



HAL
open science

Multi-Agent Based Ethical Asset Management

Nicolas Cointe, Grégory Bonnet, Olivier Boissier

► **To cite this version:**

Nicolas Cointe, Grégory Bonnet, Olivier Boissier. Multi-Agent Based Ethical Asset Management. 1st Workshop on Ethics in the Design of Intelligent Agents, Aug 2016, La Haye, Netherlands. pp.52-57. emse-01357603

HAL Id: emse-01357603

<https://hal-emse.ccsd.cnrs.fr/emse-01357603>

Submitted on 30 Aug 2016

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Multi-Agent Based Ethical Asset Management

Nicolas Cointe¹ and Grégory Bonnet² and Olivier Boissier³

Abstract. The increasing number of ethical investment funds shows how the need of ethics in asset management is growing up. In the same time, in some markets, autonomous agents are managing a larger number of financial transactions than human do. If many philosophers and economists discuss the fairness of different approaches for responsible investment, there is no strong proposition today about the implementation of autonomous agents able to take into account ethical notions in their financial decisions. This article proposes an approach to represent morals and ethics in a BDI architecture and illustrates its use in the context of ethical asset management. An analysis of a first experimentation on a simulated market is given.

1 INTRODUCTION

The increasing use of IT technologies in today financial markets is no more limited to the use of communication and automatic matching mechanisms but is invading also the decision layer where autonomous algorithms make decisions. In this paper, we are interested in asset management domain where such a transformation in the trading of assets generates several practical and ethical issues⁴. The objective and contribution of this article is use a BDI approach to embed autonomous trading agents' decisions with ethical considerations, regardless the speed or efficiency of the trading strategy.

Some people consider the use of automatic management decision as the origin of several bad effects such as market manipulations, unfair competition towards small investors and flash crashes by cascading effects. Others argue that it reduces volatility, increases transparency and stability with a lower execution cost [3]. As shown by some reports [5], ethical investment funds are even more growing and taking a significant position on the market. However, whereas the performance of such funds can be measured objectively, their "ethical" quality is more difficult to determine as it determines at least in part on the values of the observer.

Decisions by autonomous agents to whom human users delegate the power to sell/buy assets have consequences in real life [7] and as some investment funds are interested to make socially responsible and ethical trading, we are interested in the definition of mechanisms for making financial agents able to follow ethical principles, moral values and moral rules. In order to achieve this objective, we use a model of ethical judgment process proposed in [6] mapped into a BDI agent model. Such agents can decide to trade assets based on the moral and ethical preferences or values of their stakeholders.

The contributions of this article are the following: mapping of an ethical judgment process in a BDI architecture and instantiating the components of this model to the asset management domain. The paper is organized as follows. In Section 2, we introduce the asset management domain and present what are ethical considerations in such a domain. In Section 3, we present a BDI agent architecture in which the ethical judgment process presented in [6] is embedded. Thus an autonomous agent can decide on actions to execute based both on ethical principles and preferences and on moral values and rules. This BDI agent architecture is instantiated to the asset management domain. Finally, in Section 5, we offer an agent-based simulation to analyse the system's behavior.

2 ETHICS & ASSET MANAGEMENT

In this section we motivate and identify the needs of introducing ethical dimensions in the autonomous decision making supporting asset management. Firstly, we briefly present what is morals and ethics, then we present asset management domain. Then, we present the main concepts to understand ethics in such a domain.

2.1 Morals and ethics

Morals consists in a set of moral rules which describes the compliance of a given behavior with mores, values and usages of a group or a single person. These rules associate a good or bad value to some combinations of actions and contexts. They could be specific or universal, i.e. related or not to a period, a place, a community, etc. This kind of rules grounds our ability to distinguish between good and evil. Morals can be distinguished from law and legal systems in the sense that there is not explicit penalties, officials and written rules [10]. Moral rules are often supported and justified by some moral values (e.g. transparency, responsibility, ecology). Psychologists, sociologists and anthropologists almost agree that moral values are central in the evaluation of actions, people and events [15].

A set of moral rules and moral values establishes a *theory of the good* which allows humans to assess the goodness or badness of a behavior and *theories of the right* which define some criteria to recognize a fair or, at least, acceptable option. Indeed, humans commonly accept many situations where it is right and fair to satisfy needs or desires, even if it is not acceptable from a set of moral rules and values. Those theories are also respectively named *theory of values* and *theories of right conduct* [16].

Relying on some philosophers as Paul Ricoeur [14], we admit that *ethics* is a normative practical philosophical discipline of how humans should act and be toward the others. Ethics uses *ethical principles* to conciliate morals, desires and capacities of the agent. Philosophers proposed various ethical principles, such as Kant's Categorical Imperative [11] or Thomas Aquinas' Doctrine of Double Effect [12],

¹ Institut Henri Fayol, EMSE, LabHC, UMR CNRS 5516, F-42000, Saint-Etienne, France, email: nicolas.cointe@emse.fr

² Normandie Univ, UNICAEN, ENSICAEN, CNRS, GREYC, 14000 Caen, France, email: gregory.bonnet@unicaen.fr

³ Institut Henri Fayol, EMSE, LabHC, UMR CNRS 5516, F-42000, Saint-Etienne, France, email: olivier.boissier@emse.fr

⁴ <http://sevenpillarsinstitute.org/>

which are sets of rules that allow to distinguish an ethical option from a set of possible options.

Indeed, the core of ethics is the judgment. It is the final step to make a decision and it evaluates each choice, with respect to the agent's desires, morals, abilities and ethical principles. Relying on some consensual references [1] and previous work [6], *judgment* is the faculty of distinguishing the most satisfying option in a situation, regarding a set of ethical principles, for ourselves or someone else. Finally, if an agent is facing two possible choices with both good and/or bad effect, the ethical judgment allows him to make a decision in conformity with a set of ethical principles and preferences.

2.2 Asset management

The *asset management* is the art of selecting financial assets (e.g. equities, bonds, currencies, merchandises and so on) to be bought and be sold in order to manage a capital, respecting regulatory and contractual constraints, and applying an investment policy defined by the owner of the managed portfolio (a set of assets) in order to optimize his profit, considering a chosen level of risk.

The assets are commonly exchanged on a marketplace, i.e. a system designed to match bid and ask orders at the best price and the best frequency. Different types of matching methods are available, as auctions or order books, and those methods accept different types of orders, as cancellable or dynamic orders. Marketplaces are actually more than simple interfaces for buyers and sellers because they also provide a variety of functionalities:

1. to finance companies and institutions by the emission of bonds, warrants or equities;
2. to increase liquidity of the exchanges, i.e. minimizing the impact of a bid or ask order on the price;
3. to indicate the value of the assets in real time;
4. to increase the control and monitoring on the economy, by contributing to the transparency with the publication of detailed balance sheets and number of analyses.

Each asset manager composes the orders to put on the marketplace with a set of options as a possibility of cancellation, the duration of its validity, a limit price, an investment strategy and so on. To decide the best order to perform, the asset manager needs to be well informed on the state of the market, through raw data and various indicators.

2.3 Ethical dimensions of asset management

Ethical asset management, also called *responsible investment* or *social investment*, considers new information in the management decision process, as sectors, labels or any indicators on the impact of these assets and their underlying on the society. Thus, the morals of an agent (combination of moral values and rules) may be defined by an asset policy (e.g. trading nuclear-free assets or never trading in the defense sector). Moreover, the manner to trade is important too. In the last decade, the introduction of autonomous agents on the marketplaces comes with new harmful practices (e.g. layering, quote stuffing, spoofing). Therefore, the morals of an agent may also rely on transparency, honesty or avoidance of any manipulation of the market. Such policies are not about assets, but about the morality of agents's behaviors on the market.

For instance, an ethical asset manager in Islamic finance may both agree on the fact to "exclude stocks of companies that produce/distribute prohibited goods/services regarding the Shari'ah" [2] and the fact to "prefer to deal with other Islamic agents". The first

fact is part of an asset policy and the second one is part of a market policy. Those policies can be viewed as a set of moral rules. As moral rules cannot be satisfied in all contexts, ethical asset managers use ethical principles to make their decisions. By instance "Always execute the most profitable action which violate as few as possible rules" is an example of ethical principle for an ethical asset manager.

Finally, an asset manager needs to be able to judge that the asset exchanged and the modalities of the transaction are both compliant with his morals and ethics. To this end, some institutions as authorities, non-governmental organizations or journalists observe markets, funds, asset managers, and companies. From those observations, they provide evaluations that may be used by funds, companies and asset managers to make ethical decisions. For instance, the *ethiscore*⁵ is a tool proposed by some journalists to rank a set of hedge funds regarding a given set of values as ecological, political or social considerations. According with this tool, a company quoted on a market may satisfy some criteria as producing sustainable products, having a socially responsible management method and so on, depending on the values of the investors, to be considered in an ethical investment portfolio.

Knowing those concepts, our proposition consists in representing them explicitly (asset policies, market policies and evaluations) and integrate them in autonomous agents' decision process in terms of values, morals and ethics.

3 BDI AGENT ARCHITECTURE FOR ETHICAL ASSET MANAGEMENT

In this section, we first provide a global view of our architecture and then focus on the two main components for making agents able to produce ethical behaviours: goodness and rightness processes.

3.1 Global view

The agent architecture in which we introduce the necessary representations and mechanisms to have agents able to produce ethical behaviours is based on a BDI approach [13]. In this approach, the behaviour of an agent is the result of a deliberation designed to issue intentions, to bring about or to react to some world states with respect to the agent's evaluation of the situation (represented by a set \mathcal{B} of beliefs) and the agent's goals (represented by a set \mathcal{D} of desires).

To be able to produce an ethical behaviour, the basic BDI deliberation cycle must be enriched with a process to evaluate the goodness of a behaviour (represented by a *goodness process* named *GP*) and with another process to evaluate the rightness of a behaviour (represented by a *rightness process* named *RP*) resulting from the execution of actions. To this end, agents are equipped with an action knowledge base A and four other knowledge bases that define value supports VS , moral rules MR , ethical principles P and ethical preferences \succ_e . Moreover, agents are equipped with an ontology $\mathcal{O} = \mathcal{O}_v \cup \mathcal{O}_m$ of moral values \mathcal{O}_v (e.g. carefulness, ecology or transparency) and moral valuations \mathcal{O}_m (e.g. moral, quite good or immoral). The global architecture is given in Figure 1 and is issued of the judgment process proposed in [6].

In our agent architecture, each action of A is described as a pair of conditions and consequences bearing respectively on beliefs and desires. Perception *Per* and communication *Com* functions update beliefs and desires from, respectively, perception of the environment

⁵ <http://www.ethicalconsumer.org/buyersguides/money/ethicalinvestmentfunds.aspx>

and communication with other agents. From its beliefs \mathcal{B} and desires \mathcal{D} , an agent executes an *Evaluation Process EP* to assess both desirable actions, $\mathcal{A}_d \subseteq A$ (i.e. actions that allow to satisfy the consequences of the action), and executable actions, $\mathcal{A}_e \subseteq A$ (i.e. actions whose conditions are satisfied on the current beliefs about the world). The *evaluation process EP* produces desirable actions \mathcal{A}_d and executable ones \mathcal{A}_p from \mathcal{B} and \mathcal{D} . At the end of the process, we find a classical deliberation function that generates the intentions to execute given the right actions of \mathcal{A}_r .

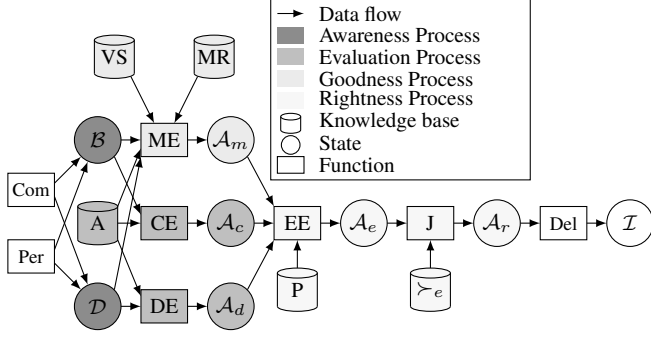


Figure 1. Ethical BDI agent architecture

3.2 Goodness process

The *goodness process GP* identifies moral actions $\mathcal{A}_m \subseteq A^6$ given the agent's beliefs \mathcal{B} and desires \mathcal{D} , the agent's actions A , the agent's value supports VS given moral values and MR moral rules knowledge base. It is defined as:

$$GP = \langle VS, MR, \mathcal{A}_m, ME \rangle$$

where ME is the moral evaluation function:

$$ME : 2^{\mathcal{D}} \times 2^{\mathcal{B}} \times 2^A \times 2^{VS} \times 2^{MR} \rightarrow 2^{\mathcal{A}_m}$$

In order to realize this goodness process, an agent uses knowledge that associates moral values to combinations of actions and situations, meaning that the execution of the actions in these situations promotes the corresponding moral values.

We represent this knowledge through value supports. A *value support* is a tuple $\langle s, v \rangle \in VS$ where $v \in \mathcal{O}_v$ is a moral value, and $s = \langle a, w \rangle$ is the support of this moral value where $a \subseteq A$, $w \subseteq \mathcal{B} \cup \mathcal{D}$. Here, the precise description of a moral value through a value support relies on the language used to represent beliefs, desires and actions. For instance, from this definition, carefulness supported by "do not buy any asset α if the volatility V is over a limit V_{limit} " may be represented by:

$$\langle \langle buy(\alpha), \{Bel(V \geq V_{limit})\} \rangle, \neg carefulness \rangle$$

where α represents any asset, $Bel(V \geq V_{limit})$ is a belief representing the context for which executing the action $buy(\alpha)$ does not support the value *carefulness*. A moral value may also be a subvalue of another more general one, i.e. all its value supports also support the more general one.

⁶ $\mathcal{A}_m \not\subseteq \mathcal{A}_d \cup \mathcal{A}_e$ because an action might be moral by itself even if it is not desired or feasible.

In addition to moral values, an agent must be able to represent and to manage moral rules. A moral rule describes the association of a moral valuation $m \in \mathcal{O}_m$ to actions or moral values in a given situation. A *moral rule* is a tuple $\langle w, o, m \rangle \in MR$ where w is a situation of the current world described by $w \subseteq \mathcal{B} \cup \mathcal{D}$ interpreted as a conjunction of beliefs and desires, $o = \langle a, v \rangle$ where $a \in A$ and $v \in \mathcal{O}_v$, and $m \in \mathcal{O}_m$ is a moral valuation that qualifies o when w holds. For instance, some rules may be represented as follows:

$$\langle Bel(sector(\alpha, medicine)), \langle buy(\alpha), _ \rangle, moral \rangle$$

$$\langle Bel(going\ down, \alpha), \langle _, carefulness \rangle, quite\ good \rangle$$

A moral rule can be more or less specific depending on the situation w or the object o . For instance "Transparency is good" is more general (having less combinations in w or o , thus applying in a larger number of situations) than "To sell an asset in a quantity superior than the available bid between the current value and the moving average minus five percent is immoral". Classically, moral theories are classified in three approaches using both moral values and moral rules as defined above, we can represent such theories.

- A *virtuous* approach uses general rules based on moral values, e.g. "Ecology is moral",
- A *deontological* approach classically considers rules concerning actions in order to describe as precisely as possible the moral behavior, e.g. "Buying an asset of an eurolabel certified company is moral"
- A *consequentialist* approach uses both general and specific rules concerning states and consequences, e.g. "Investing in an asset of an company that will practice animal testing is not moral".

3.3 Rightness Process

From the sets of possible (\mathcal{A}_p), desirable (\mathcal{A}_d) and moral actions (\mathcal{A}_m), we can introduce the *rightness process RP* aiming at assessing the rightful actions. As an ethical agent can use several *ethical principles* to conciliate these sets of actions, we consider a preference relationship between those principles. Thus, a *rightness process RP* produces rightful actions given a representation of the agent's ethics. It is defined as:

$$RP = \langle P, \succ_e, \mathcal{A}_r, EE, J \rangle$$

where P is a knowledge base of ethical principles, $\succ_e \subseteq P \times P$ an ethical preference relationship, $\mathcal{A}_r \subseteq A$ the set of rightful actions and two functions EE (evaluation of ethics) and J (judgment) such that :

$$EE : 2^{\mathcal{A}_d} \times 2^{\mathcal{A}_p} \times 2^{\mathcal{A}_m} \times 2^P \rightarrow 2^{\mathcal{E}}$$

where $\mathcal{E} = A \times P \times \{\perp, \top\}$.

$$J : 2^{\mathcal{E}} \times 2^{\succ_e} \rightarrow 2^{\mathcal{A}_r}$$

An *ethical principle* is a function which represents a philosophical theory and evaluates if it is right or wrong to execute a given action in a given situation regarding this theory. For instance "It is right to do the most desirable action which is, at least, amoral" may be a very simple principle. Formally, an *ethical principle* $p \in P$ is defined as:

$$p : 2^A \times 2^{\mathcal{B}} \times 2^{\mathcal{D}} \times 2^{MR} \times 2^V \rightarrow \{\top, \perp\}$$

The ethics evaluation function EE returns the evaluation of all desirable, feasible and moral actions (resp. \mathcal{A}_d , \mathcal{A}_p and \mathcal{A}_m) given

the set P of known ethical principles. Given a set of actions issued from EE , the judgment J selects the rightful action \mathcal{A}_r to perform, considering a set of ethical preferences (defined as a partial or total order on the ethical principles). For instance, a principle $P1 \in P$ may be “if an action is possible, desirable and motivated by a moral rule, it is right to do it” and a principle $P2 \in P$ “if an action is possible, desirable and at least not immoral, it is right to do it”. If $P1 \succ_e P2$, the agent will select a right action according with $P1$ and, if it is not feasible, a right action regarding $P2$. The right action \mathcal{A}_r is transmitted to the classic deliberation function to choose the intention I to execute.

4 AGENCY FOR ETHICAL ASSET MANAGEMENT

This section describes the experiment used to illustrate and evaluate the use of the architecture presented in the previous section. We have implemented a multi-agent system that simulates a financial market where some autonomous ethical trading agents exchange assets. This system has been implemented using the JaCaMo platform where agents are programmed using the Jason language and the market place is based on artifacts from Cartago.

4.1 Financial market modeling

We consider a marketplace where autonomous trading agents have the possibility to manage portfolio of assets and to sell or buy assets (both currencies, i.e. money, and equity securities, i.e. part of a capital stock of a company) on the market. The set of actions that an agent can execute on the market are “buy”, “sell” or “cancel” orders. They respectively correspond to the exchange of an equity for a currency, the opposite way and cancellation of a proposition of exchange if this order is not yet executed. These actions will be the ones considered in the ethical deliberation cycle of an agent. Agents can specify a limit price or can accept the current market price. Each equity is quoted in a state-of-the-art Central Limit Order Book (CLOB) [3]. A CLOB simply stores and sorts by price the set of “buy” and “sell” orders (respectively placed on bid and ask sides of the same order book) provided by the agents. When an agent put an order on the bid or ask side, the CLOB obey the following rules (see Figure 2):

- if there is no order to match with, the order is added,
- if there is an order to match with, both the incoming and the present orders are filled, and the rest of the biggest, if any, is placed in the CLOB (and may eventually match with another order).




Figure 2. Execution of a limit order added on the market

The example on the Figure 2 illustrates the addition of an ask order of thirteen assets at the price p . Before the addition, the best bid is p and the best ask is $p + 1$. The new order encounter an order on the other side during its insertion, so the biggest is splitted and the

executed parts are removed from the CLOB. At the end of the insertion, the new best bid is $p - 1$ and the new best ask is p . All these changes are perceived by the agents.

Agents get a set of beliefs describing the market and their portfolio, making them able to represent and reason on the current situation. Agents also perceive each minute a set of statistics on the activity of each asset: the volume v (the quantity of exchanged assets), two moving average prices mm and $dblmm$, respectively the average price on the last twenty minutes and on the last forty minutes, the standard deviations σ of prices, the closing prices on this period, and the up and down Bollinger bands (respectively $mm + 2\sigma$ and $mm - 2\sigma$).

The agents’ perception function provides the following beliefs from the environment:

```
indicators (Date, Marketplace, Asset, Close, Volume,
           Intensity, Mm, Dblmm, BollingerUp, BollingerDown)

onMarket (Date, Agent, Portfolio, Marketplace,
          Side, Asset, Volume, Price)

executed (Date, Agent, Portfolio, Marketplace,
          Side, Asset, Volume, Price)
```

The ethical agents are initialized with a set of beliefs about activities of the companies (e.g. EDF⁷ produces nuclear energy) and some labels about their conformity with international standards (e.g. Legrand⁸ is labeled FSC).

4.2 Ethical settings

The ethical agents know a set of organized values: for instance “environmental reporting” is considered as a subvalue of “environment”. They are declared as :

```
value ("environment" ).
subvalue ("promote_renewable_energy", "environment" ).
subvalue ("environmental_reporting", "environment" ).
subvalue ("fight_climate_change", "environment" ).
```

They also have a set of value supports as “trading assets of nuclear energy producer is not conform with the subvalue *promotion of renewable energy*”, “trading asset of an FSC-labeled company is conform with the subvalue *environmental reporting*” and “trading assets of nuclear energy producer is conform with the subvalue *fight against climate changes*”. Some examples of value supports are:

```
~valueSupport (buy (Asset, _, _, _),
               "promote_renewable_energy" ) :-
  activity (Asset, "nuclear_energy_production" ).

valueSupport (sell (Asset, _, _, _),
              "environmental_reporting" ) :-
  label (Asset, "FSC" ).
```

Agents are also equipped with moral rules stating the morality of environmental considerations. For instance, “It is moral to act in conformity with the value *environment*” is simply represented as:

```
moral_eval (X, V1, moral) :-
  valueSupport (X, V1) & subvalue (V1, "environment" ).

moral_eval (X, "environment", moral) :-
  valueSupport (X, "environment" ).
```

⁷ The French national energy producer.

⁸ A French electric infrastructure producer.

In this example, an ethical agent is now able to infer for instance that, regarding its belief, trading Legrand is moral regarding this theory of good, and that trading EDF is both moral and immoral. Finally, ethical agents are equipped with simple principles, such as “It is rightful to do a possible, not immoral and desirable action”. The implementation of this principle and some preferences is:

```
ethPrinciple("desireNR", Action):-
    possible_eval(Action, possible) &
    desire_eval(Action, desired) &
    not desire_eval(Action, undesired) &
    not moral_eval(Action, _, immoral).

prefEthics("perfectAct", "desireNR").
prefEthics("desireNR", "dutyNR").
```

4.3 Families of agents for asset management

Each agent receives a portfolio (a set of equities and currencies) at the beginning of the simulation and may exchange it on the market. Three types of agents are considered in this system: zero-intelligence, zero-ethics and ethical agents.

- *Zero-intelligence agents* are making random orders (in terms of price and volume) on the market to generate activity and simulate the "noise" of real markets. Each zero-intelligence agent is assigned to an asset. Their only desire and ethical principle are the application of this random behaviour. In this experiment, they are used to generate a realistic noisy activity on the market in order to create opportunities for the other agents.
- *Zero-ethics agents* only have a simple desirability evaluation function to speculate: if the price of the market is going up (the shortest moving mean is over the other one), they buy the asset, otherwise, they sell it. If the price goes out of the bollinger bands, these rules are inverted. This strategy is also used by the ethical agents to evaluate the desirable actions.
- *Ethical agents* implements the ethical decision process to take their decisions. An ethical agent implementing the ethical decision process without any moral value or moral rule and an single ethical principle that simply considers desirability are also ethical agents, more precisely hedonic agents. It is different from a zero-ethics agent because this agent still has all the ethical decision process and explicitly believes that its action are not moral or immoral. In this experience, ethical agents have the three following principles (by order of preferences): “It is rightful to do a possible, moral, not immoral and desirable action”, “It is rightful to do a possible, not immoral and desirable action” and “It is rightful to do a possible, moral, not immoral and not undesirable action”.

5 EXPERIMENTAL RESULTS

This section details and analyzes the results of a simulation executed with ten zero-intelligence agents per asset, eight zero-ethics agents and two ethical agents to illustrate the impact of the ethics described previously on the behavior of an ethical agent. This quantity of agents was the optimal one to generate enough opportunities in the simulations with the limited performances of a laptop. You can download this experience on the internet⁹.

At initialization, each agent receives a portfolio containing a random set of assets for a total value of 500€ more or less.

Figures 3 and 4 show the results of the experiment. Figure 3 shows all volume and price information made accessible by the market to

the agents. They concern the equities “LEGRAND”. The main data (represented by a candlestick chart) show us the evolution of the price on the market and the two moving averages mentioned in section 4.1 (the line charts in the middle of the candlestick chart) are slowly moving up and down. They are used by the desirability evaluation to detect opportunities according to the rules detailed in Section 4.3. The candlestick chart does not break often the Bollinger bands, but these breaks may happen sometimes. We observe some peaks in the volume barchart when the moving averages are crossing each other. This is due to the number of exchanges performed by the zero-ethics and ethical agents because their desirability function is triggered by this event.




Figure 3. Evolution of the asset LEGRAND. The candlestick chart represents the evolution of the price, with the moving averages in the middle and the up and down Bollinger bands on each side.

Figure 4 represents the evolution of the portfolio of an ethical agent during the experiment. It provides information on the behavior of this agent and it was chosen because it is quite representative of the portfolios of the other ethical agents. The y-axis shows the value of the portfolio and the colors depend on the assets placed in the portfolio.




Figure 4. Evolution of the portfolio of an ethical agent

Firstly, we can notice that the number of EDF equities in the portfolio never changes during the simulation. We can easily explain that by the agent’s ethical settings given in Section 4.2: it is never rightful to trade this equity because the agent thinks that EDF is a nuclear energy producer and no ethical principle provided to the agent considers an immoral action as rightful.

⁹ <https://cointe.users.greyc.fr/download/experience-EDIA2016.zip>

Secondly, we can also observe many periods where the portfolio contains the “Legrand” asset. In fact, trading this asset is the only action judged as moral due to its label. So to buy and to sell this asset is the only way to satisfy the most preferred principle and obviously, they are here the most executed actions.

Finally, we can notice different stages where the agent put in its portfolio various equities. These equities are bought or sold due to the desirability of these trades and the impossibility to execute a moral action.

6 CONCLUSION

This paper presents an ethical BDI architecture for agents in a multi-agent system. This architecture is not designed to only implement a given ethics in the decision process, but also to integrate different moral rules and values, or ethical principles as parameters of a generic architecture.

The paper also presents an experiment that illustrates, in a particular use case, how to represent and to use moral values, moral rules and ethical principles in a BDI agent in order to describe a rightful behavior. The experiment highlights how a few and simple values, moral rules and ethical principle can influence the behavior of an agent in order to incite it to prefer a set of rightful actions when they are available.

Of course, we cannot yet answer some interesting issues such as how to evaluate the cost of this ethical behavior in terms of financial performance with a real state-of-the-art trading strategy, or what is the impact of a given population of ethical agents on a market behavior. To answer those questions, we need to enrich the knowledge bases of the ethical agents with some logical models of several famous available principles in the literature (such those modeled in [4, 8, 9]) and complete the definition of moral values and rules.

Even if the morals and ethics of the agents are only used in this experiment to guide their own decisions, we intend in a future work to use them to evaluate the behavior of the other agents. Indeed, several usecases can need this kind of abilities, for instance when an authority wants to monitor the actors on a market, or when an hedge funds expresses the policy to only cooperate with other trading agents that satisfy an ethical behavior.

ACKNOWLEDGMENTS

The authors acknowledge the support of the French Agence Nationale de la Recherche (ANR) under reference ANR-13-CORD-0006. They also express their gratitude to Professor Jomi Fred Hübler for his theoretical and practical support during the design of this experiment.

REFERENCES

- [1] Ethical judgment. Free Online Psychology Dictionary, August 2015.
- [2] O.B. Ahmed, ‘Islamic equity funds: The mode of resource mobilization and placement’, *Islamic Development Bank*, (2001).
- [3] I. Aldridge, *High-frequency trading: a practical guide to algorithmic strategies and trading systems*, volume 459, John Wiley and Sons, 2009.
- [4] F. Berreby, G. Bourgne, and J.-G. Ganascia, ‘Modelling moral reasoning and ethical responsibility with logic programming’, in *Logic for Programming, Artificial Intelligence, and Reasoning*, pp. 532–548. Springer, (2015).
- [5] S. Bono, G. Bresin, F. Pezzolato, S. Ramelli, and F. Benseddik, ‘Green, social and ethical funds in europe’, Technical report, Vigeo, (2013).

- [6] N. Cointe, G. Bonnet, and O. Boissier, ‘Ethical judgment of agents’ behaviors in multi-agent systems’, in *15th International Conference on Autonomous agents and multi-agent systems*, (2016).
- [7] Directorate-General for Economic and Financial Affairs, ‘Impact of the current economic and financial crisis on potential output’, Occasional Papers 49, European Commission, (June 2009).
- [8] J.-G. Ganascia, ‘Ethical system formalization using non-monotonic logics’, in *29th Annual Conference of the Cognitive Science Society*, pp. 1013–1018, (2007).
- [9] J.-G. Ganascia, ‘Modelling ethical rules of lying with Answer Set Programming’, *Ethics and information technology*, **9**(1), 39–47, (2007).
- [10] B. Gert, ‘The definition of morality’, in *The Stanford Encyclopedia of Philosophy*, ed., Edward N. Zalta, fall edn., (2015).
- [11] R. Johnson, ‘Kant’s moral philosophy’, in *The Stanford Encyclopedia of Philosophy*, ed., Edward N. Zalta, summer edn., (2014).
- [12] A. McIntyre, ‘Doctrine of double effect’, in *The Stanford Encyclopedia of Philosophy*, ed., Edward N. Zalta, winter edn., (2014).
- [13] A.S. Rao and M.P. Georgeff, ‘BDI agents: From theory to practice’, in *Proceedings of the First International Conference on Multiagent Systems, June 12-14, 1995, San Francisco, California, USA*, eds., V.R. Lesser and L. Gasser, pp. 312–319. The MIT Press, (1995).
- [14] P. Ricoeur, *Oneself as another*, University of Chicago Press, 1995.
- [15] S.H. Schwartz, ‘Basic human values: Theory, measurement, and applications’, *Revue française de sociologie*, **47**(4), 249–288, (2006).
- [16] M. Timmons, *Moral theory: an introduction*, Rowman & Littlefield Publishers, 2012.