

Going Public and the Sale of Shares with Heterogeneous Investors: Agent-Based Computational Modelling and Computer Simulations

Pietro Panzarasa¹, Nicholas R. Jennings², Timothy J. Norman³

¹ Dept. of Electronics and Computer Science, University of Southampton,
Southampton SO17 1BJ, UK.

`pp@ecs.soton.ac.uk`

² Dept. of Electronics and Computer Science, University of Southampton,
Southampton SO17 1BJ, UK.

`nrj@ecs.soton.ac.uk`

³ Dept. of Computing Science, University of Aberdeen,
Aberdeen AB24 3UE, UK.

`tnorman@csd.abdn.ac.uk`

Abstract

In this paper we use agent-based computational modelling and computer simulations to examine the interrelationship between different selling strategies for going public. A great deal of recent empirical evidence suggests that to maximise the revenue raised from the shares sold in the public offering, it is fundamental to choose the appropriate design for the sale which, in turn, reflects the final ownership structure. This literature establishes that the market for shares is segmented and, particularly, that firms manage the sale of shares with the purpose of discriminating between relatively small and passive investors and applicants for large potentially controlling blocks. One of the key questions in this area, then, is: How and to what extent should this heterogeneity among potential investors influence the firm's strategy for selling shares? Here we attempt to address this question from the standpoint of using agent-based computational modelling and computer simulations. Results show that the design of the sale is an important determinant of the performance of the negotiation process through which the firm is sold. A sequential sale beginning with an initial public offering of dispersed shares, followed by a negotiated sale of a controlling block is, in general, more effective than other alternative selling strategies. Changing the negotiation protocol itself can act as an effective way of impacting upon the revenue raised and the length of the process. The interrelationship between the method of sale and the performance may also depend on the degree of cognitive accuracy that characterises the negotiating agents' mental representations of their physical and social environment.

Key Words: agent-based modelling and simulation; sale of shares; ownership structure; initial public offering.

1 Introduction

In this paper we use agent-based computational modelling and computer simulations to examine alternative selling strategies for going public. When a privately held firm goes public, the volume of shares transferred to new owners lastingly impacts upon the firm's ownership structure and, therefore, the firm's value will be influenced. Much literature suggests that to maximise the revenue raised from the shares sold in the public offering, it is fundamental to choose the appropriate design for the sale which, in turn, reflects the final ownership structure (Hanley and Wilhelm 1995; Mello and Parsons 1998; Mikkelsen and Ruback 1985; Shleifer and Vishny 1986; Barclay and Holderness 1989). On the one hand, the public offering of shares will mainly involve relatively *small investors* who will just remain passive holders of the firm's shares. On the other, there might well be investors who are interested in a large block of shares and are willing to ultimately play an active role in the firm's management. These *active investors* are individuals or organisations that, by purchasing a controlling interest, will eventually act either as monitors of the firm's current management or as proponents of new directions in the firm's long-term strategies and decisions via changes in the management team or simply by bringing expertise or other resources to the firm.

Active investors seeking controlling blocks can be put into competition with small, passive investors seeking the same shares in dispersed allotments. This can be done, for example, by initially selling a portion of the shares to passive investors and then later putting a controlling block up for sale on terms determined in part by the price of the outstanding shares (Mello and Parsons 1998). This highlights the importance of treating the issuance of shares as a process incorporating transactions over time, instead of a single event independent of the firm's plans for subsequent financing as has often been the case (Brennan and Franks 1995; Klein *et al.* 1991; Mikkelsen *et al.* 1995; Shipper and Smith 1986). The results of earlier sales may affect the terms of later sales, and the terms of earlier sales may be determined in part by the expected beneficial impact on ownership structure and the terms of later sales (Mello and Parsons 1998).

These observations are supported by a great deal of recent empirical evidence about the segmentation of the market for shares (Barry *et al.* 1990; Brennan and Franks 1995; Hanley and Wilhelm 1995; Holderness and Sheehan 1988; Pagano *et al.* 1996; Rydqvist and Högholm 1994). This literature shows that firms manage the sale of shares with the purpose of discriminating between relatively small and passive investors and applicants for large blocks (Brennan and Franks 1995). This engenders the belief that going public is a complex and extended process with distinct and heterogeneous markets for dispersed shareholdings and potentially influential blocks. Thus some of the key questions that need to be addressed

are: How and to what extent should this heterogeneity influence the firm's strategy for selling shares? Why do some privately held firms go public instead of selling control exclusively to another set of private investors, despite the apparent value loss associated with the free rider problem (Weston *et al.* 1990)? How should a firm design a sale of shares to maximise expected revenues? Should the marketing of dispersed shareholdings and that of potentially controlling blocks occur separately in selling a firm? In particular, does it matter that the sale disperses the shares through an initial public offering (IPO) before negotiating with an investor who is interested in buying a controlling block? Or is it better firstly to pass on a block to someone who wants a controlling stake, and then organise a subsequent IPO? In this paper, we attempt to shed some light on these substantive questions from the standpoint of using an agent-based computational approach that is rich enough to do justice to a broad range of phenomena and empirical evidence, yet precise enough to support a rigorous theoretical analysis of the fundamental mechanisms and properties that underpin the sale of a firm.

To this end, a computational agent-based model of going public will be developed using the UM-PRS architecture (see Lee *et al.* 1994 for details), which, in turn, will be extended and enriched by new computational components and operations performed upon them. In particular, a specification of the negotiating agent's mental state will be developed (Section 3.1) and individual social behaviour will be specified in terms of rules for local social interaction (Sections 3.2 and 3.3), and rules for global social interaction (Section 3.4). Furthermore, a higher-order multi-agent level will be introduced, and a description of the joint behavioural components of negotiation will be outlined (Section 3.5). The model will then be used as a test-bed for undertaking a series of computer simulations. Two groups of virtual experiments¹ will be introduced. The first group is intended to examine the individual and combined effect that the design for the sale of shares and the economic actors' negotiation behaviour have upon the performance of the going public process. The second group is undertaken to examine the individual and combined effect that the design for the sale of shares and the economic actors' cognitive accuracy have upon the performance of the process. These illustrative virtual experiments will show how the study of the financial decision of going public can be advanced by adopting methods and ideas from the areas of computational modelling and Distributed Artificial Intelligence (DAI) theory and practice. In particular, we will show how and to what extent DAI computational models provide deep insights into the going public process by permitting researchers to explore different combinations of modelling factors easily and by allowing them to analyse these factors systematically.

The remainder of the paper is structured as follows. Section 2 makes the case as to why an agent-oriented approach is well suited to modelling and reasoning about the financial process of going public. In Section 3 our computational model will be developed. The model

comprises: the economic actors' mental state specifications (Section 3.1); a set of evaluation parameters that the agents use to assess the messages received (Section 3.2); a negotiation protocol (Section 3.3); a set of selling strategies for going public (Section 3.4); and a negotiation cycle (Section 3.5). Section 4 introduces two experimental studies. For each study, we will provide a description of the experimental design and a discussion of the results. Section 5 places our work in the context of the related literature. Finally Section 6 will summarise the major findings and provide avenues for future work.

2 The Case for an Agent-Based Computational Approach

In this paper, it is argued that agent-based computing and virtual experiments offer an invaluable tool for modelling and reasoning about the properties of a process of going public. The first step in seeking to provide a justification for this claim is to informally tease out the key concepts that together make up the tapestry of the agent-based approach, and show to what extent these concepts can offer an effective way of modelling and simulating the process of going public. A more detailed account of the match between the approach and the domain will be given in the following sections where our computational model will be developed and its properties discussed.

There is no real agreement on the core question of exactly what the common threads of agent-based computing are. Nevertheless, an increasing number of researchers find the following concepts as inherently underpinning an agent-based approach: agenthood, inter-agent social behaviour, cognition, and organisational design (Jennings 2000). In what follows, each of these notions will be elucidated, and the degree of match between each of them and the nature of our research domain will be briefly discussed. As expected, a high degree of match would be indicative of the suitability of the agent-based approach to the domain.

Within the Artificial Intelligence (AI) community there is increasing agreement that the essential properties of *agenthood* are autonomy, situatedness, reactivity, proactiveness, and social ability (Franklin 1997; Wooldridge 1997, 2000). Given this, the key question is to determine how and to what extent these properties of agenthood can be seen to be a useful means of characterising the behaviour and properties of the economic actors who are typically involved in a process of going public. Let us restrict our evaluation to three basic types of actors: the seller of a firm; an active investor, and a passive investor². First, these actors have control over their actions and internal states. More specifically, although these actors may well be expected to take others' interests into account (e.g., the seller of the firm may have to consider the impact that the sale may have upon other stakeholders, such as the suppliers, the customers, the employed staff, etc.), they typically have the ability to act

independently and make independent decisions with no a priori need of the direct interventions of others (autonomy). Second, economic actors are noticeably embedded in a specific physical (e.g., a region within a country), social (e.g., a web of investors, suppliers, customers), economic and financial (e.g., a system of economic and financial practices) environment (situatedness). Third, they receive inputs from their environment (e.g., market rules; legal arrangements) and they are responsive to these perceived inputs (reactiveness). Fourth, in trying to pursue their goals (e.g., maximisation of expected utility), economic actors exhibit state-directed behaviour and are expected to strategically exploit the opportunities they have to make their goals fulfilled (proactiveness). Finally, economic actors are social actors, in that they are often expected/required to engage in a wide range of social interactions (e.g. negotiation, cooperation) with others (social ability). Drawing these points together, the essential features of agenthood seem to be particularly well suited to represent and reason about what constitutes, on a more general level, the main facets of the economic actor, and in particular of the actor whose aim is either to sell or to purchase shares of a firm.

Moving onto the second property of agent-based computing – *inter-agent social behaviour*. Most real-world problems to which an agent-based approach has been applied are decentralised, and can be typically represented in terms of systems comprising multiple agents, multiple loci of control, decentralised data, asynchronous computation, and multiple perspectives (Bond and Gasser 1988; Jennings *et al.* 1998). Traditionally, research into all types of such systems composed of multiple (semi-) autonomous components has been carried out under the banner of Multi-Agent System (MAS) theory and practice. From this perspective, agents need to interact in order to cope with their social dependencies and to have their individual objectives achieved (Jennings 1993). Establishing social interactions and exhibiting social behaviour is feasible as agents are precisely modelled as entities endowed with social ability. Moreover, agents are assigned the computational apparatus that allows them to handle their interactions with one another in a flexible manner. That is, agents make decisions about the scope and nature of their interactions at run-time, and this, in turn, allows them to deal with interactions occurring at unpredictable times, for unpredictable reasons, and between unpredictable components.

As with agenthood, we now need to examine whether and to what extent the above properties associated with the inter-agent social behaviour perspective can offer useful tools for modelling our domain, and in particular the social relationships occurring between the economic actors involved in it. The evaluation is quite straightforward. The final sale of the firm's shares is the result of a complex negotiation process in which the owner of the firm tries to come to some form of agreement with various potential investors. In this process, data is decentralised, information is incomplete and variously distributed between the seller and the investors. The complexity of the process may be further increased by the presence of

distinct markets for the shares of a firm – one for dispersed shares and another for potentially controlling blocks. This, in turn, turns out to be detrimental to the negotiating actors' ability to control and observe their environment, and subsequently to the quality and accuracy of the information they can gather through a monitoring activity. As it is impossible to know a priori what interaction strategy is suited to what counter-part, both the seller and the potential investors need to handle their interactions between each other in a flexible manner. This allows them to make run-time decisions about who to negotiate with and what selling strategies to adopt in what circumstances, as soon as new pieces of information become available. From this backdrop, the suitability of a MAS approach for coping with the decentralised nature of our domain is self-evident, as is the match between the repertoire of tools and techniques offered by this approach and the need for flexibility that must be met for managing the social relationships occurring within the domain.

Another key feature of an agent-based approach is the recognition of the primacy of *cognition* in both individual and collective rational action. Typically, theories of cognition make use of the repertoire of notions and concepts offered by folk psychology, whereby behaviour is understood in terms of a number of mental attitudes and the interaction between them (Carley 1989; Dennet 1987). Probably the best-known and most influential model of cognition that has been associated with an agent-based approach is the Belief-Desire-Intention (BDI) model (Bratman 1987; Cohen and Levesque 1990; Jennings *et al.* 1998; Rao and Georgeff 1991; Singh 1995; Wooldridge 2000). This model gets its name from the fact that it represents individual cognition as a mental state comprising such mental attitudes as beliefs, desires, and intentions. Intuitively, an agent's beliefs correspond to information the agent has about the world. Desires represent the set of states of affairs the agent wants to bring about. Finally, intentions refer to those states of affairs the agent is committed to bringing about (Rao and Georgeff 1991; Wooldridge 2000). Further, attempts have been made to explain how a BDI agent's mental state lead it to select and perform rational actions (Bratman 1987; Rao and Georgeff 1991). Along these lines, a number of theories of practical reasoning have been developed in terms of schemas that axiomatise interrelationships between beliefs, desires and intentions (Kraus *et al.* 1998; Panzarasa *et al.* 2001; Wooldridge 2000).

The BDI model is particularly well suited to reasoning about the mentalistic properties of the economic actors involved in a process of going public, and to modelling them as intentional agents endowed with a three-fold set of mental attitudes. First, each actor maintains a set of beliefs about its own environment. These beliefs are incomplete and may be updated as new pieces of information become available. Further, the actor's behaviour is state-directed, that is, it is regulated by the mental representation of some states of the world. This state-directedness may be thought of in terms of either the actor's motivation to get its

behaviour to adhere to a top-level agenda (i.e. goals) or in terms of the actor's self-commitment to achieving some states of the world (i.e. intentions). For example, the seller of a firm can be regarded as motivated by the goal to get the maximum amount of revenue from the sale of the firm (Mello and Parsons 1998). His behaviour will then be triggered by the intention to perform some actions that are instrumental to the attainment of that goal, for instance engaging in a negotiation process with potential investors, receiving messages from them, and evaluating these messages (Section 3.1). Also, the BDI model offers a way for describing the economic actor's decision-making process in terms of transitions between mental attitudes. Upon receiving new pieces of information, beliefs will be updated and intentions generated that do not contradict the actor's beliefs. For example, upon receiving a message indicating the rejection of a previous offer, the seller may get more realistic beliefs about his opponent and generate a new offer to be forwarded. Intentions, in turn, will have to match against top-level goals, which may also be subject to modifications as a result of information-gathering processes. For example, given the seller's goal to maximise the revenue from the sale of the firm, each intention to offer a specific price must be consistent with this goal, thus ensuring that, all the available information considered, the offer made conveys the best price that could have been proposed.

Finally, let us consider the fourth concept of agent-based computing - *organisational design*. The agent-based mindset provides suitable organisationally oriented abstractions for flexibly treating a collection of components either as a higher-level conceptual unit or as a pluralistic and decentralised set of interrelated units. Moreover, MAS theory and practice offer a rich set of tools and mechanisms for explicitly representing a number of organisational structures and phenomena (e.g., roles (Barbuceanu 1997), plans (Kinny *et al.* 1992), norms (Shoham and Tennenholtz 1992)). Given this, we now need to evaluate whether and to what extent this organisational design perspective of the agent-based approach can offer an adequate means for modelling our domain, and in particular its structural and organisational variables and properties. Going public is a complex process that involves varying structures for coordinating human activity, ranging from various forms of markets to flat or hierarchical organisational structures. Both intra- and inter-organisational relationships are involved. Both the seller and the investor may be either an individual economic actor or a company with its own organisational structure. The shares being sold may be those of either a stand-alone firm or of a division of a public corporation (Klein *et al.* 1991; Schipper and Smith 1986). Passive investors may be thought of as acting either in a "pure form" of decentralised market, where all potential buyers are in contact with all possible suppliers and they each make their own decision about which transaction to accept (Baligh and Richartz 1967; Malone 1987). Or, more plausibly, they may act in a centralised market where there is no need to contact all possible sellers because a broker is already in contact with those who want to sell the stock

(Malone 1987). Also, our research issue itself as to whether the design of the sale of shares may increase the value of the firm has a straightforward organisational interpretation. It emphasises the potential of the capital market to establish an optimal ownership structure for the firm (Bebchuck 1994; Bebchuck and Zingales 1995; Kahan 1993), and ultimately it evokes the well-known trade-off between markets and hierarchies which has been analysed by a number of economic and organisational design theorists (Berle and Means 1932; Williamson 1975, 1981).

Having made the case that an agent-based approach is well suited to the *theoretical* purpose of representing and reasoning about the going public process, the next step is to determine whether it can also be usefully exploited for the *practical* purpose of developing implementation architectures. More specifically, how and to what extent can an agent-based approach be adopted in an attempt to help practitioners to realise software systems for managing negotiation processes in varying real-world domains? There have already been a number of attempts at generating implemented artificial systems with negotiation capabilities (Jennings *et al.* 1996; Rosenschein and Zlotkin 1994; Sycara 1988). However, there is currently a large chasm between the theoretical work which examines and specifies various models of negotiation and the agent-based artificial systems which often use *ad hoc* models for intelligent behaviour and social interaction. In this paper we try to bridge this gap by developing a computationally grounded model that will be expressed *directly* in terms of the language, syntax, data structures and mechanisms of an agent-based architecture. Our approach, therefore, enables a straightforward mapping between a theoretical model of going public and an implementation architecture, thus providing direct assistance to practitioners who are concerned with developing artificial systems comprising negotiating economic actors variously involved in selling or buying the shares of a firm.

3 A Model of Going Public with Heterogeneous Investors

In this section we describe an agent-based computational model of going public and its fundamental assumptions. The model is in the tradition of behavioural models in the finance literature associated with work by Bebchuk and Zingales (1995), Kahan (1993), and Mello and Parsons (1998). It is intended as an elaboration of their ideas, with some variations on the theme. Its major contributions are both methodological and theoretical. Methodologically, our agent-based approach allows us to develop a fully explicated account of the negotiating agents' mental states and social behaviour (Jennings 2000). In doing so, our work features a different perspective on the role of cognition and sociality than other papers on IPOs and equity financing. Theoretically, the model extends the state of the art in that it shows how the performance and value of a firm's IPO may be determined by the ownership structure

resulting from the selling strategy and the socio-cognitive nature of the negotiating agents (Stoughton and Zechner 1998; Mello and Parsons 1998). The main claim for the model is that it provides a link between speculations about offering mechanisms and observed patterns of IPOs and processes of going public.


The model focuses on the negotiation process between the seller of the firm and the potential investors. The seller may be either the owner of the firm or the representative of a company that owns the firm. In either case, he is expected to be motivated by the goal of getting the maximum amount of revenue from the sale of the firm. The basic premise of our argument is that investors are not homogeneous in their ability to monitor the newly public firm's management (Mello and Parsons 1998; Stoughton and Zechner 1998). More specifically, our model identifies two classes of investors. First, there is a population of small and passive shareholders with virtually no monitoring capabilities. Second, there are large and active investors who seek a controlling block in order to actively shape or control the firm's management. For simplicity, small and active investors are assumed to be risk neutral (Mello and Parsons 1998). Therefore, they only differ in their demands for the shares of the firm and, consequently, in their degree of influence they are expected to exert on the firm once they become shareholders. We abstract from the population of the small investors and consider only one passive investor who represents the pool to which he belongs³. Furthermore, for simplicity, we will concentrate on the negotiation process between the seller and *one* active investor, and therefore we will not address such issues as how a number of potentially interested active investors should be put into competition with one another, and how the seller should choose a potential active investor to negotiate with⁴.

Drawing on these assumptions, we will explore the impact that the heterogeneity between active and small investors may exert upon the negotiation process underpinning the sale of a firm, and we will examine to what extent this impact can ultimately shape the final ownership structure of the firm and the performance of the sale.

As shown in Figure 1, the model exists at two main interrelated levels. At the top, the individual level is shown for each of the three economic actors that may be involved in the sale of the firm - the seller, the active and the small investor. Here the negotiating agent is modelled as a cognitive agent capable of individual social behaviour. Correspondingly, the structure of the individual agent is organised into two interrelated levels: cognition and behaviour.

For an agent to be termed a *cognitive agent* it must be endowed with mental attitudes representing the world and motivating action (Panzarasa *et al.* 2001; Rao and Georgeff 1991; Singh 1995; Wooldridge 2000). Further, for a cognitive agent to be a *social cognitive agent* it must not only have an intentional stance towards the world, but also represent other agents as

Figure 1. The Model and its Components.



cognitive agents similarly endowed with mental attitudes for representational and motivational purposes (Castelfranchi 1998; Panzarasa *et al.* 2001). We capture and incorporate these properties of cognitive and social agenthood into our model by modelling the agent as having a *cognitive structure*. Crudely, according to mainstream cognitive science and expert systems research, the agent's cognitive structure can be defined as a set comprising basic mental attitudes, cognitive principles for modifying these attitudes, operators for adding mental attitudes, the cognitive frame, the language, control procedures, social interaction propensities, and principles and mechanisms for generating inter-agent social behaviour (Carley 1989; Davis and Lenat 1980; Dennet 1987; Hayes-Roth *et al.* 1983; Waterman 1986). The cognitive structure changes as the agent interacts with its physical and social environment and acquires new information. This, in turn, affects the agent's behaviour and its interaction with the environment (Carley 1989).

The two main components of each negotiating agent's cognitive structure are assumed to be same - the mental state and a set of principles of social behaviour. The agent's mental state is the set of the basic mental attitudes that are processed by the agent to undertake both theoretical reasoning (i.e., reasoning undertaken to determine what is the case) and practical reasoning (i.e., reasoning undertaken to determine what is to be done) (Panzarasa *et al.* 2001). These mental attitudes include beliefs (pieces of information about the physical and social environment), goals (what the agent wants to bring about), and intentions (what the agent is committed to bringing about) (Cohen and Levesque 1990; Kraus *et al.* 1998; Rao and Georgeff 1991; Singh 1995). Agents are assumed to differ in terms of the instantiation of these mental attitudes, that is, they are assumed to differ in terms of what they believe about the world, and what they want and intend to bring about (Section 3.1).

Principles of social behaviour include rules for *local* social interaction and higher-order rules for *global* social interaction. The former are procedures for generating microprocesses of inter-agent social behaviour at the individual level (e.g. rules for determining when to start and to exit a social relationship; rules for communication). The latter are high-level specifications of procedures for generating macroprocesses of inter-agent social behaviour at a group level (e.g. rules for determining how to manage a number of complex social relationships involving different agents; rules for determining whether or not to engage in simultaneous or sequential social relationships). Each agent is assumed to have both types of rules for social interaction. However, sub-components and values for these sub-components may be different for each agent.

Among its rules for local social interaction, each negotiating agent is modelled as having cognitively represented mechanisms for generating and controlling negotiation-oriented social behaviour. Drawing on these mechanisms, the agent determines what conditions need to be established before negotiation can proceed, how and when interaction

with others should take place, and when negotiation should terminate. The agent's mental state impacts upon these mechanisms by affecting their two main sub-components: the evaluation parameters (Section 3.2) and the negotiation protocol (Section 3.3). On the one hand, the evaluation parameters refer to the key data structures used by each negotiating agent in order to assess the messages received and produce new messages to be forwarded to other agents during negotiation. On the other, the negotiation protocol points to the behavioural procedures and methods that are instantiated with the evaluation parameters and then used by each agent in order to assess and generate messages. Both the evaluation parameters and the negotiation protocol are affected by what the agent believes about its physical and social environment, and by what it wants and intends to achieve. Further, the evaluation parameters and the negotiation protocol co-evolve by affecting each other in a two-way interaction. Precisely, co-evolution occurs as the evaluation parameters provide the specific values for instantiating the protocol and, in turn, the protocol provides the interaction procedures that determine what evaluation parameters are needed to engage in negotiation and what their structure should be.

The rules for global social interaction govern the agent's social behaviour at a group level. Negotiating agents differ in terms of the instantiation of these rules. Let us distinguish between the seller and the investor (either small or active). On the one hand, the seller is modelled as having a set of selling strategies (Section 3.4) that are affected by his beliefs about the world, and by his goals and intentions. These strategies can be seen as a set of high-level specifications of negotiation mechanisms that represent alternative methods for going public (Mello and Parsons 1998). Specific issues covered by these strategies are: what types of investors the seller should send an offer to; what type of investor should be contacted first; whether or not different types of investors should be contacted simultaneously or sequentially. Selling strategies affect and are affected by the seller's rules for local social interaction. First, evaluation parameters provide the data structures that will be used by the seller to select and use a selling strategy, whereas the selling strategy that is chosen and used to offer shares determines what evaluation parameters are needed. Second, the negotiation protocol provides the procedural mechanisms to be used within a selling strategy and therefore affects the choice of the most appropriate strategy, whereas the selling strategy chosen and used by the seller provides indications as to how offers and counter-offers should be evaluated during negotiation.

On the other hand, the investor (either small or active), drawing on his mental state, derives a set of investment strategies that govern his decision-making processes centred around his investment portfolio. Indicative issues covered by these strategies include: what shares should be bought and when; how and when to buy a controlling block; whether or not to rely on a tender offer market for the accumulation of a controlling block or for the

liquidation of an inefficiently large block (Mello and Parsons 1998); whether or not to trade shares in a secondary market. As with selling strategies, investment strategies affect and are affected by the investor's rules for local social interaction. Since we are mainly interested in studying the impact of different offering mechanisms on the going public process, in this paper we will not be concerned with agents' investment strategies and their role in the sale of a firm, and we will therefore limit our attention to the seller's global interaction rules (Section 3.4).

In the middle of Figure 1, the social behavioural dimension of the individual level is portrayed. Here, the agent is modelled as being capable of *individual social action* aimed at producing effects on its social environment. Specifically, for a cognitive social agent a social action is an action that is regulated by the agent's cognitive structure, and that deals with another agent who is believed to be a cognitive agent whose behaviour is regulated by a cognitive structure (Castelfranchi 1998). The agent's individual social behaviour rests on the agent's cognitive structure, in that it is controlled by the principles of social behaviour which, in turn, are affected by the agent's mental state. Thus, the information the agent gathers and maintains about the world as well as the states of affairs it wants and/or intends to bring about are the ultimate causal forces on the top of the cognitive chain governing the agent's interaction with its social environment.

So far the model has been couched in terms of the individual agent's cognition and behaviour. Individualistic analytical concepts (e.g. individual mental attitudes) and cognitive and behavioural microprocesses (e.g. the generation of behavioural rules and their use for controlling individual social behaviour) have been introduced that enable us to model some of the key properties of cognitive social agenthood. However, when observed from a higher-order multi-agent perspective, two or more social cognitive agents may be described in terms of their *cognitive and social relationships* in such a way that a new collective entity (e.g. a team, a group, an organisation) can be thought of as emerging from these relationships (Panzarasa *et al.* 1999). Researchers have paid much attention to the cognitive and behavioural properties of systems composed of multiple agents, in which both cognition and action are seen as distributed among the individual agents (Levesque *et al.* 1990; Wooldridge 2000). Along these lines, attempts have been made to extend the BDI model to a social setting in which collective behaviour is modelled as grounded on mutual beliefs, and motivated by joint goals, joint intentions and joint commitments (Castelfranchi 1995; Panzarasa *et al.* 2001). At this level, sociality is described in terms of a higher-order joint mental state and joint behavioural processes that emerge from and transcend the mental states and behaviours of cognitively and socially interconnected individual agents.

The higher-order collective level has been portrayed at the bottom of Figure 1. As shown, we have two possible groups with joint cognitive and behavioural capabilities: a

group comprising the seller and the active investor, and a group comprising the seller and the small investor. Both groups can be characterised in terms of their members' being cognitively and socially connected in an attempt to reach an agreement, via negotiation, about a share price. As with the individual agent, we have two fundamental components of the collective level: joint cognition and joint behaviour. On the one hand, the mental states of the members of these groups are meshed together in such a way that a joint mental state will ensue. This joint mental state may comprise varying joint doxastic, motivational and deontic mental attitudes. Among these, a key role in guiding collective activity is played by joint commitment (Panzarasa *et al.* 2001; Wooldridge 2000). More specifically, before negotiation can commence, a joint commitment must ensue among the members of each group to achieving a state of the world (i.e. the sale of the firm) and reaching an agreement as to how to obtain that state (Panzarasa *et al.* 2001). This ensures that the negotiating agents are persistent in their efforts towards the joint activity, which, in turn, guarantees that they do not exit the negotiation process without any valid reason for doing so (Levesque *et al.* 1990).

Joint commitment affects and is affected by the performance of joint behavioural processes. First, joint behavioural processes find their roots in joint commitment, as they are generated by two or more jointly committed agents' performing individual social actions. Second, the performance of joint behavioural processes impacts back onto joint commitment by affecting the jointly committed agents' mental states. A fundamental property of joint behavioural processes is that they are executed by agents who communicate with and try to exert social influence upon each other (Panzarasa *et al.* 2001; Wooldridge 2000). At this level, both the seller and the investor send messages to each other with the purpose of impacting upon each other's mental state and generating an agreement about a share price. This gives rise to a joint activity that we propose to call "negotiation cycle" (Section 3.5). More specifically, a negotiation cycle can be defined as the sequence of the computational and behavioural processes into which the social interaction between two or more negotiating agents can be decomposed (Kraus *et al.* 1998). Since the negotiation cycle stems from the agents' individual social behaviours, it ultimately rests on the agents' cognitive structures, and more specifically on the agents' doxastic, motivational and deontic attitudes and their principles of social behaviour. Further, the negotiation cycle impacts back onto the agents' cognitive structures by affecting their mental attitudes and, through these, their local and global rules for social interaction. As the agent engages in the negotiation process, it may acquire new beliefs, fulfil its goals and intentions, or modify some of them. Experiential wisdom accumulates as a result of positive and negative reinforcement of prior mental attitudes and rules for local and global interaction (Carley 1989; Panzarasa *et al.* 2001). Mental attitudes and interaction rules that have led to what are encoded as positive outcomes are reinforced, while those that have led to negative outcomes are modified or discharged.

Finally, as a result of its impact upon the agent's mental state, the negotiation cycle affects joint commitment as this is cognitively grounded on the individual agents' mental attitudes. For example, failure to reach an agreement may lead the active investor to generate the belief that no agreement can reasonably be made with the seller of the firm. Or, negotiation may lead the investor to drop his motivation to buy shares in the firm or to adopt an intention that is incompatible with the purchase of these shares (e.g., to make other investments). In all these circumstances, the active investor will exit negotiation and the joint commitment he maintained with the seller will be dropped.

Drawing all the above points together (Figure 1), the essential components of our model can be seen to be (i) the agent's mental state; (ii) the evaluation parameters; (iii) the negotiation protocol; (iv) a set of high-level selling strategies; and (v) the negotiation cycle. In what follows, each of these components will be dealt with in turn.

3.1 The Agent's Mental State

As shown in Figure 1, the structure of the agent's mental state is organised into three mental attitudes: beliefs, goals and intentions. Agents are assumed to have the same structure of their mental states, but they differ in terms of the instantiation of the mental attitudes. In this section, we will provide a specification of content of the mental attitudes for each of the three agent types specified by our model, i.e., the small investor, the active investor, and the seller.

The Small Investor

For the small investor, i_s , the value of a share in the firm can be modelled as the sum of two components. One component is related to the "book value", BV , of the firm. BV reflects the turnover and assets of the firm calculated under a common metric, and is information in the public domain about which there is no uncertainty. Added to this common component is an idiosyncratic component, IC_s , that is private information to the small investor and that, together with BV , determines the actual price the small investor is prepared to pay for shares in the firm. IC_s may reflect, for example, the small investor's tax status or liquidity preference, which, in turn, might affect the small investor's valuation of the firm's expected cash flows (Mello and Parsons 1998).

Against this background, the small investor's mental state can be modelled as comprising a belief about the idiosyncratic component, as well as a belief about the book value of the firm:

$$Bel(i_s, the_value_of_IC_s_is_...);$$

Bel (i_s, the_value_of_BV_is_...).

The small investor's goal is assumed to be the maximisation of his expected utility by purchasing shares of the firm at the lowest possible price:

Goal (i_s, purchase_shares_at_lowest_price).

In order to have his goal fulfilled, the small investor maintains the intention to negotiate with the seller and, more specifically, to receive an offer from the seller and to evaluate it:

Int(i_s, get_ready_for_message_from_seller);

Int(i_s, evaluate_message_from_seller).

The Active Investor

The active investor, i_a , seeks a controlling block in order to actively influence or monitor future management's decisions. This investor might be interested in acquiring control because of information about a strategy for using the assets of the firm that could increase the value of the firm's cash flows. As with the small investor, the actual price that the active investor is prepared to pay for shares in the firm will vary from the book value by an idiosyncratic component, IC_a . In turn, IC_a can be split into two components: a "discount", Dis , and a "premium", Pre . On the one hand, the active investor, upon achieving control of the firm, will be in a position to implement changes in future management's decisions that may increase the value of the firm's expected cash flows. Since all shareholders will benefit from these changes, the seller can discriminate in favour of the active investor by offering the controlling block at a discount (based upon the public benefits of control) from the price paid by small investors (Mello and Parsons 1998). On the other, as long as the active investor can use the controlling block to extract private benefits, the seller can raise the price at which a controlling block is offered and, as a result, the active investor will pay a premium over the price paid by small investors who do not obtain the private benefits of control. Whether the final price at which the controlling block is offered reflects a (net) discount or premium depends upon the relative significance of the public and private benefits associated with the controlling block.

In the light of these observations, the active investor's mental state will include the following mental attitudes. First, he will hold a belief about his control discount and premium, as well as a belief about the book value of the firm:

Bel (i_a, the_value_of_Dis_is...);

Bel (i_a, the_value_of_Pre_is...);

Bel (i_a, the_value_of_BV_is...).

Second, the active investor's goal is assumed to be the maximisation of his expected utility by purchasing a controlling fraction m of the shares of the firm at the lowest possible price:

Goal (i_a, purchase_m_of_shares_at_lowest_price).

where m represents the fraction of shares that provides the active investor with a controlling interest⁵.

Third, as with the small investor, in order to have his goal fulfilled, the active investor will intend to receive an offer from the seller and to evaluate it:

Int(i_a, get_ready_for_message_from_seller);

Int(i_a, evaluate_message_from_seller).

The Seller

There may be a number of reasons that could explain the decision of selling a firm, including liquidity preferences, the realisation of gains from selling to better-positioned parties, exploiting favourable market conditions, gains from focus, and so forth (Stoughton and Zechner 1998; Zingales 1995). In what follows, we take the decision to put the firm on the market as given and concentrate on the issues surrounding the implementation of the sale. We assume that the seller of the firm, s , does not have complete information about the state of the market for dispersed shares. Nor does he have complete information about the attitude of potential active investors toward his firm. This implies that the seller can only take a guess at the small investors' idiosyncratic components. Further, he will also have to guess at the discount that should be offered to an active investor, and at the premium an active investor should pay for the privilege of being involved in the firm. The seller's incomplete information about both the small and active investors will be operationalised by assuming that the seller takes the small investor's idiosyncratic component and the active investor's control premium and discount to be random variables with normal density probability distribution.

Against this background, the seller's mental state can be modelled in the following way. First, he will maintain beliefs about: the mean and standard deviation of the active investor's control premium; the mean and standard deviation of the active investor's control

discount; the mean and standard deviation of the small investor's idiosyncratic component; and the book value of the firm. Formally, we have:

Bel (s, the_value_of_μ_{Pre}_is_...);
Bel (s, the_value_of_σ_{Pre}_is_...);
Bel (s, the_value_of_μ_{Dis}_is_...);
Bel (s, the_value_of_σ_{Dis}_is_...);
Bel (s, the_value_of_μ_{ICs}_is_...);
Bel (s, the_value_of_σ_{ICs}_is_...);
Bel (s, the_value_of_BV_is_...).

Second, the seller's goal is assumed to be to release the maximum amount of revenue from the sale of the firm:

Goal (s, release_maximum_revenue_from_sale_of_shares).

Third, in order to have this goal fulfilled, the seller maintains the intention to engage in negotiation with either small investors or an active investor. Specifically, during negotiation he holds the intention to: (i) receive a message from the active investor and evaluate it; and (ii) receive a message from the small investor and evaluate it. We have:

Int(s, get_ready_for_message_from_active_investor);
Int(s, evaluate_message_from_active_investor);
Int(s, get_ready_for_message_from_small_investor);
Int(s, evaluate_message_from_small_investor).

3.2 Evaluation Parameters

During negotiation, both the seller and the investors typically receive, generate, and send messages. In order to successfully engage in such activities, the agents need to refer to a set of rules for local inter-agent social behaviour (see Figure 1). These rules provide the agents with a negotiation mechanism that enables them to evaluate the messages received and to generate the messages to be forwarded to other agents. More specifically, such a mechanism will consist of the following two main components: (i) a set of evaluation parameters; and (ii) a procedure for generating messages. In this section, we will concentrate on the evaluation parameters, and we will provide a specification of their structure and content for each of the

three agent types here considered. Section 3.3 will focus on the negotiation protocol that provides the agents with the procedures for generating messages during negotiation.

The Small Investor

In order to establish a share price with the seller, the small investor evaluates proposals and counter-proposals based on a range of acceptable prices. This range is limited by a maximum value per share that the small investor is willing to pay, $MaxP_{is}$, and a minimum price, $MinP_{is}$, that represents the first counter-offer that the small investor will make to the seller upon receiving an unacceptable offer from him. We take these two values to represent the small investor's evaluation parameters used during negotiation. We assume that $MaxP_{is}$ is given by the sum of the book value of the firm and the small investor's idiosyncratic component. Thus,

$$MaxP_{is} = BV + IC_s$$

On the other hand, since $MinP_{is}$ is supposed to represent the minimum price that can be *reasonably* counter-offered, it will be taken to reflect the information that is publicly available and about which there is no uncertainty, i.e., the firm's book value:

$$MinP_{is} = BV$$

The Active Investor

As with the small investor, the active investor is guided by a price range limited by a maximum value that he is willing to pay, $MaxP_{ia}$, and a minimum price that represents his first counter-offer to the seller, $MinP_{ia}$. We assume that $MaxP_{ia}$ depends on the book value of the firm and the difference between the active investor's control premium and discount. That is,

$$MaxP_{ia} = BV + Pre - Dis$$

On the other hand, $MinP_{ia}$ will be taken to reflect the firm's book value:

$$MinP_{ia} = BV$$

The Seller

Finally, we have to determine the maximum and minimum values that limit the seller's price range used during negotiation. The maximum value represents the maximum price that the seller can reasonably ask, whereas the minimum value is the lowest price he will be willing to accept. As expected, both values depend upon whether the seller negotiates with the active or the small investor. Therefore, in order to determine the seller's price range, we need to distinguish between two cases: (i) negotiation with the active investor; and (ii) negotiation with the small investor.

A) Negotiation with the Active Investor

As mentioned in Section 1, an active investor can be put into competition with the small investors by putting a controlling block for sale on terms that are contingent on the market price paid by the small investors (Mello and Parsons 1998). In the light of this, in order to determine the seller's price range, we need to further distinguish between two cases: (i) the negotiation with the active investor takes place when a number of shares have already been sold to the small investors; and (ii) the negotiation with the active investor takes place with no prior knowledge of the market.

In the former case, the maximum and minimum limits of the seller's price range are given as follows. The maximum price he may ask during negotiation, $MaxP_{s/ia}$, depends on the book value of the firm as well as on a subjective component, $MaxIC_a$, reflecting the seller's belief about what the maximum value of the active investor's idiosyncratic component might be. Further, the impact of $MaxIC_a$ on $MaxP_{s/ia}$ will be taken to depend on the ratio between the (market) price at which shares have already been sold to the small investor, $P_{s/is}$, and the book value of the firm. This condition implies that as the market price increases (decreases) with respect to the book value, the seller will increase (decrease) the maximum price that he may be willing to ask for a controlling block. The reason for this is two-fold. First, as the small investor buys shares at a market price that exceeds (is lower than) the book value, a new piece of information becomes available in the public domain conveying a new higher (lower) value for the shares in the firm. This will enhance (depress) the market sentiment and, consequently, the active investor's willingness to accept a higher price for the shares will become stronger (weaker). The probability for the seller to extract more surplus from the active investor will therefore increase (decrease), which ultimately triggers an increase (decrease) in the maximum price the seller may be willing to ask for a controlling block (Mello and Parsons 1998). Second, as a result of the improved (depressed) market sentiment, the seller becomes more (less) confident in his negotiation power and interaction

abilities. Again, this leads the seller to increase (decrease) the maximum price that may be asked for a controlling block by increasing (decreasing) the impact of the subjective component (i.e., $MaxIC_a$) on the pricing of the shares. Formally, we have:

$$MaxP_{s/ia} = BV + (MaxIC_a (P_{s/is}/BV))$$

In turn, we assume that $MaxIC_a$ depends on the seller's beliefs about the active investor's control premium and discount. These beliefs are inherently uncertain as both the control premium and discount are private information only to the active investor. The seller will have to guess at the density probability distributions of these two (random) values. As we assumed that both values are normally distributed (Section 3.1), $MaxIC_a$ will ultimately depend on (the seller's guess at) the means and standard deviations of the two probability distributions. Specifically,

$$MaxIC_a = (\mu_{Pre} + b\sigma_{Pre}) - (\mu_{Dis} - b\sigma_{Dis})$$

where b is a constant factor that reflects the sensitivity of the seller's evaluation parameters to changes in the degree of accuracy of his cognitive representations.

On the other hand, the minimum price at which the seller is willing to sell shares to the active investor, $MinP_{s/ia}$, depends on the book value of the firm as well as on a subjective component, $MinIC_a$, reflecting the seller's belief about what the minimum value of the active investor's idiosyncratic component might be. Again, the impact of $MinIC_a$ on $MinP_{s/ia}$ is assumed to depend on the ratio between the market price at which shares have already been sold to the small investor, $P_{s/is}$, and the book value of the firm. As with $MaxIC_a$, we assume that $MinIC_a$ depends on (the seller's guess at) the means and standard deviations of the active investor's control premium and discount. Also, $MinIC_a$ depends on a constant sensitivity factor b defined as above. Finally, we assume that in no circumstances will the seller accept a price that is lower than the book value of the firm. Thus, we have:

$$MinP_{s/ia} = \begin{cases} BV + (MinIC_a (P_{s/is}/BV)) & \text{if } [BV + (MinIC_a (P_{s/is}/BV))] > BV \\ BV & \text{otherwise,} \end{cases}$$

where

$$MinP_{s/ia} = BV + (MinIC_a (P_{s/is}/BV))$$

When there is no knowledge about the market, that is, no negotiation has already taken place with the small investor, the seller's price range will be determined in a

straightforward way. The maximum price he may happen to ask, $MaxP^*_{s/ia}$, will be:

$$MaxP^*_{s/ia} = BV + MaxIC_a,$$

where $MaxIC_a$ is defined as above.

On the other hand, the minimum price the seller will accept from the active investor, $MinP^*_{s/ia}$, will be:

$$MinP^*_{s/ia} = BV + MinIC_a,$$

where $MinIC_a$ is defined as above.

B) Negotiation with the Small Investor

As with the active investor, the seller needs to evaluate the small investor's proposals and counter-proposals based on a price range. Also, the small investor can be put into competition with the active investor by putting dispersed holdings for sale on terms that are contingent on the share price paid by the active investor (Mello and Parsons 1998). Therefore, in order to determine the seller's price range, we need to distinguish between two cases: (i) the negotiation with the small investor takes place when shares have already been sold to the active investor; and (ii) the negotiation with the small investor takes place with no prior knowledge of the market.

In the former case, the maximum price the seller may happen to ask during negotiation, $MaxP_{s/is}$, depends on the book value of the firm as well as on a subjective component, $MaxIC_s$, reflecting the seller's belief about what the maximum value of the small investor's idiosyncratic component might be. The impact of $MaxIC_s$ on $MaxP_{s/is}$ will be taken to depend on the ratio between the price at which shares have already been sold to the active investor, $P_{s/ia}$, and the book value of the firm. This condition implies that as the price paid by the active investor increases (decreases) with respect to the book value, the seller will increase (decrease) the maximum price that he may be willing to ask for dispersed shareholdings. As before, the reason for this is two-fold. First, as the active investor buys a controlling block at a price that exceeds (is lower than) the book value, a new piece of information becomes available in the public domain that will eventually enhance (depress) the market sentiment. As a result of this, the active investor's willingness to accept a higher price for the shares will become stronger (weaker), which ultimately leads the seller to increase (decrease) the maximum price he may be willing to ask for dispersed shareholdings (Mello and Parsons 1998). Second, as the price paid for a controlling block is higher (lower)

than the book value, the seller becomes more (less) confident in his negotiation power and interaction abilities, and he will therefore be prone to increase (decrease) the maximum price that may be asked for dispersed shares. Formally, we have:

$$MaxP_{s/is} = BV + (MaxIC_s (P_{s/ia}/BV)),$$

where

$$MaxIC_s = (\mu_{IC_s} + b\sigma_{IC_s}).$$

In turn, $MaxIC_a$ depends on (the seller's beliefs about) the small investor's idiosyncratic component. Again, as we assumed that this component is normally distributed around its mean (Section 3.1), $MaxIC_a$ will ultimately depend on (the seller's guess at) the mean and standard deviation of the normal density probability distribution of the small investor's idiosyncratic component. Specifically,

$$MaxIC_s = (\mu_{IC_s} + b\sigma_{IC_s}),$$

where b is defined as above.

Similarly, the minimum price the seller is willing to accept from the small investor, $MinP_{s/is}$, is given by:

$$MinP_{s/is} = BV + (MinIC_s (P_{s/ia}/BV)), \text{ where}$$

$$MinP_{s/is} = \begin{cases} BV + (MinIC_s (P_{s/ia}/BV)) & \text{if } [BV + (MinIC_s (P_{s/ia}/BV))] > BV \\ BV & \text{otherwise,} \end{cases}$$

where

$$MinIC_s = (\mu_{IC_s} - b\sigma_{IC_s}).$$

On the other hand, where the seller has no knowledge about the market (i.e., no negotiation process has already taken place between the seller and the active investor), the maximum price the seller may ask, $MaxP^*_{s/is}$, and the minimum price he will be willing to accept from the small investor, $MinP^*_{s/is}$, will be respectively:

$$MaxP^*_{s/is} = BV + MaxIC_s, \text{ and}$$

$$MinP^*_{s/is} = BV + MinIC_s,$$

where $MaxIC_s$ and $MinIC_s$ are defined as above.

3.3 Negotiation Protocol

While the evaluation parameters provide the agents with the key data structure for assessing and generating messages, the negotiation protocol can be seen as the set of behavioural procedures that need to be instantiated with the evaluation parameters before the agents can engage in negotiation. In this section, we will outline the basic structure of the negotiation protocol and, where appropriate, we will provide a more detailed specification of its content for the three agent types here considered.

During the negotiation underpinning the sale of shares in a firm, the agents typically send offers and counter-offers regarding the share price. Upon receiving an offer, an agent will evaluate it. An offer can be either accepted or modified and sent back. Ideally, negotiation can continue until an agreement is made. However, in most cases it is not reasonable to assume that negotiation will continue for ever. Although persistent in keeping their joint commitment to negotiating, the agents are not fanatical and, after a certain number of messages, they will give up their joint activity (Levesque *et al.* 1990; Panzarasa *et al.* 2001). In such a case, negotiation will fail and no final agreement will ensue. Given this, we need to distinguish between three different types of response that the agents may give after receiving a message: (i) acceptance; (ii) modification; and (iii) rejection. In what follows, we will focus on the negotiation procedures that the agents use to generate each of these responses.

Acceptance

We need to distinguish between the following two situations: (i) the seller sends an offer to the small or active investor; and (ii) the small or active investor sends an offer to the seller. In the former case, if the offered price, $p_{s/is}$ or $p_{s/ia}$, is lower than the maximum price the investor is willing to pay, $MaxP_{is}$ or $MaxP_{ia}$, then the offer will be accepted⁶. The small or active investor will thus generate the intention to buy shares at that price:

$$Int(i_s, buy_shares_at_p_{s/is}=P_{s/is});$$

$$Int(i_a, buy_shares_at_p_{s/ia}=P_{s/ia}).$$

On the other hand, when an investor sends an offer to the seller, if the offered price, $p_{is/s}$ or $p_{ia/s}$, is greater than the minimum price the seller is willing to accept, then the offer will

be accepted. As a result, the seller will generate the intention to sell shares to the small or active investor at the offered price:

$$\begin{aligned} &Int(s, sell_shares_at_p_{is/s} = P_{s/is}); \\ &Int(s, sell_shares_at_p_{ia/s} = P_{s/ia}). \end{aligned}$$

Modification

If the price offered is greater than the maximum the investors are prepared to pay or is lower than the minimum the seller is willing to accept, then it will not be accepted. Offers, if unacceptable, will be modified and counter-offers will be sent back to the other part. As already mentioned above, since we assume the agents are not fanatical, this process will not continue for ever and modifications will not be infinite. Specifically, we assume that negotiation will continue if: (i) an acceptable offer has not been made yet; and (ii) the number of messages (offers and counter-offers) between the agents is less than an integer value n^* .

Modifications of messages are made according to a *price movement strategy*. Each price movement strategy depends on the interplay between the offer sent originally and the modification received. Let us distinguish between two cases.

Firstly, should the seller receive an unacceptable offer from an investor at time t' , $p'_{is/s}$ or $p'_{ia/s}$, he will modify it and send back a new counter-offer, $p^t_{s/is}$ or $p^t_{s/ia}$, at time t ($t > t'$). Specifically,

$$\begin{aligned} p^t_{s/is} = p''_{s/is} - c(p''_{s/is} - p'_{is/s}) &\text{ if } [p''_{s/is} - c(p''_{s/is} - p'_{is/s})] > MinP^*_{s/is} \text{ (or } MinP_{s/is}) \\ &MinP^*_{s/is} \text{ (or } MinP_{s/is}) \text{ otherwise, and} \end{aligned}$$

$$\begin{aligned} p^t_{s/ia} = p''_{s/ia} - c(p''_{s/ia} - p'_{ia/s}) &\text{ if } [p''_{s/ia} - c(p''_{s/ia} - p'_{ia/s})] > MinP^*_{s/ia} \text{ (or } MinP_{s/ia}) \\ &MinP^*_{s/ia} \text{ (or } MinP_{s/ia}) \text{ otherwise,} \end{aligned}$$

where $p''_{s/is}$ and $p''_{s/ia}$ are the offers sent at time t'' ($t'' < t' < t$) by the seller to the small and active investor respectively, and c ($0 \leq c \leq 1$) is a concession rate reflecting the seller's price movement strategy. Thus, if the modification turns out to be lower than the minimum price that the seller is willing to accept, then the seller will offer back his minimum price. Otherwise, he will send back the modification.

Secondly, should an investor receive at time t' an unacceptable offer from the seller, $p^t_{s/is}$ or $p^t_{s/ia}$, he will modify the offer received in the following way:

$$p_{is/s}^t = \begin{cases} p_{is/s}^{t''} - c(p_{is/s}^{t''} - p_{s/is}^t) & \text{if } [p_{is/s}^{t''} - c(p_{is/s}^{t''} - p_{s/is}^t)] < MaxP_{is} \\ MaxP_{is} & \text{otherwise, and} \end{cases}$$

$$p_{ia/s}^t = \begin{cases} p_{ia/s}^{t''} - c(p_{ia/s}^{t''} - p_{s/ia}^t) & \text{if } [p_{ia/s}^{t''} - c(p_{ia/s}^{t''} - p_{s/ia}^t)] < MaxP_{ia} \\ MaxP_{ia} & \text{otherwise,} \end{cases}$$

where $p_{is/s}^{t''}$ and $p_{ia/s}^{t''}$ are the offers sent at time $t'' < t' < t$ to the seller by the small and active investor respectively, and c is a concession rate reflecting the investor's price movement strategy. Similarly, if the modification is greater than the maximum price that the investor is willing to accept, then he will counter-offer his maximum price. Otherwise, he will send back the modification.

Rejection

As the agents are taken to be persistent but not fanatical, after a certain number n^* of offers and counter-offers they will give up negotiation. In such situations, no final joint agreement will ensue. Specifically, if the seller receives from an investor an unacceptable offer, $p_{is/s}$ or $p_{ia/s}$ and if the number n of messages received is greater than n^* , then he will reject the offer and form the intention to exit negotiation:

$$Int(s, \text{exit_negotiation}).$$

The rejection and exit conditions corresponding to the small and active investors are similarly intuitive.

3.4 Selling Strategies for Going Public

Much of the mainstream literature on IPOs and staged equity financing emphasises that the choice of method of sale influences whether a transfer of control will take place, which in turn affects the proceeds from the sale and how the surplus is divided between the seller and the various interested buyers (Milgrom and Weber 1982; Barclay and Holderness 1989; Brennan and Franks 1995). As shown in Figure 1, selling strategies for going public represent instantiations of rules for global social interaction that are developed and used by the seller in order to successfully manage the complex web of social relationships typically involved in the sale of a firm. Based on the basic framework of analysis developed by Mello and Parsons (1998), this section will introduce the following four alternative selling strategies for going public that have been most used in countries with developed capital markets:

- (i) a sequential sale in which the controlling block is sold before the IPO (strategy A);
- (ii) a sequential sale in which the IPO takes place before the sale of the controlling block (strategy B);
- (iii) a single-period parallel sale in which shares are offered both to small investors and to the active investor at the same time (strategy C);
- (iv) a public offering in which all shares are sold to small investors at a uniform price, without involving an active investor (strategy D).

In what follows, each selling strategy will be discussed and formalised in terms of our framework of analysis.

Selling the Controlling Block before the IPO (Strategy A)

The seller of the firm interacts with two types of potential investors: the active investor, and the small investor. The first task within this strategy is to establish a price for the sale of shares to the active investor. If this stage is successful, both the seller and the active investor will update their mental states by generating the belief that a share price has been established and the intention to sell/buy shares at that price. For example, should price $p_{s/ia}$ be accepted by both parties, the following mental attitudes would be generated:

$$\begin{aligned}
 &Bel(s, P_{s/as}=p_{s/ia}); \\
 &Bel(i_a, P_{s/ia}=p_{s/ia}); \\
 &Int(s, sell_shares_at_P_{s/ia}=p_{s/ia}); \\
 &Int(i_a, buy_shares_at_P_{s/ia}=p_{s/ia}).
 \end{aligned}$$

As a result, the seller will then proceed to sell shares at $p_{s/ia}$. Negotiation with the active investor may well be unsuccessful and no price may be agreed upon. In such a case, no transaction will take place between the two parties.

After negotiating with the active investor, the seller then proceeds to establish a share price with the small investor. Negotiation with the small investor is made contingent to negotiation with the active investor as the seller's evaluation parameters depend on whether a successful negotiation has previously taken place (see Section 3.2). If the seller is successful in conducting negotiation with the small investor, then a price will be agreed upon, and both the seller and the small investor will update their mental states as above. As a result, an economic transaction will occur. Otherwise, should negotiation with the small investor fail, no ownership transfer would take place.

Previous research suggests that selling the controlling block first avoids the free rider problem (Grossman and Hart, 1980) and can also help reduce the winner's curse problem faced by small investors (Rock, 1986). Most importantly, this strategy might have the following advantage. By selling shares using a mixed offer in which a tender offer is made first to the large investor and the tender price is then used to set the fixed price in the offer for sale to small investors, the seller might extract surplus from small investors. Along these lines, this strategy has been advocated by Stoughton and Zechner (1998) as their main prescription. While recognising that simply selling a fraction of the company may not always be an optimal mechanism, Stoughton and Zechner highlight that with the addition of some clauses and by following the right sequence of offerings the seller can adjust the ownership structure to obtain a higher sale price.

Since in Stoughton and Zechner's model there is no uncertainty about the demand of small investors, there is no loss from selling to the active investor prior to learning about aggregate demand and estimates of value. However, as long as knowledge about passive investors' demand is not perfectly accurate, selling the controlling block before the IPO may have the following disadvantage. By selling the controlling block first, the active investor's allocation cannot be contingent on the parameter of aggregate demand by the small investors, and therefore the seller cannot extract surplus from the active investor (Mello and Parsons 1998). Conversely, selling some amount of non-controlling shares first allows the seller to obtain information about the market for dispersed shares on which the price of the controlling block can be made contingent. In this case, although the sale occurs in stages and the active investor never directly bids against the small investors, competition between them may arise through the conditions imposed on the sale to the large investor based on the results of the first sale to the small investors.

In the light of these observations, a number of open research questions remain: To what extent can the potential benefits of selling a controlling block prior to selling dispersed shareholdings be greater than the concomitant drawbacks mentioned above? Are the net benefits of this selling strategy greater than the net benefits associated with selling control after the sale of dispersed allotments? We will address such issues by running a number of virtual experiments premised on our computational model (see Section 4).

Selling the Controlling Block after the IPO (Strategy B)

As with the previous selling strategy, the seller of the firm interacts with two types of potential investors: the small investor, and the active investor. The first task within this strategy is to establish a price for the sale of shares to the small investor. If this stage is successful, both the seller and the small investor will update their mental states by generating

the belief that a market share price has been established and the intention to sell/buy shares at that price. A transaction will then take place and, as a result, a non-controlling fraction of the firm will be transferred from the seller to the new owners. Otherwise, should negotiation with the small investor be unsuccessful, no ownership transfer will occur.

The seller then proceeds to establish a share price with the active investor. Again, selling a controlling block is made contingent on previous economic transactions as the seller's evaluation parameters during negotiation with the active investor depend on whether or not a successful negotiation with the small investor has previously taken place. If negotiation with the active investor turns out to be successful and a price for a controlling allotment is agreed upon, both the seller and active investor will update their mental states in a similar manner as before. A transaction will then take place. Conversely, should negotiation be unsuccessful, no transfer of control will occur.

The view that the IPO of shares may not be an isolated step but the first stage of a more elaborate process for selling shares seems to be confirmed by the existing empirical evidence. Barry *et al.* (1990) report that IPOs represent the most frequently used method of selling a fraction of the shares by venture capital firms and that during this first transaction the equity holdings of venture capital investors do not change much. Later on these investors sell a significant portion of their stakes, either to another active investor or to another company or through a follow-on offering. Control turnover subsequent to the IPO is also found by Holderness and Sheehan (1988) for the U.S., Rydqvist and Högholm (1994) for Sweden and the U.K., and Pagano *et al.* (1998) for Italy. This also seems to be the case in many equity carve-outs, according to Schipper and Smith (1986) and Klein *et al.* (1991).

Single-Period Parallel Sale (Strategy C)

The seller interacts *simultaneously* with the active investor and the small investor. Two prices therefore need to be established for the sale of shares to both investors, $P_{s/as}$ and $P_{s/ai}$. If negotiation is successful, each agent will update its mental state by generating the belief that a price has been agreed upon. We assume that the large investor gets control only with a bid that is higher than the bid of the small investor competing for the shares (Mello and Parsons 1998). Specifically, if the price established with the active investor, $p_{s/as}$, is greater than the market price for dispersed shareholdings, $p_{s/ai}$, then the seller proceeds to sell the controlling block, whereas the remaining allotment of shares will be sold to the small investors. On the other hand, if the price for dispersed shares is greater than the price for the controlling block, then the seller will sell all shares to the small investors. In terms of the seller's mental state we have:

$$\begin{aligned}
&Int(s, \text{sell_}m\text{_of_shares_to_active_investor_at_}P_{s/ia}=p_{s/ia}) \wedge \\
&Int(s, \text{sell_}(1-m)\text{_of_shares_to_small_investor_at_}P_{s/is}=p_{s/is}) && \text{if } p_{s/ia} > p_{s/is} \\
&Int(s, \text{sell_all_shares_to_small_investor_at_}P_{s/is}=p_{s/is}) && \text{otherwise}
\end{aligned}$$

where m represents the fraction of shares that provides the active investor with a controlling interest.

Perhaps because of the elaborate procedures it involves, a simultaneous sale of shares with a discriminatory allocation of a potentially controlling block is not very common. However, Brennan and Franks (1995) refer to a method that has recently been experimented within the U.K. and Australia and has features similar to what we are describing here. This method combines a private placement, targeting large investors, with a simultaneous public offering. As suggested by Brennan and Franks (1995), the effectiveness of this method depends on the allocation and pricing rules used for investors of different types. In addition, whenever it has been used, this method usually includes a clawback provision (Mello and Parsons 1998). This makes the allocation to a large investor dependent on the demand by small investors, therefore creating competition among investors of different types.

A Public Offering of All Shares at a Uniform Price (Strategy D)

As opposed to the selling strategies outlined above, where the seller interacts with both the small and the active investor in an attempt to establish two possibly different prices, in a public offering of all shares at a single fixed price, the seller interacts only with one type of investor, i.e., the small investor. Negotiation is therefore aimed at establishing a market price for dispersed shareholdings. If successfully conducted, negotiation will eventually lead the seller and the small investor to update their mental states in the usual way. A transaction will then occur. We assume that if a price for dispersed shares is established, all shares are sold at that price to the small investors.

An example of this selling strategy is the offer for sale by tender, employed by the U.K. and in France, where investors place bids for the shares indicating both quantity and price (Mello and Parsons 1998). After the bids have been received, a single share price is established and all investors buy shares at that price. As opposed to this strategy, where the seller is unable to discriminate among the small investors, it may be interesting to consider the sealed-bid discriminatory auction. This is a sale in which the investors tender bids and those bidding the highest buy shares at the price bid. There are a number of ways in which discrimination can be accomplished in an IPO. For example, empirical evidence suggests that in the U.S. the allocation of oversubscribed issues is a means for discriminating among

buyers (Hanley and Wilhelm 1995). Evidence of discrimination in the U.K. market has also been found (Brennan and Franks 1995). For the sake simplicity, in our model we abstract from specific discriminatory practices and allocation rules, and consider only IPOs in which shares must be offered at a uniform price.

3.5 The Negotiation Cycle

So far, we have glossed over the problem of exactly how an agent's cognitive structure might generate the agent's individual social behaviour (see Figure 1). More specifically, we shed some light on the role that the agent's mental attitudes (Section 3.1) and principles for social behaviour (Sections 3.2, 3.3, and 3.4) have in determining how the agent might go about negotiating a share price. In this section, we will move from the individual level to a higher-order multi-agent level. As shown in Figure 1, this collective level is organised into two components: a joint mental state and joint behavioural processes. Within the former cognitive component, joint commitment represents the key joint mental attitude that triggers and governs subsequent collective behaviour. For reasons of space, we do not consider the process of generating a joint commitment nor do we explore the different views that have been proposed in the literature about the complex web of individual and joint mental attitudes on which a joint commitment rests (Castelfranchi 1995; Levesque *et al.* 1990; Panzarasa *et al.* 2001). Our focus in this section will be limited to joint behavioural processes, and in particular to what we call the negotiation cycle.

By negotiation cycle we mean the sequence of the computational and behavioural processes into which the social interaction between two or more negotiating agents can be decomposed (Kraus *et al.* 1998). As shown in Figure 1, since the negotiation cycle stems from the agents' individual social behaviours, which in turn rest on the agents' cognitive structures, it ultimately can be explained in terms of the agents' doxastic, motivational and deontic attitudes as well as their local and global rules for social interaction. Drawing on Kraus *et al.* (1998), we assume that the negotiation cycle includes the following functional components: (i) generating a message; (ii) sending that message; (iii) receiving incoming messages; (iv) evaluating incoming messages; (v) responding to incoming messages; and (vi) continuing from step (i).

Based on its beliefs, the negotiating agent will attempt to achieve its (top-level) goals by fulfilling its own intentions (see Section 3.1). In turn, fulfilment of these intentions requires communication with other agents. The agent will typically send messages to others, and wait for a response. We assume that the agent will remember what message has been sent and who it has been forwarded to⁷. Depending on the response received, negotiation may fail (if the counter-part rejects the message), or an agreement may ensue (if the counter-part

accepts the message), or negotiation may continue (if the counter-part modifies the message). Should negotiation continue, the agent will receive a counter-offer, evaluate it, and, if appropriate, generate a new modified message to be forwarded. This cycle will continue until either an agreement is reached or one of the negotiating agents exits the process⁸.

Figure 2 is intended to portray the main functional components of the negotiation cycle and their interrelationships. Our account of the negotiation cycle is premised on the assumption that investors do not negotiate with each other; they only negotiate with the seller. Negotiation begins with the seller asking an investor (either small or active) the maximum share price he can ask, i.e., the upper value of his price range (Section 3.2).

Upon receiving the offer, the investor will evaluate it and respond in one of the following ways:

- (i) The investor may accept the seller's offer. As a result, he will update his beliefs and commit itself to act accordingly. Further, he will send a message of acceptance back to the seller. In turn, upon receiving this message, the seller will update his beliefs and commit itself to act accordingly. For example, should the small investor accept price $p_{s/is}$, he will generate the belief that the agreed-upon price is $p_{s/is}$, and form the intention to buy shares at that price:

$$\begin{aligned} &Bel(i_s, P_{s/ia}=p_{s/is}); \\ &Int(i_s, buy_shares_at_P_{s/is}=p_{s/is}). \end{aligned}$$

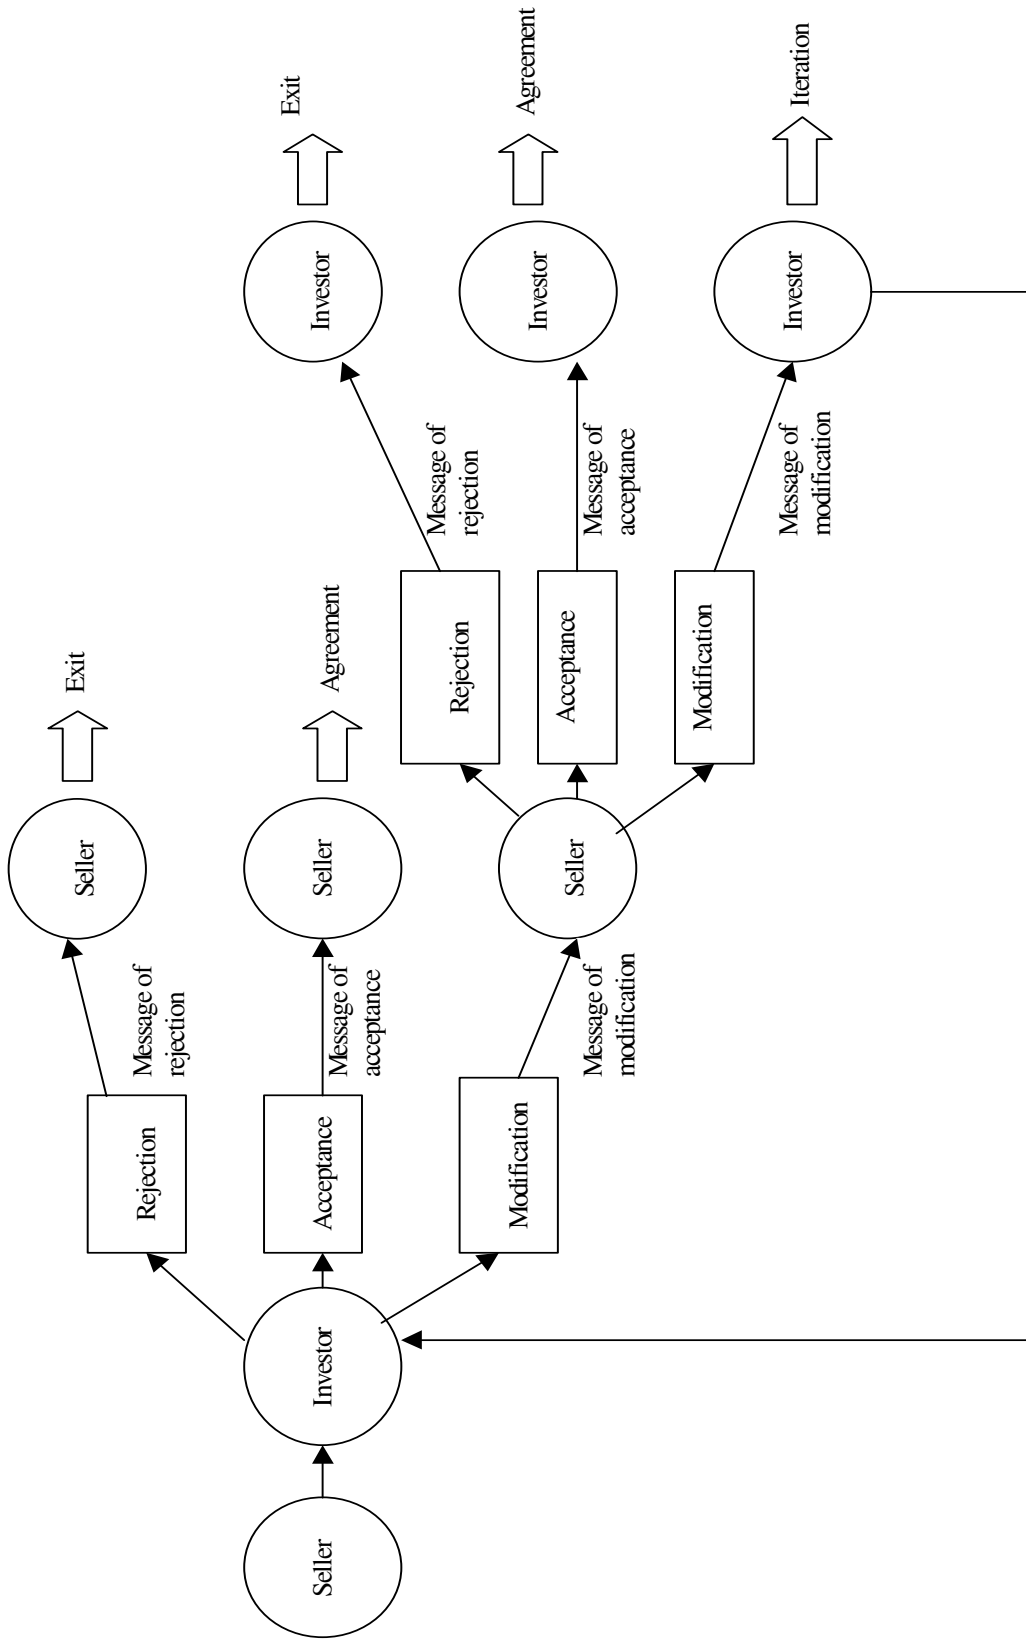
Similarly, once a message of acceptance has been received, the seller will update his mental state by generating the following mental attitudes:

$$\begin{aligned} &Bel(s, P_{s/is}=p_{s/is}); \text{ and} \\ &Int(s, sell_shares_at_P_{s/is}=p_{s/is}). \end{aligned}$$

- (ii) The investor may reject the investor's offer. This happens if: (i) the seller's offer does not match against the investor's evaluation parameters, and (ii) the number n of the messages forwarded between the seller and the investor reaches the threshold n^* ⁹. Should the investor reject the seller's offer, he will update his mental state by generating the intention to exit negotiation:

$$\begin{aligned} &Int(i_s, exit_negotiation) \text{ or} \\ &Int(i_w, exit_negotiation). \end{aligned}$$

Figure 2. The Negotiation Cycle.



Also, a message of rejection will be forwarded to the seller, who, in turn, will update his mental state in a similar way:

Int(s, exit_negotiation_with_small_investor) or
Int(s, exit_negotiation_with_active_investor).

In this case, negotiation fails and no final agreement about the price will ensue.

- (iii) Provided that $n^* > 1$, the investor may modify the seller's offer and send back a counter-offer. We assume that the first counter-offer the investor (either active or small) makes to the seller is the minimum value of his price range, that is, the book value of the firm (see Section 3.2).

Should negotiation proceed (point iii), the seller will have to evaluate the investor's counter-offer and determine whether it is acceptable or not. If acceptable, a price will be agreed upon by both parties. Otherwise, the seller may either reject the counter-offer (if $n^* = 2$) or modify it and send a modification back to the investor. The process continues iteratively until a price is either accepted (point i) or rejected (point ii) (Figure 2).

4 Analysis

The previous section introduced the main assumptions and properties of our computational model of going public. In this section, we will use this analytical structure as a test-bed to engage in two sets of analyses. The first examines the interrelationships between selling strategies, the negotiation protocol (price movement strategies), and negotiation performance, while the second set extends this comparison to a setting in which the economic actors may have varying degrees of accuracy of their cognitive representations.

In order to explore the basic interrelationships between selling strategies, negotiation protocols (price movement strategies), and performance, we carry out a first set of computer simulations in which we examine to what extent the choice of a method of selling shares and a negotiation protocol may impact upon the effectiveness and efficiency of a process of going public. Furthermore, we undertake a higher-level form of analysis by varying the degree of accuracy of the economic actors' cognitive representations. To this end, we undertake a second set of computer simulations in which we examine to what extent the choice of a method of selling shares and the accuracy of the seller's cognitive representations impact on the effectiveness and efficiency of the negotiation process.

In undertaking both sets of simulations, we make the following common assumptions. First, the total number of shares the seller can transfer either to the active investor or to the small investors is set at 100^{10} . Second, it is assumed that the book value (per share) of the firm is 10 Great Britain Pounds (GBP)¹¹. Third, control is reached with 50% of the shares, i.e., $m = 0.5^{12}$. Fourth, the number of messages n^* that represents the threshold after which agents will give up negotiation is set at 40^{13} . Fifth, for the sake of simplicity, both groups of our simulations rest on the assumption that, during negotiation, the three agent types use the same price movement strategy, that is, a negotiation protocol characterised by the same value for the concession rate c (Section 3.3)¹⁴.

4.1 Study I

One benefit of using MAS models for studying negotiation processes is that they can provide researchers with data on both the individual agent's and the group's behaviour. Such models therefore enable researchers to examine cognitive properties and behavioural processes at both the individual and collective level (see Figure 1). Study I will demonstrate the power of MAS models to provide these kinds of data.

Experimental Design

The purpose of Study I is to examine the relationship of the selling strategy and the negotiation protocol with negotiation performance. Study I uses a 4 x 6 design where the selling strategy and the negotiation protocol are manipulated. We take the selling strategy to be one of the four alternative selling strategies for going public discussed in Section 3.4: (i) selling the controlling block before the IPO (strategy A); (ii) an IPO followed by the sale of the controlling block (strategy B); (iii) a single-period parallel sale of dispersed shares and the controlling block (strategy C); and (iv) a public offering of all shares at a uniform price (strategy D). On the other hand, the negotiation protocol is manipulated through its correspondent price movement strategy, which, in turn, can be operationalised through the concession rate c ($0 \leq c \leq 1$) that agents use to modify unacceptable offers received during negotiation (see Section 3.3). We assume that, during negotiation, the agents use the same price movement strategy. That is, the concession rate c will be the same for each agent involved in the process. In order to allow for a wide range of negotiating protocols that become progressively more conciliatory, we will use the following six values for c : 7.5; 10; 12.5; 15; 17.5; and 20.

Negotiation performance is rated by two measures of negotiation effectiveness and efficiency: (i) the revenue raised in the aggregate sell of shares; and (ii) the length of time the

agents take to get to an agreement. A selling strategy turns out to be more effective than alternative strategies if it maximises the revenue raised. Conversely, a selling strategy is more efficient than alternative ones if it is associated with a lower amount of time it takes the agents to reach a decision. The revenue raised in the aggregate sell of shares is calculated by multiplying the number of shares sold by the price per share that has been agreed upon by the agents. The length of time it takes the agents to make an agreement is measured by the number of messages that the agents must send to each other before they can reach a final agreement.

To ensure that the results of our simulations reflect the underlying structure of the model and not merely particular realisations of a highly stochastic process, a Monte carlo approach is used to average out differences arising not only from different instantiations of the agents' mental states (Section 3.1), but also from instantiations of the seller's cognitive structure that share the same mental attitudes but vary in the sensitivity of the seller's evaluation parameters to changes in the accuracy of his cognitive representations (Section 3.2). As a result, our analysis is based on the average negotiation process resulting from independent runs of the simulation model. For each of these runs, distinct values for the agents' beliefs and the seller's sensitivity factor b are independently specified.

Hypothesis Tests

To examine the effects of the selling strategy and the negotiation protocol on negotiation performance, the following six null hypotheses are tested:


- (i) The four selling strategies affect the revenue raised in the aggregate sell of shares equally.
- (ii) The six price movement strategies affect the revenue raised equally.
- (iii) There are no interaction effects of the selling strategy and the price movement strategy on the revenue raised.
- (iv) The four selling strategies affect the number of messages sent during negotiation equally.
- (v) The six different price movement strategies affect the number of messages sent during negotiation equally.
- (vi) There are no interaction effects of the selling strategy and the price movement strategy on the number of messages sent during negotiation.

Results and Discussion

Figures 3 through to 6 present the results of Study I. Figures 3 and 4 show the impact that each combination of selling strategy and price movement strategy has upon the revenue raised in the aggregate sell of shares. In particular, Figure 3 shows a graph of the impact of the selling strategy on revenue at different price movement strategies. Figure 4 shows a graph of the impact of the price movement strategy on revenue when different selling strategies are used. In general, for each selling strategy, revenue decreases as the concession rate c increases. Here an obvious question is: What is driving the curves down as the concession rate increases? The answer is found in the role of the concession rate within the agents' negotiation protocol. The agents use the concession rate to modify the unacceptable offers they receive. The higher this parameter is, the more relevant the modification that is made (see Section 3.3). This, in turn, might have two alternative consequences: i) the seller makes more concessions and the firm is sold at a lower price; or ii) the investors make more concessions and the firm is sold at a higher price. Therefore, the decreasing revenue occurring in Figures 3 and 4 at higher concession rates expresses the fact that the impact of the modification made by the seller on revenue is stronger than the impact of the modification made by the investors.

Throughout all movement strategies, public offering at a uniform price (strategy D) is the selling strategy that keeps revenue at the lowest levels. This suggests that involving an active investor in the sale of the firm is, in general, a more effective strategy than a sale of all shares to small investors. There appears to be two major reasons for such a result. First, a large investor can use a controlling block to obtain private benefits and is, therefore, willing to pay a price that reflects such private benefits of control (Section 3.1). This, in turn, allows

Figure 3. The impact of the selling strategy on revenue at different price movement strategies.

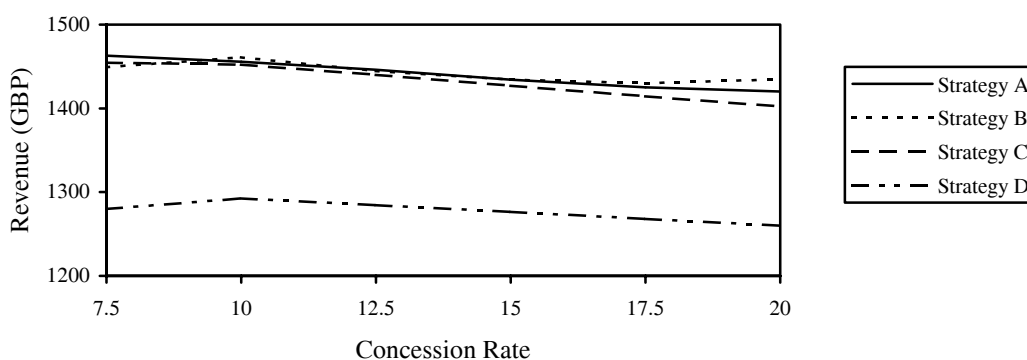


the seller to raise the price asked for a controlling block by a premium over the price at which dispersed shareholdings can be offered (Section 3.2). Second, because a large shareholder, upon achieving control of a firm, is in a position to shape future management's decisions and increase the firm's expected cash flows, all shareholders, and in particular the small investors, can benefit from the participation of a large investor in the sale. As a result of this, the seller is in a position to ask the small investors a higher price that in turn reflects the public good associated with the active investor's monitoring activities.

Not only do Figures 3 and 4 suggest that the sale of a firm should be designed to promote the participation of an active investor. Also, our results offer an indication as to how different ways of involving an active investor in the sale of a firm impact upon the effectiveness of the sale. As shown in Figure 4, strategy A is especially effective when the agents use low concession rates. At medium and high levels of this parameter, strategy B results in higher revenue than the alternative selling strategies. It produces the best results between 10 and 12.5 and between 15 and 20. It has a local maximum at 10 and a local minimum at 17.5. However, even at 17.5, strategy B is more effective than the other strategies. Furthermore, strategy B is the only one that results in an improvement of revenue as the concession rate increases from medium to high levels. Finally, the parallel strategy (strategy C) gives lower revenue than both strategies A and B (except for the first value of the concession rate), but higher than the simple public offering (strategy D).

An inspection of these results not only suggests that it is crucial that the method of sale should promote the participation of an active investor, but also that the involvement of the active investor in the sale should occur after a transaction with the small investors has taken place. This is necessary because in many instances the seller can extract the surplus associated with the active investor by making his allocation and payment contingent on the demands of the small investors (Mello and Parsons 1998).


Figure 4. The impact of the price movement strategy on revenue at different selling strategies.



Figures 5 and 6 show the impact that each combination of selling strategy and price movement strategy has upon the number of messages sent. In particular, Figure 5 shows a graph of the impact of the selling strategy on the number of messages at different price movement strategies. Figure 6 shows a graph of the impact of the price movement strategy on the number of messages at different selling strategies. Throughout all price movement strategies, public offering keeps the number of messages at the lowest levels. As in public offering the seller communicates only with the small investor, this result suggests that the number of messages is positively correlated with the number of agents among which agreement must be made. Judging from the experimental results, it also appears that, when the seller adopts a sequential strategy (either A or B) the number of messages decreases as the concession rate increases. That is, the more relevant the concessions the agents make to one another, the lower the number of messages that have to be sent in order for an agreement to be reached. Finally, throughout all price movement strategies, the parallel strategy is better than both sequential strategies A and B, but worse than the simple parallel strategy.

In combination, the four graphs indicate that strategy B is, in general, more effective than other alternative strategies (except for the lowest level of the concession rate), whereas simple public offering is strictly more efficient than alternative strategies. All our null hypotheses are rejected, as the test results support the existence of main and interaction effects of the selling strategy and the negotiation protocol on negotiation performance. Firstly, as is shown in Figures 3 and 5, the choice of a selling strategy has an impact on the revenue raised and the number of messages sent. Secondly, as is shown in Figures 4 and 6, the choice of a price movement strategy impacts upon revenue and number of messages. Finally, different combinations of selling strategy and price movement strategy have different impacts upon revenue and number of messages. This emphasises the fact that the design of the sale is an important determinant of the performance of the negotiation process through which the firm is sold.

Figure 5. The impact of the selling strategy on the number of messages at different price movement strategies.




4.2 Study II

Study I explored negotiation effectiveness and efficiency by examining external and internal measures of the group process: the final outcome of the process (measured by the revenue raised in the aggregate sale of shares) and the time it takes to get an outcome (measured by the number of messages sent during negotiation). The focus was on the impact that the selling strategy and the negotiation protocol have on these measures, under the assumption that no shift occurs in the agents' cognitive representations. Study II extends the analysis carried out with Study I to a setting in which the selling strategy and the agents' cognitive accuracy are used as independent variables. The purpose is to explore the relationship of the selling strategy and the degree of accuracy of the agents' beliefs with the revenue raised and the time that the agents take to reach an agreement. The combination of these independent variables sheds light on what are normally hidden cognitive processes. In particular, as we shall see, cognitive accuracy has an impact on the value of the available information, and this in turn affects the effectiveness of various selling strategies that are different in terms of the amount of information conveyed.

Experimental Design

Study II is undertaken to examine the effect of the selling strategy on negotiation performance at different levels of cognitive accuracy, and the effect of cognitive accuracy on negotiation performance when different selling strategies are adopted. Study II uses a 5 x 6 design where the selling strategy and the agents' cognitive accuracy are manipulated. In order to address some important questions about the role of cognitive accuracy in negotiation, we enrich the range of alternative selling strategies discussed in Section 3.4 by introducing a

Figure 6. The impact of the price movement strategy on the number of messages at different selling strategies.



a variant of sequential strategy B, which we call "strategy B *with motivation*" (strategy B*). When this strategy is adopted, the seller provides the active investor with more information about the offer being made. We assume that the active investor believes that the seller is *trustworthy*. In general, an agent will be most likely to accept what another agent says if it considers that agent trustworthy (Galliers 1988). Therefore, if the trustworthy seller provides the active investor with more detailed information on which the price being offered is made contingent (e.g. the estimated control premium mean or standard deviation, the price already established with other investors), then the active investor will take this additional piece of information to be true and will be more motivated to accept the offer received and act accordingly than would be the case if the seller had forwarded only a price with no additional information.

In undertaking our computer simulations, we assume that the additional piece of information conveyed with strategy B* is the market price that has been established with the small investor. That is, when the seller adopts strategy B*, he will inform the active investor about the market price established with the small investor at an earlier stage. As the active investor believes that the seller is trustworthy, he will believe that the market price communicated is true. The role of this additional piece of information (i.e., the market price) communicated by the seller is to justify the price offered to the active investor. As a result of this, the active investor's motivation to accept the offer received from the seller is expected to be enhanced. Consequently, the active investor is expected to accept the price offered less reluctantly than he would do in the case where he receives an offer without any further information that might justify it¹⁵. In the light of this, it is expected that strategy B* leads the active investor to become more conciliatory, thus allowing the seller to get the most out of the negotiation process, in terms both of the revenue raised and the length of time needed to reach an agreement. Our computer experiments are intended to examine whether these expectations about strategy B* can be confirmed or rejected, that is, at a more general level, whether and to what extent sharing additional information may impact upon the outcome of the negotiation process.

Our second independent variable – the agents' cognitive accuracy - will be manipulated through the quality of the seller's beliefs about the active and small investors' demands for the shares of the firm. In turn, the quality of the seller's beliefs is measured by the value the seller assigns to the standard deviation of the active investor's control premium and control discount, σ_{Pre} and σ_{Dis} respectively, and of the small investor's idiosyncratic component, σ_{Ics} (see Sections 3.1. and 3.2). For simplicity, we assume that $\sigma_{Pre}=\sigma_{Dis}=\sigma_{Ics}$ ¹⁶, and we will use the following six values for them: 0; 0.1; 0.2; 0.3; 0.4; and 0.5. These values

allow us to cover varying degrees of the seller's cognitive accuracy, from perfectly accurate beliefs ($\sigma = 0$) through to progressively more inaccurate beliefs.

As in Study I, our dependent variable, negotiation performance, is measured by: (i) the revenue raised in the aggregate sell of shares; and (ii) the length of time the agents take to reach an agreement. Also, the revenue raised is given by the number of shares sold times the agreed-upon price per share. Finally, the length of time is measured by the number of messages that the agents have sent to each other during negotiation.

In addition, as in Study I, to ensure that the results of this group of simulations reflect the underpinning structure of the model and not merely stochastic processes, we adopt a Monte Carlo approach to average out differences arising from different instantiations of the initial structural configuration. More specifically, in addition to differences arising from different values for the agents' mental states and for the sensitivity of the seller's evaluation parameters to changes in his cognitive accuracy, in this study we also need to average out differences arising from different values for the concession rate. As a result of this, the following analysis is based on the average negotiation process that results from a number of independent runs of the simulation model, each of them characterised by distinct independently specified values for the agents' beliefs, the sensitivity factor b , and the concession rate c .

Hypothesis Tests

To examine the effects of the selling strategy and the seller's cognitive accuracy on negotiation performance, the following six null hypotheses are tested:

- (i) The five selling strategies affect the revenue raised in the aggregate sell of shares equally.
- (ii) The six different degrees of the seller's cognitive accuracy affect the revenue raised equally.
- (iii) There are no interaction effects of the selling strategy and the seller's cognitive accuracy on the revenue raised.
- (iv) The five selling strategies affect the number of messages sent during negotiation equally.
- (v) The six different degrees of the seller's cognitive accuracy affect the number of messages sent during negotiation equally.
- (vi) There are no interaction effects of the selling strategy and the seller's cognitive accuracy on the number of messages sent during negotiation.


Results and Discussion

Figures 7 through to 10 show the results of Study II. Figures 7 and 8 focus on the impact that each combination of selling strategy and degree of the seller's cognitive accuracy has upon the revenue raised in the aggregate sell of shares. Figure 7 shows a graph of the impact of the selling strategy on revenue at different degrees of the seller's cognitive accuracy. Figure 8 shows a graph of the impact of the seller's cognitive accuracy on revenue for different selling strategies.

In general, as shown in Figure 7, the higher the degree of the seller's cognitive accuracy, the better the negotiation performance. One way of interpreting these results is that the agents are willing to make more concessions to each other when information of worse quality is available. In particular, the lower the degree of cognitive accuracy, the lower the minimum price that the seller is willing to accept from either the small or the active investor. This, in turn, implies that the lower the degree of cognitive accuracy, the lower the price that may be agreed upon during negotiation.

However, at high degrees of the seller's cognitive accuracy, strategies A and B do not work well. For example, when standard deviation is 0.1, strategy A is not fully successful as the seller makes an agreement only with the active investor, but not with the small investor. The revenue raised is therefore relatively low. On the other hand, with a standard deviation of 0.1, strategy B fails in that the seller and the active investor cannot reach an agreement. These results contrast with the performance of strategy B* that, at high degrees of cognitive accuracy, gives the highest revenue. How can this phenomenon be explained? The answer is found in the relationship that information has with different degrees of cognitive accuracy.

Figure 7. The impact of the selling strategy on revenue at different levels of cognitive accuracy




The higher the degree of the seller's cognitive accuracy, the lower the maximum price that the seller is willing to offer and the higher the minimum price the seller is willing to accept (Section 3.2). Therefore, the higher the degree of the seller's cognitive accuracy, the narrower the seller's price range. This, in turn, might have two consequences:

- (i) it might well be the case that the seller's price range will not eventually overlap with either the active investor's or the small investor's price range; or
- (ii) the overlapping area of the agents' price ranges becomes narrower as the spread between the minimum the seller is willing to accept and the maximum the investors are willing to offer becomes lower.

In the former case, negotiation fails and no agreement is made. In the latter, the agents must communicate a greater number of messages before they can reach an agreement. In such situations, agreement may be reached in a quicker and more effective way if the agents exchange more information so as to support their requests during negotiation. Figure 7 indicates that the more informative strategy B* is more effective than alternative selling strategies at high degrees of the seller's cognitive accuracy. For example, when the seller's cognitive representation is perfectly accurate (i.e. the standard deviation is zero) strategy B* gives the highest revenue, about 1394.98 (GBP), whereas strategy B gives only 600 (GBP) as no agreement is made between the seller and the active investor. This is because, when strategy B* is adopted, the trustworthy seller communicates an additional piece of information (i.e., the market price established with the small investor) to the active investor so as to justify the offer made. In turn, a greater amount of information helps the two agents to reach an agreement.

Figure 8. The impact of cognitive accuracy on revenue at different selling strategies.




Therefore, these results imply that the impact of having greater amounts of information on negotiation effectiveness becomes stronger as cognitive accuracy increases. Another result that is evident from Figure 7 is that, at medium and low degrees of the seller's cognitive accuracy, a public offering of all shares at a uniform price is not optimal. In such situations, a sequential sale involving an active investor is more effective as it allows the seller to obtain more information about the aggregate demand on which subsequent price can be made contingent. As Figure 8 indicates, at high and medium degrees of the seller's cognitive accuracy, strategy A dominates strategy B. However, at the same levels of accuracy (except for a standard deviation of 0.3), strategy A is dominated by strategy B*. On the other hand, at low levels of the seller's cognitive accuracy, strategy A is dominated by both strategy B and strategy B*.

In combination, these results suggest that, in general, a sequential sale beginning with an IPO, followed by a negotiated sale of a controlling block is more effective than alternative selling strategies. Further, the effectiveness of this sequential selling strategy over alternative strategies may require additional amounts of information to be exchanged between agents, depending on what is the degree of cognitive accuracy that characterises the negotiating agents' mental representations. In particular, since the impact of information on negotiation performance becomes stronger as the degree of cognitive accuracy increases, the more informative strategy B* becomes the dominant strategy at high levels of cognitive accuracy.

Figures 9 and 10 show the impact that each combination of selling strategy and degree of the seller's cognitive accuracy has upon the number of messages sent. Figure 9 shows a graph of the impact of the selling strategy on the number of messages sent at different levels of the seller's cognitive accuracy. Figure 10 shows a graph of the impact of the seller's cognitive accuracy on the number of messages sent when different selling

Figure 9. The impact of the selling strategy on the number of messages at different levels of cognitive accuracy.




strategies are used. At high levels of the seller's cognitive accuracy, strategies A and B are associated with the highest number of messages sent. This is because the agents' price ranges are narrower and it takes more time for the agents to reach the overlapping area. Therefore, the agents will keep communicating with each other, at least until a certain number of messages have been sent. In general, strategy B* requires fewer messages than strategy B. This result suggests that the additional piece of information sent together with the main message helps the agents (the seller and the active investor) to reach an agreement more quickly.

As is evident in Figures 9 and 10, throughout all levels of the seller's cognitive accuracy the lowest number of messages is associated with the public offering of all shares at a uniform price. As in Study I, this result suggests that the number of messages is positively correlated with the number of agents among which agreement must be made (public offering involves only two agents). Also, strategy B* works quite well in terms of keeping the number of messages reasonably low. In particular, strategy B* results in fewer messages than the other two sequential strategies, A and B, and at medium and low levels of the seller's cognitive accuracy it also generates fewer messages than the parallel strategy.

In combination, the four graphs of Study II indicate that our null hypotheses are rejected, as the test results support the existence of main and interaction effects of the selling strategy and the seller's cognitive accuracy on negotiation performance. Our results highlight the fact that it is important to design the sale of shares with the final ownership structure in mind. Like Study I, Study II suggests the idea that the choice of the selling strategy has an impact on negotiation effectiveness, and that such impact depends on the degree of cognitive accuracy that characterises the negotiating agents' mental representations. In particular, Study

Figure 10. The impact of cognitive accuracy on the number of messages at different selling strategies.



II supports the idea that an IPO followed by a sale of the controlling block dominates alternative strategies, in terms of generating higher revenue. However, the effectiveness of this sequential strategy depends on whether or not additional information is communicated and, in turn, on what is the degree of the agents' cognitive accuracy. Furthermore, according to Study II, the choice of the selling strategy has an impact upon negotiation efficiency. In particular, a sale in which only two agents are involved (parallel strategy) dominates alternative strategies in terms of keeping the number of messages sent at low levels. Also, the degree of cognitive accuracy has differing effects upon negotiation efficiency, depending on what selling strategy is used. For example, strategy A and strategy B generate fewer messages as the degree of the seller's cognitive accuracy decreases. As said, the higher the degree of cognitive accuracy, the narrower the overlapping area of the agents' price ranges, and the greater the number of messages that are sent before an agreement is reached.

5 Related Work

This paper draws upon agent-based computing and MAS theory and practice to provide a formal model of the financial process of going public and to undertake a series of computer simulations. Therefore, our work shares common research issues with two related fields: (i) research on the computational modelling of automated negotiation; and (ii) research on IPOs and staged equity financing. In what follows each area will be briefly dealt with and links will be discussed with our work.

In DAI approaches to automated negotiation, negotiation is usually referred to as a communication-based coordination mechanism used to resolve conflicts within a group of agents (Sycara 1988). As opposed to this functional conception, in this paper we adopted a socio-cognitive approach to negotiation, here conceived of as a joint behavioural process in which socially interconnected cognitive agents take an intentional stance towards each other (Dennet 1987; Panzarasa *et al.* 2001). Along these lines, we adhered to the definition of negotiation given by Pruitt (1981), who states that "negotiation is a process by which a joint decision is made by two or more parties. The parties first verbalise contradictory demands and then move towards agreements by a process of concessions or search for new alternatives". There have been several streams of research in DAI that have approached the problem of developing an agent-based architecture for automated negotiation in specific domains. The *Persuader* system by Sycara helps a company and its union to find a compromise concerning labour-conditions (Sycara 1989). In distributed transportation scheduling, shipping companies negotiate the exchange of orders and their prices (Fisher *et al.* 1995). The Diplomat system by Kraus and Lehman (1995) develops a model for a negotiating agent in the game Diplomacy domain, and addresses the issues of who to negotiate with, how

to generate proposals, how to evaluate counter-proposals, and how to form coalitions among the players. In a distributed business process management domain, agents discuss service level agreements, the conditions under which they provide services to each other (Jennings *et al.* 1996). Research by Durfee and Lesser (1991) addresses the issue of agents' communication in distributed problem-solving systems, where information from several agents is combined in order for a joint solution to be reached.

In all the above work, the focus is mainly on the process of negotiation among agents and not only on the outcome. Our approach is consistent with this point of view, and complements it in that we concentrated predominantly on the decision-making apparatus of the negotiating agents rather than on the negotiation protocol or the mechanism design. Along these lines, our work is most closely related to the formalism described in Parsons *et al.* (1998), where a model of argumentation-based reasoning and negotiation has been proposed for autonomous BDI agents. Like in their work, we modelled negotiation in terms of transitions between mental attitudes. However, our focus has been more on the practical instantiation of the model within a specific real-world scenario, and we showed how this domain can shape the content of the negotiating agents' mental states. As to our formalisation of the agents' cognition, two further comments are worth mentioning. First, our assumption that agents are persistent but not fanatical in their attempts to reach an agreement draws on our previous work on collaborative decision-making (Panzarasa *et al.* 2001) and is consistent with the formalisation of joint commitments developed in Levesque *et al.* (1990). Second, our assumption that agents are expected-utility maximisers is consistent with Rosenschein and Zlotkin's theory of mechanism design for automated agents (1994). In contrast to Rosenschein and Zlotkin's work, which is mainly based on the assumption that agents have complete knowledge of the payoff matrix, and hence of the counterparts' preferences, we made the more realistic assumption that agents have incomplete and not perfectly accurate beliefs about their physical and social environment. In our social simulations, we further examined the role of this assumption by analysing the impact that changing the degree of cognitive accuracy of the agents' beliefs may have upon the performance of the negotiation process.

Moving on to the financial modelling aspects of this work. In this paper, we chose to evaluate and compare alternative strategies for going public by running a number of computer simulations premised on an agent-based computational model. A more traditional method for pursuing this type of research is to analyse each alternative individually, solving for the equilibrium market price and for equilibrium investor strategies and then calculating the revenue raised by the seller of the firm (Chemmanur 1993; Schipper and Smith 1986; Stoughton and Zechner 1998). However, the enormous range of alternatives available for the sale of shares has led the majority of researchers to restrict their consideration to pairs of

alternative selling strategies and to establish the relative benefits of one over the other. This may represent a serious limitation as the choice of the optimal selling strategy is based only on a relatively restricted number of alternative strategies. Alternatively, another research option to evaluate strategies for going public is to adopt the analytical tools and methods employed in the mechanism design literature, centred upon the revelation principle (Mello and Parsons 1998; Milgrom and Weber, 1982; Barclay and Holderness, 1989; Brennan and Franks, 1995). The advantage of using this approach is that it allows researchers to characterise the maximum revenue that can be raised from *any* alternative selling strategy, given a set of informational assumptions. After characterising the maximum revenue raised from any selling strategy, various specific alternatives can be compared and evaluated, and the optimal strategy is eventually identified.

In contrast to both of these research approaches, we used computational modelling and computer simulations to evaluate different strategies for going public. A downside of this approach is that it is not grounded on a fully explicated mathematical model of the process, whereby properties can be formally proved as a function of the underlying assumptions and parameters. However, this disadvantage is more than compensated for by the way that our approach offers new insights into data structures and cognitive and behavioural processes both at the individual and group level that would not otherwise have been captured by the more traditional analytic methods used in the financial literature.

Matching computational models and simulation results with the behaviour of real modelling subjects can provide some degree of validation of the models (Burton and Obel 1995; Carley and Prietula 1994). As Shown in Section 4, our findings support the hypothesis that an IPO aimed at selling shares to passive and dispersed investors followed and by a sale of the controlling block to an active investor dominates alternative strategies of going public, in terms of generating higher revenue. Our results correspond with the findings from previous empirical research on IPOs. For example, it has been observed that the IPO of shares is not an isolated step but part of a more elaborate process for selling shares (e.g. Barry *et al.* 1990; Holderness and Sheehan 1988; Rydqvist and Högholm 1994, Pagano *et al.* 1998). Also, control turnover subsequent to the IPO is found to be the case in many equity carve-outs (Schipper and Smith 1986; Klein *et al.* 1991). This comparison of the existing empirical evidence on IPOs with our simulation results strongly demonstrates the usefulness of agent-based computational modelling for understanding the process of going public.

6 Conclusions and Future Work

This paper has sought to justify the claim that the use of computational modelling, and in particular MAS theory and practice, can provide both analytical tools and a conceptual

framework for studying the financial process of going public. In making this claim, we developed an agent-based model of the process, offered a behavioural explanation of negotiation, examined the cognitive structure of the negotiating agents, and structured a set of operations in computational terms (Yost and Newell 1993). In addition, two groups of social simulations were undertaken to ascertain the benefits and properties of our agent-based model. We found first, that simulation is valuable as a research tool for reasoning about our domain (Section 4) and second, that some of the results of studies I and II can be matched with existing empirical evidence (Section 5). Even though only a few design factors were considered, the two studies highlighted that ownership structure matters for the value of a corporation and thereby for the revenue raised in the aggregate sell of shares.

Furthermore, we examined how different methods for the sale of shares fare in establishing the appropriate ownership structure and maximising revenue. Along these lines, our results may be seen as a contribution to the ongoing debate over the importance of treating controlling blocks distinctively in selling a firm. Viewing the IPO as a step in a more complete process of selling the firm is the result of considering the inherent asymmetry of investors together with the strategic behaviour on the part of the seller. In particular, because large shareholders can provide the public good associated with monitoring activity, we have shown that the revenue raised in the aggregate sale of shares is maximised when an active investor participates in the sale, and if this depends on his chance of getting control, the sale should be designed to benefit him. To this end, we modelled the active investor's idiosyncratic component as incorporating a discount which, in turn, reflects the public benefits of control. However, guaranteeing the large shareholder a controlling stake would eliminate the competitive pressure to bid aggressively. Therefore, it is crucial that the method of sale promotes the participation of a potential large shareholder and at the same time makes his allocation and payment contingent on the demands of the small investors. To this end, we modelled the seller's evaluation parameters used during negotiation with the active investor as contingent on whether or not a sale of shares to the small investors had already taken place (Section 3.2).

On a more general level, this paper may also be seen as an attempt at illustrating the power of agent-based computational modelling in providing a large number of non-trivial insights, and developing theories about social and organisational processes (Carley and Wallace 1995). Computational simulations may be seen as useful precursors to laboratory experiments with human subjects when the latter require a lot of money and time (Stasser 1988). The complexity of real subjects and the limits of human cognitive capability prevent researchers from doing systematic and thorough studies of multi-agent social processes. Our research domain is an eloquent example of the advantage of virtual experiments in situations where real-life experimentation with real subjects would be prohibitively expensive, if not

impossible, and the complexity of the social and cognitive processes involved too high to be successfully handled.

Furthermore, it has been argued that computational simulations may often generate unanticipated results, which may provide new insights that are not readily perceived by direct observation of the research subjects (Carley and Wallace 1995). Our work provided justification for this claim. For example, in Study II, one such unexpected result was that the impact of having greater amounts of information on negotiation effectiveness becomes stronger as the agents' cognitive accuracy increases. As a result, our simulation provided an opportunity to analyse systematically the interrelationship between the value of a new piece of information and the degree of the cognitive accuracy of the mental representations of the agent who uses that information. This enabled us to integrate a new theoretical proposition into the set of predictions implied by our model, with the possibility of adding new insights for future research (Abelson 1968).

The computational model we developed is undoubtedly incomplete and a number of issues raised in this paper require further investigation. For example, we have shown first how the seller may adjust his offer on the basis of the counter-offer received, and second how the offer made to one investor may be made contingent on the offer made to the other (Section 3.2). However, there are other mechanisms that could be used to help improve the seller's knowledge about both the demand and the valuation of different investors. For example, should it be possible to open a when-issued market in the new stock, then the price in this market would provide crucial information to the seller that would make it possible to extract the maximum surplus given the need to assure a successful sale (Mello and Parsons 1998). Also, in the paper we have modelled an IPO where all shares are sold at a uniform price (Section 3.4). However, there are also IPOs in which investors tender bids and those bidding the highest pay the price bid. Future work may therefore be concerned with extending the set of selling strategies we considered in this paper so as to include different allocation rules and discriminatory practices among buyers (Brennan and Franks 1995; Hanley and Wilhelm 1995). Finally, our model may be further enriched if it included more than one potential large shareholder. In such situations, it would be necessary to run an auction or contest among all large shareholders. Subsequently, the advantage of selling some shares to the public ahead of time could be evaluated.

Notes

1. The term "virtual experiment" refers to an experiment in which the agent whose performance is being monitored is modelled computationally. As virtual reality is the computationally generated mirror of reality, a virtual experiment can be viewed as the computationally generated mirror of a human laboratory experiment or human field (natural) experiment.

2. The analysis could be extended to a setting in which other actors may well be involved. For example, in most real-world scenarios an investment banker or underwriter typically acts as a broker with an active and continuing relationship with the institutional investment community. This relationship may be beneficial to the seller of the firm for two main reasons. First, the investment banker may be able to identify those investors with monitoring capabilities and provide them with a favourable treatment in terms either of a price discount or of priority. Second, the investment banker may negotiate directly with the active investor, thus allowing the seller to extract more surplus (Stoughton and Zechner 1998). For simplicity, in this paper we will limit our attention to the seller of the firm and the active and small investors. We leave the refinement and enrichment of the set of economic actors for future work.

3. Crudely, there are two main reasons that justify this assumption. Firstly, in most real-world scenarios, small agents can be seen as homogeneous and, therefore, can be treated as endowed with the same mental states, attitudes, interest in the firm, negotiation abilities and power, and the same capabilities of inter-agent social behaviour. Indeed, one single small agent interested in dispersed allotments of shares can be considered as *price-taker*, in that he has virtually no power of influencing the price at which he can buy the firm's shares. The final share price can be seen as emerging from an agreement between the seller of the firm and the *whole* population of the small investors. As a result, the analysis of the negotiation process can conveniently abstract from each single small investor's idiosyncratic cognitive and/or social feature, and model just one of them as *representative* of the whole population. Secondly, our choice reflects the assumption that the seller is unable to discriminate among the small investors, even in those real-world scenarios in which they might have significant idiosyncratic differences. This assumption of absence of discrimination will be discussed later in the paper (Section 3.4).

4. Note that the choice of modelling only one active investor cannot be seen as an abstraction from the population of potential active investors, along the lines of the abstraction that we made from the population of the small investors. The reason is that each potential active investor has his own idiosyncratic capability of influencing the final price at which a controlling block can be bought (i.e., an active investor is a *price-maker*). These capabilities

are related to the public and private benefits of control (Section 3.1). Correspondingly, the seller is expected to be in a position to discriminate among potential active investors, in order to obtain greater extraction of surplus.

5. In our virtual experiments (Section 4), we will assume that control is reached with 50% of the shares (i.e., $m = 0.5$), although a lower proportion can be used to illustrate the problem.

6. Note that this is a strong assumption as it would be reasonable for the agent to keep negotiating even when an "acceptable" price has been offered and keep trying to obtain a lower one. In the interest of simplicity, we have not attempted to represent this aspect in our model.

7. This information is used by the agent to respond to incoming messages (see the specification of the negotiation protocol in Section 3.3).

8. For simplicity, at this stage we allow the agent to send only one argument (i.e., a price) in each message. However, it is possible to relax this restriction, but this will require the development of rules for deciding when to send more than one argument, which arguments to send in such situations, and how to evaluate messages which consist of more than one argument. In the second group of our virtual experiments (Section 4), we will provide an example of negotiation with messages containing two arguments (the price offered to the active investor and the market price established with the small investor).

9. Clearly, as only one message has been forwarded at this stage of negotiation, n^* is to be set at 1 in order for rejection to occur at this stage.

10. The results of both groups of simulations are rather insensitive to changes in this assumption, that is, to the precise value chosen for the total number of shares. For example, increasing the number of shares that can be sold has a negligible impact on negotiation performance, in terms both of effectiveness and efficiency.

11. As with the assumption about the number of shares, the results of our simulations are rather insensitive to changes in the precise value chosen for instantiating the book value of the firm.

12. For simplicity, the analysis does not consider more realistic cases in which control can be reached with less than 50% of the shares, such as economic transactions in which voting rights can be isolated from cash flow rights (Mello and Parsons 1998).

13. We consider this value a reasonable compromise between the case in which the agents are highly reluctant to concede and, therefore, are prone to exit negotiation "too early", and the case in which the agents are "fanatical" and keep negotiating until they reach an agreement.

14. A more in-depth analysis should take into account those scenarios in which the seller, the active investor and the small investor use different concession rates and, therefore, differ in terms of using more or less conciliatory negotiation protocols. Along these lines, the impact

of different combinations of the agents' negotiation protocols upon negotiation performance can be studied. We leave such refinements for future work.

15. Technically, in our computational model, the active investor's more conciliatory attitude towards the offers received from the seller is modelled via the impact that the market price communicated by the seller has upon the active investor's price range. More specifically, when strategy B* is adopted, if $P_{s/is}$ is the market price at which shares have already been sold to the small investor, the seller will forward a message to the active investor comprising both an offer, say $p_{s/ia}$, and the market price $P_{s/is}$. As the active investor believes that the seller is trustworthy, he will also believe that $P_{s/is}$ is true. As a result of this, he will update his evaluation parameters in the following way. His minimum price, $MinP_{ia}$ will be set at $P_{s/is}$ instead of the book value BV , whereas the maximum price, $MaxP_{ia}$, will be set at $(P_{s/is} + Pre - Dis)$ instead of $(BV + Pre - Dis)$ (see Section 3.2). As the market price will never be lower than the book value of the firm (Section 3.2), the above modifications imply that both the minimum counter-offer the active investor will make and the maximum price he will be willing to accept from the seller will increase. Against this, the ultimate performance of negotiation is expected to improve, both in terms of the final agreed-upon price and the amount of time it takes to get to an agreement.

16. A more accurate and realistic analysis should study the impact that a mix of different beliefs that vary in their cognitive accuracy have upon negotiation performance. We leave such refinements for future work.

References

Abelson, R. P. (1968). "Simulation of Social Behaviour", in C. Lindzey and E. Aronson (eds.), *Handbook of Social Psychology*, Vol. 2, Reading, MA: Addison-Wesley, pp. 274-356.

Baligh, H. H., and L. Richartz. *Vertical Market Structures*, Boston, MA: Allyn and Bacon, 1967.

Barbuceanu, M. (1997). "Coordinating Agents By Role-Based Social Constraints and Conversation Plans", *Proceedings of the Fourteenth National Conference on Artificial Intelligence*, 16-21.

Barclay, M., and C. Holderness. (1989). "Private Benefits of Control of Public Corporations", *Journal of Financial Economics* 25, 371-395.

Barry, C., C. Muscarella, J. Peavy, and M. Vetsuypens. (1990). "The Role of Venture Capital in the Creation of Public Companies", *Journal of Financial Economics* 27, 447-471.

Bebchuk, L. (1994). "Efficient and Inefficient Sales of Corporate Control", *Quarterly Journal of Economics* 109, 957-994.

Bebchuk, L., and L. Zingales. (1995). "Private Versus Social Optimality", Discussion Paper Series, Harvard Law School, Cambridge, MA.

Berle, A., and G. Means. (1932). *The Modern Corporation and Private Property*, New York, NY: World Inc.

Bond, A., and L. Gasser (eds.). (1988). *Readings in Distributed Artificial Intelligence*, San Mateo, CA: Kaufmann.

Bratman, M. E. (1987). *Intention, Plans, and Practical Reason*, Cambridge, MA: Harvard University Press.

Brennan, M., and J. Franks. (1995). "Underpricing, Ownership and Control in the Initial Equity Public Offering of Equity Securities in the U.K.", Working Paper 12-95, University of California, Los Angeles, CA.

Burton, R. M., and B. Obel. (1995). "The Validity of Computational Models in Organisation Science: From Model Realism to Purpose of the Model", *Computational and Mathematical Organisation Theory* 1(1), 57-71.

Carley, K. (1989). "The Value of Cognitive Foundations for Dynamic Social Theory", *Journal of Mathematical Sociology*, 14(2-3), 171-208.

Carley, K., and W. A. Wallace. (1995). "Editorial", *Computational and Mathematical Organisation Theory* 1(1), 5-7.

Carley, K., and M. J. Prietula (eds.). (1994). *Computational Organisation Theory*, Hillsdale, NJ: Lawrence Erlbaum Associates.

Castelfranchi, C. (1995). "Commitments: From Individual Intentions to Groups and Organisations", in V. Lesser (ed.), *Proceedings of the First International Conference on Multi-Agent Systems*, San Francisco, CA: AAAI Press and MIT Press, pp. 41-48.

Castelfranchi, C. (1998). "Modelling Social Action for AI Agents", *Artificial Intelligence* 103, 157-182.

Chemmanur, P. (1993). "The Pricing of Initial Public Offerings: A Dynamic Model with Information Production", *Journal of Finance* 48, 285-304.

Cohen, P. R., and H. J. Levesque. (1990). "Intention is Choice With Commitment", *Artificial Intelligence* 42(3), 213-61.

Davis, R., and D. Lenat. (1980). *Knowledgebase Systems in Artificial Intelligence*, New York, NY: McGraw Hill.

Dennet, D. C. (1987). *The Intentional Stance*, Cambridge, MA: The MIT Press.

Durfee, E. H. and V. R. Lesser. (1991). "Partial Global Planning: A Coordination Framework for Distributed Hypothesis Formation" *IEEE Trans. Systems Man Cybernet.* 21(5), 1167-1183.

Fisher, K., J. J. Mueller, M. Pischel, and D. Schier. (1995). "A Model for Cooperative Transportation Scheduling", *Proceedings of the International Conference on Multi-Agent Systems (ICMAS-95)*, AAAI/MIT Press, pp. 109-116.

Franklin, S., and A. Graesser. (1997). "Is it an Agent, or Just a Program? A Taxonomy For Autonomous Agents", in J. P. Mueller, M. J. Wooldridge, and N. R. Jennings (eds.), *Intelligent Agents III*, Berlin: Springer Verlag, 21-35.

Galliers, J. R. (1988). "A Strategic Framework for Multi-agent Co-operative Dialogue", *Proceedings of the Eighth European Conference on Artificial Intelligence*, London: Pitman, pp. 415-20.

Grossman, S., and O. Hart. (1980). "Takeover Bids, the Free-Rider Problem, and the Theory of the Corporation", *Bell Journal of Economics* 11, 42-64.

Hanley, K., and W. Wilhelm. (1995). "Evidence on the Strategic Allocation of Initial Public Offerings", *Journal of Financial Economics* 37, 239-257.

Hayes-Roth, F., D. A. Waterman, and D. R. Lenat (eds.). (1983). *Building Expert Systems*, Reading, MA: Addison-Wesley.

Holderness, C., and D. Sheehan. (1988). "The Role of Majority Shareholders in Public Held Corporations: An Exploratory Analysis", *Journal of Financial Economics* 20, 317-347.

Jennings, N. R. (2000). "On Agent-Based Software Engineering", *Artificial Intelligence* 117, 277-296.

Jennings, N. R., P. Faratin, M. J. Johnson, T. J. Norman, P. O'Brien and M. E. Wiegand. (1996). "Agent-Based Business Process Management", *International Journal of Cooperative Information Systems* 5(2&3), 105-130.

Jennings, N. R., K. Sycara, and M. Wooldridge. (1998). "A Roadmap of Agent Research and Development", *Autonomous Agents and Multi-Agent Systems* 1, 275-306.

Kahan, M. (1993). "Sales of Corporate Control", *Journal of Law, Economics and Organisation* 9, 368-379.

Kinny, D., M. Ljungberg, A. S. Rao, E. Sonnenberg, G. Tidhar and E. Werner. (1992). "Planned Team Activity", in C. Castelfranchi and E. Werner (eds.), *Artificial Social Systems – Selected Papers from the Fourth European Workshop on Modelling Autonomous Agents in a Multi-Agent World, MAAMAW-92, Vol. 830*, Berlin: Springer-Verlag, pp. 226-256.

Klein, A., J. Rosenfeld, and W. Beranck. (1991). "The Two Stages of an Equity Carve-Out and the Price Response of Parent and Subsidiary Stock", *Managerial and Decision Economics* 12, 449-460.

Kraus, S. K., and D. Lehmann. (1995). "Designing and Building an Automated Negotiating Agent", *Computational Intelligence* 11(1), 132-171.

Kraus, S., K. Sycara and A. Evenchil. (1998). "Reaching Agreements Through Argumentation: A Logical Model and Implementation", *Artificial Intelligence* 104, 1-69.

Lee, J., M. J. Huber, E. H. Durfee, and P. G. Kenny. (1994). "UM-PRS: An Implementation of the Procedural Reasoning System for Multirobot Applications", in *Conference on Intelligent Robotics in Field, Factory, Service, and Space (CIRFFSS '94)*, Houston, TX, pp. 842-849.

Levesque, J., P. R. Cohen, and J. H. T. Nunes. (1990). "On Acting Together", *Proceedings of the Eighth National Conference on Artificial Intelligence, AAAI-90*, pp. 94-99.

Malone, T. W. (1987). "Modeling Coordination in Organizations and markets", *Management Science* 33(10), 1317-1332.

Mello, A. S., and J. E. Parsons. (1998). "Going Public and the Ownership Structure of the Firm", *Journal of Financial Economics* 49, 79-109.

Milgrom, P., and R. Weber. (1982). "A Theory of Auctions and Competitive Bidding", *Econometrica* 50, 1089-1122.

Mikkelson, W., M. Partch, and K. Shah. (1995). "Performance of Companies around Initial Public Offerings", Working Paper, University of Oregon, Eugene.

Mikkelson, W., and R. Ruback. (1985). "An Empirical Analysis of the Interfirm Equity Investment Process", *Journal of Financial Economics* 14, 523-553.

Pagano, M., F. Panetta, and L. Zingales. (1998). "Why Do Companies Go Public? An Empirical Analysis", *Journal of Finance* 53, 27-64.

Panzarasa, P., T. Norman and N. R. Jennings. (1999). "Modeling Sociality in the BDI Framework", in *Intelligent Agent Technology: Systems, Methodologies, and Tools. Proceedings of the First Asian Pacific Conference on Intelligent Agent Technology (IAT-99)*, World Scientific Publishing, pp. 202-206.

Panzarasa, P., N. R. Jennings, and T. Norman. (2001). "Formalising Collaborative Decision-Making and Practical Reasoning in Multi-Agent Systems", *Journal of Logic and Computation*, forthcoming.

Parsons, S., C. Sierra and N. R. Jennings. (1998). "Agents that Reason and Negotiate by Arguing", *Journal of Logic and Computation* 8(3), 261-292.

- Pruitt, D. (1981). *Negotiation Behaviour*, New York, NY: Academic Press.
- Rao, S., and M. P. Georgeff. (1991). "Modeling Agents Within a BDI Architecture", in R. Fikes, and E. Sandewall (eds.), *Proceeding of the Second International Conference on Principles of Knowledge Representation and Reasoning (KR '91)*, Cambridge, MA: Morgan Kaufmann, pp. 473-484.
- Rock, K. (1986). "Why Are New Issues Underpriced?", *Journal of Financial Economics* 15, 187-212.
- Rosenschein, J., and G. Zlotkin. (1994). *Rules of Encounter: Designing Conventions for Automated Negotiation among Computers*, Cambridge, MA: The MIT Press.
- Rydqvist, K., and K. Högholm. (1994). "Going Public in the 1990s: Evidence from Sweden", *European Financial Management* 1, 287-316.
- Shleifer, A., and R. Vishny. (1986). "Large Shareholders and Corporate Control", *Journal of Political Economy* 94(3), 461-488.
- Shoham, Y., and M. Tennenholz. (1992). "On the Synthesis of Useful Social Laws For Artificial Agent Societies", in *Proceedings of the Tenth National Conference on Artificial Intelligence, AAAI-92*, San Jose, CA, pp. 276-280.
- Shipper, K., and A. Smith. (1986). "A Comparison of Equity Carve-Outs and Seasoned Equity Offerings", *Journal of Financial Economics* 15, 153-186.
- Singh, M. P. (1995). *Multi-Agent Systems: A Theoretical Framework for Intentions, Know-how, and Communications. Lecture Notes in Computer Science, Vol. 799*, Berlin: Springer-Verlag.
- Stasser, G. (1988). "Computer Simulation as a Research Tool: The DISCUSS Model of Group Decision-Making", *Journal of Experimental Social Psychology* 24, 393-422.
- Stoughton, N., and J. Zechner. (1998). "IPO-Mechanisms, Monitoring and Ownership Structure", *Journal of Financial Economics* 49, 45-77.

Sycara, K., (1988). "Multi-Agent Compromise Via Negotiation", in A. Bond and L. Gasser (eds.), *Readings in Distributed Artificial Intelligence*, San Mateo, CA: Kaufmann,

Sycara, K. (1989). "Argumentation: Planning Other Agents' Plans", in *Proceedings of the Eleventh International Joint Conference on Artificial Intelligence*, pp. 517-523.

Waterman, D. A. (1986). *A Guide to Expert Systems*, Reading, MA: Addison-Wesley Publishing Co.

Weston, F., K. Chung, and S. Hoag. (1990). *Mergers, Restructuring, and Corporate Control*, Englewood Cliffs, NJ: Prentice Hall.

Williamson, O. E. (1975). *Markets and Hierarchies*, New York, NY: Free Press.

Williamson, O. E. (1981). "The Economics of Organization: The Transaction Cost Approach", *American Journal of Sociology* 87(1981), 548-575.

Wooldridge, M. (1997). "Agent-Based Software Engineering", in *IEE Proceedings on Software Engineering* 144(1), 26-37.

Wooldridge, M. (2000). *Reasoning About Rational Agents*, Cambridge, MA: The MIT Press.

Yost, G. R., and A. Newell. (1993). "A Problem Space Approach to Expert System Specification", in P. S. Rosenbloom, J. E. Laird, and A. Newell (eds.), *The Soar Papers: Research on Integrated Intelligence* 2(46), Cambridge, MA: The MIT Press, pp. 982-988,

Zingales, L. (1995). "Insider Ownership and the Decision To Go Public", *Review of Economic Studies* 62, 425-448.