Towards a qualitative model of everyday political reasoning

Kenneth D. Forbus **Sven E. Kuehne**
Qualitative Reasoning Group, Northwestern University
1890 Maple Avenue, Evanston, IL, 60201, USA
{forbus, kuehne}@northwestern.edu

Abstract

Everyday political reasoning seems to have many properties in common with everyday physical reasoning: it involves continuous parameters, numerical values are typically unavailable, and causal reasoning about qualitatively distinct behaviors is important. Can the techniques developed by the QR community be used to develop qualitative models of everyday political reasoning? This paper explores that question, using understanding texts about world history and current events as a focusing task. We outline some of the ontological issues involved, including modeling of emotions. We dissect a sample text to illustrate how QR ideas can be applied to understanding it, and discuss a large-scale corpus analysis in progress. Our conclusion is that QR techniques show promise in capturing important aspects of everyday political reasoning.

1 Introduction

Open any newspaper, and one finds statements like

"Mr. Koizumi says he understands the growing calls for sanctions against North Korea, but he believes a combination of dialogue and pressure is the best way to proceed." [1]

"Tired of the conflict, Russian commanders had already reduced the scale and tempo of offensive operations in the Caucasus, and the Chechens have done likewise in the last 10 days." [20]

Understanding statements such as these requires reasoning about continuous properties. The number of calls for sanctions, pressure, scale, and tempo can be thought of as continuous parameters and changes in these cause changes in the situations that they are part of. For instance, more calls for sanctions could lead to more pressure, and reducing the scale and tempo of offensive operations reduces casualties, materiel used, and tensions provoked, and the last of these might be one of the reasons why the opponents are doing the same. This kind of everyday reasoning about politics shares many properties with everyday physical reasoning. It involves continuous properties, but not detailed information about their numerical values: such information is either nonexistent or unavailable. Causal reasoning is crucial, to understand the consequences of events and to understand what about a situation might be changed to make it better. In other words, everyday political reasoning may be understandable, at least in part, in terms of qualitative modeling techniques. This paper explores that hypothesis.

It is often argued that everyday language involves metaphors, often with the physical world as the chosen base domain (e.g., "pressure" in the first quote above) [13]. This is indeed a fascinating question. One alternative is that we are actively computing comparisons constantly, applying our physical knowledge analogically in political thinking. Another possibility is that such language represents "frozen" metaphors, essentially alternate meanings, which do not require cross-domain comparison. Fortunately, for our purposes, this question is currently a side-issue. We are interested in how well qualitative representations can be used to model everyday political reasoning. The only thing that would change between the on-line mapping and frozen hypotheses is the domain that qualitative model is stated in. We believe that the best way to explore this issue is to tackle everyday political reasoning directly, to see how far we can use qualitative modeling in it. As we understand how far we can go, and what the theory looks like, we will have a better basis on which to evaluate whether or not there are physical metaphors that could lead to these kinds of models, and how much of it physical metaphors can explain.

In previous work [12], we argued that qualitative representations play an important role in natural language semantics because they capture the conceptual structure of many human mental models. An implication of our hypothesis that QR is applicable to political reasoning is that QR has an even larger role in natural language semantics than just helping to express knowledge about the physical world. If correct, understanding how to apply QR ideas to everyday political reasoning constitutes an important step towards being able to create software that automatically understand texts about world history and current events.

The rest of this paper outlines how we are using qualitative modeling, specifically QP theory [1], to capture the meaning of everyday statements such as those above. We begin by examining the ontological assumptions required, including the idea of a *political agent*, the kinds of quantities and processes involved, and the need for modeling emotions. Next we show how these ideas can be applied in a detailed analysis of an example. Then we summarize a corpus analysis in progress. Finally, we discuss related work not mentioned elsewhere, and planned future work.

2 Ontology

Everyday reasoning about politics can involve aspects of reasoning about the physical world (e.g., the impacts of geography on options) and the economic world (e.g., trade imbalances leading to tariffs). Since both of these areas have already been studied successfully in prior QR work (cf. [1,19] for economics), we ignore them here. Instead we focus on interactions between *political agents*. A political agent is a person, group, or organization which interacts with other political agents. Examples of political agents include individual people and politicians, interest groups (e.g., Basque separatists), political action groups and parties (e.g., MoveOn, the Libertarian Party), and individual nations and trans-national organizations (e.g., Spain, the UN, the Red Cross). Such agents are often hierarchical, e.g. cities can be part of counties, which can be parts of states or provinces, which in turn are parts of nations. Individuals are often members of multiple larger-scale political agents which cross-cut hierarchies, e.g., members of the program committee of an international conference are typically members of multiple distinct countries.

Each political agent can be thought of as having some of the properties of a person. For instance, the people of a nation can be thought of as angry if there is a widespread perception of being lied to by their government, even though not every person living there may be angry. By treating groups and organizations as if they were people, we can use our mental models of people as models for political agents. This is a powerful analogy for people, given that we have reasonably good models of interactions with other people, due to sustained experience with them. But it is a source of complexity for qualitative modeling, since we now have to model at least some aspects of human relationships.

There are several other sources of complexity that make political reasoning more ontologically challenging than physical reasoning. First, the laws of combination for properties of a group political agent based on the properties of its members are complex. For example, in physical modeling one can define aggregate extensive parameters via sums of the corresponding extensive parameters of the constituents, and the aggregate parameters are governed by the same laws as the parameters for the constituents. By contrast, the actions that can be taken on individuals, corporations, and governments when they are in debt are quite different. Moreover, inferring properties of members based on properties of aggregates is tricky: it can be the case that "North Korea wants unification with the South." is true, even if that same statement might not be true for every individual in North Korea.

The second source of complexity is that the number of types of potential interactions is significantly larger. In the physical world, there tend to be a small number of primitive types of processes (e.g., flows, transformations, creation/destruction). In the political world, there are diplomatic interactions (e.g., negotiating, exchanges, fact-finding), economic interactions (e.g., trade, embargos, tariffs), and military interactions (e.g., staging exercises, invasion, insurgency). This panoply of alternatives suggests that our job is best thought of as defining how QR is combined with other kinds of representations to support causal reasoning about continuous aspects of interactions, rather than attempting an exhaustive catalog of interactions. Something this open-ended is best constructed via machine learning, especially since the set expands as people invent new ways to interact (e.g., outsourcing financial services, using hijacked airplanes as large suicide weapons). Corpus analyses like those in Section 4 provide a starting point, setting the stage for co-training [1,17]

The third source of complexity is that interactions are often best viewed as discrete. For example, the imposition of a tariff, an insult, and an invasion are all best thought of as events without internal temporal structure for many purposes. This suggests that, in addition to the kind of integral relationship provided by QP theory's direct influences, we also need to model discrete changes. For this we can use Kim's *discrete process* representation [10], which provides the following operators:

```
(increaseQuantity <qty>)
(decreaseQuantity <qty>)
(increaseQuantityBy <qty> <amt>)
(decreaseQuantityBy <qty> <amt>)
```

Unlike I+/I-, the derivative of <qty> is not defined during the time the process is active. Discrete processes are defined the same as physical processes in QP theory, except that instead of direct influences, there is an *Effects* field that uses the operators above. The action of a discrete process is considered to be atomic, in that there can be no temporal sub-events that occur during it. If there are such events, the analysis needs to be broken down to a finer level of detail until the discrete processes at that level can be reasonably viewed as atomic.

2.1 Emotions

Political agents are treated as if they have emotions. For instance, countries that are afraid of something might exhibit irrational behavior (e.g., the Patriot Act in the US, which was passed just after 9/11). This means that we need a model of emotions that can be formalized in QP theory. We use the Ortony, Clore, & Collins [16] model of emotions, hereafter OCC, for this purpose. We will not describe a complete implementation of OCC in QP theory, both because this is a complex topic in itself, and it will require extensions to OCC. However, we can go far enough to provide confidence that this would be a reasonable endeavor.

OCC decomposes the appraisal process of generating emotions according to whether the emotions concern events, agents, or objects. Events are appraised with respect to whether or not they are desirable or undesirable with respect to one's goals. For example, one feels joy when winning a lottery, and distress when injured, because these events are evaluated as desirable with regard to a standing goal of maintaining economic well being and undesirable with regard to a standing goal of maintaining health, respectively. Agents (more exactly, the actions taken by an agent) are evaluated with respect to one's standards, the expectations about how that agent should behave. For example, one might feel pride in having written a program particularly well, and shame when a program turns out to have flaws that, in retrospect, one should have detected. Objects are evaluated with respect to one's attitudes. For example, one might love the interior of a new car, but find the color that its body is painted disgusting.

Each category of emotion is governed by a pair of parameters: *desirability* for event-based emotions, *praiseworthiness* for agent-based emotions, and *appealingness* for object-based emotions. In each pair, one parameter represents the sum of the positive contributions for that dimension, and the other parameter represents the sum of the negative contributions for that dimension. This lets us distinguish between situations where the emotions are consistent in direction but very weak versus where they are very strong but in conflict (as with the car example above). For brevity we will refer to these parameters by *desire+, desire-, praise+, praise-, and appeal+, appeal-*.

In [16] a scheme involving production rules and numerical parameters is proposed as a computational scheme for the theory. For example, when the sum of *desire+* and *desire-* is positive, a parameter JOY-POTENTIAL is assigned to a value determined by two unspecified functions:

If DESIRE(p, e, t) > 0
THEN *set* JOY-POTENTIAL(p, e, t) =
$f_j[|\text{DESIRE}(p, e, t)|, I_g(p, e, t)]$
where |DESIRE(p, e, t)| is the absolute value of a function that returns the degree of desirability that a person, *p*, assigns to some perceived event, *e*, at time, *t*, under normal conditions, and where $I_g(p, e, t)$ is a function that returns the value of the combined effects of the global intensity variables.

Another rule uses JOY-POTENTIAL to determine whether or not an emotion actually occurs:

If JOY-POTENTIAL(p, e, t) > JOY-THRESHOLD(p, t)
THEN *set* JOY-INTENSITY(p, e, t) = JOY-POTENTIAL(p, e, t) – JOY-THRESHOLD(p, t)
ELSE *set* JOY-INTENSITY(p, e, t) = 0

It is straightforward to translate this description into QP-style model fragments:

```
ModelFragment Possible-Joy
  :participants ?p a person
                ?e an event
  :conditions desire+(?p, ?e) > zero
  :consequences
   joy-potential(?p, ?e)
         ∝Q+ desire+(?p, ?e)
   ;; Global variables in OCC
   ;; affecting potential joy
   joy-potential(?p, ?e)
         ∝Q+ sense-of-reality(?p, ?e)
   joy-potential(?p, ?e)
         ∝Q+ proximity(?p, ?e)
   joy-potential(?p, ?e)
         ∝Q+ unexpectedness(?p, ?e)
   joy-potential(?p, ?e)
         ∝Q+ arousal(?p)
```

```
ModelFragment Joy
  :participants ?p a person
                ?e an event
  :conditions
   joy-potential(?p, ?e)
       > joy-threshold(?p)
  :consequences
   joy-intensity(?p, ?e)
     = joy-potential(?p, ?e)
        - joy-threshold(?p)
```

These model fragments introduce a number of new parameters. Where do the values for sense-of-reality, proximity, and unexpectedness come from? In a model that was created with a first-person perspective, these would be computed as part of an event perception system, which organizes incoming information into instances of known categories of events. In a model that was created from a third-person, predictive perspective, these would be computed from models that describe how another agent would perceive the event. (One of the classic failures of understanding between cultures is misjudging how others will perceive an event.) The contents of the event perception system (or the contents of another agent's event perception system) would be modeled via discrete processes, whose consequences include `increaseQuantityBy` and `decreaseQuantityBy` statements that set these parameter values.

What about arousal, the other parameter introduced above? In [16] arousal is described as partly physiological, somewhat dependent on non-emotional causes, and something that decays very slowly. Presumably arousal is a directly influenced parameter, with a simple decay process that directly influences it towards zero. Emotions fade with time, is decay in arousal sufficient for this? No, since arousal could rise or fall depending on all sorts of other things that can happen to an agent. Proximity, which [16] describes as a combination of both temporal and psychological closeness, provides a more appropriate mechanism. (Although proximity, too, can be influenced via arousal.) Again, a decay process that diminishes proximity towards zero will drain the emotion of its force over time, eventually making that instance of it inactive as it falls below the `joy-threshold`.

There is a symmetric quantity `distress-potential`, conditioned on desire- being greater than zero, which helps govern `distress`. Similar model fragments can be written for the other pairs of parameters governing the primitive emotions, which suggests that OCC can indeed be formalized in QP theory.

2.2 Processes and limit points

In the physical world, limit points are often relatively simple. Intensive parameters are compared against each other to condition flows. Parameters that govern existence (e.g., AmountOf or Mass) are compared against zero. Specialized values, such as boiling points and freezing points, mark phase transitions. In the political world, some limit points are straightforward, but many are not. For instance, whether or not a nation has nuclear weapons boils down to a comparison of the number of them that they have (if any) to zero. But there isn't a name for the level of disenchantment, anger, disillusionment, or whatever else is involved that marks the beginning of a civil war.

Sometimes physical metaphors are used to introduce political limit points. For instance, "flash point" is defined by dictionary.com as both "the lowest temperature at which the vapor of a combustible liquid can be made to ignite momentarily in air." and "The point at which eruption into significant action, creation, or violence occurs: "The shootdown did not increase international tensions to the flash point" (Seymour M. Hersch)." But this is more the exception that the rule. The different stages of warming or cooling in diplomatic relationships between two countries, for instance, do not to our knowledge have a well-defined set of names for the points at which transitions occur.

3 Extended Example

To see what an analysis of an everyday political argument might look like, we dissect a paragraph from an article in Stratfor, a private intelligence service [21], which analyzes some of the geopolitical consequences of US actions in Iraq as of June, 2004.

> As US-Iranian relations became increasingly strained during the winter, the Saudis increased their cooperation with the United States.

The most robust way to model relationships between political agents seems to be reification. By reifying them, we can add parameters as needed (e.g., intensity, strain, cooperation on particular topics) and they can be participants in other model fragments (e.g., embargo). From the context of the article, it is clear that "cooperation" refers to the war on terror. Given the range of issues and activities that political agents might or might not cooperate on, it seems wiser to introduce parameters for cooperation on specific areas than to model this via a single parameter. In addition to setting out the parameters to be thinking about, this sentence is also stating net changes in both over the period (i.e., increase in relationship strain for US-Iran, increase in cooperation regarding war on terror for US-Saudi Arabia).

> They [Saudi Arabia] also made it clear to the Americans that they were in danger of losing their balance as the pressures on them mounted.

Losing balance here is a disturbance from equilibrium that is sufficient to cross some (tacit) limit point, changing an aspect of their qualitative state that is normally constant. The pressures were radical Islamists working harder against the Saudi government, in response to the Saudi crack down on them, which in turn was caused by US pressure.

> The United States liked what it saw in the Saudi intensification of the war effort, even in the face of increased resistance.

Here the US, as a political agent, is experiencing an emotion. Since this pertains to the action of an agent, Saudi Arabia, by the OCC account Saudi actions were better than expected by the standard used by the US. The relationship between the Saudi government and the Saudi radical Islamists is described elsewhere in the article as "an incipient civil war", which is what the resistance parameter (whose value increased over the same (tacitly specified) time interval) is a parameter of.

> The United States did not like what it saw in Tehran, concerned that the relationship there was getting out of hand.

Again another emotion for a political agent. It is not Iran which is disliked here, but particular events (i.e., actions taken to become the dominant power in the Persian Gulf). "out of hand" here is a metaphor for "out of control", but we do not see any continuous parameters directly implicated here.

> Finally, in April, it [US] became completely disenchanted with the Shiite leadership of Iraq.

We believe this is best modeled as a limit point being reached in the quantity space of a quantity belonging to the relationship entity involving the US and the Iraq Shiite leadership. Further analysis would be required to figure out exactly which parameter: To us, either something like "degree of interest alignment" or "trust" would be equally reasonable, given the rest of the article. As noted in Section 2.2, this is probably an example of a tacit limit point.

4 Corpus Analysis

Analyzing a small example by hand is a good way to glean some initial insights. One tool for examining how far qualitative modeling ideas can go in capturing political reasoning is performing a corpus analysis. For example, [12] describes the results of analyzing a small corpus of sentences from an explanatory text about heat, temperature, and types of heat flow. This analysis was done by hand, since it was only 216 sentences. The syntactic realizations of QP theory constructs discovered during this analysis formed the basis for a controlled language, QRG-CE. QRG-CE is a subset of English with a restricted grammar and a set of interpretation rules that enables the construction of process instances from paragraphs of text. This raises several interesting questions:

1. Do the syntactic patterns that we found for explanatory physical texts apply to everyday political arguments?
2. If they do, what is their coverage? Put another way, how many more patterns are there?

To address these questions, we have moved to computer-assisted corpus analysis. We are currently working with the entire 1999 volume of the New York Times, consisting of 6.4 million sentences[1]. This is a daunting amount of data, requiring substantial automatic pre-filtering to reduce the data considered to a manageable level. In our preliminary experiments, we have used a multi-stage filtering process. The first stage uses regular expressions based on the vocabulary and syntactic patterns we found in the previous corpus analysis. The sentences that survived this first stage of filtering were then run through a modified version of our Explanation Agent NLU system [11], which broadened its criterion on what would be considered a possible quantity (specifically, we generalized to the Cyc concept `ScalarInterval`, which subsumes temperament, monetary values, feeling attributes, formality/politeness of speech, and a number of other concepts in addition to `Quantity` and `PhysicalQuantity`, which is what we restricted it to previously.).

Given that unrestricted newspaper text is far richer syntactically than QRG-CE, we did not expect to find much. However, since our parser is a bottom-up chart parser, it returns information about partial parses, including fragmented phrase information. Therefore we expected to find a large number of successfully parsed references to quantities. This was not the case. There certainly were quite a lot of them, but many fewer than we expected. While the EA NLU system can handle a variety of quantity references (including possessive phrases, adjectives, and spatial references), there are clearly other patterns that we need to capture.

What is the source of this difference? In addition to size, there are several other factors that make this corpus challenging. First, it covers all of the New York Times for 1999, not just politics. Continuous parameters and metaphors are used in sports, entertainment, and many other aspects of life, so filtering down to political arguments is tricky. Second, our previous corpus was explanatory text, designed for learners. By contrast, newspaper readers are assumed to have a broad background already.

One factor is that our set of words indicating directions of change was too small. The first stage of our initial analysis revealed that references to increases (by the noun and verb 'increase' and their morphological variants) were much more frequent than references to decreases by a factor of more than 22. For a follow-up analysis we analyzed the same corpus material for two separate, expanded lists of synonyms for increases and decreases. Instead of filtering for a list of references to known quantity types, the hypothesis is that references to quantities should occur in the context of increases and decreases. Especially in newspaper articles, changes are interesting to the reader and worth reporting.

Out of a corpus of 6.4 million sentences our filtering stage found 62,117 candidate sentences (~1% of the corpus material) mentioning decreases (based on a list of 66 items) and 195,482 candidate sentences (~3% of the corpus) mentioning increases (based on a list of 89 items). The ratio of 22:1 in favor of increases drops to about 3:1 if we allow synonym terms instead of just direct references. Overall, references to increases and decreases can be found in roughly 4 percent of the corpus material. This is a sharp contrast to the earlier analysis of a science text, where 43% of the material could be captured via QP theory. To be sure, we have not yet expanded the analysis to include qualitative relationships and processes. But the results so far indicate that qualitative representations may well play a smaller role in understanding political texts versus physical texts.

5 Related Work

In [4], Kamps and Peli argue that applying QR to social science domains requires shifting focus from simulation to model-building, and that model fragments are less likely to be broadly reusable. We agree with that assessment. One advantage of the hybrid similarity-based/first principles model of qualitative reasoning we assume [3] is that the model fragments can be only partially abstracted, remaining partially defined in terms of concrete situations and applied by analogy to new situations.

A number of models of emotion have been proposed. Our use of the OCC model is due to its focus on appraisal. For instance, Frijda's theory [5] has been implemented in two computer models [6,15]. The EMA model [7] has perhaps the most advanced implementation, driving virtual humans in an interactive simulation environment that uses speech recognition, natural language processing, gestures, and facial expressions to communicate with human users [8]. By contrast, our goals so far have focused on prediction, so appraisal seems like the natural place to begin. It would be useful to see if qualitative models could also be used to productively model other aspects of emotions as well.

[1] This is part of the AQUAINT corpus [14].

6 Discussion

Can qualitative reasoning techniques be applied to qualitative modeling of everyday political reasoning? While the final answer cannot be known at this point, we think the answer will turn out to be yes. While the ontology is much broader than the physical world, the same QR ideas appear to be applicable to important parts of it. Even emotion, which is perhaps the farthest topic from previous qualitative modeling efforts, seems to have important components which can be implemented directly in QP theory. The results of our corpus analysis to date suggest that this is indeed a complex endeavor, but only that the road is long, not that there are impassable obstacles.

As noted above, the large number of interactions and political agents makes by-hand construction of domain theories unlikely to succeed. Consequently, we are exploring two approaches. The more conservative is using pattern-mining techniques on our large corpora (cf. [22]). We can use this to answer some of the open questions raised by our analyses so far, e.g., whether or not the differences in quantity references we are finding come from being explanatory versus non-explanatory texts, or does it come from being political versus physical texts? The other is creating an expanded controlled language for expressing the kind of knowledge found in everyday political reasoning and world history texts, and having a system learn by reading [4].

Acknowledgments

This research was supported by DARPA IPTO and the Artificial Intelligence Program of the Office of Naval Research. We thank Dedre Gentner and Andrew Ortony for valuable comments.

References

1. Califf, M.E. and Mooney, R.J. 1999. Relational learning of pattern-match rules for information extraction. In *Proceedings of AAAI-99*.
2. Forbus, K. D. (1984). Qualitative Process Theory. *Artificial Intelligence*, **24**, 85-168.
3. Forbus, K., Gentner, D. 1997. Qualitative mental models: Simulations or memories? *Proceedings of QR97*, Cortona, Italy.
4. Forbus, K., Riesbeck, C. and Birnbaum, L. *Learning by Reading*. http://www.qrg.northwestern.edu/projects/LearningReader/index.htm
5. Frijda, N.H. (1986). *The emotions*. Cambridge, U.K., Cambridge University Press.
6. Frijda, N.H. und Swagerman, J. (1987). *Can computers feel? Theory and design of an emotional system.* In: *Cognition and Emotion*, 1, S. 235-258.
7. Gratch, J. and Marsella, S. "A Domain-independent framework for modeling emotion, " *Journal of Cognitive Systems Research,* Volume 5, Issue 4, 2004, pp. 269-306
8. Gratch, J. and Marsella, M. "Evaluating a computational model of emotion, " *Journal of Autonomous Agents and Multiagent Systems (Special issue on the best of AAMAS 2004)*, (to appear)
9. Kamps, J. and Peli, G. 1995. Qualitative reasoning beyond the physics domain: The density dependence theory of organizational ecology. *Proceedings of QR95.*
10. Kim, H. 1993. Qualitative reasoning about fluids and mechanics. Ph.D. dissertation and ILS Technical Report, Northwestern University.
11. Kuehne, S. E. (2004). *Understanding Natural Language Descriptions of Physical Phenomena*. Ph.D. thesis, Northwestern University, Evanston, IL.
12. Kuehne, S. E., & Forbus, K. D. (2002). Qualitative Physics as a component in natural language semantics: A progress report. In *Proceedings: 24th Annual Conference of the Cognitive Science Society*, George Mason University, Fairfax, VA.
13. Lakeoff, G. and Johnson, M. 1980. *Metaphors we live by*. University of Chicago Press.
14. Linguistic Data Corsortium (2002). *AQUAINT corpus*, Trustees of the University of Pennsylvania, Philadelphia, PA.
15. Moffat, D. und Frijda, N.H. (1995). *Where there's a Will there's an Agent*. In: M.J. Woolridge und N.R. Jennings (eds.): *Intelligent Agents - Proceedings of the 1994 Workshop on Agent Theories, Architectures, and Languages*. Springer.
16. Ortony, A., Clore, G. and Collins, A. 1980. *The Cognitive Structure of Emotions*. Cambridge University Press, Cambridge, England.
17. Pierce, D. and Cardie, C. 2001. Limitations of co-training for natural language learning from large datasets. *Proceedings of the 2001 Conference on Empirical Methods in Natural Language Processing*.
18. Shayman, S. 2005. Japan Resists Formal Sanctions Against North Korea. Voice Of America, 2/12/05
19. Steinmann, Christian 1997. Qualitative reasoning on economic models. *Proceedings of QR97.*
20. Friedman, G. 2004. Iraq, Iran, and Saudi Arabia. Stratfor, Inc., www.stratfor.com, 6/3/04.
21. Stratfor, Inc. 2005. Russia: Chechen Cease Fire and Pressure to Negotiate. 2/11/05
22. Yangarber, R., Grishman, R., Tapanainen, P., & Huttunen, S. (2000). Automatic Acquisition of Domain Knowledge for Information Extraction. In *Proceedings: 18th International Conference on Computational Linguistics (COLING-2000)*, Saarbruecken, Germany.