Advances in Hypergame Theory

Russell R. Vane III
General Dynamics Advanced Information Systems
1400 Key Blvd., Suite 700
Arlington, VA 20170
1-703-516-6376

russ.vane@gd-ais.com

ABSTRACT
This articles reprises findings from hypergame work in Vane's doctoral dissertation [Vane 2000] that are directly applicable to decision-theoretic (DT) and game-theoretic (GT), or simply, DTGT agents. Hypergame theory meets many of the methodological requirements for bounded rationality agent research. While this early hypergame work considered expected value and worst case scenarios (from adversarial uncertainty), additional work for US government stimulated scenarios added considerations for luck (best case scenarios). Additionally, the effectiveness measure was refined to account for differences in the inherent robustness of each option/plan. Both of these extensions support the needs of researchers of effective DTGT agents. Discussions on context and evidentiary mechanisms are also included.

Categories and Subject Descriptors
I.2.8 [**Problem Solving, Control Methods and Search**]: DTGT agents represent an approach to reasoning in complex problem domains to accomplish objectives and reduce risks.

General Terms
Algorithms, Economics, and Theory.

Keywords
Hypergames, hypergame expected utility, uncertainty, contexts, and belief.

1. INTRODUCTION
Hypergame theory (HT) is all about attempting to model opponents and situations well. It incorporates key features of both sides of the DT, GT debate. Any model, such as HT, can be wrong, so an effectiveness measure was proposed that accounts for the spectrum from a decision maker's complete certainty to complete uncertainty. This continuum essentially blends decision theory's strong reliance on knowledge about other players' decision processes and game theory's strong reliance on a model of the exact situation. A fully effective plan is defined as one which is better than all other alternatives over the certainty-uncertainty domain. If any option/plan were fully effective it would always be the best alternative. No plan is ever that good. However, the derived effectiveness measure, called Hypergame Expected Utility (HEU), may be a very useful DTGT plan selection mechanism.

One plus for HT is its explicit commitment to representing different views of any competitive situation for each of the players. HT explicitly breaks game theory's requirement on consistent alignment of beliefs among opponents/players, which is needed for the calculation of Nash Equilibrium Mixed Strategies, hereafter NEMS. Another plus is that plan dynamics reflected in the HEU effectiveness measure can be used to delay action and gather more information to attempt to reduce uncertainty by exploring plan vulnerabilities. The additional reasoning in HT reduces the modeling parsimony of GT normal games, which can be a minus.

Since competitive situations often include a number of factors which cause opponents to view the options and results of game situations differently, HT appears more suitable to real world situations. Some of these considerations include:

(1) Differences in player knowledge, expertise
(2) Differences in player starting situation assessment
(3) Differences in player on-going assessment capability (evidence processing)
(4) Differences in player understanding of plan projection (what beats what?)
(5) Differences in player information (both at the commitment phase and during the operations)
(6) Differences in robustness, resilience of each player's plans
(7) Player Constraints because of time
(8) Differences in player creativity (what tricks can be added) such as feints, hidden reserves, denial and deception operations

These considerations motivate HT's extension of the GT normal form to account for situational aspects that might arise in real competitive situations. They will be referred to throughout the rest of the article as C-1 for consideration 1, when discussing variations from GT. Considering all of these differences in the same way that game theory treats them would become perilously combinatoric. Instead, an approach is used that accounts for much of the variability while not burdening either the modeling or computation ability of decision-makers or DTGT agents.

Lastly, hypergame theory can be used to become more rigorous in checking hypotheses and stimulating new option generation. Because so many more of the features of the competitive situation are recorded, hypergame theory's normal form can be used to evaluate the opportunity cost, potential for gain, and risk of loss that most competitive decisions should consider to be judged as wise. More will be explained about these features in the discussion and findings sections.

2. BACKGROUND

The reader is presented the minimum number of seminal hypergame theory papers needed to understand this paper's findings. The first reference to Bennett's hypergame theory was published in [Bennett, Dando 1979] as "Complex Strategic Analysis: A Hypergame Study of the Fall of France." In this article, the authors make a compelling case that strict use of game theory would have prevented the Germans from winning the May 1940 campaign. Instead the Germans did attack where France was the weakest. Bennett and Huxham attempted to generalize HT in 1982. The second is the research of Vane and co-authors to discover the general properties and limitations of HT. All diagrams and formulae are presented in the Vane style for consistency.

2.1 Birth of Hypergame Theory

Four significant contributions occurred in Bennett and Dando's work. They are: a profound understanding of the dual nature of the perceptions of two organizations or game players, a GT mechanism for predicting each organization's behavior, a realization that one side might achieve a result that vastly exceeds the GT value of the full game, and an implicit respect for the nature of adversarial outguessing. Each of these contributions will be explained below.

Cognitive psychological literature is rife with examples of how different contexts, past experiences, and cultural factors can change the effectiveness of intelligence analysts [Heuer 1999], military planners [Tolcott et al. 1989], and people in general [*Nobel laureate* Kahnemann 2002]. Bennett and Dando showed that considerations for the subgame that France was playing were important, even if it was not the full game that the Germans were considering. They actively addressed the intuition that even fully-revealed games such as chess can be played very differently, conceptually. They also showed a way to represent it.

National defense plans are hard to create (C-7). They require many staff officers to perform months of planning and documentation. Therefore the Germans could reason that the French (they) would quickly realize that they could easily stop an attack through the Ardennes and they knew that the Germans would know that they could as well. Thus by using a two-stage consistent knowledge of rationality, the French could create a subgame, calculated by applying the Nash Equilibrium Mixed Strategy (NEMS). After all, there is a cost to developing a plan against an attack that will never come. This creates the first GT normal game. Using the German's supposition that the French might exclude seriously planning against a German attack that they could easily thwart, the Germans created a German GT normal game with the added effects of surprise allowed. Bennett and Dando used the probability calculations from the French subgame to evaluate the new German game, rather than applying the NEMS.

HT actually frees the creative planner to attempt to exploit the thought processes of opponents, as does DT. By doing this, Bennett and Dando replaced the original GT definition of player rationality, by substituting one which considers mistakes in valuation of the results, similar to C-4. But the Germans also realized the inherent vulnerability of their plan to French disruption (C-6). Thus, Bennett and Dando have the Germans calculate the probability that the French will defend in the Ardennes and prevent a breakout. By deriving a French behavior probability based on game theory (rather than using a DT approach), Bennett and Dando showed that hypergame rationality had been accomplished by the Germans. Furthermore, the Germans could anticipate a great victory if successful (C-8). But they had to keep the French in the dark (C-5).

Lastly, Bennett and Dando recognized that their model was not truth – the Germans were not guaranteed victory. They pointed out that the German reasoning and French reasoning were rational. They suggested that GT's requirement for a full game does not allow deception, trickery, feints, and lures that cause opponent play to degrade. Yet, they recognized that they had given up GT's guarantee of the value of the game as the minimum expected result. In fact the Germans could have been trounced.

Unfortunately, their hypergame theory is only referenced in the literature three times outside of the original authors. It was essentially ignored.

2.2 Re-emergence of Hypergame Theory

This section will concentrate on the results of Vane's *original* hypergame research. Hypergame theory did not catch on in the 1980's in the DT or GT community. In fact, HT may still be in the intellectual backwaters. But Vane and Lehner began to publish articles [Vane, Lehner 1990] about how HT could be used in military planning. And Vane participated in GTDT Agents Workshop in 1999 in London [Vane, Lehner 2002] which was attached to the Fifth European Conference on Symbolic and Quantitative Approaches to Reasoning about Uncertainty (ECSQARU). Around the time that Vane produced his dissertation, numerous intelligence and military planners were beginning to discuss asymmetric warfare and Osbourne and Rubinstein in "A Course on Game Theory" had penned that "Modeling asymmetries in abilities and perceptions of a situation by different players is a fascinating challenge for future research, which models of 'bounded rationality' have begun to tackle."

HT **is** a model of bounded rationality. HT is a representation of possibility and likelihood with an embedded representation of error. The background material shows how two views of the same game can be incorporated in a hypergame normal form that records player options, adversarial and nature caused situations, predicted outcomes, possible reasoning contexts (even improbable ones), current evidence assessment, and the fears and expectations of the choosing player. This section also includes a number of previous findings on different kinds of strategies that can be used to select plans using an HT representation.

Vane's dissertation research explored 14 random generated four by four games with the following properties: each four by four game had no row or column that was inferior to another row or column. This fact encompasses two truths: (1) that every option for both players was usable. (2) But the options, either row or column were not forced to be in the NEMS for the four by four, full game. Since these were zero-sum, two-player games; the value in the matrix was the value to the Row player (and hence the negative of the value to the Column Player). All sixteen, 3x3, subgames were evaluated for each, 4x4, full game. These games were to represent any *m* x *n* game and all arbitrary μ x ν subgames.

The question was: What could we say about all of Row's possible strategies? The Row player could play game theoretically: the NEMS 3x3 for each subgame, or the NEMS 4x4 for the full game, or a NEMS 4x3 (eliminating the unavailable column); or decision theoretically by choosing a single row. Because all possible 3x3 subgames were evaluated, the goal was to discover the properties of such games with an eye towards generality. In fact the evidence became so predictable that fifteen analytic results were proven [Vane 2000]. Some of the more important ones will be shared in this section. An example of such a game is included in figure 1.

It quickly became my hope that more of the information that needed to be shared with reviewers or peer decision-makers could be appended to the game, as a hypergame normal representation. Looking at the two games above, it was obvious that since they were actually both from the row player's perspective then they could be superimposed. While non-zero sum games were not envisioned in this research, subsequent decision-theoretic additions showed that any column player value system could be accommodated. Thus, the hypergame normal form is a *structured argument* with utility values and beliefs attached.

In the hypergame normal form in figure 2, both games are included in the game results portion (lower right hand). As a reminder NEMS m x n solutions are stated as two vectors, the **c**olumn player's **m**ixed **s**trategy (CMS) and the **r**ow player's **m**ixed **s**trategy (RMS). These are made explicit in hypergame normal form. Notice that the full game, 3x3, and 4x3 CMSs can be recorded as potential estimates of Column's play. Because of later work [Vane, Griffith 2005] these are now considered situational contexts in the upper right hand portion of the hypergame normal form. An informational region (in the adjacent left side of the game) associated with the RMS for each context provides the row player with information on the NEMS solutions for each situational context. These are rarely the choice that Row makes. However, the upper left hand section is used to evaluate current evidence to determine which situational contexts are likely to be in play. This probability vector also sums to one and weights how much each situational context is considered by the row player to be participating in the situational belief (context summary), C_{Σ}. Note that the RMS 3x3 is always less than the RMS 4x3, hence N/A, or not applicable.

	Col. 1	Col. 2	Col. 3
Row 1	-2	-1	3
Row 2	-1	4	-1
Row 3	2	-3	-4

Column Player Subgame

	Col. 1	Col. 2	Col. 3	Col. 4
Row 1	-2	-1	3	-2
Row 2	-1	4	-1	1
Row 3	2	-3	-4	1
Row 4	-2	-2	5	-5

Row Player Full Game

Figure 1. Subgame and Full Game Reasoning

				C_{Σ}	.58	.12	.30	0
	0			**CMS 4x3**	.6	.15	.25	
		1		**CMS 3x3**	.58	.12	.30	-
			0	**CMS 4x4**	.435	.05	.32	.195
MO EU(*,C_{Σ})	**RMS 4x3**	**RMS 3x3**	**RMS 4x4**	**Full Game**	**Col. 1**	**Col. 2**	**Col. 3**	**Col. 4**
0 (-.39)	0	.40	.31	**Row 1**	-2	-1	3	-2
0 (-.39)	.35	.26	.28	**Row 2**	-1	4	-1	1
0 (-.39)	.35	.34	.34	**Row 3**	2	-3	-4	1
1 (.08)	.3	-	.07	**Row 4**	-2	-2	5	-5
.08 (1)	-.27	N/A	-.36 (3)	**EU(*,C_{Σ})**				
-.36 (2)	-.36	N/A	-.36 (4)	**EU(*,WC)**				

Figure 2. Four by Four Hypergame in Normal Form

Hypergame theory includes a simple calculation of the worst case for any option and mixed strategy. The worst case is the least result in the row for any option and the lowest expected value in a mixed strategy evaluation. The maximum of the minimums is what GT proposes to guarantee; it raises the floor of one's expectations and reduces the effects of outguessing. Often GT is useful for pruning options that are mathematically dominated. But the rest of this article has been about how one might model and achieve better results than the best of the worst results (what maximin really means).

The following equations are used to record the options as an *m* rows by *n* columns matrix U, CMS's as probability vectors, RMS's as probability vectors, and beliefs about the applicability of each CMS to the situation as a probability vector P. These beliefs are summarized in C_Σ, also a probability vector.

$$U = \begin{bmatrix} u_{11} & .. & u_{1n} \\ .. & .. & .. \\ u_{m1} & .. & u_{mn} \end{bmatrix}$$

$$CMS_k = [c_{k1}..c_{kn}], 1 = \sum_{j=1}^{n} c_{kj}$$

$$RMS_k = [r_{k1}..r_{km}], 1 = \sum_{i=1}^{m} r_{ki}$$

$$P = [P_0..P_K], 1 = \sum_{k=0}^{K} P_k$$

$$C_\Sigma = [S_1..S_n], 1 = \sum_{j=1}^{n} S_j, S_j = \sum_{k=0}^{K} P_k \cdot c_{kj}$$

In a completed hypergame, C_Σ represents the final accumulation of what the Row player believes about the entire situation – C_Σ sits atop the hypergame as distilled competitive intelligence, factoring in the evidence and the Row player's knowledge. Row calculates which option is best in the Modeling Opponent (MO) column on the far left of figure 2 and uses values in the hypergame to compute the Row player's predictions. Four entries show what results are expected for a modeled opponent: what Row expects (1), the worst case (2), what the full game NEMS expects(3) and what the full game NEMS promised(4). The Row player also knows how much this additional thinking has improved the result expectations by adopting a non-NEMS solution or how much risk we have accepted in our pursuit for a higher than EU(NEMS(m, n)) result.

By tradition the full game NEMS is calculated as CMS_0 because it represents a starting point for this reasoning. However CMS_0 has another purpose as well. When the Row Player can not assign all of the current situational belief to subgame CMS's, then the full game receives the remaining belief (1- sum($P_1..P_K$)). This finding came from the GT definition of full game NEMS(m,n) as a "don't know" condition. Indeed any other assignment of such belief actually favors one option over another, which is counterintuitive to the spirit of uncertainty.

2.3 How Rows and Columns Affect the NEMS Expected Utilities

We should expect that adding new option represented by adding a row in the hypergame always preserves or increases the expected utility for the NEMS of the game.

EU(NEMS(m, n)) ≤ EU(NEMS(m+1, n))

If an opponent ignores/excludes a column (one of the opponent's options), the expected utility also preserves or increases for the NEMS of the new game

EU(NEMS(m, n)) ≤ EU(NEMS(m, n-1))

What can be said about adding *i* rows and subtracting *j* columns? This action is the generalized result of the two steps.

However, *adding just one column or subtracting just one row makes the expected utility of the NEMS unpredictable.*

EU(NEMS(m, n) ?? NEMS(m+i, n+1)

EU(NEMS(m, n) ?? NEMS(m-1, n-j)

The rationale is that the subtracted row may be the most significant factor in preserving the vale of the EU(NEMS(m, n)) or it may be completely insignificant.

Similarly, removing *i* rows or adding a *j* columns lowers the expected utility of the NEMS of a full m x n game.

EU(NEMS(m-i, n)) ≤ EU(NEMS(m, n))

EU(NEMS(m, n+j)) ≤ EU(NEMS(m, n))

Likewise, subtracting just one column or adding one row may change the relationship, requiring an empirical determination.

EU(NEMS(m, n) ?? NEMS(m-i, n-1)

EU(NEMS(m, n) ?? NEMS(m+1, n+j)

2.4 Hypergame Strategy Effectiveness

Before hypergame theory was invented, game theory already provided a solution to such games – the Nash Equilibrium Mixed Strategy (NEMS). The NEMS must be the "floor" for any hypergame strategy (hyperstrategy) to be deemed effective. An effective hyperstategy provides at least the expected utility of the NEMS for the full game.

A hyperstrategy, *H*, is effective if and only if the expected utility of the hyperstrategy and the summary of beliefs, C_Σ, is greater than the expected utility of the full game Nash Equilibrium Mixed Strategy (NEMS(mxn)).

$$Eff(H) \leftrightarrow EU(H, C_\Sigma) \geq EU(NEMS(m \times n))$$

Thus, an effective hyperstrategy incorporates the Row player's belief about the hypergame. Practically, effectiveness was assessed by determining whether:

EU(H, C_Σ) ≥ EU(RMS(m x n),C_Σ) and

EU(H,CMS(m x n)) = EU(RMS(m x n),CMS(m x n)) were satisfied.

If the second term is not true and becomes

EU(H,CMS(m x n)) < EU(RMS(m x n),CMS(m x n))

then the hyperstrategy ceases being effective at some level of uncertainty. *Most hyperstrategies are partially effective.* However, if any hyperstrategy is even partially effective; then the

NEMS solution is actually suboptimal as the original hypergame researchers surmised.

Four types of effective hyperstrategies exist: the *NEMS(m x n)*, a default full game strategy; *modeling opponent*, which is the hyperstrategy that closely approximates decision theory; *pick subgame*, where the Row player chooses a NEMS solution based on all available rows but using less than all available columns; and *weighted subgame*, a blend of modeling opponent and pick subgame hyperstrategies. One of the results of Vane's doctoral dissertation was to show that the weighted subgame hyperstrategy always performs worse than a combination of the other three hyperstrategies based on the player's estimate of uncertainty.

2.5 Hypergame Expected Utility and Uncertainty

A hypergame is a model of belief, but it is not a guarantor of results. This small section reviews the fact that any game theoretic or decision theoretic model can be wrong and how wrong it can be. In the explanation of the hypergame normal form, the reader was introduced to the worst case calculation.

$$HEU(s,g) = EU(s,C_{\Sigma})\cdot(1-g) + EU(s,WC)\cdot(g),$$

where s represents the strategy being plotted and g represents the uncertainty associated with the hypergame (as a normalized probability in the range 0 to 1). C_{Σ} is the summary of belief and *WC* is the worst case result.

See figure three to observe the plot of the HEU formula over uncertainty from zero to one for MO vs. NEMS full game. The shaded area shows an advantage over GT for Row.

HEU creates a mathematical bridge between DT and GT. By incorporating a variable that represents the planner's fear of being outguessed. When the planner has no fear of being outguessed, then the HEU resembles DT. When the planner is in complete fear of being outguessed then the HEU resembles GT. Please notice how the HEU plot shows the current breakpoint for uncertainty of the Row player needed to select the Model Opponent hyperstrategy. Model Opponent hyperstrategies are almost always a single row from the hypergame. When the EU(*, C_{Σ}) of two rows are exactly equal, then the Row player should choose the row with the best worst case.

3. DISCUSSION

The original work was done to explore the properties of hypergame theory, particularly about how adding a "new option" for the Row player or constraining the Column player to a proper subgame affects analytical results. Several new aspects of game theory were discovered, as well as a compact representation – hypergame normal form. Since then, I have been privileged to receive funding for several additional adversarial/competitive research projects.

This section discusses why I consider game theory as a last resort prescription; how HT supports reasoning about luck and option robustness; and lastly how HT might prevent cognitive bias by better representing situational beliefs, evidence monitoring, consensus gathering, and game eliciting. HT can become the defacto standard for structured arguments for developing trust among reasoners or facilitating dialogue about important decisions in groups. HT would ascend to the popularity of Multi-Attribute Utility Analysis (DT) for such activities.

3.1 Game Theory as a Last Resort Prescription

It was discovered that game theory is *seldom a prescriptive theory for play*. Instead it helps us to identify options that are mathematically dominated by combinations of other options. Sometimes this eliminates all other options, but one – where it becomes a prescription by elimination.

In fact any option in the mixed strategy always results in the exact value of the game if played against a player ascribing to the strategy of playing a full game NEMS. Semantically this makes

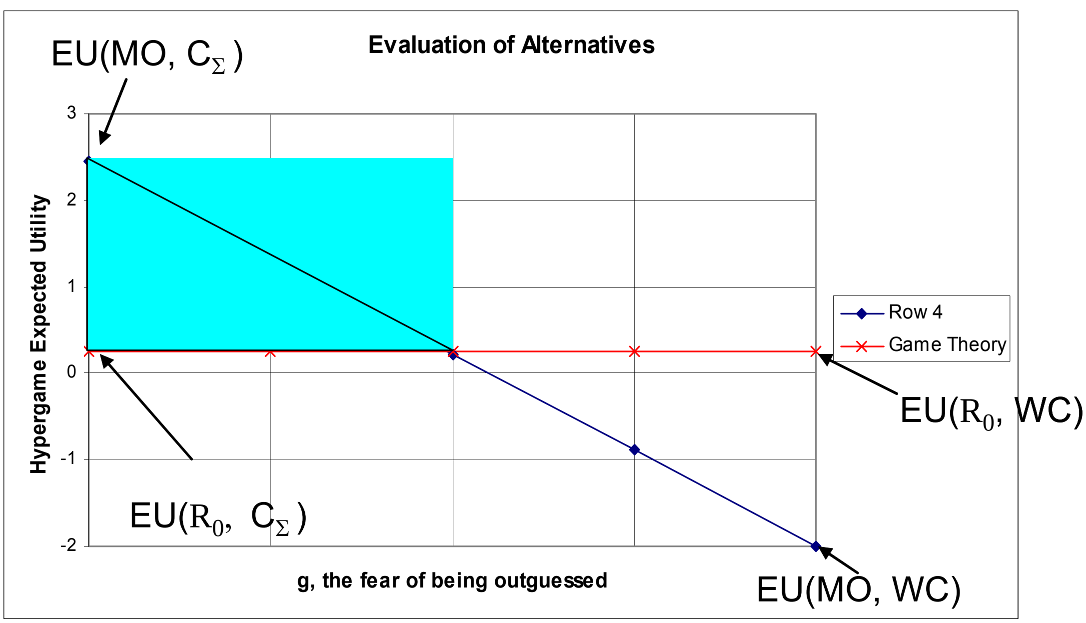

Figure 3. Hypergame Expected Utility, Original

sense. If your opponent is willing to remove his "volitional" decision-making mathematically, then playing any strategy not "ruled out" by NEMS earns the expected utility of the NEMS solution.

One example is of playing the game rock-paper-scissors with equal probability against a child who always plays "rock". Since rock is in the NEMS mix; the child "wins", "draws," and "loses" as often as the GT playing adult. The child could blithely assume that being completely transparent in a strategy game is effective, or that "rock" is a really great strategy. HT offers an alternative of using paper to win more often, that would stimulate the child to learn something about strategies.

3.2 Reasoning about Luck

In figure four, an expanded HEU plot over Vane's dissertation was developed. There are two primary regions based on the x-axis related to certainty that tend towards uncertainty to the right and serendipity to the left. Luck leads us to the best outcomes, while uncertainty should tend towards the worst outcomes. When these are zero, this indicates a complete trust in the hypergame model predicting actual operational results. The range of all of the possible results should be encompassed in this diagram. The range of evaluated outcomes (the y-axis) has a winning region (above the zero-line) and a losing region (below).

There are three strategies plotted in figure four. The current Model Opponent (MO) strategy which may be read left to right as follows: The current hypothesis has respectable utility if Row is lucky during implementation, but one plan beats it if the planner is very lucky. Otherwise, MO is the best choice until the planner is .3 uncertain about being outguessed, whereupon the safe hypothesis and its meager utility outweighs the very large negatives that accrue if the MO plan is outguessed (as shown at the far right). Both the luck breakpoint and risk breakpoint give the planner an estimate of the rationality in using MO or not.

The reader is reminded that 0 point on the x-axis should correspond to EU(s, C_Σ), complete trust of the hypergame as a model of each strategy *s* under the summary belief criteria.

3.3 Reasoning about Option Properties, such as Robustness

Some additional subtleties are introduced in figure five on the next page. In this plot of the HEU for each strategy from figure 4, the actual reliability or robustness of the plan is plotted.

As a first approximation, robust plans are ones with lots of contingencies. This is analogous to plans that have a lot or "OR'd" options. Brittle plans are those that depend on a number of parts of the plan being synchronized or occurring properly in the right order. This is analogous to plans that have a lot or "AND'd" resources and tasks. Straight lines represent those plans that appear to be neither.

Note that this allows the planner to account for really bad outcomes even in the robust, safe plan (heavy dashed line) – when "everything goes wrong." Additionally, the Row Player is aware that there should be more concern about maintaining an accurate estimate of uncertainty, for the current "best hypothesis, MO. The risk breakpoint has moved closer to the zero uncertainty line and the upside potential (luck) appears to be less likely, too.

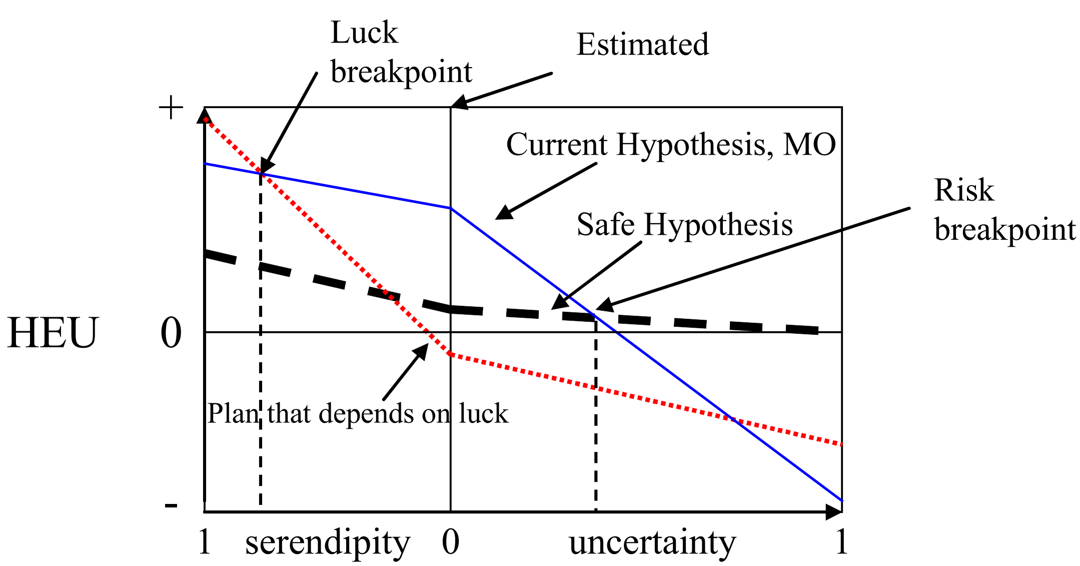

Figure 4. HEU with luck considered

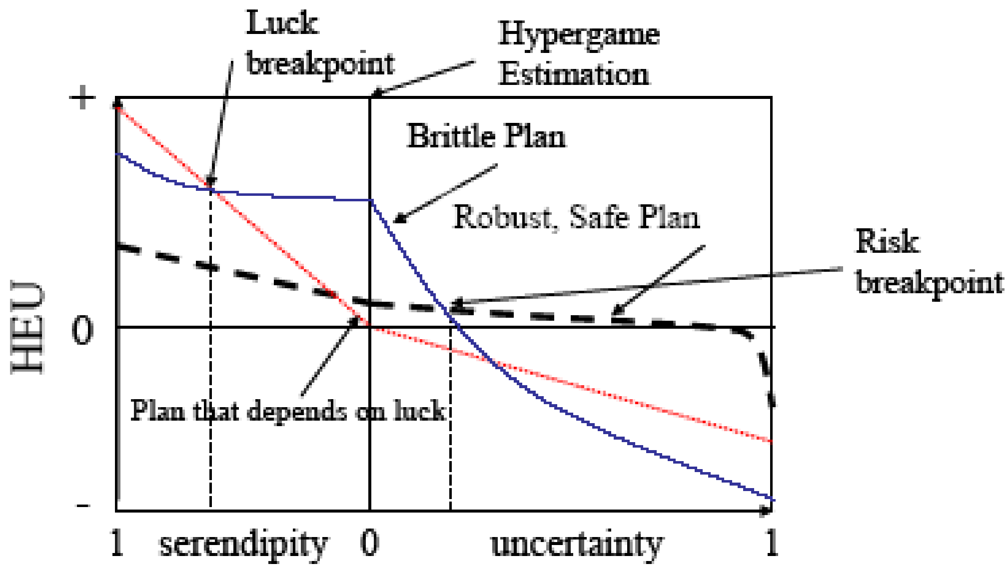

Figure 5. HEU with robustness curves

3.4 Preventing Cognitive Bias and Shortfalls

HT can directly address problems with human cognition: short term memory (often represented as the Miller Number of seven plus or minus two); complex storytelling and future projection (I call the Simon Number – three plus or minus one); and Kahneman's System1, Intuitive and System 2, Deliberative, models of human cognitive processing. The explicit recording of HT parameters aids the short term limitations, while the HT option-situation-context matrices supports enhanced course of action and situational analyses.

It will probably take DTGT agent intelligent augmentation of a human planner to mitigate the shortfalls of cognitive processing, particularly by aiding the intuitive processing. Since this processing is done almost instantaneously based on recognized patterns and slowly accumulated expertise, a lot of computer processing will need to be done to support it. My colleague, Douglas Griffith, is writing about this in his work on neosymbiosis [Griffith 2005, 2006].

3.4.1 Preventing Cognitive Bias and Shortfalls

Since HT allows the planner to record beliefs in an organized way, we believe that this addresses an attempt to make the invisible visible. By recording and labeling contexts, in the above work competing CMS's, HT enables the planners to see where a small change in evidence would result in a large swing in estimated likelihood of some enemy operation or situation.

This allows us to even reason with completely different value matrices for other players and to distill organizational behavior (as DT does) into a new situational context, as another HT CMS.

3.4.2 Monitoring Evidence

During the original hypergame work, no formalized evidentiary reasoning techniques were proposed. They were assumed to exist in the subjective probability assessments assigned in the upper left hand portion of the hypergame. It is expected that more model based approaches that allow larger data streams to be interpreted and recorded 24/7 will help determine the P values in HT. Promising work by Schum, [Schum 2001, 2003] indicates that discovery, invention, and completeness of evidential considerations are enhanced by applying varying evidential mechanisms. Furthermore, software agents may be employed to scrutinize evidence as it arrives in a decision support system, increasing both the speed and breadth of intelligence monitoring before, during, and after operations begin.

3.4.3 Gathering Consensus

We have some experience supporting computer maintained structured arguments that can be applied to a hypergame consensus tool. The idea is to allow members of a decision-making team to register different views and support them with HT. For instance a merger could be proposed by the corporation's Merger's and Acquisitions team for review and discussion via HT on future business opportunities and corporate culture by the entire management team.

3.4.4 Eliciting Hypergames

Hypergame theory can become a strategic prescription that far exceeds the usefulness of game theory. Because it addresses limited intelligence situations, plans can be made that incorporate many of the C-1 to C-8 considerations outlined in the introduction. HT supports knowledge elicitation and an approach to dynamic reasoning that is missing from the declarative aspects of game theory.

One such concept that has emerged is the concept that not only should a planner attempt to create a new, unanticipated option (row); but the planner should not feel that the task is done until the new row's most dangerous anti-plan (often called situation since it is recorded in a column) is also identified. This way new options for Row are not subject to as much confirmation bias. It has been my experience that smart people can invent useful new plans, fall in love with them, and never see the counterstroke. A balanced way to avoid this is to require a counter-plan for every new plan to be added to an HT matrix. Unexpectedly, perhaps, is that there has been a pattern of closure or completeness using this technique. Often, by forcing one to consider nasty counter plans that the process turns up one of the original enemy plans as the almost ideal counter. Often it was the unstated reason why the opponent's column was included originally.

Anecdotally, GT almost always requires a gestational period for the reasoning to settle down to a consistent alignment of beliefs. In fact both sides of the Cold War conflict wanted the other side to know all of their options to generate the kind of "lose-lose" perception needed to prevent military adventurism and nuclear exchange.

4. FINDINGS

HT appears to be a worthwhile, relatively new representation that would benefit from a larger community of practice and that HT would aid those involved in real world reasoning issues. HT's capacity to represent complex considerations such as the feint in the First Gulf War is documented [Vane 2000].

This new work helps decision makers to focus on meaningful intelligence to reduce uncertainty. The actual semantics of an option are used to determine what evidence to collect and when to change strategies. While the aspects of collection are must be based on the vulnerabilities of friendly plans and the properties of current sensor systems.

By supporting the ideas of observe-orient-decide-act mentioned in Boyd {Boyd 1988]. The following findings emerge, by referring to figure five on the previous page:

(1) The who hypergame representation and associated HEU plot meets the observe-orient phase of the problem.

(2) The HEU plot highlights when one should decide (choose a plan) and act or observe-orient again to reduce uncertainty.

(3) As a plan is executed, the continuous observe-orient loops either reduce uncertainty or show that we were fooled. Hence, we can see even in initial feedback whether our modle of the situation is valid to the degree needed.

Unlike GT which requires consistent alignment of beliefs among opponents or DT which does not force a comprehensive evaluation of low probability situations, HT appears to be a good enough approach to search an abstract decision space that can reduce the computation complexity of real world planning.

5. CONCLUSIONS

General Dynamics is implementing a network of Real-time Control System Four (RCS-4) DTGT agents in the Cognitive Agent Architecture (Cougaar) to explore how the sensing, modeling, planning, and particularly judging modules of heterogeneous agents collaborating to solve a difficult problem can be improved by applying HT. I hope to report this work next year.

Lastly, by sharing predicted outcomes between calling and called agents, we expect to fold the search space into layered utility spaces that greatly transforms the domain of planning from movement and constraint space to strategy space. HT naturally supports an idea of confidence from its representation of luck and uncertainty.

6. ACKNOWLEDGMENTS

My thanks to Mr. David Oppenheimer and Mr. Mark Kusiak of GDAIS for supporting this paper and my participation in the DTGT workshop.

7. REFERENCES

[1] Bennett, P.G., Dando, M.R. (1979) Complex Strategic Analysis: A Study of the Fall of France, *Journal of Operational Research Society* **33** 41-50

[2] Griffith, D. (2005) Beyond Usability: The New Symbiosis, *Ergonomics in Design,* **13**, 30-31.

[3] Griffith, D. (2006). Neo-symbiosis: a conceptual tool for system design. Hawaii International Conference on System Sciences, **39**: Kaui, Hawaii.

[4] Heuer, R. (1999) *Psychology of Intelligence Analysis*, Center for the Study of Intelligence

[5] Kahneman, D. 2002. *Maps of bounded rationality: a perspective on intuitive judgment and choice.* Nobel Prize lecture, December 8

[6] Schum, D.A. (2001) Evidential Completeness and Baconian Probability, *Monograph for NASA Langley Research Center*, George Mason University

[7] Schum, D.A. (2003) Boolean Functions and Some Issues in Discovery and Invention, *Monograph for Joint Military Intelligence College*, DIA

[8] Tolcott, M.A., Marvin, F.F, Lehner, P.E. (1989) Expert decision making in evolving situations, *IEEE Transactions on Systems, Man and Cybernetics* **19** 606-615

[9] Vane, R.R, Lehner, P.E. (1990) "Hypergames and AI in Planning", *Innovative Approaches to Planning, Scheduling and Control: Proceedings of a Workshop Held in San Diego, California, November 5-8, 1990*, Morgan Kaufmann

[10] Vane, R.R. (2000) Using Hypergames to Select Plans in Competitive Environments, *Doctoral Dissertation*, George Mason University

[11] Vane, R.R., Lehner, P.E. (2002) Using Hypergames to Increase Planned Payoff and Reduce Risk," *Game Theory and Decision Theory in Agent-Based Systems* (Parsons, Gmytrasiewicz, Wooldridge, eds.), Kluwer Academic

[12] Vane, R.R., Griffith, D. (2005) "Augmenting Cognition for Adversarial Reasoning", *First International Conference for Augmented Cognition*, Las Vegas, NV