CHR for Social Responsibility

Veronica Dahl[1], Bradley Coleman[2], J. Emilio Miralles[3], and Erez Maharshak[4]

1 **School of Computing Science, Simon Fraser University**
Burnaby, B.C. Canada
veronica@cs.sfu.ca
2 **School of Computing Science, Simon Fraser University**
Burnaby, B.C. Canada
bradley@proxydemocracy.org
3 **Department of Physics, Simon Fraser University**
Burnaby, B.C. Canada
emiralle@sfu.ca
4 **Cognitive Science Department and School of Computing Science, Simon Fraser University**
Burnaby, B.C. Canada
erez@proxydemocracy.org

Abstract

Publicly traded corporations often operate against the public's interest, serving a very limited group of stakeholders. This is counter-intuitive, since the public as a whole owns these corporations through direct investment in the stock-market, as well as indirect investment in mutual, index, and pension funds. Interestingly, the public's role in the proxy voting process, which allows shareholders to influence their company's direction and decisions, is essentially ignored by individual investors. We speculate that a prime reason for this lack of participation is information overload, and the disproportionate efforts required for an investor to make an informed decision. In this paper we propose a CHR based model that significantly simplifies the decision making process, allowing users to set general guidelines that can be applied to every company they own to produce voting recommendations. The use of CHR here is particularly advantageous as it allows users to easily track back the most relevant data that was used to formulate the decision, without the user having to go through large amounts of irrelevant information. Finally we describe a simplified algorithm that could be used as part of this model.

1998 ACM Subject Classification D.3.2 Language Classifications - Constraint and logic languages, H.4.2 Types of Systems – Decision support, K.4.3 [Computers and Society] Organizational Impacts

Keywords and phrases Constraint handling rules, principled decision making, informed voting, client directed voting, social responsibility.

Digital Object Identifier 10.4230/LIPIcs.ICLP.2012.370

1 Introduction

The financial crisis of 2008 showed us all how deeply corporations impact our lives, and in turn the importance of Corporate Social Responsibility (CSR). In this paper we argue that modern information technology can promote CSR, both from the perspective of investing in socially responsible companies or investment vehicles, and in using the voting rights that shareholders are given to impact companies in a way that can make them both more profitable and more socially responsible.

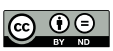

Technical Communications of the 28th International Conference on Logic Programming (ICLP'12).
Editors: A. Dovier and V. Santos Costa; pp. 370–380

Leibniz International Proceedings in Informatics
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

In particular we argue that Constraint Handling Rules (CHR) [3] is an effective tool for this, and we propose a CHR-based model of Informed Advice capable of recommending investment and voting decisions that are consistent with a user's stated preferences (values), and of explaining the reasons for these recommendations.

As well, we propose a practical and novel method of making automatic proxy voting recommendations based on a user's values. Currently, only around 5% [4] of retail investors vote their proxies, and the rate of informed participation is even lower. Our tool addresses the root causes of this dearth, namely information overload and under-load, and could increase the quantity and quality of decisions. We propose a method that only asks users to choose how aggressively they wish to vote on various classes of proposals, like sweat-shop labor, the environment, and director elections.

The paper is structured as follows: in section 2 we discuss the main issues around making effective and principled decisions. Section 3 presents the model itself, exemplified through principled decision making on investment in a company or mutual/index fund. In section 4, we describe our novel method of making automatic proxy voting recommendations. Section 5 presents our concluding remarks.

Consistent use of our proposed model would inform users who might otherwise be unaware of which decisions would clash with their value system and why. More importantly, it would spare users from going through the incredible amount of information needed in order to make informed decisions. Our model potentially allows presenting users with specific and focused information relevant to the decision at hand. Insofar as users become empowered to vote and decide according to their value systems, a wider range of stakeholders and owners would take part in the corporate governance process. This could promote a healthy evolution in management concepts and higher corporate governance standards.

2 Making effective and principled decisions

There is growing recognition of the fact that the public needs to be given choices that reflect their values and principles. For instance, ethical investment is increasing in popularity. Ethical investing indices are those that only include companies satisfying environmental or social criteria. Among these, Wikipedia names as examples those of FTSE4Good Index, and Dow Jones Sustainability Index. It also explains that strict mechanical criteria for inclusion and exclusion in such indices is important to prevent accusations of ideological bias in selection, as well as to prevent market manipulation, e.g. in Canada when Nortel was permitted to rise to over 30% of the TSE 300 index value[1].

However, it is not easy to effectively recognize which choices really do reflect one's values and principles. As also pointed out in Wikipedia, mechanical criteria can yield misleading results, since a firm can satisfy mechanical "ethical criteria", e.g. regarding board composition or hiring practices, but fail to perform ethically with respect to shareholders, e.g. Enron. There is indeed the risk that the seeming "seal of approval" of an ethical index may facilitate scams by putting investors more at ease.

For instance, self descriptions regarding sustainability or ethics could induce mechanical criteria to include the self-describing company among ethical ones, even though their ethical traits may be questioned elsewhere. Thus, a company such as Dow Chemical, which Wikipedia describes as responsible for actions that some consider unethical or against sustainability[2],

[1] http://en.wikipedia.org/wiki/Stock_market_index#Ethical_stock_market_indices
[2] Such as managing a nuclear weapons production facility that produced plutonium triggers for hydrogen

would probably be rated as ethical by an automatic rating agent looking at Dow Chemical's own website.

Even if we instruct automatic criteria to avoid using self descriptions, indices deemed to be sustainable could give high ratings even to companies that are believed not to be, and this inclusion alone could induce mechanical criteria to advise ethically-conscious investors to invest in those companies. A case in point, the Dow Jones Sustainability World Index recently rated the Dow Chemical Company as "one of the top performers in the global chemical industry", giving it the highest scores in the sector for operational eco-efficiency, customer relationship management and environmental reporting.

An alternative to mechanical criteria might be provided by market transparency and disclosure, but the problem remains how to collect the data into a knowledge base and how to use it together with users' stated principles and preferences in order to guide the search for advice that is consistent with those values. We shall argue that to solve this problem, we can resort to databases where the information is verified by specialists (e.g., international lawyers in charge of lawsuits concerning a questionable firm can verify whether the firm was found guilty), together with a CHR program which can be tailored for specialization into various applications that consult those databases (e.g. informed voting, informed investment, etc.).

3 Our proposed model for responsibly-informed decision making

As discussed, mechanical selection and disclosure and transparency are the present options for dealing with the fact that corporations are, in general, not trustworthy. None of these options is satisfactory, since mechanical selection can actually perpetuate scam by creating a false sense of security, and there always will be corporations that pretend to disclose and to exhibit transparency while not being totally sincere.

This dilemma between unsatisfying options can be overridden by a) placing relevant and reliable information in a database which can then be consulted, b) placing the description of users' values, companies' values and any other relevant information as initial constraints for a CHR program, and c) letting the CHR program run over the database when given some kind of "question", e.g. who should a given user vote for, or what decision should he/she make when faced with some specific problem.

We next introduce our proposed model through examples in the specific financial domain of choosing companies to invest on. There will be a system-defined part of the program, which the user needs not be concerned with, and which will adequately process the user's definitions.

3.1 Priority definitions

These are done through propagation rules, in which each criterion is associated with either a high, low, or medium priority[3]. To exemplify:

bombs (the Rocky Flats Plant), manufacturing napalm B, supplying the dioxin containing Agent Orange that was used as a weapon, or producing a soil fumigant, DBCP, which was responsible for sterilizing male workers in banana plantations in Latin America after most domestic uses of DBCP were banned in 1977 due to the successful lawsuits from workers at Dow's DBCP production who were made sterile by exposure to the compound.

[3] Alternatively, we could allow for numeric measures such as 0.9, or perhaps automatically translate less precise measures such as high, low, and medium into numeric values according to some algorithm that takes all the user's priorities into account.

```
priorities ==> priority(environmentalSafety,high),
 priority(humanSafety,high),
 priority(transparency,medium),
 priority(goodHistoricYield,medium), ...
```

3.2 Goal definitions

These are also described through propagation rules, e.g. the user can set goals such as to get at least 5% on average yearly, e.g.:

```
goals ==> goal(minimumAvgeYield,5), ...
```

3.3 The knowledge base: its sources, reasons, and trustworthiness

We use a knowledge base that lists for all candidate companies specific incidents that justify a given score with respect to each of the criteria. We do this through a 5-ary predicate "criterion", whose first argument names the criterion in question, whose second argument refers to a company, whose third argument rates the level at which the company satisfies the criterion (i.e., low, medium, or high), whose fourth argument summarizes the reason for said rating, and whose fifth and last argument records the URL from which this summary was extracted, e.g.:

```
criterion(humanSafety,'DowChemical',low,'because it sells
 chemicals that damage the human nervous system and have been
 banned from the US for that reason, to third world countries
 that do not yet have protective regulations',
 'http://en.wikipedia.org/wiki/Dow_Chemical_Company#DBCP').
```

The database also lists historic average yields and any other goal defined under "goals", e.g.[4]:

```
achievedHistorically(minimumAvgeYield,DowChemical,40).
```

It is important to consider from what sources the database will be constructed in each case. For instance having consulted Dow Chemical's own description, which presents itself as a sustainable company, the following contradicting information could co-exist in the same database:

```
criterion(humanSafety,'DowChemical',high,'because it addresses
 many of the world's most challenging problems such as the need
 for clean water, renewable energy generation and conservation,
 and increasing agricultural productivity',
 'http://www.dow.com/news/corporate/2011/20110908a.htm').
```

As we can see, the information to be included in the knowledge base can be unreliable in some cases, or partial, or even contradictory (as in the case of Dow Chemical listing itself as sustainable, which clashes with the belief of some that several of its actions were far from sustainable or ethical). We can adopt some criteria to decrease the risk of error in rating, such as adding a marker of trustworthiness to each of the above rules related to the quality of independent verification, e.g. the international lawyers that participated in the legal actions against Dow Chemical could give faith that the first rule is accurate.

[4] This is an example, not an actual yield.

We postulate that in the long run, the best way to achieve enough information is to allow for information to come from various sources, to allow for the potential contradictions that this can generate, and to simply output both the (perhaps contradicting) advice and its rationale in each case. The human user can then make up his mind on which of the info to follow, or in a later stage of our system, we could resort to argumentation theory [7, 1, 2] to weigh the merits of each argument and counter-argument automatically.

Our proposal is innovative because previous solutions to the problem of preventing market manipulation are based on either mechanical selection, which is always less discriminative than humans', and as we saw facilitates scams, or on just trusting the corporations, which many people believe to be imprudent at the least.

It is important to note that in a world in which web documents can contain timely information not easily found elsewhere, and in which concept extraction from web documents is becoming more and more efficient, we will likely tend to rely increasingly on them as sources of information. For this additional reason we believe that it is important to allow contradicting information to enter our database, while giving the user appropriate tools to deal with it a posteriori, such as argumentation theory. In any case, since some of the issues at hand might be controversial, we feel it is best to allow diverse points of view to be reflected, and leave it to the user to decide which one he or she wants to adhere to. The important thing for this is that we provide the rationale and the URL source which will allow them to make a truly informed decision.

3.4 The system

The system's shell itself can be created in just a few, relatively simple, CHR rules. To exemplify, the following CHR rule expresses that a Company meets a given criterion and goal if the user's values for that criterion and goal coincide with the values of the company. Now since the goal is met, the reason and its justifying link get printed as well.

```
priority(Criterion,Value), goal(G,N)
 ==> criterion(Criterion,Company,Value,Reason,Link),
 achievedHistorically(G,Company,N1),
 N1> N,
 print_reason(Company,Reason,Link)
 | ok(Company).
```

Of course, we could modify such rules in various ways as needed, e.g. demanding that the user's value either coincides with, or is less than, the value the company gets, or we could check all goals, not just the single one of the example.

The result will be a list of companies that are eligible according to each goal and criterion. This can be sorted out by the user or further processed by the system.

4 Principled and informed voting

The above will rapidly turn relevant as financial data-tagging becomes the standard practice mandated by the U.S. Securities and Exchange Commission[5]. This will allow matching various complex criteria in large data-sets. One especially practical and interesting implementation

[5] http://www.forbes.com/sites/tomgroenfeldt/2011/09/16/mandated-data-tagging-makes-sec-reports-useful-to-investors/

would be matching individuals' beliefs and voting principles with financial and proxy data. This would assist users in arriving at a voting decision when a large volume of data is available and only a tiny subset is relevant for the decision. One simplified possibility for this will be explored in the next section.

4.1 Motivation

As corporate power grows and the power of governments falls, mechanisms to govern corporations become more important. As governmental power falls, their *power* to regulate corporations falls as well. Further, as the influence of corporations over governments increases (eg. lobbying), the *will* of governments to regulate corporations also falls. This can form a positive feedback loop.

Hope is not lost though, since there is an existing structure at every publicly traded company wherein its shareholders – the corporation's actual owners – vote on high level decisions at the company. This mechanism is aptly referred to as corporate governance. Loosely, shareholders do not vote on which product the company should release, but instead they help elect a board of directors that will steer management well.

Furthermore, if a shareholder owns a fixed amount of shares, they may place a proposal of their own on the company's ballot for all of the shareholders to vote on. This is known as a shareholder proposal. As a policy, management supports management proposals and opposes shareholder proposals, because if management agreed with a shareholder proposal, then they would implement it themselves. Shareholder proposals can range in subject matter from the genocide in Darfur, to the environment, to sweat-shop labor, to executive compensation, to disclosing political contributions, to amending the corporate by-laws to, for instance, separate the Chairman and CEO positions. These proposals can call for an action, or for disclosure from the company. Since the range of management and shareholder proposals is so wide, good voting can lead to corporations that are both more profitable, and more socially responsible.

Shareholders are the literal owners of a company. Indeed, the ownership relationship is very powerful, arguably more so than the citizenship relationship between citizens and their governments, yet only about 5% of retail investors vote their shares [4], and even fewer investors cast informed votes. This is because of apathy, information overload, and information under-load.

ProxyDemocracy.org is a non-partisan, non-profit website that attempts to solve these (and related) issues, primarily in the U.S.. The website is essentially an interactive database of votes cast by institutional investors. Most of these votes were collected by scraping mandatory disclosure filings which are disclosed after the votes have been cast. However, for the purposes of this paper, we will focus on institutions who disclose votes prior to meetings. There are ten such institutions whose votes can be found on ProxyDemocracy. These institutions primarily predisclose in an effort to gain support for their positions by getting out ahead of other investors. Other investors, very large and very small, can and do see these early votes on ProxyDemocracy.org and use them to inform their own voting decisions.

4.2 Client directed voting

Client directed voting (CDV) seeks to improve voting rates and/or voting quality by automating the voting process. Mark Latham introduced this concept in [5], but [6] is more current. Still, the term "Client Directed Voting" was coined by Stephen Norman of American

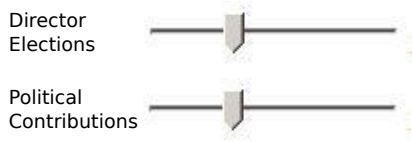

Figure 1 Left is passive, right is activist.

Express in 2006. There are many ways to envision this and to implement it. Here we propose a method which we argue will improve both participation and voting quality.

The state-of-the-art in computing sciences is such that modern information systems could aggregate the votes of institutions that publish their votes before meetings and use them to make voting recommendations congruent with a user's stated values. The user need only declare his voting preferences once and these can be used to automatically generate suggested votes on an on-going basis. The user would still likely need to approve each suggested vote for regulatory purposes. Automated voting is surely not foolproof, but the automation largely overcomes information overload, and the basis of respected institutional investors (that have the expertise and resources to vote well) largely overcomes information under-load.

We will use the term *passive* to describe a vote in favor of management, and *activist* to describe a vote against management. There is no value judgement intended by this word choice. We use these two terms to abstract the type of proposal so that we do not have to distinguish between management and shareholder proposals and votes that are "For", "Against", or "Withhold". Each vote will simply be considered to be either activist or passive.

The back-end of this CDV implementation requires predisclosing institutions, their voting histories, and a simple mechanism for automatically classifying proposals by type (e.g. director elections, political contributions). ProxyDemocracy currently shows the early votes of ten institutions, as well as their voting histories broken down by issue type, so these three requirements are straight-forward and attainable with a moderate effort.

The front-end requires only that for each issue type, the user indicates how activist or passive they want to vote. The user will indicate how activist or passive they wish to vote on each issue type by moving a slider, where the leftmost position will be as passive as possible, and the rightmost position will be as activist as possible (see Figure 1).

Voting decisions will be based entirely on the early votes of these predisclosers, their voting histories, and the user's slider positions.

This algorithm does not ask users to decide which predisclosing institutions they want factored into their voting decisions. This is a tempting, but problematic design choice, because most users are not familiar with these institutions, nor do they have opinions on their voting records. In fact, this algorithm only asks users the one question that any CDV scheme that aims to help users vote their values must ask, namely how they want to vote on the issues.

Furthermore, the algorithm we propose uses each and every predisclosing institution to make the decision. This is regardless of whether, for some issue, an institution's voting record aligns with the user's given values or not. This approach is inspired by the mathematical field of Information Theory, in that there is useful information in each and every voting decision, and thus each one should be factored into the final decision.

4.3 Our algorithm

This algorithm only decides a users' vote for one single proposal. For this vote, assume that k institutions have disclosed their vote by this time. We will use all k votes to make our decision. Let v_i be the vote of institution i, where $v_i = \{\text{activist, passive}\}$. These votes are represented by:

$$v_1, v_2, \ldots, v_k.$$

Also, let p_i, where $0 \leq p_i \leq 1$, be the historical probability of institution i casting an activist vote. More specifically, p_i is exactly the frequency with which institution i has cast activist votes on this issue in the past, which we know from this institution's voting record. For example, if some p_i is close to 0 on this issue, then institution i is passive on this issue. These probabilities are represented by:

$$p_1, p_2, \ldots, p_k.$$

For a given issue (like director elections, see figure 1), the position the user chooses will become the value s, where $-1 \leq s \leq 1$. For example, if $s = -1$, then the the user is maximally passive on this issue, if $s = 0$, then the user is neutral, and if $s = 1$, then the user is maximally activist. The users voting preference for this issue is represented by s.

Note that the the p_i for these k institutions may conform to any probability distribution. For instance, all k institutions might be very passive. We will use a weighting function $f : \{1, \ldots, k\} \to \mathbb{R}$ to attempt to mitigate any bias in this probability distribution by assigning each vote a weight. The function f is just one possible weighting function, there may be many that work well.

$$f_i = \begin{cases} 1 - p_i & \text{if } v_i = \text{activist}, \\ p_i & \text{if } v_i = \text{passive}. \end{cases}$$

The function f assigns a very small weight to a passive vote of an institution that votes passively on this issue, and similarly assigns a very small weight to an activist vote from an institution that is an activist on this issue. Conversely, f will assign a large weight to a passive vote from an activist institution, and the symmetric. Thus, if all of the predisclosers are historically passive on an issue, but the user wants to be activist, then if at least one predisclosers cast an active vote, this scheme will weight that vote high and the passive votes low.

The intuition behind this is that if an institution that is passive on directors votes against a director, then this director must be significantly unsatisfactory. Similarly, if a very environmentalist institution votes against an environmental proposal, there is a high likelihood that there is a significant problem with this proposal.

Now we will explain the second weighting function $g : \{1, \ldots, k\} \to \mathbb{R}$. This function scales some of the weightings that f created, and using the slider value s, creates new weightings, g_i:

$$g_i = \begin{cases} f_i & \text{if } v_i = \text{activist and } s \leq 0, \\ f_i(2 - (1 - s)) & \text{if } v_i = \text{activist and } s > 0, \\ f_i(2 - (1 + s)) & \text{if } v_i = \text{passive and } s < 0, \\ f_i & \text{if } v_i = \text{passive and } s \geq 0. \end{cases}$$

Again, function g is just one of many possible weighting functions. Further, it is limited in that, if the user is passive on this issue, it will only scale the passive votes by at most two. Observe that if the user picks the value -1, it will multiply the weight of each passive vote by 2. The function g can indeed scale either the passive or activist votes by at least 1 and at most 2. This scaling factor is an arbitrary choice, and can likely only be justified with rigorous experimentation with this and other weighting functions.

The intuition behind f and g is that first f helps to mitigate the bias from the vote sources, and then g applies the user's bias. Finally if $\sum_{v_i=Activist} g_i > \sum_{v_i=Passive} g_i$, then we cast an activist vote, and otherwise we vote passively. Below is an example.

Prior to demonstrating this simple approach in the next section, it is worth noting that we have tested it on a very large data set of approximately one million historical proxy votes, which has so far showed the effectiveness of the general approach. You can find a detailed account of the trial in this URL: `https://docs.google.com/open?id=0B57uHUYhCLdCRi15RGtoQlRvX2c`.

4.4 Example

In this example, we show how four institutions voted in a 2010 proposal to elect Charles Prince to serve on the board of directors of Johnson & Johnson. Some believe Mr. Prince to be a controversial director. He served as the chairman and chief executive officer of the investment bank Citigroup from 2002 until his resignation in November of 2007, shortly before the financial crisis of 2008. Thus, some claim he might have been involved in the decisions that, shortly after his resignation, brought about Citigroup's collapse (market capitalization crashed from $244B down to $20B) and subsequent bailout by the U.S. federal government. While some people may question his competence, his experience and influence as a chief executive and banker may be important to Johnson & Johnson. He was first appointed in 2006 and continues to serve on Johnson & Johnson's board.

The institutional voters for this proposal include CalSTRS, the pension fund for the teachers of California, AFSCME, a labor union pension fund, CBIS, a Catholic ministry pension fund, and Vanguard, the world's largest mutual fund company with more than $2T in assets under management. The first three institutions predisclose their proxy votes, but Vanguard does not. Hence, we would not have been privy to their vote (vi value) in advance. Yet, being a past proposal we can include it just to represent a more conservative school of thought.

In the first table, we are given the p_i's and the v_i's of the predisclosing voters (including Vanguard, which did not actually predisclose), and we compute the f_i's and the g_i's for four different values of s.

Institution	p_i	v_i	f_i	g_i (s=-.7)	g_i (s=-.3)	g_i (s=0)	g_i (s=1)
CalSTRS	.470	Activist	.530	.530	.530	.530	1.060
AFSCME	.438	Passive	.438	.745	.569	.438	.438
CBIS	.500	Passive	.500	.850	.650	.500	.500
Vanguard	.085	Passive	.085	.145	.111	.085	.085

In the second table, we show the sums of the activist and passive votes for each of the four values of s and the vote that is actually cast. With these four predisclosers, when $s = 1$, the algorithm casts an activist vote. This is because the sum of the weighted passive votes is less than the one weighted and scaled activist vote cast by CalSTRS (an activist vote is a vote against Mr. Prince). Similarly, the algorithm casts a passive vote when s is $-.7, -.3$

or 0. Note that if the f value of CalSTRS, which is .530 had been only slightly higher, the algorithm would have voted for Mr. Prince regardless of the s (slider value) chosen by the user.

s	$\sum_{v_i=Activist} g_i$	$\sum_{v_i=Passive} g_i$	vote cast
-.7	.530	1.739	Passive
-.3	.530	1.330	Passive
0	.530	1.023	Passive
1	1.060	1.023	Activist

5 Conclusion

We have presented arguments in favor of developing computerized systems for responsible decision making based on CHR, and exemplified our ideas in the context of automatically helping a user choose among investment possibilities in accordance with the user's values. We have also proposed an algorithm for client directed voting which can be readily incorporated as well into our CHR based system, thanks to its high modularity, and implemented a toy CHR program as proof of concept (see: `https://docs.google.com/open?id=0B57uHUYhCLdCRi15RGtoQlRvX2c`)The idea that the world is a symbol for thought, or thought materialized, can be postulated at various levels, from the most literal to the most metaphysical. By allowing principled thought guided by humanistic ends to become a matter of fact embedded in computer systems we can contribute both to humanize computers and to expand human consciousness in ways direly needed at the present juncture in our civilization. In particular, automating informed decision making could, through forcing corporations to choose between meeting users' values or losing their support, revolutionize the way corporations are run enough to transform them into agents of positive change. With this paper we hope to stimulate further work along these lines.

It is also important to note that, while we have focused on the specific areas of financial advice and voting, our described methodology can be readily adapted to other areas where informed decision-making can be supported by computers. Thus, other than by empowering voters and investors, the social implications of our proposed model's research could be mind-boggling from the point of view of potential systematic contributions to societal participation (through making wider consultation possible) and to the elevation of the world's educational levels, given that even poor areas of the world are gaining affordable access to mobile phones to which the needed databases could be fed and consulted. For instance our informed advice system adapted to medicine could allow patients to make more educated, conscious, and personally germane decisions on their treatments through relevant automatic expansion of the information received at a doctor's visit. In the long run, our research will make it possible to implement a model of machine informed human cognition around guidelines that consistently focus on human values and concerns, hence promoting an expansion of global consciousness around humanistic lines, with profound transformational effects.

6 Acknowledgements

This work was supported by V. Dahl's NSERC grant 31611024. We would like to thank the anonymous referees for their feedback, Henry Saint Dahl for timely advice, and Andy Eggers at ProxyDemocracy/the London School of Commerce for his work on the database of votes.

References

1 Xiuyi Fan and Francesca Toni. Assumption-based argumentation dialogues. In Toby Walsh, editor, *IJCAI*, pages 198–203. IJCAI/AAAI, 2011.

2 Xiuyi Fan and Francesca Toni. Conflict resolution with argumentation dialogues. In Liz Sonenberg, Peter Stone, Kagan Tumer, and Pinar Yolum, editors, *AAMAS*, pages 1095–1096. IFAAMAS, 2011.

3 Thom Frühwirth and Frank Raiser, editors. *Constraint Handling Rules: Compilation, Execution, and Analysis*. March 2011.

4 Chris Kentouris. Broadridge to u.s. corporations: Tell employees to vote. `http://www.securitiestechnologymonitor.com/news/-27412-1.html`, March 2011.

5 Mark Latham. The internet will drive corporate monitoring. *Corporate Governance International*, 3(2), 2000. Available at votermedia.org/publications. 1999 version at SSRN eLibrary".

6 Mark Latham. Proxy voting brand competition. *Journal of Investment Management*, 5(1), 2007. Available at votermedia.org/publications.

7 Iyad Rahwan and Guillermo R. Simari, editors. *Argumentation in Artificial Intelligence*. Springer Publishing Company, Incorporated, 1st edition, 2009.