

# Team Decision Problems

A paper by Dr. Roy Radner

Andrew Brennan

April 18, 2011

# 1 Literature

Radner's paper on team decision theory is a direct response to a 1955 paper by Marschak [14]. Marschak's paper introduces the theory of teams by providing motivation for the development of such a field. He calls for the introduction of scientific language and concepts to study ideas of human organization. He claims that technical development is driven only by precise language and well-defined concepts, and issues the challenge to apply such techniques to organizational behaviour. The paper includes a couple of solved examples for two team members with artificial cost functions and an implicit prior on the various events. The technique for solving these examples is analogous to game theory, which was more highly developed at that time.

In response to Marschak's call but before Radner published his paper on Team Decision Problems, there were discussions and some preliminary results regarding simplified team problems. Radner's first attempt at the problem [15] in 1959 considers convex polyhedral cost functions, and proves that optimal decisions are linear. It does not contain the mathematical formalism of his next paper, and was not nearly as widely received. Interestingly, in this paper he claims that "the 'character' of the decision problem is determined by the form of the function to be maximized" [15], rather than the modern belief that the information structure is the key ingredient to the difficulty and character of a decision problem.

Team Decision Problems [16] was Radner's second response Marschak's challenge. This paper uses considerably more highly developed mathematical formalism to approach more general results regarding the existence and uniqueness of optimal decisions. To do this, Radner uses ideas from measure theory, statistics, and decision theory.

Decision theory is the primary framework that Radner extended to formalize the theory of teams. Decision theory refers to making choices under uncertainty. The core of the problem proposed by Marschak is for many players with a common payoff to make a choice under uncertainty with only partial knowledge, so the extension seems natural. While the origins of decision theory are old, it was reinvigorated by a 1939 paper by Wald [19]. Radner directly compares his problem formulation with the idea of a risk function, as used slightly differently by Wald [20], Hodges & Lehmann [11], Blackwell and Girshick [3], and Chernoff [4].

In 1956, Radner got his PhD in Mathematical Statistics, so he was well prepared to frame this question in statistical language. The idea of minimax estimators in decision problems, as used in the last third of the paper, are due to his PhD supervisor Savage [18] and Radner's use of the Bayesian approach in the first part of the paper is credited to this text as well, although Savage is not a founder of the field.

One of the major contributions Radner's paper made to the emerging field of team decisions is its mathematical formalism and use of  $\sigma$ -fields to represent information. This approach seems to have been inspired by Bahadur's [2] work with the fields, subfields, and how they relate to the idea of a statistic. Halmos' recent texts on finite dimensional vector spaces [8] and Hilbert spaces [7] paved the way to the existence and uniqueness result of Theorem 2.3.

## 2 Results

### 2.1 Framework

A major contribution of this paper was simply to frame the question formally.

The state of the world is denoted  $x \in X$ , and  $\mathcal{X}$  is a  $\sigma$ -field of measurable subsets of  $X$ . The decisions are either concerned with predicting a future state of the world or estimating the distribution of a random quantity, or both. Therefore, the state is modelled as  $x = (z, p) \in Z \times P$ , where  $Z$  is a class of random events and  $P$  is a set of probability measures on  $Z$ . The prediction question refers to predicting the random event  $z$ , and estimation refers to finding  $p \in P$  to most accurately describe the distribution of the events in  $Z$ . Let  $\mathcal{Z}$  and  $\mathcal{P}$  be  $\sigma$ -fields of measurable subsets of  $Z$  and  $P$  respectively.

For  $N$  decision makers  $\{DM^1, \dots, DM^N\}$ , let a decision  $a = (a_1, \dots, a_N) \in D = D_1 \times \dots \times D_N$ , where  $D$  is the set of all possible decisions, and each decision maker  $DM^i$  chooses a component decision  $a_i$  from possible decisions  $D_i$ . Let  $\mathcal{D}_i$  be a  $\sigma$ -field of measurable subsets of  $D_i$  for each  $i = 1, \dots, N$ .

It will be necessary to specify the information available to each decision maker. The paper employs two ways of doing this. The first is to specify information subfields,  $\mathcal{Y}_i \subset \mathcal{X}$ . Then a team decision function  $\alpha = (\alpha_1, \dots, \alpha_N)$  is any function such that  $\alpha_i : X \rightarrow D_i$  is measurable with respect to  $\mathcal{Y}_i$ . The second and equivalent way of specifying available information is by a transformation. For some measurable space  $Y$  and measurable function  $T : X \rightarrow Y$ , component decision functions  $\alpha_i : Y_i \rightarrow D_i$  must be measurable. The transformation  $T$  specifies the information available to each decision maker.

The payoff function  $u$  specifies the reward  $u(a, x)$  for any decision  $a$  at any state of the world  $x$ . For a given  $p \in P$ , the expected payoff is  $U(\alpha, p) = \int_Z u[\alpha(z, p), (z, p)] dp(z)$ . By specifying a prior distribution  $G$  on  $P$ , the Bayes expected payoff is  $V(\alpha, G) = \int_P U(\alpha, p) dG(p) = \int_X u[\alpha[x], x] dp(z) dG(p)$ . Let  $A$  be the set of all allowable team decision functions, a subset of the measurable team decision functions. As an example, the team could be limited to linear policies. A Bayes team decision function minimizes the Bayes expected payoff among all decision functions in  $A$ .

The other quantity to be optimized is risk,  $\rho(\alpha, p) = \sup_{\alpha' \in A} U(\alpha', p) - U(\alpha, p)$ . Let  $\mathcal{A}$  be a  $\sigma$ -field of subsets of  $A$ , and  $A^*$  a set of probability measures on  $\mathcal{A}$ . An element  $\delta \in A^*$  is called a randomized decision policy. The randomized decision policy  $\hat{\delta}$  is called minimax in  $A^*$  if  $\sup_{p \in P} \rho(\hat{\delta}, p) \leq \sup_{p \in P} \rho(\delta, p) \quad \forall \delta \in A^*$ . A minimax policy optimizes for the worst case scenario.

## 2.2 Person-by-person maximization and stationarity in the Bayes problem

The paper discusses conditions under which a person-by-person optimal policy is Bayes. The situation is more general than discussed in course lectures because the payoff function is permitted to vary according to the state of the world.

**Definition** For a differentiable cost function  $u$ ,  $\alpha$  is called a *stationary* decision function if  $V(\alpha) > -\infty$  and  $\frac{\partial}{\partial a_i} E[u(\alpha_1(x), \dots, a_i, \dots, \alpha_N(x), x)] | \mathcal{Y}_i |_{a_i = \alpha_i(x)}$  is zero for a.e.  $x$  and each  $i$ .

**Definition** A function  $V(\alpha)$  is called *locally finite at  $\alpha$*  if  $|V(\alpha)| < \infty$ ; and for any decision function  $\delta$  such that  $|V(\alpha + \delta)| < \infty$ ,  $\exists \{k_i\}_{i=1}^N$  such that  $|V(\alpha_1 + h_1 \delta_1, \dots, \alpha_N + h_N \delta_N)| < \infty$  for all  $\{h_i\}_{i=1}^N$  such that  $|h_i| \leq k_i$  for  $i = 1, \dots, N$ .

**Theorem 2.1** *If*

1.  $u(a, x)$  is concave and differentiable in  $a$  for a.e.  $x$ ,
2.  $\sup_{\beta} V(\beta) < \infty$ ,

3.  $V$  is locally finite at  $\alpha$ ,

4.  $\alpha$  is stationary,

then  $\alpha$  is Bayes.

This theorem is important because it significantly reduces the problem to finding person-by-person optimal solutions, which is much simpler than finding Bayes decision functions.

### 2.3 Quadratic Payoff

The remainder of the paper assumes that the payoff is quadratic for any state of the world  $x$ . This dictates that  $u(a, x) = \lambda(x) + 2a'\delta(x) - a'Q(x)a$ , for some  $\lambda : X \rightarrow \mathbb{R}$ ,  $\delta : X \rightarrow \mathbb{R}^N$ , and  $Q : X \rightarrow \mathbb{R}^{N \times N}$  each measurable (with respect to  $\mathcal{X}$ ), and  $Q(x)$  positive definite for a.e.  $x$ .

This dictates that the optimal team decision function is  $\gamma(x) = Q^{-1}(x)\delta(x)$  for any given  $x$ . The loss associated with using any other decision  $a$  is

$$\begin{aligned} [a - \gamma(x)]'Q(x)[a - \gamma(x)] &= a'Q(x)a - \delta(x)'a - a'\delta(x) + \delta'(x)Q(x)\delta(x) \\ &= u(\gamma(x), x) - u(a, x) \end{aligned}$$

Since  $\gamma$  may not be in  $A$  (may not even be measurable with respect to each  $\mathcal{Y}_i$ ), the problem becomes finding the decision policy  $\alpha$  to minimize the expected loss,  $\sigma(\alpha, p) = E[(\alpha(x) - \gamma(x))'Q(x)(\alpha(x) - \gamma(x))|p]$ . The Bayes expected loss  $\bar{\sigma}$  is accordingly defined by  $\bar{\sigma}(\alpha, G) = \int_P \sigma(\alpha, p) dG(p)$ .

### 2.4 Projections

With the quadratic payoff function, an optimal team decision function can be found with use of Hilbert Spaces and projections. All expectations in this section are with respect to the prior  $G$ . Let  $H$  be the space of all measurable functions  $\alpha : X \rightarrow \mathbb{R}^N$  s.t.  $E[\alpha(x)'Q(x)\alpha(x)] < \infty$ . Under the inner product  $(\alpha, \beta) = E[\alpha(x)'Q(x)\beta(x)]$ ,  $H$  is a nontrivial Hilbert Space. Let  $A$  be the set of all measurable team decision functions. This framework is sufficient to prove the following two theorems, which together give conditions under which a unique optimal team decision function will exist.

**Theorem 2.2** For any measurable  $\gamma : X \rightarrow \mathbb{R}^N$ , the set  $F$  of  $\alpha \in A$  for which

$$E[\alpha(x) - \gamma(x)]'Q(x)[\alpha(x) - \gamma(x)] < \infty$$

is either empty or it is the closed linear subvariety  $A \cap (\gamma + H)$  of the complete linear variety  $(\gamma + H)$  under the distance function

$$d(\alpha, \beta) = \|(\alpha - \gamma) - (\beta - \gamma)\|_H$$

**Theorem 2.3** If  $F$  is not empty, then there is a unique team decision function  $\alpha$  that minimizes the Bayes expected loss  $\bar{\sigma}(\alpha) = \|\alpha - \gamma\|^2$  on  $F$ , and  $\alpha$  is the orthogonal projection of  $\gamma$  into  $F$ .

The proofs of these theorems were discussed in my presentation and rely almost exclusively on results from Hilbert Space theory.

The results regarding stationarity are extended in the quadratic case.

**Theorem 2.4** *Let  $r(x) = \min_a [a'Q(x)a] / [\sum_i q_{ii}(x)a_i^2]$ , and let  $r$  be the essential infimum of  $r(x)$ . If  $r > 0$  and if  $\alpha$  is stationary, the  $\alpha$  is Bayes.*

The proof follows from the previous theorems and some norm inequalities. Note that  $r(x) > 0 \forall x$ , so  $r \geq 0$ . In particular, if  $Q(x) = Q$  is independent of  $x$ , then  $r > 0$ .

## 2.5 $Q$ is a constant matrix

As noted by the author, the assumption that  $Q$  is a constant matrix is a significant reduction in generality. If  $Q(x)$  is permitted to change based on the state of the world  $x$ , then for each  $x$ ,  $Q$  can be chosen to be an approximation of the payoff function in the neighbourhood of  $\gamma(x)$ , the optimal decision for that state. This allows the quadratic form to be a reasonable approximation of a much larger class of payoff functions. However, if  $Q$  is constant, it has to be a global estimate of the payoff function, which is much less general.

The main result in this section is concerned with the case when the prior distribution induces Gaussian distributions for all the information variables and  $\gamma$ , the optimal decision function if all team members knew all the information. Since the system considered in the paper is static, this is essentially an LQG system with many team members – linear system with quadratic payoff, Gaussian state variables, and Gaussian uncertainty on the random quantity. As with the standard LQG setup, the result is that the optimal decision functions are linear.

Here, the information transformation  $T : X \rightarrow Y$  is done by the functions  $\eta_i : X \rightarrow Y_i$ , which represent the information available to  $DM^i$  under various states of the world.

**Theorem 2.5** *If, under a given a priori distribution  $G$ ,  $\gamma$  and the information functions  $\eta_1, \dots, \eta_N$  have a joint normal distribution, with parameters*

- $Cov(\eta_i, \eta_j) = C_{ij}$ ,  $C_{ii} = I_{K_i}$ ,  $E[\eta_i] = 0$
- $E[\delta_i | \eta_i(x) = y_i] = E[\delta_i] + d'_i y_i$ , for  $\delta(x) = Q\gamma(x)$

*then the components of the unique Bayes team decision function are linear,  $\alpha_i[\eta_i(x)] = b'_i \eta_i(x) + c_i$ , where the vectors  $b_i$  and the numbers  $c_i$  are determined by the systems of linear equations,*

$$\sum_j q_{ij} C_{ij} b_j = d_i, \quad \forall i = 1, \dots, N;$$

$$\sum_j q_{ij} c_j = E\delta_i, \quad \forall i = 1, \dots, N;$$

## 2.6 Markoff decision functions and risk

This section focuses on problems where each team member is trying to estimate the value of a linear functional of the mean of a random variable  $z \in Z$ . The loss for the team continues to be quadratic in form and is specified by a constant matrix  $Q$ . The analysis is done by considering only linear estimators with bounded expected loss. A Markoff estimator is defined to be an estimator that minimizes the expected loss for all values of the mean, among all linear estimators with bounded expected loss. Due to its similarity in form to the problem of minimum-mean-square-error linear unbiased estimation, which is exactly this problem with only one decision maker, it is known that the Markoff problem of finding such an estimator usually has a solution.

The first main result from this chapter gives a closed form for a Markoff estimator in the following situation. Assume each individual observes a different random vector, and the covariance between these vectors is known to everybody. The means of the vectors are not known, but are known to lie in a linear subspace of the direct sum of the  $N$  vector spaces. Then, a homogeneous linear form of the optimal estimator is given in terms of a projection.

In the case where there is only one decision maker, it is always true that the Markoff estimator of a linear functional of the mean of the observed vector is simply the same functional applied to the projection of the observed mean onto the subspace where the mean is known to lie. If this situation holds for multiple decision makers, the problem is called *decomposable*, and is not generally the case. The paper outlines two cases where the problem is decomposable. The first is when the covariance  $C_{ij} = 0 \forall i \neq j$ . The second occurs when each observed vector is of the same dimension  $k$ , is known to live in the same subspace  $M_i = M_0$ , and the covariance is  $C_{ij} = \sigma_{ij}I_k$ .

Finally, under several circumstances, a Markoff estimator is minimax in the set of all permissible team decision functions. In order for this result to hold, the set of all permissible covariances must be bounded in some way and the set  $P$  of probability distributions has to contain a sufficient density of normal distributions. The three cases that are proved deal with various ways of bounding this set and all require that for a given mean and covariance combination, there is a normal distribution in  $P$  with the corresponding mean and covariance. All three cases deal with the same quadratic cost function.

### 3 Critique

It would be easy to look at this paper and find its limitations. Can all payoffs be modelled as locally quadratic around the optimal decision? Is it possible for such an optimal decision not to exist so as not to be localizable around it? How is time incorporated into the model if decision makers affect the state of the world? Why was communication not incorporated?

However, such criticisms fail to view the paper in the academic setting of its time. In reality, this was the preliminary defining work of team problems in a decision theoretic environment. Decision theory until this point considered only one decision maker with full information. Using decision theory for multiple players with a common payoff but different information was introduced in this work, which has made a particularly useful framework for studying problems of decentralized decision making.

Radner started by stating a general model for static team decisions before making further assumptions regarding stationarity or quadratic payoff functions. The potential criticisms attack the assumptions that Radner made in order to gain tractability. As was soon realized, tractability would continue to be a central issue in team decision problems, as they can very easily become unwieldy. As such, many future directions were not to extend the generality of the work, but to explore various special cases. As an example, Groves [6] and Arrow & Radner [1] studied team problems with the structure that there is a central controller who allocates resources to various managers, who in turn produce goods at a rate dependent on their allocated resource and some random process. The interesting result in this case is that as the number of managers increases without bound, the expected goods production converges to the case where the centralized controller and each manager knows everything, which is essentially a one-person problem whose optimal solution is an upper bound for the team solution.

A theme in team decision problems is that the information structure determines the difficulty

of the question. Radner made no assumptions about the information structure in this work. This generality allows the results to be widely applicable, but is extremely difficult to extend. In order to simplify problems and obtain results, other researchers have assumed information structures, once again extending these results by considering special cases. This work began with researchers including Witsenhausen [21] and Ho [9],[10]. The most important classification system of the information structures describes the structure as being either centralized, which is a significant reduction of team decision problems to problems with a central controller; partially nested, which is also a simplification of the original problem but with more generality than the centralized problem while maintaining some tractability; and non-classical, which is in general extremely difficult, as evidenced by Witsenhausen's counterexample [21].

Most of the work that has built on this paper uses one of three main results. The first is the general formulation of team decision problems, which extends even to games where individuals have different information, which is a non-classical view of game theory since the traditional assumption is that all knowledge is common knowledge. The difference with the game theoretic questions is that although the payoff functions depend on all decision makers, each individual tries to optimize his/her own payoff function rather than working together to optimize a global function. Another direction in which the framework has been extended is to incorporate time-dependence into the model, where previous decisions affect future decisions [12]. This extension is natural since most teams operate for extended periods. Especially in business and economics, which is the perspective of Radner and Marschak, the teams continually make decisions based on changing information that depends on past decisions, so having an optimal policy in that setting is worth considerably more than a policy in a static setting.

The second and very widely used result of this paper is Theorem 2.1. Finding person-by-person optimal strategies is considerably easier than finding globally optimal strategies, so knowing when the simplification can be made is a very important contribution. This result has been relaxed slightly by Krainak, Speyer, and Marcus [13]. DeWaal and VanSchuppen [5] give an analogous result in the case of a discrete action space with applications in distributed computing. However, these three results remain the only known conditions whereby person-by-person optimality implies globally optimal team decision functions [5]. Because of its great simplification, the fact that these are the only existing results limits the possible directions of some modern research to cases where these theorems apply. One future direction of research could be to find alternative formulations that have the same result, or to find weaker sufficient conditions with a quest toward necessity.

A third result that is used considerably is Theorem 2.5, which makes the simplifications of a constant payoff (independent of state  $x$ ) and a prior distribution that induces normal distributions in all information variables and optimal decision functions. These assumptions reduce the problem considerably, since the ideal policy is shown to be a linear function of the observation.

Interestingly, the first portion of the paper, where Bayesian decision functions are sought, seems to have garnered a lot more attention than the last third of the paper, which considers Markoff estimators. There seems to be considerable future work in this area. For example, Radner presented only three cases in which the Markoff estimator is minimax, and made no claims of this list being exhaustive.

The assumptions that Radner made were reasonable to gain some headway into a difficult field: locally quadratic cost functions, the existence of optimal decisions around which to localize, and normal distributions on the information and optimal decisions. As a first paper, these are all very reasonable, and since future work has primarily been concerned with special cases, it seems as

though the introductory paper was at a satisfactory level of generality. However, there are many situations that are not of this form, and those have not been nearly as extensively studied. Fruitful future directions may involve a similar framework with a completely different set of assumptions. For example, Rusmevichientong and Van Roy [17] note that the Gaussian assumption is unrealistic in complex organizational structures, but locality of interaction is an unexploited realistic symmetry that helps to reduce the problem.

Finally, as noted in Marschak [14], cost functions can often be difficult to specify. What is the cost of a life versus saving a sum of money in a military setting? For economic applications, the costs and benefits are usually in terms of economic gain and market position, so are relatively easily measured. In order to better understand other organizations and biological systems though, it is necessary to have a realistic cost function. It seems as though this caveat will always be a liability with the team decision theory approach.

## References

- [1] K.J. Arrow and R. Radner. Allocation of resources in large teams. *Econometrica: Journal of the Econometric Society*, pages 361–385, 1979.
- [2] RR Bahadur. Statistics and subfields. *The Annals of Mathematical statistics*, 26(3):490–497, 1955.
- [3] D. Blackwell and MA Girshick. Theory of Games and Statistical Decisions. *New York*, 1954.
- [4] H. Chernoff. Rational selection of decision functions. *Econometrica: journal of the Econometric Society*, pages 422–443, 1954.
- [5] P.R. de Waal and J.H. van Schuppen. A class of team problems with discrete action spaces: optimality conditions based on multimodularity. *SIAM Journal on Control and Optimization*, 38(3):875–892, 2000.
- [6] T. Groves. Incentives in teams. *Econometrica: Journal of the Econometric Society*, pages 617–631, 1973.
- [7] P.R. Halmos. *Introduction to Hilbert space and the theory of spectral multiplicity*. Chelsea Pub Co, 1957.
- [8] P.R. Halmos. *Finite-dimensional vector spaces*. Springer, 1958.
- [9] Y.C. Ho. Team decision theory and information structures in optimal control problems—Part I. *Automatic Control, IEEE Transactions on*, 17(1):15–22, 1972.
- [10] Y.C. Ho. Team decision theory and information structures. *Proceedings of the IEEE*, 68(6):644–654, 1980.
- [11] JL Hodges and E.L. Lehmann. Some problems in minimax point estimation. *The Annals of Mathematical Statistics*, 21(2):182–197, 1950.

- [12] J. Marschak and R. Radner. Economic theory of teams. 1972.
- [13] J. Krainak, J. Speyer, and S. Marcus. Static team problems—Part I: Sufficient conditions and the exponential cost criterion. *Automatic Control, IEEE Transactions on*, 27(4):839–848, 1982.
- [14] J. Marschak. Elements for a Theory of Teams. *Management Science*, 1(2):127–137, 1955.
- [15] R. Radner. The application of linear programming to team decision problems. *Management Science*, 5(2):143–150, 1959.
- [16] R. Radner. Team decision problems. *The Annals of Mathematical Statistics*, 33(3):857–881, 1962.
- [17] P. Rusmevichientong and B. Van Roy. Decentralized decision-making in a large team with local information. *Games and Economic Behavior*, 43(2):266–295, 2003.
- [18] L.J. Savage. *The foundations of statistics*. Dover, New York, 1954.
- [19] A. Wald. Contributions to the theory of statistical estimation and testing hypotheses. *The Annals of Mathematical Statistics*, 10(4):299–326, 1939.
- [20] A. Wald. Statistical decision functions. *New York*, 660, 1950.
- [21] H.S. Witsenhausen. A counterexample in stochastic optimum control. *SIAM Journal on Control*, 6:131, 1968.