The Norwegian Oil Fund Investment Decider N.O.F.I.D.*

Daniel S. Valland†, Tor Aimar Carlsen, Mikkel Lea,
Arild Pensgaard, Sebastian Karlsen, Pavel Sotov and Marija Slavkovik‡

Abstract

As the Norwegian Oil Fund keeps growing, so does the ethical impact of the investments it makes. Each investment made by the fund has to be subjected to considerable ethical evaluation to ensure its suitability. The trend of automating investing imposes the question of how to ensure that the decisions made by an algorithm are still within desirable ethical boundaries. To tackle this question we propose an approach to automating the ethical evaluation of investments through an ethically constrained decision aid system. To demonstrate the feasibility of our approach, we develop a prototype for an investment decision aid system called "The Norwegian Oil Fund Investment Decider" (N.O.F.I.D). The prototype enables automated suggestions for selecting stocks for investment by the Norwegian Oil fund. N.O.F.I.D maintains laws and societal norms, adhering to the ethical theory of rule utilitarianism. Similar decision aid systems may exist today, but to the best of our knowledge, none focus primarily on the ethical aspects of investments.

1 Introduction

Decision aid systems are information systems that can be used to support different kinds of business or organizational decision making activities (Keen, 1987). Such systems have been around for many years and are used to aid professionals in many different fields in making efficient and important decisions. Within the field of economics, decision aid stock trading bots have become quite efficient (McGowan, 2010). Fully automatic systems that deal with the day to day trading of stocks and hedge fund dealings are steadily becoming the norm and not the exception (Scopino, 2015). While the main object of an investment is to make profit, the ethical and moral aspects are also becoming more

*The authors of this paper are master students at the University of Bergen. The presented work was conducted as an independent student research project on machine ethics under the supervision of M. Slavkovik.

†Contact email: Daniel.Valland@student.uib.no

‡Contact email: marija.slavkovik@uib.no

Please cite as Valland et al.: The Norwegian Oil Fund Investment Decider N.O.F.I.D. (2017). Paper presented at NOKOBIT 2017, Oslo, 27-29 Nov. NOKOBIT, vol. 25, no. 1, Bibsys Open Journal Systems, ISSN 1894-7719

important (Bauer et al., 2006). As a consequence, the role ethics plays in recommending decisions by decision aid systems is also increasing.

As there is a variety of adverse situations that a company can be related to when trading stocks, such as child labour, war crimes and corruption, some funds hire ethics specialists to evaluate the companies that are being considered for a possible investment. The Norwegian Government Pension Fund Global, more commonly known as the Norwegian Oil Fund, is one of the worlds largest funds (Global, 2017). This fund has a council with ethics specialists which give advice whether investments in financial instruments are consistent with the Fund's Ethical Guidelines (SWFI, 2017).

We are here motivated by the need to implement ethical constraints in the decision aid systems that are used for trading stocks. In particular we are looking into a hypothetical automated trading system for the Norwegian Oil Fund and ways to implement the constraints typically imposed by their Ethical Council. We want to show the feasibility of a system that will function as an ethical council for an autonomous investment system.

Concerns of ethical behaviour of automated and autonomous systems are the subject of interest of *machine ethics*. Machine ethics is a new area of artificial intelligence that is concerned with accomplishing ethical behaviour from machines (Anderson and Leigh Anderson, 2007). By implementing an ethical theory in the decision process of an autonomous system, one could be able to ensure that the system makes ethical recommendations and decisions. There exist no agreement of what is the ideal ethical theory for machine ethics in general and trading in particular. The choice of which ethical theory is optimal for a financial decision aid system is an open and difficult problem.

For the purpose of our feasibility probe, we have settled on using *rule utilitarianism* (Haines, 2015). Numerous ethical theories have been developed and some have been considered for implementing in an intelligent autonomous system (Powers, 2006; Etzioni and Etzioni, 2017). Utilitarianism discerns between 'good' and 'bad' decisions by quantifying the 'good' and 'bad' each decision brings about. Since in the context of investments options are compared quantitatively, utilitarianism appears to be a simple first approach towards incorporating ethical assessment of decisions. We do however consider it a starting point in the discussion of how to build ethical decision aid systems, rather than a conclusive recommendation.

Utilitarianism, as all consequentialist theories, estimates how ethical an action is based on the consequences of that action. It can be difficult to estimate correctly what are the consequences of an action. Rule utilitarianism somewhat circumvents this consequentialist problem by considering "heuristic" rules of ethical behaviour, estimating an action as right when it conforms to a rule that leads to the greatest good (Atkinson et al., 1969). What is the greater good is, of course, also very difficult to quantify. We chose rule utilitarianism out of practical reasons - out of all the ethical theories developed in moral philosophy, rule utilitarianism is among the least person-centric and easiest to implement as an algorithm. Thus, in this work, we concern ourselves with the following problem: "How can we embed rule utilitarianism in an automated decision-making process for a financial investments system?" In particular we use as rules the principles of investment developed by the ethical council of the Norwegian Oil Fund.

To show that our approach is feasible we have developed a function high-fidelity

prototype. The user that this system is developed for are people that trade stocks on behalf of the Norwegian Oil Fund and similar investment funds. The prototype application is created using Java[1] and has a fully functioning and simple graphic user interface that allows the user to modify variables for the presented stocks.

The main focus of this project is within ethical decision making. Thus, our system is concerned primarily with aiding in ethical investments, and does not consider many other market factors such as risk. It does a basic assessment of profitability, but focuses more extensively on the ethical evaluation of stocks based on formalized ethic regulations provided by the Norwegian Government (Lovdata, 2014). The main contribution of this paper is to show that even with simple resources, a practical ethical trading system can be developed.

This paper is structured as follows. We begin by introducing the basics of the Norwegian Oil Fund, the moral theory we implement and a decision aid system in Section 2. In Section 3 we outline the assumptions and choices we made to implement our decision aid system. In Section 4 we describe the architecture of N.O.F.I.D., our decision aid system. In Section 5 we describe the built prototype. In Section 6 we discuss related work. Lastly in Section 7 we summarise our contribution, draw our conclusions and outline open questions and directions for future work.

2 Preliminaries

We begin by giving a brief description of the Norwegian Oil Fund and the ethical constraints under which it conducts its investing. We also give an overview of rule utilitarianism and a brief description of decision aid systems.

Government pension fund global (The Norwegian Oil Fund)

The government pension fund global, commonly known as "The Norwegian Oil Fund", was established in 1990 to underpin long-term considerations when phasing petroleum revenues into the Norwegian economy. Norges Bank Investment Management manages the fund on behalf of the Ministry of Finance of Norway, which owns the fund on behalf of the Norwegian people. The ministry determines the fund's investment strategy, following advice from, among others, Norges Bank Investment Management and discussions in Parliament. The management mandate defines the investment universe and the fund's strategic reference index. The Ministry of Finance has on a regular basis transferred capital to the fund from the Norwegian state's petroleum revenues. The fund's capital is invested abroad, to avoid overheating the Norwegian economy and to shield it from the effects of oil price fluctuations. The fund invests in international equity, fixed-income markets, and real estate. The aim of the investment strategy is to have a diversified investment mix that will give the highest possible risk-adjusted return within the guidelines set by the ministry.

When the Norwegian oil fund is considering an investment, ethical criteria exist that must be followed. The first criterion is the production of weapons. If a company is producing weapons which are used to break humanitarian rights, an investment should

[1]The N.O.F.I.D Java source code can be accessed at `https://github.com/danielvalland/NOFID`

not be made. The second criterion is the production of tobacco. The third criterion is the sale of weapons and military equipment to countries that are known for humanitarian rights. If a company breaks any of these three criteria, the fund should not invest in it Lovdata (2014). Even if a company is not breaking any of the criteria above, it is still not given that an investment should be made.

The three product criteria for exclusion cover the products a company is selling or making. The Norwegian oil fund also has to consider their behaviour and ethics. A company can produce something which is considered harmless according to the product criteria, but if the way they produce it is not satisfying, the investment should not be made. The criteria for behaviour includes violation of human rights, e.g. torture, murder, deprivation of liberty and the worst forms of child labour. Furthermore, involvement in war, corruption, environmental damage and serious violations of basic ethical norms are also criteria that should stop any investment from the Norwegian Oil Fund (Lovdata, 2014).

The work of the Council on Ethics is an ongoing monitoring of the Oil Fund's portfolio of investments with a purpose of uncovering any breaches of the criteria for products or behaviour. A typical evaluation process consists of four stages which are identification of the companies, selection of the companies, an in-depth assessment, and finally, a recommendation of the companies.

Formal ethical regulations exist that govern what the Norwegian oil fund can and cannot do. The questions of whether an investment adheres to these rules are known as compliance issues. The utility returned by the decision aid system we construct should reflect the degree of compliance a given investment has in relation to these rules. We have identified the following categories of ethical concerns that are impacted by compliance constraints by the Norwegian government (Lovdata, 2014): Human rights violation (§ 3. a), Production of tobacco (§ 2. 1b), Weapon production (§ 2. 1a & 1b), Child labour (§ 3. a), Corruption (§ 3. e), War crimes (§ 3. b), Violation of ethical norms (§ 3. f), Coal & Energy production based on thermal coal (§ 2. 2), Release of CO2 gas & other climate damaging activities (§ 3. d & e)

In the above list, we generalized the ethical concerns extracted from legal-text into problem classes such that they become more usable for classification purposes. In this case, these classes are translatable into moral rules or rules in accordance with rule utilitarianism in the form "The fund should not invest in companies engaged in X", where X is a given problem class. This results in the following rules:

- The fund should not invest in companies engaged in Human rights violation (§ 3. a)
- The fund should not invest in companies engaged in Production of tobacco (§ 2. 1b)
- The fund should not invest in companies engaged in Weapon production (§ 2. 1a & 1b)
- The fund should not invest in companies engaged in Child labour (§ 3. a)
- The fund should not invest in companies engaged in Corruption (§ 3. e)
- The fund should not invest in companies engaged in War crimes (§ 3. b)
- The fund should not invest in companies engaged in Violation of ethical norms (§ 3. f)
- The fund should not invest in companies engaged in Coal & Energy production based on thermal coal (§ 2. 2)

- The fund should not invest in companies engaged in the release of CO2 gas & other climate damaging activities (§ 3. d & e)

In addition to these rules, we assume that profit leads to good (pleasure), and so, there is a linear relation between the amount of profit an investment generates and the amount of pleasure produced in the Norwegian society. We therefore add the following rule (Lovdata, 2014):

- The fund should invest in stocks that have the highest possible expected return

Rule utilitarianism

Problems of recommending concepts of right and wrong conduct, as well as categorising actions as right or wrong are traditionally studied in moral philosophy. Many moral theories have been proposed, with *utilitarianism* being perhaps one of the best known among them. Utilitarianism is developed around the idea that an action is good if its consequences increase the well being, or utility, in the world.

Traditionally, utilitarianism has been interpreted as *act utilitarianism*, which holds that the utilitarian criterion should always be applied directly to individual acts (Harsanyi, 1977). This means that in any particular situation, an action is morally right when here and now it appears to yield the highest social utility. However, this form of utilitarianism has many practical implications which are inconsistent with commonly held moral convictions.

Harsanyi (1977) gives some examples of these practical implications which are at odds with our common sense morality. Consider for instance that our common sense morality says that we should keep our promises, except if this would impose disproportionally great hardships on us or on someone else. Here act utilitarianism would imply that we should always break our promises whenever this will have better, even if only slight better, direct social consequences than keeping our promise would. To amend these shortcomings of act utilitarianism (Harrod, 1936) proposed a revised form of utilitarianism, which is now commonly called *rule utilitarianism*. Harrod (1936) suggested that the utilitarian criterion should be applied in the first instance, not to each individual act, but instead to the moral rule governing this act (and to the general social practice that would result from abiding by the relevant rule).

According to Harsanyi (1977) an individual act should be considered to be morally right if it conforms to the correct moral rule which applies to the type of situation in question, regardless of whether it is an act that will or will not yield the highest possible social utility on that particular occasion. The correct moral rule on the other hand should be defined as that particular rule of behaviour that would return the highest possible social utility in the long run, if it was followed by everyone in the same situation. For example, except for certain emergency situations, it is our moral duty not to steal. Even if stealing would give you an increase in individual utility short term, in the long term, society will be better of if people can trust that their private property is safe.

When making an investment for the Norwegian Oil Fund, the investment team has to consider different facts regarding whether the final decision of making the investment is

right. On the one hand one has the economical perspective where one wants to accumulate the maximal amount of money for the fund. On the other hand one will have to look at the moral behind making certain investments. For example, it is our moral duty not to increase the amount of suffering in the world or its uneven distribution. This duty may be in contrast to making investments that give an increase in societal utility in terms of financial gain for the fund and thus improve the well being of the Norwegian people. In the context of the decision aid system for the Norwegian Oil Fund, the moral rules we use are the decisions made by the government and ethical council, regarding what types of investments are labelled as unethical and as such to be avoided.

Decision aid systems

Decision support systems (DSS) or Decision aid Systems are systems that can be used as support for different kinds of business or organizational decision-making. These systems typically produce a list of available alternatives, sort that list in order of desirability, or choose an option from the list as a recommendation of which action is best with respect to provided criteria. According to Keen (1987), a DSS is defined in terms of the structure of the task it addresses and ta DSS reflects an implementation strategy for making computers useful to managers. DSS systems are often used by mid and higher levels of management in order to make decisions about problems that are rapidly changing and which are not easily specified in advance.

The three fundamental components of a DSS architecture are: the *database (or knowledge base)* that is the area of where the program collects data and what data it takes into consideration when assisting in decision making, *the model* that is the context of what decisions are going to be made, and different kinds of user criteria for the domain of which it operates. Finally *the user interface* used to interact with the DSS.

Shim et al. (2002) outline that DSS researchers and developers should: (i) identify areas where tools are needed to transform uncertain and incomplete data, along with qualitative insight, into useful knowledge, (ii) be more prescriptive about effective decision making by using intelligent systems and methods, (iii) exploit advancing software tools to improve the productivity of working and decision making time, and (iv) assist and guide DSS practitioners in improving their core knowledge of effective decision support.

3 Assumptions and choices

Our system must be able to reflect ethical laws and norms present in the Norwegian society. We have chosen to build our system with the moral framework of rule-utilitarianism. Having made this choice we reflect laws and norms of ethical behaviour in the form of rule utilitarianism style rules. These laws and norms are those defined in Norwegian law and society, which have been formulated in a democratic process. Because of this, we can expect that the decisions our decision aid system arrives at would be acceptable to the Norwegian society.

Since our focus is to implement ethical constraining to investments, and not design a full automated investment system, it is necessary to make several assumptions in regard to basic market behaviours associated with the stock market. We assume that profit from a given investment translates to utility of well being of society within the scope

of our utilitarian theory. Our system is concerned with balancing the dilemmas that can exist between choosing between increasing profit and abiding by ethical rules. In our implementation rule utilitarianism, both profit and obeying the ethical rules lead to increase of utility.

Certain moral rules are intuitively more grave to violate than others. For example, tobacco production can be perceived as comparatively far less serious than war crimes. To capture this difference among constraints we chose to use numerical values to represent the priority of ethical constraints, namely we cardinally rank the ethical constraints. Alternatively, the decision aid system can be supplied with an ordinal ranking in the form A $>$B $>$C (where A,B and C are rules) by sorting the constraints in descending order according to their priority value.

In our prototype the user can edit the priority values of the ethical constraints manually. In principle, the decision of how the ethical constraints are ranked in the decision aid system is a decision that has to be made by ethical experts.

We consider having priorities between the ethical constraints to be advantageous for several reasons. The priorities of constraints can be adjusted to reflect the considerations of ethics experts employed by the fund. The ranking can be used to reflect changes in law, such as the passing of new bills of legislation. Furthermore, the ranking can also be used to reflect political changes in the world that necessitates a change in the priorities of constraints in the system.

We need some way of numerically representing a stock's expected profit. We elect to use a stock's Beta value (Investopedia, 2017) as an indication of expected profit. Beta indicates how a stock usually responds to market movements. A stock with a Beta of 1 moves exactly like the market, whereas a stock with a beta less than 1 is more stable than the market. Further, a stock with a Beta greater than 1 is more volatile than the market to which it is compared. Following from this, a stock with a Beta greater than 1 can be expected to outperform the market during market growth (Bull markets), and under perform the market during market downturns (Bear markets) (Investopedia, 2017).

We assume that the users of our system will be able to acquire or generate the data that the system needs to perform decision aid. We imagine that in a production environment, other systems can be used to gather and curate the needed data (such as through API's), and provide this to our system much like a software pipeline (where the output of one component is directly used as the input of another).

Lastly, we need a way to represent the notion of uncertain knowledge in the stock market, that is, a belief exists that a company is engaged in a given unethical activity. In the real world this information of ethical dubiousness is usually not certain, nonetheless an estimate can be given of it's certainty. We represented the notion of uncertain knowledge of company activities by use of *rumours*. Each stock has a set of rumours, each of which has a given probability of being true, and assigned to each rumour are the rules affected by the given rumour, if true.

4 Decision aid system architecture

The architecture of our ethically constrained decision aid system is fairly simple. It consists of an environment, an ethical evaluator that identifies the ethical cost of a decision

based on rumours, and a decision recommendation mechanism.

We defined the following components that constitute the environment the agent operates in: a stock market, a set of stocks, a set of rumours for each stock, and a set of ethical rules that the system must follow. The difference between our environment and the real world is the data that is provided as input to the system (test data versus real data).

From the rules formulated in the Section 2, we can observe that the rule of increasing profitability exists in potential conflict with the ethical rules derived from the constraints on investment we extracted from (Lovdata, 2014). This forms an ethical dilemma which we solve by allowing each ethical rule to be weighted (each rule has a weight constituting the importance/priority of following that rule). This way, the dilemma between rules (a stock is both ethical and unethical from different rules) are solved by prioritizing decisions that follow the rules with the highest priority/weight. The system does not resolve the case when two investment options are of the same utility, since the system produces a ranking of options rather than a choice of one investment option.

A stock market (Java) class is used to provide a container interface from which to select the most appropriate investments. To create a decision evaluation mechanism, we define a heuristic function that considers all relevant attributes of a given stock and returns a total utility value for that stock. The heuristic function considers the following attributes to compute the utility of a given stock: expected return of the stock, probability of each rumour associated with the given stock, and priority/severity of each rule violated by rumours.

The decision aid system starts by considering each rumour associated with a given stock. Each rumour has a given ethical constraint that is violated by the rumour if true. This is combined as:

$$\text{Cost} = \text{Probability} \times \text{Priority},$$

and is considered as a negative utility due to the violation of constraints. The decision aid system then considers the expected return of a stock as positive utility. The total utility of the stock is computed by subtracting the total negative utility incurred by the violation of constraints from the total positive utility incurred by expected profit. The decision of how to normalise the cost of a violated ethical constraint with the amount of earning in money by an investment is intended to be made by ethics experts.

After evaluating each stock using the heuristic function, the list of stocks can then be sorted according to their total utility, and the appropriate amount of stocks given as a final result of the decision aid. That is, the output of this process is a sorted list of stocks which the system recommends as ethical investments according to the defined rules.

5 Prototype

We implement a functioning high-fidelity decision aid system prototype in Java that uses our interpretation of rule-utilitarianism together with heuristic evaluation. Our prototype is able to rank a collection of company stocks with respect to their desirability for investment, based on rumours on whether the companies are behaving ethically and whether their products are ethical or ethically sourced.

The first stage in using the system is data gathering. Here, the system receives several

forms of input from the user (or in principle from another investment system) in order to gather enough data to be able to produce recommendations. The system receives the following input (all as CSV Files): a list of stocks, a list of ethical rules, and a list of rumours.

After the input is provided, data adjustment can be done to allow the system to reflect changes in the market, or the priorities of the system users. The user can manually adjust the priority of each ethical constraint, and also the probability of each market rumour. Clearly, the manual adjustment is only a possibility in the prototype. In a fully operational system, there would be considerations of access control to ensure that adjustments are made responsibly.

Stocks are given as input in the form of a CSV file containing the following information: Stock Ticker, Name of the stock, Last Sale Price, Market Cap of stock, IPO Year, Sector, Industry, Summary Quote, and The Stocks Beta Value. This schema is directly compatible with CSV files downloaded from `NASDAQ.com` (2017) that contains the needed information from all companies on US stock exchanges. The only additional column needed is a stocks Beta value. Here, we assume that other systems can provide this information if used in a real scenario (such as API's).

Rumours are given as rows of Stock Ticker, Label, Affected Rule and Probability. Where: Label is a textual description of the rumour, "Affected Rule" is a given rule that is violated should the rumour be true, and the probability is an estimation of how likely the rumour is to be true in the form of percentage (decimal number from 0-1 (0-100%)).

Ethical constraints are given as pairs of (Label, Priority) where label is a textual description of the constraint. Priority is a numerical value representing how ethically important the given constraint is compared to the other constraints considered.

When data entry and adjustment is complete, the user can click "Run Decision Aid", and the system evaluates the stocks based on all information given using heuristic evaluation. The user is presented with a recommended ethical stock portfolio ordered from most to least desirable stock investment. The system also provides some over-all statistics such as average ethical score (utility score) achieved by the selected portfolio, and average expected return based on Beta value. Further, the system also presents a benchmark, which is exactly the same information but considering only profitability, allowing us to compare the financial gain of the ethical with the non-ethical portfolio selection.

To verify the accuracy of our system, we have designed several test cases where the optimal choices for the system (if any) are known in advance, and we can compare the results suggested by the system, and our manual assessment of the optimal output. The objective of these tests are to verify that the rumours and rules defined effect the result in an expected manner. We start each case with an hypothesis of the output we expect based on the parameters (rumours and constraints) we have given as input to the system. Each test case is carried out using the same 5 random stocks, the only difference being the rules and rumours we specify for the system.

Case 1: Control. We start by running a control test with no rumours specified. Since no ethical constraints are violated by any of the stocks, we expect the ordering of the stocks given by the system as output to be unaffected by the evaluation process. An example of a run in Case 1 is given in Table 3 bellow.

Table 1: Case 1 Input

Stock Ticker	Expected Return
DDD	1.0
WUBA	1.0
WBAI	1.0
MMM	1.0
AHC	1.0

Table 2: Case 1 Output

Stock Ticker	Utility
DDD	1.0
WUBA	1.0
WBAI	1.0
MMM	1.0
AHC	1.0

Table 3: **Results of Case 1.** the output matches the hypothesis.

Case 2: Differing expected return In this case, we alter the expected return of each stock. Still there are no rumours defined. Here we expect the output to be sorted in descending order according to their expected return (Return = Utility). An example of a run in Case 2 is given in Table 6 bellow.

Table 4: Case 2 Input

Stock Ticker	Expected Return
DDD	1.0
WUBA	0.9
WBAI	1.5
MMM	1.1
AHC	1.15

Table 5: Case 2 Output

Stock Ticker	Utility
WBAI	1.5
AHC	1.15
MMM	1.1
DDD	1.0
WUBA	0.9

Table 6: **Results of Case 2.** the output matches the hypothesis.

Case 3: Rumours In this Case, we introduce a rumour to Case 2, and try to effect a change in the results. In Case 2, WBAI was selected by the system as the best stock to invest in. In this case, we add a rumour that WBAI is engaged in some activity that violates the rule of "War Crimes". We set the probability to be 0.8, and the rule priority remains unchanged at 1.0. We now expect the stock to appear much further down the list due to this violation. An example of a run in Case 3 is given in Table 10.

Table 7: Case 3 Rumours

Stock Ticker	Description	Affected Rule	Probability
WBAI	Test Rumour	War Crimes	0.8

Table 8: Case 3 Input (same from Case 2)

Stock Ticker	Expected Return
DDD	1.0
WUBA	0.9
WBAI	1.5
MMM	1.1
AHC	1.15

Table 9: Case 3 Output

Stock Ticker	Utility
AHC	1.15
MMM	1.1
DDD	1.0
WUBA	0.9
WBAI	0.7

Table 10: **Results of Case 3.** The system places the stock at the bottom of the list, and the output matches the hypothesis.

Case 4: Rule Priority In Case 3, the system regarded the severity of the rule violations to justify placing WBAI at the bottom of the list. In this case, we alter the priority of the rule "War Crimes" from the previous 1.0 down to 0.45, and run the test again. Now, we expect WBAI to appear much higher on the list, since the priority of the rule it is violating has reduced. An example of a run in Case 4 is given in Table 14.

Table 11: Case 4 Rumours

Stock Ticker	Description	Affected Rule	Probability
WBAI	Test Rumour	War Crimes	0.8

Table 12: Case 4 Input (same from Case 2)

Stock Ticker	Expected Return
DDD	1.0
WUBA	0.9
WBAI	1.5
MMM	1.1
AHC	1.15

Table 13: Case 4 Output

Stock Ticker	Utility
AHC	1.15
WBAI	1.14
MMM	1.1
DDD	1.0
WUBA	0.9

Table 14: **Results of Case 4.** The system places the stock 2nd in the list, and the output matches the hypothesis.

6 Related work

The field of ethical decision making within financial decision aid is a relatively new area of research, and few existing solutions consider ethics of financial decisions directly. Meredith and Arnott (2003) highlight the problem of lack of research on the ethics of decision aid systems.

Adam et al. (2012) explore the potential impact of decision support systems on the ethical challenges of public purchasing decision makers. Thus this work is concerned with the problem of should ethics play a role rather then how can ethics be implemented in public purchasing. They find that ensuring that decisions made in public purchasing are ethical is not addressed in the law, but that this is an important problem that should be addressed. Adam et al. (2012) conclude that future research is needed in the use of decision support systems as compliance tools.

An example of a clinical decision aid system developed with the purpose of enabling ethical behaviour is the MedEthEx system (Anderson et al., 2006). This system is constructed to evaluate ethical dilemmas in medical decisions (Anderson et al., 2006). While medical decisions and financial decisions may have little in common with regards to application, the ethical impact of such systems is comparable, especially with regard to the decision maker. Both systems incorporate some notion of principles or rules that should be followed. While we construct our system to implement rule utilitarianism, MedEthEx uses *prima facie duties* (Beauchamp and Childress, 2009; Ross and Stratton-Lake, 2002) as parameters in a machine learning approach towards training the system to discern between more and less ethical actions.

Moor (2006) distinguishes four types of moral artificial agents with respect to their abilities to make ethical decisions. Among these, *ethically explicit agents* are programmed

to distinguish more from less ethical actions. In contrast, *ethically implicit agents* are agents that are not programmed to distinguish right from wrong actions, but which have their unethical actions constrained with constraints supplied from the programmer. With respect to Moor's taxonomy, our system can be seen as an ethically implicit agent, while MedEthEx would be an example of an en ethically explicit agent.

Dennis et al. (2016) implement an ethically implicit artificial agent. They implement their agent using externally supplied prima facie inspired duties and a priority order over them to constrain the unethical actions of the agent. Theirs is not a decision aid system, but an intelligent agent and they are concerned particularly with verifying the ethical behaviour of the implemented agent.

Since our system implements an existing ethical theory, it can be considered to implement a *top-down ethical approach* (Wallach et al., 2008). The agent reasoning system (Dennis et al., 2016) is also an instance of the top-down approach. In contrast the MedEthEx system implements a *bottom-up approach to ethical reasoning* since the system is not supplied with an ethical theory, but develops one on its own using machine learning.

Other decision aid systems such as Land Allocation Decision Support System (LADSS) (Matthews et al., 1999) and Multi-level and Interactive Stock Market Investment System (MISMIS) (Cho, 2010) do incorporate the fundamental components of decision aid systems like our system, but use different and more complex methodologies. LADSS addresses issues regarding climate change, agricultural sustainability and rural development. MISMIS on the other hand is a decision aid system for stock market investment, focusing on maximizing the utility of an investment. The systems do address some social impacts in their respective fields, but both fail to explicitly implement an ethical component of the systems, like we do in N.O.F.I.D.

7 Summary

In this paper we considered the problem of enabling ethical sensitivity in decision aid systems specialised in making financial investments. Our goal was to demonstrate that this is a feasible task by creating a functioning prototype based on rule utilitarianism. In particular we used ethical investment constraints developed by the Ethical Council of the Norwegian Oil Fund. In general our work contributes to raising the very important issue of considering the ethical impact of the decisions that are outsourced to machines.The decision may be made by a machine, but it is has a real impact in the real world. Our work offers an initial strategy on enabling financial decision aid systems to make ethical decisions, of course within a very limited sense of the term.

To better understand the relevant motivation and approaches for financial decision making, we contacted the Council on Ethics, and received documentation that helped us understand how they conduct their work. We used the feedback to adapt and match the most important constraints, laws and norms to rule-utilitarianism which we implemented in our prototype. We represent the different ethical undesirability of particular rules with a numerical priority scale. The uncertain information of the unethical conduct of some products and companies we model as rumours.

What we implement is a very simple decision aid system. None the less, even this simple system shows that an ethical governor can be implemented in financial decision aid

system. Our test results show that the system is able to respond appropriately to a range of ethical scenarios. The results suggest that the system is able to properly prioritize and solve the dilemma of profit versus ethics. Clearly, the system also has numerous shortcomings. For example, we abstracted from the problem of creating the rumours and the ethical dimension they introduce to the decision making. Furthermore, unlike full ethical reasoning approaches in moral philosophy, we do not offer to solve dilemmas. A open problem in machine ethics is how much ethical reasoning is sufficient for an intelligent autonomous system. This question merits further discussion than the scope of this paper can offer.

Ethical decision aid systems are an emerging field, and many issues concerning them remain to be explored. Focusing in particular on automated trading, extensive future work remains to be done on the automatic gathering and generation of data that can be used by N.O.F.I.D and systems like it. We imagine that N.O.F.I.D ideally would be a part of a larger pipe-line system (one systems output is directly used as the input of another) of systems where earlier systems in the chain can automate the task of gathering and curating rumours and stock data. Further, the ethical decision aid provided by our system could also be acted upon by autonomous agents such as automated trading bots, at which point a fully automated system could be achieved.

In our work we assumed that the topic of the rumours of unethical conduct are given before hand - they are provided by the Ethics council. We conveniently assumed that the stocks information is available pre-annotated with rumours. Rumours however, can in theory also be automatically generated. Can one establish, and if yes how, possibly in real time, which ethical rules is a company rumoured to violate, and the credibility of those rumours, by processing the online information and data on that company?

We found that the use of rule utilitarianism served our purpose, but is this moral theory sufficient to capture all the real time flexibility and efficiency of the Ethical Council? For ethical automated trading to become a truly viable reality, the implementability of this and other ethical theories in this context must be meticulously examined.

References

Adam, F., Csaki, C., Prier, E., and Bufacchi, V. (2012). Ethical decision making and decision support systems in public procurement - a theoretical discussion. *Supply Chain Forum: An International Journal*, 13(3).

Anderson, M., Anderson, S. L., and Armen, C. (2006). An approach to computing ethics. *IEEE Intelligent Systems*, 21(4):56–63.

Anderson, M. and Leigh Anderson, S. (2007). Machine ethics: Creating an ethical intelligent agent. *AI Magazine*, 28(4):15.

Atkinson, R. F., Garner, R., and Rosen, B. (1969). Moral philosophy: A systematic introduction to normative ethics and meta-ethics. *Philosophical Quarterly*, 19(75):181.

Bauer, R., Otten, R., and Tourani Rad, A. (2006). Ethical investing in australia: Is there a financial penalty? *Pacific-Basin Finance Journal*, 14(1):33 – 48.

Beauchamp, T. L. and Childress, J. F. (2009). *Principles of Biomedical Ethics*. Oxford University Press, sixth edition.

Cho, V. (2010). MISMIS – A comprehensive decision support system for stock market investment. *Knowledge-Based Systems*, 23(6):626 – 633.

Dennis, L. A., Fisher, M., Slavkovik, M., and Webster, M. (2016). Formal verification of ethical choices in autonomous systems. *Robotics and Autonomous Systems*, 77:1–14.

Etzioni, A. and Etzioni, O. (2017). Incorporating ethics into artificial intelligence. *The Journal of Ethics*, pages 1–16.

Global, N. G. P. F. (2017). Council's Activities. `http://etikkradet.no/en/councils-activities/`. Online, Accessed: 2017-06-05.

Haines, W. (2015). Consequentialism. In *The internet encyclopaedia of philosophy*. http://www.iep.utm.edu/conseque/.

Harrod, R. (1936). *Utilitarianism Revised*. Oxford University Press.

Harsanyi, J. (1977). Rule utilitarianism and decision theory. *Erkenntnis (1975-)*, 11(1):25–53.

Investopedia (2017). Investopedia - Beta. `http://www.investopedia.com/terms/b/beta.asp`. Online, Accessed: 2017-04-13.

Keen, P. (1987). Decision support systems: The next decade. *Decision Support Systems*, 3(3):253–265.

Lovdata (2014). Retningslinjer for observasjon og utelukkelse fra Statens pensjonsfond utland. `https://lovdata.no/dokument/INS/forskrift/2014-12-18-1793?q=retningslinjer+%2B+utland`. Online, Accessed: 2017-04-13.

Matthews, K. B., Sibbald, A., and Craw, S. (1999). Implementation of a spatial decision support system for rural land use planning: integrating geographic information system and environmental models with search and optimisation algorithms. *Computers and Electronics in Agriculture*, 23(1):9 – 26.

McGowan, M. J. (2010). The rise of computerized high frequency trading: Use and controversy. *Duke Law and Technology Review*, (Issue 16).

Meredith, R. and Arnott, D. (2003). On ethics and decision support systems development. In *Proceedings of Pacific Asia Conference on Information Systems(PACIS)*, pages 1562–1575. Association for Information SystemsAIS Electronic Library (AISeL) PACIS 2003. `http://aisel.aisnet.org/pacis2003/106`.

Moor, J. H. (2006). The nature, importance, and difficulty of machine ethics. *IEEE Intelligent Systems*, 21(4):18–21.

NASDAQ (2017). Company List: NASDAQ. `http://www.nasdaq.com/screening/company-list.aspx`. Online, Accessed: 2017-06-05.

Powers, T. M. (2006). Prospects for a kantian machine. *IEEE Intelligent Systems*, 21(4):46–51.

Ross, W. and Stratton-Lake, P. (2002). *The Right and the Good*. British moral philosophers. Clarendon Press.

Scopino, G. (2015). Do Automated Trading Systems Dream of Manipulating the Price of Futures Contracts? Policing Markets for Improper Trading Practices by Algorithmic Robots. *Florida Law Review*, (67):221–293.

Shim, J. P., Warkentin, M., Courtney, J. F., Power, D. J., Sharda, R., and Carlsson, C. (2002). Past, present, and future of decision support technology. *Decision Support Systems*, 33(2):111–126.

SWFI (2017). Tracking the Activity of Sovereign Wealth Funds, Pensions and other Public Funds. `http://www.swfinstitute.org/fund-rankings//`. Online, Accessed: 2017-06-05.

Wallach, W., Allen, C., and Smit, I. (2008). Machine morality: bottom-up and top-down approaches for modelling human moral faculties. *AI Soc.*, 22(4):565–582.