

**COMPROMISE REACHING MECHANISMS
IN MULTI-GROUP/MULTI-PLAYER NEGOTIATION PROCESSES**

**Shimon Schocken
Department of Information Systems
Stern School of Business
New York University**

**Robert A. Hummel
Courant Institute of Mathematical Sciences
New York University**

Working Paper Series

STERN IS-93-14

Compromise Reaching Mechanisms
in Multi-Group/Multi-Player Negotiation Processes

*Shimon Schocken*¹

*Robert A. Hummel*²

June 16, 1993

¹ Department of Information Systems, Stern School of Business, New York University, 44 W. 4th Street, New York, NY 10003. Internet: schocken@stern.nyu.edu.

² Courant Institute of Mathematical Sciences, New York University, 251 Mercer Street, New York, NY 10012. Internet: hummel@cs.nyu.edu

Compromise Reaching Mechanisms in Multi-Group/Multi-Player Negotiation Processes

We consider a situation in which multiple decision-makers who are partitioned into two or more distinct groups are asked to recommend a uniform course of action which is drawn from a finite and explicit set of potential alternatives. We present group-level and player-level mechanisms to reach a compromise decision under such circumstances. The group-level mechanism is based on the Dempster-Shafer theory of evidence; The player-level mechanism employs a set-product operation that aggregates the individual decisions over a certain space of committees. Previous research established that the two mechanisms are isomorphic, which, in the context of the present paper, implies that they yield the same compromise decision. However, unlike the Dempster-Shafer theory, which was criticized for lack of external validity, the set-product mechanism has plausible properties in the context of group decision making. With that in mind, the paper seeks to (i) report about an interesting relationship between group decision research and AI methods to manage uncertainty, and (ii) build a foundation for an inter-disciplinary research that exploits this linkage.

1 Introduction

Negotiation is an iterative and interactive process in which two or more decision-makers, henceforth referred to as *players*, revise their initial decisions as they obtain new information and go through successive negotiation rounds. The process ends when all players decide to stick to their present decisions, which are then said to be *final*. Assuming that a consensus decision has not been reached, and that the group *must* select or recommend one course of action, there are two possible ways to resolve the final stage of the negotiation process. The group can either choose one of the final decisions and forsake all others, in which case one player has ‘won’ the process, or it can seek a compromise decision that takes into consideration all the players’ final decisions, in one way or another. The latter strategy characterizes a cooperative group whose “decision processes involve *unitary* tasks that usually require the combination of all group members’ outputs into a single output” (Morrison, Morrison, & Vogel, 1992). Needless to say, such a compromise reaching mechanism must be fair and plausible, or else some players will refuse to use it or abide to its outcome. Further, a mechanism that does not lend itself to a logical justification will make it difficult for the group to defend the compromise decision if and when the need arises.

Put figure 1 around here

The various stages of the decision process that we postulate are depicted in figure 1. Several methodologies and systems were developed to support such cooperative group processes. For example, Bui (1987) has integrated a variety of methods, including an extended version of ELECTRE, Delphi and Nominal Group Techniques, and computer-based conferencing

and mediating tools, to create a decision support system for cooperative groups, called Co-oP. By eliciting aspiration and reservation levels from group members, Lewandowski (1989) has calculated proxi importance weights which he then used to support cooperative committee decision making. These ideas were implemented in the SCDAS system, which also helps individual decision-makers to assess degrees of disagreement and distance from the aggregated group decision. It is important to emphasize at the outset that unlike these approaches, the present paper focuses only on the very last stage in the negotiation process – reaching a compromise – whereas the other stages are mentioned here in order to put the paper’s contribution in a proper perspective. In a nutshell, then, our objective is to describe a plausible mechanism for generating a group compromise decision, as the very last stage of an iterative and possibly long negotiation process. Further, we constraint our analysis to a subset of multi-player decision processes that have the following characteristics:

1. All players face the same exhaustive and mutually-exclusive set of alternatives, among which one or more alternatives must be chosen.
2. Each player makes a decision by selecting a subset of alternatives. A singleton subset is interpreted as “the single alternative in the subset must be pursued,” whereas non-singleton subsets are interpreted as “one of the alternatives in the subset must be pursued.” At the final stage of the negotiation process (see figure 1), each player is characterized by a single decision, i.e. a final subset of alternatives.
3. The players fall into different groups that are characterized by different objectives, constrains, affiliations, sources of evidence, or any other property that serves to partition the players into distinct categories in a non-ambiguous way.

Examples of decision settings that are consistent with the above assumptions abound in practice. For instance, the players can be physicians, executives, or professors. The set of possibilities can represent different diagnoses, products, or candidates for a dean position, respectively. The different groups into which the players fall can be medical specialties, business divisions, and academic departments, respectively. In each of these cases, it is assumed that a large group of decision-makers is asked to choose, or recommend, one or more courses of action out of many possibilities. The choice can be highly focused, as in “we should make an offer to this particular candidate,” or it can be less committed, as in “we should develop either product A or product B.”

We see that the decision tasks that we focus on consist of *discrete alternative problems* – problems in which the players face an explicit and fixed set of possible alternatives (Ziont’s, 1989). The players are not restricted to choosing a single alternative, though, making the process akin to Brams & Fishburn’s (1992) *Approval Voting* concept, in which members of a group can vote for as many candidates as they like in a multicandidate election. Other related research includes the work of Oral, Kettani, and Lang (1991), whose decision support methodology is designed to assist groups in selecting a subset from a finite list of alternatives.

Formally, we assume that a group of n players must select one or more alternatives from an exhaustive and mutually-exclusive set of alternatives, denoted θ . The players belong to m groups, denoted $g_1 \dots g_m$, each consisting of n_i players, $i = 1, \dots, m$, such that $n = n_1 + \dots n_m$. The m groups are characterized by different objectives and sources of information regarding θ . After evaluating the available information, each player selects independently a subset of θ which, in his opinion, contains the best possible course of

action. At that point, a negotiation process ensues, during which some players revise their initial decisions. When the process terminates, the group faces a set of n final individual decisions, or subsets of θ , from which a group decision must be selected. How can the group reach a compromise decision under such circumstances? For practical as well as analytic reasons, we divide the answer into two variants, which are depicted symbolically in figure 2:

- Group-level compromise
- Individual-level compromise

Put figure 2 around here

The group-level mechanism (top of figure 2) assumes that each group conducts its deliberations independently, producing a single group decision. That is, once the group deliberations are over, the individual opinions of the group members are replaced with a single group decision which is broadcast to the other groups, at which point a negotiation round takes place. In the final stage of the process, a compromise reaching mechanism is applied to the final m group decisions, producing a compromise decision.

The individual-level mechanism (bottom of figure 2) assumes that the various players interact with each other directly and across group boundaries, using a certain committees structure that will be discussed shortly. In the final stage of the negotiation process, a compromise reaching mechanism is applied to the entire gamut of n individual decisions. Although the mechanism takes into consideration the players' different group affiliations,

it operates over the larger space of individual, rather than group, decisions. One attractive feature of this mechanism is that individual opinions are tracked throughout the process, rather than replaced with group opinions, a process which entails a loss of information.

Both compromise reaching mechanisms are subtly related to the Dempster-Shafer theory of evidence. This theory concerns the elicitation and manipulation of degrees of belief rendered by multiple sources of evidence to a common set of propositions. Recently, Hummel and Manevitz (1992) observed that the evidential reasoning that the Dempster-Shafer model attempts to manage is characterized by two types of uncertainty: *intrinsic* and *extrinsic*. In the context of the present paper, these correspond to the notions of *players* and *groups*, respectively. The source of *intrinsic* uncertainty is pluralism, or diversity of opinions: when two or more players are presented with the same information, the conclusions that they draw may well be different. In contrast, *extrinsic* uncertainty occurs because different groups, by virtue of their different affiliations, training, etc., may also lead to different conclusions. The Dempster-Shafer theory can be viewed as an attempt to manage both types of uncertainty simultaneously – a point which is taken up later in the paper.

The paper builds on previous work on uncertainty management in AI. The set-product mechanism that we describe was used by Hummel & Landy (1988) in their probabilistic analysis of Dempster's rule. This work was continued by Hummel & Manevitz (1992), who extended it to an axiomatic foundation for probabilistic reasoning in AI. On the applied side, Schocken & Pyun (1990) presented a Dempster-Shafer relevance calculus with a multi-player semantics, where the players corresponded to searchers and catalogers of keyword-based information. This work was extended considerably by Schocken & Hummel (1993), who used the multi-player semantics to expose several limitations in the Dempster-Shafer

model and propose ways to resolve them. The present paper makes the first attempt to link this analytic work to the applied context of negotiations and group decision making – creating a bridge between two rather remote areas of research.

The plan of the paper is as follows. The next section specifies the *decision space* from which players are asked to draw their individual decisions. This sets the stage for the next section, which describes several means to measure the *intensity* of those decisions. The next two sections describe group-level and individual-level compromise reaching mechanisms. A discussion section concludes the paper with implications and future research directions.

2 The Decision Space

Suppose that a group of stock market analysts is asked to predict tomorrow's stock market direction. To keep the example simple, we take the set of alternatives (in this example – predictions) to be $\theta = \{\text{up, same, down}\}$, representing the three mutually-exclusive propositions “*the market will go up more than 5 points*”, “*the market will remain at today's index plus/minus 5 points*”, or “*the market will go down more than 5 points*”. We define the *decision space* to be the power-set of θ minus the empty set, which in this case is $2^\theta = \{\{\text{up}\}, \{\text{same}\}, \{\text{down}\}, \{\text{up, same}\}, \{\text{up, down}\}, \{\text{same, down}\}, \{\text{up, same, down}\}\}$.

Each element in the decision space 2^θ is a possible decision, which is interpreted as a disjunction of alternatives. Thus, an analyst would select the decision $\{\text{up, same}\}$ if she believes that the market will either remain the same, or will go up (which is the same as saying “*the market will not go down*”), whereas the decision $\{\text{up}\}$ characterizes an analyst

who is not only more bullish, but also more confident in his prediction than the previous analyst. As a rule, since decisions are interpreted as disjunctions of alternatives, the smaller is the number of the disjuncts in a decision (viz: the cardinality of the subset), the higher is the confidence of the holder of that decision. Thus, singletons represent highly focused decisions, whereas θ – the maximal element in 2^θ – represents the “least committed” decision. For example, the decision {up, same, down} stands for the vacuous forecast “the market will either go up, down, or remain the same” – a valid but obviously useless opinion. Actually the opinion is not completely useless, because it indicates that its holder was unable to make a more focused decision. If the player happens to be an experienced analyst, it might mean that more discerning information must be gathered before a reasoned prediction can be made.

Also note in passing that even though the number of potential decisions in 2^θ grows exponentially with the cardinality of θ , the semantics of the decision process will usually render most of these decisions as non-sensical. For example, it might be argued that the decision {up, down} makes no sense in a market analysis scenario, in which case the decision can be eliminated from the decision space available to the analysts. In any given application, domain-specific arguments can be used to prune out numerous decisions (arbitrary subsets of θ) that imply inconsistent or otherwise unreasonable courses of action for the group, keeping the *effective* decision space (set of options available to the players) much smaller than 2^θ .

We conclude this section with another example. in which θ stands for a set of potential candidates for a job, and 2^θ represents different groupings of candidates. In this case, elements of 2^θ can represent Hispanic candidates, female candidates, candidates with Ph.D.’s,

candidates from Indianapolis, or simply the set of candidates Bill, Lisa, and Jane. In a typical hiring situation, different groups or players in the organizations will select different subsets from this power-set, and these decisions will ‘shrink’ and ‘expand’ considerably as the negotiations process unfolds (similar to VanGundy’s (1988) characterization of *convergent* and *divergent* group decision processes). However, when all is said and done, the organization will have to make an offer to a *single* candidate, at which point a compromise reaching mechanism may be required.

3 Decision Intensity Measures

Having discussed the space from which individual decisions are drawn, we now turn to describe several means to estimate the degrees of support that such decisions gain in the group process. Let G be a group of n players, each making an individual decision by selecting a subset of alternatives from θ . We take the *group decision* to be a collection of pairs of the form:

$$d_G = \{(X, m) | X \in 2^\theta \text{ and } 0 \leq m \leq 1\} \quad (1)$$

In each pair, X represents a particular decision (a subset of θ), and the m parameter that accompanies it represents the *intensity* of that decision, which varies from 0 to 1. Thus, the intensity is essentially a mapping $m : 2^\theta \rightarrow [0, 1]$ and the m values mentioned in (1) are shorthand for $m(X)$. We deliberately leave the semantics of intensity at a somewhat abstract level, because different applications can call for different definitions which will then require different intensity calculi. However, it is quite natural to begin with a probabilistic interpretation of intensity – either subjective or frequentist. A subjective interpretation

views $m(X)$ values as degrees of support that vary from 0 (X is an impossible decision) to 1 (X is an inevitable decision). A frequentist interpretation can view $m(\cdot)$ as some statistic of the spread of decisions within the group. In particular, we consider the following definition of $m : 2^\theta \rightarrow [0, 1]$:

$$m(X) = \begin{cases} \#X/n & X \neq \emptyset \\ 0 & X = \emptyset \end{cases} \quad (2)$$

where $\#X$ is the number of players in the group whose decision is X

Under this interpretation, $m(X)$ is the fraction of the players in the group whose individual decisions were X . This fraction can serve as an estimate of the probability that a player who is randomly sampled from that group will hold the decision X .

Of particular interest is the intensity of the vacuous decision – $m(\theta)$, i.e. the fraction of players whose decision was “choose any one of the available alternatives.” This value can be viewed as a first-order measure of a group’s indecisiveness. For example, if $m(\theta) = 1$, not even one player in the group was willing to make a concrete decision. On the other extreme, $m(\theta) = 0$ characterizes a group in which all players made decisions that had at least *some* discerning power. The distinction is a bit tricky because even though $m(\theta) = 1$ represents a strictly indecisive group, $m(\theta) = 0$ does not necessarily imply a decisive group. For example, let $\theta = \{\text{Nicks, Bulls, Suns, Sonics}\}$, and suppose that at the beginning of the NBA playoff a group of sports writers decides unanimously that this year’s champion will be either the Nicks, the Bulls, or the Suns. The group will be characterized by $m(\theta) = 0$,

but it is hardly decisive; For that reason, we view $m(\theta)$ as merely a first-order measure of group indecisiveness.

Returning to our stock market example, where $\theta = \{\text{up, same, down}\}$, consider a research analysts group which is required to produce a buy/sell/hold recommendation for the company's stockbrokers regarding, say, investments in index-based securities. We assume that the group consists of eight analysts, which, for the sake or future reference, are referred to it as *group number 1*. Suppose that the group dynamics were such that four players predict that the market will remain more or less the same, one player predicts that the market will go down, two players predict that the market will either remain the same or will go down, and one player refuses to commit himself to any concrete prediction. Using (3), group number 1 will be characterized by the following $m(\cdot)$ function:

$$\begin{aligned}
 m_1(\text{same}) &= 4/8 \\
 m_1(\text{down}) &= 1/8 \\
 m_1(\text{same, down}) &= 2/8 \\
 m_1(\theta) &= 1/8
 \end{aligned}
 \tag{3}$$

Now, taken as support function, $m(X)$ is kind of "hollow," because it measures only the fraction of players whose decision was *exactly* X . Recalling that X represents a *disjunction* of alternatives, e.g. $\text{up} \vee \text{same}$, it is reasonable to argue that players who made more focused decisions like up or same (as singletons) provide indirect support to $\text{up} \vee \text{same}$ as well. Likewise, it can be argued that players whose decision was Y where $X \cap Y \neq \emptyset$ also provide a potential support to X . For example, if $X = \{\text{up, same}\}$ and $Y = \{\text{same, down}\}$ ¹, players who decide Y essentially say that they cannot discriminate between same and down ,

¹We intentionally use the proposition $a \vee b$ and the subset $\{a, b\}$ interchangeably, to remind the reader that the former is the semantic interpretation of the latter.

i.e., that their support floats somewhere in the subset Y , but that they cannot specify exactly *where*. In the extreme case, the players' support can concentrate exclusively on the alternative *same*, in which case they also end up supporting (indirectly) the decision $X = \{\text{up, same}\}$.

Hence, we see that using the function $m(\cdot)$ as raw input, we can compute some 'high-level' measures of support. In particular, in addition to tracking the fraction of players who decided X *exactly* (that is what $m(\cdot)$ does), we wish to track the fraction of players whose decisions were subsumed by X , and the fraction of players whose decisions overlap X . These two measures are denoted $\text{Bel}(X)$ and $\text{Pl}(X)$, respectively, and are defined as follows:

$$\text{Bel}(X) = \sum_{A \subseteq X} m(A) \quad (4)$$

$$\text{Pl}(X) = \sum_{A \cap X \neq \emptyset} m(A) \quad (5)$$

In words, the $\text{Bel}(X)$ function measures the *total support* that X has received from the group, whereas $\text{Pl}(X)$ measures the upper-bound of that function. In the case of the support function given by (3), definitions (4) and (5) yield the following measures:

$$\begin{aligned} \text{Bel}_1(\text{same}) &= 4/8 \\ \text{Bel}_1(\text{down}) &= 1/8 \\ \text{Bel}_1(\text{same, down}) &= 2/8 + 4/8 + 1/8 = 7/8 \\ \text{Bel}_1(\theta) &= 1/8 + 4/8 + 1/8 + 2/8 = 1 \end{aligned} \quad (6)$$

$$\begin{aligned}
Pl_1(\text{same}) &= 4/8 + 2/8 + 1/8 = 7/8 \\
Pl_1(\text{down}) &= 1/8 + 2/8 + 1/8 = 4/8 \\
Pl_1(\text{same, down}) &= 2/8 + 4/8 + 1/8 + 1/8 = 1 \\
Pl_1(\theta) &= 1/8 + 4/8 + 1/8 + 2/8 = 1
\end{aligned}
\tag{7}$$

Readers who are familiar with the Dempster-Shafer theory of evidence have probably recognized $m(X)$, $Bel(X)$, and $Pl(X)$ as the theory's mass, belief, and plausibility functions, respectively. The novel aspect of the present analysis is that it uses a plausible semantics – the multi-player decision setting – to justify these functions by construction. Note that since $Bel(X)$ and $Pl(X)$ are completely defined in terms of $m(X)$, it is sufficient to continue the analysis by focusing only on the latter function.

4 Group-level Compromise

The stock market example was based on the assumption that all the research analysts were members of a homogeneous group. In reality, this is rarely the case. For example, it might be that different groups of analysts have access to different sources of information – a situation which occurs frequently in Wall Street. Specifically, let us assume that our second group of analysts consists of five players, of whom two think that the market will either go up or remain the same, while three think that the market will either go down or remain the same.

Using definition (3), the decision of group number 2 is characterized by:

$$\begin{aligned} m_2(\text{up, same}) &= 2/5 \\ m_2(\text{same, down}) &= 3/5 \end{aligned} \tag{8}$$

To recapitulate, we now have two groups of analysts who try to predict the market's direction. In both groups, all the players draw their decisions from the common set of possibilities $\theta = \{\text{up, same, down}\}$. After deliberating their decisions *within* each group, two group decisions emerge, as given by (3) and (8). Taken as a whole, however, the two groups of analysts must endorse a single prediction, or recommendation, which will then be broadcast to the company's brokers. What should the research analysts do? At the group level, they can either adopt one group decision and ignore the other, or they can seek a compromise decision, i.e., a decision that takes both groups into consideration.

One can think of many different mechanisms to reach such a compromise. For example, the two groups can be simply collapsed together, and the combined group of $8+5=13$ analysts can use (3) to compute their pooled decision. However, such a mechanism will be quite naive, because (i) it will not utilize the group differences; and (ii) it will not allow the groups to interact in any meaningful way. Hence, a more sophisticated mechanism is called for. One such mechanism is available in the form of Dempster's rule, which can be used to combine $m(\cdot)$ functions that come from different sources. Formally, Dempster's rule computes a pooled mass function $m = m_1 \oplus m_2 : 2^\theta \rightarrow [0, 1]$ as follows:

$$m'(X) = \sum_{\forall A, B \subseteq \theta \text{ s.t. } A \cap B = X} m_1(A) \cdot m_2(B), \quad (9)$$

$$m(X) = \begin{cases} \frac{1}{1-m'(\emptyset)} \cdot m'(X) & X \neq \emptyset \\ 0 & X = \emptyset \end{cases} \quad (10)$$

In words, (9) computes the overall support that a subset X collects by summing up the products $m_1(A) \cdot m_2(B)$ over all the subsets $A, B \in 2^\theta$ whose intersection gives X . One adverse side-effect of this computation is that the pooled group decision can contain the result $m'(\emptyset) > 0$, which is disallowed by m 's definition (3). This anomaly is corrected by (10), which normalizes $m'(X)$ in such a way that $m(\emptyset)$ becomes 0 by definition and all the other $m(X)$ for $X \neq \emptyset$ sum up to 1. The need for this manipulation can be justified by our multi-player semantics, which views $m(X)$ as the fraction of players whose decision was X . Recalling that the players were not allowed to choose the decision \emptyset , we see that $m(\emptyset)$ must be 0 by definition. Therefore, if we wish the $m(\cdot)$ functions defined by (3) to be closed under Dempster's rule, we must normalize them accordingly.

To illustrate the operation of Dempster's rule in the context of our stock market analysis, consider its application to the group decisions described by (3) and (8). The computation is illustrated through an intersection table, following Shafer (1976):

	$m_1(\text{up, same}) = 2/5$	$m_1(\text{same, down}) = 3/5$
$m_2(\text{same}) = 4/8$	$m'(\text{same}) = 8/40$	$m'(\text{same}) = 12/40$
$m_2(\text{down}) = 1/8$	$m'(\emptyset) = 2/40$	$m'(\text{down}) = 3/40$
$m_2(\text{same, down}) = 2/8$	$m'(\text{same}) = 4/40$	$m'(\text{same, down}) = 6/40$
$m_2(\theta) = 1/8$	$m'(\text{up, same}) = 2/40$	$m'(\text{same, down}) = 3/40$

$$\begin{aligned}
m'(\text{same}) &= 8/40 + 12/40 + 4/40 = 24/40, \\
m'(\text{down}) &= 3/40, \\
m'(\text{same, down}) &= 6/40 + 3/40 = 9/40, \\
m'(\text{up, same}) &= 2/40, \\
m'(\emptyset) &= 2/40.
\end{aligned} \tag{11}$$

And, after multiplying by $\frac{1}{1-m'(\emptyset)} = 20/19$:

$$\begin{aligned}
m(\text{same}) &= 24/40 \cdot 20/19 = 24/38, \\
m(\text{down}) &= 3/40 \cdot 20/19 = 3/38, \\
m(\text{same, down}) &= 9/40 \cdot 20/19 = 9/38, \\
m(\text{up, same}) &= 2/40 \cdot 20/19 = 2/38, \\
m(\emptyset) &= 0,
\end{aligned} \tag{12}$$

Verbally, the compromise decision can be stated as follows: “*after reviewing the market conditions, our research analysts believe quite strongly that tomorrow’s market will either remain the same or will go down. Recommendation: stay foot*”.

Note that (9-10) is commutative and associative. Hence, the formulae can be used to reach a compromise decision among more than two groups, in any desired order. Therefore, the compromise reaching mechanism can be applied in a cumulative fashion, accommodating new group decisions as they become available. In spite of these desired properties, though, Dempster's rule is considered a highly controversial operator, and many have tried to justify it on mathematical as well as on cognitive grounds (Baron, 1987, Kyburg, 1987, Lindley, 1987, Schocken and Kleindorfer, 1989, and Zadeh, 1986 – to give only a few representative examples of this line of research). As we will see shortly, one advantage of the multi-player semantics is that unlike the above interpretations, it can be used to justify Dempster's rule in a remarkably intuitive and simple way, in the context of certain applications.

5 Individual-level Compromise

The previous section described a *group-level* compromise mechanism. In that scenario, each group deliberated independently to produce a group decision, and the group decisions were then combined to produce a global compromise. This section describes a *player-level* compromise mechanism that operates over the individual decisions of the group members, rather than over the summarized group decisions. The two mechanisms have a rather different dynamics, but as we will see later they are tightly interrelated.

We continue with our stock market example and the same groups of analysts described in the previous section. As we assumed before, each player is presented with the same set of possibilities $\theta = \{\text{up, same, down}\}$. However, we now add a subtle twist to the decision elicitation process. Instead of asking the players to *select* subsets of 'good,' or 'likely,'

alternatives from θ , we ask them to *rule out* subsets of ‘bad,’ or ‘unlikely,’ alternatives. Two restrictions are placed on the players’s decisions. First, they must rule out alternatives categorically. Second, the players are not allowed to rule out *all* the alternatives.

For notational convenience, we represent each individual decision as a binary vector in which 0 codes that the respective alternative has been ruled out by the player and 1 otherwise. For example, given that $\theta = \{\text{up, same, down}\}$, the decision $(0, 1, 1)$ codes that the up direction has been ruled out by the player, implying the player’s opinion that the market will either remain the same, or will go down. Since the player’s decision must imply at least one ‘good’ alternative, the decision $(0,0,0)$ is disallowed, and the state of insufficient reason, or total ignorance, is modeled by the decision $(1,1,1)$. This particular decision is consistent with a player who is unable to rule out *any* alternative.

With this notation, the individual decisions of the players in group numbers 1 and 2 of the previous section are tabulated at the left side of figure 3. As the tables indicate, in the first group four analysts rule out up and down, one analyst rules out up and same, two analysts rule out up, and one analyst rules out no alternatives. Similarly, in the second group two analysts rule out down and three analysts rule out up. Given that all players announce their decisions as final, how can we reach a compromise decision that takes all the individual decisions into consideration?

As in the group-level case, one can think of a variety of such multi-player compromise reaching mechanisms. The mechanisms will differ in terms of players matching rules and decision combination operators, but their overall goal will be the same: reducing a multitude of possibly conflicting and redundant decisions into a uniform group decision. As a

rule, though, we believe that the only mechanisms which are worth researching are those that have (i) a solid normative justification, and (ii) a practical face validity. Figure 1 illustrates the operation of one such mechanism, the *set-product* model – which is based on Hummel & Landy’s (1988) probabilistic interpretation of Dempster’s rule. This simple model can be viewed as a point of departure towards developing more sophisticated mechanisms that can be customized for different group decision situations.

Put figure 3 around here

The set-product mechanism implements what may be termed a *cartesian consensus operation*. First, one constructs all the committees of two that can be formed by matching each analyst from the first group with each analyst from the second group. In the above example, there are $8 \times 5 = 40$ such pairs of analysts. The decision of each committee is then taken to be the Boolean conjunction of the decisions of the committee’s members. For example, the two group 1 analysts whose individual decisions were (0,1,1) and the two group 2 analysts whose individual decisions were (1,1,0) yield four committees whose joint decision is (0,1,0) – the fifth tuple in the set-product table. Since the only thing that distinguishes one committee from another (in this particular model) is the committee’s joint opinion, the committees that have identical opinions are collated, leading to the third column in figure 3. Finally, note that the Boolean conjunction of valid decisions may well generate an invalid decision, which happens whenever the operator produces vectors of the form (0,0,0). This is an anomaly, since null decisions are disallowed in our model. One way to resolve the problem is to simply disregard all the committees that produced null decisions, leading to the right-most and final table in figure 3.

The set-product compromise mechanism is of special interest to us because of its unique relationship to the Dempster-Shafer model. Specifically, we have already pointed out that the two left-most tables in figure 3 induce the mass functions m_1 and m_2 , as given earlier in the paper in (3) and (8). As the reader can verify, the right-most table in figure 3 – the normalized outcome of the set-product operation on the above two tables – induces the mass function $m = m_1 \oplus m_2$ given by (12), where \oplus stands for the standard Dempster rule. This is no coincidence: if we denote the mass function that captures a set of decisions in a group G by m_G , the set-product compromise mechanism by \otimes , and Dempster’s rule by \oplus , we have the following isomorphism: $m_{G_1 \otimes G_2} = m_{G_1} \oplus m_{G_2}$.

6 Discussion

Proofs of the isomorphism result given above can be found in Hummel and Landy (1988), where it was established in the context of a general probabilistic analysis, and in Schocken and Hummel (1992), where it was established in the applied setting of an information indexing and retrieval application. As the previous section illustrated, the isomorphism is particularly compelling in the present context of a multi-player decision process: unlike the rather cryptic Dempster-Shafer model – which is normally justified on the basis of heuristic grounds – the components of the multi-player model are straightforward, intuitive, and most importantly – explicit. Further, the set-product mechanism has several desired properties in the context of group decision making.

First, recall that the compromise reaching mechanism that we sought was supposed to combine individual decisions that are drawn from two distinct groups consisting of equally-

qualified players (this restriction can be lifted, leading to weighted versions of our compromise calculus). Thus, we had assumed that each and every one of the players in both groups was supposed to influence the compromise decision with the same impact. Under such circumstances, the set-product mechanism provides the largest possible set of pairings of players; any other combination scheme will restrict the set of pairings based on some additional structure which was not part of our original assumptions. Further, the set-product mechanism has the following advantages:

- It is commutative, treating each group of players on an equal basis;
- It is associative, so that the order and precedence of the compromise reaching process have no influence on the final outcome;
- Each player in each group has an equal role, and is not thwarted by pairing with an unrepresentative set of players from other groups;
- The Boolean conjunction operator that is used to produce the decision of each committee is plausible, because it rules out the decisions that were ruled out by *both* members of the committee.

To conclude, the contribution of this paper can be seen from two different perspectives – that of the Dempster-Shafer model, and that of multi-player negotiations theories. Beginning with the former perspective, we observe that the set-product mechanism, along with the isomorphism result, imply that the Dempster-Shafer model can be viewed and justified as a means to resolve different decisions expressed by domain experts that are

partitioned into different groups. From a negotiations research perspective, we have presented a particular compromise reaching mechanism which happens to be isomorphic to the Dempster-Shafer model.

This reservation is quite important, because different assumptions about the nature of the negotiation process may require different ways to summarize and combine the decisions of the individual players. For example, the players may be given different weights, the individual decisions can be allowed to be probabilistic rather than Boolean, and the committees structure can be altered or constrained in many different ways. In general, one can envision a whole parametric gamut of multi-player/multi-group compromise reaching mechanisms, of which the Dempster-Shafer model is merely one special instance. In particular, the different objectives of the various groups – an important consideration which was not part of the present analysis – should be taken into consideration in formulating such mechanisms. This would entail an ambitious research program that integrates two rather remote areas of research – group decision theory, on the one hand, and AI-based methods to manage uncertainty, on the other. We hope that this paper provides the first step in that direction.

References

- [1] J. Baron. Second-order probabilities and belief functions. *Theory and Decision*, 22, 1987.
- [2] S.J. Brams and P.C. Fishburn. Approval voting in scientific and engineering societies. *Group Decision and Negotiation*, 1(1):41-55, 1992.
- [3] T.X. Bui. *Co-oP - a Group Decision Support System for Cooperative Multiple Criteria Group Decision Making*. Berlin: Springer-Verlag, 1987.
- [4] R.A. Hummel and M.S. Landy. A statistical viewpoint on the theory of evidence. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 10(2):235-247, 1988.
- [5] R.A. Hummel and L. Manevitz. Foundations of reasoning with uncertainty. Technical report, Courant Institute of Mathematical Sciences, New York University, 1992.
- [6] H.E. Kyburg. Bayesian and non-Bayesian evidential updating. *Artificial Intelligence*, 31:271-293, 1987.
- [7] A. Lewandowski. SCDAS - decision support system for group decision making: Decision theoretic framework. *Decision Support Systems*, 5(4):403-423, 1989.
- [8] D.V. Lindley. The probability approach to the treatment of uncertainty in artificial intelligence and expert systems. *Statistical Science*, 2(1):17-24, 1987.
- [9] J. Morrison, M. Morrison, and D. Vogel. Software to support business teams. *Group Decision and Negotiation*, 1(2):91-116, 1992.
- [10] M. Oral, O. Kettani, and P. Lang. A methodology for collective evaluation and selection of industrial R&D projects. *Management Science*, 37(7):871-885, 1991.
- [11] S. Schocken and R.A. Hummel. On the use of the Dempster-Shafer model in information indexing and retrieval applications. Technical report, Information Systems Department, Stern School of Business, New York University, 1992. STERN IS-92-27.
- [12] S. Schocken and P. R. Kleindorfer. Artificial intelligence dialects of the Bayesian belief revision language. *IEEE Transactions on Systems, Man, and Cybernetics*, 19:1106-1121, 1989.

- [13] S. Schocken and J. Pyun. A Dempster-Shafer model of relevance. In *Proc. of the 23rd Hawaii Int. Conf. on System Sciences*, pages 544–551, 1990.
- [14] G. Shafer. *A Mathematical Theory of Evidence*. Princeton University Press, 1976.
- [15] A.B. VanGundy. *Techniques of Structured Problem Solving*. New York: Van Nostrand Reinhold, 1988.
- [16] L.A. Zadeh. A simple view of the Dempster-Shafer theory of evidence and its implication on the rule of combination. *The AI Magazine*, pages 85–90, 1986.
- [17] S. Zionts. Multiple criteria mathematical programming: an updated overview and several approaches. In B. Karpak and S. Zionts, editors, *Multiple Criteria Decision Making and Risk Analysis Using Microcomputers*. Berlin: Springer Verlag, 1989.

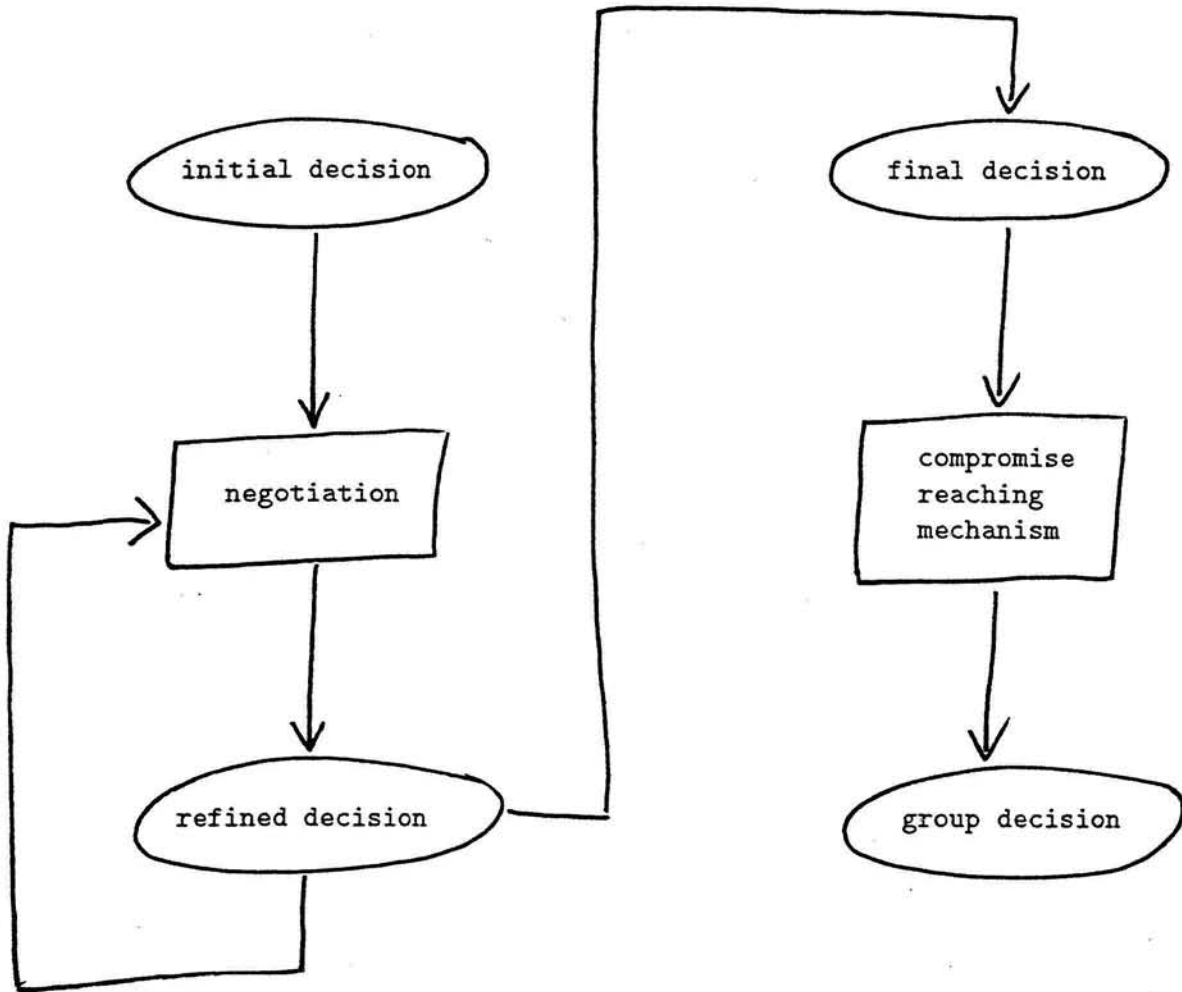


Figure 1: The various stages of a cooperative group negotiations process

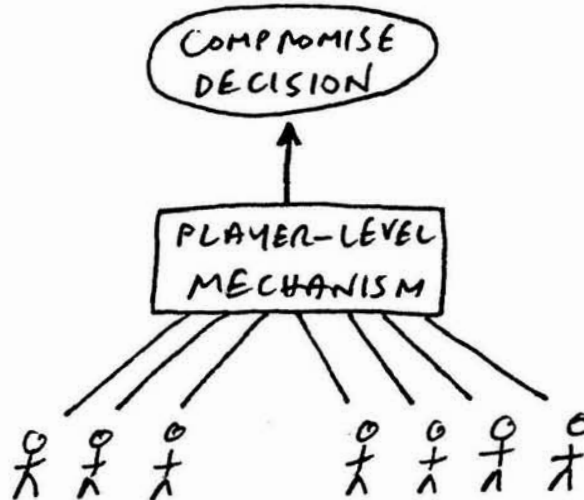
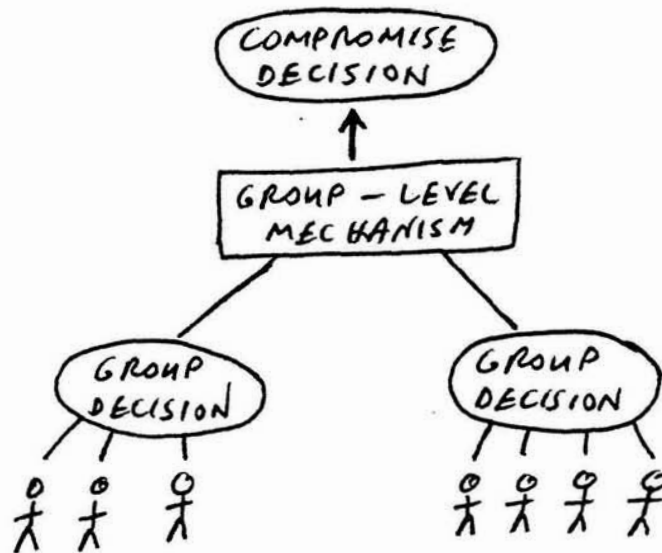


Figure 2: Group-level (top) and individual-level (bottom) compromise reaching mechanisms.

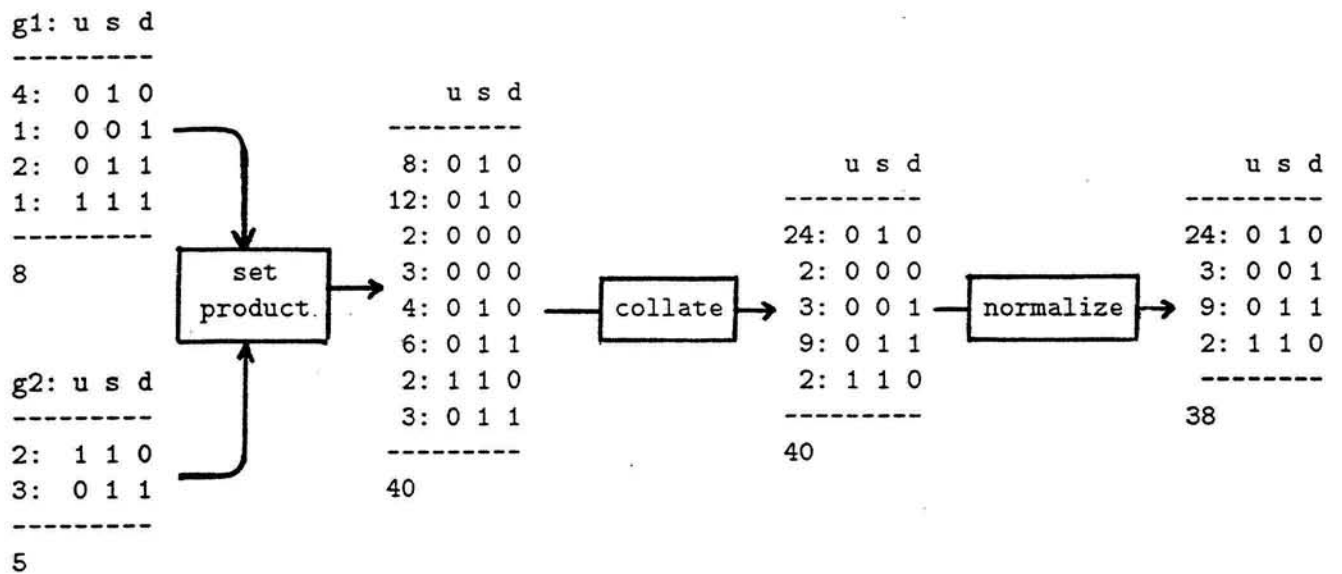


Figure 3: The set-product compromise reaching mechanism and its various stages