

Designing Heuristics: Hybrid Computational Models for Teaching the Negotiation of Complex Contracts

*Gregory Todd Jones**

Editors' Note: Jones asks, can we learn something useful about negotiation by taking the people out of it? Paradoxical as this may sound, only inside a computer is it possible to run enough variations of some scenarios to generate believable outcome ranges. Surprisingly, Jones shows how easily real negotiators might find themselves grappling with one of these "complex" contracts.

Introduction

The goal for most of us who teach negotiation is to equip our students with skill sets that will make them more effective at reaching agreements that not only improve their own outcomes, but that ideally maximize social welfare. Tactics for doing so tend to be focused on the behaviors of individual negotiators, whether that is adherence to standards of rationality or, more often, strategies for accounting for departures from rationality – on both sides of the table. While teaching styles and substantive content may vary, the centrality of the actors makes it very rare that time is devoted to the consideration of structural mechanisms that may have every bit as much influence on the course of negotiations. The consequence of this gap is that negotiation students are not as equipped as they could be, and if there is an argument to be made for the benefits of integrative negotiation strategy, as most teachers of negotiation would concur, we are all worse off as a result.

* **Gregory Todd Jones** is director of research at the Interuniversity Consortium on Negotiation and Conflict Resolution at the Georgia State University College of Law in Atlanta, Georgia. His email address is gtjones@gsu.edu.

In what follows, I present results of simulated negotiations between computer agents which by definition lack the behavioral repertoire of human negotiators. I suggest that there is much to learn from the impact of structural constraints on the possible paths of these idealized negotiations, and that these lessons have the capacity to make our students better negotiators.

The negotiation of even the most straightforward real-world contracts tends to be quite complex. A contract with only 25 distinct issues with two alternatives each presents the parties with more than 33 million possible contracts,¹ far too many to be evaluated exhaustively within feasible time constraints. Further, while most work related to the analysis of negotiation assumes these issues to be independent and therefore the utility functions used in evaluating possible contracts to be linear (Ehtamo, Kettunen, and Hamalainen 2001; Faratin, Sierra, and Jennings 2002; Cheng, Chan, and Lin 2005), it is more typical that contract issues exhibit high levels of interdependence that result in nonlinear utility functions with the possibility of multiple local optima (Klein, Faratin, and Bar-Yam; Bar-Yam 1997; for a thorough treatment of multi-attribute utility, see generally Keeny and Raiffa 1993).

For example, imagine the negotiation with your partner about how the household discretionary income is to be spent. Each possible purchase represents a binary issue in the "contract space," with that issue set to 1 if the item is to be purchased and to 0 if not. On your agenda is upgrading your computer. Buying that new wide flat screen monitor you have been admiring would provide you with a certain amount of independent utility, but it would be made even better if you could also replace your CPU with a machine with ample memory and a high resolution video card to fully support the capabilities of the new monitor. Buying the new CPU by itself would offer some independent benefit. You can never have too much computer memory. But you would really get the most utility out of the new CPU if you could also buy the new monitor that would fully leverage the capabilities of the new CPU. Across all of the possibilities on the combined wish lists for you and your partner, there would likely be many such interdependencies, and that makes arriving at agreement about how to spend the money much more complex. Indeed, once a decision is reached, for example, to buy the new CPU, you may be trapped in the "local optima" of buying the new monitor. And, by doing so, you might fail to consider the possibility of forgoing both the CPU and the monitor in favor of another basket of purchases that might offer higher aggregate utility for you and your partner. As a consequence, even simple negotiations frequently result in sub-optimal, Pareto inferior agreements (Raiffa 2003).

At their essence, all but the most trivial negotiations introduce a search problem not unlike those faced by computer scientists when combinatorial explosion forces the consideration of a constrained set of possible solutions to complex problems.² This paper seeks to address the challenges of intractably large contract spaces and utility functions with multiple local optima by relying on two well known computational models for nonlinear optimization, simulated annealing and tabu search optimization. Annealing will suggest that negotiators should be “encouraged to consider contracts that may be inferior to previous proposals, at least early in the negotiation.” And tabu lists will counsel that “if a particular contract instantiation has been recently proposed, and agreement was not reached, consider other proposals for a while.” The results of such computational simulation can offer new advice for negotiators and negotiation teachers alike.

Several scholars, particularly Faratin and Klein (for a representative work among more than a dozen similar publications, see Klein and Faratin 2003) have applied simulated annealing, a probabilistic algorithm for locating approximations to global optima in large search spaces, to the mediated negotiation of complex contracts. The algorithm owes its name, as well as its substance, to the process of annealing in metallurgy, in which materials are heated and gradually cooled under controlled conditions in order to favorably impact such characteristics as strength and hardness (Kirkpatrick, Gelatt, and Vecchi 1983; Cerny 1985). However, with simulated annealing-based negotiation strategies alone, there remains the danger that small sections of the contract space can be cycled through, leaving parties trapped in the local optima that we seek to avoid. Your partner, who favors a new audio system, suggests that it only makes sense to upgrade the audio system if you are to buy new speakers and suggests holding off on the new computer monitor. You respond that you could use the current stereo equipment for a while longer. You modify the current proposal by suggesting going ahead with the purchase of the speakers – as well as the monitor. Your partner agrees to purchase the monitor, but proposes that you wait on the CPU and get the stereo instead. And so forth. While incremental agreements are being reached, the same territory will keep being rehashed.

In response to this difficulty, this paper proposes hybrid models that integrate tabu lists as well. Tabu lists are a simple form of tabu search optimization methodology which address the problems associated with local optima by employing memory structures to force exploration of regions of the search space that may otherwise go un-

explored (Glover and Laguna 1997). Alright, you point out, we have already considered the possibility of waiting on the monitor and buying both the stereo and the speakers, and failed to agree. Let's table that proposal for a while and consider other possibilities.

In addition to these novel models, my work can be distinguished from previous research by virtue of its motivation. While in most previous research, including that of Faratin and Klein, "bargaining automated agents are programmed with rules-of-thumb distilled from intuitions about good behavioral practice in human negotiations" (Binmore and Vulkan 1999: 3), with the objective of developing more effective autonomous agents, my objective is just the opposite. I am using computational simulations to aid in the design of social heuristics and institutional mechanisms that will make humans more effective in social exchange. My criteria include not only that the generated negotiation mechanisms produce more socially optimal negotiation outcomes, but also that they can be relatively easily translated from the abstract digital world of autonomous agents into the negotiation classroom.

Defining the Problem of Contract Complexity

The exponential complexity of multi-issue contracts can defy a negotiated agreement in at least two ways. First, when pre-existing circumstances or preliminary negotiations (Contract 1 in Figure One)

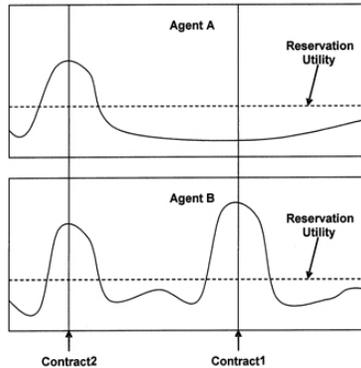


Figure One – Failure to Reach Agreement

place one party (Agent B) at an optimum that is above her reservation utility while the other party (Agent A) remains below his reservation utility, and where the first party (Agent B) only considers alternative proposals that are strictly better than the current proposal (a protocol known as "hill climbing"), there is no opportunity to reach alternative contracts (Contract 2) that would be acceptable to both parties. Second, where both parties find themselves at local

optima and both are above their respective reservation utilities (Contract 1 in Figure Two), if even one of the parties engages in hill climbing, Pareto superior agreements (Contract 2) can never be reached.³

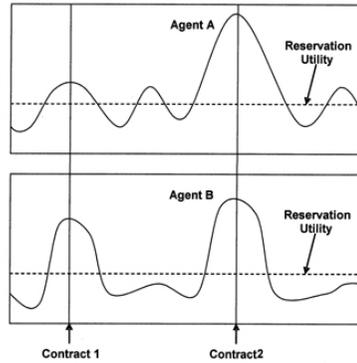


Figure Two – Pareto Inferior Agreement

Simulated Annealing

Simulated annealing deals with this complexity and the prospect of local optima in the following way. Using the context of negotiation to illuminate, you start with an arbitrarily high “temperature” (in the sense related to annealing and not to be confused with the level of emotions in negotiation) and evaluate the other side's proposed alternative contracts by comparing the utility of the alternative with the current contract “on the table.” Where the utility of the alternative is greater from your perspective, you agree that this contract should become the new working solution. Where the alternative represents a reduction in utility from your perspective, you are more likely to accept the contract if the difference in utility is small and the temperature is high. The “temperature” of the negotiation declines steadily according to some prescribed cooling schedule, which in my simulations is linear. As a result, you are more likely to accept large utility reductions (at least for the sake of discussion) early in the negotiations and less and less likely to accept even small reductions as the negotiations wear on.⁴ At the limit, annealers become hill climbers, accepting alternative proposals only if they are strictly superior in utility from the individual's perspective. Where $P(\delta E, T)$ is the probability of accepting an alternative contract given a difference in utility (δE) and a temperature (T),

$$P(\delta E, T) = 1$$

where the alternative contract is better than the contract "on the table."

$$= e^{-\left[\frac{\delta E}{T}\right]}$$

otherwise

Tabu Lists

While simulated annealing can make it more likely for a party to accept an inferior contract as a working solution, particularly early in the negotiation, there remains an additional danger: parties may trade acceptances of proposals that mirror the same issues, in cycles that continue until the temperature cools, to the point that even these proposals are rejected and the parties are left with no agreement, or at best, a substantially Pareto inferior agreement. To address this possibility, my simulations implement hybrid protocols that integrate the very simple idea of tabu lists, a sort of memory for contract proposals that have not produced agreement. The search procedure moves from solution x to solution x' , in the neighborhood of x , $N^*(x)$. The tabu list contains solutions recently attempted (less than n proposals ago, where n is known as the tabu tenure), and proposals on the tabu list are excluded from $N^*(x)$. Put more simply, once we have considered a particular contract and put it aside, we will not return to the consideration of that contract for a specified number of proposals. This prevents the possibility that parties become trapped in "back-and-forth" proposals that limit consideration to a small area of the contract space. Or as a mediator might say "we are beating a dead horse here...let's talk about something else for a while."

The Simulations

To experiment with the usefulness of these computational models in simulated negotiations, I developed an "in silico" negotiation laboratory (see Figure Three) for a two-party contract negotiation with 25 interdependent issues, presenting the possibility of a highly complex, non-linear contract space. The software makes it quite easy to iterate simulations, varying parameters for initial temperature (the higher the initial temperature, the more open negotiators are to alternative proposals at the outset), cooling schedule (the more rapidly the temperature cools, the more rapidly negotiators become less flexible, ultimately being willing to consider only those alternative contracts that are strictly improvements from their point of view), tabu list length (the number of contract proposals that are tabled

once they are considered and agreement is not reached), and the extent to which parties cognitively “bundle” issues together during the course of the negotiation (bounded rationality suggests that a human negotiator can only consider a relatively small number of issues at any one time).

Experiments proceeded in three stages. First, utility schedules for the 25 issues were independently randomized for each party, with each utility derived from a uniform distribution between 1 and -1 (utilities can be negative as well as positive). Note that the diagonal of these matrices represent the independent utility of each contract issue, while the symmetrical, off-diagonal values represent the interdependent utility of each issue pair. Using an annealing algorithm, the optimal contract for each agent was determined for purposes of comparison; based on modifiable assumptions of salience and cognitive bundling, an initial proposal was generated from Agent A. Contracts were made up of 25 boolean integers, where 1 indicated the inclusion of each particular issue and 0 indicated its absence.

Finally, the two agents, each with modifiable negotiation strategies (hill climbing or hybrid annealing), engaged in a negotiation consisting of a pre-determined number of alternating proposals. Figure Four reports mean joint utility and mean agent utilities after 30 negotiations in each of four conditions: 1) A hill climbs and B hill climbs; 2) A employs a hybrid annealing strategy and B hill climbs; 3) A hill climbs and B employs a hybrid annealing strategy; and 4) both A and B employ a hybrid annealing strategy. Rearranging these results (see Figure Five) makes it clear that, like Klein, Faratin, and Bar-Yam, my simulations produce a prisoner’s dilemma where: 1) hill climbing is dominant; that is, each agent is privately better off hill climbing regardless of what the other agent does; and 2) there is a deficient equilibrium at mutual hill climbing; that is, aggregate social welfare is maximized with mutual annealing and is at its lowest with mutual hill climbing. However, there is no individual incentive for annealing.

A	B	% of Optimal Utility		
		Joint	A	B
Hill	Hill	0.435	0.386	0.251
Anneal	Hill	0.556	0.051	0.777
Hill	Anneal	0.694	0.895	0.229
Anneal	Anneal	0.947	0.570	0.704

Figure Four – Hill Climbing vs. Hybrid Annealing

		B	
		Hill	Anneal
A	Hill	0.386	0.895
	Anneal	0.777	0.704
		0.251	0.229
		0.435	0.694
		0.051	0.570
		0.556	0.947

% Optimal Utility:

A
B
Joint

Figure Five – A Prisoner's Dilemma

Discussion & Conclusions

Klein, Faratin, and Bar-Yam go on to try and develop institutional mechanisms that will avoid the deficient equilibrium of the prisoner's dilemma with particular application to autonomous agents. However, this is not my concern. Preliminary results of my research have demonstrated that the hybrid computational models considered are successful in improving social welfare and this prisoner's dilemma, like others, can be overcome in repeated negotiations with some level of "correlated association" (Skyrms, 1996) wherein agents are at least slightly more likely to interact with other agents of their own kind, by means of signaling, reputation mechanisms, or proximity in space.

Further, it is no accident that both simulated annealing and tabu lists can be translated into very straightforward suggestions for how to conduct human negotiations. Annealing suggests that negotiators should be "encouraged to consider contracts that may be inferior to previous proposals, at least early in the negotiation." And tabu lists counsel that "if a particular contract instantiation has been recently proposed, and agreement was not reached, consider other proposals for a while." My intent was to use the computational simulations to aid in the design of social heuristics that may improve the outcomes of human negotiations.

Future research will examine more closely the role of cooling schedules and the length of tabu lists. I also plan to add the possibility of both qualitative and quantitative issues to the boolean issues that are now considered. Higher order issue interdependence and multi-party ($n > 2$) negotiations that give rise to coalitions would also both be of interest. Most importantly, however, the next phase of my research will be conducted in human subjects laboratories to

determine if the insight provided by the computational simulations is translatable to the context of actual human exchange, and therefore to the teaching of negotiation.

Notes

Much of this work was completed while a Visiting Researcher at the Max Planck Institute for Research on Collective Goods, Bonn, Germany. Their kind hospitality is gratefully acknowledged. I would like to thank Doug Yarn, Reidar Hagtvedt, Christoph Engel, Arndt Bröder, Martin Beckenkamp, Stephan Tontrup and the members of the Max Planck Institute for Research on Collective Goods Rationality Working Group for stimulating discussions and many helpful comments. I am also thankful to Christopher Honeyman, James Coben, and other unknown reviewers for many excellent suggestions.

¹ The actual number is $2^{25} = 33,554,432$.

² Indeed, the negotiation problem would be made much worse if there were more than two parties such that the formation of coalitions was possible (Krippendorff 1986). The role of coalitions is not formally dealt with here.

³ The pragmatic consequences of these circumstances are well known in the mediation literature. For example, Christopher Honeyman notes

There is almost always more than one set of terms that can make for a mutually acceptable settlement; "give" on one item may compensate for "take" on an apparently unrelated one. However, many mediators are under some degree of time pressure, from the press of other work, the need to show progress to the mediator's appointing agency or peers, and for many other reasons. This pressure encourages a mediator to search for the first mutually agreeable settlement package rather than for some conception of the best agreement. Often the parties are aware of this. Despite the theoretical existence of an agreement, it is not unusual for negotiators to meet afterwards, with or without the mediator, and modify the original set of terms to their mutual benefit. It is as if these parties needed the mediator's assistance to get within hailing distance of each other, but once that is accomplished, the parties' superior knowledge of their own needs makes direct negotiation fruitful again (Honeyman 2006: 588).

⁴ It is important to point out that this metaphor runs counter to the way "heat" is generally thought of in the context of negotiation. As expressed in the phrase "more heat than light," we would expect less agreement when heat, meant to represent ire, is high. The notion that more concession should be the goal when the heat is high, is a function only of the sequence of annealing in metallurgy, and should not be a source of confusion.

References

- Bar-Yam, Y. 1997. *Dynamics of complex systems*. Reading, MA: Addison-Wesley.
- Binmore, K. and N. Vulkan. 1999. Applying game theory to automated negotiation. *Netomics* 1(1): 1-9.
- Cerny, V. 1985. Thermodynamical approach to the traveling salesman problem: An efficient simulation algorithm. *Journal of Optimization Theory and Applications* 45(1): 41-51.
- Cheng, C., C. H. Chan, and K. Lin. 2005. Intelligent agents for e-marketplace: Negotiation with issue trade-offs by fuzzy inference systems. *Decision Support Systems* 42(2): 626-638.
- Ehtamo, H., E. Kettunen, and R. P. Hamalainen. 2001. Searching for joint gains in multi-party negotiations. *European Journal of Operational Research* 130(1): 54-69.
- Faratin, P., C. Sierra, and N. R. Jennings. 2002. Using similarity criteria to make issue trade-offs in automated negotiations. *Artificial Intelligence* 142(2): 205-237.
- Glover, F. and M. Laguna. 1997. *Tabu search*. Norwell, MA: Kluwer.
- Honeyman, C. 2006. Understanding mediators. In *The negotiator's fieldbook: The desk reference for the experienced negotiator*, edited by A. K. Schneider and C. Honeyman. Washington, DC: American Bar Association.
- Keeney, R. L. and H. Raiffa. 1993. *Decisions with multiple objectives: Preferences and value trade-offs*. Cambridge: Cambridge University Press.
- Kirkpatrick, S., C. D. Gelatt, and M. P. Vecchi. 1983. Optimization by simulated annealing. *Science* 220: 671-678.
- Klein, M. and P. Faratin. 2003. Negotiating complex contracts. *Group Decision and Negotiation* 12(1): 111-125.
- Klein, M., P. Faratin, and Y. Bar-Yam. Using an annealing mediator to solve the prisoners' dilemma in the negotiation of complex contracts. (Unpublished working paper).
- Krippendorff, K. 1986. A dictionary of cybernetics. (Unpublished report, The Annenberg School of Communications, University of Pennsylvania).
- Raiffa, H. 2003. *Negotiation analysis: The science and art of collaborative decision making*. Cambridge, MA: Belknap Press.
- Skyrms, B. 1996. *Evolution of the social contract*. Cambridge: Cambridge University Press.

