because real-world decision making is already pretty good, and *irrelevant* because real-world decision making is intuitive rather than reflective.

I think each one of these claims is wrong or misleading. In this chapter, I will describe a naturalistic training strategy for improving decision making skills. This

strategy will serve as a counterexample to all three of the objections itemized above:

- The training strategy is premised on the importance of critical thinking skills that complement and go beyond pattern recognition (point 3 above). These skills monitor, verify, and improve the results of recognition in high-stakes and novel situations, when immediate action on a recognized response is not necessary. Critical thinking skills are inextricably tied to the recognitional processes they regulate, however, and do not represent an analytical *alternative* to recognition-based processing. Similarly, Klein (1993) describes how people use mental simulation to verify actions when recognition of the appropriate response is uncertain.

- The model that underlies the training is based on interviews with and observations of real-world decision makers. But it does not involve the naturalistic fallacy, because it does not indiscriminately infer what is desirable from what exists (point 1). Instead, the model of critical thinking skills is based on (a) *comparisons* of more and less experienced, and more and less proficient, real-world decision makers, and on (b) *correlations* between the use of the critical thinking skills (e.g., after training) and successful performance in real-world tasks, such as better situation assessments and better decisions. The model derives additional tentative normative force from the face validity and plausibility of the decision strategies it describes (Shafer & Tversky, 1988; Cohen & Freeman, 1996).

- Finally, the naturalistic approach does not imply that real-world decision makers never make errors (point 2). Errors can be identified by examining discrepancies between more and less experienced, or more and less proficient, decision makers as identified by peers, or aspects of decision processes that are correlated with *less* successful performance in real-world tasks, where success involves the achievement of explicit organizational or personal goals. Rather than denying the existence of errors, the naturalistic approach provides a more useful way of looking at errors. For example, they are not defined as deviations from the purely formal constraints of decision theory. (Such definitions prove unexpectedly slippery in any case, since deviations from one formal model may be consistent with some other model; see Smithson, 1989; Cohen, 1993.)

A more thorough discussion of these and other issues may be found in Cohen (1993).

A Naturalistic Model of Critical Thinking

Proficient decision makers are *recognitionally skilled*: that is, they are able to recognize a large number of situations as familiar and to retrieve an appropriate response. Recent research in tactical decision making suggests that proficient decision makers are also *meta-recognitionally skilled* (Cohen, Freeman, & Wolf, 1996). In novel situations where no familiar pattern fits, proficient decision makers supplement recognition with processes that verify its results and correct problems.

Based on critical incident interviews with active-duty naval officers, we developed a framework for decision making, called the Recognition/Metacognition (R/M) model (Cohen et al., 1996). The model describes a set of critical thinking strategies that supplement recognitional processes. Structured situation models (i.e., *schemas*), often in the form of *stories* about enemy intent, causally organize information about a situation and provide a basis for metarecognitional processes. Metarecognitional processes determine when it is worthwhile to think more about a problem; identify evidence-conclusion relationships within the story; critique the story for incompleteness, conflict, and unreliability; and attempt to improve it, by collecting or retrieving new information and revising assumptions. At a somewhat more detailed level, meta-recognitional processes include:

1• Identification of evidence-conclusion relationships (or *arguments*) within the evolving situation model and plan. This is simply an implicit or explicit awareness that cue A was *observed* on this occasion, while intent to attack along with expectations of observing cue B were *inferred*. On some other occasion cue B might be observed and cue A inferred.

2• Processes of *critiquing* that identify problems in the arguments that support a conclusion (e.g., hostile intent) within the situation model or plan. Critiquing can result in the discovery of three kinds of problems: *incompleteness, unreliability*, or *conflict*. An argument is incomplete if it does not provide support either for or against a conclusion of interest (e.g., the kinematics of the track suggest only that it is a military aircraft, but say nothing about hostile intent; this conclusion is too general for deciding whether to engage). Two arguments conflict with one another if they provide support both for and against a conclusion, respectively (e.g., the heading of a track toward own ship suggests hostile intent, while its slow speed argues for routine patrol). Finally, an argument is unreliable if it provides support for, but not against, a conclusion, but the support depends on unexamined assumptions. Unreliable support may shift or vanish when its premises are further considered.

3• Processes of *correcting* that respond to these problems. Correcting

steps may instigate external action, such as collecting additional data, and two kinds of internal actions, attention shifting and assumption revision, that regulate the operation of the recognitional system. Shifting the focus of attention stimulates retrieval of new, potentially relevant information in long-term memory and brings additional arguments into view for meta-recognitional critiquing. Adding or dropping assumptions permits what-if reasoning, queries for alternative causes and effects, and adoption a single coherent model or plan. These processes in combination help to fill gaps in models or plans, resolve conflict among arguments, and search for more reliable arguments.

4• A higher-level process, called the *quick test*, which controls critiquing and correcting. Metarecognitional strategies, like other actions, are shaped in part by past experiences of success and failure. Meta-recognitional processing occurs when the benefits associated with critical thinking outweigh the costs. This is likely to be the case when the costs of delay are acceptable (i.e., time is available for critical thinking), the situation is uncertain or novel (i.e., recognitional conclusions are subject to improvement), and the costs of an error in acting on the current recognitional conclusion are high. The quick test considers these three factors and, if conditions are appropriate, inhibits recognition-based responding and interposes a process of critical thinking. When these conditions are not satisfied, the quick test allows immediate action based on the current best response.

Figure 1 summarizes the relationships among these processes. It highlights the functional distinction between recognitional processes (at the top of the figure) and metacognitive ones (the shaded boxes). The recognitional level provides information to the metacognitive level, while the metacognitive level exerts control over the recognitional level. In the R/M model, metacognition monitors the recognitional processing, maintains a model or description of it (i.e., identifies arguments and problems of incompleteness, conflict, and unreliability), and modifies recognitional activity by inhibiting overt action, shifting attention, and adopting or dropping assumptions. These functional differences may or may not correspond to structural or physiological ones (see Nelson & Narens, 1994). A more detailed description of the R/M model may be found in Cohen et al. (1996).

The R/M model explains how experienced decision makers are able to exploit their experience in a domain and at the same time handle uncertainty and novelty. They construct and manipulate concrete, visualizable models of the situation, not abstract aggregations (such as 70% chance of hostile intent, 30% chance non-hostile). Uncertainty is represented explicitly at the metacognitive level, by "annotating" the situation model or plan to highlight points of incompleteness, conflict, and unreliability. In response to specific problems of this kind,

metacognitive strategies try to improve the current situation model and plan or find better ones.

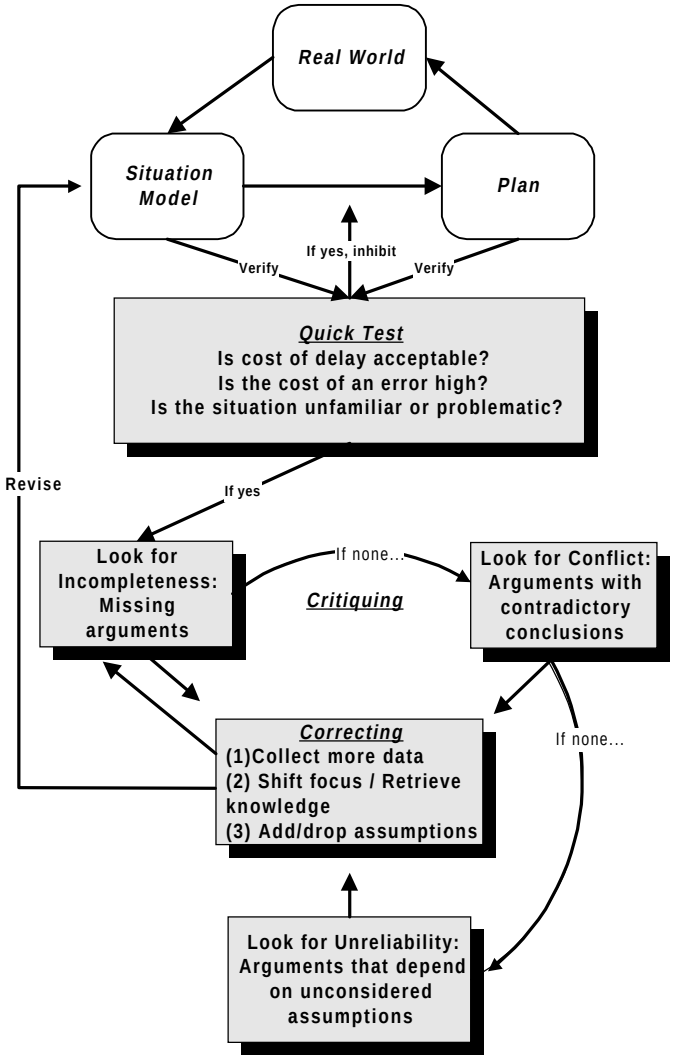

Figure 1. Components of the Recognition/Metacognition Model.

Metarecognitional processing is highly dynamic and iterative. The next processing step is determined locally by the results of earlier steps, rather than by a global, fixed procedure (as in Bayesian inference). Correcting for one problem may sometimes (but not always) lead to identification and correction of another problem. For example, a gap in an argument may be filled by collecting further data or remembering previously known information, or, if these fail, by making assumptions. The resulting more specific argument may then turn out to conflict with other arguments. Such conflict may then be addressed by looking for unreliability in one of the conflicting arguments. In doing so, metarecognitional processing might shift focus from the conclusion to the grounds of the argument. This may result in retrieval of previous experiences with the source of the information that is the grounds for the conflicting argument. Such experiences may suggest that the source is not to be trusted. The conflict, which arose because

of the implicit, or unconsidered, assumption that this source was accurate, is now resolved. (Alternatively, what if no relevant information were retrieved about the source? A new cycle of critiquing would identify this gap in knowledge, and it might be corrected, for example, by adopting the explicit assumption that the unfamiliar source is not trustworthy. Conflict would be eliminated, but the story now depends on the potentially unreliable assumption about this source. Attention might now be shifted to the other conflicting arguments.) This process stops when the quick test indicates that the benefits of further metarecognitional actions are likely to be outweighed by the risks of delay, and that action on the basis of the current best model or plan is called for. The output is a coherent, consistent model or plan together with an understanding of its strengths and weaknesses.

A Naturalistic Strategy for Critical Thinking Training

Training based on the Recognition/Metacognition model has been developed, focusing on the decision of whether or not to engage an approaching air or sea contact whose intent is unknown, under conditions of undeclared hostility. The training is based on interviews with active-duty Naval officers, in which they described experiences of this kind in the Persian Gulf, the Gulf of Sidra, and elsewhere (Kaempf, Klein, Thordsen, & Wolf, 1996). Many aspects of the training are based on differences in the way that more and less experienced officers handled similar situations.

We call the training *critical thinking* because it is designed for situations where familiar patterns or rules do not fit. For example, some features of the situation may match the standard hostile intent pattern (e.g., an aircraft turning toward own ship from a hostile country), but others do not (e.g., its speed is slower than expected) and may even match parts of another pattern (e.g., commercial airliner). The training is divided into four segments: (1) An overview of the cycle of creating, testing, and evaluating stories to improve situation understanding; (2) a particular kind of story based on hostile intent; (3) strategies for finding and correcting problems with stories; and (4) guidelines for when critical thinking is appropriate and when immediate action is necessary. In each of these segments, officers listen to a brief verbal presentation of the concepts, followed by questions and discussion. They then participate in interactive scenario-based exercises designed to provide practice in the relevant skill. Feedback during these exercises is provided by group discussion and by the instructor. The basic concepts of the four training segments are as follows:

- *Creating, testing, and evaluating stories.* This section provides an overview of the critical thinking process, called *STEP*. When an assessment is uncertain, decision makers can take it seriously by constructing a *S*tory around it. The story includes the past and future events that would be expected if the assessment were true. Decision makers use the story to *T*est the assessment, by comparing expectations

to what is known or observed. When evidence appears to conflict with the assessment, they try to patch up the story by explaining the evidence. They then *E*valuate the result; if the patched up story involves too many unreliable assumptions, they generate alternative assessments and begin the cycle again. In the meantime, they *P*lan against the possibility that their current best story is wrong.

- Figure 2 summarizes the *STEP* process.

- *Hostile-intent stories*. Stories contain certain typical components. Knowledge of these components can help decision makers notice and fill gaps in the stories they construct. A particularly important kind of story is built around the assessment of hostile intent. For example, a complete hostile intent story explains why an attack is taking place against a particular target by a particular platform, and accounts for how that platform has localized the target and is arriving at a position suitable for engaging it (see Figure 3). The training teaches officers by practice and example how to discover story components and to let them guide decision makers to relevant evidence regarding assessments of intent.

- *Critiquing stories.* After a story is constructed, decision makers step back and evaluate its plausibility. This segment of the training introduces a devil's advocate technique for uncovering hidden assumptions in a story and generating alternative interpretations of the evidence. An infallible crystal ball persistently tells the decision maker that the current assessment is wrong, despite the evidence that appears to support it, and asks for an explanation of that evidence. Regardless of how confident decision makers are in their assessments, this technique can successfully alert them to significant alternatives. It can also help them see how conflicting data could fit into a story. In each case, the technique helps decision makers expose and evaluate assumptions underlying their reading of the evidence.

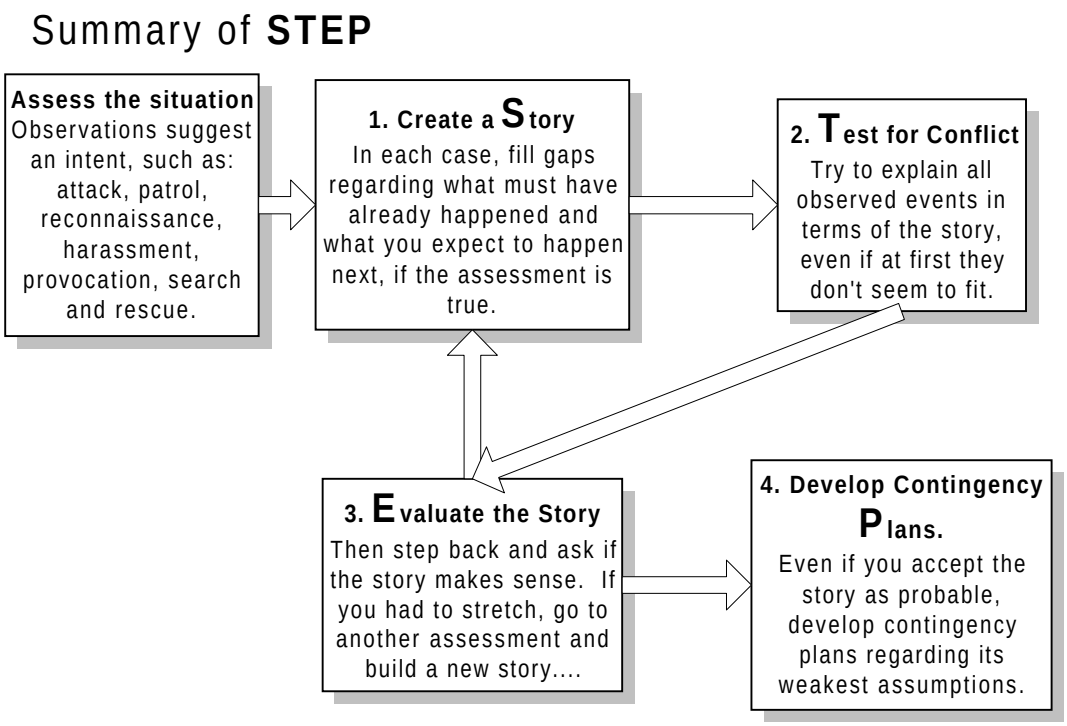

Figure 2. A cycle of four steps for critical thinking.

- *When to think more.* Critical thinking is not always appropriate. Unless three conditions are satisfied, the decision maker should probably act immediately: (1) The risk of delay must be acceptable. (2) The cost of an error if one acts immediately must be high. And (3) the situation must be non-routine or problematic in some way. Training focuses on the way experienced decision makers apply these criteria. For example, they tend to utilize more precise estimates of available time based on the particularities of the situation, a longer-term outlook in estimating the cost of an error, and greater sensitivity to the mismatch between the situation and any familiar pattern.

Critical thinking training has now been tested with active-duty officers at two Navy training facilities (Cohen, Freeman, and Thompson, in press). The evaluation examined the effects of training on critical thinking processes, its effects on performance (i.e., assessments and actions), and participants' own evaluations of the training. The results were encouraging. For example, in one study training increased the number of factors officers considered in assessing the intent of a track by 30%, increased the amount of conflicting evidence they noticed by 58%, increased the number of assumptions they identified underlying the interpretation of that evidence by 27%, and increased the number of alternative assessments they generated by 41%. Critical thinking training can also improve the accuracy of assessments. Agreement with a subject matter expert increased significantly in two out of four test scenarios in the two studies, by 79% and 35%, respectively. At the same time, the training (non-significantly) increased officers'

confidence in their assessments in both studies. Subjective evaluations of the training were generally positive.

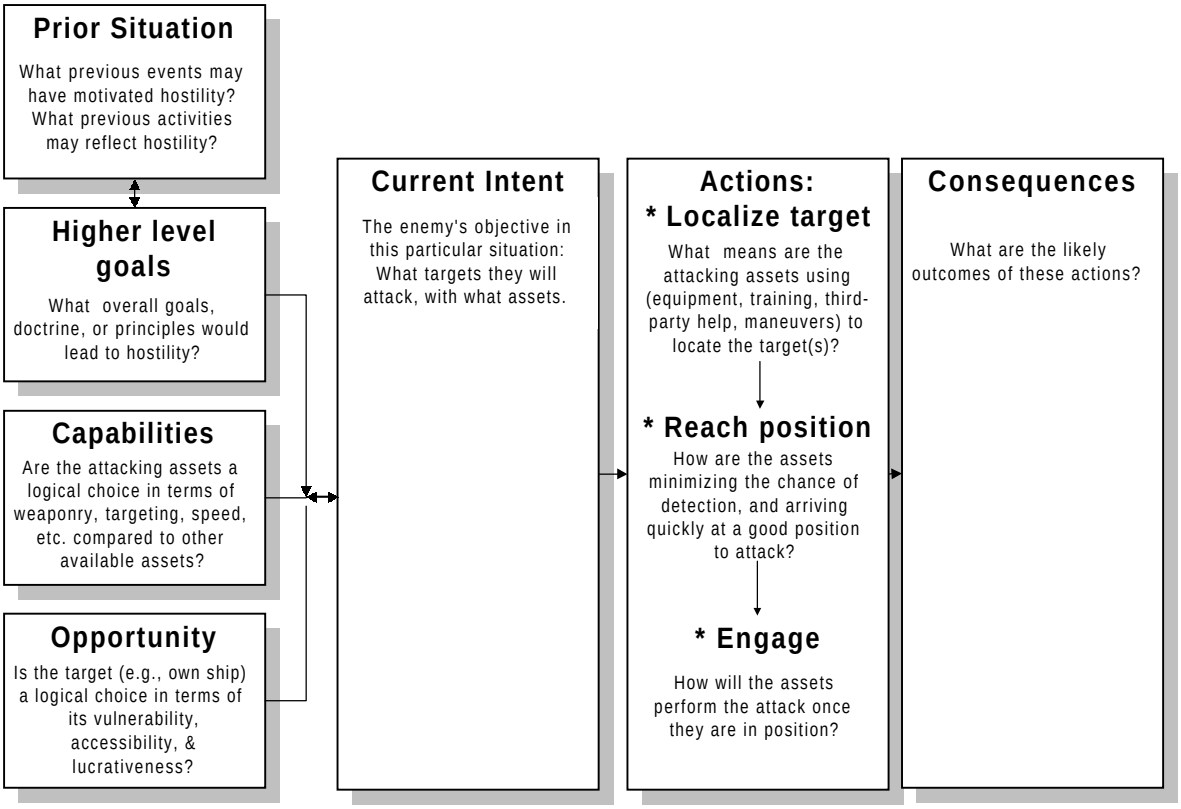

Figure 3. Elements of a hostile intent story.

These tests strongly suggest that meta-recognitional skills can be taught effectively, that officers will use them in relatively realistic tactical situations, and that enhanced meta-recognitional skills will lead to improved performance.

Conclusion

The research reported here illustrates the transition of a naturalistic model of decision making to training. This transition naturally implies a prescriptive commitment, a claim that the strategies to be taught will lead to successful decision making performance. The initial justification for the cognitive model and the training based on it was a comparison of strategies used by more and less experienced decision makers in critical incidents elicited in interviews. A step further in confirming the prescriptive validity of the training was reflected in demonstrations of the effectiveness of the training. In particular, the training not only increases the frequency of critical thinking strategies, but also appears to improve both the accuracy of assessments and the appropriateness of actions.

References

Cohen, M.S. (1993), 'The naturalistic basis of decision biases.' In G.A. Klein, J. Orasanu, R. Calderwood & C.E. Zsambok (Eds.), *Decision Making in Action: Models and Methods*, Norwood, NJ: Ablex Publishing Corporation, pp. 51-102.

Cohen, M.S., & Freeman, J.T. (in press), 'Critical thinking in tactical decision making: A model and a training strategy.' In J.A. Cannon-Bowers & E. Salas (Eds.), *Decision Making under Stress: Implications for Training and Simulation*.

Cohen, M.S., & Freeman, J.T. (1996), 'Thinking naturally about uncertainty,' *Proceedings of the Human Factors & Ergonomics Society, 40 Annual Meeting*, Santa Monica, CA: Human Factors Society.

Cohen, M.S., Freeman, J.T., & Wolf, S. (1996), 'Meta-recognition in time stressed decision making: Recognizing, critiquing, and correcting,' *Human Factors, 38*(2), pp. 206-219.

Doherty, M.E. (1993), 'A laboratory scientist's view of naturalistic decision making.' In G.A. Klein, J. Orasanu, R. Calderwood & C.E. Zsambok (Eds.), *Decision Making in Action: Models and Methods*, Norwood, NJ: Ablex Publishing Corporation, pp. 362-388.

Kaempf, G.L., Klein, G., Thordsen, M.L., & Wolf, S. (1996), 'Decision making in complex command-and-control environments,' *Human Factors, 38*(2), pp. 206-219.

Kahneman, D., Slovic, P., & Tversky, A. (Eds.), (1982), *Judgment Under Uncertainty: Heuristics and Biases*, Cambridge, UK: Cambridge University Press.

Klein, G.A. (1993), 'A Recognition-Primed Decision (RPD) model of rapid decision making.' In G.A. Klein, J. Orasanu, R. Calderwood & C.E. Zsambok (Eds.), *Decision Making in Action: Models and Methods*, Norwood, NJ: Ablex Publishing Corporation., pp. 138-147.

Nelson, T.O., & Narens, L. (1994), 'Why investigative metacognition?' In J. Metcalfe & A.P. Shimamura (Eds.), *Metacognition*, Cambridge, MA: The MIT Press, pp. 1-25.

Shafer, G., & Tversky, A. (1989), 'Languages and designs for probability judgment.' In D.E. Bell, H. Raiffa & A. Tversky (Eds.), *Decision Making: Descriptive, Normative, and Prescriptive Interactions*, Cambridge, UK: Cambridge University Press, pp. 237-265.

Smithson, M. (1989), *Ignorance and Uncertainty: Emerging Paradigms*, NY: Springer-Verlag.