Johannes Castner, Dissertation Description

My dissertation work is concerned with the structures of American political causal reasoning and the mechanisms by which the variation thereof gives rise to the observed variation of political action. How do causal beliefs influence how people come to vote for politicians and how politicians, in turn, come to vote for bills; and how do people and politicians come to change their minds (if and when they do)? "Structure of Decision: The Cognitive Maps of Political Elites" (1976), henceforth "SD", edited by the political scientist Robert Axelrod, provides a framework by which to, quite naturally, study people's causal models of the world. As far as I know, it is the only existing framework that formulates models of deliberation and decision processes in terms of causal belief systems. The central claim in SD is that, having correctly elicited a person's relevant causal beliefs, we can make inferences about the rules by which (s)he will form new beliefs about the future or explain the past, and by which (s)he might choose between feasible political actions. The notion that causal beliefs are important in some way seems to be well accepted, judging from the enormous efforts devoted to establishing identification of causes and their effects in the policy oriented social sciences. Moreover, the strong normative language that often accompanies the identified causal effect in this literature makes no secret of its purpose to change some political actor's mind about both a causal belief and some implied political action. However, these causes and effects are most often treated in isolation and it is rarely asked how this new information about a causal effect, if it is found to be persuasive, would integrate with the actor's other beliefs. New knowledge of a causal effect, may, for example, create inconsistencies in the actor's other beliefs, which could lead to a radical remaking of her entire belief system. Strong changes in her belief system, in turn, could have stark political and social consequences for the political actor, because the new configuration of her beliefs may generate such a high degree of cognitive dissonance within her existing social and political network that the actor is either forced to leave her social surroundings altogether, or accept to be severely marginalized. On the other hand, as long as the differences in beliefs between members of a party are not greatly conflicting (there is little cognitive dissonance), greater diversity in beliefs (higher cognitive diversity) may be beneficial to the political success of the group. Indeed, it may allow the group to respond to a greater variety of current political needs and changes and appeal to more constituencies. The cognitive diversity of a party may enable its members to coordinate their discussions in such a way as to make them intelligible to a greater amount of voters, thereby securing a greater share of the votes.

To explore how cognitive diversity and dissonance of American parties have evolved over time, how they are related to a party's success and how they have influenced the nature of the debate on some key historic issues, I make use of a newly assembled longitudinal data set that contains the entire Congressional Record since 1873 (Jensen, Kaplan, Naidu, and Wilse-Samson, September 2012). Where the empirical component of the foundational work in SD was only able to analyse small samples of transcribed political discussions[1], thanks to recent innovations in computational linguistics, I will be able to perform similar analysis on a much

[1] I am gathering as much of this expert hand-coded data as I can find and I will use it as the gold standard, against which the automated extraction of causal assertions, as well as the hand-coding by Amazon Turk workers, can be judged (the automated coding should be in very close agreement with the coding by Axelrod and co-authors).

larger scale and extend this analysis to the "cognitive maps" of thousands of non-elite Americans (as described below).

My dissertation advisors, Peter Bearman and Suresh Naidu, are helping me in my quest to fill important theoretical holes, as well as to understand the limitations and advantages of my data. For the part of my work that involves computational linguistics tools and techniques, I am being advised by the computational linguists Nizar Habash and Chris Callison-Burch. Additionally, the computer scientist and graph theorist, Jonathan Gross, is advising me in my efforts to create a typology of belief systems and to construct meaningful measures of cognitive diversity and cognitive dissonance from the data structures that I obtain.

In what follows, I discuss in more detail the data structures describing causal beliefs and the methods I use to derive and analyse them. The unit of analysis, in SD, is a directed, labelled graph, which is a mathematical object that is obtained from a speaker's transcribed assertions and that represents the given speaker's causal belief system. As a practical matter, in order to distil such objects from text, it is necessary to find and to correctly decode the causal assertions of each speaker contributing to the text. As part of its analytical framework, SD provides me with the types of relation that I should be looking for in the text, which fall into seven categories (labels) of causal relation (positive, negative, zero, completely indeterminate, non-zero, non-positive and non-negative). The newly developed method of frame-semantic role labelling (Smith et. al. 2012) provides "a tool for automatic analysis of the frame-semantic structure of English text". By combining it with some additional tools [2], I have put it to use to detect and correctly label causal frames and their constituents (cause, relation, effect). While most of these tools have recently been built, tested and made available as open source software by computational linguists, up until now they have been extremely cumbersome to use and to apply to social science problems (by non-programmers and non-linguists) and I believe that this has prevented their widespread use in the social sciences. In the process of making these tools useful for my own purposes, as described above, I have teamed up with Evan Misshula (a programmer and PhD student in Criminal Justice at CUNY). Together, we have bundled these new tools and we will soon be able to make them available as one easy to use open-source package of computational linguistics tools with a convenient graphical user interface, so that social scientists can apply them to the problems that interest them without first needing to venture into linguistics to learn the details of their implementation and without needing to know how to program.

For a bounded empirical application of my new approach, I have teamed up with Marion Dumas, a fellow Ph.D student in Sustainable Development, to study American causal belief systems surrounding climate change; in Congress, from 1960 to today, (using the Congressional Record) and in today's public at large (using Amazon Mechanical Turk, henceforth AMT[3]). Nizar Habash originally advised me to use AMT as a means by which to test my new causal

[2] For example, I am using a powerful Noun Phrase Coreference System (Haghighi and Dan Klein, 2009) to make concept mentioning in texts commensurate ("tax hikes" and "tax increases" etc. are mapped to the same concept, and pronouns are replaced by the noun phrases that they stand in for). Note that this tool has only recently been developed and this kind of analysis would not have been considered possible by most computational linguists, three years ago.

[3] Amazon Mechanical Turk is a crowd-sourcing Internet marketplace that allows so called Requesters to coordinate the use of human intelligence to perform diverse tasks, such as choosing, among several, the most visually attractive photograph of a store-front, writing product descriptions, or identifying performers on music CDs.

parsing tool, but by discussing this issue with Chris Callison-Burch, I realized that with a few well-designed additional questions, I could also create a very useful data set, containing the causal assertions of a large sample of non-elite Americans.

To test my new causal-phrase parsing method[4] we will instruct AMT workers to label causal phrases and their constituents (causes, relations and effects). In order to simultaneously elicit the workers' judgements about politicians' statements about climate change, we will select phrases from the Congressional Record (from 1960 onward) at random, conditional only on belonging to a climate change related discussion[5]. Additionally, we will ask AMT workers to write very short essays about climate change and about how they think that it relates to a set of important variables. We will then parse these essays using the automatic causal-phrase parsing method that I have developed and analyse them according to the SD framework. We will specifically design our questions so as to elicit assertions about the causes of climate change and all factors that may be able to halt or reverse it, as well as the expected effects that climate change might have on the workers' own lives and on sub-national, national and inter-national political, economic and social conditions and processes. Of course, we will also present the workers with a host of more traditional questions which are essential to our work (their geographic location is of particular interest).

Using these exciting new tools and data sources, we will be able to address questions such as, how did American politicians' causal argumentation, surrounding climate change, evolve structurally since 1960 (as politicians have become more aware of climate change, have their causal belief systems become systematically more complex, have they become more cohesive), do today's citizens' causal argumentation structures, surrounding climate change, vary significantly in space and in particular, how do the citizen's causal belief systems compare to those of the politicians who represent them? Our research aim is to understand the mechanisms and reasons by which citizens and politicians come to argue and vote for certain policies, and our work is not laying claim to a particular set of policy implications. However, we do suggest that a deeper understanding of causal reasoning and deliberation in America, surrounding climate change, could aid the design of more equitable, effective, as well as efficient climate policy.

References

Aria Haghighi and Dan Klein. Simple Coreference Resolution with Rich Syntactic and Semantic Features. EMNLP 2009.

[4]This was recommended to me by both, Nizar Habash and Chris Callison-Burch. With the help of professor Callison-Burch, who is an expert on effectively and efficiently using AMT for NLP purposes, we are currently designing our questionnaire for this work. We are also working on an application to the Russel Sage Foundation Small Grants Program in Behavioral Economics in order to fund the part of our work that will make use of AMT. A recent study (Berinsky, Huber, and Lenz 2012) has shown that, while Amazon Turk workers are not fully representative of the American public (they over-represent the political left), they nevertheless do respond similarly to a representative sample on most political questions and in most experimental settings.

[5]For the purposes of testing my causal parsing tool, the topic of discussion should not make a difference, as long as causes, relations and effects are semantically and syntactically encoded in the same way across topics (most of the magic lies in deep syntactic and shallow semantic parsing).

Axelrod, R.M. (1976). “Structure of decision: the cognitive maps of political elites”, Princeton, N.J.

Bonham, G. Matthew, Michael J. Shapiro & George J. Nozicka. 1976. “A Cognitive Process Model of Foreign Policy Decision-Making”, Simulation and Games, vol. 7, no. 2, pp. 123-152.

Collin F. Baker, Charles J. Fillmore, and John B. Lowe. 1998. The Berkeley FrameNet Project. In Proceedings of the 36th Annual Meeting of the Association for Computational Linguistics and 17th International Conference on Computational Linguistics - Volume 1 (ACL ’98), Vol. 1. Association for Computational Linguistics, Stroudsburg, PA, USA, 86-90.

Dan Klein and Christopher D. Manning. 2003. Accurate Unlexicalized Parsing. Proceedings of the 41st Meeting of the Association for Computational Linguistics, pp. 423-430.

Dipanjan Das, Andre Martins, Nathan Schneider, Desai Chen, & Noah A. Smith. 2012. Language Technologies Institute, Carnegie Mellon University http://www.ark.cs.cmu.edu/SEMAFOR