Volume 1, No. 6 | November–December 2018

The Journal of Robotics, Artificial Intelligence & Law

fastcase FULL COURT PRESS

Enhancing Contract Playbooks with Interactive Intelligence—Part II

Marc Lauritsen*

Contract playbooks are rich sources of knowledge and guidance. Yet, even when digitally delivered, in few organizations are playbooks interactive or "smart." This article explores two technological fronts upon which to change that situation: one (document automation) that is mostly well established and another (preference management) that is relatively novel. The first part of this two-part article, which appeared in the September-October 2018 issue of The Journal of Robotics, Artificial Intelligence & Law, *summarized current ways in which technology is being applied to contracting processes, described playbooks and their limitations, reviewed how well-established document automation techniques can be applied to address some of those limitations and took up ways in which contracting is a process of choice management. This second part of the article describes a system and method for managing choices, applies that method to contracting, and contains concluding thoughts.*

This second part of a two-part article describes a system and method for managing choices, applies that method to contracting, and contains concluding thoughts.

A Tool of Choice

A method called "choiceboxing" is designed for decision support via online environments.[34] It enables collaborative deliberation and relies on interactive visualization and social production techniques. In a nutshell, choiceboxing involves a systematic process of:

- Identifying aspects of candidate solutions that you *might* care about, with special attention to those aspects that involve true "make or break" requirements;
- Casting a wide net for candidates, and culling out those that satisfy essential requirements;

Robotics, Artificial Intelligence & Law / November–December 2018, Vol. 1, No. 6, pp. 385–401.

ISSN 2575-5633 (print) /ISSN 2575-5617 (online)

- Specifying how much you care about candidate features (organizing them in at least relative importance, so you can deal meaningfully with the inevitable trade-offs);
- Assessing candidates on the important features, in a way that permits each to be given a roughly quantitative score on its relative goodness for each feature;
- Using a weighted factors approach to judge which candidate does best, all things considered;
- Iterating this process through group discussion and continuing deliberation.

Choiceboxing uses an interactive three-dimensional representation of the decision in process, in which all options, factors, and evaluative perspectives can be accommodated. It can be used within a single organization for individual decisions, or across organizations as a kind of Wikipedia of Choice whereby collective wisdom can be accumulated and shared.

By convention, options are positioned left to right, factors top to bottom, and perspectives front to back (see Figure 1). There is a column for each option, a row for each factor, and a layer for each perspective. Each cell at the intersection of such a column, row, and layer represents the characterization of some option in terms of some factor according to some perspective.

An example should help make this clearer.

Figure 1. A choicebox

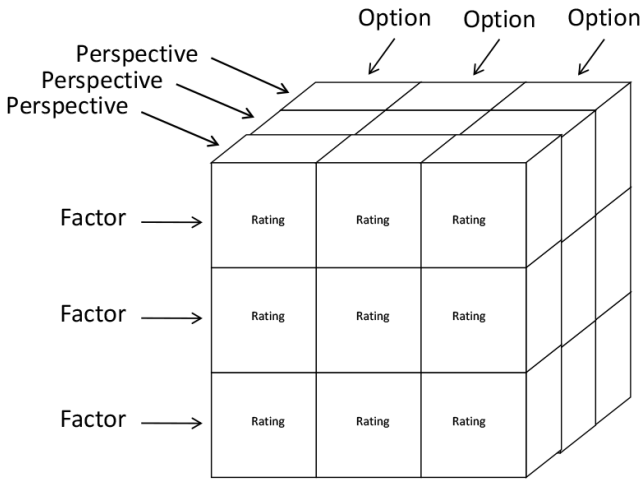

Figure 2. A choice of contract management software

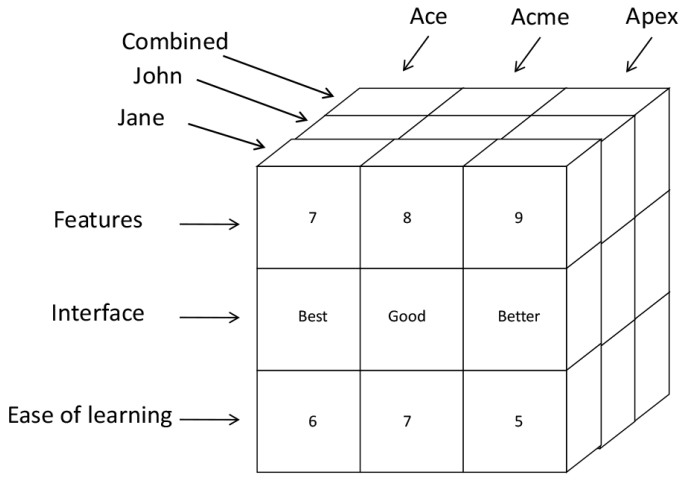

Example: Choosing a Technology

Choiceboxing is most naturally applied to decisions that involve selections from discrete sets of options, like the choice of a particular technology.

Imagine that Jane and John work in an organization that is deciding which contract management system to buy. They have narrowed it down to three products: Ace, Acme, and Apex. After much discussion, the choice seems to hinge on three factors: completeness of features, quality of interface, and ease of learning. Figure 2 depicts how this matrix of options, factors, and perspectives might be represented in a choicebox. We are seeing Jane's perspective up front. The factors are matters of opinion, so her ratings and those of John may well differ.

Figure 3 makes the separate perspective layers clearer. Now we can see some of John's different ratings, as well as average ratings on the combined layer. (The box can also usefully be sliced along the other axes, showing one option or factor at a time.)

Taking this a couple of steps further, one can express each assessment of each option from each perspective in a separate block of goodness like that shown in Figure 4.

One can then position such blocks within the overall framework of a choicebox, with associated totals as shown in Figure 5. Note that the each person can set different relative importance heights for the rows.

Figure 3. Perspective slices

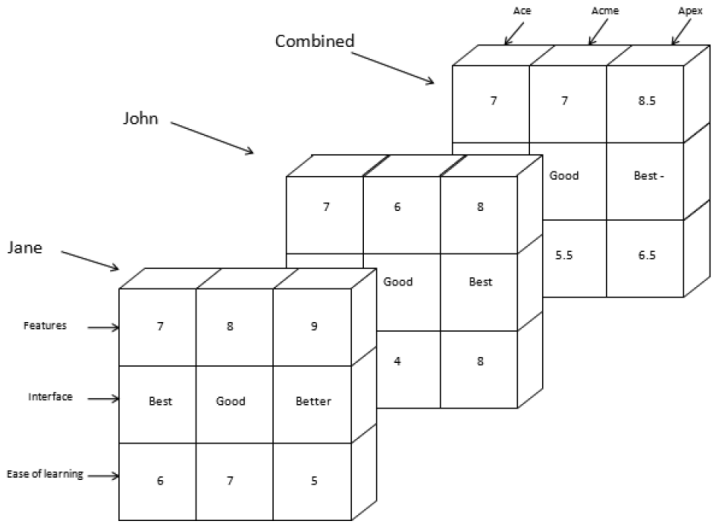

Figure 4. Value boxes

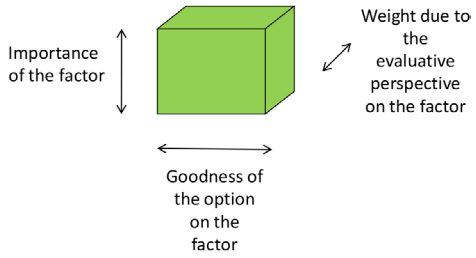

You can imagine the total boxes at top as having been formed by melting down, combining, and reshaping the ingots of goodness in the columns beneath them.

Many real-world choices, of course, involve more options and considerations. Usually, though, the ultimate choice boils down to a small subset of each.

Preferences, Circumstances, and Choice Spaces

An important aspect of choiceboxing not shown above is the use of *preference functions* to map between possible features within given factors and their relative acceptability or goodness. Often

Figure 5. Value boxes with totals

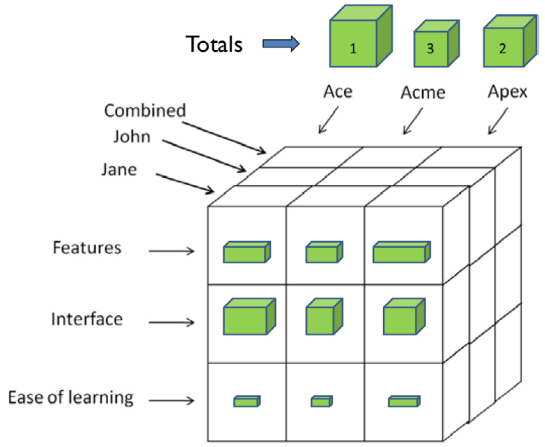

Figure 6. Dinner companion preferences

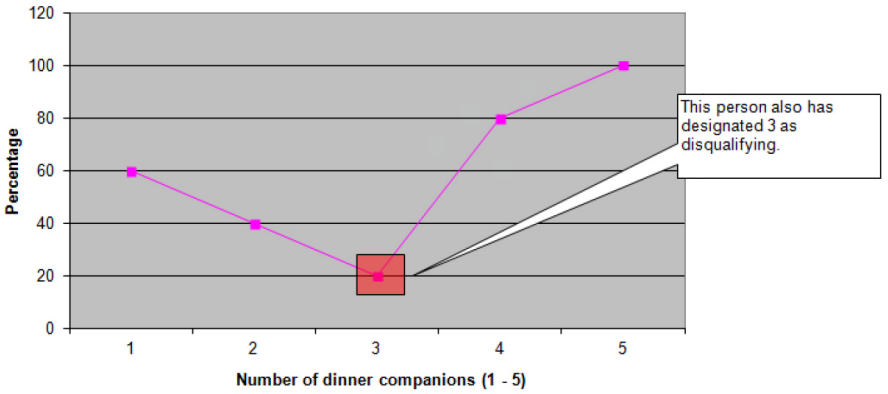

option features are non-numeric (like the hand feel of a camera or the color of a car), and need to be quantified in order to be combined with other features. Even numeric features (like pixels or miles per gallon) may have non-linear implications in terms of preference. (Twice as many pixels may not be considered twice as desirable.) Moreover, the ranges of possible preferences for all factors need to be normalized to a common scale so as to be fairly commensurable in proportion to the explicit weights given by a decision maker.

Imagine someone who likes having dinner with one person better than with two, but likes four or five even better. And refuses to eat with three. A possible preference function, shown in Figure 6,

is expressed in terms of percentages of optimality and shown graphically.

The choicebox model goes on to take into account the *circumstances* of the parties and other facts that have a bearing on the relative goodness of options on different factors from the various perspectives engaged in a decision.

Another important concept is that of a *choice space*. That is the virtual space in which a number of decisions and decision makers—along with suppliers, helpers, and other would-be influencers—are in play. Such environments can yield powerful network effects.

Wider Applications

In addition to its use in product selection processes as just outlined, the machinery of choiceboxing can help in a wide variety of decision processes. It can be used to compare candidates for a job, arguments that should be emphasized in an appellate argument, treatment options for a disease, or hypotheses about the cause of a health issue or machine malfunction.

In business, choiceboxing can help organizations choose trading partners, negotiate agreements, and resolve disputes.[35] Each of these contexts involves alternatives as to which there are competing trade-offs and multiple evaluative perspectives. Tools that leverage interactive visualization and social production to enable collaborative deliberation may be conspicuous in next generation business ecosystems.

In law, choiceboxing has promising applications for client counseling (e.g., deciding whether to treat a staff member as an employee or contractor),[36] legal services for the poor,[37] and judicial decision making.[38] This author has taught several law school courses in which students build such applications.[39] It was used by a group of lawyers, judges, professors, and technologists to prioritize potential legal technology investments by the federal Legal Services Corporation.[40] The views of 32 participants on the relative efficacy of 10 possible activities in achieving 11 desired objectives were elicited and consolidated in a structure that looks like that shown in Figure 7.

So what about applying choiceboxing to the drafting and negotiation of contracts?

Figure 7. A stack of evaluative perspectives

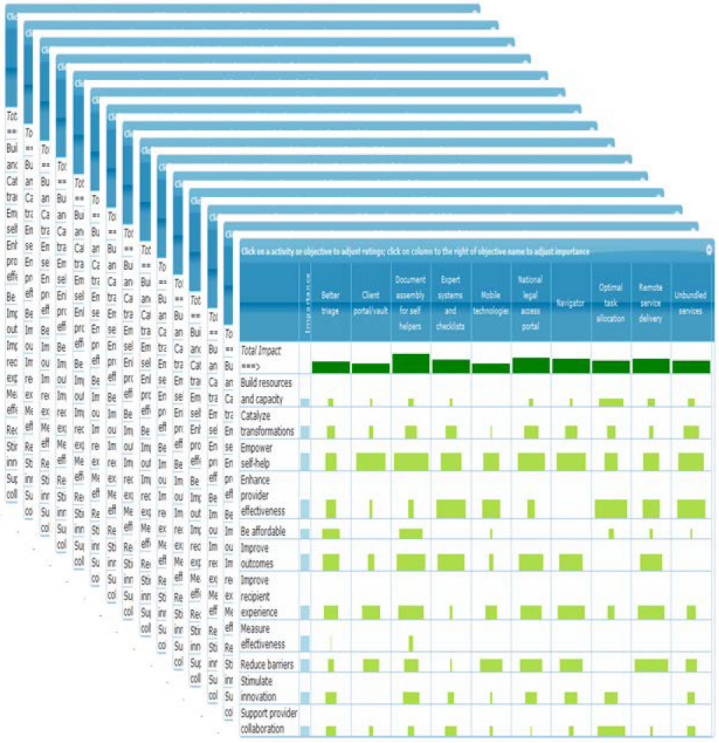

Choiceboxing a Contract

The moving parts of a contract drafting or negotiation exercise map quite readily into the choicebox framework:

- The *options* (columns) are possible or proposed versions of the contract.
- The *factors* (rows) are the terms, provisions, and other aspects that might vary from one version to another. These aspects of the options are a subset of the variables at play in a deal, including not only clauses that deal with different issues, but also discrete terms like price and duration that might be internal to a given clause. In addition to specific terms and provisions, versions can differ in global aspects such as plainness of language and consistency of voice and terminology.

- The *perspectives* are those of the parties or particular team members within a party (and perhaps other entities and/or reference standards).

Each participant can thus express (to itself or each other) views about the relative acceptability of each aspect in each contract version under consideration, *and* about the relative importance of different aspects. Specific aspects of particular versions may be judged not only in terms of their substantive acceptability, but also in terms of attributes such as industry standardness, clarity, brevity, and novelty.

The options are the contending versions or proposed contracts, including the one ultimately agreed to, each being a different package of aspect handlings. Unlike choices of the kind described earlier, where the options have relatively fixed sets of features, in contracting new combinations are often being introduced and assessed. The very attributes of options are subject to case-specific crafting and negotiation. Contract negotiation thus is a kind of configuration exercise.

Suppose for instance that a negotiation has come down to differences over a material adverse change clause and a closing condition in a preferred stock agreement between Newco and Ventura. Members of the negotiating teams have different perspectives because they assign different weights to the goals of, for example, (a) preserving good will and the prospect of future dealings with the counterparty, (b) avoiding risk, and (c) minimizing the prospect of protracted litigation.[41]

Figure 8 shows what Newco's options might look in a choice-box with four contract versions—"Newco ideal," "Ventura ideal," "Draft 13," and "Signed contract." Preferences among the various provisions are expressed by the weights at left. (Weights are expressed on a scale of zero to 10.) The numbers in the cells represent the percentage of goodness a party gets under the scenario. So, for example, for the "Newco ideal" version, Newco gets its ideal outcome (represented as 100 percent) on every issue, and obviously would find that most attractive. Likewise, it might find the "Ventura ideal" *least* attractive (represented as zero percent). Compromise versions like the penultimate Draft (13) and the final contract give each party some of what it most wants.

Figure 9 shows what the corresponding sheet for Ventura might look like. Notice that it has different factor importance weights;

Figure 8. Newco's ratings and weights

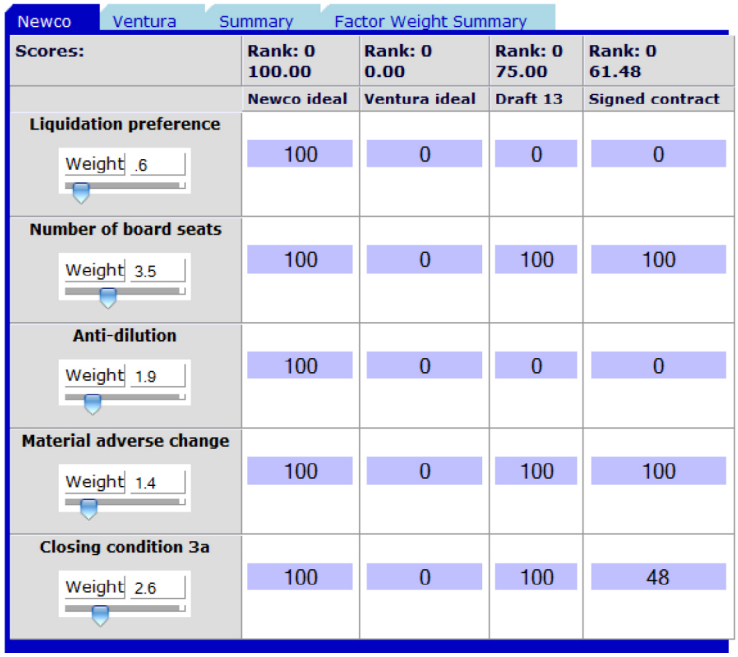

Newco | Ventura | Summary | Factor Weight Summary

Scores:	Rank: 0 100.00	Rank: 0 0.00	Rank: 0 75.00	Rank: 0 61.48
	Newco ideal	Ventura ideal	Draft 13	Signed contract
Liquidation preference Weight .6	100	0	0	0
Number of board seats Weight 3.5	100	0	100	100
Anti-dilution Weight 1.9	100	0	0	0
Material adverse change Weight 1.4	100	0	100	100
Closing condition 3a Weight 2.6	100	0	100	48

Figure 9. Ventura's ratings and weights

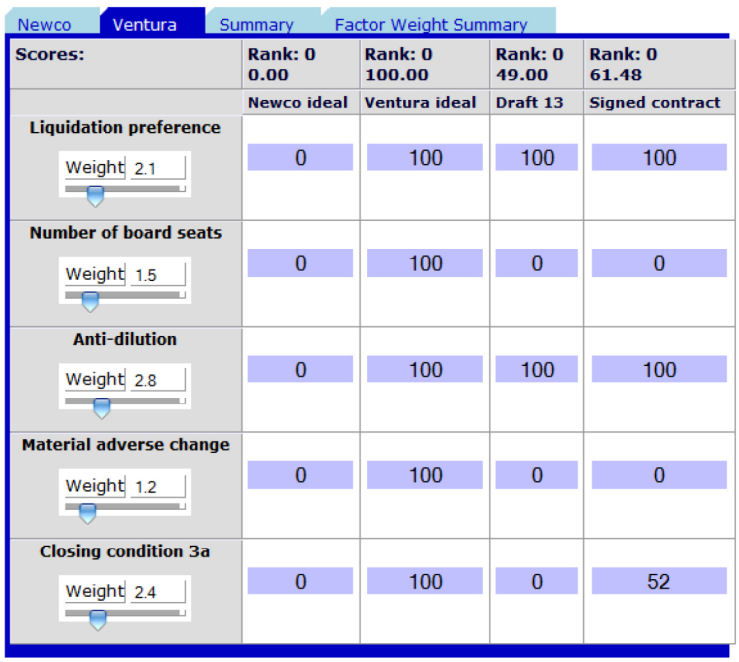

Newco | Ventura | Summary | Factor Weight Summary

Scores:	Rank: 0 0.00	Rank: 0 100.00	Rank: 0 49.00	Rank: 0 61.48
	Newco ideal	Ventura ideal	Draft 13	Signed contract
Liquidation preference Weight 2.1	0	100	100	100
Number of board seats Weight 1.5	0	100	0	0
Anti-dilution Weight 2.8	0	100	100	100
Material adverse change Weight 1.2	0	100	0	0
Closing condition 3a Weight 2.4	0	100	0	52

Figure 10. A three-dimensional rendering

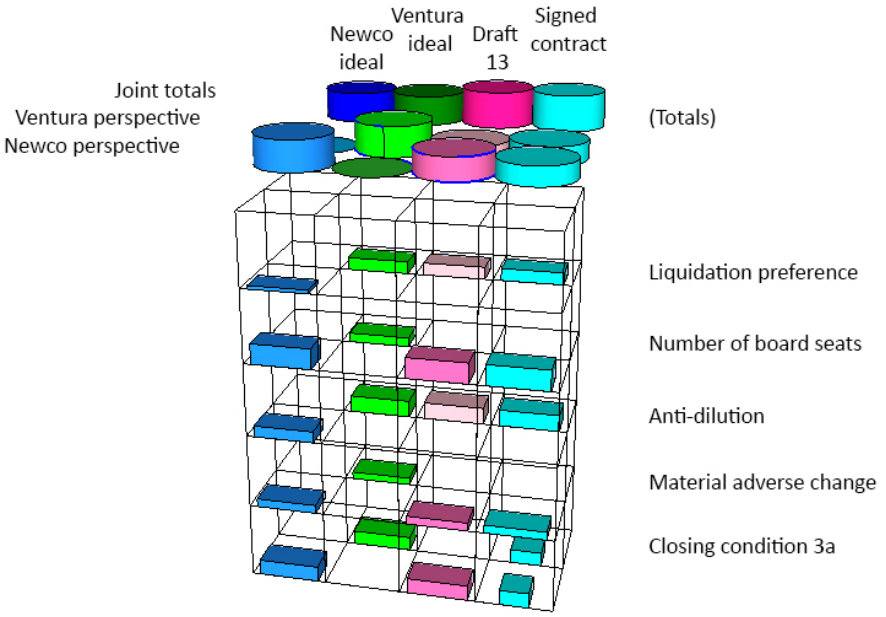

for instance, caring much more about liquidation preferences than Newco.

In this imaginary case, each party gets either all (100) or none (0) of what it wants on most of the issues in controversy. Most negotiations of course involve sub-"ideal" resolutions of multiple issues for all parties, but also many that even the "losing" party finds more than minimally preferable. The 48 and 52 in the bottom right cells of the parties' perspectives above represent how they respectively judge the desirability of the compromise resolution reached over a tricky closing condition issue in the final contract.

When rendered three-dimensionally, this configuration of weights and preferences might look like that shown in Figure 10.

Here the rectangular blocks represent the value of deal terms as allocated to the two parties under the various scenarios, and the cylindrical shapes at top the total goodness in each vertical column. Thus, Newco and Ventura get everything respectively in the first two scenarios, and the total utility for these two scenarios (back row of cylinders) is the same. Draft 13 produces a higher joint utility, but it is unevenly distributed between them. The executed contract—reflecting the closing condition compromise—delivers just slightly less joint utility, but it is evenly divided.

Uses

Rolling out the full machinery of a choicebox will make most sense when the stakes in a contract are high enough to engender prolonged deliberation or negotiation, and when neither party is in position to dictate terms yet both have strong and conflicting preferences for particular terms and provisions. Reaching a mutually acceptable outcome in such a situation can present a bewildering array of entangled choices. Current technologies provide little support for managing the needed judgments and trade-offs. Choiceboxing offers a structured way to capture the relative acceptability and preferability of contending versions, both granularly and in the aggregate. It enables participants not just to see how the alternatives differ, but to express how much they care about those differences.

Preference functions can be configured to reflect not only zones of absolute and relative acceptability, but aspect handlings that may require particular levels of executive approval. The acceptability of particular handlings will often depend on circumstances such as type of contract, the ways in which *other* aspects are handled in a given version, and the particular transaction or relationship being contemplated. They may also depend on the current business climate in general. They can encompass as wide spectra of variations as an enterprise may encounter.

An enterprise of any size should have a comprehensive set of "care curves" (codified preferences among potential term variations). Guidance and other resources can be made available right in the context of a specific issue within a particular kind of contract. Such collections can serve to codify and communicate standards and practices.

Since any responsible contract manager will give at least glancing attention and rough valuation to all aspects of every contract seriously under consideration, that evaluative information is available to be preserved and harvested, which can help with both the construction and maintenance of such preference function collections. When a contract has been proposed by another party, analytical tools of the sort described in Part I of this article can be used to do a first characterization of its terms in relation to a reference standard by finding the nearest matches to formulations already present in that standard's preference functions, subject to human confirmation.

Benefits

As a tool for contract drafting and negotiation, a choicebox-enhanced playbook would:

- Enable people to more reliably judge the relative preferability of proposed drafts and access a tidy overview of issues, settled and outstanding;
- Help elevate substance over syntax, by enabling participants to decide *what* they want to agree to, independent of the sequence and structure of words;
- Promote clarity about the difference between how much better a draft is on a particular point versus how much one cares about that kind of betterness;
- Bring focus to the things people care enough about to argue over, or fight for;
- Supply a framework within which the known or inferred preferences of counterparties can be tracked;
- Provide a record from which participants can readily retrieve *reasons* for having taken or accepted certain positions;
- Offer fuller expressiveness of company standards and policies because it not only captures preferred language but also the relative preferability of particular variants and the relative importance of types of terms and conditions;
- Help parties maximize their joint utility in a deal by suggesting trades by which each would gain net value;
- In lawyer/client relationships, be a medium through which the advisability of various terms and provisions can be communicated and stored.

A choicebox-like visualization, once understood, provides a convenient way to grasp the dynamics of the contracting process, especially for those more visually than textually or numerically inclined. Parties can directly interact with such models to express preferences and explore solutions, perhaps in mutually invisible ways that a neutral (human or machine) accesses to suggest collectively optimizing moves. (An adequately trustworthy intermediary that knew both parties' preference profiles could compute and propose Pareto optimizing packages of terms.)

Visual depiction can help, for instance, in anticipating ways in which one party might try to "game" the other, such as by exaggerating the degree to which it cares about an issue, so as to wring an unjustified concession in exchange for forgoing its ostensibly preferred way of resolving it.

Visual depiction can also remind people that value considerations are almost always in play other than those directly associated with the terms about which one may be bargaining. For example, in a merger there are costs of negotiation or litigation that might be incurred or avoided under various scenarios, and benefits in terms of public relations and industrial peace that may accrue. Also, even if the participants' decisional frameworks are necessarily entangled, there are often considerations that are peculiar to one side or the other, providing asymmetries that can be exploited for mutually optimizing results.

While this article has focused on contexts in which people frequently engage in high-stakes contract drafting, many of the above benefits could be made available to consumers, small businesses, and other parties that only contract occasionally, via commercial, non-profit, or governmental implementations of playbooks that are not designed to just serve the purposes of a single organization. Such resources could efficiently communicate what is best (or at least realistic) contract-wise for a given party in given circumstances.

Challenges

Research by Ap Dijksterhuis and his colleagues in Holland (such as their paper on the rational subconscious[42]) casts doubt on whether "weighted adding" yields better results than intuitive, gist-based approaches to complex decisions. Our unaided minds are just not very good at it. Cognitive prostheses like choiceboxes will help counter that deficit. More importantly, they enable more open and productive deliberations *across* minds by externalizing our thinking.

Still, aspect judgments and factor weightings can be thorny, especially when cast in misleadingly precise number formats. The choiceboxing interface allows ratings and weightings via graphical sliders, which are more inviting for approximations. That approach

will likely prove essential and readily accepted once understood. It seems clear that *some* quasi-quantitative delineation of the relative goodness of options and the relative importance of factors at play is inescapably part of responsible decision making.[43]

Doing justice to the underlying complexity of effective playbooks requires going beyond two-dimensional grids. And three-dimensional matrices on the rows of each layer of which are hung preference functions (containing not just the relative acceptability of alternative provisions from a given perspective, but review thresholds and escalation paths) will not initially be intuitive to most lawyers. But with a good interface and decent orientation playbook creators and managers should find such devices quite powerful and satisfying. Having the knowledge and policy they encode then be operationally available to drafters will be a major payoff.

In actual review and negotiation of contracts, parties will often encounter versions of provisions that don't exactly match those in their playbooks. As the new forms are reviewed and scored, though, the playbook can expand. Who gets to define a company's official assessment of a variant, and what processes go with that assessment, will be part of an implementation.

Playbooks will, of course, evolve, so there are version control issues. Which version was used to assess a given contract is important metadata.

While it should be relatively simple to total up the overall goodness of a contract version from each party's perspective (as the weighted sum of the goodness of each version's aspects, where the weight is the relative importance of that aspect), just as in document assembly there are situations in which inter-provision logic may add wrinkles. For example, there may be versions of two provisions that independently are optimal to one party, but sub-optimal when both are included. Preferences with respect to how one issue is handled can, however, be conditional on how other issues are being handled.

For some parties, in some circumstances, not addressing an issue is to their advantage. An effective way to model the required-ness of a provision is to encode the unacceptability (to one or more potential participants) of it being omitted.

Average contract professionals will not take to this approach unless it somehow *de*complexifies their work. The approach does, in several dimensions. It does so, for example, by codifying the

ranges of features about which "we don't care," quickly and accurately flagging ones that do pose non-negligible costs, and providing an apparatus within which individuals and groups can efficiently perform and track their decisional efforts.

Besides these practical complexities, computer scientists have identified serious computational complexity in such efforts as programming software agents to construct optimal allocations of resources in the context of e-commerce and e-trading environments. Dunne et al.[44] conclude that the problem is NP-hard. Ito et al.,[45] Klein et al.,[46] and Zhang et al.[47] also document the enormous problem and solution spaces that contract drafting can present.

In theory, contracts can be fiendishly complex, although in practice, mere mortals settle them all the time. Often many terms are uncontroversial or quickly resolved to the parties' mutual satisfaction. In practice, contracting parties only need to "box" the differences they care about. Being able easily to inspect versions, identify differences, and characterize the difference those differences make can render the entire process much more manageable. Even in situations of uneven bargaining power, the better organized and more determined party often prevails.

Keeping Score

Well-chosen contracts are those in which all parties have made informed and voluntary decisions about what is best, taking each other's interests into account. (Getting everything your way is not ordinarily ideal—a dispirited counterparty may not make for optimal performance.)

Organizations that find themselves contracting frequently—and thus recognize sufficient stakes to justify the expense of creating and maintaining playbooks—may find that it is cost-effective to take the next steps and inject one or both of the kinds of interactive intelligence described in this article. Such automation techniques improve quality, efficiency, and job satisfaction.

Individuals and entities that do not have the luxury of custom playbooks could greatly benefit from commercial, nonprofit, or governmental systems that alert them to contractual arrangements that best serve their interests.

Linking the rich representation of preferentiality that the framework described here offers with analytical tools and automated

assembly interfaces would deliver contract management functionality not presently seen elsewhere.

While largely a conceptual and interface innovation, choiceboxing in particular presents new opportunities to encode and disseminate knowledge. As applied to contracting, there can be considerable institutional value in codifying preferences in forms that can be visually inspected. A mixed-initiative system that presents a dynamic view of where contending drafts of a contract stand in relation to such preferences, while intelligently surfacing alerts and composing appropriate texts, could significantly improve outcomes.

Musical scores serve well to notate the interplay of melodies, whether consonant or dissonant. Deals are dances. We could use a graphical notation for the choreography of contracting.

Notes (continued from Part I)

* Marc Lauritsen, president of Capstone Practice Systems, is a lawyer, educator, and software developer. Mr. Lauritsen taught in and directed the clinical program at Harvard Law School, and directed Project PERICLES, a research program in law and computers. Mr. Lauritsen may be reached at marc@capstonepractice.com.

34. This section is repetitive of ones in recent papers, but is included here for easy reference by those unfamiliar with the method.

35. For some thoughts on these possibilities, see Lauritsen, M. "Boxing" Choices for Better Dispute Resolution. International Journal of Online Dispute Resolution (1) 1: 70-92 (2014); *available at* http://www.elevenjournals.com/tijdschrift/ijodr/2014/1/IJODR_2014_001_001_005.pdf.

36. http://www.americanbar.org/publications/law_practice_today_home/law_practice_today_archive/december11/dancing-in-the-cloud.html.

37. http://jolt.law.harvard.edu/assets/misc/Lauritsen-DecisionSpace.pdf.

38. *See* Lauritsen, M. On Balance, In Proceedings of the Fourteenth International Conference on Artificial Intelligence and Law, 83-91 (2013).

39. Lauritsen, M. 2017. The Centrality of Choice in Legal Work. *Suffolk UL Rev.*, 50, p.447, *available at* http://ssrn.com/abstract=3049297.

40. *See* Collaborative Deliberation through Interactive Visualization, A Choiceboxing Case Study, *available at* https://t.co/UOCw1nTr1V.

41. I owe this example to George Triantis.

42. Dijksterhuis, A., Bos, M., Nordgren, L., & Van Baaren, R. On making the right choice: The deliberation-without-attention effect. Science, 311 (5763), 1005-1007 (2006).

43. Lauritsen, M. On Balance, In Proceedings of the Fourteenth International Conference on Artificial Intelligence and Law, 83-91 (2013) and

Sartor, G. The Logic of Proportionality: Reasoning with Non-Numerical Magnitudes. German LJ, 14, 1419 (2013).

44. Dunne, P. E., Wooldridge, M., & Laurence, M. The complexity of contract negotiation. Artificial Intelligence, 164 (1), 23-46 (2005).

45. Ito, T., Hattori, H., & Klein, M. Multi-issue Negotiation Protocol for Agents: Exploring Nonlinear Utility Spaces. In IJCAI (Vol. 7, pp. 1347-1352) (2007).

46. Klein, M., Faratin, P., Sayama, H., & Bar-Yam, Y. Negotiating complex contracts. Group Decision and Negotiation, 12 (2), 111-125 (2003).

47. Zhang, X., Klein, M., & Marsa-Maestre, I. Scalable Complex Contract Negotiation With Structured Search and Agenda Management. In Twenty-Eighth AAAI Conference on Artificial Intelligence (2014).