

7.1 Description of the Full-Proposal Research Project

1. Project Abstract (max. 1 page)

The proposed project aims to apply the paradigm of agent-based computational economics (Tesfatsion and Judd, 2006) to model complex reasoning processes in a market setting. Basically we assume a market in which the main product is information and complex analysis, on an issue that does not provide immediate feedback from the objective world. An example would be long-term financial advice; other examples would be analysis of world energy resources or the impact of climate change. The absence of immediate feedback means that the quality of the analysis of the various information providers (which we will simply refer to as “consultants”, even though information providers can also be institutions) cannot be assessed by simply observing who is “right” the most often times. Instead, one needs to assess the quality of the analysis itself, by carefully examining whether all factors have been properly taken into account. Such an assessment can be troublesome from the perspective of the client, who is after all not an expert on the topic himself (otherwise there would be no reason to avail of the services of the consultant). Nevertheless, there are a few heuristics available to the clients to make such an assessment. These heuristics are imperfect, however, and it can be that in some situations the clients will make a wrong assessment of the expertise of a particular consultant. On the collective level, this can lead to consultants trying to optimize not their real expertise (which would cost significant resources to do so) but merely their perceived expertise. We assume that the reputation of the consultants is subject to a reputation system. The situation here is similar to that of, for instance, eBay, where the buyer provides an assessment of the performance of the seller. However, where on eBay it is usually immediately clear to the buyer whether the product fulfills the advertised specification, this is much more difficult to assess if the product being sold consists of information and analysis, whose objective truth cannot immediately be assessed. Examples would be claims and analysis about issues like the viability of a particular long-term investment strategy, the magnitude of world energy resources or the economic impact of climate change of a particular region. One cannot directly assess whether positions on these issues are *true*, the best one can do is to assess whether they are *well-informed*. This (imperfect) assessment is then the basis of the feedback to the reputation system.

We are interested in examining under which conditions the information providers (“consultants”) have sufficient incentives to provide good quality analysis to their clients. The preliminary results of our software simulator (Staab and Caminada, 2010) as well as other research (Mathis et al, 2009) indicate that these incentives are not always strong enough to rule out providing low quality information. If such becomes the pervasive strategy of the consultants, there are consequences regarding the informedness not only of individual information consumers (“clients”) but also for the system as a whole.

An example of the failure of collective reasoning in a market setting can be found in the recent credit crisis. Until roughly three years ago, even the possibility of such the recent events would not have been considered in any serious way by most market participants. How is it possible that so many reputable experts (including news media and government agencies) were wrong? Was this because (1) the individual market participants were plain irrational, (2) the relevant information that could have predicted the crisis was simply not available or (3) because of inherent flaws in market-based the collective reasoning process? Without downplaying the results in the field of behavioral finance, we think that investor psychology alone is insufficient to explain the magnitude of the crisis. Similarly, it can be observed that information about what is now perceived to be one of the sources of the crisis (like mortgages being sold to people who could not afford home ownership) was indeed available. We will therefore focus on the third

factor: the inherent flaws in market-based collective reasoning processes. How exactly to model these flaws is the topic of the currently proposed project.¹

2. Accordance with FNR CORE thematic programme description (max 0.5 page)

The LAAMI project is situated within the domain of **Development and Performance of the Financial System**, in the thematic research priority **Innovation in Services (IS)** of the CORE 2010 call. The basic theme of the proposed project is how to explain market imperfections from the perspective of flaws in collective reasoning. For this, we look at how providers of information and analysis are assessed in market settings. In particular, we examine markets in which the main product is information and analysis (like for instance financial advisers, rating agencies, etc). That the collective opinions that are generated by these markets have a thorough impact on other markets (like the stock market) is beyond discussion. Our aim is to help explain the behavior of the latter type of markets by examining the imperfections in the former type of markets.

3. Definition of Proposed Project (max. 10 pages for 3.1 – 3.3)

3.1 Current state of the art including your relevant previous work

The Efficient Market Hypothesis states that markets are rational in the sense that all publicly available information has already been taken into account regarding price formation. In ideal circumstances, this would lead to a durable equilibrium that only gradually slides as new information comes in. However, past experiences regarding bubbles, crashes and even the instability of the entire financial system cast doubt on the concept of the rational market.

Explaining market irrationality poses a problem for classical economic theories. Some approaches seek the source of market irrationality in the irrationality of the individual market participants, as is done in the field of behavioural finance, where the aim is to explain market behaviour by taking into account various psychological factors that can cause market participants to deviate from what is considered as the most rational behaviour. Examples of such factors are *mental conservatism* (Phillips and Edwards, 1966), *loss aversion* (Kahneman and Tversky, 1979 and 1991) and even factors related to weather (Saunders, 1993, Hirshleifer and Shumway 2003) and season (Bouman and Jacobsen 2002).

Without downplaying the results in the field of behavioural finance, the current project seeks the source of inefficient markets not so much in investor psychology, but in fundamental problems of bounded rationality and strategic self-interested behaviour of those who supply the market with analysis, opinion and advice.

An example of a market party whose main role is to provide analysis and opinion is that of a credit rating agency. These play a pivotal role in market perception of credit worthiness, and it has been observed that any additional information that is provided by these has a significant impact on the credit markets, even in case this additional information is simply the result of more refined ratings becoming available (Tang, 2009). Nevertheless, the role of rating agencies

¹ It should be mentioned that when we discuss the "flaws" in the collective reasoning process, we take the collective perspective and not the individual perspective. What might be optimal behavior of individual participants aiming at maximizing their own utility might not yield a socially desirable outcome at the collective level. In cases where the result at the collective level is clearly suboptimal, despite the fact that the individual market participants are (bounded) rational, we say that the process that yields the collective outcome is flawed.

has received heavy criticism, even before the recent credit crisis. Partnoy, for instance, argues against the dominant “reputational capital” view of credit rating agencies, in which these prosper “based on their ability to accumulate and retain reputational capital (i.e. good reputations) by providing valuable information about the bonds they rate” (Partnoy, 1999). Instead, Partnoy argues that the main role of a rating agency is merely to allow investors to satisfy certain regulatory requirements. Sinclair argues in a similar way that markets and government actors take account of rating agencies not because they are right, but because they are thought to be an authoritative source of judgements (Sinclair, 2005). Ferri et al. examine the role of rating agencies during the Asian crisis of 1997. They observe that after having failed to predict the emergence of the crisis, rating agencies responded by becoming excessively conservative and giving lower ratings than was actually deserved, therefore contributing to amplifying the crisis (Ferri et al, 1999). These findings are in line with those of Amadou (2004) who observes that “ratings do not predict currency crisis and are instead downgraded ex-post.” Mathis et al. claim that considerations about their reputation might not be sufficient to dissuade rating agencies from giving a too positive rating to certain structured products. They observe that “the reputation argument only works when a sufficiently large fraction of the credit rating agency income comes from other sources than rating complex products. By contrast, when rating complex products becomes a major source of income for the credit rating agency, we show that it is always too lax with a positive probability and inflates ratings with a probability one when its reputation is good enough” (Mathis et al, 2009). Apart from this, there is also the issue of ratings competition and ratings shopping (Skreta and Veldkamp, 2009) which is sometimes even raised by those having worked in the ratings industry themselves (Fons, 2008).

Despite the fairly horrible performance of the credit ratings agencies, as witnessed in the recent credit crisis, it is striking to see that the ratings industry has continued to operate, while still being relied on both by regulators and by the public. (Spatt, 2009). In fact, “causal empiricism suggests that even after the dramatic weakening the credit rating agency reputations that investors continue to view somewhat seriously the information reflected in ratings changes” (Spatt, 2009).

Although the role of credit rating agencies has received a lot of attention, the problem of institutions providing low quality analysis and opinion to the market is not limited to credit rating agencies alone. Another interesting example can be found in the International Energy Agency (IEA), which supplies markets and OECD governments with information and analysis regarding medium-term, and long-term energy resources and demand. For its analysis of world oil supplies, the IEA relies on official data, as is reported by OPEC. However, as is pointed out by Legett (2006) and Campbell (2009) that these data are likely to be distorted. The reason for this is that in the 1980's the OPEC production quota system came into force and the production quota of the individual members became dependent on the size of the reported reserves. This provided a clear incentive to over report, especially since there was no independent auditing of reported reserves. Yet, despite of these shortcomings, these data became used by the IEA as the basis for their estimates on how long world oil reserves would last.

The example of the International Energy Agency makes clear that in order to evaluate the trustworthiness of an information source; one cannot always rely on past empirical observations regarding the accuracy of its predictions. Such a strategy might work when observations are frequent (like is for instance the case when evaluating the performance of a weather forecaster) but would be hard to maintain when it comes to long-term issues, like the duration of world energy resources, the amount of future climate change, or even the viability of the investment strategy of a particular pension fund. Such instances, where short-term feedback is unavailable, will be referred to as *non-immediate feedback* issues.

When it comes to non-immediate feedback issues, one can assess the trustworthiness of an information provider not so much by the correctness of its (previous) analysis, but by looking at

the quality of the analysis itself, in order to evaluate whether for instance all relevant factors have been properly taken into account, something that for instance was not the case in the IEA example discussed earlier. Carrying out such an assessment, however, can require significant resources since in order to do so, one would often need to have considerable expertise in order to determine whether the analysis of the information provider makes sense. In the remainder of this document, we refer to issues of non-trivial reasoning whose execution as well as verification costs significant efforts and/or resources as *complex* issues.

In order to understand the problematic nature of complex non-immediate feedback issues, consider the example of risk assessment in financial markets. Especially in the derivatives market, the methodology has become so complex that only an elite group of specialized mathematicians is able to master it. To become member of this group requires significant effort and resources, and once this investment has been made it might not in one's best interest to point limitations of the newly acquired skills. It should therefore not have come as a surprise when these limitations were detected not in advance but when finally feedback did arrive, that is, during the peak of the recent crisis, when for many institutions that relied on this methodology, it was already too late.

We consider problems like the quality of credit ratings, the viability of world energy forecasts, and the perceived expertise of risk assessment specialists not as isolated incidents, but as symptoms of fundamental problems with respect to collective epistemic and market informedness on complex non-immediate feedback issues. These are the deeper problems that are the topic of the current project proposal. In order to study them, we propose to apply techniques from the domain of knowledge representation and reasoning, which will be discussed in the following paragraphs.

One of the currently prominent techniques for knowledge representation and reasoning can be found in the field of formal argumentation (Rahwan and Simari, 2009). The basic idea is to regard defeasible entailment in terms of abstract or structured entities called *arguments* which interact with each other in either a positive (support) or negative (attack) way. Although the origin of formal argumentation theory can be traced back to the early nineties with the work of Simari and Loui (1993) and Vreeswijk (1993), it has gained a major impulse with the landmark paper of Dung (1995) which introduced the concept of an *argumentation framework* and shows how some earlier approaches of non-monotonic reasoning, like default logic (Reiter, 1980), logic programming (Gelfond and Lifschitz, 1988, 1991) and Pollock's OSCAR system (Pollock, 1992), can be seen as special instances of argumentation theory.

The work of Dung has led to significant follow-up research, which includes issues like alternative semantics (Caminada, 2006, Dung, Mancarella and Toni 2007), structured arguments (Caminada and Amgoud, 2007, Prakken, 2009), defeasible priorities and higher order argumentation (Prakken and Sartor, 1997, Modgil 2009) assumption based argumentation (Bondarenko, Dung, Kowalski and Toni, 1997), argumentation schemes (Prakken 2005a). The connection between argumentation theory and formal models of dialogue has been studied by Prakken (2005b), who has also examined various procedural aspects of argumentation theory. Experimental psychological research on the extent to which formal argumentation theory coincides with human reasoning has been carried out by Rahwan et al. (2010). Proof procedures for various approaches of abstract argumentation have been stated by (Prakken and Sartor, 1997, Vreeswijk and Prakken, 2000, Modgil and Caminada, 2009, Caminada and Wu 2009) and related issues of computational complexity have been studied by Dunne and Wooldridge (2009) and Dunne and Caminada (2008).

Purposes for which argumentation theory has been applied include non-monotonic inference (Caminada and Amgoud, 2007), inter-agent communication (McBurney and Parsons, 2009, Prakken 2009), agent-based decision making (Amgoud, 2009) and reasoning in the context of

the semantic web (Rahwan et al, 2007, Rahwan and Reed, 2009). Overall, it is fair to say that over the past years, argumentation theory has become a mainstream approach for various applications in the field of knowledge representation and reasoning.

Our proposal is to apply argumentation theory in the context of agent-based computational economics (Tsfatsion, 2002, Tsfatsion and Judd, 2006). Agent-based computational economics (ACE) can be described as “the computational study of economics modelled as evolving systems of autonomous interacting agents” (Tsfatsion, 2001). The idea is that “economic worlds can be computationally constructed that are populated with heterogeneous agents who determine their interactions with other agents and with the environment on the basis of internalized social norms, internal behavioural rules, and data acquired on the basis of experience” (Tsfatsion, 2001). Also, one can allow for agent heterogeneity by having the market be populated by agents with different characteristics, such as differences in access to and interpretation of available information, different expectations, or different trading strategies (Alfarano et al, 2005). “The traders interact, for example, by exchanging information or they trade by imitating the behaviour of others. The market possesses, then, an endogenous dynamics, and the strict one-to-one relationship with the news arrival process [as assumed by EMH] does not hold any longer” (Alfarano et al, 2005).

The paradigm of agent-based computational economics has been applied to study a range of different issues. The phenomenon of *herding*, for instance, has been studied by Lux and Marchesi (2000) and by Kirman and Teyssi re (2002). The issue of learning and adaptive behaviour has been studied by Lin (2005), who models an agent-based system in the context of inflation forecasting. In this work, each agent gathers information from his contacts and forms an inflation forecast based on this information, using belief generation procedures. “When the actual inflation is realized, an agent is in a position to learn, i.e. adjust his own network strategy and beliefs” (Lin, 2005). Hegselmann and Krause (2002, 2005) study the issue of opinion dynamics, when agents have only a limited confidence in their own opinion and are therefore willing to adjust it when becoming aware of the opinions of other agents. In the particular context of financial markets, Colucci and Valori (2008) examine market stability in the presence of heterogeneous agents. They observe situations in which heterogeneity provides stability and situations in which it does not. Chiarella et al (2008) study markets for risky securities where trading agents have different strategies (fundamentalist, chartist or noise trader), but are also aware of the beliefs of others. Similar work, in the context of stock option markets, has been done by Ecca et al (2008).

To give an impression of the particular type of research we are interested in, we now describe recent, though preliminary, work of Staab and Caminada (2009) in the field of computational economics and social epistemology. The idea is to consider a multi-agent system consisting of two types of agents, information providers (called “consultants”) and information consumers (called “clients”). We assume that there is a particular (single) topic that is important to the clients. Although the precise topic is left abstract, it is assumed to be a complex non-immediate feedback issue, as described earlier in this document. For simplicity, we assume that there is no feed-back loop between prediction and outcome. That is, the outcome is determined purely by external factors and does not depend on the agents' beliefs. Since the topic is important to the clients, they are willing to pay the consultants for information and analysis, which the consultants provide under a non-disclosure agreement that rules out sharing this information among the clients. The consultants' analysis comes in the form of a set of arguments that are relevant to the main issue. Standard argumentation theory is then applied in order to draw a conclusion on the main issue, given the set of arguments made available by the consultant.

The system operates in a number of rounds. In each round, new information (arguments) becomes available, and it is up to the consultants to decide whether or not to procure these arguments. However, procurement is assumed to cost resources, since the consultants have to

invest time and/or money for the necessary research to obtain them. For simplicity, we assume a fixed and purely financial cost for a consultant procuring an argument, and we assume a fixed number of arguments becoming available each round.

The community of consultants is assumed to be heterogeneous in the sense that there exist two types of consultants: *well-informed* and *ill-informed*. The well-informed consultants try to give the best possible quality of analysis to their clients, basically by procuring all relevant arguments as soon as they come available. The ill-informed consultants, on the other hand, have a more refined strategy which will be described in more details later.

Although clients are not allowed to share the information provided by the consultants among each other, they do however share their (positive or negative) experiences with each individual consultant using a *reputation system* (Resnick et al, 2000). If a client finds out it knows at least as much as the consultant (has a superset of relevant arguments) it reports a bad experience. However, if the consultant was able to provide additional relevant information (arguments) the client reports a positive experience.

In every round, the clients decide which consultant to hire, in a weighted probabilistic way based on the consultant's price (which is a negative factor) and reputation (which is a positive factor). Where the well-informed consultants procure all relevant information (arguments) as soon as it becomes available, the ill-informed consultants follow the more sophisticated strategy of trying to procure information only if they meet a client that is at least as informed as they are. That is, the ill-informed consultants' main consideration is the *perception* of being informed, rather than actually *being* informed. As long as this perception is warranted, the ill-informed consultant is willing to continue to operate based on outdated or shallow knowledge. Ill-informed consultants run a higher risk of causing bad experiences with clients (which then lowers their reputation) but are able to offer a lower price due to their decreased costs of information procurement. At the other hand, even well-informed consultants run a (small) risk of causing bad experiences, in case they meet a client whose knowledge is already state of the art

The aim of the simulator is to examine the circumstances under which the well-informed or the ill-informed consultants are economically the most successful. That is, as a consultant, under which circumstances does it pay off to use a strategy of well-informedness or of ill-informedness? In particular, three factors are being considered: (1) the speed with which new information (modelled as arguments) becomes available, (2) the relative importance which clients attribute to price (negative factor) in relation to reputation (positive factor) and (3) the fraction of well-informed consultants in relation to ill-informed consultants. One of the limitations of the current simulator is that this fraction is assumed to be fixed. That is, if a consultant applies the well-informed (or ill-informed) strategy, then it keeps on applying this strategy throughout the entire run of the simulation.

In general, the current simulator allows one to observe that the ill-informed strategy tends to economically outperform the well-informed strategy in situations where (1) the speed at which new information arrives is high, making procurement of all available information relatively costly, (2) the clients attribute a relatively high importance to price in relation to reputation, and (3) there exists a critical mass of ill-informed consultants. Although points (1) and (2) are relatively straightforward, point (3) requires some additional explanation. When the percentage of ill-informed consultants is high, it becomes hard for the clients to obtain the kind of information needed to detect that the ill-informedness of a particular consultant. Therefore, it will be less likely for an ill-informed consultant to be caught on this, meaning that the reputation will suffer relatively little damage from applying the ill-informed strategy. This together with the fact that ill-informed consultants can offer their services for lower prices (due to the decreased costs of information procurement) leads to their economic outperformance at the expense of the well-informed consultants.

In the current treatment, we have only provided a brief overlook of the current simulator. More details, including various charts of economic performance of the well-informed and ill-informed consultants, under various conditions, can be found in (Staab and Caminada 2010), which builds on earlier work described in (Staab and Caminada, 2009a, 2009b).

One last point of attention is whether one assumes a forecast to be able to influence the object of the forecast. That is, does one want to take into account self-fulfilling or self-denying prophecies?

When designing and implementing the prototype simulator (which the current project aims to re-implement using a more extended and realistic model) it was assumed that there would be no feedback loops between forecasts and the object of prediction. An example of an economic highly relevant issue where such feedback loops are absent can be found in natural resources analysis; the amount of remaining recoverable oil or rare earth elements does not depend on one's knowledge about it. However, there also exist domains where a forecast does have effects on its object, that is, where the forecast has self-fulfilling or self-denying properties. If time and resources allow, we would like to implement a simulation model that explicitly takes into account self-fulfilling and self-denying prophecies, which would have the advantage that the simulator becomes relevant for a greater class of problem domains.

3.2 Project objectives and contribution to knowledge development

Our aim is to construct a simulator in the area of agent-based computational economics that can produce insight in the dynamics of knowledge generation in a market setting. We assume the presence of different strategies at the side of the information providers ("consultants"). Some of these strategies might be aimed at spending the necessary resources in order to provide a high-quality analysis, whereas other strategies might be aimed at economizing on these resources and provide an analysis that is aimed merely to be *perceived* as being of high (or at least sufficient) quality. In particular, we would like to implement a simulator that overcomes the limitations of the currently implemented version in the following ways:

1. agents have access to more refined strategies and are able to change strategy during the course of the simulator, if they have reason to believe that doing so would be beneficial to them
2. agents can perceive the validity of pieces of information with various levels of confidence, influenced by factors of coherence and the opinion of other agents
3. the environment can give various levels of feedback of the actual truthfulness of the information traded in the simulation
4. the issue at stake can be of various levels of complexity, and evaluating the correctness of the arguments, as well as the precise way in which they interact, might require resources itself

We now treat each of these points in more detail. In the current version of the simulator, the strategy of a consultant is fixed. Nevertheless we observe that the ill-informed strategy outperforms the well-informed strategy if the percentage of ill-informed consultants has passed a certain threshold. If consultant agents are able to switch strategies, based on what provides them with the best economic performance, our hypothesis is that there will be two equilibria: one in which all consultant agents adopt the well-informed strategy, and one in which all consultant agents adopt the ill-informed strategy. In both cases, there is a critical mass of consultant agents that adopt a certain strategy, so it is in each consultant agent's best interest also to adopt it.

Our aim is then to experiment with how stable these equilibriums are. For instance, if the system is in the ill-informed equilibrium, then one may ask how strong the client agent's preference of reputation over price has to improve in order to make it worthwhile for the consultant agents to break out of the equilibrium and adopt the well-informed strategy instead. That is, we would like to examine how strong the force from specific factors (like price/reputation preference) has to be in order to go from one equilibrium to the other.

To some extent, the phenomenon that we would like to study is related to *information asymmetry*. While each consultant agent knows which strategy it is applying, the clients sometimes have difficulty detecting this. However, the (hypothesized) existence of two equilibriums (one well-informed and one ill-informed) means that the situation differs from, for instance, the traditional example of the second-hand car market where just one equilibrium exists: the lemon one. Furthermore, it is not a-priori clear whether clients can detect in which of the two equilibriums the system finds itself. One of the research questions is therefore to examine to which extent the system provides subtle hints to the clients to help them determining this.

As for the second point, we expect that once factors of coherence and peer opinion are factored into the system, there will be observable occurrences of *herding*, especially when relatively few agents are actually assessing the validity and strength of the arguments (in the current simulator it is simply assumed that each argument is valid). Collective reasoning of acceptance or rejection of arguments has previously been examined by Caminada and Pigozzi (2009), and the proposed project will take these results into account.

As for the third point, it can be observed that many problems discussed in this paper specifically occur in the absence of immediate feedback from the objective world, that is, they occur in complex non-immediate feedback domains. If frequent and directly observable feedback is available, then even for issues involving complex reasoning, assessing which consultant to trust is simply a matter of examining who is right the most often. One does not have to have expertise on meteorology in order to assess which weather forecasting agency is best. However, if this directly observable feedback is absent, then one has to assess the internal structure of the available analysis (the forecast). As was argued before, our hypothesis is that imperfections in this assessment could lead to two equilibriums, one well-informed and one ill-informed. The current simulator assumes no objective feedback at all, but it would be interesting to adjust the simulator, allowing for various levels of feedback. It could be possible that, if feedback comes in only in the really long term, the system has already entered the ill-informed equilibrium. Great is then the surprise when the objective feedback finally comes in, and the existing consensus turns out to be fatally wrong, even though it was supported by many information suppliers (consultants) that were at least perceived as reputable.

As for the fourth point, the level of complexity has to some extent been modeled in the current simulator by the number of arguments becoming available in each round. However, one could also imagine other ways of modeling complexity, possibly in the internal structure of the arguments, using approaches like (Caminada and Amgoud, 2007) or (Prakken, 2010). In particular, we are interested in modeling complexity in such a way that complex analysis requires more efforts to verify.

Overall, what we are interested in is coming up with a credible model that can explain the flaws of collective reasoning in a market setting. In particular, we would like to investigate the conditions under which these flaws can occur. That is, which are the factors under which there is a risk for the market to be ill-informed. We believe that catastrophic failures of collective reasoning, as were witnessed with the recent credit crisis, require new methods of analysis,

which aim at modeling and simulating in which way and under which circumstances these failures can occur.

3.3 Methods and approach

Like mentioned earlier, our approach is to use an agent-based simulator to examine possible flaws of collective reasoning in a market setting. Thus, our method can be seen as an instance of agent-based computational economics (Tesfatsion, 2001). However, our main interest is not so much in traditional issues like price formation or profitability, but in the resulting level of informedness in the overall system.

Our approach is to use argumentation theory as one of the basic components of the system. An advantage of using argumentation theory is that it can model positive as well as negative factors, due to the defeasible nature of its entailment. Although defeasibility could also be implemented using theories like default logic (Reiter, 1980) or logic programming (Gelfond and Lifschitz, 1988, 1991) argumentation theory offers the additional advantage that it can facilitate different levels of abstraction, by allowing arguments to be seen as purely abstract entities, or as having an internal structure that resembles individual reasoning steps. Arguments with internal structure also allows argumentation theory to implement the earlier mentioned formalisms of default reasoning and logic programming, as well as other formalisms such as (Pollock, 1992). Hence, argumentation theory can be seen as a more general approach than some of the earlier work regarding defeasible (non-monotonic) entailment. Furthermore, recent work has brought several types of additional functionality, like defeasible priorities (Modgil, 2009) and accrual (Lucerno et al, 2009). Overall, we consider argumentation theory to be sufficiently mature and expressively rich in order to serve as one of the basic components of the formal model underlying the proposed simulator.

As for the development platform, the choice would be either to apply a general software development platform (such as Java, in which the current version of the simulator has been developed) or to apply a platform that is specifically meant for implementing multi-agent systems (such as 3APL (Dastani et al, 2003) or the ART testbed (Fullam and Barber, 2006)). The choice of the particular development platform will be made in an early stage of the project, after the general architecture of the simulator has been specified. The various platforms will then be evaluated to which extent they are able to facilitate the implementation of the software architecture.

The project will start, however, with an extensive literature review regarding the role of information and common belief with respect to the recent credit crisis. The aim is to examine what are perceived to be the factors that caused fundamentally wrong opinions and beliefs in the market, resulting in the near collapse of major parts of the financial system. Mathis et al, (2009), for instance, claim that rating agencies had a clear incentive to be over positive when rating certain complex products, an incentive that was not neutralized by considerations of reputation. The current simulator does not implement these kinds of incentives (although its outcome suggests that even in their absence, the system can still enter an ill-informed state) and it would be interesting to examine how these could be formalized in the underlying formal model of the proposed simulator.

4. Project Plan and Work Package Description

4.1 Project plan (1 page)

The proposed project consists of four work packages. The first three packages contain the actual research that needs to be performed, whereas the fourth work package is aimed at publishing the overall results of the project. The work packages are non-overlapping. This is

necessary, because each successive work package builds on the results of the previous work package. Work package 1 aims to provide a comprehensive overview of the current state of the art regarding formal models of information quality and dishonesty. This outcome is then fed into work package 2, which aims at providing an overall formal model, taking into account factors that have previously been identified in the literature, as well as factors that are identified in addition. This formal model will then be implemented in a software simulator that is to be developed in work package 3. Developing this simulator is done in two separate tasks. The first task consists of design and selection of a development platform. The second task consists of the actual implementation and testing. Implementation and testing is to be performed by one of the students (Bachelor or Masters) at the Computer Science and Communications Research Group of the University of Luxembourg, whereas the design and selection of the development platform will be done by a postdoc to be employed on the project. Once the simulator will be up and running, experimenting with different parameters will be relatively straightforward. This will be done in work package 4, which also aims to disseminate the results of the project as a whole.

There will be one person (Martin Caminada) to be full-time employed on the project, throughout its entire lifetime. Caminada has also been involved in developing a previous prototype simulator, that was discussed earlier in this document. Although the previous prototype can be characterized as quite primitive and lacking various forms of functionality, it can still be seen as a proof of concept, namely that it is possible to apply techniques in knowledge representation and reasoning to simulate processes reasoning and information dissemination in a market setting, and that using this simulation one is able to gain insight regarding the conditions under which these processes do or do not lead to well-founded common beliefs. The currently proposed project will draw upon the experiences gained by developing the previous prototype.

Two other researchers to be involved in the project are Björn Ottersten and Tibor Neugebauer. In his role as director of the Interdisciplinary Center for Security and Trust, Ottersten has wide experience with managing interdisciplinary research projects. His role will be mainly monitoring and project management. This includes tasks such as helping to select the student to be involved in the software development in work package 3. The connection with economic research will be warranted by Neugebauer, who is director of research at the Luxembourg School of Finance.

Apart from the formal participants, we also plan to profit from the expertise available at the Individual and Collective Reasoning group (ICR) of the University of Luxembourg, in particular from that of ICR members Leon van der Torre and Emil Weydert.

4.2 Work packages and tasks (1 page/work package suggested)

Work package number	1	
Work package title	Review Formal Models of Information Quality and Dishonesty	
Work package leader	Start date	End date
Martin Caminada	Month 1	Month 4
Objective		
The aim of work package 1 is to provide an overview of the state of the art regarding formal models of quality of complex analysis (possibly, but not exclusively, using formal logic) and possibilities for dishonesty. In particular, we are interested in models that can explain strategies of how an agent can pretend to have a greater expertise on a particular topic than is actually the case, as well as in work that explains under which circumstances considerations of reputation are insufficient to neutralize incentives for dishonesty (like for instance has been studied in the work of Mathis et al (2009)). Other relevant work includes that of Rahwan, Larson and Tohmé (2009) who study incentives for dishonesty using abstract argumentation theory, as well as work		

<p>that is currently being performed by Pigozzi and others. Our aim is to come up with a comprehensive survey of the state of the art regarding dishonesty and strategic behavior of information providers, including our own previous work.</p> <p>The resulting document then allows us to analyze which parts of existing theories we would like to incorporate in our own framework, and which aspects have not yet been treated in any previous work and therefore need to be formalized and implemented from scratch.</p>			
Tasks			
<p>The main task here consists of performing a literature review, both in the field of (computational) economics and in the field of knowledge representation and reasoning. We aim to explicitly ask other relevant researchers in these fields for input, in order to make sure that the final document is as comprehensive as possible.</p>			
Interdependence with other work packages			
<p>The results of work package 1 will be applied in work package 2, where an overall formal model is to be developed.</p>			
Deliverables and Milestones			
<p>The main deliverable is a technical report containing the overview. Depending on the precise result, a shortened version of this report could be considered for submission to the ECSQARU 2011 conference (European Conferences on Symbolic and Quantitative Approaches to Reasoning with Uncertainty).</p>			
Resources			
Name of researcher	Partner	Qualif. Level	Person Months
Martin Caminada	SnT	Postdoc	4
Björn Ottersten	SnT	Professor	1
Tibor Neugebauer	LSF	Professor	1

Work package number	2		
Work package title	Development of a Formal Model		
Work package leader	Start date	End date	
Martin Caminada	Month 5	Month 12	
Objective			
<p>The aim of work package 2 is to construct a formal model of reasoning, in a market setting, about complex non-immediate feedback issues. Some of the aspects that need to be addressed in this framework are how to model complex reasoning, how information consumers can assess the quality of the analysis provided to them, and what are the possibilities of the information providers to exploit inherent weaknesses in this assessment. Apart from this, we also aim to model the relevant aspects identified in work package 1. The overall aim is to provide a model that takes into account both existing and additionally identified aspects, and can serve as the basis for implementation in a software simulator.</p>			
Tasks			
<p>The main task here consists of providing a formal model (possibly using formal logic and/or formal argumentation theory) that can serve as a basis for the implementation of a software simulator.</p>			
Interdependence with other work packages			
<p>Work package 2 builds on the review carried out in work package 1, and will serve as the basis for the software simulator to be developed in work package 3.</p>			
Deliverables and Milestones			
<p>The main deliverable is a technical report containing the formal model. A shortened version of this technical report is to be submitted to a relevant workshop, such as the yearly <i>Multi-Agent-Based Simulation</i> workshop. The aim of this submission is to receive early stage feedback of the research community, which could then still be implemented in work package 3.</p>			
Resources			

Name of researcher	Partner	Qualif. Level	Person Months
Martin Caminada	SnT	Postdoc	8
Björn Ottersten	SnT	Professor	1
Tibor Neugebauer	LSF	Professor	1

Work package number	3		
Work package title	Design and Implementation of Software Simulator		
Work package leader	Start date	End date	
Martin Caminada	Month 13	Month 21	
Objective			
<p>The aim of work package 3 is to be able to study the simulated behavior of a market in which the main product is information and in which there are possible incentives for dishonesty. We want to understand under which conditions the individual and collective beliefs of the simulated market participants are in line with what would have been the beliefs if dishonesty would not play a role. This issue is relevant from the perspective of <i>mechanism design</i> (Rosenschein and Zlotkin, 1994), which is the field in which it is studied how to design a market which satisfies some global desirable properties (in this case, a low impact of dishonesty). On the other hand, the simulator should also be able to identify those settings in which the market results in providing collective beliefs that fundamentally deviate from what seems to be the most rational, given all information that is available in the system. The preliminary simulator that we have developed before provides indications that such situations do exist, and the aim of the newly developed simulator in the current work package is to precisely pinpoint these situations.</p>			
Tasks			
<p>Work package 3 is mainly involved with software development. The definition of requirements is provided by the formal model developed in work package 2. Based on this, a design of the simulator will be developed, which serves as a basis for implementation. One of the things that needs to be examined is whether implementation will take place using a general software development platform (like Java or Netlogo) or using a specialized agent-based development platform. Part of work package 3 will be a quick evaluation of the different possibilities for choosing a software development platform. Once the design is ready and the actual development platform has been chosen, the actual programming task of developing the simulator is to be proposed as a bachelor or master thesis project to one of the students at the Computer Science and Communications Research Unit (CSC) of the University of Luxembourg.</p>			
Interdependence with other work packages			
<p>Work package 3 uses the results of work package 2 (development of a formal model) as the requirements analysis. The results of work package 3 will be among the main results of the project as a whole, and will be the basis of publications and other dissemination activities in work package 4.</p>			
Deliverables and Milestones			
<p>Work package 3 has two milestones. The first one, after 3 months, will be the design of the system and the selection of the software development platform. The second milestone will be the finishing of the implemented simulator, which is scheduled to take place after 9 months.</p>			
Resources			
Name of researcher	Partner	Qualif. Level	Person Months
Martin Caminada	SnT	Postdoc	9
Björn Ottersten	SnT	Professor	1
Tibor Neugebauer	LSF	Professor	1
(student programmer)	SnT	Student	6

Work package number	4
Work package title	Experimentation and Dissemination

Work package leader	Start date	End date	
Tibor Neugebauer	Month 19	Month 24	
Objective			
Where the aim of work package 4 was the development of the simulator itself, the aim of work package 5 is to disseminate (publish) the results of the simulator, as well as providing a general overview of the simulator itself. These results are interesting for two research communities. Since the topic is individual and collective reasoning, it seems reasonable to aim to disseminate the results at conferences like KR (Knowledge Representation) or AAMAS (Autonomous Agents and Multi-Agent Systems). The fact that this reasoning takes place essentially under market conditions, makes the results particularly interesting for those working in computational economics, and one of the main aims of work package 4 is to write a paper that is to be submitted to the <i>Computational Economics</i> journal.			
Tasks			
The main task is to write a number of papers to be submitted both in the field of knowledge representation and reasoning, and in the field of computational economics.			
Interdependence with other work packages			
Work package 4 uses the results of work package 3 (design and dissemination of software simulator) as well as of work package 2 (development of a formal model). Work package 4 is executed partly in parallel with work package 3 to allow for the results of the experimentation to feed back into the development process.			
Deliverables and Milestones			
Milestones will be the submission of the two main papers (one for the field of knowledge representation and reasoning and one for the field of computational economics). Once the actual software simulator has been finished (at the end of work package 3) we expect the experimentation to be relatively straightforward, since it essentially consists of running the simulator, which will also recurrently take place during the testing phase at the end of work package 3..			
Resources			
Name of researcher	Partner	Qualif. Level	Person Months
Martin Caminada	SnT	Postdoc	3
Björn Ottersten	SnT	Professor	1
Tibor Neugebauer	LSF	Professor	1

4.3 Risk management and quality assurance (max. 1 page)

One of the assumptions in the project is the availability of a bachelor or master student who is willing to do part of the software development in work package 3 for his/her bachelor or master thesis. Although we cannot completely take this availability for granted, experience has shown that finding a student to perform programming is not too troublesome, especially since there is a remuneration associated to this student task.

Since normally student thesis projects are started in February, work package 3 has been scheduled to start in Month 12 in order to be able to select a student who could start working on the project right away. This, however, assumes that the project is able to start in January 2011, and that no delays take place in work packages 1 and 2, whose progress will regularly be monitored.

Since the outputs of work packages 1 and 2 will be vital for the success of work package 3, our aim is to obtain early stage feedback on the results of these work packages by submitting to workshops like MABS (Multi-Agent Based Simulation). Hence, the aim of the publications mentioned in the above descriptions of work packages 1 and 2 is not so much to have a high-level publication, but more to have early-stage feedback of work that is essentially still in

progress. Higher-ranked publications are to be submitted at the end of the project, in work package 4.

5. Project Outputs (max. 3 pages for 5.1 – 5.3)

5.1 Contributions to advancement of knowledge, people and institutions

In general, the project aims to increase understanding of how knowledge is generated when the individual agents involved in the knowledge generation are subject to market-based incentives. This subject has, as far as we are aware, never been studied, at least not when taking into account complex forms of reasoning. We believe this to be the blind spot when trying to explain instances of market failure, since if the market is ill-informed due to flaws in the collective reasoning process, market imperfection is a logical result.

On the individual level, the project is expected to advance the position of the main researcher to be hired (Martin Caminada). Caminada has done some high-ranking publications in the past (including in top-journals such as *Artificial Intelligence*) mainly in the topic of knowledge representation and reasoning. His current research interest is in reasoning about and reasoning in the possible presence of dishonesty and strategic communication.

The project would also be a good opportunity for one of the students at the CSC research group of the University of Luxembourg to become familiar with advanced techniques of multi-agent systems, computational economics, simulation and knowledge representation techniques. Knowledge of simulation techniques, especially in the context of economics, would be beneficial for various other purposes when choosing a career in the financial sector when having an IT background.

The project would also be the first form of collaboration between the Interdisciplinary Center for Security and Trust (SnT) and the Luxembourg School of Finance. During earlier informal discussions it became clear that there exists a significant overlap in research interests, and we think that the current project is a good way of trying to build on these overlaps. It is quite possible that the project will serve as a initiator for further collaboration, even after the project itself has been finished.

5.1 bis Advancement of the young PI's research career

Not applicable.

5.2 Intended long-term impact and contribution to FNR CORE programme objectives

The project is expected to gain understanding of how and under which conditions markets will start to have warranted or unwarranted beliefs regarding complex issues. The importance of this topic is illustrated by the recent credit crisis. If the market would have been sufficiently aware of the inherent risks related to the mortgage backed securities, then some of the events that we have experienced over the last few years might have been avoided. We regard the recent credit crisis not as an isolated incident but as the result of fundamental flaws in collective reasoning under market conditions. How exactly to model these flaws is one of the main tasks of the proposed project.

Although the currently proposed project can be seen as one of the first attempts to model these flaws (as well as the conditions under which they occur and their impact) we believe that any insight in this issue would be highly beneficial from the perspective of public policy. If, for instance, the finding would be that flaws in collective reasoning are more likely to occur in the presence of highly complex products, then this could be a reason to have at least a reasonable

level of prudence when allowing critical financial institutions (such as banks with customer deposits) to trade in these products. We hope that our findings will contribute to the debate among economists precisely about these issues. This is why we aim to disseminate the findings of the project not only in scientific publications, but also for a more general audience. As an example, one of the proposed collaborators of the project (Martin Caminada) has recently been invited to give a presentation on this subject for a student symposium in the Netherlands, on artificial intelligence models of deception. Caminada has also published previously articles for a more general audience (like on Japanese robotics and the IT-industry of emerging economies) and intends to do the same for the results of the current project, although a precise forum still needs to be identified.

The modeling of reasoning processes under market conditions, as well as studying the flaws in these processes, is still a relatively new and unexplored topic. Once the basic approach and results of the project become known, it is quite likely that this will result in various forms follow-up research of other researchers. This would be in line with some previous work of Caminada², which is highly cited and is now used as the basis of various work of other researchers.

The domain described in the current research proposal would also be an opportunity for Luxembourg to initiate and internationally acknowledged line of research, which would increase the visibility of artificial intelligence and economics research being carried out in Luxembourg. Furthermore, when successfully performed, the research would provide some insight in when to trust collective opinions (which are generated in the kind of market-based conditions described earlier in this document) and when not. This kind of insight would be highly beneficial for the Luxembourg finance sector, and we expect part of the dissemination of the results to take place through the contact network of the Luxembourg School of Finance.

5.3 PhD student supervision and research lines

No PhD student is to be assigned on the proposed project, which will employ only one postdoc and one student programmer.

Current research that is taking place within the Computer Science and Communications (CSC) research group, and especially within its Interdisciplinary Laboratory for Intelligent and Adaptive Systems (ILIAS) includes mathematical models of reasoning, both in quantitative (numerical) and qualitative (symbolic) forms. This research includes the construction and analysis of various formal logics for epistemic reasoning (that is, for reasoning about knowledge), as well different projects (such as AASTM and DYNARG) that are related to formal models of argumentation. The currently proposed project would therefore fit very well in some of the current research lines of CSC.

Most current economic models assume that financial markets are efficient, yet the existence of durable disequilibria and the increasing set of paradoxes rising from the recent past necessitate the adoption of new tools describing the bounded rationality of human agents. The LSF develops theoretical and experimental works in this field in collaboration with external institutions, and the current project would therefore fit into this particular research line of LSF as well.

² Like (Caminada and Amgoud 2007), (Caminada, 2006a) and (Caminada, 2006b)

6. Project Participants and Management (max. 2 pages for 6.1 – 6.4)

6.1 Description of the consortium

The research is aimed to be a collaboration between the Interdisciplinary Center for Security, reliability and Trust (SnT), and the Luxembourg School of Finance (LSF). The project will be administered by SnT, with a strong input from LSF.

The Interdisciplinary Centre for Security, Reliability and Trust launched by the University of Luxembourg (UL) targets research and PhD education in Security, Reliability and Trust. SnT provides a platform for interaction and collaboration between university researchers and external partners (industry/institutes/government). To create an impact, an interdisciplinary research approach is essential, taking not only technical aspects into account but also addressing business/organisational, human/user, and legal/regulatory issues.

The Luxembourg School of Finance is the Department of Finance of the Faculty of Law, Economics and Finance of the University of Luxembourg. LSF's mission is to offer post-graduate education and to foster research in the areas of banking and finance. The core objective of the Luxembourg School of Finance is to be an academic leader in the creation and dissemination of an integrated body of knowledge in banking and finance. In order to fulfill this mission, LSF places a high priority on research. A permanent research program in the areas of asset and risk management, fund industry, behavioral finance, law and finance, and quantitative finance has been set up at the end of 2005.

SnT as a partner in the project has a wide experience with formal methods and formal models of reasoning, and is hence a logical partner in the project. Input and feedback from the Luxembourg school of Finance will ensure that the project is in line also with economic theory. This includes aspects of price formation, and pricing strategies of the information providers.

6.2 List of total resources

Name of researcher	Partner institution	Position	Qualification Level	Gender	Person*m onths
Martin Caminada	SnT	main researcher	postdoc	M	24
Björn Ottersten	SnT	PI	professor	M	4
Tibor Neugebauer	LSF	adviser	professor	M	4
Student	SnT	programmer	MsSc	M/F	6

6.3 Communication and decision-making

Since the project is relatively small, we do not expect decision-making to be a major issue. In principle, all necessary decisions are made by the PI (Björn Ottersten) after consulting the other researchers involved in the project. Coordination with LSF will take place through frequent meetings with Tibor Neugebauer. At the beginning of the project (work packages 1 and 2), these meetings are scheduled to be twice a month, in order to obtain input regarding existing research and the formation of an overall formal model. Somewhat less input of LSF will be required during the software development phase (work package 3), although it would still be useful to obtain feedback during the testing phase, to determine whether the results are in line with one would expect from general economic theory. More intense collaboration is expected to

take place at the end of the project, at the dissemination phase (work package 4), when the obtained results need to be embedded in and related to existing lines of research in economic theory.

6.4 Consortium agreement and Intellectual Property Rights management

Both SnT and LSF are part of the University of Luxembourg, so no consortium agreement or intellectual property rights agreement is necessary. The University of Luxembourg will have intellectual property of all resulting documents and software developed during the course of the project. The software simulator is to be licensed under an open source license (such as GPL). Since the main purpose of the simulator is to generate theoretical insight (rather than being directly applicable in everyday situations), the advantage of making the simulator available under an open source license is that this would allow for other researchers to verify our results, as well as to apply the simulator as a basis for subsequent research.

7. Justification of project budget and requested funding (2 pages)

The main part of the funding requested is necessary for employing a postdoc who will work full-time on the project, during its lifespan of two years, as well as a student programmer who will work on the technical implementation of the simulator. The two other researchers involved in the project (Björn Ottersten and Tibor Neugebauer) are already employed by the University of Luxembourg (by SnT and LSF, respectively) and their salaries are paid from existing university budget.

The results of work packages 1 and 2 are to be submitted and presented at an international workshop, in order to receive early stage feedback. However, since the relevant workshops are being held at different locations, it is not yet certain that they will be held in Europe at the time the results from work packages 1 and 2 become available. Therefore, we have requested additional travelling budget for the year 2011, in case it is necessary to present the results outside of Europe.

As for the results of the project as a whole, to be published in work package 4, we explicitly want to keep open the possibility of presenting these at one of the top-conferences (like AAMAS or AAAI) in the field, which are often outside of Europe. Hence, we have requested additional travel budget for 2012 as well.

No funding is requested for the expenses of providing a workplace (including that of providing a PC) since these will be covered by existing budget of SnT, in case no existing equipment from previous projects is available. Similarly, no funding is requested for purchasing literature, since this is expected to be covered by existing SnT budget.

8. Bibliography / References

S. Alfarano, Th. Lux and F. Wagner
Estimation of Agent-Based Models:
The Case of an Asymmetric Herding Model
Computational Economics (2005) 26: 19-49

L. Amgoud
Argumentation for Decision Making
Book chapter in I. Rahwan and G. Simari (eds.) Argumentation in Artificial Intelligence, 2009
(pages 301- 320)

A. Bondarenko, P.M. Dung, R.A. Kowalski and F. Toni
An abstract, argumentation-theoretic approach to default reasoning
Artificial Intelligence Vol 93, pp. 63-101, 1997

S. Bouman and B. Jacobsen
The Halloween Indicator, "Sell in May and Go Away": Another Puzzle
The American Economic Review, Vol. 92, No. 5 (Dec., 2002), pp. 1618-1635

M.W.A. Caminada.
On the Issue of Reinstatement in Argumentation. JELIA 2006, pages 111-123.

M.W.A. Caminada.
Semi-stable semantics. In P.E. Dunne and T.J.M. Bench-
Capon, editors, Computational Models of Argument; Proceedings of COMMA
2006, pages 121–130. IOS Press, 2006.

M.W.A. Caminada and L. Amgoud.
On the evaluation of argumentation formalisms. Artificial Intelligence 171(5-6):286-310 (2007)

M.W.A. Caminada and Y. Wu.
An Argument Game for Stable Semantics. Logic Journal of the IGPL 17(1):77-90 (2009)

C. Campbell
Colin Campbell's Response to the Guardian IEA Reporting, The Oil Drum, 2009
(<http://www.theoil Drum.com/node/5970>)

D. Colucci and V. Valori
Asset Price Dynamics When Behavioural Heterogeneity Varies
Comput Econ (2008) 32:3–20

M. Dastani, F. Dignum, J.J. Meyer
3APL: A Programming Language for Cognitive Agents. ERCIM News, European Research
Consortium for Informatics and Mathematics, Special issue on Cognitive Systems, No. 53,
2003.

P.M. Dung.
On the acceptability of arguments and its fundamental role in non-monotonic reasoning, logic
programming and n-person games. Artificial Intelligence, 77:321–357, 1995.

- P. M. Dung, P. Mancarella and F. Toni.
Computing ideal sceptical argumentation. *Artificial Intelligence*, 171(10-15):642–674, 2007.
- P.E. Dunne and M.W.A. Caminada.
Computational Complexity of Semi-Stable Semantics in Abstract Argumentation Frameworks.
JELIA 2008, pages 153-165.
- P. E. Dunne and M. Wooldridge.
Complexity of Abstract Argumentation. Chapter 5 of: *Argumentation in Artificial Intelligence* (Ed: I. Rahwan and G. Simari), Springer, pages 85-104, 2009
- S. Ecça, M. Marchesi, A. Setzu
Modeling and Simulation of an Artificial Stock Option Market. *Comput Econ* (2008) 32:37–53
- J.S. Fons
Rating Competition and Structured Finance
The Journal of Structured Finance Fall 2008, Vol. 14, No. 3: pp. 7-15
- G. Ferri, L.G. Liu, J.E. Stiglitz
The Pro-cyclical Role of Rating Agencies: Evidence from the East Asian Crisis
Economic Notes, Volume 28, Number 3, November 1999 , pp. 335-355
- K. Fullam and K.S. Barber
Learning Trust Strategies in Reputation Exchange Networks," *The Fifth International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS-2006)*, Hakodate, Japan, May 8-12, pp. 1241-1248.
- M. Gelfond and V. Lifschitz.
The stable model semantics for logic programming.
In Proceedings of the 5th International Conference/Symposium on Logic Programming, pages 1070–1080. MIT Press, 1988.
- M. Gelfond and V. Lifschitz.
Classical negation in logic programs and disjunctive databases.
New Generation Computing, 9(3/4):365–385, 1991.
- D. Hirshleifer and T. Shumway
Good Day Sunshine: Stock Returns and the Weather
The Journal of Finance, Vol. 58, No. 3 (Jun., 2003), pp. 1009-1032
- T. Lux and M. Marchesi
Scaling and criticality in a stochastic multi-agent model of a financial market. *Nature*, 397, 498–500. (1999)

D. Kahneman and A. Tversky

Prospect Theory: An Analysis of Decision under Risk
Econometrica, Vol. 47, No. 2 (Mar., 1979), pp. 263-291

A. Kirman and G. Teyssi re, G.

Microeconomic models for long memory in the volatility of financial time series. *Studies in Nonlinear Dynamics & Econometrics*, 5, 137–156. (2002)

J. Leggett

What they don't want you to know about the coming oil crisis
The Independent, Friday, 20 January 2006

J.S. Lin

Learning in a Network Economy
Computational Economics (2005) 25: 59-74

M.J. G mez Lucero, C.I. Ches nevar and G.R. Simari

On the Accrual of Arguments in Defeasible Logic Programming
Proceeding of the 21st International Conference on Artificial Intelligence (IJCAI 2009) pp. 804-809

J. Mathis, J. McAndrews, J.C. Rochet

Rating the raters: Are reputation concerns powerful enough to discipline rating agencies?
Journal of Monetary Economics Volume 56, Issue 5, July 2009, Pages 657-674

P. McBurney and S. Parsons

Dialogue Games for Agent Argumentation

book chapter in I. Rahwan and G. Simari (eds.) Argumentation in Artificial Intelligence, 2009
pages 261-280

S. Modgil

Reasoning About Preferences in Argumentation Frameworks,
Artificial Intelligence Volume 173, Issue 9-10, pp 901-1040, 2009

S. Modgil and M.W.A. Caminada

Proof Theories and Algorithms for Abstract Argumentation Frameworks book chapter in I.
Rahwan and G. Simari (eds.) Argumentation in Artificial Intelligence, 2009

F. Partnoy

The Siskel and Ebert of Financial Markets: Two Thumbs Down for the Credit Rating Agencies.
Washington University Law Quarterly, Vol 77, pp. 619-712, 1999

J. L. Pollock

How to reason defeasibly. Artificial Intelligence, 57:1–42, 1992.

H. Prakken

AI & Law, logic and argument schemes. *Argumentation* 19 (2005): 303-320. (special issue on The Toulmin model today)

H. Prakken

Coherence and flexibility in dialogue games for argumentation. *Journal of Logic and Computation* 15 (2005): 1009-1040.

H. Prakken

Models for Persuasion Dialogue

book chapter in I. Rahwan and G. Simari (eds.) *Argumentation in Artificial Intelligence*, 2009
pages 281-300

H. Prakken

An abstract framework for argumentation with structured arguments. *Argument and Computation* Vol 1, No 2, to appear 2010

H. Prakken and G. Sartor

Argument-based extended logic programming with defeasible priorities,
Journal of Applied Non-Classical Logics, Volume 7, pp. 25-75, 1997

I. Rahwan, F. Zablith and Ch. Reed

Laying the Foundations for a World Wide Argument Web. *Artificial Intelligence*, Vol 171, No 10-15, pages 897-921 (2007)

I. Rahwan, K. Larson and F. Tohmé

A Characterisation of Strategy-Proofness for Grounded Argumentation Semantics.

In: Semantics of the 21st International Joint Conference on Artificial Intelligence (IJCAI), Pasadena, California, USA (2009)

I. Rahwan and Ch. Reed

The Argument Interchange Format. Book chapter in I. Rahwan and G. Simari (eds.) Argumentation in Artificial Intelligence, 2009

pages 383-402

I. Rahwan and G. Simari

Argumentation in Artificial Intelligence, Springer Verlag, 2009

I. Rahwan, M. I. Madakkatel, J. F. Bonnefon, R. N. Awan and S. Abdallah.

Behavioural Experiments for Assessing the Abstract Argumentation Semantics of Reinstatement. Cognitive Science. (accepted) 2010

R. Reiter

A logic for default reasoning. Artificial Intelligence, 13:81–132, 1980.

P. Resnick, K. Kuwabara, R. Zeckhauser, E. Friedman

Reputation systems. Commun. ACM 43(12) (2000) 45–48

J.S. Rosenschein and G. Zlotkin

Rules of Encounter; designing conventions for automated negotiation among computers
MIT Press, 1994

E.M. Saunders Jr.

Stock Prices and Wall Street Weather

The American Economic Review, Vol. 83, No. 5 (Dec., 1993), pp. 1337-1345

G. Simari and R. Loui. A mathematical treatment of defeasible reasoning
and its implementation. Artificial Intelligence, 53:125–157, 1992.

T.J. Sinclair

The New Masters of Capital; American Bond Rating Agencies and the Politics of
Creditworthiness. Cornell University Press, 2005

Ch. Spatt

Discussion of "Ratings Shopping and Asset Complexity: A Theory of Rating Inflation"

Journal of Monetary Economics. Volume 56, Issue 5, July 2009, Pages 696-699

V. Skreta and L. Veldkamp

Ratings shopping and asset complexity: A theory of ratings inflation

Journal of Monetary Economics. Volume 56, Issue 5, July 2009, Pages 678-695

E. Staab and M.W.A. Caminada
On the Profitability of Incompetence
Proceedings of the 11th International Workshop on Multi-Agent-Based Simulation (MABS 2010)

E. Staab and M.W.A. Caminada.
Assessing the Impact of Informedness on a Consultant's Profit.
Technical Report University of Luxembourg. ISBN 978-2-87971-027-3 (2010)

E. Staab and M.W.A. Caminada.
Simulating Knowledge and Dishonesty in a Client-Consultant Setting.
Proceedings of the 21st Benelux Conference on Artificial Intelligence (BNAIC 2009) pages 397-398

A.N.R. Sy
Rating the rating agencies: Anticipating currency crises or debt crises?
Journal of Banking & Finance
Volume 28, Issue 11, November 2004, Pages 2845-2867

T.T. Tang
Information asymmetry and firms' credit market access:
Evidence from Moody's credit rating format refinement
Journal of Financial Economics. Volume 93, Issue 2, August 2009, Pages 325-351

L. Tesfatsion
Introduction to the Special Issue on Agent-Based Computational Economics
Computational Economics 18: 1–8, 2001.

L. Tesfatsion
Agent-Based Computational Economics: Growing Economies From the Bottom Up
Artificial Life Vol. 8, No. 1, Pages 55-82 (2002)

L. Tesfatsion and K. Judd (eds.)
Handbook of Computational Economics, Volume 2, Agent-Based Computational Economics
Elsevier/North-Holland, 2006

A. Tversky and D. Kahneman
Loss Aversion in Riskless Choice: A Reference-Dependent Model
The Quarterly Journal of Economics, Vol. 106, No. 4 (Nov., 1991), pp. 1039-1061

G. Vreeswijk
Studies in defeasible argumentation. PhD thesis at Free University of Amsterdam, 1993.

G.A.W. Vreeswijk and H. Prakken
Credulous and sceptical argument games for preferred semantics
Proceedings of the 7th European Workshop on Logic for Artificial Intelligence (JELIA-00)
Springer Verlag, Springer Lecture Notes in AI 1919, pp. 239-253 (2000)