



ELSEVIER

Available at
WWW.MATHEMATICSWEB.ORG
POWERED BY SCIENCE @ DIRECT®

Simulation Modelling Practice and Theory 10 (2002) 417–453

**SIMULATION
MODELLING**
PRACTICE AND THEORY

www.elsevier.com/locate/simpat

Social influence, negotiation and cognition

Pietro Panzarasa *, Nicholas R. Jennings

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

Received 1 November 2001; accepted 1 March 2002

Abstract

To understand how interpersonal agreements can be generated within complexly differentiated social systems, we develop an agent-based computational model of negotiation in which social influence plays a key role in the attainment of social and cognitive integration. The model reflects a view of social influence that is predicated on the interactions among such factors as the agents' cognition, their abilities to initiate and maintain social behaviour, as well as the structural patterns of social relations in which influence unfolds. Findings from a set of computer simulations of the model show that the degree to which agents are influenced depends on the network of relations in which they are located, on the order in which interactions occur, and on the type of information that these interactions convey. We also find that a fundamental role in explaining influence is played by how inclined the agents are to be conciliatory with each other, how accurate their beliefs are, and how self-confident they are in dealing with their social interactions. Moreover, the model provides insights into the trade-offs typically involved in the exercise of social influence.

© 2002 Elsevier Science B.V. All rights reserved.

Keywords: Agreement; Mental states; Social structure; Negotiation protocol; Cognitive accuracy

1. Introduction

Social influence has long been an object of popular fascination and scientific research in such fields as social psychology [55,80] structural and network analysis [30,37,45], sociology [19,21,31] administrative science [75], organisation theory [7,16,59], and distributed artificial intelligence [41,83]. Across all these disciplines, influence researchers are concerned with the same challenging puzzles. How can one

* Corresponding author.

E-mail addresses: pp@ecs.soton.ac.uk (P. Panzarasa), nrj@ecs.soton.ac.uk (N.R. Jennings).

persuade another individual or group to change their mental attitudes? How can one convince an acquaintance that what they believe is correct is in fact wrong, and what they used to consider wrong is right? When and how do people become more flexible and change their beliefs and goals, and when do they become more committed and stick to their attitudes? When do we want people to agree with us, and how do we go about reaching an agreement?

These are all issues of social influence. What they have in common is the problem of how socially and cognitively differentiated individuals can become socially and cognitively integrated. As pointed out by Friedkin [30], this represents an old sociological problem that is rooted in Durkheim's insight that social differentiation does not necessarily lead to discordant actions and attitudes. There are forms of social differentiation that entail an integrative social structure in which interpersonal agreement can be generated [25]. However, elucidating mechanisms by which people in different positions come to be coordinated still remains a key theoretical problem that has long been studied in research areas such as persuasion and attitude change [10], power and authority relations [52], group polarisation [58], social conformity [55], minority influence [57] and the generation of social norms [72].

A first line of enquiry has been concerned with the structural conditions that foster and impede the development of interpersonal agreements. On the one hand, some theorists have been primarily interested in the social control implications of social structure and institutionalised arrangements within organisations and social institutions [16,75]. Institutionalised status characteristics and power structures have been regarded by these theorists as the main sources of social integration. In this view, social structure influences the individual actors in two main ways [75]. Firstly, it permits stable expectations to be generated by each actor as to the behaviour and mental states of the other actors. Secondly, it provides the general stimuli and attention-directors that shape the actors' attitudes, habits, states of mind and, ultimately, behaviour. Along these lines, the main function of social structure is to set the conditions for the development of coordinated activity and interpersonal agreement. On the other hand, other theorists have argued that even strongly constraining social structures do not subsume the whole individual, and actors who occupy identical social positions might exhibit substantially different behaviour and attitudes [48,49,85]. Personal goals, beliefs, and values introduce highly idiosyncratic and diverse responses to the circumstances of identical social positions. In this case, when extraneous conditions seep into the determination of the behaviour and mental states of actors in social positions, the generation of interpersonal agreements and coordinated activity cannot be regarded simply as mechanistically determined by institutional arrangements.

The emphasis on the individuals and their idiosyncratic attitudes underscores the importance of a second line of inquiry concerned with the role that cognitive and behavioural micromechanisms play in the attainment of social integration in complexly differentiated social structures. In order to address this issue, theorists have variously emphasised either social interaction procedures or cognitive processes that have a bearing on the generation of coordinated activity and interpersonal agreements. Some researchers, like for instance symbolic interactionists, have been

concerned with the interaction and influence processes leading to the formation of various types of agreements such as shared understandings, definitions of the situations, and interpersonal “contracts” [46,77]. Related lines of inquiry in social psychology developed theoretical models of social influence that describe the emergence of collective opinions in natural settings by focussing on the underlying network of social interactions [22,29,42]. Finally, other theorists have concentrated on individual decision rules, problem-solving procedures and local rationality mechanisms [66,70]. In an effort to investigate the aggregate properties of an organisation or a social system generated by repeated interactions among the actors involved, assumptions have been developed about the cognitive processes and strategies by which decisions are made by the individual actor [39]. Along these lines, a great body of research has concentrated on the emergence of coordinated activity and agreements from a set of individual motives, values, preferences, objectives, decision strategies and micromechanisms of social behaviour [2,3,67].

All the foregoing research avenues provide an analysis of the direct impact that structural, social or cognitive factors have on the generation of interpersonal agreements. However, with only few exceptions [30], most of these studies rarely look at the effect that combinations of different factors have on the final agreement. As a result, the problem of elucidating mechanisms of social influence that are based on a fully explicated model of the interconnections among social structure, cognition and social interaction, remains relatively unexplored. This article presents one possible approach to this very general problem by postulating a model of negotiation and agreement generation that is rooted in the interactionist approach, but, nonetheless, will take into account both the contextual features of the set of social relations in which the influence process unfolds, as well as the actors’ cognitive structures and their social abilities to initiate and maintain social interactions with one another.

To this end, included in the model are the following components of the influence process: (i) the actors’ mental states; (ii) principles for social interaction; and (iii) structural patterns of social relations. An actor’s mental state is the set of the basic mental attitudes that are processed to undertake both theoretical and practical reasoning [61,62]. Principles for social interaction refer to the procedures that actors use to initiate and maintain social processes with others (e.g. rules for determining when to start and when to exit a social relationship; rules for communication). Structural patterns of social relations reflect the social context in which an actor operates and are concerned with the articulation of the structural conditions in which social processes are generated (e.g. who interacts with whom; how social relationships involving more than two actors are organised; in what order actors are involved in social relationships, and whether they are involved simultaneously or sequentially). Our interests lie in the main and combined effects of these factors on social influence, here regarded as the process through which an actor, located in a social context and capable of cognitive and social processes, affects, and thereby alters, the mental states of other actors. Thus, the model reflects a conception of social influence as a socio-cognitive process whose typical outcome is the adoption of socially motivated mental attitudes and, ultimately, when played out in a network of social relations, the generation of interpersonal agreements among different actors.

In this paper, we elaborate on this idea and we conduct simulation studies in which we vary systematically the key components of the model. Results allow us to explore the dynamic implications of social influence over time, and provide us with new insights into the attainability of higher-level mentalistic phenomena in complexly differentiated communities [30,34]. In particular, we run a set of virtual experiments to investigate the role of social influence in the negotiation process underpinning the sale of a privately held company through the public offering of shares and, possibly, through the involvement of an active investor. We use this domain in our study because of its suitability to a formal modelling approach. Building on mainstream literature in behavioural finance [47,60], we develop a model of negotiation by deriving its main assumptions from extant theory and established empirical studies of going public processes. Using this model, we examine to what extent the performance of the sale of a firm is affected by the actors' negotiation protocols and cognitive structures as well as by the selling strategies which, in turn, reflect structural patterns of social relationships in which social influence unfolds. By exploring the relative effects of these factors on the performance of negotiation, an attempt is made here to determine, on theoretic grounds, the role that social influence processes play in the attainability of higher-order mentalistic phenomena such as interpersonal agreements.

In the following section, the process of going public is presented by describing its articulation, the various economic actors involved, the range of possible strategies that can be used for selling a firm, and the debate that exists in the literature as to the relative benefits of these strategies. This is followed by a detailed description of the model of negotiation and its main components. The model will then be used to run a series of simulations, for which we will provide a description of the experimental design and a discussion of the results. Simulations will be organised into two groups. The first group is intended to examine the individual and combined effect that structural patterns of social influence and the economic actors' negotiation behaviour have on the performance of the sale. In the second group, we look at the role of actors' cognition more closely, and we explore the individual and combined effect that structural patterns of social influence and the actors' cognitive accuracy have upon the performance of the process. Finally, we summarise and discuss our major findings, evaluate the model's scope and provide some avenues for future work.

2. Framework description

The domain in which we will analyse the issues raised above is the sale of a privately held company through a number of alternative selling strategies, each reflecting a different structural pattern of social influence that the seller can exercise upon the potential investors [63]. We envisage examining whether and to what extent the choice of a selling strategy, in combination with other social and cognitive factors, impacts upon the negotiating agents' mental states and, ultimately, upon the generation of an agreement and the performance of the going public process.

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. A great deal of literature and 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 [4,33,47,51,74]. On the one hand, the public offering of shares may mainly involve relatively *small investors* who will just remain passive holders of the firm's shares. On the other, there might 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 [47]. 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 [8,40,50,73]. 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 [47]. Accordingly, what types of investors are involved in negotiation, in what order they are influenced and using which method represent fundamental decisions that are likely to impact upon the final outcome of the sale.

These observations are supported by a great deal of recent empirical evidence about the segmentation of the market for shares [5,8,33,35,60,71]. 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 [8]. This engenders the belief that going public is not an isolated step, but a complex and extended process with distinct and heterogeneous markets for dispersed shareholdings and potentially influential blocks.

Even though this view of the sale of a firm seems to be confirmed in the literature, nonetheless debates exist over the relative benefits of different selling strategies involving different actors in a different order. Some theorists argue that selling the controlling block first avoids the free rider problem [32] and can also help reduce the winner's curse problem faced by small investors [69]. Along these lines, it has been suggested that 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 [76]. By contrast, other researchers claim that by selling the controlling block first, the active investor's allocation cannot be made contingent on the parameter of aggregate demand by the small investors, and therefore the seller cannot extract

surplus from the active investor [47]. In this vein, it is argued that 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. Barry et al. [5], for example, report that control turnover subsequent to the initial public offering (IPO) represents the most frequently used method for sales in which venture capital firms are involved. In this case, during the 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. Along these lines, control turnover subsequent to the IPO has been found by Holderness and Sheehan [35] for the US, Rydqvist and Högholm [71] for Sweden and the UK, and Pagano et al. [60] for Italy. Furthermore, it has been suggested that this strategy has been adopted in many equity carve-outs [40,73].

In addition to these studies that explored selling strategies in which different investors are approached by the seller sequentially in a different order, other studies have focused on strategies in which all the investors interact with the seller simultaneously. There are researchers who suggest evidence of a simultaneous sale of shares to small investors with a discriminatory allocation of a potentially controlling block. For instance, Brennan and Franks [8] refer to a method that combines a private placement, targeting large investors, with a simultaneous public offering. It has also been argued that, whenever used, this selling strategy usually includes a clawback provision [47]. This makes the allocation to a large investor dependent on the demand by small investors, therefore creating competition among investors of different types. Finally, other strategies have been reported in which the seller interacts only with the small investors. An example of this selling strategy is the offer for sale by tender, employed in the UK and in France, where investors place bids for the shares indicating both quantity and price [47]. After the bids have been received, a single share price is established and all investors buy shares at that price. Accordingly, with this strategy the seller is unable to discriminate among the small investors, and all shares are offered at a uniform price. However, there is an interesting variant of this strategy, in which the seller is allowed to adopt discriminatory practices and allocation rules: 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 US the allocation of oversubscribed issues is a means for discriminating among buyers [33]. Evidence of discrimination in connection with an IPO has been found also in the UK market [8].

The foregoing studies indicate that there is debate in the literature as to the relative impact of the design of a sale on performance. Therefore, pulling all the available findings into a single coherent model represents a difficult task. There are two main limitations to these studies. First, their results are based on a series of empirical and case studies that rarely look at the same or comparable variables. As a result, it is often difficult to directly compare and interrelate their findings. Second, these studies rarely describe the process underlying the relationship among variables. Hence, the problems of analysing the appropriate variables that are likely to affect perfor-

mance, and how they interact with each other over time in generating their effects remain relatively unexplored. More specifically, some of the key questions that need to be addressed are: How and to what extent should the heterogeneity among potential investors affect the firm's strategy for selling shares? How should a firm approach and influence potential investors in order to maximise its 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 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? This paper presents one possible approach to these problems by proposing an agent-based computational model of negotiation and agreement generation in connection with the process of going public. Using this model, we will address the above questions by undertaking a comprehensive and systematic process-based analysis of the relative impact that different ways of exercising social influence have on the performance of the sale of the firm.

3. Basic model

In this section we describe our model of negotiation and agreement generation, its fundamental assumptions, components and properties (see Panzarasa et al. [63] for a more technical description). The model draws on and extends the tradition of behavioural models in the finance literature associated with work by Bebchuk and Zingales [6], Kahan [38], and Mello and Parsons [47]. It focuses on the negotiation process between the seller of the firm and potential investors. The seller may be either the owner of the firm or the representative of a company that owns the firm. The main assumption on which the model is premised is that potential investors are not homogeneous in their ability to monitor the newly public firm's management [47,76]. We identify 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 [47]. Therefore, they only differ in their demands for the shares of the firm and, consequently, in the degree of influence they are expected to exert on the firm once they become shareholders. For simplicity, we assume that the seller is unable to discriminate among the small investors, and accordingly the model components will be couched in terms of only one small investor who is assumed to represent the whole population. Similarly, interactions with only one potential active investor will be examined, and therefore the model 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 with.

Both the seller and the investors are conceptualised as *cognitive* agents endowed with mental attitudes representing the world and motivating action [61,62,84].

Furthermore, not only does each agent have an intentional stance towards the world, but it also represents the other agents as cognitive agents similarly endowed with mental attitudes for representational and motivational purposes [23]. The cognitive characterisation of the agents' mental states and decision-making apparatus is here formalised building on a fairly standard belief–desire–intention (BDI) framework [17,68,84]. Accordingly, each agent's mental state is seen as a set of interrelated mental attitudes comprising beliefs, goals and intentions. First, beliefs correspond to the information each agent has about its environment. Beliefs are incomplete and may be updated as new pieces of information become available. Second, goals represent the top-level agenda or the set of states of affairs the agent wants to bring about. Finally, intentions refer to those states of affairs that the agent commits itself to bring about in order to achieve (some of) its goals [68,84]. Furthermore, the model draws on the BDI framework for describing each agent's decision-making process in terms of transitions between mental attitudes. Upon receiving new information, beliefs will be updated and intentions generated that do not contradict the agent's beliefs. 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.

Finally, both the seller and the investors will be modelled as autonomous, reactive, proactive, and socially capable agents [28,36,84]. First, they have control over their own tasks and resources and will take part in cooperative activities only if they choose to do so (autonomy). Second, each agent is assumed to respond to any perceived change that takes place within its environment and that affects its mental state (reactiveness). Third, agents do not simply act in response to their environment, but they exhibit opportunistic behaviour, take the initiative where appropriate, exploit and create serendipity (proactiveness). Fourth, agents can initiate social relationships and therefore engage in a wide variety of social interactions with others in a flexible manner (social ability).

Drawing on these assumptions, in what follows we will present a description of the model. For ease of communication, the model will be organised into four fundamental components: the agents' mental state specifications (Section 3.1); the evaluation parameters used by the agents to assess the messages received (Section 3.2); the negotiation protocol (Section 3.3); and a set of selling strategies for going public (Section 3.4).

3.1. The agents' mental states

Even though the structures of both the investors' and the seller's mental states are assumed to be organised into the same mental attitudes (beliefs, goals and intentions), nonetheless each type of agent differs in terms of the instantiation of the mental attitudes.

Small investors have beliefs about the value of the firm. We assume that, for each 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 a share 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 [47]. The small investor's beliefs are therefore assumed to reflect both BV and IC_s . Based on these beliefs, the small investors generate the goal to maximise their expected utility by purchasing shares of the firm at the lowest possible price. Finally, in order to have their goal fulfilled, the small investors maintain the intention to negotiate with the seller and, more specifically, to receive an offer from the seller and to evaluate it.

As with the small investors, the active investor's beliefs reflect his evaluation of the firm. The active investor, i_a , seeks a controlling block in order to actively influence or monitor future management's decisions. He might be interested in acquiring control for instance because of information he has about a strategy that would allow him to use the assets of the firm and increase the value of the firm's cash flows. As with the small investors, the actual price that the active investor is prepared to pay for a share 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 reason for a discount is that the active investor, upon achieving control of the firm, is expected to 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 [47]. On the other hand, 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. Given this, the active investor's beliefs are assumed to reflect both the book value, BV, and the idiosyncratic components Dis and Pre. Furthermore, the active investor generates the goal to maximise his expected utility by purchasing a controlling fraction m of the shares at the lowest possible price. Finally, as with the small investors, in order to have his goal fulfilled, the active investor will intend to receive an offer from the seller and to evaluate it.

On to the seller's mental state. 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 [76,86]. In this paper, we will 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, as well as 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. The seller's beliefs are therefore assumed to reflect: the book value of the firm, BV ; the mean and standard deviation of the active investor's control premium, μ_{Pre} and σ_{Pre} respectively; the mean and standard deviation of the active investor's control discount, μ_{Dis} and σ_{Dis} respectively; the mean and standard deviation of the small investor's idiosyncratic component, μ_{IC_s} and σ_{IC_s} respectively. Furthermore, the seller's goal is assumed to be to release the maximum amount of revenue from the sale of the firm. Finally, in order to have his goal fulfilled, the seller maintains the intention to receive and evaluate the messages forwarded by the investors.

3.2. Evaluation parameters

During negotiation, all the agents will typically receive, evaluate, generate, and send messages. The execution of such activities is based on a set of rules for inter-agent social behaviour. These rules consist of the following two main components: (i) a set of evaluation parameters; and (ii) a protocol for generating messages. In this section, we will concentrate on the evaluation parameters, whereas in Section 3.3 our focus will be on the negotiation protocol.

In order to establish a share price with the seller, the small investors evaluate proposals and counter-proposals based on a range of acceptable prices. For each small investor, this range is limited by the maximum value per share that he is willing to pay, $MaxP_{i_s}$, and the minimum price, $MinP_{i_s}$, that represents the first counter-offer that he will make to the seller upon receiving an unacceptable offer. These two values represent the small investors' evaluation parameters used during negotiation. We assume that $MaxP_{i_s}$ is given by the sum of the book value of the firm and the small investor's idiosyncratic component, that is: $MaxP_{i_s} = BV + IC_s$. On the other hand, $MinP_{i_s}$ is assumed to reflect the information that is publicly available and about which there is no uncertainty, i.e., the firm's book value: $MinP_{i_s} = BV$.

As with the small investor, the active investor's behaviour is guided by a price range limited by a maximum value that he is willing to pay, $MaxP_{i_a}$, and a minimum price that represents his first counter-offer to the seller, $MinP_{i_a}$. We assume that $MaxP_{i_a}$ depends on the book value of the firm and the difference between the active investor's control premium and discount, that is: $MaxP_{i_a} = BV + Pre - Dis$. On the other hand, $MinP_{i_a}$ will simply reflect the firm's book value, that is: $MinP_{i_a} = BV$.

Let us now turn to the seller's price range. Both the maximum value (i.e. the maximum price that the seller can reasonably ask) and the minimum value (i.e. the lowest price he will be willing to accept) depend upon whether the seller negotiates with the active or the small investors. Furthermore, both types of investors (small and active) can be put into competition with each other by approaching one type with an offer

that is contingent on the price paid by the other type at an earlier stage [47]. Thus, the price range used by the seller to negotiate with one type of investors will also reflect whether or not negotiation with the other type has already taken place.

When the seller negotiates with the active investor and shares have already been sold to the small investors, the maximum and minimum limits of the seller's price range, $\text{Max}P_{s/i_a}$ and $\text{Min}P_{s/i_a}$, are determined as follows. $\text{Max}P_{s/i_a}$ is assumed to depend 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 $\text{Max}P_{s/i_a}$ will depend on the ratio between the price established with the small investors, P_{s/i_s} , and the book value of the firm. Thus, 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:

$$\text{Max}P_{s/i_a} = \text{BV} + (\text{MaxIC}_a(P_{s/i_s}/\text{BV}))$$

We assume that MaxIC_a depends on the seller's beliefs concerning the probability distributions of the active investor's control premium and discount. More specifically, MaxIC_a will reflect the seller's guess at the means and standard deviations of both distributions:

$$\text{MaxIC}_a = (\mu_{\text{Pre}} + b\sigma_{\text{Pre}}) - (\mu_{\text{Dis}} - b\sigma_{\text{Dis}})$$

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

The minimum price at which the seller is willing to sell shares to the active investor, $\text{Min}P_{s/i_a}$, 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 $\text{Min}P_{s/i_a}$ is assumed to depend on the ratio between the market price at which shares have already been sold to the small investors, P_{s/i_s} , and the book value of the firm. As with MaxIC_a , we assume that MinIC_a reflects the seller's guess at the means and standard deviations of the active investor's control premium and discount as well as a constant sensitivity factor b . 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:

$$\text{Min}P_{s/i_a} = \begin{cases} \text{BV} + (\text{MinIC}_a(P_{s/i_s}/\text{BV})) & \text{if } [\text{BV} + (\text{MinIC}_a(P_{s/i_s}/\text{BV}))] > \text{BV} \\ \text{BV} & \text{otherwise} \end{cases}$$

where $\text{MinIC}_a = (\mu_{\text{Pre}} - b\sigma_{\text{Pre}}) - (\mu_{\text{Dis}} + b\sigma_{\text{Dis}})$.

When no negotiation has already taken place with the small investors, the seller's price range is determined as follows. The maximum price he may happen to ask will be: $\text{Max}P_{s/i_a}^* = \text{BV} + \text{MaxIC}_a$, where MaxIC_a is defined as above. On the other hand, the minimum acceptable price will be: $\text{Min}P_{s/i_a}^* = \text{BV} + \text{MinIC}_a$, where MinIC_a is defined as above.

When the seller negotiates with the small investors, his price range is determined as follows. We assume that, when a controlling block has already been sold, the

small investors are 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 [47]. In this case, the maximum price the seller may happen to ask, $\text{Max}P_{s/is}$, 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 investors' idiosyncratic component might be. The impact of MaxIC_s on $\text{Max}P_{s/is}$ will 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. Thus, 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:

$$\text{Max}P_{s/is} = \text{BV} + (\text{MaxIC}_s(P_{s/ia}/\text{BV}))$$

where $\text{MaxIC}_s = (\mu_{\text{IC}_s} + b\sigma_{\text{IC}_s})$.

Similarly, the minimum price the seller is willing to accept from the small investors, $\text{Min}P_{s/is}$, is given by:

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

where $\text{MinIC}_s = (\mu_{\text{IC}_s} - b\sigma_{\text{IC}_s})$.

Finally, when no negotiation has already taken place with the active investor, the seller's maximum and minimum prices used during negotiation with the small investors will be respectively:

$$\text{Max}P_{s/is}^* = \text{BV} + \text{MaxIC}_s \quad \text{and} \quad \text{Min}P_{s/is}^* = \text{BV} + \text{MinIC}_s$$

where MaxIC_s and MinIC_s are defined as above.

3.3. Negotiation protocol

Based on their beliefs, the agents will attempt to achieve their (top-level) goals by fulfilling their own intentions. In turn, fulfilment of intentions generates social interaction. The negotiation protocol refers to the social interaction mechanisms that, once instantiated with the evaluation parameters, are used by the agents to evaluate incoming messages and forward new ones. The protocol builds on the following main assumptions: (i) each agent will remember what message has been sent and who it has been forwarded to; (ii) investors do not negotiate with each other, but only with the seller; (iii) negotiation begins with the seller asking an investor (either small or active) the maximum value of his price range; and (iv) the first counter-offer the investor (either active or small) makes to the seller is the minimum acceptable value of his price range.

After a message with an offer has been sent, agents will typically wait for a response. Upon receiving an offer, each agent will evaluate it. An offer can be either accepted or modified and sent back. Ideally, negotiation can continue until an agreement is reached. However, in most cases it is not reasonable to assume that negotiation will continue forever. Although persistent in keeping their joint commitment to

reaching an agreement, the agents are not fanatical and, after a certain number of messages, they will give up negotiation [18,62]. Given this, we can distinguish between three different responses that the agents may give after receiving a message: (i) acceptance; (ii) modification; and (iii) rejection.

We assume that a message is accepted in the following two situations: (i) the seller's offer to the small investors, $p_{s/is}$, or to the active investor, $p_{s/ia}$, is lower than, respectively, $\text{Max}P_{is}$ or $\text{Max}P_{ia}$; ¹ and (ii) the counter-offer forwarded to the seller by the small investors, $p_{is/s}$, or by the active investor, $p_{ia/s}$, is greater than the correspondent minimum price the seller is willing to accept.

Messages that are not accepted are usually modified and returned to the other party. However, under our assumption of limited persistence, unacceptable messages can not only be modified, but also rejected. The agent who rejects a message will generate the intention to exit negotiation. In this case, negotiation fails and no final agreement is reached. We assume that a message will be rejected if: (i) it is not acceptable; and (ii) the number of messages (offers and counter-offers) that have already been exchanged among the agents is greater than an integer value n^* .

However, when the number of messages exchanged is less than n^* , the agents will keep negotiating and will modify the unacceptable offers received. Modifications are made according to a *price movement strategy*. Each price movement strategy depends on the interplay between the offer sent originally and the counter-offered received. Firstly, should the seller receive at time t' an unacceptable offer from the small or active investor, $p'_{is/s}$ or $p'_{ia/s}$ respectively, he will modify this offer and send back at time t ($t > t'$) a new counter-offer, $p^t_{s/is}$ or $p^t_{s/ia}$. Specifically,

$$p^t_{s/is} = \begin{cases} p''_{s/is} - c(p''_{s/is} - p'_{is/s}) & \text{if } [p''_{s/is} - c(p''_{s/is} - p'_{is/s})] > \text{Min}P^*_{s/is} \text{ (or } \text{Min}P_{s/is}) \\ \text{Min}P^*_{s/is} \text{ (or } \text{Min}P_{s/is}) & \text{otherwise} \end{cases}$$

and

$$p^t_{s/ia} = \begin{cases} p''_{s/ia} - c(p''_{s/ia} - p'_{ia/s}) & \text{if } [p''_{s/ia} - c(p''_{s/ia} - p'_{ia/s})] > \text{Min}P^*_{s/ia} \text{ (or } \text{Min}P_{s/ia}) \\ \text{Min}P^*_{s/ia} \text{ (or } \text{Min}P_{s/ia}) & \text{otherwise} \end{cases}$$

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 the small or active investor receive at time t' an unacceptable offer from the seller, $p'_{is/s}$ or $p'_{ia/s}$ respectively, he will modify the offer received in the following way:

$$p^t_{is/s} = \begin{cases} p''_{is/s} - c(p''_{is/s} - p'_{is/s}) & \text{if } [p''_{is/s} - c(p''_{is/s} - p'_{is/s})] < \text{Max}P_{is} \\ \text{Max}P_{is} & \text{otherwise} \end{cases}$$

¹ 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. For the sake of simplicity, we have not attempted to represent this aspect in our model.

and

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

where $p_{is/s}''$ and $p_{ia/s}''$ are the offers sent at time t'' ($t'' < t' < t$) to the seller by the small and active investor respectively, and c is a concession rate reflecting the investors' 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.

3.4. Selling strategies for going public

Building on mainstream literature on staged equity financing and behavioural finance (see Section 2), we assume that the sale of the firm can be organised according to one of the following selling strategies for going public that have been most used in countries with developed markets [4,8,47,53]:

- (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 our second set of simulations (Section 5.2), our main goal will be to gain a better understanding of how information and the agents' cognition interact to impact on the performance of social influence.

To this end, we will introduce the following variant of strategy B, in which more information is used to negotiate with the active investor:

- (v) a sequential sale in which the IPO is followed by an informationally enriched sale of a controlling block (strategy B*).

In what follows, we will briefly describe each of the above five strategies.

(i) *Strategy A.* The seller interacts with two types of potential investors: the active investor and the small investors. The first task is to establish a price for the sale of shares to the active investor. To this end, the seller will negotiate with the active investor by letting him have an offer with a proposed price without any further information. 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. However, 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

investors. To this end, he will send them an offer with a proposed price and will let them know whether or not a controlling block has been sold at an earlier stage. Furthermore, through his evaluation parameters (Section 3.2), the seller will be in a position to make the price offered to the small investors contingent on prior negotiation with the active investor. If the seller is successful in negotiating with the small investors, then a price will be agreed upon, and both the seller and the small investors will update their mental states. As a result, a transaction will occur. Otherwise, should negotiation fail, no ownership transfer would take place.

(ii) *Strategy B.* The seller interacts with both the small investors and the active investor. The first task is to establish a price for the sale of shares to the small investors. To this end, the seller will typically negotiate with the small investors by forwarding them a price with no additional information. If this stage is successful, both the seller and the small investors will update their mental states by generating the belief that a market 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 be unsuccessful, no ownership transfer will occur. The seller then proceeds to establish a share price with the active investor by offering him a price and letting him know whether or not an IPO has occurred. Again, selling a controlling block is made contingent on previous economic transactions since the seller's evaluation parameters for negotiation with the active investor depend on whether or not prior negotiation with the small investors has already 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. A transaction will then take place. Conversely, should negotiation be unsuccessful, no transfer of control will occur.

(iii) *Strategy C.* The seller interacts simultaneously with the active investor and the small investors. Thus, two prices need to be established for the sale of shares to both types of investors. To this end, the seller negotiates and tries to influence the investors by approaching them simultaneously, sending them an offer, and letting them know of the ongoing parallel negotiation for both a controlling block and dispersed shareholdings. 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 active investor gets control only with a bid that is higher than the bid of the small investors competing for the shares [47]. Specifically, if the price established with the active investor is greater than the market price for dispersed shareholdings, then the seller proceeds to sell the controlling block, whereas the remaining allotment of shares will be sold to the small investors. Conversely, 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.

(iv) *Strategy D.* With this strategy the seller interacts only with the small investors. Negotiation is therefore aimed at establishing only a market price for dispersed shareholdings. The seller will typically negotiate with the small investors by letting them have an offer with a proposed price and no further information. If successful, negotiation will eventually lead the seller and the small investors 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.

(v) *Strategy B**. This is a variant of strategy B, in which the seller tries to more convincingly exert social influence upon the active investor by providing him with more detailed information on which the price being offered is made contingent. More specifically, the seller will approach the active investor with an offer comprising not only a price for the controlling block but also the market price established with the small investors. We assume that, once the active investor has updated his mental state with the market price for dispersed shareholdings, he will also modify his evaluation parameters in the following way. If $P_{s/is}$ is the market price for dispersed shareholdings communicated by the seller, the active investor will set his minimum acceptable price $\text{Min}P_{ia}$ at $P_{s/is}$ rather than at the book value BV. On the other hand, the maximum price, $\text{Max}P_{ia}$, will be set at $P_{s/is} + \text{Pre} - \text{Dis}$ instead of $\text{BV} + \text{Pre} - \text{Dis}$. Since the market price will never be lower than the book value of the firm (Section 3.2), these 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.

In what follows, we are not concerned with issues related to the seller's freedom and ability to choose what selling strategy to use. Nor are we interested in exploring the extent to which the market determines what strategy is actually adopted. Our interest in the above five selling strategies simply lies in their potential for offering a straightforward way to manipulate the structural components of social influence. In fact, they provide the structural patterns of social relations in which social influence can be exercised and agreements achieved among actors in different social positions. Moreover, each selling strategy, by reflecting a distinct pattern of social relations and interactions, is associated with a distinct communication network. In this respect, different strategies reveal different communication patterns in which actors are connected through different pieces of information, of different quality, and forwarded in different temporal orders. Building directly on the proposed taxonomy of strategies, we are therefore in a position to elucidate the systematic implications of the social structure in terms of the exercise of social influence and the generation of agreements.

4. Virtual experiments: design and methodology

The model introduced in the previous section will now be used to undertake a series of virtual experiments. Our aim is to examine the individual and combined effects of social structure, cognition and social interaction on social influence and the generation of interpersonal agreements. We organise our simulations into two sets. In the first set, we use a 4×6 design where influence structural patterns and social interaction are manipulated. On the one hand, influence patterns are manipulated using selling strategies A, B, C and D. On the other hand, social interaction is operationalised through the negotiation protocol and its correspondent concession rate c ($0 \leq c \leq 1$) that agents use to instantiate their price movement strategy for modifying the unacceptable offers received during negotiation (Section 3.3). To allow

for a range of 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.

In the second set of simulations, we use a 5×6 design where influence structural patterns and cognition are manipulated. Here, we extend the range of possible influence patterns by including strategy B* into the set of alternative selling strategies that can be used. This will enable us to get a deeper insight into the role that information plays in negotiation. More specifically, we will investigate whether, when strategy B* is adopted, the active investor becomes more motivated to accept the offer received than would be the case if the seller had forwarded only a price without any further information that might justify it. Thus, a comparative analysis of strategies B and B* will lead to an evaluation of the extent to which sharing additional information makes the exercise of social influence more effective and efficient. The 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_{IC_s} (Section 3.2). The higher the standard deviation, the less accurate the seller's beliefs. The standard deviation will be either 0, 0.1, 0.2, 0.3, 0.4, or 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.

In both sets of simulations, the dependent variable—influence performance—is operationalised through the performance of the sale of the firm, which in turn is rated by two measures expressing, respectively, the effectiveness and efficiency of the sale: (i) the revenue raised; and (ii) the length of time the agents take to get to an agreement. A particular way of exercising social influence, defined in terms of a combination of an influence structural pattern with a protocol or with a degree of cognitive accuracy, turns out to be more effective than alternative ways if it allows the seller to raise higher revenue. Conversely, influence becomes more efficient if it is associated with a lower amount of time it takes the agents to reach an agreement. 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.

Negotiation is generated using the UM-PRS architecture² (see Lee et al. [43] for details), which, in turn, we extended and enriched by adding the computational

² The University of Michigan Procedural Reasoning System (UM-PRS) is a high-level agent architecture in which: (i) the agents are represented using notions of mental states; (ii) the agents' actions are a function of these mental states; (iii) the agents' mental states may change over time; and (iv) mental state changes are formalised in terms of axioms and inference rules. UM-PRS allows the user to define and control the agents' mental states. In addition, the user can test different patterns of interagent social behaviour and argumentation in the negotiation process.

components introduced in Section 3. This simulation test-bed offered us the chance to examine systematically the interactions among, and the effects of, a variety of structural, social and cognitive factors on negotiation performance that have long concerned researchers. Nevertheless, in this paper our focus will be limited to structural patterns of social relations, the negotiation protocol and agents' cognitive accuracy. Multiple simulations were conducted for each set of parameter values investigated. However, for ease of communication, we will present below selected results in graphic form based on the following set of common assumptions: (i) the total number of shares the seller can transfer either to the active investor or to the small investors is set at 100; (ii) control is reached with 50% of the shares, i.e., $m = 0.50$; ³ (iii) the number of messages n^* that represents the threshold after which agents will give up negotiation is set at 40; ⁴ (iv) during negotiation, all the agents use the same protocol, and therefore the concession rate c is assumed to be the same across the three types of agents; ⁵ and (v) we assume that the seller's degree of cognitive accuracy is reflected by the same value for the standard deviation across all the random variables, namely $\sigma_{\text{Pre}} = \sigma_{\text{Dis}} = \sigma_{\text{IC}_s}$. ⁶

To ensure that the findings reported here are not idiosyncratic and do not depend primarily on the above assumptions, we conducted a stability analysis for a variety of instantiations of the model. This effort included a sensitivity analysis to determine that performance does not vary substantially across changes in those parameter values held constant in the results reported here, and across changes in many combinations of them. However, even with a predefined set of initial structural conditions (number of shares, control percentage, commitment threshold, homogeneity of protocols and standard deviations across agents), there still exists considerable variation among the possible negotiation processes that match these characteristics. Thus, to ensure that the results reported here reflect the underlying structure of the model and not merely particular realisations of a highly stochastic process, we used a Monte Carlo approach to average out differences arising from distinct independently specified instantiations of all the remaining structural parameters, including those characterising the agents' mental states (e.g. beliefs; sensitivity factor). Our findings

³ For simplicity, the analysis does not consider more realistic situations 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 [47].

⁴ 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.

⁵ A more in-depth analysis should take into account those scenarios in which the seller, the active investor and the small investors use different concession rates and, therefore, differ in terms of using more or less conciliatory negotiation protocols. In this case, the impact of different combinations of the agents' negotiation protocols upon negotiation performance could be systematically studied. We leave such refinements for future work.

⁶ Obviously, the seller could have different degrees of cognitive accuracy depending on the object of his beliefs. For example, he might be perfectly accurate about the active investor's position, but highly inaccurate when it comes to expressing a belief about the small investors. The impact of different combinations of standard deviations is another topic we leave for future work.

are based on 300 simulations for each combination of the independent variables. This amounts to 7200 runs for the first set of simulations, and 9000 for the second one.

5. Results and analysis

One benefit of using agent-based computational models for studying organisational processes is that they can provide researchers with data at both the individual agent's and the group's level. Here, we use our agent-based model of negotiation to explore the role of social influence in the generation of interpersonal agreements. We focus on two questions. First, to what extent does social structure interact with social behaviour to impact on influence? To address this question, we undertake a first set of simulations in which we examine the individual and combined impact that the selling strategy and the concession rate have on the effectiveness and efficiency of the sale (Section 5.1). The second question we explore is the impact of social structure and cognition on social influence. To analyse this issue, we undertake a second set of simulations in which we examine the individual and combined effect that the selling strategy and the accuracy of the seller's beliefs have on the effectiveness and efficiency of the sale (Section 5.2). This will shed light on what are normally hidden cognitive processes. Results will allow us to investigate the relation between cognitive accuracy and information and, in turn, how this relation affects the relative performance of various structural patterns of social influence that differ in terms of the type and amount of information conveyed.

5.1. Negotiation protocol and social interaction

We begin by exploring the interrelationship between selling strategy and negotiation protocol, and their impact on the performance of the sale. In so doing, we will firstly undertake a comparative analysis of the selling strategies across different protocols, and secondly we will determine the impact that the protocol has on the effectiveness and efficiency of the sale when different selling strategies are used.

Fig. 1 shows the impact that each combination of selling strategy and concession rate has upon the revenue raised in the aggregate sell of shares. In terms of a comparative analysis between selling strategies across different protocols, we find that, throughout all protocols, public offering at a uniform price (strategy D) is the selling strategy that keeps revenue at the lowest levels. Conversely, 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. This suggests that social influence becomes more effective as the heterogeneity of the actors involved increases. In terms of our model, there appear 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 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




Fig. 1. The impact of structural patterns of social influence on negotiation effectiveness with different protocols.

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, the seller is in a position to ask the small investors for a higher price that in turn reflects the public good associated with the active investor's monitoring activities.

Not only does Fig. 1 suggest that the sale of a firm should be designed to promote the participation of an active investor. It can also be used to derive interesting implications as to how the active investor should be involved in the sale when different protocols are adopted. As it turns out, strategy A is especially effective when the agents use low concession rates. At medium and high levels of this parameter, strategy B generates higher revenue than the alternative selling strategies. It produces the best results between 0.10 and 0.125 and between 0.15 and 0.20. It has a local maximum at 0.10 and a local minimum at 0.175. However, even at 0.175, strategy B is more effective than the other strategies. Finally, the parallel strategy (strategy C) gives lower revenue than both strategies A and B (except for the lowest value of the concession rate), but higher than the simple public offering (strategy D).

An inspection of these results, therefore, 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. In this case, the sale is, on average, more effective because 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 [47]. Overall, what these results suggest is that the effectiveness of social influence depends not only on the heterogeneity of the agents involved, but also on the temporal order in which the agents are approached and on the type of information that is used to exercise influence. This emphasises the impact of the social structure

on the performance of social influence. How effectively an agent is influenced depends on the network of social relations in which it is located, on the order in which interactions occur within the network, and on the type of information that these interactions convey.

Fig. 2 shows the impact that each combination of selling strategy and concession rate has upon the number of messages sent. In terms of a comparative analysis between strategies, these results suggest that, throughout all negotiation protocols, public offering keeps the number of messages at the lowest levels. Since in public offering the seller communicates only with the small investors, we find that the number of messages is positively correlated with the degree of heterogeneity of the agents among which an agreement must be made. Thus, there is a trade-off between effectiveness and efficiency as the investors' heterogeneity is manipulated. In fact, on average, an increase in effectiveness resulting from the participation of an active investor in the sale can be obtained only at the cost of lower efficiency resulting from the increased heterogeneity of the agents involved.

Moreover, Fig. 2 shows that, throughout all negotiation protocols, parallel strategy C works better than both sequential strategies A and B, but is worse than strategy D. This means that the negative effects (in terms of more messages) of an increase in the heterogeneity of the agents involved can be mitigated if the seller interacts simultaneously with all the investors rather than with sequential negotiations. Finally, at low and medium values of the concession rate, strategy B is more efficient than strategy A, whereas at high values of the concession rate strategy A is better than strategy B. This suggests that the relative benefits (in terms of efficiency) of different ways of interacting sequentially depend on the agents' social behaviour and interaction style. When they are conciliatory and interact sequentially, an agreement can be reached more quickly if negotiation with the small investors is subsequent to




Fig. 2. The impact of structural patterns of social influence on negotiation efficiency with different protocols.

negotiation with the active investor. Conversely, when the agents decide to be more persistent and stick to their own initial positions, making information contingent on prior negotiation with the small investors and using this information to influence the active investor allows the seller to reach an agreement more quickly.

In combination, Figs. 1 and 2 reveal the existence of individual and combined effects of selling strategy and negotiation protocol on negotiation performance. Under each of the protocols here considered, the choice of a selling strategy has an impact on the revenue raised and the number of messages sent. In particular, strategy B, in general, turns out to be 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. Furthermore, under each of the selling strategies, the choice of the negotiation protocol impacts upon revenue and number of messages. When the agents try to influence one another by adhering to a specific structural pattern of interaction, performance is ultimately affected by how persistent the agents are in sticking to their initial views. To gain a deeper insight into the relative impact of negotiation protocol on performance when different selling strategies are used, we performed a regression analysis on the simulation data shown in Figs. 1 and 2.

The regressions in Table 1 indicate that, on average, the protocol has a strong negative effect on effectiveness: the more conciliatory the agents are, the lower the price at which the firm is sold. Here an obvious question is: What is driving revenue 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 significant the modification that is made (Section 3.3). This, in turn, might have two alternative consequences: (i) the seller makes more concessions and the firm is sold at

Table 1
Regression analysis for performance on negotiation protocol (concession rate)

Independent variables	Effectiveness (revenue)			Efficiency (number of messages)		
	Standardised regression coefficient	R ²	Adjusted R ²	Standardised regression coefficient	R ²	Adjusted R ²
Protocol ^a	−0.965* (−7.334)	0.931	0.913	−0.927** (−4.954)	0.860	0.825
Protocol * Strategy A	−0.995*** (−19.064)	0.989	0.986	−0.949* (−6.036)	0.901	0.876
Protocol * Strategy B	−0.808* (−2.704)	0.652	0.565	−0.945* (−5.774)	0.893	0.866
Protocol * Strategy C	−0.988*** (−12.742)	0.976	0.970	−0.655 (−1.732)	0.429	0.286
Protocol * Strategy D	−0.831* (−2.992)	0.691	0.614	−0.645 (−1.722)	0.529	0.386

Note: For each cell, this is a regression of the six means in Figs. 1 and 2. Numbers in parentheses are *t*-statistics. **p* < 0.05; ***p* < 0.01; ****p* < 0.001.

^a Regression coefficients for the average performance across the four selling strategies.

a lower price; or (ii) the investors make more concessions and the firm is sold at a higher price. Therefore, the decrease in revenue as the concession rate increases expresses the fact that, on average, the impact of the modification made by the seller on revenue is stronger than the impact of the modification made by the investors.

When interaction effects between protocol and selling strategy are controlled, Table 1 allows for a comparative analysis of the impact of the negotiation protocol on the sale across different selling strategies. In particular, the negative impact of protocol on revenue is stronger if strategies A and C are used, whereas it turns out to be weaker under strategies B and D. Furthermore, an increase in the concession rate generates the most significant decrease in effectiveness in connection with strategy A, and the least significant one with strategy B. Interacting first with the small investors and then with the active investor enables the seller to exercise social influence in the most beneficial manner as his social behaviour becomes more conciliatory. This means that, as the seller becomes more conciliatory, the overall benefits of exercising influence through the exchange of information are greater when the information is made contingent on prior interactions with the small investors (strategy B) than when it is made contingent on prior interactions with the active investor (strategy A) or on simultaneous interactions (strategy C). Finally, when only the small investors are involved in the sale (strategy D), the negative effects of protocol on revenue are less significant than those produced when strategies A and C are used, but stronger than the effects generated in connection with strategy B. This suggests that, in general, as the concession rate increases, exercising influence on less heterogeneous agents allows the seller to decrease the loss in revenue associated with his more conciliatory behaviour. However, as the concession rate increases, this loss can be further reduced if an active investor is involved in negotiation and the seller decides to influence him using information contingent on prior interactions with the small investors (strategy B).

Perhaps the most interesting pattern in Table 1 comes from comparison of the same variable across outcomes. Whereas the protocol has negative effects on effectiveness, it has a positive impact on efficiency. On average, the number of messages decreases as the concession rate increases. That is, the more significant the concessions the agents make to one another, the lower the number of messages that need to be sent in order for an agreement to be reached. This suggests that there is a trade-off between effectiveness and efficiency. On average, when the agents choose to be more conciliatory with each other, they can reach an agreement more quickly, but the firm is sold at a lower price. However, as with effectiveness, when interaction effects between protocol and selling strategy are controlled, a comparative analysis shows that the relative impact of protocol on efficiency changes across selling strategies. Strategies A and B emphasise the role of the protocol, whereas strategies C and D turn out to mitigate the effects of a change in the concession rate. Strategy A is the strategy under which an increase in the concession rate generates the most significant increase in efficiency. This suggests that interaction with the small investors subsequent to negotiation with the active investor allows the seller to obtain the greatest benefits from a more conciliatory social behaviour in terms of reaching an agreement more quickly. However, this selling strategy will also generate the most significant

loss in revenue. The seller is therefore confronted with a trade-off. When he decides to improve efficiency with a more conciliatory behaviour, he must choose whether to get the most from this decision in terms of reducing the time to agreement but at the same time give up the most in terms of loss of revenue (by using strategy A), or to give up some benefits in terms of time to agreement but at the same time get the most in terms of the lowest possible loss of revenue (by using strategy B). Finally, the positive effects of protocol on efficiency are the least significant when only the small investors are involved in the sale (strategy D). This suggests that the more heterogeneous the negotiating agents are, the more beneficial the overall effects of being more conciliatory on the time it takes to make an agreement.

5.2. Cognitive accuracy and the value of information

So far, we have analysed the impact that selling strategy and negotiation protocol have on performance, under the assumption that no shift occurs in the agents' cognitive representations of their environment. In the following set of simulations, we extend our analysis to a setting in which influence is exercised under different degrees of cognitive accuracy.

Fig. 3 shows the impact that each combination of selling strategy and degree of cognitive accuracy has upon the revenue raised in the aggregate sell of shares (negotiation effectiveness). We begin the analysis of these findings by investigating the effects that an increase in the amount of the information used to exert social influence has on effectiveness, and to what extent these effects depend on the agents' cognitive accuracy.

At high degrees of cognitive accuracy, strategy B does not work well. This result contrasts with the performance of strategy B* that, at high degrees of cognitive accuracy, gives the highest revenue. How can this phenomenon be explained? The an-




Fig. 3. The impact of structural patterns of social influence on negotiation effectiveness at different levels of cognitive accuracy.

swer is found in the relationship that the value of an increase in information has with the agents' self-confidence during negotiation, which in turn depends on how accurate the agents' cognitive representations are. At higher degrees of cognitive accuracy, the agents are more self-confident and, therefore, become more reluctant to make proposals and counter-proposals that deviate from the average value of their price range. This, in turn, might have two consequences: (i) it might be the case that the seller's price range will never overlap with either the active investor's or the small investors' 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, it may become more difficult to settle for a price that is acceptable to both parties. In such situations, agreement may be reached in a more effective way if the agents influence one another by exchanging more information so as to support their requests during negotiation.

Fig. 3 indicates that, at high degrees of cognitive accuracy, the more informative strategy B* is more effective than its variant B as well as strategy C. For example, when the seller's beliefs are perfectly accurate (i.e. the standard deviation is zero), strategy B* gives the highest revenue, about 1394.98, whereas strategy B gives only 600 since no agreement is made between the seller and the active investor. This is because, with strategy B*, the seller influences the active investor by communicating an additional piece of information so as to justify the offer made. In turn, a greater amount of information helps the two agents to reach an agreement. These results have a two-fold implication. Firstly, on average, increasing the amount of information for social influence purposes improves effectiveness. Secondly, the impact of increasing the amount of information on effectiveness depends on the agents' cognitive accuracy. In particular, our findings suggest that the benefits (in terms of more effectiveness) of using more information are emphasised as cognitive accuracy increases.

Another result that emerges from Fig. 3 is that, at medium and low degrees of cognitive accuracy, a public offering of all shares at a uniform price is associated with the worst performance. At such degrees of accuracy, promoting the participation of an active investor in the sale and, therefore, increasing the heterogeneity of the agents involved improve the effectiveness of negotiation. In particular, at medium and low degrees of cognitive accuracy, strategy C mostly dominates strategy D. Thus, with strategy C, the small investors, who are informed of the ongoing simultaneous negotiation with an active investor, are more effectively influenced to accept the offer than they are with strategy D where social influence is exerted only by offering a price with no additional information. On the other hand, strategy A dominates strategy C: providing the small investors with information about the controlling block sold at an earlier stage is therefore more effective than providing them with information about a parallel negotiation simultaneously carried out with the active investor. The impact of social influence that the seller exerts on the active investor shows a similar pattern. Most of the time, at medium and low levels of cognitive accuracy, strategy B is more effective than strategy C. Again, this result is explained by the additional information conveyed by strategy B and used to persuade the active investor to accept the price offered.




Fig. 4. The impact of structural patterns of social influence on negotiation efficiency at different levels of cognitive accuracy.

Fig. 4 shows the impact that each combination of selling strategy and degree of cognitive accuracy has upon the number of messages sent (negotiation efficiency). Throughout all levels of cognitive accuracy, the lowest number of messages is associated with strategy D. Once again, this result suggests that the number of messages is positively correlated with the heterogeneity of the agents among which an agreement must be made (public offering involves only small investors). Thus, at medium and low degrees of cognitive accuracy, we find a trade-off between effectiveness and efficiency as the investors' heterogeneity is manipulated. In fact, at such degrees of accuracy, the participation of an active investor in the sale brings about an increase in effectiveness only at the cost of lower efficiency.

Not only does Fig. 4 suggest that, to improve efficiency, the participation of an active investor in the sale should be promoted. It also allows for a comparative analysis between alternative strategies in which an active investor is involved, thus suggesting a set of implications as to how his participation should be designed. On average, strategy B* requires fewer messages than strategy B. This suggests that, when the active investor is approached after negotiation with the small investors, increasing the amount of information helps the seller to influence the active investor and reach an agreement more quickly. However, not only does Fig. 4 show that using more information speeds up negotiation; it also reveals that these positive effects of information are emphasised at high degrees of cognitive accuracy. Furthermore, we find that, on average, strategy B* also works better than strategy A in terms of keeping the number of messages at a lower level. Thus, social influence is more efficient when information is made contingent on prior negotiation with the small investors rather than on prior negotiation with the active investor. Since what information is used and in what order it is forwarded depend on the pattern of interaction in which influence unfolds, once again this result indicates the existence of an interaction between social structure and social influence. In addition, at medium and low

levels of cognitive accuracy, strategy B* generates fewer messages than the parallel strategy. Interacting sequentially rather than simultaneously has different implications in terms of the information that can be produced and exchanged, and this in turn generates differing results in terms of how quickly influence can be exercised. Finally, Fig. 4 shows that, on average, strategy A is less efficient than strategy C: here, the effect of influencing the small investors with information contingent on prior negotiation with the active investor seems to be weaker than the effect of influencing the active investor with information contingent on simultaneous negotiation with the small investors.

To gain a deeper insight into the relative effect of cognitive accuracy on performance across all five selling strategies, we conducted a regression analysis on the simulation data shown in Figs. 3 and 4. Table 2 suggests that, on average, cognitive accuracy has a positive effect on effectiveness: the higher the degree of cognitive accuracy, the higher the revenue raised. We noted above that, as their cognitive accuracy increases, the agents become more self-confident and increase the concentration of the possible prices they could offer during negotiation around an average value. In the light of this, one way of interpreting the result in Table 2 is that, the higher the degree of cognitive accuracy, the higher the minimum price that the seller is willing to accept from either the small investors or the active investor and, hence, the higher the price that may eventually be agreed upon. Thus, on average, as the seller's cognitive representations become more accurate, the influence that he can exercise on his counter-parts becomes more effective in terms of a higher price at which the firm is sold.

Table 2
Regression analysis for performance on cognitive accuracy

Independent variables	Effectiveness (revenue)			Efficiency (number of messages)		
	Standardised regression coefficient	R ²	Adjusted R ²	Standardised regression coefficient	R ²	Adjusted R ²
Cognitive accuracy ^a	0.344 (0.733)	0.118	−0.102	0.783 (2.519)	0.613	0.517
Cog. Acc. * Strategy A	−0.665 (−1.779)	0.442	0.302	0.912* (4.440)	0.831	0.789
Cog. Acc. * Strategy B	−0.721 (−2.081)	0.520	0.400	0.905* (4.252)	0.819	0.774
Cog. Acc. * Strategy B*	0.987*** (12.120)	0.973	0.967	0.115 (0.232)	0.013	−0.233
Cog. Acc * Strategy C	0.861* (3.378)	0.740	0.676	−0.680 (−1.857)	0.463	0.329
Cog. Acc. * Strategy D	0.851 (3.368)	0.780	0.696	−0.670 (−1.867)	0.473	0.339

Note: For each cell, this is a regression of the 6 means in Figs. 3 and 4. Numbers in parentheses are *t*-statistics. **p* < 0.05; ***p* < 0.01; ****p* < 0.001.

^aRegression coefficients for the average performance across the five selling strategies.

This result refers to the average performance across the five selling strategies. However, when we look at the interaction effects between accuracy and selling strategy, Table 2 complements our findings in three important respects. First, the benefits of an increase in accuracy are emphasised when strategy B* is used. As noted above, strategy B* is associated with the case where more information is exchanged among the agents. When more information is exchanged, the effects of an increase in cognitive accuracy can be fully exploited, and this amplifies the gain that can be obtained in effectiveness. Thus, information and cognitive accuracy complement each other in their impact on effectiveness: whereas the effects of using more information are emphasised when accuracy is high, the effects of an increase in accuracy are emphasised when more information is used. Second, having more accurate beliefs has positive effects on effectiveness also when strategies C and D are used, but in this case the effects are lower than with strategy B*. This suggests that, as cognitive accuracy increases, the seller is most successful in increasing the revenue obtained from the sale when he influences the active investor by letting him know the market price previously established with the small investors. In contrast, as cognitive accuracy increases, the increase in revenue becomes lower if the seller negotiates simultaneously with the active investor and the small investors, or if only the small investors are involved in negotiation.

Finally, Table 2 shows that the impact of cognitive accuracy on effectiveness is negative when strategies A and B are used. Under these strategies, increasing the accuracy of the seller's beliefs generates lower revenue. As Fig. 3 suggests, this result is determined by the fact that at high levels of accuracy, both strategies do not work well. When the amount of the information exchanged is not sufficiently high and social influence is exercised in sequential interactions using information contingent on previous negotiations, increasing accuracy may not be enough to improve the agents' self-confidence. In this case, an increase in cognitive accuracy may even make the seller aware that the information he can use to influence his counter-parts is inadequate, and this can make him inclined to make more concessions and accept a lower price. This negative effect of accuracy on effectiveness is emphasised particularly when strategy B is used. This means that as accuracy increases the seller becomes less self-confident when he has to interact first with the small investors and then with the active investors rather than in the opposite order. Self-confidence is therefore most undermined precisely under that selling strategy that, when the right amount of information is exchanged, allows an increase in cognitive accuracy to generate the most significant increase in revenue.

When we look at the impact of cognitive accuracy on efficiency, Table 2 shows that, on average, the more accurate the agents' beliefs are, the longer it takes to reach an agreement. As we noted above, the reason for this is that, as cognitive accuracy increases, the agents become more self-confident and, accordingly, more reluctant to offer prices that deviate from an average one. This makes it more difficult for the agents to converge on a final price, thus increasing the number of messages that need to be exchanged before an agreement is reached. Thus, Table 2 shows that there is a trade-off between effectiveness and efficiency as accuracy is manipulated. In fact, on average, an increase in effectiveness as a result of an improvement in cognitive accuracy can be obtained only at the cost of lower efficiency.

To explore this trade-off, we now look at how cognitive accuracy and selling strategy interact to impact on efficiency. Table 2 shows that an increase in accuracy has the strongest negative impact on efficiency when sequential strategies A and B are used. However, things are different when strategy B* is used. Here, a change in accuracy does not have significant effects on efficiency. Therefore, not only does strategy B* require, on average, fewer messages than strategies A and B, as shown in Fig. 4, but it also turns out to be the most appropriate selling strategy under which accuracy can be increased to obtain the maximum increase in revenue, yet minimising the increase in the time it takes to reach an agreement. This result suggests that the positive effects on efficiency of using more information to exercise social influence are likely to mitigate the negative effects of having more accurate beliefs and therefore being more self-confident in negotiation. When additional information is exchanged, the seller can easily impact on the active investor's mental state: even though he becomes more self-confident as a result of more accurate beliefs, nonetheless it takes only a few more messages to effectively influence the active investor to accept a higher price. In contrast, when the amount of information is limited, like with strategies A and B, increasing accuracy not only generates lower revenue, but also has a strong negative impact on efficiency. The reason is that, as accuracy increases, the seller adjusts his price range by decreasing the maximum price and increasing the minimum price he will ask in negotiation. Whereas the former adjustment determines acceptance of a final lower price, the latter generates an increase in the number of messages that need to be exchanged before an agreement is reached.

Finally, Table 2 shows that when strategies C and D are used, the impact of accuracy on efficiency becomes positive. Under both strategies, increasing accuracy determines an improvement in efficiency. However, the reasons for this pattern are different depending on what strategy is used. Under strategy C, the adjustments in the seller's price range induced by an increase in his accuracy are such that the effects of decreasing the maximum acceptable price (in terms of less time to agreement) are stronger than the effects of increasing the minimum acceptable price (in terms of more time to agreement). Thus, when the seller chooses a simultaneous pattern of interaction rather than a sequential one, the effects of an increase in cognitive accuracy are two-fold. On the one hand, the increase in the minimum acceptable price eventually determines higher revenue; on the other, the decrease in the maximum price speeds up the process by making the seller's possible offers more likely to be accepted. This explains the increase in both effectiveness and efficiency in strategy C as accuracy increases.

The same pattern emerges with strategy D. However, in this case the positive effects of increasing accuracy on efficiency are connected to the fact that the seller is interacting only with one class of investors, namely the small investors. In general, interacting with fewer classes of investors reduces the opportunities of selling the firm. Thus, when an increase in accuracy is combined with a decrease in the heterogeneity of the agents involved in negotiation, the seller's price adjustments are made in such a way that his disposition to require a higher price as a result of his increased self-confidence is mitigated by the need to reduce the risk of jeopardising the whole operation by making too high offers to only one class of investors. This explains the

two-fold effect of an increase in accuracy under strategy D: while the increase in the seller's minimum acceptable price guarantees a sufficiently high final price, the decrease in his maximum price speeds up the process by inducing the seller to offer a sufficiently low price that will be accepted more quickly.

6. Discussion and conclusions

Research on social influence can typically be organised into a variety of distinct problem areas, such as social conformity [24,26,27], group polarisation [9,56,81], minority influence [54,57], power [20,52] and persuasion [10,58]. Despite the differences in their theoretical and methodological orientations, scholars across all these research areas can be unified by a widespread acceptance of the same paradigmatic model for conceptualising social influence. At the heart of this model, usually referred to as the “dual-process model” [24], is the recognition that there are two conceptually distinct processes of social influence that often co-occur in every-day life: informational and normative social influence. *Informational* influence is defined as influence to accept information from another as evidence about reality. *Normative* influence is influence to conform to the positive expectations of another in order to gain approval or avoid rejection. Thus, whereas informational influence refers to cognitively motivated processes affecting the “private” side of an individual's identity, normative influence is concerned with socially motivated processes affecting the “public” side of an individual's identity. As it stands, this dichotomy poses many problems since it is based on the assumption first that informational influence cannot be socially mediated, second that social norms have no effect on individual attitude change [80]. This assumption reflects an antagonism between individual and society. It suggests that the role of social structure and interaction is only to constrain individuals' behaviour, whereas change in individuals' mental attitudes results only from information-based, cognitive and asocial processes.

Attempts have been made to address the problems posed by the distinction between informational and normative influence. A variety of theories have been proposed that, while retaining the insights of the dual-process model, try to integrate influence phenomena within a unified conceptual scheme more consistent with empirical data and endowed with more explanatory power [1,82]. From a social psychological perspective, the most influential of these attempts has been the *self-categorisation* theory [78,79]. The hallmark of this theory is the idea that the individual, the social context and social norms are strictly intertwined, and that society affects and shapes the cognitive activities of the individual and vice versa. In organisation studies, this idea of co-evolution between the social world and the individuals' cognitive worlds has been further developed by the constructural theory [12,13]. According to this theory, self construction and social reconstruction are parallel processes: social changes result from changes in the distribution of knowledge as individuals interact, acquire and communicate information.

Both the self-categorisation and the constructural theory pave the way for a unified perspective in our understanding of social influence in several respects.

The distinction between normative and informational influence can be replaced by the idea that the basic influence process is one where the normative positions of people tend to be cognitively accepted as valid, and the validity of information is cognitively established by the social context and social norms. Individuals shift towards persuasive material and change their cognitive worlds accordingly, but what is persuasive is not a matter of information that can be abstracted from the social context in which interactions occur. In contrast, it can be argued that the degree to which information is persuasive depends on the degree to which it has been validated through its connections with an underpinning social structure in which individuals engage in social interactions, communicate messages, acquire and disseminate knowledge.

The model we proposed in this paper is an attempt to refine thinking on these matters. By developing a view of social influence that is predicated on the interrelations among individuals' cognition, principles of social behaviour and structural patterns of social relations, an effort has been made to integrate cognitively oriented informational influence and socially based normative influence into a unified theoretical perspective. In our model, the social world and the individual cognitive processes are brought together in an integrated explanation in which information has validity and impacts on agents' mental states to the degree that it reflects the social structure and the social interactions through which it is generated and exchanged. For example, the role of the additional piece of information included in strategy B* depends on a combination of social and cognitive factors. It impacts on the active investor's mental state as long as it is forwarded to the active investor by the seller, and provided that a social relationship with the small investors has already taken place. The strength of its impact also depends on whether or not the active investor believes that the seller is trustworthy which, in turn, determines whether or not the active investor believes that the message received from the seller is correct.

The idea of interdependence between cognitively and socially motivated processes of influence is also reflected in our conceptualisation of the agent who exerts, and is subjected to, social influence. The agent is here regarded as a cognitive associative social entity, engaged in an iterated series of social actions and interactions aimed at *completing* its mental state [61]. In our model, besides providing new reasons for keeping individually motivated attitudes, the social world in which the agent is located also offers mental attitudes that can be adopted to complement or merely to change its individual mental state. The complex interplay between the cognitive and the social world turns out to be a process in which social influence plays a key function in complementing and augmenting the agent's bare individual mental attitudes with socially motivated ones. In this view, social influence is a socio-cognitive process that impacts on the agents' cognition through the social context in which the agents are located and the social interactions in which they are involved. Embedded in a structural pattern of social relations, agents interact and by interacting they persuade each other of the validity of each other's views. Informational validity, in turn, has a social-normative aspect: it is affected by social interactions between agents within varying social contexts, embodied in the agents' social identity and group memberships. Influence generates attitude change in the agents' mental states

because in their interactions agents acquire information that they accept as valid, and in turn they accept information as valid precisely because it is generated and gathered within a specific social context and as a result of a range of social interactions with other agents.

In this paper, we explored the properties of our conception of social influence by focusing on the negotiation process that underlies the sale of a privately held company through the public offering of shares and through the involvement of an active investor. We then ran a series of simulations to examine the effects of social influence in terms of the performance of the sale. In doing this, we systematically analysed social influence in terms of its underpinning cognitive, social and structural components, and we studied the individual and combined effects of these components on the performance of the sale. Agents' cognition has been formalised drawing on traditional BDI architectures [17,68,84], and the role that cognition has in determining how effectively and efficiently social influence is exercised has been examined by manipulating cognitive accuracy. We also concentrated on social interactions among agents and we formalised them in terms of the negotiation protocol that the agents use to initiate and maintain social processes with others. Finally, we focused on a set of stable patterns of social relationships that provide the agents with the structural context in which they can exercise social influence.

The proposed model has considerable explanatory power. Many of its implications are consistent with known findings both in research on staged equity financing and, more generally, in research on social influence undertaken from a variety of theoretical perspectives. New theoretical insights can also be derived from the model. Overall, our results indicate that social influence matters in the generation of interpersonal agreements between the seller of a firm and potential investors. Different ways of exercising influence have differing impacts on the effectiveness and efficiency of the process of agreement generation. More specifically, we found that the effects of social influence on performance depend on how persistent the agents are in sticking to their initial views. Additionally, our results show that social influence is affected by the agents' cognitive accuracy. For example, we found that, on average, the impact that social influence has on effectiveness becomes weaker as the seller's beliefs become less accurate. Our results also shed light on what are normally hidden cognitive processes. In fact, we found that cognitive accuracy and information reinforce each other in their impact on performance. Whereas high accuracy emphasises the positive effects on effectiveness of an increase in the amount of information, using more information amplifies the possible gains in effectiveness that result from an increase in accuracy. In turn, this interaction between cognitive accuracy and information affects the relative performance of the proposed range of structural patterns of interaction that differ in terms of the amount and type of the information conveyed. More interestingly, our findings indicated that when it comes to deciding how to exercise influence, agents are confronted with trade-offs. For example, on average, an increase in effectiveness as a result of an improvement in cognitive accuracy can be obtained only at the cost of lower efficiency. In addition, when the agents choose to be more conciliatory with each other, they can exercise influence more efficiently, but only to the detriment of effectiveness.

Our study has a number of methodological implications regarding some of the advantages of using computational techniques and simulations for studying social and organisational processes [11,14,44]. First, it allowed us to simultaneously explore a range of structural, social and cognitive factors and their relative impacts on a complex process such as social influence. In our effort to build a plausible model of negotiation in which social influence is exercised in connection with the process of going public, we tried to rely on established theories and available evidence in behavioural finance. However, even though the individual components of the model are fairly well understood, we suspected that their combined effect was not. Thus, using computational techniques, we simulated these components together over time, and this allowed us to yield valuable insights into different combinations of modelling factors. Second, using simulations, we could study social influence, in terms of its antecedents, components and consequences, more easily, more systematically and with less cost than would be the case if we had conducted surveys, case or field studies. Our research domain is a good example of the relative advantages of simulations in situations in which 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. Third, by operationalising an important but difficult-to-measure determinant of performance, such as cognitive accuracy, we could explore unanticipated properties of cognitive processes at both the individual and collective level [15]. One of these properties, for example, is the relation between self-confidence, information and cognitive accuracy, and the role that this relation has in affecting the performance of social influence.

Finally, it is important to point out that the model we developed is undoubtedly incomplete and a number of issues raised in this paper require further investigation. Our findings are best regarded as hypotheses, and considerably more analysis and research need to be done before they can be considered as anything more than mere conjectures. We developed our computational model building on extant empirical evidence; even so, however, the model did not prove any of the hypotheses it helped develop. It merely allowed us to derive empirically verifiable predictions by answering logical “what if” questions. To prove these predictions, future empirical studies are needed for validating our model and testing our findings against real data. Furthermore, even though our main interests here lie in the generation of interpersonal agreements in complexly differentiated social systems, nonetheless our modelling strategy could also be applied to the study of the generation of other higher-order forms of cognition, such as organisational culture [34], organisational competence [65], and organisational commitment [64]. In all these areas of organisational studies, social influence can be regarded as the key socio-cognitive process that allows researchers to predicate theories of “group mind”-like constructs upon theories of individual cognition, social structure and social interaction. Clearly, a more analytical understanding of the underpinning factors that determine how influence unfolds over time would sharpen the debate centred around the attainability and evolution of collective forms of mental models and cognitive architectures that cut across multiple agents within and among organisations.

References

- [1] D. Abrams, M.A. Hogg, Social identification, self-categorization and social influence, in: W. Stroebe, M. Hewstone (Eds.), *European Review of Social Psychology*, vol. 1, John Wiley, Chichester, 1990, pp. 195–228.
- [2] K.W. Arrow, *Social Choice and Individual Values*, Wiley, NY, 1951.
- [3] R. Axelrod, *The Evolution of Cooperation*, Basic Books, NY, 1984.
- [4] M. Barclay, C. Holderness, Private benefits of control of public corporations, *Journal of Financial Economics* 25 (1989) 371–395.
- [5] C. Barry, C. Muscarella, J. Peavy, M. Vetsuypens, The role of venture capital in the creation of public companies, *Journal of Financial Economics* 27 (1990) 447–471.
- [6] L. Bebchuk, L. Zingales, Private versus social optimality, Discussion Paper Series, Harvard Law School, Cambridge, MA, 1995.
- [7] D.J. Brass, Being in the right place: a structural analysis of individual influence in an organization, *Administrative Science Quarterly* 29 (1984) 518–539.
- [8] M. Brennan, J. Franks, Underpricing, ownership and control in the initial equity public offering of equity securities in the UK, Working Paper 12-95, University of California, Los Angeles, CA, 1995.
- [9] H. Brown, *Social Psychology*, The Free Press, NY, 1965.
- [10] E. Burnstein, A. Vinokur, Persuasive argumentation and social comparison as determinants of attitude polarization, *Journal of Experimental Social Psychology* 13 (1977) 315–332.
- [11] R.M. Burton, B. Obel, The validity of computational models in organisation science: from model realism to purpose of the model, *Computational and Mathematical Organisation Theory* 1 (1) (1995) 57–71.
- [12] K.M. Carley, The value of cognitive foundations for dynamic social theory, *Journal of Mathematical Sociology* 14 (2–3) (1989) 171–208.
- [13] K.M. Carley, A theory of group stability, *American Sociological Review* 56 (1991) 331–354.
- [14] K.M. Carley, M.J. Prietula (Eds.), *Computational Organization Theory*, Lawrence Erlbaum Associates, Hillsdale, NJ, 1994.
- [15] K.M. Carley, W.A. Wallace, Editorial, *Computational and Mathematical Organisation Theory* 1 (1) (1995) 5–7.
- [16] D. Cartwright, Influence, leadership, control, in: J.G. March (Ed.), *Handbook of Organizations*, Rand-McNally, Chicago, IL, 1965, pp. 1–47.
- [17] P.R. Cohen, H.J. Levesque, Intention is choice with commitment, *Artificial Intelligence* 42 (3) (1990) 213–261.
- [18] P.R. Cohen, H.J. Levesque, Teamwork, *Noûs* 25 (4) (1991) 487–512.
- [19] J. Coleman, *Foundations of Social Theory*, Harvard University Press, Cambridge, MA, 1990.
- [20] B. Collins, B.H. Raven, Group structure: attraction, coalitions, communication, and power, in: G. Lindzey, E. Aronson (Eds.), *Handbook of Social Psychology*, vol. 4, Addison-Wesley, Reading, MA, 1969, pp. 102–204.
- [21] J.A. Davis, Structural balance, mechanical solidarity, and interpersonal relations, *American Journal of Sociology* 68 (1963) 444–462.
- [22] J.H. Davis, Group decision and social integration: a theory of social decision schemes, *Psychological Review* 80 (1973) 118–121.
- [23] D.C. Dennet, *The Intentional Stance*, The MIT Press, Cambridge, MA, 1987.
- [24] M. Deutsch, H.B. Gerard, A study of normative and informational social influences upon individual judgement, *Journal of Abnormal and Social Psychology* 51 (1955) 629–636.
- [25] E. Durkheim, *The Division of Labour in Society*, The Free Press, NY, 1933.
- [26] L. Festinger, Informal social communications, *Psychological Review* 57 (1950) 271–282.
- [27] L. Festinger, A theory of social comparison processes, *Human Relations* 7 (1954) 117–140.
- [28] S. Franklin, A. Graesser, Is it an agent, or just a program? A taxonomy for autonomous agents, in: J.P. Mueller, M. Wooldridge, N.R. Jennings (Eds.), *Intelligent Agents*, vol. III, Springer Verlag, Berlin, 1997, pp. 21–35.
- [29] J.R.P. French Jr., A formal theory of social power, *The Psychological Review* 63 (1956) 181–194.

- [30] N.E. Friedkin, *A Structural Theory of Social Influence*, Cambridge University Press, Cambridge, MA, 1998.
- [31] A. Giddens, *The Constitution of Society: Outline of the Theory of Structuration*, University of California Press, Berkeley, CA, 1984.
- [32] S. Grossman, O. Hart, Takeover bids, the free-rider problem, and the theory of the corporation, *Bell Journal of Economics* 11 (1980) 42–64.
- [33] K. Hanley, W. Wilhelm, Evidence on the strategic allocation of initial public offerings, *Journal of Financial Economics* 37 (1995) 239–257.
- [34] J.R. Harrison, G.R. Carroll, Keeping the faith: a model of cultural transmission in formal organizations, *Administrative Science Quarterly* 36 (1991) 552–582.
- [35] C. Holderness, D. Sheehan, The role of majority shareholders in public held corporations: an exploratory analysis, *Journal of Financial Economics* 20 (1988) 317–347.
- [36] N.R. Jennings, On agent-based software engineering, *Artificial Intelligence* 117 (2000) 277–296.
- [37] E.C. Johnsen, Structure and process: agreement models for friendship formation, *Social Networks* 8 (1986) 257–306.
- [38] M. Kahan, Sales of corporate control, *Journal of Law, Economics and Organisation* 9 (1993) 368–379.
- [39] A. Kirman, Whom or what does the representative agent represent? *Journal of Economic Perspectives* 6 (2) (1992) 117–136.
- [40] A. Klein, J. Rosenfeld, W. Beranck, The two stages of an equity carve-out and the price response of parent and subsidiary stock, *Managerial and Decision Economics* 12 (1991) 449–460.
- [41] S. Kraus, K. Sycara, A. Evenchil, Reaching agreements through argumentation: a logical model and implementation, *Artificial Intelligence* 104 (1998) 1–69.
- [42] B. Latane, The psychology of social impact, *American Psychologist* 36 (1981) 343–356.
- [43] J. Lee, M.J. Huber, E.H. Durfee, P.G. Kenny, 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, 1994, pp. 842–849.
- [44] A. Lomi, E.R. Larsen (Eds.), *Dynamics of Organizations. Computational Modeling and Organization Theories*, AAAI Press/The MIT Press, Menlo Park, CA, 2001.
- [45] P.V. Marsden, N.E. Friedkin, Network studies of social influence, *Sociological Methods & Research* 22 (1993) 127–151.
- [46] G.H. Mead, *Mind, Self and Society*, University of Chicago Press, Chicago, IL, 1962.
- [47] A.S. Mello, J.E. Parsons, Going public and the ownership structure of the firm, *Journal of Financial Economics* 49 (1998) 79–109.
- [48] R.K. Merton, The role-set: problems in sociological theory, *British Journal of Sociology* 7 (1957) 106–120.
- [49] R.K. Merton, E. Barber, Sociological ambivalence, in: R.K. Merton (Ed.), *Sociological Ambivalence and Other Essays*, The Free Press, NY, 1976, pp. 3–31.
- [50] W. Mikkelsen, M. Partch, K. Shah, Performance of Companies around Initial Public Offerings, Working Paper, University of Oregon, Eugene, 1995.
- [51] W. Mikkelsen, R. Ruback, An empirical analysis of the interfirm equity investment process, *Journal of Financial Economics* 14 (1985) 523–553.
- [52] S. Milgram, *Obedience to Authority*, Tavistock, London, 1974.
- [53] P. Milgrom, R. Weber, A theory of auctions and competitive bidding, *Econometrica* 50 (1982) 1089–1122.
- [54] S. Moscovici, *Social Influence and Social Change*, Academic Press, London, 1976.
- [55] S. Moscovici, Social influence and conformity, in: G. Lindzey, E. Aronson (Eds.), *The Handbook of Social Psychology*, third ed., vol. 2, Random House, NY, 1985, pp. 347–412.
- [56] S. Moscovici, M. Zavalloni, The group as a polarizer of attitudes, *Journal of Personality and Social Psychology* 12 (1969) 125–135.
- [57] G. Mugny, *The Power of Minorities*, Academic Press, London, 1982.
- [58] D.G. Myers, H. Lamm, The group polarization phenomenon, *Psychological Bulletin* 83 (1976) 602–627.

- [59] G. Ouchi, The relationship between organizational structure and organizational control, *Administrative Science Quarterly* 22 (1977) 95–113.
- [60] M. Pagano, F. Panetta, L. Zingales, Why do companies go public? An empirical analysis, *Journal of Finance* 53 (1998) 27–64.
- [61] P. Panzarasa, N.R. Jennings, T.J. Norman, Social mental shaping: modelling the impact of sociality on the mental states of autonomous agents, *Computational Intelligence* 17 (4) (2001) 738–782.
- [62] P. Panzarasa, N.R. Jennings, T.J. Norman, Formalising collaborative decision-making and practical reasoning in multi-agent systems, *Journal of Logic and Computation* 11 (6) (2001) 1–63.
- [63] P. Panzarasa, N.R. Jennings, T.J. Norman, Going public and the sale of shares with heterogeneous investors: agent-based computational modelling and computer simulations, *International Journal of Group Decision and Negotiation* 10 (5) (2001) 423–470.
- [64] R.T. Pascale, The paradox of ‘corporate culture’, *California Management Review* 27 (2) (1985) 26–41.
- [65] C.K. Prahalad, G. Hamel, The core competence of the corporation, *Harvard Business Review* (May–June 1990) 79–91.
- [66] R. Radner, Dynamic games in organization theory, in: M. Majumdar (Ed.), *Organizations with Incomplete Information: Essays in Economic Analysis*, Cambridge University Press, NY, 1998, pp. 110–158.
- [67] H. Raiffa, *The Art and Science of Negotiation*, Harvard University Press, Cambridge, MA, 1982.
- [68] S. Rao, M.P. Georgeff, Modeling agents within a BDI architecture, in: R. Fikes, E. Sandewall (Eds.), *Proceedings of the Second International Conference on Principles of Knowledge Representation and Reasoning (KR’91)*, Morgan Kaufmann, Cambridge, MA, 1991, pp. 473–484.
- [69] K. Rock, Why are new issues underpriced? *Journal of Financial Economics* 15 (1986) 187–212.
- [70] A. Rubinstein, *Modeling Bounded Rationality*, The MIT Press, Cambridge, MA, 1998.
- [71] K. Rydqvist, K. Högholm, Going public in the 1990s: evidence from Sweden, *European Financial Management* 1 (1994) 287–316.
- [72] M. Sherif, *The Psychology of Social Norms*, Harper and Brothers, NY, 1936.
- [73] K. Shipper, A. Smith, A comparison of equity carve-outs and seasoned equity offerings, *Journal of Financial Economics* 15 (1986) 153–186.
- [74] A. Shleifer, R. Vishny, Large shareholders and corporate control, *Journal of Political Economy* 94 (3) (1986) 461–488.
- [75] H.A. Simon, *Administrative Behavior. A Study of Decision-Making Processes in Administrative Organizations*, fourth ed., The Free Press, New York, NY, 1997.
- [76] N. Stoughton, J. Zechner, IPO-mechanisms, monitoring and ownership structure, *Journal of Financial Economics* 49 (1998) 45–77.
- [77] S. Stryker, A. Staham, Symbolic interaction and role theory, in: G. Lindzey, E. Aronson (Eds.), *Handbook of Social Psychology*, third ed., vol. 1, Random House, NY, 1985, pp. 311–378.
- [78] J.G. Turner, Towards a cognitive redefinition of the social group, in: H. Tajfel (Ed.), *Social Identity and Intergroup Relations*, Cambridge University Press, Cambridge, MA, 1982, pp. 15–40.
- [79] J.G. Turner, Social categorization and the self-concept: a social cognitive theory of group behavior, in: E.J. Lawler (Ed.), *Advances in Group Processes*, vol. 2, JAI Press, Greenwich, CT, 1985, pp. 77–122.
- [80] J.G. Turner, *Social Influence*, Open University Press, Milton Keynes, 1991.
- [81] J.G. Turner, M.A. Hogg, P.J. Oakers, S.D. Reicher, M.S. Wetherell, *Rediscovering the Social Group: A Self-Categorization Theory*, Basil Blackwell, Oxford, 1987.
- [82] J.G. Turner, P.J. Oakes, Self-categorization theory and social influence, in: P.B. Paulus (Ed.), *The Psychology of Group Influence*, second ed., Lawrence Erlbaum Associates, Hillsdale, NJ, 1989, pp. 233–275.
- [83] E. Werner, Cooperating agents: a unified theory of communication and social structure, in: L. Gasser, M.N. Huhns (Eds.), *Distributed Artificial Intelligence*, Pitman/Morgan Kaufmann, London, 1989, pp. 3–36.
- [84] M. Wooldridge, *Reasoning About Rational Agents*, The MIT Press, Cambridge, MA, 2000.

- [85] D.H. Wrong, *The Problem of Order: What Unites and Divides Society*, Harvard University Press, Cambridge, MA, 1961.
- [86] L. Zingales, Insider ownership and the decision to go public, *Review of Economic Studies* 62 (1995) 425–448.