

OR/MS Today - October 2006
Computerized Third Party

WHEN THREE'S NOT A CROWD

John Kettelle



e-mail: kettelle@earthlink.net

From peace treaties to labor contracts: The case for employing a computerized third party for group decision-making and negotiations.

One of the most frustrating things about a negotiation is that each party has private information that is central to a sound agreement, but it can at the same time be folly for either party to share that information directly with the other party. A simple example is a sales transaction. The seller should not start by telling his prospect the lowest price he is prepared to accept, nor should the buyer start by telling the seller the highest price he is prepared to pay. Clearly, if there is, in fact, room for a deal, whoever first conveys this information is conceding the entire room to the other party.

In many negotiations, there will be an array of issues and corresponding sets of values placed by either side on the associated possible deals that might be struck. For either side to convey its version of these values directly to the other can be similarly damaging. Obviously, some kind of a third party can be helpful.

A traditional solution to this frustrating situation is for the two principals to find a human third party (HTP) with whom they can both confide. But HTPs have four well-known problems:

- Either side has to be concerned about the risk that the HTP may be a secret ally with the other side.
- Even if that risk can be accepted, there is a hazard that the HTP, who needs to be intimately familiar with the relevant secrets of both sides, may reveal clues about those secrets by accident.
- If the issue has any real complexity, the HTP will need a corresponding amount of prior expertise, not only on how to negotiate, but also on details of the issues peculiar to that particular problem.
- Real complexity can put a serious burden on the HTP, and any large staff that he may require multiplies all the above problems.

Before we attempt a general description of the CTP process, let us start with a rather simple example — what we call the pizza problem.

Suppose two nerds (Tom and Dick) enter a pizza parlor and thoughtlessly order a pizza, specifying that one side be mushrooms and the other side pepperoni. When they sit down to eat, they decide to make a real problem of just how to share it. They start by agreeing to require an "equi-max" solution, to wit:

1. Each man should get the same fraction of his own perceived value of the whole pizza, and
2. This common fraction should be maximized.

Now suppose their values about mushrooms and pepperoni are as follows:

Individual Values of Ingredients

	Mushrooms	Pepperoni
Tom	.6	.4
Dick	.1	.9

If Tom and Dick were normal, they would likely end up with Tom simply getting the mushroom side and Dick the pepperoni. But they are indeed nerds, and they never leave home without a computer that can satisfy the above equi-max sharing conditions. It gives them the following result:

Equi-Max Sharing Conditions

	Mushrooms	Pepperoni	Total Value
Tom	1.0	.23	.69
Dick	None	.77	.69
	Total	1.38	

Comments:

- It turns out that a computer can solve this problem, for any number of different flavors, and for any number of people who have a claim on the pizza, with a typical linear programming package. For the two-person problem, the computer can generate what is called a Pareto frontier — a plot of the most value Tom can get as a function of what Dick would get ([see box, page 23](#)).
- The "normal" resolution, Tom getting all the mushrooms and Dick all the pepperoni, gives a larger total (but obviously not one equally shared) value of 1.5. This suggests the issue of side-payments.
- If Dick really enters into the competitive spirit of this game, and furthermore knows Tom well enough to predict exactly what values Tom will enter, Dick can mimic Tom's preferences, except with an insignificantly larger value for the pepperoni (and a correspondingly smaller value for the mushrooms). The computer's solution will then be to give Dick all the pepperoni (which the computer thinks is worth only .6 to Tom, instead of the actual .9), plus one-sixth of the mushrooms, for a total treat worth over 90 percent to Dick, and only 33 percent to poor Tom.
- It is amusing what happens if *both* players use this gambit. The computer then dutifully gives each what the other really should have received, thereby providing what might be called the "pessimum" (worst possible) solution. Of course, even though they are both nerds, they would be likely to trade.

If one player happens to be a lot hungrier — or richer — than the other, there may be room for a side payment.

Operation of the CTP

First phase — the feasibility mode:

1. Each side submits a list of issues to be addressed.
2. The computer then reports the combined list to both sides. Because learning what one side deems relevant may be exploited by the other, each side will have been free to include some factors that were in fact unimportant to him.
3. Each side then submits to the computer his value function defined on all the factors in the combined list (the irrelevant ones being ignored, because *this* set of inputs will *not* be automatically revealed to the other side), and also enters the minimum acceptable value that it requires for a deal. Defining this value function is indeed a logical component of any negotiation.
4. The computer then derives the Pareto frontier — the set of deals in which any payoff for one side is accompanied by the largest possible payoff for the other.
5. It then reports to both sides whether there is "room for a deal" — whether there are *any* deals on the frontier that meet the minimum requirements of both sides.

No room for a deal. If the computer reports there *isn't* room for a deal, then the sides can either break off the negotiation altogether or cycle back to step 3, presumably with lower requirements. (A central point here is the possible inclusion of an additional party — one who has an exogenous stake in the success of an agreement. This issue can be dealt with either by including him in the negotiation from the beginning or by appealing to him *post hoc* — only after this current step. In the *post hoc* case, the two principals may form an alliance, and negotiate as a team with the newcomer.)

If we have departed the above steps via the "No" exit, this will indeed be bad news. (If it *weren't* news, there would have been no point in starting the negotiations.) The obvious question will be whether the parties could have learned this more easily, without all the detailed homework defined above. There are indeed buffered ways to do this, typically based on each side supplying *boundaries* of acceptability (below which the value of an agreement would be worse than no agreement at all), rather than the more complex value functions. This is normally done, at least implicitly, in "usual" situations. Unfortunately, this very "implicit" decision is often flawed, for the simple reason that rhetoric often misleads the other party about the real levels of acceptability. This is even further complicated by there being two kinds of "rhetoric" (in the misleading sense of the word): 1. rhetoric that is truly public, meant to gather public support; and 2. direct "private" communications with the other party, in which unfiltered truth-telling can still be counter-productive.

Room for a deal. But suppose the answer is "Yes." If we have indeed left via the "Yes" exit, the first thing is to recognize a positive change of mood for both parties — the news that there is at least one deal preferable to "no agreement." They are both left with the cheerful project of sharing the excesses over their minimum requirements. (In the negotiation community a popular acronym is BATNA — the "best alternative to no action." In the context of the CTP, BATNA gets parsed in two — the existence of at least one deal better than no action, and a logical cooperative search for the sides to choose one.

Clearly, among an array of alternatives both sides might prefer to "no action"; what is "best" for one is likely to be "worst" for the other, leaving BATNA poorly defined.)

Collaborative phase. The objective in this phase is for the two players to agree on how to share the excess beyond the minimally acceptable payoffs that have already been assured. On the way to this objective, either side may request information outputs (and who should get them) from the computer. The request itself (but not the requested output) will be automatically relayed to the other side, and will be granted only with his concurrence. A search for good strategies for the two sides promises to be fascinating, but here are a few thoughts:

- By now, the only relevant data may be the points on the feasible portion of the Pareto frontier.
- The two sides could agree to share that entire set of Pareto points (that is, the sets of conditions that generate the payoffs, not the payoffs themselves).
- Each could allow the other side to see only the array of its own payoffs, reducing what might some call the "jealousy" factor.
- The two sides could start by asking to see only one candidate point, for example the one somehow closest to the midpoint of the acceptable solutions, and by mutual agreement explore the frontier step by step.

The "digesting" phase. If things have gone well, this "mating dance" will by now have generated what might be called a "presumable" deal. Before this deal is finally ratified, you can make a case for what might be called a *digesting* period. Whether it deserves minutes, hours, days or weeks depends on the circumstances (minutes for the response to an accidental nuclear missile launch, months for a deal on non-proliferation); but in any case the deal's payoff, and its very structure, may have turned out to be somewhat of a surprise. This may warrant some additional thoughts that could lead to an even better deal (or avoidance of an unforeseen flaw in somebody's logic).

There may be a modification that would provide a very large gain for one side in exchange for a much smaller loss for the other. This might be addressed by some kind of a side payment beyond those already embedded in the frontier, or simply by giving the conceding party a chance to be generous.

Implementation and publication. If security permits, the negotiation can be completed remotely and secretly — perhaps to be executed only when something else happens. A secret alliance.

Once the basic agreement has been made, the parties can then use the CTP to work out second-order details of implementation — details not worth the effort before the basic agreement was assured.

Penalties for backing out. It is conceivable that one side could simply be using the buffer negotiation as a fishing expedition — simply to learn more about the other side's attitudes. An inherent protection against this will be the cost and time expended on the sham negotiation itself. If this is not enough deterrent, each side might be required to put in some kind of an initial deposit, which would be forfeited if one person arbitrarily rejected an agreement which had met all of the criteria that he himself had entered during the course of the process.

Applications

For us, the most exciting thing about the CTP concept has been the apparent payoffs of specific applications. Here is a sampling of governmental and commercial applications. Personal ones, like managing romantic relationships are left as exercises for the reader.

• **The Israeli-Palestinian problem**

This is a predicament replete with the kind of issues we have already mentioned. A partial list of issues includes:

- the boundary between Palestine and Israel;
- the fate of various settlements;
- sharing of water resources;
- convenient routes between Gaza and the West Bank, and (presuming that Israeli settlements in the West Bank survive), convenient connections between those settlements and Israel; and
- limits on the military capabilities of Palestine, and corresponding cooperation from Israel in case Palestine gets threatened by its other neighbors.

A frequent issue in a negotiation between two parties is the interest of another significant party or stakeholder. In this case the United States is clearly such a party. For that matter, the Middle East Arab community (somehow defined) is another. Theoretically these stakeholders can be, from the beginning, an inherent part of the computer-oriented negotiation. Or they can be added after the two principals have completed a preliminary round.

• **Military conflict terminations**

The termination of a war is typically a two-step process:

1. The cessation of fighting, typically called an armistice, which is to a first approximation often surrender by one side.
2. A follow-on detailed set of actions, a subset of which is sometimes called a peace treaty.

As the CTP process gets established it can help with each step, and provide a smoother connection between them. Examples from both World Wars suggest payoffs that are so dramatic that even a small chance of success can argue for a current "war avoidance" CTP system.

WWI — Only a few months after the U.S. intercession in April 1917, it should have been clear to both sides that an Allied victory was only a question of time. A short set of agreed conditions could very likely have led to an armistice a full year before November 1918. The nature of this armistice agreement might well have included:

- details about boundaries;

- first-order agreements about the nature of the surviving German government; and
- some preliminary provision for reparations (and post-war reconstruction).

In fact, a CTP might well have made U.S. intervention unnecessary, and even averted the war itself — a war that was a totally irrational response to the assassination of Archduke Ferdinand by a Bosnian Serb in 1914.

WWII — A RAND study ["Strategic Surrender," by Paul Kecskemeti, Stanford University Press, 1958] examined in detail the three successive WWII surrenders — of the Italians, the Germans and the Japanese. What follows draws heavily on that excellent book.

Italy: It may be an open question whether the armistice with Italy could have been with Mussolini or only with the putative post-Mussolini Italian government. In either case, the armistice needed to have been negotiated more quickly. It might well have provided for an effective "flip-flop" of the Italian military (including its battleship "Roma") to join the Allies in the continuing war against Germany.

Germany: With respect to the fall of Germany, there would have been a number of dramatic uses for the CTP:

- There just *might* have been an opportunity to make a deal with the general staff, and for example with the adherents of General Rommel, that could have obviated the Normandy invasion.
- Perhaps even more important, a viable CTP could have been used in the Yalta negotiations, for an agreement among the Western allies and the Soviet Union regarding the fate of the eastern European countries and of Germany itself. One "detail" of this could have been a better partition of Germany or some alternative to *any* partition.

Japan: Our own war began on Dec. 7, 1941, when the Japanese bombed Pearl Harbor — "A day that will live in infamy." As early as five months later (just after the Battle of Midway), Japan sent an envoy to Switzerland (a kind of no-man's-land), looking for early opportunities to end the war. In November 1943, Roosevelt, Churchill and Chiang Kai-shek published the Cairo declaration — that all Japan's recent conquests would have to be forfeited, but that its prior territory would survive — *better for Japan than the eventual result*.

As early as 18 months before Hiroshima, the Japanese military had recognized that Japan was going to lose the war. They approached Stalin to act somehow as a third party — a classic example of the hazards of using an HTP! Stalin just kept his mouth shut, never relaying Japan's interest in any kind of a surrender.

In the months before Hiroshima, the firebomb raids on Tokyo, coupled with our (submarine) naval blockade, were making the Japanese increasingly frustrated. In necessarily (lacking a CPT!) awkward ways, Japan conveyed directly to Washington its interest in some kind of a peace treaty, indicating that the survival of the emperor was a central issue.

Communications were so bad that Truman's reply, which tried to imply that this issue was indeed negotiable, was awkwardly delivered by an airdrop (over the royal palace) of a cluster of leaflets. The

Japanese reply (in spite of their desperate desire for some kind of a surrender) was misinterpreted as haughty defiance, prompting Truman to abandon the negotiation and immediately drop the bomb. The most elementary version of a CTP could well have avoided Hiroshima. (Some have made the case that the world has been better off to have been so dramatically confronted with the consequences of nuclear warfare, but it was not a demonstration that Truman would have made just for that reason.)

MAD (mutually assured destruction) — As one survivor of the Cold War, this concept is still a central component of our military posture. But it probably still has a significant chink in it — an accidentally launched single missile from one side might trigger the carefully planned MAD event by the other. The travel time of a ballistic missile between the United States and Russia is only about 70 minutes. If an accidental launch (from Russia, China or even the United States) is still an even remotely conceivable tragedy — and the possibility that a CTP just might avert it — could be well worth such an arrangement. The computer itself might be kept in a stationary orbit. Each country would maintain continuous contact with the CTP, and exercise frequent drills.

A well-thought out structure for making some deal will have been established within the CTP. An awkward (but far better to the ultimate "no deal"!) might be that the launching side will accept a single retaliatory missile from some list of possible targets of the other side. If a quick deal is made, wonderful. Otherwise the two sides could "destroy each other with a clear conscience."

Peace Treaties — The WWI Treaty of Versailles is normally considered to have been a primary condition that led to the rise of Hitler and the consequent initiation of WWII. The peace arrangements as WWII ended had differences and similarities. The differences included the cooperative restructuring of both Germany and Japan by the Western Allies and particularly by the United States. One can make a case that these successful arrangements (such as the Marshall Plan in Europe and benevolent administration of Japan by McArthur's temporary government) could well have been further enhanced by a CTP arrangement in which the resulting deals could have even more cooperatively fashioned.

A closely related issue, involving the United States and the Soviet Union right after Hiroshima, might have been some kind of a quick negotiation between them that could have led to a really great non-proliferation treaty, where the United States would have been left as the *only* country in the world with nuclear weapons.

• **Tactical Negotiations**

There is a tragic supply of battles that should never have been fought, or at least should have been terminated before their predictable conclusions. One set would include some of the WWII land battles in the Pacific, such as Iwo Jima and Okinawa. There is another set of examples whose resolution seems less feasible but is exceedingly frustrating. It is typified by the remarkably accurate loss predictions for deaths in both sides in the half dozen large battles of the Civil War, and also for much of the trench warfare in WWI. Instead of completing such battles whose outcomes are so predictable, the opposing forces might use the CTP to compute the expected deaths for each side (deaths typically very painful and agonizing) and then let each side randomly choose the same numbers of soldiers and render them more mercifully *hors de combat* (for example in some "joint" POW camp).

•

Non-Government

Labor-management contracts. These are gold mines for non-zero-sum elements. For each element of a deal, the values to the union typically have a very different structure from the costs to management. One central set of issues — salaries and hourly rates — can be an exception to this, but retirement benefits, vacations and holidays, sick leave, health benefits and overtime premiums typically have significant differences in costs to management versus values to labor. We have already noted that it can be counterproductive for one side to tell the other the values it assigns to various components of a deal, so here is indeed a good role for CTP.

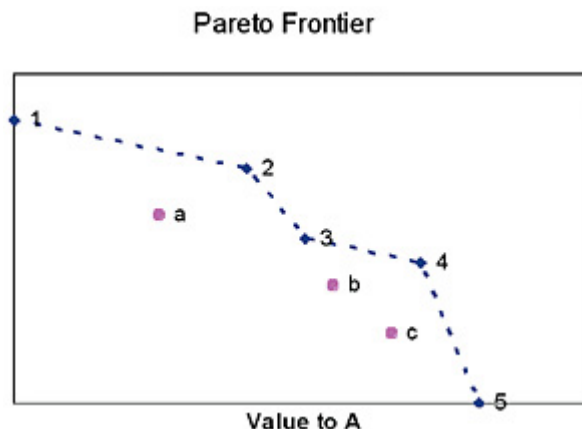
Explicitly facing this labor issue in the CTP context highlights another well-known aspect of labor-management contracts — mainly, different contracts for different categories of labor. For example, retirement benefits versus current income may have different relative values for airline flight personnel versus maintenance workers. This illustrates the need for a many-sided (rather than two-sided) negotiation structure.

Acquisitions and mergers. There are a many obvious components of a merger transaction:

- How much of the stock of the acquired company will be paid for in cash or in stock of the acquiring company?
- What happens to key figures of the acquired company? (For example, the very people negotiating the purchase.)
- A central element of a merger discussion is each side's estimate of the synergy of the merger — the extent to which the new entity will be worth more than the sum of the original parts. The CTP may provide an efficient way for the two companies to share enough information so that each can separately form its own estimate of the overall payoff of the merger, without directly giving each other information that can hurt them if the merger doesn't go through. The formation of coalitions among several companies (consistent with monopoly problems) would be a fascinating CTP application.

Pareto Frontier

Possible outcomes of a negotiation between parties A and B are plotted in the figure below.



For each lettered point, there is at least one numbered point that is better for both A and B. For example, 2-5 is better than "a," and 4 is better than both "b" and "c." The numbered points are called "undominated," and constitute the Pareto Frontier. To the extent assets can be continuously divided (such as in the Pizza Problem), the Pareto Frontier can be a continuous curve. When there are only two parties, the resolution obviously wants to be on the Pareto Frontier. When there are more than two, there can be all kinds of ways that a subset of the parties can cooperate to the detriment of the other parties.

John D. Kettelle (kettelle@comcast.net) has been an active member of the operations research/management science community since 1953. In addition to his work in negotiation, he has been examining ways to apply operations research to whole industries (rather than just individual companies). In his work on computerized third parties, the author is particularly indebted to professors Howard Raiffa, Steven Brams and Melvin Shakun.