PITFALLS IN THE USE OF IMPERFECT INFORMATION

William E. Dean

April 1988

P-7430

PITFALLS IN THE USE OF IMPERFECT INFORMATION

INTRODUCTION

High-intensity conflict has come a long way since the days when a proverbial Napoleon could stand on a hillside and take in the full scope of the battle. Commanders at the corps or theater level do not rely on their own senses for information, but on a group of specialists coordinated by the intelligence officer. That officer serves as a conduit to the commander, and is his source of information about the enemy. The quality of the commander's understanding of enemy intentions depends on the quality of the communication between the intelligence officer and the commander.

Between the two extremes -- either the commander has no intelligence regarding enemy intentions, or enemy intentions are obvious -- lies the case of imperfect information, where there is non-trivial yet uncertain knowledge about the other side's intentions. In this kind of situation, breakdowns in communication can produce two possible results:

- Under-confidence: the intelligence officer communicates the uncertainty, and the commander disregards the information because it is uncertain, and
- Over-confidence: the intelligence officer suppresses the uncertainty, and the commander takes the information at face value.[1]

This paper considers the impact of these pitfalls by using game theory and decision theory to show how the outcome of "faulty" decisions compares with the outcome from "good" decisions.

[1]Of course, the commander could suppress the uncertainty all by himself.

AN EXAMPLE

Consider a zero-sum game between two opponents, designated Red and Blue. The range of strategies and their payoffs are common knowledge, but suppose that Blue has some knowledge of which strategy Red may choose; that is, Blue has a probability distribution for Red strategies.

At the lowest level of information, Blue merely calculates the rational strategy that Red would adopt using game theory. In particular, if the solution is a mixed strategy, Blue remains uncertain about Red's choice of strategy.

At an intermediate level, Blue might have some intelligence on Red, such as observations of preparations for Red's next move. This information, though imperfect, enables Blue to employ decision theory.

Figure 1 shows a simple example in which Red has two strategies, R1 and R2. Blue's strategy B1 is excellent as a counter to R1, so the payoff equals 1, the best that can happen (victory). Likewise B2 is a counter to R2. However, B1 is a dismal match to R2, so the payoff equals 0, the worst that can happen (defeat). Likewise, B2 is a mismatch for R1. Blue also has a hedging strategy B3, with which Blue can muddle through no matter what Red may choose. The payoff for B3 equals u, some value between 1/2 and 1.

		Red Strategies		
		R1	R2	
	B1	1	0	
Blue Strategies	B2	0	1	
	B3	u	u	where 1/2 < u < 1

Fig. 1--Payoff Matrix for Blue

For instance, suppose that in a conventional war in Europe, Red has an operational maneuver group (OMG) to insert. Let R1 and R2 correspond to the possible places to insert the OMG. Blue wants to interdict the OMG and also send ground forces to reinforce the anticipated insertion point. Let B1 and B2 correspond to the possible places to anticipate the OMG. Let B3 be the strategy of sending half of the forces to each place.

The outcome of B1 against R1 (B2 against R2) is the successful "chewing up" of the OMG. The outcome of B1 against R2 (B2 against R1) is a Blue attack on the wrong target, and Red achieves a breakthrough.

The outcome of B3 leads to neither the destruction of the OMG nor a nasty breakthrough, but something in between, not of a spectacular nature. The "value" of B3, u, is not calculated from objective results of battle (how many km of Red advance, how much attrition), but rather u indicates the commander's preference for muddling through for sure compared to risking defeat. A lot of subjective considerations get wrapped up in this little number. As will be shown below, u is a *preference-probability*[2] because it states how the commander values the outcome of B3 compared to the best and worst possible outcomes.

Here are examples of ideas that might be rattling around inside the head of a risk-averse (high-u) commander:

- "We only get one chance for defense. If they get a breakthrough, it's all over!"
- "If I lose this battle, my reputation is ruined!"

A less risk-averse (moderate-u) commander might be thinking:

- "This is just one of many battles."
- "If I win this battle, I'm a hero and I get a rapid promotion."

---

[2]Robert D. Behn and James W. Vaupel, *Quick Analysis for Busy Decision Makers*, Basic Books, 1982, pp. 26-46.

Specifically, if the commander *could* choose between muddling through for sure on the one hand, and a lottery like that represented in Figure 2, he would have three choices which depend on a comparison of a value u and a probability X:

$X < u$, choose to muddle through for sure.

$X = u$, be indifferent between the options.

$X > u$, choose the risky lottery.

Hence u is not only the outcome of strategy B3, it is also a quantitative indication of the commander's willingness to take a risk. The more risk-averse the commander, the closer u gets to 1. If the commander thinks that defeat is awful, that attitude tends to raise u, but if he thinks that victory is far better than muddling, that thought tends to lower u.

```
                             prob = X
                           +----------+ Victory
                           |
            Lottery     +------+
   +--------------------+chance|
   |                    +------+
   |                       |prob = 1-X
   |                       +----------+ Defeat
+--+---+
|      |
|choice|
|      |
+--+---+
   |
   |
   |        Certainty
   +----------------------------------+ Muddle Through
```

X = probability of victory if the lottery is chosen

Fig. 2--Hypothetical Lottery

GAME THEORY

Game theory applies to the case in which Blue has no information about Red intentions. For the two-player zero-sum game, two points of view are possible.[3] One view looks at the opponent's capability; this leads to the "maxmin" or "minimax" principle. The other approach looks at the opponent's intentions. Both approaches lead to the same solution. I adopt the latter approach here, because it serves as a ground-work for the imperfect-information case.

Let b_1, b_2, b_3 be the frequency with which Blue selects strategy B1,B2,B3, respectively. Let r_1, r_2 be the frequency with which Red selects strategy R1,R2. The value v of the game to Blue is given by

$$\begin{aligned} v = {} & (b_1\, r_1\, 1) + (b_1\, r_2\, 0) \\ & + (b_2\, r_1\, 0) + (b_2\, r_2\, 1) \\ & + (b_3\, r_1\, u) + (b_3\, r_2\, u). \end{aligned}$$

By symmetry it follows that

$$b_1 = b_2 = 1/2(1 - b_3), \quad \text{and} \quad r_1 = r_2 = 1/2.$$

Thus the formula boils down to

$$v = 1/2(1 - b_3) + (b_3\, u) = 1/2 + (u - 1/2)\, b_3.$$

Clearly, for $u > 1/2$, the game-theoretic solution consists of *always* choosing B3, the muddle-through strategy, for the value

$$v = u.$$

---

[3]See, for example, Jonathan Cave, *Introduction to Game Theory*, The RAND Corporation, P-7336, April 1987.

IMPERFECT INFORMATION

Now suppose that Blue has some intelligence about what Red plans to do, that based on some sort of analysis Blue estimates the probability of Red choosing R1 to be p, and the probability of Red selecting R2 to be $1 - p$. The symmetry is broken, so the preceding equations no longer apply. The expected outcome to Blue is now given by

$$v = (b_1\ p\ 1) + [b_1\ (1-p)\ 0]$$

$$+ (b_2\ p\ 0) + [b_2\ (1-p)\ 1]$$

$$+ (b_3\ p\ u) + [b_3\ (1-p)\ u].$$

This simply boils down to

$$v = (b_1\ p) + [b_2\ (1-p)] + (b_3\ u).$$

The choice of strategy, as well as the expected outcome, depends on p *and* u:

$p > u$, choose B1 for an expected outcome of p.
$(1-p) > u$, choose B2 for an expected outcome of $1 - p$.
$p < u$ *and* $(1-p) < u$, choose B3 for an outcome of u.

Note that in some cases Blue can do better than game theory, that is, to obtain an expected outcome in excess of u, as indicated in Figure 3. For low p, Blue has good reason to believe that Red will choose R2, so Blue is willing to take a small risk in the face of a good opportunity for victory. Likewise, for high p, Blue expects R1 and so chooses B1. For intermediate values of p, the Blue commander decides that the risk is too great, so he chooses to muddle through with B3.

Note that the decision-making depends critically on u. The more risk-averse the commander, the higher his level of u, the less likely he will make use of the available imperfect information.

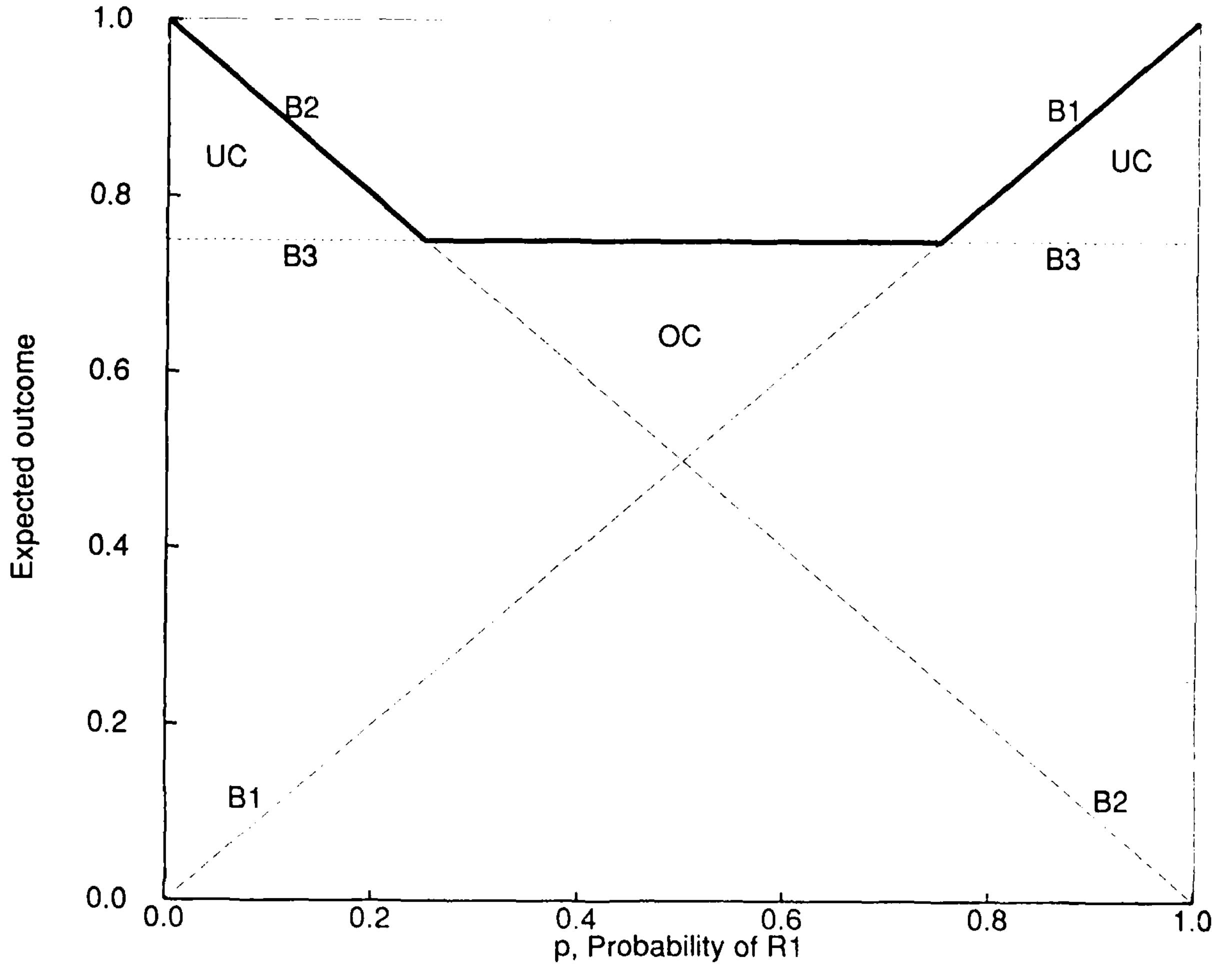

UC - Triangle of under-confidence

OC - Triangle of over-confidence

Fig. 3--Imperfect Information and Game Outcome

PITFALLS

Two pitfalls can occur when communicating and using imperfect information. One stems from under-confidence, the other from over-confidence.

Under-confidence

Suppose that the commander says, "Where do these probabilities come from anyway? Why should I believe them?" Indeed, if p has an intermediate value, such a response is harmless. However, if p or $1 - p$ exceeds u, then the commander loses something by throwing away the information. The loss depends on *both p and u*, as indicated in Figure 3.[4] The more risk-averse the commander, the less likely that he loses any opportunity for victory, and if so, the loss is slight (according to his risk preference).

Over-confidence

Suppose that the intelligent officer subjectively resolves his uncertainty and presents his best estimate of Red's choice.[5] The unwitting commander can then say, "Aha! I have just the strategy for victory."

If p is close to 0 or close to 1, such confidence entails little risk. For intermediate p, the risk is excessive. Again, the loss depends on *both p and u*. The more risk-averse the commander, the greater his chagrin if the best estimate turns out to be wrong.

---

[4]The *area* between the curves has no obvious interpretation, *unless* all values of p are equally likely.

[5]See for example a discussion of the subjective resolution of uncertainty in John D. Steinbruner, *The Cybernetic Theory of Decision*, Princeton University Press, 1974, pp. 109-124.

DISCUSSION

The risk-averse commander does not have much use for imperfect information. He throws away information unless it is nearly certain. If a commander is willing to take risks, he can make use of imperfect information in a broader range of situations.

If a commander does not trust the information, and throws it away, he can deprive himself of some good opportunities for victory. He hedges even when the probability of defeat is small.

If a commander receives only best estimates, and *knows* that the estimates are over-confident, he may throw them out. In this case, the commander over-compensates by becoming under-confident.

However, if a commander accepts best estimates, and acts on them, he can take on high risks due to the self-deception. This is an abuse of information that "just isn't there."

The intelligence officer and the commander need to communicate clearly with each other. The former needs to know the latter's preference for boldness and for hedging. The latter needs to know the former's tendency to emphasize or ignore uncertainty. Terms such as "possibly," "likely," and "probably" may serve these needs. However, explicit quantification of p and u provides a much-needed transparency.[6]

---

[6] In practice both parties may often provide interval estimates of p and u rather than point estimates.