ABOUT PON
NEWS
CLEARINGHO

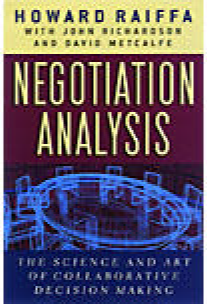

See Also:
- About the Book

Negotiation Analysis Supplement

by Howard Raiffa

Welcome!

This **Supplement** presents an extension of the material in the parent book, **Negotiation Analysis (NA).** When my publisher balked at the size of the first draft of **NA** — then known as *Collaborative Decision Making* — I was adamant that all must stay or else the picture would be incomplete. But I was torn: a book of 800 pages was a bit much. The compromise solution, once imagined, was acceptable to author and publisher: make all excess material including mathematical appendices available on the Web. So here you are and I am.

There is an added reason for this Supplement: adaptation of the book *plus* Supplement for classroom use. Let me refer to the book cum supplement as **Negotiation Analysis Supplemented** and abbreviated as **NAS.** I did not intend it this way, but I'm now of the opinion that **NAS** when coupled with *Smart Choices,* a book written by myself and two former Ph.D. students of mine, provides surprisingly good coverage of the burgeoning new field called *Decision Science.*

I had several venues for a Web site to choose amongst, but the most appealing was through the auspices of the Program on Negotiation, which I helped co-found with Roger Fisher in the early '80s.

The separate parts of the *Supplement* are given here in PDF format (and require a PDF Reader). Click the links below each of the following synopses to download that piece of the *Supplement.*

Table of Contents of the Supplement

A Detailed Table of Contents of NA.

This detailed Table of Contents takes several pages to present and according to my publisher is redundant with a detailed index, which is included in the book. My druthers, space permitting, would have been for the inclusion of both; but, happy compromise, **NAS** includes both.

PDF: 51 KB / 18 pages

NA Part Introductions and Chapter Summaries

PDF: 41 KB / 9 pages

My Short Vita

The Harvard University Press does not include much info about its authors on the jacket cover. I would have liked more. So, here is some more.

PDF: 5 KB / 1 page

The Role of a Neutral Joint Analyst

The Neutral Joint Analyst (NJA) is special type of intervenor. At times he (or she or a team) acts as facilitator, or as (proactive) mediator, or as (non-evaluative) arbitrator; but the NJA is also part decision analyst and game theorist who adds analytical skills in helping the parties realize joint gains. The NJA may be an invitee or invitor to the negotiation table. This short letter from an NJA to possible negotiating parties describes what the NJA has to offer. I'm intrigued with this intervenor. Perhaps it's because I identify so openly with this character.

PDF: 12 KB / 3 pages

Chapter 10B: More on Auctions and Bids.

The material in Chapter 10 on Auctions and Bids is rather extensive but stops at an inopportune place just when the material gets really interesting. This Supplement continues to examine the case when the prize is uncertain — like it would be if the prize was a glass jar filled with coins and where the bidding parties have different subjective probability assessments about the value of this prize. It also describes why in some cases the winner of certain types of auctions tend to lose money overall — the so-called "Winner's Curse." The Supplement also includes the case where one of the bidding parties has privileged information about the objectivre common value of the prize. This is known in the literature as the case with an "Unknown common prize with asymmetric information." It illustrates the power of possessing good information: you can intimidate other bidders.

The presentation is quite engaging, I would say, for those of you who relish mathematical reasoning. Try it. Some, with little math, have registered surprise at themselves: they actually enjoyed struggling with more mathematics than they thought they could tolerate.

PDF: 102 KB / 34 pages

An Appendix on Financial Risk Analysis.

It was not clear to me just where this material belonged. Somewhere in Part III.

It examines the case where negotiating parties have an opportunity to squeeze out joint gains by cleverly sharing financial obligations and rewards — perhaps over time. The parties may differ on probabilistic perceptions, on utilities —

e.g., on willingness to assume financial risks — and on the time value of money. The parties could struggle to do this by themselves, but how much more efficient this can be done through the interventions of a neutral joint analyst. The message rings clear: *differences beget joint gains.* The topic has a vitality of its own and if the reader becomes intrigued and wants more, then he would be well-advised to read some of the seminal papers of Robert Wilson on Syndicates and beyond.

PDF: 206 KB / 7 pages

The Two-Person, Zero-Sum Game

This could be an appendix to either Chapter four or ten. In **NA** I spend quite some time on game theory but there is a glaring void in the book on two important topics that were historically central to the development and acceptance of the theory: (1) the famous minimax theorem of the two-person, zero-sum game, and (2) the power of bluffing. I try in this and the next section to partially fill this void.

Consider the extreme case in which you are playing against a strict adversary who has diametrically opposing interests than you: if you prefer A over B, he, or she, prefers the opposite. This is the essence of the zero-sum case. Now if both parties have a finite number of strategies, then, you, by playing cleverly, can guarantee a return of at least v and your adversary, by playing cleverly, can guarantee you can get at most v. The mathematics is profound and beautiful and its inventor, John von Neuman, a genius. There is another, non-recognized reason for the importance of this type of security analysis. Consider the case where the row player, say, knows his own payofs but knows precious little about the payoffs of the other player,. The game cannot be analyzed for equilibria because the players don't have complete information. This happens quite often in practice and doing security analysis on one's own payoffs may be a starting point for further subjective analysis.

PDF: 51 KB / 12 pages

Security Analysis (as Done with EXCEL)

Shows how to find optimal, randomized strategies in a zero-sum game by using a special, linear programming, computer program, SOLVER, an add-in for EXCEL. In the game context the game problem can be approached from either the perspective of the maximizing row player or the minimizing column player. These two formulations are said to be dual to each other. The theory of linear programming got a shot in the arm when von Neumann, on the basis of the obvious duality of the special case of the zero-sum game, posited a general theory of duality for all LP problems. It has something to do with equilibrium pricing and immediately LP became the darling of the mathematical economist.

PDF: 26 KB / 5 pages

Simulations

If the book is to be used as a text in a course, I strongly recommend the instructor use several simulated negotiation exercises to be suggested in the Supplement. Let me explain how this might work.

For each of several chapters of the book, the materials in the Supplement will suggest appropriate motivating simulations. Suppose one of these simulations involve two parties — say labor and management — and you are assigned to play the role of labor; you will then be paired with a management player and arrange for negotiation time. Both of you will be expected to read a couple of pages or so of general information and, given your role assignment, you will have to read and study some confidential information for "labor eyes only." You then negotiate with your management counterpart in either face-to-face negotiations (the most preferred) or by telephone or by e-mail and keep records of what occurred. The controller of this simulation collects the data from many labor/management pairs and prepares a briefing report giving a statistical analysis of what worked and what did not. As in duplicate bridge, you should now ponder how well you did versus how well other labor players did when playing the identical simulation. Now you are ready to read the appropriate chapter in **NA** covering some of this material in a more active manner.

Download the Entire Current Version of the Supplement

This file includes the preceding *Table of Contents*

PDF: 93 pages

Problems downloading or opening these documents? Please contact Ed Hillis at hillis@law.harvard.edu.

(Modifed June 27, 2003)

Negotiation Analysis: The Science and Art of Collaborative Decision Making by Howard Raiffa (with contributing authors David Metcalfe and John Richardson), Harvard University Press, 2003.

DETAILED TABLE of CONTENTS

Chapter 4: Game Theory

Chapter 5: Negotiation Analysis

PART II: TWO-PARTY DISTRIBUTIVE NEGOTIATIONS

Chapter 6: Elmtree House

Chapter 10 Auctions and Bids

PART III: TWO-PARTY INTEGRATIVE NEGOTIATIONS

Chapter 11: Preparing for Negotiations

Chapter 12: Template Evaluation

Tradeoffs, BATNAs, and RVs

Chapter 13: Efficiency, Equity and Fair Division

Chapter 14. JOINT TEMPLATE ANALYSIS (*The General Case*)

Chapter 15: Real Behavior in Negotiations

Chapter 16: Negotiating without Reciprocity

PART IV: INTERVENTIONS IN TWO-PARTY DISPUTES

Chapter 17: Mostly Facilitation and Mediation

Chapter 18 : Mostly Arbitration

Chapter 19: WHAT IS FAIR?

Chapter 20: PARALLEL NEGOTIATIONS

PART V: MANY PARTIES

Chapter 21: Collaboration with Many Actors

Chapter 22: CONSENSUS

Chapter 23 : COALITIONS

Chapter 24: Voting

Chapter 25: Pluralistic Parties

26: Interventions in Multiparty Negotiations

Chapter 27: Social Dilemmas

Negotiation Analysis: The Science and Art of Collaborative Decision Making by Howard Raiffa (with contributing authors David Metcalfe and John Richardson), Harvard University Press, 2003.

Part Introductions

Part I:	Decision-Making Perspectives
Part II:	Two-Party Distributive Negotiations
Part III:	Two-Party Integrative Negotiations
Part IV:	Interventions in Two-Party Disputes
Part V:	Many-Party Negotiations

Part I: Decision-Making Perspectives

Part I comprising chapters one to five is an extended introduction to the book which develops analysis *for* negotiation -- not analysis *of* negotiation -- drawing extensively on:

- Decision Analysis -- individual prescriptive decision making
- Behavioral Decision Theory -- judgmental decision making
- Game Theory -- interactive decision making.

Chapter one contrasts three perspectives of individual decision making:

- Normative analysis -- how ideally decisions *should* be made (by super rational individuals without psyches);
- Descriptive analysis -- how real people actually *do* behave (very often at variance with the normative abstraction);
- Prescriptive Analysis -- how real people *could* behave better (more advantageously) with some systematic reflection.

In group decision making this threefold breakdown becomes more complicated. Much of the book gives partisan (prescriptive) advice to one decision maker in a group assuming the other individuals are not privy to that advice and where the others behave in a more descriptive fashion. In contrast to this orientation, the theory of games takes a jointly normative approach.

Chapter 2 is a tutorial that develops enough of decision analysis to be of use in the rest of the book, but it leaves for later, just-in-time development of additional topics.

Chapter 3 examines real behavior and indicates that many individuals, some of the time, do not act in conformity with rational behavior -- especially when uncertainties loom large. We call such examples of "deviant" behavior *anomalies, biases, errors,* or *traps* and we try to rationalize why and how they might occur.

Chapter 4 develops many of the concepts of the non-cooperative theory of games in a tutorial, self-discovery, pedagogical style; but as in chapter 2 it leaves special topics in the theory for just-in-time, later development where needed.

Chapter 5 provides an outline of the rest of the book. It introduces the concept of idealized, joint behavior in which all protagonists in a negotiation agree to negotiate in a truly collaborative manner agreeing to tell the truth and all of the truth. We call this joint FOTE analysis -- Full, Open, Truthful Exchange. We claim that often in actual *deal* making in contrast to *dispute* settling that this ideal is often approximated and the FOTE examination sets up an ideal against which other approaches can be compared. In the sequel, we back away from FOTE analyses to consider POTE-like behavior in which some protagonists tell the truth but not the whole truth -- *Partial,* Open Truthful Exchange -- leaving their bottom lines hidden from view of others; and back away further to NOTE behavior -- No Open, Truthful Exchange.

There is a tension in negotiations between *creating* actions designed to build a bigger pie and *claiming* actions designed to get a large share of the pie so created. The main concern of this book is involved in the creating function but the claiming function cannot be ignored. Part II of the book (chapters 6 -10) is mostly about claiming behavior for two-party negotiations; Part III (chapters 11- 17) mostly about *creating* joint gains; Part IV (chapters 18-20) about external interventions -- how facilitators, mediators, and arbitrators could help; and Part V (chapters 21-27) complicates all by introducing more parties into the fray.

PART II: Two-Party Distributive Negotiations (Chapters 6-10)

Chapter six presents a case study (Elmtree House) of the prototypical negotiation problem of this part of the book. A seller (in this case an institution that owns a halfway house) wishes to sell an asset (its residence, the house) to a buyer (a developer). The seller wants more; the buyer less; and they haggle. It's all a matter of claiming a larger part of a fixed size pie. There are two parties that negotiate over one issue: money – but it could be time or any other single commodity.

In the next chapter the problem is abstracted and the discussion addresses such questions as: How to prepare? Who should declare first? Why? How extreme a first offer? What's a reasonable counter-offer? What is the pattern of concessions? And so forth. The chapter discusses a related double-auction game in which the seller and buyer simultaneously offer sealed bids and a transaction takes place if and only if the bids are compatible. This game is analyzed normatively, descriptively and prescriptively and we conclude that good old-fashioned haggling might be better.

Chapter eight complicates the picture by introducing uncertainty, like the uncertainty of a trial outcome. We analyze the case of negotiating an out-of-court settlement in a liability case. The dominating uncertainties are who would win the court battle and the size of the jury award if the plaintiff prevails. We examine the problem as an individual decision problem under uncertainty and illustrate and extend the methodology introduced in Chapter two.

Chapter nine introduces the complication of time. It starts off by a seller dynamically searching for buyers. When to quit searching is another one of those decision problems under uncertainty. The central case of this chapter is the discussion of a hypothetical strike game and this is analyzed from the perspective of the individual decision maker and of game theory. Empirical laboratory observations demonstrate the potential for dysfunctional escalatory, vindictive behavior. Strike mechanisms have their problems but sometimes negotiating parties may wish there were a strike mechanism. Well that’s an opportunity for negotiation.

An Appendix to Chapter 9, presented on the Web, considers risk sharing and contingent contracting. We illustrate a case where the parties engage in a distributive bargaining problem with no zone of possible agreement. The negotiating parties, however, may disagree on probabilistic assessments, on attitudes towards risk and on time value of money (i.e., different discount factors) and these differences can be exploited in seeking joint gains through contingent contracting and risk sharing. There are now possibilities of creating a bigger pie. But still the problem is mainly one of claiming a larger share of the pie. This appendix serves as segue into Part III of the book.

Chapter ten considers the case where there are many buyers and just one seller. The seller could negotiate sequentially with each potential buyer or engage in a competitive bidding or auction procedure. The chapter compares several different types of auction and competitive bidding procedures and we use individual decision analysis to give partisan prescriptive advice to one of the players. We end the chapter by considering the all-too-common case when two partners want to dissolve their partnership by having one partner buy out the other.

We supplement the chapter by considering some advanced topics on the Web: such problems as bidding for the rights for a commodity whose worth is very uncertain – like bidding for rights to drill for oil. Empirically, the winners of such auctions end up by losing money -- the so-called winner’s curse. Why is this so? Life becomes even more complicated when there are asymmetries in knowledge about the uncertainty in question – like when one party has confidential exploratory information about the potential size of an oil field.

Part III: Two-Party Integrative Negotiations (Chapters 11-16)

This part of the book, comprising chapters 11 to 16, deals with Two-Party *Integrative* Negotiations. By "integrative negotiations" we mean: negotiations having the potential of resulting in joint gains. Part II dealt with Two-Party *Distributive* Negotiations which involved partitioning a fixed-sized pie. Whereas Part II was mostly about *claiming* tactics, Part III will be mostly about *creating* tactics -- how to create a bigger pie. But there is a tension between tactics used to create a larger pie and tactics used to claim a large portion of the pie created. How to balance this tension is part of the art and science of negotiation.

Much of negotiation involves the settlement of *disputes*. The bulk of our attention in Part III, however, will be with examining *deals* as presenting opportunities for joint gains (accompanied, of course, by some distributional strains) in contrast to disputes or problems that have to be resolved. Behaviorally speaking, negotiations involved in deal-making tend to be more collaborative than those involving the settlement of festering disputes. But still, much of the advice given for deals are relevant also for disputes. Part of this advice is to try to convert a dispute into a deal and to prevent a deal from becoming a dispute.

Chapters 11 and 12 concentrate on preparing for negotiations after the potential for such deal-making has already been flagged. We don't examine how we might find promising potential deals. We propose three phases in preparing for negotiations. In Phase 1, the parties deliberate separately and alone and consider their interests, objectives, visions; their alternatives to the present negotiation under review; their options; and gather information about the other side about their interests, alternatives and about their negotiation veracity and style. In Phase 2, the parties engage in an informal dialogue in which they feel each other out; they selectively and adaptively share information; they brainstorm together; they try to establish an appropriate relationship; and they decide whether they should proceed further. If they see a green light ahead, we suggest that they jointly decide just what must be decided. We strongly suggest that they construct a *template* to guide their further negotiations. The template specifies the issues that have to be decided and possible resolutions for each issue. In Phase 3 of their preparation, they individually and confidentially evaluate the template by clarifying their preferences for different contract outcomes and by exploring their Best Alternative To a Negotiating Agreement (BATNA). Some analytically-minded party might go so far as quantitatively scoring the template and establishing a numerical Reservation Value (RV) which specifies the minimum value that would be acceptable in its negotiations. Chapter 11 deals with Template *Construction*, 12 with Template *Evaluation*, and Chapters 13and 14 with Template A*nalysis*. Chapters 11 to 14 have a jointly normative flavor; 15 is mostly jointly descriptive, 16 mostly asymmetric prescriptive/ descriptive -- i.e., we give partisan prescriptive advice to Party A based on a descriptive analysis of the behavior of Party B.

Chapter 13 examines an idealization: both parties negotiate agreements using Full, Open, Truthful Exchange (FOTE) -- they tell each other the truth and the whole truth. [It was a surprise to me to find that in real-world, deal-making negotiations, this ideal is often the reality.] The chapter discusses efficient and equitable contract outcomes for the special case in which two parties have to share a set of indivisible items, the so-called Fair Division Problem. The problem is special in two ways: (1) the ith issue of the template involves what should be done with the ith item and there are just two possible resolutions (give it to A or to B -- no sharing allowed); and (2) there is no alternative to no agreement -- they must agree. This problem, interesting enough in its own right, sets up the analytical agenda of the next chapter on efficiency and equity of the scored template when some issues have more than two resolutions and where BATNAs and RVs play an important role. Chapter 13 deals with Template Analysis for the special template of the Fair Division problem and Chapter 14 generalizes this analysis to the general template and where there are alternatives to negotiations present. In Chapter 13 we learn that the source of joint gains stem mainly from the differences between the player's evaluations of the template and this insight carries over into 14. The flavor of 14 is still FOTE.

Chapter 15 deals with behavioral realities. We completely disengage from FOTE and comment on laboratory behavior with student-subjects. What happens when subjects without any training act as agents for principals who have provided them with a scored template and precise reservation values stating the minimum acceptable contract values? It took me by surprise when I first performed such simulations how very varied were the final contracts. It's important to examine what seems to work and what doesn't.

Still in chapter 15, we next consider anomalies, biases and errors of behavior in real-world negotiation settings. In an interactive setting misinterpretations beget misinterpretations and a dynamic may ensue where the parties spiral downwards in their pursuit of joint harms rather than joint gains. Cultural differences make it harder to establish a constructive negotiation style even though cross-cultural differences in interests are often the source for potential joint gains.

Chapter 16 takes a prescriptive/descriptive orientation. It offers partisan coaching advice to an analytically inclined collaborator suggesting how he might negotiate with a negotiator, who is less analytically inclined, who falls short of the FOTE ideal -- perhaps falls in the POTE camp -- who claims prematurely and excessively. We suggest how to negotiate in a hostile dispute with a disputatious adversary who needs to emote about your client's past despicable behavior. Part of the strategy is to try to foster more collaboration by effectively changing the others negotiation style.

PART IV: Interventions In Two-Party Disputes

Our first chapter in this part starts out by considering conventional facilitation and mediation reserving the next chapter for conventional arbitration. Rather than getting involved in disputes about nomenclature, we adopt the neutral language of “external helper “ or simply “helper” or “3rd party intervenor”. In our abstraction, we talk about the {A, B; H} dynamic where A and B are the negotiating parties and H the helper. We start by drawing up lists of conventional roles that H might perform, saving those that have a more analytical flavor for later development. . What role H plays depends on how H gets involved (as invitee or invitor or a mixture of these), on the context of the problem, on tradition, on the particular temperaments of the three actors involved, and on what is to be negotiated.

We talk about the interests of H as another player -- but a different type of player. We review the many reasons why external helpers are not used when perhaps they should.

The chapter then considers more active roles for the mediator by involving H more in generating proposals (e.g., in the form of Single Negotiation Texts) for the consideration of A and B. We discuss President Carter's role at Camp David as a proactive mediator with clout in the negotiations between Sadat of Egypt and Begin of Israel. We also consider, as a second case study of proactive mediation, my intervention in helping divide an art collection between two brothers who knew ahead of time that they could not do it alone without jeopardizing their relationship.

We comment on the rapid growth of FMA (facilitation, mediation, and arbitration) interventions and on the surge of interest in Alternate Dispute Resolution. ADR constitutes a spectrum of interventions ranging from weak facilitation to strong mediation and arbitration. In practice, a negotiated impasse often triggers mandatory arbitration where H, in an evaluative mode, is expected to impose a "solution" that is binding on the parties. There are all sorts of variations: arbitration may not be mandatory but a voluntary option on the part of the negotiators; the proposal of H may not be binding but suggestive; and in some circumstance H may be expected not only to be neutral initially but to be non-evaluative throughout.

In the standard distributive bargain where A wants a higher value and B a lower value, the conventional arbiter, after fact finding, proposes a final solution. In contrast, in final-offer arbitration, A and B are required to submit sealed final offers and then the arbiter, H, is required to select one of these offers. H has no other choice! We examine a game-theoretic treatment of this type of final-offer arbitration and surprisingly find it somewhat flawed from a normative perspective even though it seems to work in practice.

The latter part of the chapter deals with cases of complex integrative negotiations in which H, acting as a neutral joint analyst (NJA), helps the parties achieve an ideal collaborative compromise solution described at length in chapters 11 to 14. Such an NJA acts as a special kind of non-evaluative arbiter. The possible reluctance of the parties to truthfully reveal their reservation values (RVs) complicates the analysis and a double auction bidding system is introduced to help resolve this complication.

After the completion of an unassisted negotiation, the parties realizing that there may be still joint gains to be had, might choose to invite in an NJA to try to *embellish* their agreement in a so-called *post-settlement settlement*. In this case the parties do not have to reveal their RVs to the NJA -- the previously negotiated compromise acts as a pseudo joint reservation value.

There remains the nagging question for the NJA: which point on the efficient frontier should be selected as the most equitable. What is fair? The next chapter discusses this question in some detail. In Chapter 13 and 14 we already introduced some candidate rationales for "fairness" such as Nash's solution that maximizes the product of excesses over RVs and the one that maximizes the minimum of the two proportions of potentials (POPs).

The last chapter in this part examines the common nature of intractable disputes (mainly between feuding contiguous countries or national entities or cultural groupings), identifies a set of barriers that prevents constructive negotiations from taking place, and suggests how parallel negotiations (track II) could help.

The chapter documents a successful intervention in the Peru/Ecuador conflict. Influential, non-official surrogates from the belligerent countries were brought to Cambridge, Massachusetts by the not-for-profit Conflict Management Group to engage in a week-long activity called Facilitated Joint Brainstorming (FJB). The facilitator, Roger Fisher, (backed by a team of helpers including myself) engaged the group in generating creative options for the controversy, ruling out all claiming tactics. The surrogates were evidently inspired enough at week's end to return to their countries and to become change agents back at home. This is the way it was supposed to work and in interviews with the subjects it seemed to have worked that way. Fisher was publicly commended by the President of Ecuador and Foreign Minister of Peru for his contributions to peace in the region.

This FJB exercise involved a logistical feat: getting 11 distinguished surrogates (e.g., former foreign minister, or retired military chiefs of staff) around a table in a neutral spot. As variant of the FJB which uses face-to-face discussion of surrogates, we introduce Facilitated, Shuttle, Pre-Negotiations with Surrogates. The facilitator (perhaps a dissertation writer) shuttles back and forth interviewing key surrogates on their perceptions of the controversy, on their interests, on eliciting creative options, on brainstorming in serial fashion, on obtaining their value tradeoffs, on determining their BATNAs, and on identifying several -- not one! -- possible feasible contracts. The facilitator can then prepare a pre-mediation briefing report (a PMBR) describing the problem and hopefully some potential "solutions," -- i.e., some equitable and efficient compromises The hope is that the PMBR may trigger a reconsideration of the merits of trying to negotiate a solution to a festering dispute or motivate the consideration of an opportunity that might be profitably negotiated.

PART V: MANY PARTIES (Chapters 21-27)

In this part we extend all fronts into many-party lands. By “many parties”, as in some primitive tribes, we mean *more than two.*

Everyone knows that reaching a decision becomes harder the more people are involved. Negotiations among multiple parties – or even decisions, which must be jointly taken by a group which is ostensibly on the same side – are often long and unhappy affairs.

Up to now we have been talking about relatively simple bilateral negotiations (with the addition of a third party helper, who is not also a party at interest in the negotiation). Once we move out of that setting, there are many different geometries we could consider:

- A Group of Separate, Individual Negotiators
- Bilateral Negotiations with Multiple Participants on Each Side
- A Group of Advisors Preparing One Side for Negotiations
- A Permanent Decision-Making or Advisory Group
- An Ad Hoc Decision-Making or Advisory Group.

Each of these groups pose particular challenges. But there is also a core of similar difficulties that will have to be managed in all of those contexts. Chapter 21 starts by listing in what ways some groups do better than individuals and then why some groups do worse. We examine some anomalies of group behavior. It then attempts to understand the underlying reasons for this behavior – *why* do some groups behave so poorly. After examining losses due to cognitive overload, poor coordination, poor communication, and poor motivation, the chapter considers general prescriptive advice designed to partially ameliorate some of these causes. The advice considers such matters as: membership in the group; use of a facilitator or chair; need for an ongoing visual documentation of the deliberations; role of brainstorming and devising; focus on purpose and choice of a problem solving framework; decomposition of tasks and formation of sub-committees; and the allocation of time.

The next chapter (22) considers groups that strive for consensual agreement. Each of the members wishes to act in unison but there are strains in the community. Still they don't resort to voting or coalition formation. To a large extent this discussion generalizes the two-party material in chapters 11-13 which has a strong FOTE flavor. As was the case for two parties, we once again examine the fair division problem – but now with more than two parties – before launching into our full analysis of feasible, efficient, and equitable contracts for the many-party problem.

Chapter 23 has a game-theoretic flavor. It starts out by discussing the dynamics of coalition building and I recount some of my experiences at IIASA where I often engaged in that practice. We next consider the highly structured problem where every coalition is valued at some specific, fully known, total monetary worth (with common knowledge of this input data) and where players must then join coalitions and decide how to divide the resulting joint revenue. It's a wild game that defies prediction of outcome. Still some modest advice can be given. We examine principles that can be employed (a favorite one by Lloyd Shapley) to divide up the joint proceeds of the coalition of the whole in ways that reflect the power of various sub-committees. We end our consideration of coalitions by giving partisan advice to one party engaged in two separate intertwining negotiations; each sets up a BATNA for the other.

Chapter 24 examines the use of voting procedures for group action. Such common schemes as majority rules are fundamentally flawed since it can lead to intransitivities of group preferences. We review Arrow's famous Impossibility Theorem that proves that there is no way for three or more (i.e., many) individuals to combine their individual, ordinal preferences to obtain a group ordinal preference without violating some appealing desiderata. Many party negotiations often result in the identification of a few viable contracts for adoption, and guess what? They vote. One way out of the dilemma is to demand a richer set of inputs from the voters. They are asked not only about their ordinal preferences but the intensities of their preferences. They are asked for cardinal orderings. However, easier said than done – especially with the realities of insincere voting. We introduce a case study that examines how a group of scientists selected a trajectory for the Voyager mission to outer space.

The next chapter examines cases where at least one side in a negotiation is non-monolithic. We imagine a two-party (external) negotiation across a table that is complicated by the existence of an intense internal negotiation on each side of the table. Contracts that are negotiated across the table may result in winners and losers on one side of the table and the losers may try to block the negotiations unless they are duly compensated by internal transfers. Some call it "bribes." The challenge is to synchronize internal and external negotiations.

In the next chapter we extend the material in Part IV on interventions in two-party deals and disputes to many parties. We examine a hypothetical case of a community that has to negotiate with a family that wants to use part of their extensive land-holdings to erect a shopping mall and some high-rise office space. The quid pro quo is that the town will get money from a higher tax base and some needed land that can be used in many different ways. A facilitator/mediator is more active with the many-sided interests *within* the town than with the division *between* the town and the developer. We conclude with a case study of energy negotiations in the European Union. There were a series of facilitator/mediators who vainly sought a feasible solution (for all the participating countries), before the Luxenburgers stepped up to the plate and found the way toward mutual acceptance.

We end our book on a philosophical note by going back to the two-person prisoner's dilemma game (or social trap) and introducing more players. The headings tell the story: the commons problem (from sheep grazing to over-population of our planet); on over-utilization of resources (like fishing and forests); on polluting the atmosphere and despoiling the seas; on free riders (both individual, non-charity givers, as well as countries that want everybody else to curtail their green-house emissions); on the Not-In-My-Back-Yard (NIMBY) syndrome; on environmental injustice and on why disadvantaged communities could exact enormous compensations for the acceptance of an undesirable facility. And finally, all this is brought down to the personal level: "How much are you willing to sacrifice for the good of the whole?"

Howard Raiffa: Short Vita (January, 2003)

Howard Raiffa is the Frank P. Ramsey Professor (Emeritus) of Managerial Economics, a joint chair held by the Business School and the Kennedy School of Government at Harvard University. He held professorial positions in the Department of Economics and in the Department Statistics and chaired the inter-university doctoral program in Decision Sciences. He was the co- founder with Roger Fisher of the intra-university Program on Negotiation located at the Harvard Law School.

As a scientific advisor to McGeorge Bundy and later to Philip Handler, President of the U.S. National Academy of Sciences, Raiffa helped in the negotiations ('67 - '72) leading to the creation of the International Institute for Applied Systems Analysis (IIASA). The Institute, supported by 16 Academies of Sciences from the East (led by the Soviet Union) and the West (led by the United States), was designed as a confidence-building measure during the cold war. Raiffa was the first Director of the Institute ('72 - '75).

Raiffa received his Ph.D. in mathematics in 1951 from the University of Michigan. He has worked in operations research, game theory, statistical decision theory, decision analysis, risk analysis, behavioral decision theory, and in conflict resolution and mediation and has received lifetime achievement awards from each of those societies representing those fields. He has supervised close to 100 doctoral dissertations. He holds honorary degree from Harvard, Carnegie Mellon, Michigan, Ben Gurion, and Northwestern Universities. In year 2000 he received the prestigious Dickson Prize for Science (conferred annually to one recipient by the University Professors of Carnegie Mellon University).

Supplement to *Negotiation Analysis: The Science and Art of Collaborative Decision Making* by Howard Raiffa (with contributing authors David Metcalfe and John Richardson), Harvard University Press, 2003.

The Neutral Joint Analyst

Dear XXX and YYY;

I have been asked by ZZZ who just happens to know that you are about to enter into very complicated negotiations. ZZZ thought I might help you both and suggested I write to you and explain my craft. I bring to the table skills of the intervenor (facilitator, mediator, and arbitrator), coupled with the skills of a decision analyst. I call myself a Neutral Joint Analyst – an NJA if you will – I would like to help you to find both an efficient and equitable outcome of the negotiation opportunity you collectively confront. By *efficient* I mean that you squeeze out the full potential of this negotiation opportunity and not leave potential joint gains go unrealized sitting on the table. By equitable I mean that I will treat you equally in seeking a resolution of your joint decision problem – a resolution that gives each of you a comparable fair gain over what each of you could gain acting separately. We can think of this as *aided collaborative decision making.*

If you were to use my services – which because of my starving family don't come free – here is the way I propose to help you.

Stage 1. Formulating Your Joint Problem or Opportunity

I'll help you jointly to formulate your collective problem and opportunity. We jointly must decide just what has to be decided and decide on the alternatives for resolution. In joint and individual meetings with the two of you I will get to know the details of your problem and opportunity and create a framework or template which specifies the issues to be addressed and possible resolutions of each of these issues.

Along the way, from each of you, I might learn confidential material that you would prefer not to disclose to the other. My reputation would be destroyed if I were not to honor your confidential reflections most responsibly. Absolutely no unauthorized leaks. Part of my comparative advantage is that you might choose to tell me things that you would rather not disclose to the other side so that in devising cooperative solutions of your joint problem (or opportunity) I may be in the best position to devise creative alternatives. I don't want to come close to jeopardizing my comparative advantage by telling them what's for my ears only.

Stage 2. Evaluating What You Individually Want and Need

Working separately with each of you, I try to get you to express what you really want and desire – your basic tradeoffs. Helping you to grapple with tough questions like: how much of X do I need to justify giving up an amount of delta on Y. Not easy, but, in a negotiation context, important enough to have someone who can help you to think about horrendous tradeoffs.

For efficiency and equity analysis to come, I'll probe what might happen if you can' t come to an agreement with the other side, You'll have to think hard about what decision problem you might face if negotiations are of no avail.

I'll act under the assumption that what you tell me is the truth – no use paying me if you tell me false info – but you may wish not to disclose some particularly sensitive info. What I want from you is the truth but not necessarily the whole truth. You'll have to tell me just what I can disclose to the other side. You are in control of the sharing of information. Most of my clients are willing to share more info if they feel the other side is also forthcoming.

If I'm to help you fully I would like to learn just what are your possibilities if negotiations were to be aborted. I'll encourage you, with my help if you so desire, to explore your Best Alternative to No Agreement (BATNA) to help you assess the minimum you should be willing to accept in the actual negotiations.

Stage 3. Analyzing Your Joint Problem

I want to help you find a suitable compromise contract that is perceived as fair and squeezes out the joint potential gains. We don't want you to agree on a contract when there is another one that you both prefer. If you both are so inclined to tell me the whole truth, (confidentially) then I'll be in a position to suggest a compromise solution. From my perspective finding such a "solution" is a mathematical programming problem. You could jointly agree ex ante to bind yourselves to my suggested 'solution" or I could play the role of a post-settlement analyst: both of you come to a compromise solution and then let me see if I can embellish it by finding an alternate solution that you both might prefer. I can do all this not because I'm smarter than you but I'm privy to information that each of you lack and I'm a professional problem-solver.

To build up trust and confidence, why don’t you try me out on a easy problem that is not critically important to you. If you want I could suggest a problem that you both could practice on with me

. Your friendly neutral joint analyst,

Norman J. Altman

Supplement to *Negotiation Analysis; the Science and Art of Collaborative Decision Making* by Howard Raiffa (with contributing authors David Metcalfe and John Richardson), Harvard University Press, 2003.

Ch 10b Auctions And Bids B

Bidding Traps Arising from Ignoring Cues from the Actions of Others.

The case studies in this part of the chapter deal with the subtleties of subjective probabilities. Player A might have an intitial probabilistic assessment of some uncertainty, and after observing cues from the behavior of player B, A should modify this distribution. Sometimes these cues are subtle and empirically ignored and individuals are prone to fall into various auction or bidding traps. This will be exemplified by three problems.

From an analytical perspective, the material of this chapter is related to the discussion of Conditional Ambiguities and the Monty Hall conundra in Chapter 3.

PROBLEM 1: THE WINNER'S CURSE. *COMMON UNCERTAIN VALUE WITH DIFFERING PROBABILISTIC ASSESSMENTS*

We’ll start this more involved analysis with special case. This time we’ll assume that the bidders are involved in an auction where the prize has an uncertain but commonly known value and that the bidders must grapple with privately held probability distributions. In other words, the post-auction value of the prize to each and every bidder is exactly the same. However, each bidder has a different probabilistic assessment of what the post-auction value of the prize will be. They keep this information secret and act on their beliefs. To set the stage for further discussion we use two examples, one is frivolous and simple, the other real and complex.

Two Examples

Example 1. A gallon-sized glass jar is full of pennies, nickels, dimes, and quarters. The jar is transparent for all to see but no one knows the proportions of different coins or the total number of coins. Each of 30 bidders, say, has his or her own guesstimate of the value of the jar. The bidders can lift the jar if they wish. The prize is the contents of the jar.

Example 2. The prize is the oil deposits under a given site up for competitive bid. All the bidders have some information about the deposits but different guesses about the value. Let's simplify a bit by saying with full information the value of the oil would be the same for all bidders.

What's the strategic essence in these two problems? First, this is the case of a prize that has a common objective value, but--and it is a big but--the common value is uncertain and each of many bidders have differing *subjective* perceptions of this value.

The Coin-Packed Transparent Jar

Suppose that you have looked at the jar containing the coins and have assessed your own subjective probability distribution of its true monetary value. Your best guess is $200 but your distribution is quite spread out.[1] You are bidding against 29 other bidders. The number 29 is not critical, but the number should be sizeable, not 3 or 4. Well, what do you bid? For the purpose of this exercise let's assume that you are risk neutral. Let me invite Jay and Ann to join this discussion.

Jay: It seems to me that if I'm risk neutral my RV should be $200. That's as high as I would bid in the English auction, and in the Vickrey auction I would bid my RV of $200. Now for the English auction -- or for the standard competitive bid -- I would shade down from my RV.

Ann: I agree so far with Jay. And since there are so many other bidders, I wouldn't shade down too much from the RV of $200 -- say $185.

---

[1] Assume that $200 is the Expected Value of your assessed distribution.

Jay: This seems like an easy case. Why single it out?	Exp: Well you should suspect by now that there will be a special message of earth-shaking import.

Suppose the potential bidders recorded their single best guess at the uncertain but common value of the prize. There would be 30 guesstimates and these would have some distribution.

We still don't know where the true value is. It could be below, or close to, or above the average of these 30 best guesses. Suppose for now that the true value is close to the average of these 30 best guesses. Now consider three of the bidders: Alicia, Bernie, and Charlie. Alicia's best guess is on the low end of the distribution of best guesses. Bernie's best guess is slightly above the average of the best guesses and Charlie's is extreme on the high side. Their subjective, assessed probability perceptions of the true value are shown in Figure 10b.1.

Figure 10b.1 Three Subjective Probabilistic Perceptions of the True Value

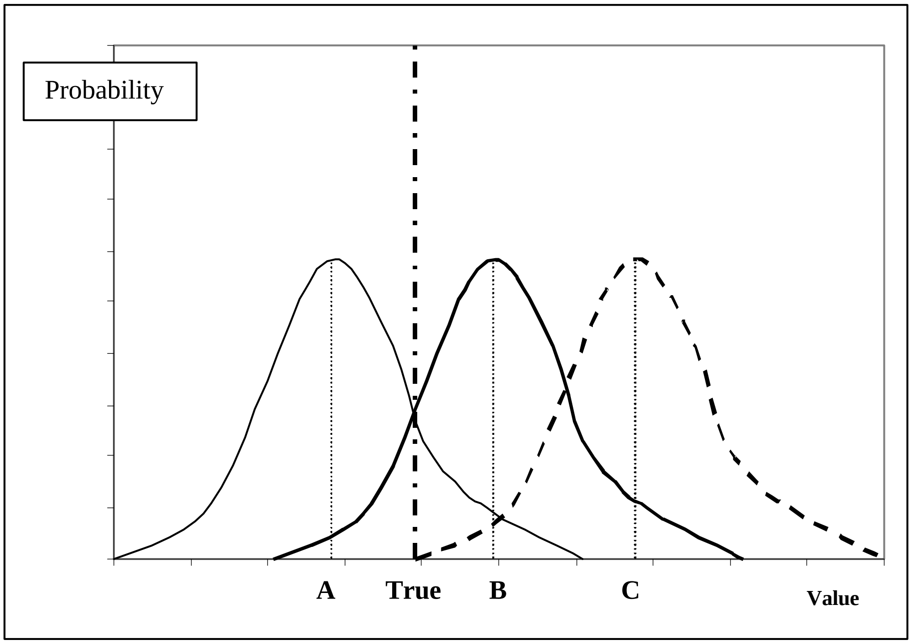

The trouble is that Alicia, Bernie and Charlie do not know where their best guesses lie among the distribution of best guesses. Now do you, Ann or Jay, see a problem lurking here?

Jay: Well, if the bidders follow our advice and bid close to their RVs, then Charlie will probably win the bid, or someone even more way out on the right tail will be the winner of the prize.

Ann: Hey, this is a good deal for the auctioneer. The more bidders we have the more chance there is for a right tail outlier and if bidders bid close to their RVs then the winner of the uncertain prize will be paying too much.

Jay: The winner will be a loser.

Exp: That's the message! *If you win, you lose.*

THE WINNER'S CURSE

The strong desire to win a prize in the face of competition often predominates over the rational analysis of what the prize is worth. Bidders ascribe the act of winning a greater value than the assessed value of the winnings. For example, a lot of oil companies that win highly contested bids, end up by paying too much. This anomaly is encountered in innumerable auctions.

Exp: Well, what should you do about this sad complication?

Jay: It's not sad from the auctioneer's point of view.

Exp: Agreed. I posed two questions to you for the coin-packed transparent jar problem. Question A asks for your best guess of the value of the prize before knowing any of the other bids. And Question B asks for your revised best guess knowing now that your best bid of $185 is the top of the 30 bids.

Ann: Ah, that's tricky. I would now know that most of the others would perceive the value of the common prize as below mine. If I thought that their opinions were somewhat informative, I would now want to bias my best guess downwards.

Jay: I would want to bid an amount "X" say, so that if I were told my X value were the highest bid, my revised best guess at the common value would still be higher than X. That's the way I would get around the winner's curse.

Ann: This really is complicated. Let's assume that we originally had a distribution centered around $200. That's our best guess without worrying about the strategic bidding problem. We then said earlier that the more bidders there are the closer we should bid to our RV. But now comes the winner's curse. The larger the number of other bidders the more worried we should be if we win. So with 29 other bidders we should be bidding much lower than our $200 original RV. We need a strategic RV like Jay just said.

Jay: Did I say that?

Ann: Sure, Jay, you said that you would want to believe that you could expect a profit <u>after</u> knowing that your bid was highest. Can you analyze this problem further?

Exp: Sure, but it's the qualitative message that's important. The game theorists have a ball with this problem. But in order for the game theorists to do their thing which is equilibrium analysis -- the problem has to be modified so that each bidder gets confidential information and each is given an objective probability distribution of the original RV's of the others. It is too complicated for me to discuss in the elementary presentation.

PROBLEM 2: AQUIRING A COMPANY; *IGNORING INFORMATION*

The "Acquiring a Company" exercise is the creation of Max Bazerman and William Samuelson. Please read the caselet and after a few minutes of reflection, answer the questions on the bottom of the page. When you're ready, read on.

ACQUIRING A COMPANY

In the following exercise you will represent company A (the Acquirer), which is currently considering acquiring Company T (the Target) by means of a tender offer. You plan to tender in cash for 100% of Company T's shares but are unsure how high a price to offer. The main complication is this: the value of Company T depends directly on the outcome of a major oil exploration project it is currently undertaking. Indeed, the very viability of the Company T depends on the exploration outcome. If the project fails, the company under current management will be worth nothing - $0/share. But if the project succeed, the value of the company under current management could be as high as $100/share. All share values between $0 and $100 are equally likely.

By all estimates, the company will be worth considerably more in the hands of company A than under current management. In fact, whatever the ultimate value under current management, the company will be worth fifty per cent more under the management of Company A rather than Company T. If the project fails, the company is worth $0/share under either management. If the exploration project generates a $50/share value under current management, the value under company A is $75/share. Similarly, a $100/share value under Company T implies a $150/share value under Company A, and so on.

The board of directors of Company A has asked you to determine the price they should offer for Company T's shares. This offer must be made now, before the outcome of the drilling project is known. From all indications, Company T would be happy to be acquired by Company A, provided it is at a profitable price. Moreover, Company T wishes to avoid, at all cost, the potential of a takeover bid by any other firm. You expect Company T to delay a decision on your bid until the results of the project are in, then accept or reject your offer before the news of the drilling reaches the press.

Thus, you (Company A) will not know the results of the exploration project when submitting your offer, but Company T will know the results when deciding whether or not to accept your offer. In addition, Company T is expected to accept any offer by company A that is greater than the (per share) value of the company under current management.

As the representative of Company A, you are deliberating over offers in the range of $0/share (this is tantamount to making no offer at all) to $150/share. What offer per share would you tender for Company T's stock?

My tender price is: $_______ per share.

****************************** *(Pause)*

Acquiring a Company: Simulation

Let’s thrash about a bit before we become systematic. How good would it be to tender an offer of $60 per share? Well let’s see what happens. Since all values between $0 and $100 are equally likely for Company T’s shares we ran 100 simulated trials based on values churned out by a random number generator. In this experiment we varied the value of Company T to see what the expected return to Company A will be when it offers T $60/share. The first twenty results of this experiment are shown in Figure 10B.2. Column 2, under the heading, “Value to Target” contains the random numbers between 0 and 100, column 3, “Value to Acquirer” shows the value of T under the management of the Company A (e.g., 18.29 x 1.5 = 27.42), column 4, “Acceptance” indicates whether T accepts A’s offer, and column 5, “Payoff to Acquirer” logs the payoff to Company A from acquiring (or not) Company T at the given price.

Figure 10B.2 Acquiring a Company Simulation With $60 Offer

TRIAL	**VALUE TO TARGET**	**VALUE TO ACQUIRER**	**ACCEPT**	**PAYOFF TO ACQUIRER**	**CUMULATIVE AVERAGE**
1	18.29	27.44	TRUE	-32.57	-32.57
2	71.89	107.84	FALSE	0.00	-16.28
3	53.62	80.43	TRUE	20.43	-4.05
4	53.31	79.97	TRUE	19.97	1.96
5	50.09	75.14	TRUE	15.14	4.59
6	30.7	46.05	TRUE	-13.95	1.50
7	80.45	120.68	FALSE	0.00	1.29
8	79.66	119.49	FALSE	0.00	1.13
9	18.23	27.35	TRUE	-32.66	-2.63
10	8.21	12.32	TRUE	-47.69	-7.13
11	37.32	55.98	TRUE	-4.02	-6.85
12	84	126.00	FALSE	0.00	-6.28
13	71.65	107.48	FALSE	0.00	-5.80
14	80.52	120.78	FALSE	0.00	-5.38
15	39.26	58.89	TRUE	-1.11	-5.10
16	19.11	28.67	TRUE	-31.34	-6.74
17	57.45	86.18	TRUE	26.18	-4.80
18	4.08	6.12	TRUE	-53.88	-7.53
19	1.45	2.18	TRUE	-57.83	-10.17
20	55.44	83.16	TRUE	23.16	-8.51

The first trial is a disappointing value to the target of $18.29 and the value to the Acquirer is 3/2 x $18.29 or $27.42. A bid of $60 would yield an acceptance (TRUE) and the net proceeds to the Acquirer would be $27.42 - $60.00 or -$32.58. Not starting off so well. But on the third trial the Acquirer makes $20.43. Over 100 trials the average return is -$9.13. Not a good deal.

We also ran the same experiment but inserted an offer of $50 instead of $60. We generated 100 new random numbers and the average return with this tender offer is still disappointing: a dismal -$6.20 per share return. What’s cooking here? The Acquirer should be able to make some money. Or should she?

Intuitive Thinking

Let’s examine how most people think about this problem. They start off asking the natural question: What is the expected value per share of the company to the Target? And the answer is $50. Next they ask: What is the expected value (per share) to the Acquirer? And their correct answer is 3/2 x $50 or $75. Next they say if their RV is $75. what should they bid? They shade down a bit from $75 to $60 say.

Jay: Can I interrupt here? That's just what I did. I thought I was clever. I bid $51. Now what's wrong with that logic?

Exp: Ann, what's wrong?

Ann: I did something similar. But when you posed the question, "What's wrong?" and after seeing the simulations, I think I see what's wrong. I should have guessed this from the beginning. I should have suspected. Why did you introduce this problem? It smacks of the winner's curse. For really high values of the company to Mr. T., Ms. A is not going to sell the company. Once Ms. A knows that Mr. T accepts the tender offer her RV changes.

Jay: Oh, that's clever. It is a disguised version of the Winner's Curse.

Exp: O.K. Let's examine what happens with a tender bid of $60. If Mr. T accepts the bid of $60, what is the expected value per share to Ms. T?

Jay: Well that will be $30 and not $50. That was my mistake. I didn't condition my probability assessment on Mr. T accepting my bid.

Exp: Now if Mr. A accepts your tender bid of $60, what is the expected value of the firm to you?

Ann: Well that would be 3/2 x $30 or $45 per share. So a tender bid of $60 would result in an expected loss of $60 - $45 or $15. That's even worse than the simulation indicated.

Exp: Your analysis is still not quite right Ann. The simulation is better than you're giving it credit for.

Acquiring a Company: Decision Tree Analysis

It's instructive to draw a decision tree of the problem from the Acquirer's perspective (see Figure 10B.3). At the first move, Bidder A must tender an offer price per share. In the diagram we follow the branch corresponding to a $60 offer. At the next move, Mr. T accepts or rejects the offer. As far as A is concerned this is a chance move: T accepts if the value of the firm to A is less than or equal to $60; and rejects the offer if x > $60. The chance of acceptance is the chance that x $\geq$ 60 or 60/100 or 3/5.

If Mr. T accepts the offer of $60, then as far as A is concerned the next move is at the control of chance and all values from $0 to $60 are equally likely. I approximate this chance mode with 6 equally likely values of 5, 15, 25, 35, 45, and 55. The expected value to Ms. A if she offers $60 and T accepts is -$15. But now to evaluate the offer of $60 we must move backwards in the tree to the node following the branch where A chooses $60. The expected value at being at this node is $\frac{3}{5}x(15)+\frac{2}{5}x(0),or-\$9.$ Now that agrees pretty closely with the simulated result of -$9.13.

Figure 10B.3 Decision Tree Analysis for Acquiring a Company

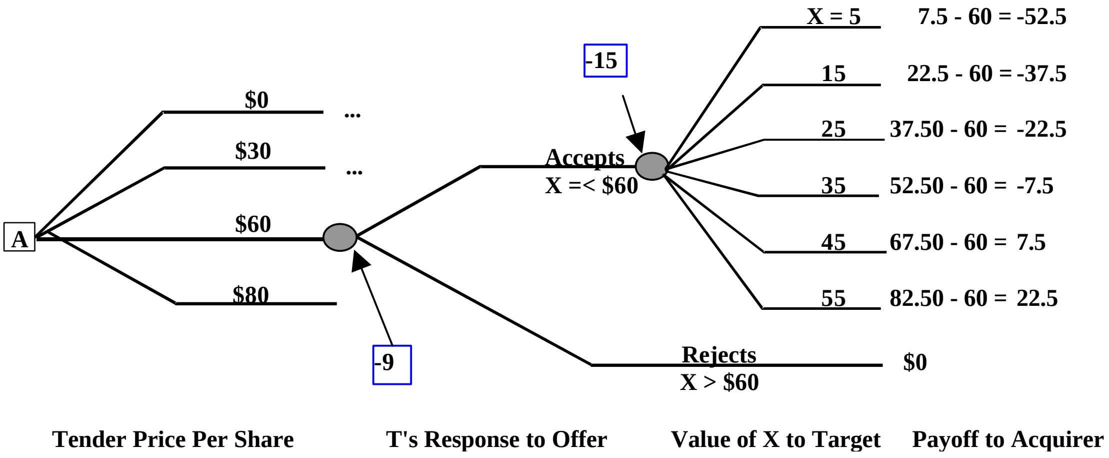

We can perform the same analysis for any bid b. Here's the argument. Keep the decision tree in mind. If b is greater than the value of the company, call that V, then the offer will be accepted. <u>But</u> if b < V the company will reject your offer. Whenever your bid is accepted your return is $\left(\frac{3}{2}V - b\right)$, where V is still uncertain. Since your bid was accepted any value in the range from 0 to b is possible and all are equally likely. So on average the company is worth $\frac{b}{2}$ whenever your bid b is accepted, and we get $\left[\frac{3}{2}\left(\frac{b}{2}\right) - b\right] \text{or} \left(\frac{3}{4}b - b\right)$ *so* already we see that this is a bad, bad deal. There is an analytical lesson here. Unaided intuition has limitations. Simulating an actual play of the game can provide fundamental insights. Drawing a symbolic decision tree can also help straighten out your thinking. Get used to these techniques.

Acquiring A Company: Empirical Results

Bazerman and Samuelson asked graduate MBA for their tender offer bids per share. The results are shown in Figure 10B.4. One hundred and twenty-three students responded when no money was involved. About three-quarters of them offered prices over $50 per share for a negative return. The situation changed only a bit when money was exchanged as demonstrated with the experiment with 66 other students.

Figure 10B.4 Empirical Results for Acquiring a Company

BIDS	NO INCENTIVES (n = 123)	MONETARY INCENTIVES (n = 66)
0	9%	8%
1-49	16%	29%
50-59	37%	26%
60-69	15%	13%
70-79	22%	20%
80+	1%	4%

Below we plot the expected payoff as a function of the tender bid offer. It's a deal that Ms. A should shun.

Figure 10b.5 Expected Return To Acquirer as a Function of Bid

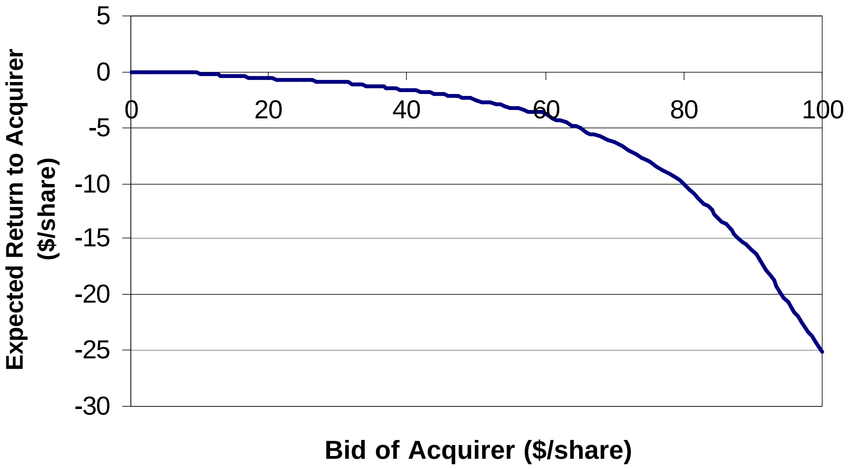

As Bazerman points out, not only do subjects do poorly on this exercise but when the game is repeated over time they persist in their irrational behaviors. That's because as shown in the simulations (Table 10B.2) some returns are positive.

PROBLEM 3: ASYMMETRIC INFORMATION

We're now going to consider another special case and use some game theoretic techniques to improve the value of our prescriptive advice. The structural characteristics, or rules of the game, for this auction are that the prize is of uncertain but common value (as it was with the transparent jar), that all bidders share the same

objective probabilities about the ultimate value of the prize, but, and here's the twist, there are asymmetries in the information held by different bidders. At least one of the bidders has "inside information" about the likely value of the prize.

Maxco-Gambit: Two Oil Companies Bidding for Drilling Rights

Let's motivate this bidding case by bringing in two oil companies, Maxco and Gambit, as our protagonists. Both oil companies have to submit sealed bids for the rights to drill at a given site. Gambit, however, has privileged, confidential information. Gambit owns the adjoining sites and has dug exploratory wells near the boundaries of this site. So Gambit knows more than Maxco about the chances of hitting oil and, if so, how much there might be. Gambit is careful not to share this information with its competitor, Maxco.

We'll abstract and simplify this situation into laboratory experiment. Imagine that the prize is equally likely to be $0, $100, or $200. In the motivating oil problem these are in units of millions or tens of millions of dollars. To be even more graphic suppose that I, as the experimenter, have three sealed envelopes containing respectively prizes of $0, $100, and $200. I choose one of these envelopes at random; put it on the table in front of you and put aside the other two envelopes. Now for the gimmick.

Assume that the bidders, labeled Maxco and Gambit, have differential information. Gambit is shown the contents of the selected envelope. Gambit now has perfect information about the value of the prize. Maxco knows that Gambit knows and Gambit knows that Maxco knows that Gambit knows and so on. Think hard about this one.

What would you bid as Maxco? And what would you bid as Gambit conditioned on what you might learn about the contents of the envelope?

*********************(Pause)*

Let's play the conventional sealed-bidding game: high bidder wins at the high price. When you are ready please read on. O.K. Let's resume. Let me call in Jay and Ann again.

Exp: What did you bid as Maxco?

Jay: As Maxco I bid $63.

Exp: And Ann, what do you do as Gambit?

Ann: My strategy was to bid $0 for the $0 prize; $51 for the $100 prize, and $112 for the $200 prize.

Exp: So in this case Jay playing Maxco would win the prize if it were $100 (since his bid of $63 is higher than her bid of $51), but he would lose the $200 prize (since his bid of $63 is less than her bid of $112).

Jay: Good. I would capture the $100 prize.

Ann: Why are you so pleased with yourself? You would win $37 if the prize were $100 but you would lose $63 if the prize were zero.

Jay: You're right. But you were not so smart, either. You would bid $112 for the $200-prize. You could have squeezed harder and come in with some value like $65 say.

Exp: Well, what do we see here? If Gambit is very conservative -- i.e., bids close to $100 for a $100 prize and close to $200 for a $200 prize -- then Maxco will only win when the actual prize is zero.

Jay: This is really the winner's curse once again. Maxco wins the prize only to lose money. So the conclusion seems to be that Maxco shouldn't bid. Especially against a conservative Gambit.

Ann: But if Maxco thinks this way, then I, as Gambit, shouldn't be conservative. I should bid $30 or a $100 prize and say $40 for a $200 prize.

Jay: But be careful Ann. If I suspect that you will be greedy and bid so low, I, as Maxco, will be able to sneak in with a bid of $50 and get both the $100 and $200 prize.

Jay: It is important for us as a player, to try to out-psych the other player. But where does this end?

Exp: Well, one tack would be to examine equilibrium behavior. The idea is to give advice to both players that will be stable. The advice should be such that each player would want to follow the advice if the other player follows the advice. The other tack is to examine empirically how people like yourselves behave so that you could behave best against this empirical mix.

Ann: And how should I behave if I don't have this empirical information about others?

Exp: That's the usual case, of course. You just have to imagine yourself in the role of the other party; try to think as you believe he or she might think; and heroically assess your probabilistic betting distribution of what the other party will do. I like to think of l00 different people each playing the role of the other party and I assess the distribution of what I think they will do.

Jay: But if you're lousy at this guessing game, you won't do well will you?

Exp: That's right. But I think you are better off thinking this way than not. Of course, some laboratory experimental results bolster my impressions.

Before continuing I want to recount an experience I had in a large MBA class in Negotiation Analysis. Early in the semester I described the Maxco-Gambit game and chose two students at random to play the game for money in the class. I was the banker. I'll use fictitious names: Max for Maxco, and Gambriel for Gambit.

I shuffled three envelopes with prizes $0, $10, and $20, chose one at random, and showed it to Gambriel. Then Max and Gambriel wrote down their bids. "Well, what did you bid Max?" I asked. "Fifteen dollars," replied Max and there were impolite guffaws from the 100 onlookers. “And you Gambriel?” I continued. “I bid $12,” said Gambriel, “because the envelope indicated the $20 prize. It turns out to have been a stupid bid.”

Max was delighted. He netted $5 and snickered back at his disapproving audience. This shows the danger of a sample of size one. Here is a case of a very bad decision for Max that turned out to have a good outcome. Some students took Max’s side and said that if the outcome was good, the decision was good. Max said, “I’m lucky in gambles; I had an intuitive flash that the chosen card was $20. So I gambled and won.” I thought to myself, “I wouldn’t hire this guy!”.

Think carefully about the Maxco-Gambit game described above, then answer the following five questions:

1. *If you were Gambit, would you prefer to obtain your confidential information* *secretively* *(so that Maxco would not know you had this privileged information) or* *openly* *(so that Maxco knew that you knew the value of the prize and he or she didn’t)?*
2. *If you were Maxco and by counter-intelligence, say, you could also learn the value of the prize, would you want to learn this information* *secretively* *(so that Gambit would not know you also knew the value of the prize) or* *openly* *(so that he or she would know you knew)?*
3. *If you were Gambit and you were not given the confidential information about the value of the prize but could buy this information (in a way that Maxco would know you knew), up to how much would you pay?*
4. *Neither Maxco nor Gambit knows the value of the prize. But this information is available on a confidential basis to the highest bidder. What would be your RV for this auction?*
5. *There is a sequence of repeated Maxco-Gambit games. At each trial the high bidder at that trial gets the privileged confidential information at the next trial.*

> *This means, for example, that if Maxco outbids Gambit on trial 4, then Maxco learns the value of the prize at trial 5 and Gambit now knows at trial 5 that Maxco knows the value of the prize and she doesn't. How would you bid?*

****

Ann: I, Gambit, would want Maxco to know I have the information. I would want to intimidate him to bid zero so that I could bid low. Indeed, even if I could not get the information, I would want to lead him to believe I had the information.

Exp: Good. Now for question 2.

Jay: I, as Maxco, would want my counter-intelligence to be secret. I would like her to believe I was in the dark, so she would be greedy and bid low. Then, if I learned the prize were $200, and if she didn't know I knew, then I could steal the $200 prize for $70 or $80.

Exp: I agree. Now for question 3.

Ann: I'm Gambit so I suppose you want me to respond. But I don't have the faintest idea what the answer is. Is the answer obvious?

Exp: No, it's not. But the line of attack should be clear. Here's what I would do to answer the problem. I would first consider the problem when I, Gambit, did not know the value of the prize and Maxco knew I didn't know. To answer that problem I would assess a probability distribution for the bid of Maxco. Then I could figure out my best retort and calculate my expected return. Next I would consider the case where I learn the value of the prize and I know Maxco knows that I know. Once again I would assess a probability distribution for Maxco bids and then calculate my best $100 and $200 responses and finally calculated the value of that game. The difference would inform me about the value of this information. Go on to question 4.

Jay: My turn. It's a little different from question 3 because if I don't win the bid, she will win it and clobber me. I have to figure out my RV for this information with the understanding that if I have it I can intimidate her while if she has it she can intimidate me.

Exp: That's right. Go on to question 5.

Ann: I would want to bid higher than usual because if I win, I could intimidate Maxco at the next round. It would depend on the number of rounds to be played. This would be fun to play.

Exp: All right, if you want to do this, I'll be the bank, but let's play it for pennies instead of dollars.

Asymmetric Imperfect Information

These Maxco-Gambit bidding exercises were motivated by the real-world problem of bidding for oil: the right to drill at a given site. As you recall Gambit had privileged information, but not necessarily perfect information. In the exercises we just covered, we assumed the case of the polar extreme where Gambit learns the exact value of the prize. We now back away from that extreme assumption and to capture this complexity we assume that Gambit doesn't learn the contents of the envelope that determines the prize, but learns something less precise. We now reduce the size of the prizes tenfold and use three envelopes with prizes $0, $10, or $20; and one is taken at random to determine the actual prize. Now assume that Gambit learns the contents of one of the two discarded envelopes. And assume, as before, that Maxco sees that Gambit has this additional information. For example, if Gambit learns that one of the discarded envelopes is $20, then she knows that the prize is equally likely to be $0 or $10. The information is valuable, but how valuable? How intimidated should Maxco be in this case? What's the expected value of the game now to Maxco?

Some of you, intrepid souls, might want to try an analysis yourselves. Others, and I guess the majority of you, might want to take a respite from these Maxco-Gambit games and push on with the fascinating material still to come.

Maxco-Gambit: Illustrative Bids

Let's explore some possible bids. See Figure 10B.6. Let m represent Maxco's bid and G_0, G_{10} and G_{20} represent Gambit's bids if she learns the contents of the prize are $0, $10, and $20 respectively. Case 1 examines the results when m

= 10, G_0 = 0, G_{10} = 8, and G_{20} = 11. What happens to Maxco in this case? Well Maxco would lose his bid of $10 if the prize is 0; wins $0 if the prize is $10; and he does not win the bid if the prize is $20 -- that is because Gambit's bid G_{20} is 11 which is higher than m = 10. So Maxco's expected winnings is: $\frac{1}{3}x(-10)+\frac{1}{3}(0)+\frac{1}{3}x(0)$ or

-10/3. Now glance down the seven illustrative cases and reflect a bit on what you see happening.

Figure 10b.6 Illustrative Bids for Maxco-Gambit

CASES	**m**	**G_0**	**G_{10}**	**G_{20}**	**EV(m)**
1	10	0	8	11	-10/3
2	10	0	7	9	0
3	5	0	7	8	-5/3
4	5	0	3	7	
5	5	0	3	4	15/3
6	5	0	3	4	-2/3
7	5	0	1	1.5	24/3

From this simple analysis we can already see that if Gambit is conservative (i.e., G_{10} near $10, and G_{20} near $20) Maxco loses money when he or she has bid. By contrast, if Gambit squeezes hard (i.e., G_{10} and G_{20} are very low), then Maxco can make some money by bidding 5 or so.

Maxco-Gambit Decision Trees

We now proceed to a decision tree analysis of the bidding problem with asymmetric information. First we examine Maxco's decision tree (see Figure 10B.7).

Figure 10B.7 Maxco's Decision Tree

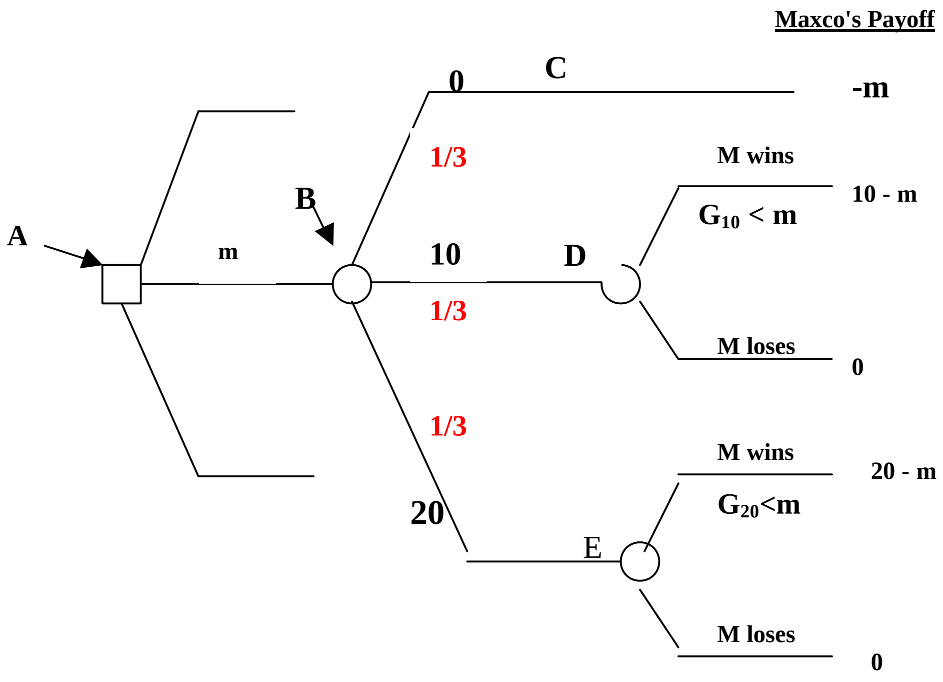

Let’s consider Maxco’s mathematical machinations. Maxco must submit a bid m. Then, at chance node B, there is a one-third probability of yielding a prize of \$0, \$10, or \$20. Let’s examine node D. We imagine that Gambit knows the prize is \$10 and submits a bid of G_{10}. If that bid, G_{10} is less than m, Maxco wins l0-m; if $G_{10} > m$, then Maxco loses and gets a zero pay off. The expected payoff at node D is therefore

$$(10-m) \text{ x } \text{Prob } [G_{l0} < m] .$$

Arguing in a similar fashion, the expected payoff at node E is

$$(20-m) \text{ x } \text{Prob } [G_{20} < m] .$$

The expected payoff for Maxco is the average of the payoffs at nodes C, D, and E. For various m-values, Maxco must think of the possibilities that G_{10} and G_{20} are below m.

Now let's look at Gambit's decision tree (see Figure 10b.8). At node A, chance yields prizes of \$0, \$10, and \$20. At node B, Gambit, knowing the prize is worth \$10, must choose a bid G_{10}. At node D, Gambit wins if Maxco's bid m is less than G_{10}. A similar story for nodes D and E. Gambit must assess probabilities that Maxco's bid is less than G_{10} and G_{20}.

Figure 10B.8 Gambit's Decision Tree

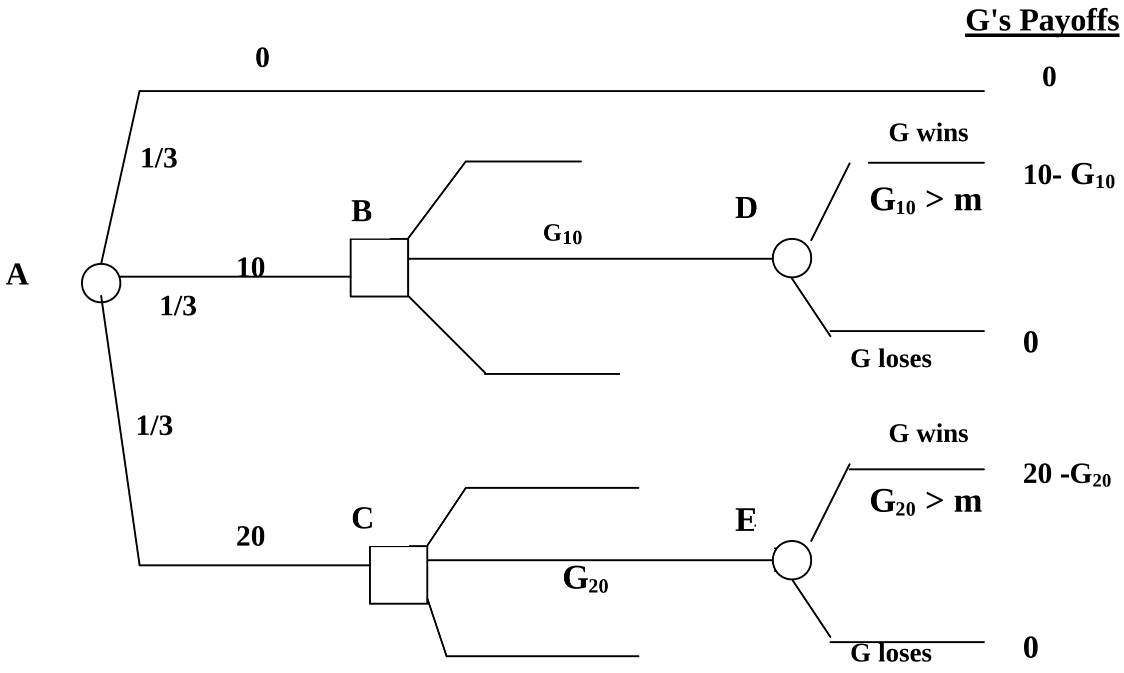

Maxco's uncertainty involves the unknown values of G10 and G20. Should Maxco work through Gambit's tree to figure out what Gambit will do? This is not easy because Gambit's tree depends on the unknown Maxco bid m. How far should we go thinking about what each of us is thinking? Game theorists want to go the whole way. To think what the other is thinking about what you're thinking about ... etcetera. In practice most people have trouble going even the first step in new complicated problems. Of course, if the problem is very clear cut, important and repetitive, it may be necessary to do more interactive, reflection analysis and then equilibrium analysis becomes of real practical importance. Let's see what happens empirically with the Maxco-Gambit game in the laboratory.

Empirical Distribution of Maxco and Gambit Bids

We can learn something by studying the way people bid in practice - even if its not a real bidding situation. Here are the results from a class of eighty-six Harvard MBAs. The average Maxco bid was \$4.82 with a standard deviation of \$0.84. Note that 25 students bid \$0 and 4 bid over \$10. The center graph depicts the distribution of Gambit bids after Gambit learns the prize is worth \$10. The modal bid is \$9; the average bid was \$6.73 with a standard deviation of \$1.04. One confused soul bid \$11. The bottom graph depicts the distribution of bids of Gambit after Gambit learns the prize is \$20. In this case \$10 was the modal bid and 34 students bid higher than \$10. The average bid was \$10.14 with a standard deviation of \$2.12.

Figure 10B.9 Empirical Results of Maxco-Gambit Bids (Harvard MBAs)

Maxco Bids with Unknown Prize

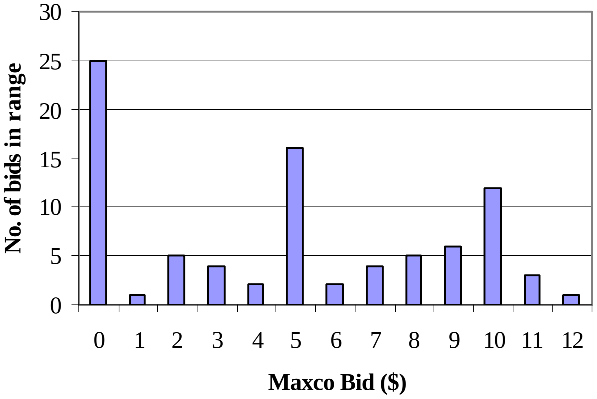

Average bid $4.82 Standard deviation $0.84

Gambit Bids with a $10 Prize

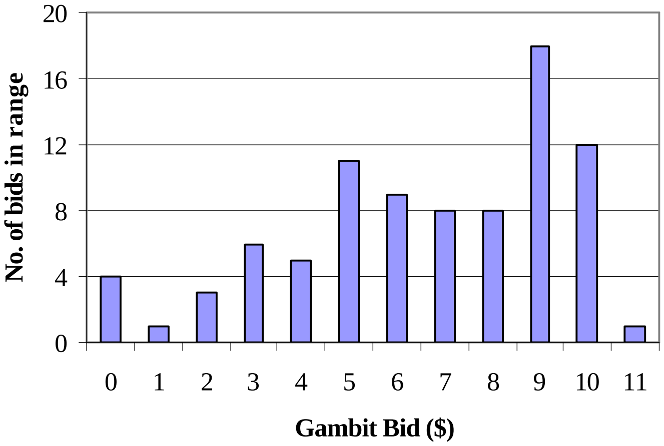

Average bid $6.73 Standard deviation $1.04

Gambit Bids with a $20 Prize

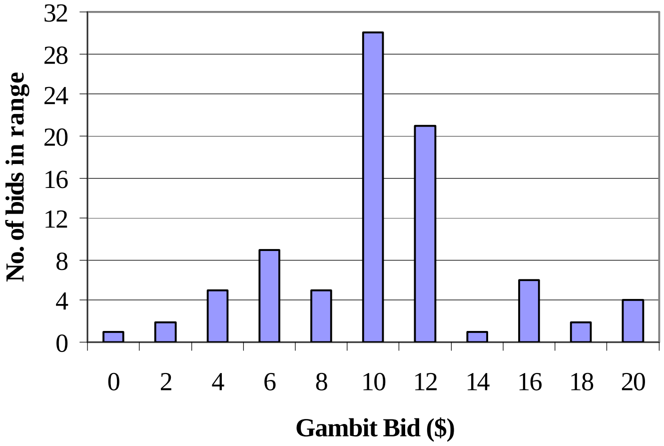

Average bid $10.14   Standard deviation $2.12

Expected Payoff To Maxco

Figure 10b.10 plots the expected return to Maxco against the empirical distribution of Gambit bids. The curve is a bit irregular but the message is clear. Positive Maxco bids yield negative expected payoffs! Against the empirical mix of Gambit bids, Maxco should bid $0.

Figure 10B.10 Expected Payoff to Maxco

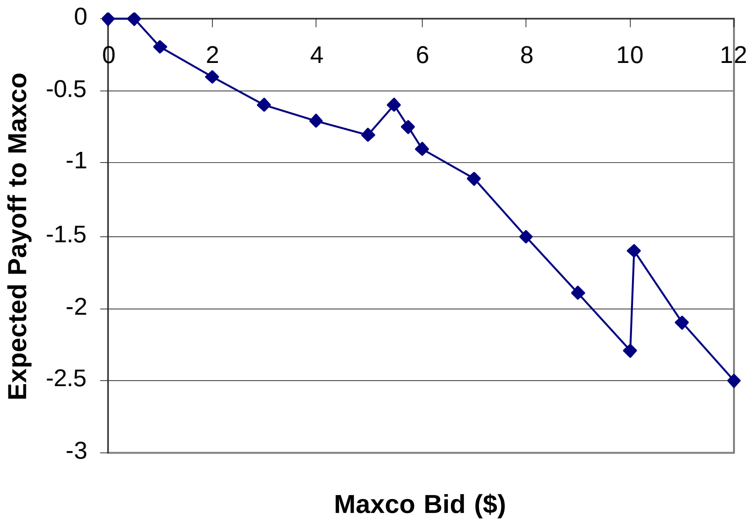

Expected Return to Gambit

Figure 10b.11 plots the expected return to Gambit when Gambit learns that the prize is worth $10 of Maxco bids. It is computed against the empirical mix of Maxco bids. The expected return to Gambit is rather flat for bids of $6 or less -- it peaks about $5.50. The $0 bid for Gambit is really more like $0.10; it beats all strictly zero bids of Maxco. Bids above $6 rapidly deteriorate. The actual modal bid of $9 for Gambit yields an expected return (against these Maxco bids) of $0.75.

Figure 10B.11 Expected Payoff to Gambit With a $10 Prize

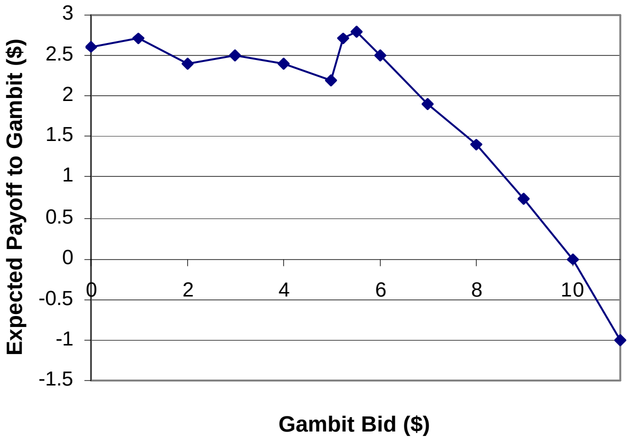

Expected Return to Gambit When the Prize is $20

Figure 10b.12 plots the expected return to Gambit when Gambit learns that the prize is worth $20. It is computed against the empirical mix of Maxco bids. The best bid for Gambit is a little over $10 -- just enough to outbid a lot of those misguided Maxco bids of $9 and $10. Very conservative bids of Gambit above $11 rapidly lose expected payoff for Gambit.

Figure 10B.12 Expected Payoff to Gambit with a $20 Prize

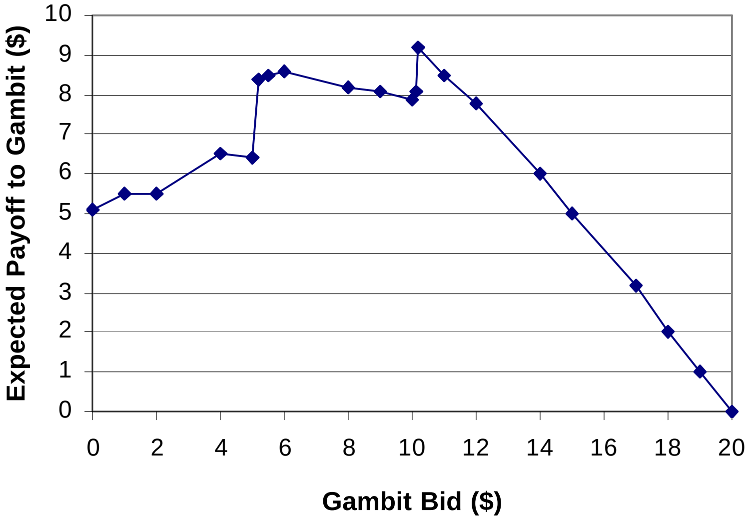

Now let’s reflect about what all this means. Let me start a dialogue with Ann and Jay.

EXP: Well, what do you think? What do these last three pages, depicting the empirical results, tell us?

Ann: Some of your students were confused.

EXP: Granted. You always have to expect some modicum of confusion and misunderstanding. But that happens in the real world as well as in the laboratory.

What do you think happens after I show these results to the bidders and then they play the game a second time? Why don’t you think about this before reading on.

After Feedback

EXP: I had my students replay this game and what do you expect happened?

Ann: A lot more Maxco players would bid $0.

EXP: That’s absolutely right. And Gambit players?

Jay: I would bet they would bid a lot lower than they did before, especially for the $20 prize.

EXP: That’s also true. And it also turns out that some Maxco players who live dangerously and bid about $5.50 steal a few big prizes from very aggressive Gambits.

Let’s speculate. What kind of operational advice can we give to the players that would be in equilibrium? The advice to each must be such that neither bidder would want to change knowing the other would follow our advice.

Ann: I don’t see how you could do this. If you tell Maxco to bid any amount like $3.00, then Gambit would want to let g_{10} and G_{20} equal to $3.10. But if Gambit were told to bid $3.10, then Maxco would want to bid $3.20.

Jay: It seems to me any advice to Maxco would be unstable if Gambit knew it.

EXP: So far I agree with you, but you haven’t thought hard enough yet. The trick is to tell Maxco and Gambit to use randomized procedures. It can be shown, but I won't do it here, that it is possible to give randomized advice to each that yields an equilibrium situation.

Jay: I'm not sure I follow you. You would tell Maxco to base his bid on a random drawing from some distribution. Like all bids should be equally likely between $0 and $10. Is that it?

EXP: Or to bid $0 with probability 0.4 say, and the spread the other 0.6 probability uniformly over the interval of $0 to $10. That's the idea. Jay.

Ann: And you're saying that if Gambit learns about this advice, but doesn't learn about the particular random drawing, then she should also follow your advice. But your advice to her must also involve random elements or else Marco could sneak in with a non-randomized response.

EXP: That's the idea. Such equilibrium advice is shown in an appendix.[2] It's not so easy to derive.

Jay: You indicate on the bottom of the page that if Maxco follows your advice, then his expected return is zero. That's scary. Why bother? Why not just bid zero?

EXP: Yup, you're right. But if Maxco were advised to bid zero, then Gambit could bid nice and low and we would go around in circles once again. The only equilibrium situation is as I indicated. Look, we've gone into this equilibrium behavior more than I think it warrants, so let's push on.

Jay: I'm not satisfied

EXP: Let's consider the use of <u>shut-out</u> bids. Would it ever be advantageous for Gambit to announce her strategy to Maxco before the game starts? By announcing her strategy, she would declare exactly what she would bid if the prize V equals $0, equals $10, equals $20.

Ann: I can't see how this could help Gambit.

EXP: Suppose Gambit announces that she will bid $0 if V=0; will bid a bit over $5.00, say $5.10, if V=$l0; and will bid a bit over $l0, say $10.10, if V=$20.

Jay: Well, if Maxco believed Gambit, then Maxco could bid just higher than $5.10, say $5.20.

EXP: What would Maxco's EV be in that case?

Jay: Well, if V were $0, he'd lose $5.20; if V were $10, he would win $4.80; if V were $20, he would lose the

---

[2] Not included in the preliminary version.

auction and get zero. Since each of these is equally likely, I guess his EV would be negative.

Ann: Maxco would be best off bidding zero.

EXP: Gambit's publicly announced strategy is called a *shut-out bid*. It is meant to intimidate Maxco.

Ann: But if the intimidation were to work, Gambit could surreptitiously bid just above $0.00, say 10¢ if V were $10 or $20.

Jay: Ah yes, but Maxco might expect that Gambit is unprincipled and conniving and bid say 40¢.

EXP: If Gambit threatened you, playing the role of Maxco, with intimidating shut-out bid, and if you could have a pre-play discussion with Gambit what would you say to him?

Ann: I don't see what Maxco could do.

EXP: How about if Maxco threatens to bid $10.20? If Gambit were locked in to his shut-out bid, then Maxco would steal the bid.

Ann: But why would he do that? He would lose $10.20, if V were 0; lose 20¢, if V were $10; and gain $9.80, if V were $20; which would yield a negative EMV.

EXP: How negative?

Ann: Not much. About negative 20¢.

EXP: And what happens to Gambit?

Jay: He would get nothing. Ah, I see what's going on. If Gambit can publicly announce a shut out bid Maxco can really sock it to Gambit and lose only 20¢ in the deal. And perhaps Maxco can demand some of the payoff to behave nicely.

EXP: All I'm trying to do is to stir up the pot a bit. Shut-out bids should not be taken lying down by the Maxcos of the world allowing all the gain to go to the intimidating party. Finally you might want to recall the discussion of the Ultimatum Game of Chapter 4.

An exercise for the mathematically inclined: Consider the Maxco-Gambit Problem with just two equally likely payoffs: $0 and $100. As before Gambit has privileged information and knows the value of the prize. There's common knowledge. Find the pair of randomized equilibria strategies. Hint: Try the randomized strategy for Maxco that puts some weight on a bid of $0 and distributes the rest in a way that makes any pure strategy of Gambit optimal in the range of $0 to $50. Try a similar trick for Gambit. The value of the game for Maxco turns out to be a disappointing $0. But nevertheless bidding $0 for Maxco cannot be a part of an equilibrium pair.

Core Concepts

We consider the distributive negotiation problem when a single seller is confronted with several potential buyers. (Or a single buyer with several sellers.) Face-to-face negotiations can now be replaced by several types of auctions or competitive bidding mechanisms. We introduce several of these designs and examine them from both a game-theoretic, decision analytic, and behavioral perspective and both from the orientation of a bidder and the auctioneer.

In the open, ascending outcry auction a bidder should (but often doesn't) prepare for the auction by determining her breakeven value. There is no need for an analytically inclined bidder to assess a probability distribution of the Maximum Bid Of Others -- a so-called MBOO analysis for this case. A MBOO analysis is central to the decision analytic approach - but not game theoretic approach -- of the Dutch (descending) auction and the competitive, sealed-bid mechanism.

The high-bidder-wins-at-the-second-high-price auction (or philatelist or Vickery auction) has the nice feature that it is optimum to bid one's breakeven value regardless of what other bidders choose to do. It does not require a MBOO analysis to act optimally.

The chapter then considers two special cases: (1) reciprocal buy/sell bids that are especially useful when two business partners seek to dissolve their partnership; (2) combinatorial bids such as the FCC auction when the federal government auctioned off 2500 licenses for radio frequencies. The complicating feature is that for many bidders the value of a given lease depends on the other leases that bidder wins.

The central theme of Part B deals with subtle issue of conditional probabilities. There is often information to be gleaned from the prior choices of others.

Problem 1 considers the case where the prize is common but uncertain. In bidding for a transparent jar filled with coins, the winner of the bid would very often be someone who had an extreme perception of the value of the jar. The winner often finds out that he or she has paid too much. Thus the winner (of the prize) turns out to be a financial loser -- thus the "Winner's Curse."

Problem 2 considers the (Bazerman-Samuelson) case where an *Acquiring* firm submits a bid to a *Target* firm that can either accept the offer or reject it. The target firm has privileged information about the value of the firm. Thus, the acquirer's conditional probability distribution for the value of the firm should depend on whether the target accepts or rejects the offer. In the case discussed, the acquirer should not want any deal the target would accept. This subtlety is missed by most subjects.

It's another case of subjects becoming confused about conditional probabilities and is somewhat related to the Monty Hall example in Chapter 3.

Problem 3 is abstracted from the case of two oil wildcatters bidding for a lease when one bidder has privileged information about the chances and extent of the oil deposits.
In the Maxco Gambit Cases both bidders have identica1 probability distributions of the value of the prize but before the bidding takes place Gambit gets privileged information about the value of the prize. This raises the question of the value of information in a game setting. Part of the value comes from an intimidation factor: the less knowledgeable player may be wise not to bid or to bid low and the party with the added information can exploit this knowledge.

One Appendix does an equilibrium Analysis of the both-pay competitive sealed-bid problem. The both-pay ascending auction, or the so-called escalation game, was featured in Chapter 9. Another Appendix does equilibrium analysis for a slight simplification of the Maxco-Gambit Problem. In this case it is equally likely that the prize is \$0 or \$100 and Gambit, not Maxco, has this privileged information.

Appendices

A. On Choosing a Random Quantity with a Given Cumulative Distribution Function

Let X denote the random quantity (variable) with CDF F(.). By definition, for any constant x,

$$F(x) = \text{Probability}\ \{ X \leq x \}\ .$$

In order to select a random value for X governed by F, we proceed as follows:

Let R be a random number on the interval from 0 to 1.0 – all values on the unit interval being equally likely. Independent R-values can be generated in EXCEL by the command +RAND. For any single R drawn,
say r, obtain corresponding x from the equation

$$r = F(x) = \text{Probability}\ \{ X \leq x \}\ ..$$

Let

$$F(x) = x^2, \quad \text{for } 0 \leq x \leq 1.$$

To find a random number, x, governed by F, we solve

$$r = x^2$$

or

$$x = r^{½},$$

where r is a random number on the interval from 0 to 1.

B. The All-Pay Competitive Bid with a Common Known Prize

Rules of the Game. Each of n + 1 players must submit a sealed bid for a $100 prize. The top bidder wins the $100 but all bidders sacrifice the value of their bid.

Illustration: Let n =2 so that there are three bidders. Assume the bids are as shown :

Bidder	Bid	Payoff
A	$23	- $23
B	$48	+$52
C	$19	- $19

What would you bid ?

Equilibrium Analysis.

A little reflection indicates that the players must choose a randomized (mixed) bid on the interval from $0 to $100. We seek a symmetric solution where each bidder uses a common Cumulative Distribution Function F. If the other n players each choose F and you bid an amount b, your expected value payoff is

$$EV(b \mid F) = -b + 100 \ \text{Prob}\{ \text{all n bidders choose values} \leq b \}$$

$$= -b + 100\,[F(b)]^{n} = -b + 100\,F^{n}(b).$$

Now any bid b in the interval from $0 to $100 must be a good bid against this common F, or else you yourself would not use F; hence it must be that

$$EV(b \mid F) = \text{constant, } K \text{ (say)},$$

so that

$$-b + 100\,F^{n}(b) = K, \quad \text{for all } 0 \leq b \leq 100 .$$

Now since, by definition, F(100) = 1, we must have

$$-100 + 100(1) = K, \ \text{or } K = 0,$$

and substituting this value for K we get

$$F^{n}(b) = b/100$$

or

$$F(b) = [\,b/100\,]^{1/n} .$$

To use the mixed (randomized) strategy F, you would generate the random number r and solve

$$r = [\,b/100\,]^{1/n} \quad \text{or} \quad b = 100\,r^{n} .$$

So we see that as the number of bidders increase, you should bid more cautiously but your equilibrium expected payoff remains 0 for all n.

For n =1 , i.e., you are bidding against one other bidder, you should choose equally likely bids from $0 to $100. If you choose bid b, you get a b/100 chance of winning 100 so that your expected up side just cancels the downside.

This does not look like a good game to play. So many bidders will either not bid (i.e., choose b = 0), or choose $b < 100\ r^n$, and in this case perhaps you should bid even more aggressively than the equilibrium theory suggests. But if this logic appeals to you, it might appeal to others and therefor perhaps you shouldn’t bid after all. But we can go on … but we won’t.

Supplement to *Negotiation Analysis; the Science and Art of Collaborative Decision Making* by Howard Raiffa (with contributing authors David Metcalfe and John Richardson), Harvard University Press, 2003.

Ch 14 Addendum On Risk Sharing

Mr. George is an oil wildcatter. His past diligence and business acumen have assured him a good reputation, and he now enjoys the right to drill for oil at a given site. The trouble is that he has liquidity problems: most of his money is tied up in other risky ventures and his credit rating at the bank is not favorable. The cost of drilling is uncertain, but he has the possibility of taking seismic soundings at the site which will yield some information—but not perfect information—about the possibilities of finding oil. He could plunge all his financial resources into this deal and go it alone; or he could borrow more money at the bank; or he could cut others into the deal, either by means of a straight proportional sharing of profits and losses or, perhaps, a different proportional sharing on the up and down sides of the deal. Let's assume for the moment that all contingent, financial sharing arrangements are secure—that all contracts are inviolable both in law and in the intent of the protagonists—and look at one way in which this problem can be abstracted into a risk-sharing negotiation problem. The terrain we're about to enter into is so vast, including as it does financial markets, equity financing, insurance, and reinsurance, that we must be careful not to get lost in its intricate byways.

Mr. George approaches Mr. Lloyd, a speculator, to share his risky venture with him. They examine their options and identify one strategy that appears promising, but the payoffs are uncertain: these depend on the (uncertain) cost of drilling, on how much oil is down there, on how easy the oil is to recover, on future regulations, on future oil prices, and on a lot more. To simplify, we'll say that they depend on which one of five states of the world—A, B, C, D, or E—will prevail. (If we were to be more realistic, we might use something like five thousand states of the world.) Mr. Lloyd consults his own experts and obtains probabilistic assessments of the five potential outcomes; these differ from Mr. George's assessments as shown in Table 14A.1. All assessments are kept confidential.

TABLE 14A.1. ***Potential outcomes and sharing rules.***

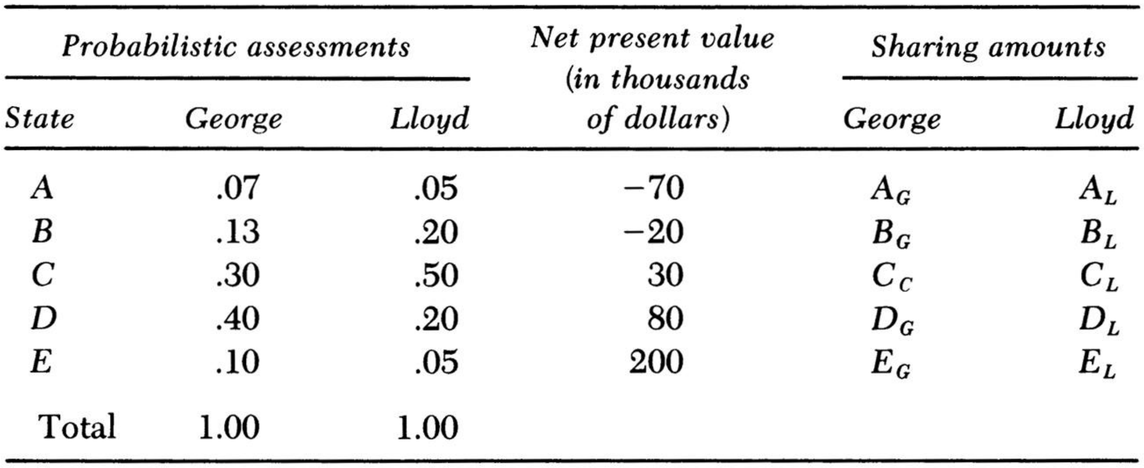

Probabilistic assessments			*Net present value (in thousands of dollars)*	*Sharing amounts*	
State	*George*	*Lloyd*		*George*	*Lloyd*
A	.07	.05	−70	A_G	A_L
B	.13	.20	−20	B_G	B_L
C	.30	.50	30	C_C	C_L
D	.40	.20	80	D_G	D_L
E	.10	.05	200	E_G	E_L
Total	1.00	1.00			

George and Lloyd, however, do agree on the financial implications of the deal, conditional on a given future state prevailing.[1] A dry hole (states), for instance, would lead to a loss of $70,000, or an abortive attempt after a negative seismic sounding (state B) would lead to a loss of $20,000.

George and Lloyd have to agree on how to share the financial proceeds in each of the five states. If state A unfortunately occurs, then the team will lose $70,000. George, who is short of funds, will want Lloyd to assume most of this loss. But, of course, Lloyd is not going to agree with this unless his own shares are sufficiently high for the states C, D, and E. He also might want George to share in some of the losses if A or B occurs, just to keep George honest—or, more felicitously put, to give George the right incentives.

George and Lloyd have to decide how to share the loss of $70,000 if state A occurs. In order to keep our notation symmetric (which makes it easier to generalize to more than two risk sharers at a later stage), let's suppose that George and Lloyd have to select two numbers: A_G, the payoff to George if A occurs, and A_L, the payoff to Lloyd if A occurs (all payoffs

---

[1] If they were to disagree on the financial consequences associated with a given state, then they could decompose that state into two or more states with differing probabilities. Our present format (including additional states) is thus quite general. For example, if George thinks the payoff in state C is $30,000 and Lloyd thinks it is $50,000, then state C could be split into two states, C' and C'', with payoffs $30,000 and $50,000, respectively. George may assign probabilities of .3 and zero to C' and C'', whereas Lloyd may assign probabilities of zero and .5 to C' and C'', respectively.

are in thousands of dollars). We then require that $A_G + A_L = -70$. They have to decide analogously on the splits in cases B. C, D, and E. SO overall, George and Lloyd have to decide on ten numbers: A_G, A_L,…, E_G, E_L (see Table 14A.1), subject to the following set of five constraints:

$$\begin{aligned} A_G + A_L &= -70, \\ B_G + B_L &= -20, \\ C_G + C_L &= 30, \\ D_G + D_L &= 80, \\ E_G + E_L &= 200 . \end{aligned}$$

For any determination of these ten numbers, George and Lloyd will each be confronted with a risk profile. George's risk profile will yield financial prizes A_G, B_G, . . . , E_G with probabilities .07, .13, . . . , 10, respectively; Lloyd's risk profile (see chapter 2) will yield financial prizes A_L, B_L, . . . , E_L with probabilities .05, .20, . . . , 05, respectively. Their reactions to these lotteries will depend on their attitudes toward risk taking. It could be that a specific risk-sharing plan (determined by a specific setting of the ten numbers)[2] is inefficient in the sense that the ten risk-sharing numbers could be changed to improve the risk profile for each party (in that party's subjective opinion). In other words, there may be

---

[2] Because of the five financial constraints, these ten numbers have really five degrees of freedom: once we determine what Lloyd gets in each state, George gets the complement.

opportunities for joint gains. Figure 14A.2 depicts graphically what could occur. For a specific risk sharing plan, Q (which arises through the specification of ten legitimate numbers), George might assign a certainty equivalent to his resulting risk profile of $5,000, and Lloyd might assign a certainty equivalent to his resulting risk profile of $13,000. However, as depicted, the risk-sharing plan Q is not efficient: they both can improve, since there are joint evaluations of risk-sharing deals that fall northeast of Q. George controls the ownership of the deal and can remind (threaten?)

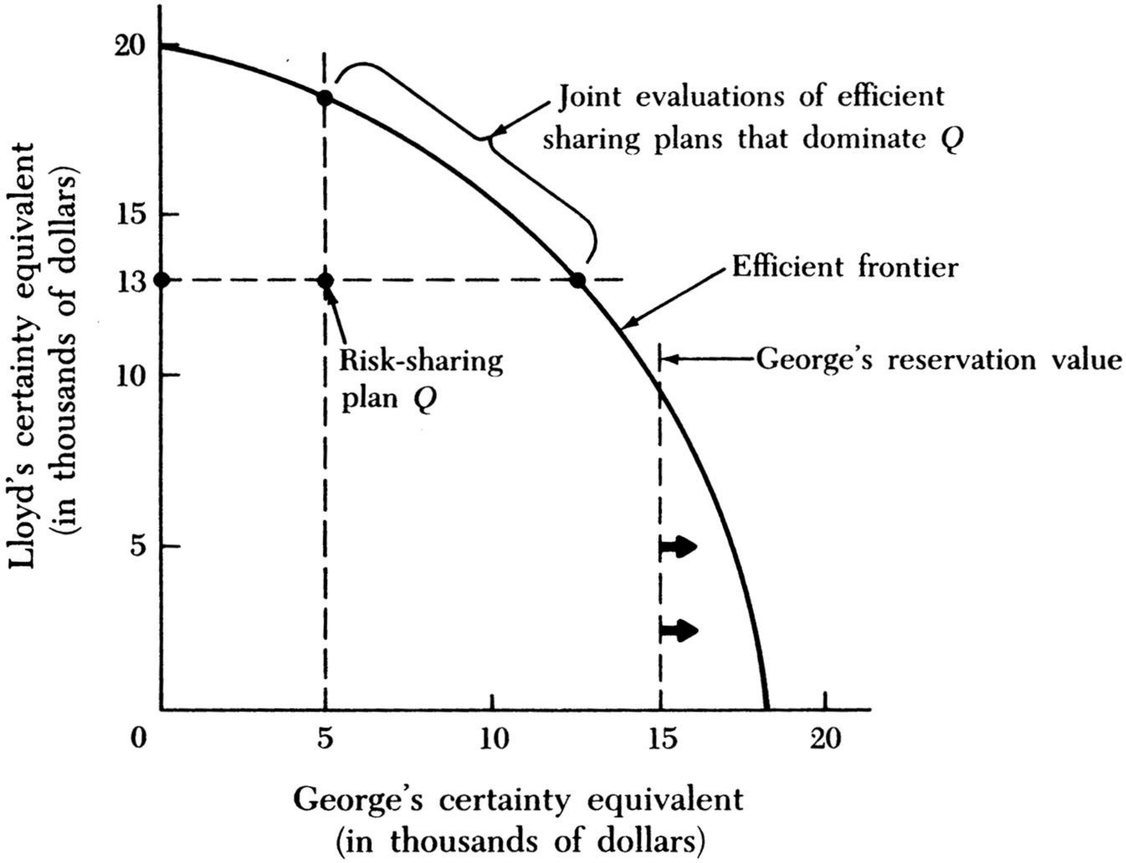

Figure 14.2. Set of joint evaluations of risk-sharing agreements.

Lloyd that there are other speculators who would love to join in the venture. Lloyd could counter that he, too, has a choice of other potential drilling deals. They also can remind each other about the transaction costs of starting negotiations with other partners and how nice it would be to work on other deals together in the future. The point of all this is that Lloyd and George are involved in a negotiation that bears strong similarities to other negotiations we have considered. This is not the place to discuss details about how such sharing procedures are made or could be made more efficiently. As in most negotiation processes, the protagonists have to worry about their alternatives if they find it impossible to come to an agreement. Each must consider the other external opportunities available to him before he can arrive at a reservation price for the present set of negotiations.

Suppose that George will deal with Lloyd only if he can get a certainty equivalent of at least $15,000 from a mutually agreed upon risk sharing deal; in other words, George's reservation price is $15,000. As shown in Figure 14A.2, it may be possible to satisfy George and still get a positive return for Lloyd, but there's not much leeway. They may never find sharing arrangements that are mutually acceptable, even though such agreements might exist.

The more structured the market, the easier it becomes to assess objectively these externally driven reservation prices. We do not engage in much haggling with our insurance companies when we obtain collision insurance for our automobiles. But the owner of an oil tanker entering into troubled waters might have some negotiating leverage with his insurance suppliers, and vice versa. By and large, we can think of these negotiations as collaborative decision making.

Supplement to *Negotiation Analysis: The Science and Art of Collaborative Decision Making* by Howard Raiffa (with contributing authors David Metcalfe and John Richardson), Harvard University Press, 2003.

Strictly Opposing Two-Player Games

by Howard Raiffa

Strategy of Presentation

We don't seek generality but present enough of the theory in a special case to enable the reader to generalize to more general cases.

We'll continue with the development in Chapter 4 of *Negotiation Analysis.* Consider the problem of two protagonists: Rowena and Colin who are strictly opposing. If Rowena prefers A to B then Colin prefers B to A.

Let Rowena have two pure strategies R1 and R2 and Colin have strategies C1 and C2 as was the case for the bi-matrix games of Ch. 4. For any (Ri, Cj) pair we need only display Rowena's utility payoffs since Colin's motivation is to minimize Rowena's payoff from the game.

To emphasize the point that the payoffs are in utility units and that utilities are essentially breakeven probabilistic judgments we'll interpret the payoffs as follows: the payoff to Rowena for the pair (Ri, Cj) is a utility number Aij between 0.0 and 1.0 which gives Rowena an Aij standard probability at some desirable prize W and a complementary probability at the status quo. Think of Colin supplying this prize to Rowena so that the higher the number Aij the better for Rowena and the worse for Colin.

The initial case we examine has the structure

	C1	**C2**
R1	A11	A12
R2	A21	A22

Rowena must choose either strategy R1 or R2; Colin has the choice of C1 or C2. The four Aij numbers are known fully by the players; in the parlance of game theory, the payoff numbers are common knowledge.

We consider two qualitatively different games:

Table 1

Game 1	C1*	C2	Security Level
R1*	.4	.5	.4*
R2	.3	.2	.2
Sec. Loss	.4*	.5	Maximin = Minimax

Table 2

Game 2A	C1	C2	Security Level
R1	.1	.9	.1
R2	.7	.2	.2*
Sec. Loss	.7*	.9	Maximin < Minimax

Analysis of Game 1

In Game 1 there is an equilibrium pair (R1, C1): If Rowena holds fixed at R1 there is no motivation for Colin to move from C1 to C2; and if Colin holds fixed at C1 there is no motivation for Rowena to move from R1. No other pair is in equilibrium. In Game 1, by playing R1 Rowena can guarantee herself a return of .4 units. By playing C1 Colin can guarantee that Rowena can get at most .4 units. No other pair is in equilibrium. Also note that the .4 entry is the max of its column and the min of its row.

In Game 1, Rowena, by playing R1, has a security level of .4; that .4 is the minimum of row R1. By playing R2 Rowena has a security level of .2 which is the minimum of row R2. The two row minima are .4 and .2 and the maxima of these two row minima is .4. Turning to Colin’s viewpoint: the maximum that Colin can lose playing C1 is .4 (the max of column C1; and the max he can lose with C2 is .5. Thus .4 and .5 are Colin’s security losses and if he minimizes his security losses (i.e., if he minimizes his maximum potential losses of .4 and .5 , he’ll choose C1 which guarantees maximum loss .4 units.

Game 1 has a value of .4 and Rowena by playing the strategy R1 that maximizes the row minima – her so-called maximin strategy – can guarantee herself at least .4; and Colin, by playing the strategy C1 that minimizes the column maxima – his so-called minimax strategy – can guarantee that Rowena gets at most .4. Nice neat little reinforcing bundle! How much different things are going to be when we turn our attention to Game 2.

Analysis of Game 2

Check the following observations:

- There are no equilibria pairs in pure strategies. No payoff entry for Rowena is both the maximum of its column and the minimum of its row.

- The minima in each row (labeled security levels) are .1 and .2 and the maximum of these minima is .2. Thus the maximin value is .2 and the maximin choice for Rowena is R2.

  The maxima in each column (labeled security losses) are .7 and .9 and the minimum of these maxima is .7. Thus the minimax value is .7 and the minimax choice for Colin is C1.

- The maximin value of .2 – the value that Rowena can secure for herself – is less than the minimax valueof .7 – the value that Colin can secure for himself.

  Figure 1:

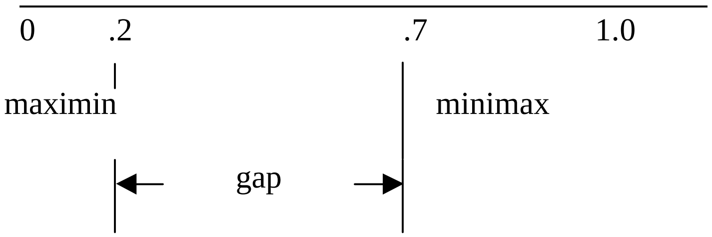

Randomized (or Mixed) Strategies

Suppose that Rowena, instead of choosing a pure strategy of R1 or R2, tosses a coin: heads she chooses R1 and tails R2. This is a so-called randomized (or mixed) strategy which chooses R1 with standard probability of x = .5 and complementary probability of 1 – x or also .5. Shortly we'll consider the class of mixed strategies for x values between.

If Rowena uses Mixed Strategy with x = .5, which we label MS(.5), then her payoff against C1 is .5 x .1 + .5 x .7 = .4 ; and her payoff against C2 is .5 x .9 + .5 x .2 = .55. Note – and this deserves some thought – that all probabilities are standard or what we have called canonical. We have studiously avoided, thus far, in any use of subjective probabilities.

Table 3

Game 2B

	C1	C2	Security Level
R1	.1	.9	.1
R2	.7	.2	.2
Toss Coin MS(.5)	.4	.55	.55*
Security Loss	.7*	.9	

By introducing MS(.5) we have raised Rowena's maximin value (her security level) from .2 to .55 and partially closed the gap between the minimax and the maximin values.

We now have some work cut out for ourselves and some conjectures to be made:

- We have introduced for Rowena a new class of mixed strategies – the MS(x) strategies. We now can manipulate x to achieve the optimal maximin value.
- We can do the same trick for Colin: let him use C1 with probability y and C2 with probability 1- y trying to improve his security losses. We should suspect that by randomization he could improve (i.e., lower) his minimax value.
- Wouldn't it be "loverly " if the gap between minimax and maximin values vanished with proper choices of x and y. In this eventuality, the nice story for game 1 would be repeated for Game 2 once mixed strategies are introduced. There will be some value, v*, such that Rowena, by choice of x*, can guarantee herself at least v* and Colin, by choice of y* can guarantee that Rowena get at most v*. Furthermore the optimal x – let's call it x* – and the optimal y* are in equilibrium. We'll see that's the case.

Geometry

Now to explain Fig. 1. The horizontal axis depicts the value of x, the probability of the mixed strategy choosing R1. Hence as x goes from 0.0 (choice of R2) to 1.0 (choice of R1) the payoff against C1 goes from .7 on the left to .1 on the right; and against C2 the payoff goes for .9 on the right to .2 on the left.

Now suppose Rowena were to announce she is going to use MS(.5) – i.e., she is going to toss a fair coin. So we go to .5 on the horizontal axis and Colin has the choice the C1 or the C2 lines. As the minimizing player, Colin will choose the C1 line against MS(.5).But against x = .2 (say), Colin would prefer C2. Hence the broken darkened line represents the payoff to Rowena as a function of x assuming that Colin chooses the best retort from him. Thus the darkened line is the minimum function (or the security function) that Rowena can achieve as a function of her chosen x. We see, and a little algebra shows that Rowena should choose x* = 5/13 for her maximin value that will guarantee her a return of 4.692. These numbers can be read off a precisely drawn figure but they also can be derived algebraically as follows: We require

$$.1\, x^* + .7\,(1- x^*) = .9\, x^* + .2\,(1 - x^*)$$

and solving, we get x* = 5/13 = .385. The maximin value v*, is thus

$$v^* = .1\,(5/13\,) + .7\,(8/13\,) = 6.1/13 = .4692.$$

End of story for Romena’s analysis and now on for more about Colin’s analysis.

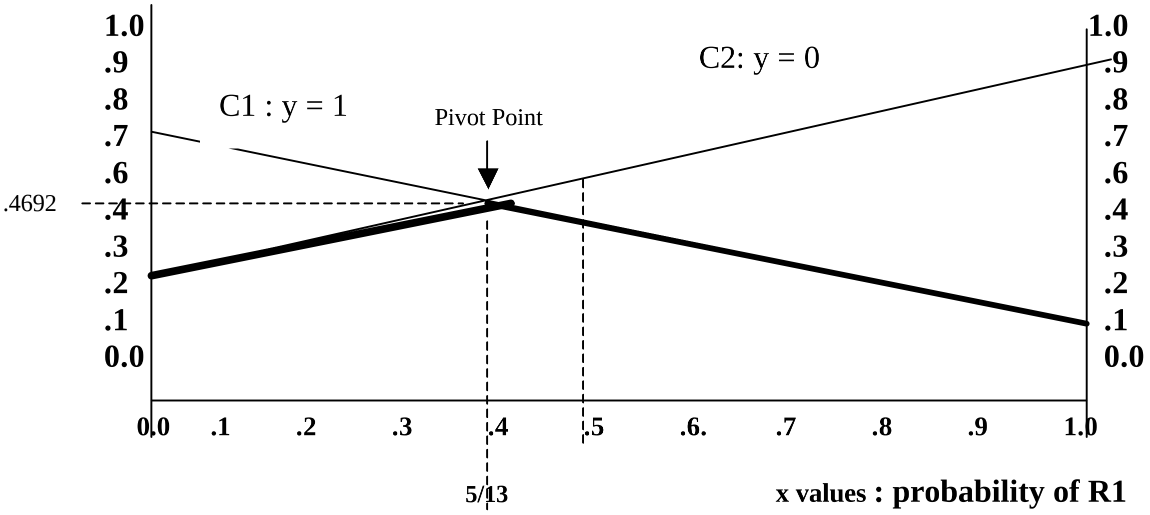

Figure 2:

Colin’s Analysis

We could, but we won’t, present a figure for Colin analogous to Rowena’s. Instead of the broken linear minimum function, Colin would be presented with a maximum function and he would choose the minimum of the maximum function. Instead we use Rowena’s geometry to develop Colin’s analysis.

Let y designate the probability that Colin uses C1, letting 1 – y be the probability of C2. Now here is the key: for any specific y Colin uses, the return to Rowena can be shown as a line in Fig.1 that goes through the pivot point. Indeed the lines for y = 0 and for y = 1 are already shown on that diagram. Let’s examine the line we would get for y = .5. It would be midway between lines for y = 0 and y= 1 but not quite horizontal. For y = .5 the intercept on the left (the R2 side) would be .5 x .7 + .5 x . 2 = .45; and the right intercept (the R1 side) would be at .5 x .1 + .5 x .9 = .5. Hence the line for y =.5 would start on the left at a height of .45 and tilt slightly upwards , passing through the pivot point and reaching a right of .5 at the right. If Colin were to announce y = .5, Rowena would choose the right end point (corresponding to x =1 or to R1). Colin would like to choose y such that the line it offers to Rowena is horizontal going through the pivot point (as all lines for different y values do). To find this y we set

$$.1\,y + .9\,(1 - y) = .7\,y + .2\,(1 - y)$$

or

$$y = 7/13 = .538$$

and the height of the line is

$$.1\,(7/13) + .9\,(6/13) = 6.1/13 = .4692$$

The following display is helpful in keeping these calculated numbers in mind:

Table 4

		7/13	6/13	Colin’s minimax strategy
Rowena’s maximin strategy	5/13	.1	.9	6.1/13
	8/13	.7	.2	6.1/13
		6.1/13	6.1/13	

The calculations say: Rowena can guarantee a payoff of 6.1/13 by playing her maximin strategy. Colin can guarantee that Rowena gets at most 6.1/13 by playing his minimax strategy. Rowena's maximin and Colin's minimax strategies are in equilibrium in the sense that there is no motivation for either to change if the other holds fixed. End of story.

Extensions to Bigger Games

[Use of an EXCEL Linear programming package to find Rowena's maximin strategy and Colin's minimax strategies and the resulting value of the game with more than two strategies for each player.]

Let m and n represent respectively the number of pure strategies for Rowena and Colin. We have just completed the story for m = n = 2. We now investigate the analysis for more general m and n.

First let us illustrate a game that has an equilibrium pair in pure strategies. See Table 5. We want to continue interpreting the payoff numbers as utilities so think of the payoffs as chances out of ten that Rowena will receive a desirable prize from Colin.

Table 5

	C1	C2	C3	C4	Security Level
R1	1	3	4	1	1
R2	6	**5**	8	7	**5** maximin
R3	4	4	0	2	0
Security Loss	6	**5** minimax	8	7	

In the game of Table 5, m = 3 and n = 4. We assume the players have complete knowledge of the structure and payoffs; and each knows the other knows etcetera. That's what we mean by common knowledge.

Rowena's security levels for her pure strategies are the minima of the rows and the maximizer of these row minima is R2. Her maximin value in pure strategies is 5. She can guarantee at least 5 units for herself by playing R2. Colin's security losses are the maxima of each of the columns and the minimizer of these column maxima is C2; he can guarantee, by playing C2 that Rowena gets no more than 5 units. In this case the maximin equals the minimax.

The pair (R2, C2) is in equilibrium since R2 is best for Rowena against C2; and C2 is best for Colin (the minimizing player) against R2. Observe also that the (R2, C2) entry of 5 units is the maximum of its column and the minimum of its row. This summarizes what need be said about strictly competitive games with pure strategy equilibria.

Games with no equilibrium pair in pure strategies. Mixed strategies to the rescue.

Consider the game depicted in Table 6. The row minima are 1, 1, 2 and the maximum of these minima is 2, Rowena's maximin value in pure strategies. The maxima in each Column are respectively 5, 3, 4, 6 and the minimum of these column maxima is 3.Thus Colin can secure, by choice of C2, not to lose more than 3 units but Rowena can only command 2 units for herself. There is a gap between the maximin value and the minimax value. We also note that there is no equilibrium pair – no entry is both max in its column and min in its row.

Table 6 No Equilibrium Pair

	C1	C2	C3	C4	Row Minima	
R1	1	3	4	1	1	
R2	5	2	1	6	1	
R3	4	3	3	2	2	Maximin
Column Maxima	5	3	4	6		
		Minimax				

We will now show that Rowena can increase her security level – i.e., increase her maximin value by use of randomized strategies.

Consider for example the randomized strategy that assigns equally likely values to R1, R2, and R3 – i.e., assigns a .333 prob to each pure strategy as shown in Table 5.

Table 7

Randomized Strategy		**C1**	**C2**	**C3**	**C4**	Row Security
0.333	**R1**	1	3	4	1	1
0.333	**R2**	5	2	1	6	1
0.333	**R3**	4	3	3	2	2
random strategy		3.333	2.667	2.667	3.000	2.667

If Rowena uses this randomization, then her expected return against any column chosen by Colin is the simple average of the entries in that column. Hence her payoffs in utility terms are as shown in the last row of Table 7 and the security level of that randomized (or mixed) strategy is the minimum of those entries, a whopping 2.667 value. Can she do better by using a different randomization? How much can she aspire to? Certainly not above Colin's minimax value of 3? Can she aspire to get a secutity level of 2.9? This becomes now a mathematical programming problem. First let us note that the value for Rowena's entry against C1, which we already know is the average of the C1 column of numbers, is also the sumproduct of the strategy column and the C1 column. Indeed the payoff of any legitimate randomization against any column is merely the sumproduct of those two columns of numbers.

Use of EXCEL to Compute Rowena's Highest Security Level Using Mixed Strategies.

We first set up Table 8. We choose a ballpark aspiration level of 2.5. That number is used as a convenience to set the calculations and is arbitrary at this time. Our aim is to find a legitimate randomized strategy – a set of three non-negative numbers that sum to 1.0 – that will yield expected payoffs against C1 to C4 that are each in excess of some pre-specified aspiration level. In this case we see that the randomized strategy as shown yield numbers that are each in excess of the aspiration value of 2.5. These excesses are shown in the last row of numbers in Table 6. The problem is: Given the setup in Table 8, maximize the aspiration-level cell and find a set of legitimate randomization values for which each of the excesses is non-negative. This is a linear programming problem ideally setup to use SOLVER, an attachment on EXCEL.

Table 8

Aspiration **2.5**

Randomized Strategy		**C1**	**C2**	**C3**	**C4**
0.333	**R1**	1	3	4	1
0.333	**R2**	5	2	1	6
0.333	**R3**	4	3	3	2
Random strategy		3.333	2.667	2.667	3.000
Excess		0.833	0.167	0.167	0.500

Sumproduct of randomized strategy
column and C1 column

When SOLVER was queried it responded by showing Table 9, which shows how Rowena could choose a randomized strategy yielding her maximum possible aspiration of 2.75.

Table 9

Aspiration **2.75**

Randomized Strategy		**C1**	**C2**	**C3**	**C4**
0.250	**R1**	1	3	4	1
0.250	**R2**	5	2	1	6
0.500	**R3**	4	3	3	2
1.000					
random strategy		3.500	2.750	2.750	2.750
Excess		0.750	0.000	0.000	0.000

sumproduct of randomized strategy
column and C1 column

Colin's Analysis

We start the process in Table 11. We choose a representative random strategy for Colin as shown in B3:E3. The sum of those four probabilities is shown in cell F3, which in this case is 1.0. The expected payoffs of this randomized strategy against R1 is 2.3;

against R2 is 3.6; against R3 is 2.7. These are sumproducts of the randomized strategy and the row payoffs. With a representative arbitrary aspiration level of 2.9 (pulled out of the air) we see that Colin cannot achieve this aspiration level using the displayed random strategy (B3:E3) in because the excess with R2 I s negative. Remember that Colin is the minimizing player and the excesses are the aspiration value minus the EV columns. We see that as pitted against R2 the resulting EV of 3.6 exceeds the aspiration level of 2.9.

	A	B	C	D	E	F	G
1	**Table 11**		Aspiration	2.9			
2							
3	C's Rand. Str.	0.1	0.2	0.3	0.4	1	
4		C1	C2	C3	C4	EV	Excess
5	R1	1	3	4	1	2.3	0.6
6	R2	5	2	1	6	3.6	-0.7
7	R3	4	3	3	2	2.7	0.2

Colin's Linear Programming Problem:

Given the payoff table (in B5:E7);

find

legitimate values for the four adjustable numbers in B3:E3

– – i.e., numbers that are non- negative and whose sum appearing in cell F3 is 1.0

and an aspiration level in D1 ;

to maximize

the entry in D1

subject to

keeping the calculated values (the excesses) in F5:F7 non-negative.

SOLVER takes a second or so to produce the answer shown in Table 12. Colin can hold Rowena down to at most 2.75 by playing his minimax randomized strategy.

Table 12 Aspiration **2.75**

Colin's minimax Rand. Str.	0	0.5	0.25	0.25	1	
	C1	C2	C3	C4	EV	Excess
R1	1	3	4	1	2.75	0
R2	5	2	1	6	2.75	0
R3	4	3	3	2	2.75	0

The dual solution – Rowena's as well as Colin's – appears in Table 13.

Table 13

	A	B	C	D	E	F	G	H
1								
2								
3		0	0.5	0.25	0.25	1	Row's EV	excess
4	**0.25**	1	3	4	1		2.75	0
5	**0.25**	5	2	1	6		2.75	0
6	**0.5**	4	3	3	2		2.75	0
7	1							
8	Col's EV	3.5	2.75	2.75	2.75		**2.75**	

That completes the story except for a proof that argues that the best aspirations for Colin and Rowena are always the same. The gap with randomization between the maximin and minimax values is zero. We know how to make this result convincing but it's too specialized for our needs and is thus omitted.

Supplement to *Negotiation Analysis; the Science and Art of Collaborative Decision Making* by Howard Raiffa (with contributing authors David Metcalfe and John Richardson), Harvard University Press, 2003.

Strictly Opposing Two-Player Games: Security Analysis

Fig. 1 : (R2,C2) in equilibrium

	C1	**C2**	**C3**	**C4**
R1	**1**	**3**	**4**	**1**
R2	**6**	**5**	**8**	**7**
R3	**4**	**4**	**0**	**2**

Max of Col. and Min of Row

Fig 2 :

	C1	**C2**	**C3**	**C4**	Row Security	
R1	1	3	4	1	1	
R2	5	2	1	6	1	
R3	4	3	3	2	2	Maximin
Column Security	5	3	4	6		
		Minimax				

Fig. 3

	C1	**C2**	**C3**	**C4**	Row Security	
R1	1	3	4	1	1	
R2	5	2	1	6	1	
R3	4	3	3	2	2	Maximin
Column Security	5	3	4	6		
		Minimax				

No equil. pair in pure strategies

Fig. 4 : Intro of Randomized str.

Strategy		C1	C2	C3	C4
x1	R1	1	**3**	**4**	**1**
x2	R2	**5**	**2**	**1**	**6**
x3	R3	**4**	**3**	**3**	**2**
1	EV				

Sum products of Xs and Cs

Fig. 5:

Table 6

Aspiration **2.75**

Randomized Strategy		**C1**	**C2**	**C3**	**C4**	Row Security
0.250	**R1**	1	3	4	1	1
0.250	**R2**	5	2	1	6	1
0.500	**R3**	4	3	3	2	2
1.000						
random strategy		3.500	2.750	2.750	2.750	
Excess		0.750	0.000	0.000	0.000	

sumproduct of randomized strategy column and C1 column

How far can we push up the aspiration of 2.5 and still keep all excesses positive?

Aspiration **2.75**

Randomized Strategy		**C1**	**C2**	**C3**	**C4**
0.250	**R1**	1	3	4	1
0.250	**R2**	5	2	1	6
0.500	**R3**	4	3	3	2
Randomized	Strategy	3.500	2.750	2.750	2.750
Excess		0.750	0.000	0.000	0.000

Table 9 Aspiration 2.75

	0	0.5	0.25	0.25	1	
	C1	**C2**	**C3**	**C4**	EV	Excess
R1	1	3	4	1	2.75	0
R2	5	2	1	6	2.75	0
R3	4	3	3	2	2.75	0

Table 9 Aspiration 2.75

Rand. Str.	0	0.5	0.25	0.25	1	
	C1	**C2**	**C3**	**C4**	EV	Excess
R1	1	3	4	1	2.75	0
R2	5	2	1	6	2.75	0
R3	4	3	3	2	2.75	0

	A	B	C	D	E	F	G
1	Table 11		Aspiration	2.9			
2							
3	Rand. Str.	0.1	0.2	0.3	0.4	1	
4		C1	C2	C3	C4	EV	Excess
5	R1	1	3	4	1	2.3	0.6
6	R2	5	2	1	6	3.6	-0.7
7	R3	4	3	3	2	2.7	0.2

	A	B	C	D	E	F	G	H
1								
2								
3		0	0.5	0.25	0.25	1	rand. str.	excess
4	0.25	1	3	4	1		2.75	0
5	0.25	5	2	1	6		2.75	0
6	0.5	4	3	3	2		2.75	0
7	1							
8	random str	3.5	2.75	2.75	2.75		2.75	
9	excess	0.75	0	0	0			
10								

	A	B	C	D	E	F	G	H
1								
2								
3		0.1	0.3	0.4	0.2	1	rand. str.	excess
4	0.1	1	3	4	1		2.8	-0.3
5	0.5	5	2	1	6		2.7	-0.2
6	0.4	4	3	3	2		2.9	-0.4
7	1							
8	random str	4.2	2.5	2.1	3.9		2.5	
9	excess	1.7	0	-0.4	1.4			

	A	B	C	D	E	F	G	H
1								
2								
3		0	0.5	0.25	0.25	1	rand. str.	excess
4	0.25	1	3	4	1		2.75	0
5	0.25	5	2	1	6		2.75	0
6	0.5	4	3	3	2		2.75	0
7	1							
8	random str	3.5	2.75	2.75	2.75		2.75	
9	excess	0.75	0	0	0			