Ten-Year History of Social Network Logics in China

*LIU Fenrong**
*LI Dazhu***

Abstract

The paper presents a ten-year history of social network logics in China. It tells the story of how this new research area was started, how its research agenda was extended, and, in particular, how a focus on graph games developed. Important ideas and research results are summarized, with an emphasis on the connections between them. An important aspect of this history is the successful collaboration between Chinese and international researchers.

Keywords: social network logics, peer pressure, graph games, dynamics, consensus

Desetletna zgodovina logike družbenih omrežij na Kitajskem

Izvleček

Članek predstavlja desetletno zgodovino logike družbenih omrežij na Kitajskem. Pripoveduje zgodbo o tem, kako je to novo področje nastalo, kako se je potem širil njegov osnovni raziskovalni program in, še posebej, kako je znotraj področja vzniknil fokus na igre z grafi. Članek povzema pomembne ideje in rezultate raziskav, s posebnim poudarkom na povezavah, ki obstajajo med njimi. Pomemben vidik te zgodovine pa predstavlja uspešno sodelovanje med kitajskimi raziskovalci in njihovimi kolegi v mednarodnem prostoru.

Ključne besede: logika družbenih omrežij, vrstniški pritisk, igre grafov, dinamika, konsenz

* LIU Fenrong, Department of Philosophy, Tsinghua University, Beijing, China.
Email address: fenrong@tsinghua.edu.cn

** LI Dazhu, Institute of Philosophy, Chinese Academy of Sciences and Department of Philosophy, University of Chinese Academy of Sciences, Beijing, China.
Email address: lidazhu@ucas.ac.cn

How Social Network Logics Got Started

In 2010 Jeremy Seligman, a logician at the University of Auckland, was spending two months of his sabbatical at Tsinghua University in Beijing, teaching a seminar on "situation theory and channel theory", and he had several research discussions with his host Fenrong Liu. Both of them were trained in modern logic, but they shared a strong interest in Chinese philosophy. The question they asked back then was: Can we develop a logic with features that are important to reasoning in the Chinese tradition? In other words, can we incorporate some interesting notions of Chinese philosophy in modern logic? For instance, Confucianism has profound ideas about social roles, relationships, and hierarchy that have influenced the nation and its people for many years. Certainly, they play a role in people's reasoning and social interactions. This led to their first joint paper "Logic in the Community", which was also written with Patrick Girard, a colleague of Seligman's at Auckland (Seligman, Liu and Girard 2011).

The paper "Logic in the Community" was written much like a research proposal, laying out the problems that the authors wanted to study in the following years. It started with a section "Reasoning About Social Relations" with a description of this project (ibid., 178):

> Communities consist of individuals bounds together by social relationships and roles. Within communities, individuals reason about each other's beliefs, knowledge and preferences. Knowledge, belief, preferences and even the social relationships are constantly changing, and yet our ability to keep track of these changes is an important part of what it means to belong to a community. In the past 50 years, our patterns of reasoning about knowledge, beliefs and preferences have been extensively studied by logicians, but the way in which we are influenced by social relationships has received little attention.

From the above, a rough picture of social network logic emerges. It is built up on the tradition of modal logic, but with a new focus on people and their social relationships. In addition to dynamic logics for reasoning about knowledge, beliefs and preferences, social network logic introduces a second dimension, social relation, to the framework. The paper highlighted many intriguing issues such as "Facebook friends", "deference to expert opinion", "peer pressure", "community norms", etc. It also gave the building blocks of a two-dimensional approach: one dimension standing for each person's epistemic space—the range of situations (or "worlds") that person considers possible; the second for each person's community—those other people with whom they may have closer or more remote social relationships.

An outline of social network logic, or, more specifically social *epistemic* logic, will be given in the next section. In third and fourth section and, we will review subsequent research on this logic. Fifth section will consider a number of closely related research directions, and the penultimate section introduces the main results of a newer direction inspired by the study of games. In the last section, we will briefly discuss research on social network logic internationally and then conclude this work.

Introducing a Social Dimension to Epistemic Logic

Epistemic logic involves the addition of a modal operator K to standard propositional logic, with the formula $K\varphi$ interpreted either as "I know that φ" or the more objective "it is known that φ". Growing from the work of the pioneering logicians of the early 20th century, such as Rudolph Carnap and Arthur Prior, the seminal paper on epistemic logic was by G. H. Von Wright (1951). It was first given book-length treatment in Jakko Hintikka's *Knowledge and Belief: An Introduction to the Logic of the Two Notions* (1962). Much of the early work focused on logical validities involving the K operator. Since knowledge implies truth, it was standard to regard $Kp \to p$ as valid, but the "introspective" principles that something known is known to be known ($Kp \to KKp$) or that something unknown is known to be unknown ($\neg Kp \to K\neg Kp$) were more contentious. The similar operator of belief, B, was naturally distinguished by different logical properties.

In the 1980s and 1990s interest in epistemic logic accelerated and expanded largely because of its relevance to computer science and the representation of "knowledge" in software. An important development was the indexing of the knowledge operator with the name of the knower: $K_a\varphi$ is then interpreted as "a knows that φ". Because of application to computer systems, the knowers were referred to more generally as "agents", and that is the terminology we will use here.

A formal semantics for epistemic logic can be given in the style of Kripke. A set W of points (usually called "worlds") represents the different ways the relevant facts could be, not only objective facts about objects and their properties, but also epistemic facts about who knows what. This representation is achieved through a function V mapping propositional variables to subsets of W and a binary relation R between the points of W. $V(p)$ is interpreted as the set of worlds in which p is true and Ruv is interpreted as meaning that in world u, the knower has not yet ruled out v as an epistemic possibility. In other words, the agent does not know whether she is in world u or world v (or any of the other worlds in the R relation to u). Together these elements comprise an *epistemic model* $M = \langle W, R, V \rangle$. Typically, further constraints are imposed on the R relation: reflexivity, transitivity and sometimes symmetry. A recursive definition of a *satisfaction* relation between models, worlds and formulas is then defined as follows:

$M, w \vDash p$ iff $w \in V(p)$

$M, w \vDash \neg\varphi$ iff $M, w \nvDash \varphi$

$M, w \vDash (\varphi \wedge \psi)$ iff $M, w \vDash \varphi$ and $M, w \vDash \psi$

$M, w \vDash K\varphi$ iff $M, v \vDash \varphi$ for each v such that Rwv

Here $M, w \vDash \varphi$ is read "w satisfies φ in M" and this is interpreted to mean that φ expresses a proposition that would be true in the world represented by w in model M. So, for example, if $M, w \vDash (p \wedge \neg Kp)$ then w represents a world in which p is true but not known to be true. A formula is logically valid only in cases when it is satisfied by every world in every model. And that's where the restrictions on R come in. In order for $Kp \to p$ to be logically valid, the relation R must be reflexive. This is standard material in the study of modal logic, much of which will be assumed in what follows. (Readers unfamiliar with modal logic should consult a suitable textbook, such as Blackburn, de Rijke and Venema 2002).

The extension to "multi-agent" epistemic logic, with an operator K_a for each agent a of a given set A is straightforward. The models $M = \langle W, R, V \rangle$ now consist of a family of relations R_a, one for each agent a, and the satisfaction definition is almost identical:

$M, w \vDash K_a\varphi$ iff $M, v \vDash \varphi$ for each v such that $R_a wv$

Nonetheless, multi-agent epistemic logic is considerably more interesting than its single-agent ancestor. That's because it is possible to express "higher-order" epistemic facts: what one agent knows about what another agent knows, or doesn't know. For example, $K_a \neg K_b p$ represents a's knowing that b does not know that p. Moreover, extensions of the language allow reasoning about what is *commonly known* to a group of agents (Fagin et al. 2004) and the addition of a range of model-changing "dynamic" operators extends all this to the logic of how knowledge changes under various acts of communication (Baltag, Moss and Solecki 1998).

It is with this background of research in epistemic logic that the development of social network logic, and specifically the social epistemic logic of Seligman, Liu and Girard (2011) must be seen. The innovation of that paper was to add a new dimension to the models: the social dimension. Instead of evaluating formulas based on worlds, the new idea is to evaluate them on pairs (w, a) consisting of a world w and an agent a. In this new system, a formula expresses an "agent-indexical" proposition: the satisfaction of φ by the pair (w, a) is interpreted to mean that φ is true in w from the point of view of agent a. For example, take p to express the agent-indexical proposition "I'm in danger". Then $M, w, a \vDash p$ is interpreted to mean that in world w, agent a is in danger. The definition of satisfaction given earlier is thus subtly modified to this two-dimensional setting:

$M, w, a \vDash p$ iff $(w, a) \in V(p)$

$M, w, a \vDash \neg\varphi$ iff $M, w, a \nvDash \varphi$

$M, w, a \vDash (\varphi \wedge \psi)$ iff $M, w, a \vDash \varphi$ and $M, w, a \vDash \psi$

$M, w, a \vDash K\varphi$ iff $M, v, a \vDash \varphi$ for each v such that $R_a wv$

Notice in particular the last clause, for *K*. There is now no need to index the *K* operator: the relation of knowledge to a knower is a consequence of using agent-indexical propositions.

This logical shift was accompanied by two main additions to the language. The first is an operator *F* which corresponds to a relation *S* between agents, that is, a *social* relation. As in Seligman, Liu and Girard (2011), we will interpret this as the "friendship" relation, although this stands as proxy for any number of social relationships, or indeed any relation between people. (Later applications involve taking *F* to be the "seeing" relation.) The interpretation of formulas using *F* is probably best understood by contrast: with *p* interpreted as above ("I'm in danger"), *KFp* means that I know that all my friends are in danger, whereas *FKp* means that all my friends know that they are in danger. The De Morgan dual of *F* (i.e., $\neg F \neg$) is written $\langle F \rangle$, so that $\langle F \rangle\, p$ is interpreted to mean that I have a friend who is in danger. The clause in the definition of satisfaction corresponding to *F* is the following:

$M, w, a \vDash F\varphi$ iff $M, w, b \vDash \varphi$ for each b such that $S_w ab$

As you can see, the "relation" *S* is in fact a family of relations, S_w, for each w in W. That's because social facts may vary between worlds, or, in other words, they may be known by some agents but not by others.

The second addition concerns the reference to agents. In standard epistemic logic, there is no need to distinguish between agents and their names. In fact, it cannot be done. The agents are only represented syntactically, as indices to the *K* operator. In the models of social epistemic logic, however, they are "in the model" and so a distinction can be made. The language is therefore enriched to contain a number of ways of managing references to agents:

(a) There are names *n*, *m*, etc. which refer to agents, but need not do so rigidly: they may refer to different agents in different worlds, so allowing for the ignorance about who is named what. In fact, no distinction is made between the name n and the agent-indexical proposition "I am n", so, for example, $\langle F\rangle n$ means that n is my friend.

(b) Names are also used to shift perspective, as indices to a new operator: $@_n\varphi$ means that the agent-indexical proposition expressed by φ holds not of me but of the agent named n. So, for example, $@_nKp$ means that n knows that p (the same as K_np in standard epistemic logic).

(c) Variables are also used to refer to agents, in a way that can be bound to the indexical subject of the proposition (just like the first-person pronoun in natural language). This is done with the "down-arrow" operator. $\downarrow x.\varphi$ is interpreted just like φ but with any free variable x it contains acting as a name for the indexically determined agent.

All three of these referential devices are adaptations from hybrid logic (Blackburn and Seligman 1996), an extension of modal logic using ideas originally developed by Arthur Prior and later reinvented in many places. Together they determine a richly expressive language in which many propositions about the social and epistemic properties of agents can be stated.

There are a number of equivalent ways of implementing the semantics. Here we choose one that is hopefully easy to understand. Since the agent names are a special kind of propositional variable, we allow the valuation V to determine for each name n a set of world-agent pairs, with the interpretation that (w, a) is in $V(n)$ just in case n refers to a in world w. (And so there is a restriction on V that there is a unique a such that (w, a) is in $V(n)$.) With this in place, satisfaction for names and the @ operator is defined as follows:

$M, w, a \vDash n$ iff $(w, a) \in V(n)$

$M, w, a \vDash @_n\varphi$ iff $M, w, b \vDash \varphi$ for the unique b such that $(w, b) \in V(n)$

The handling of variables also presents alternatives, just as for predicate logic. Here, mainly for completeness of this introduction, we will follow the standard approach of using an *assignment function* g as a parameter to the definition of satisfaction. The assignment function assigns agents to variables, and can be altered using the $\downarrow$ binder:

$M, g, w, a \vDash x$ iff $g(x)=a$

$M, g, w, a \vDash \downarrow x.\varphi$ iff $M, g', w, a \vDash \varphi$ where $g'(y) = a$ if $y=x$; $g(y)$ otherwise.

(The g must of course be added as a parameter to all the above clauses.)

This completes the outline of social epistemic logic, as conceived in Seligman, Liu and Girard (2011). We have not covered details of the intended areas of application, but many of them will be covered in what follows. In the next section we will focus on developments of this logic by the authors of that paper and their students, as a result of research collaborations mostly conducted in China, but also as a result of Chinese students studying overseas.

The Development of Social Epistemic Logic

Shortly after "Logic in the Community", Liu, Seligman and Girard turned to dynamic extensions of social epistemic logic. These has been sketched briefly in the 2011 paper, but received more attention at the 2013 TARK conference: "Facebook and the Epistemic Logic of Friendship". In dynamic epistemic logic (DEL), the then standard approach to the epistemic logic of communication, communicative actions are modelled as operations that change the structure of the epistemic models, typically by adding and removing worlds and links in the R relation. This captures both the effect of the communication on the epistemic states of the agents, updating those who receive it, as well as the additional uncertainty created for those agents who did not, or who have only partial knowledge of who did. The central concept of the TARK paper is the *social announcement*. This is an action that accommodates the agent-perspectival aspects of communication, from the points of view of both the sender and receiver. The 'friendship' relation is taken as the channel. For example, I might broadcast to my friends that I am in danger. This is indexical information about me (from my perspective), sent to my friends whose knowledge is thereby updated with non-indexical information about me ("he is in danger"). Various kinds of social announcement were defined in the paper, and these were modelled using a powerful extension of DEL developed by the authors elsewhere (Girard, Liu and Seligman "General Dynamic Dynamic Logic" (2012)). Operators for changing social relationships, such as dropping and adding friends, were also considered, as were indexicalized variants of the concept of common knowledge, and the dynamics of questions and answers.

Meanwhile, some of the basic work on social epistemic logic was yet to be done. In particular, there was no complete axiomatization. In early conversations in Beijing, Katsuhiko Sano, a researcher from Japan's Advanced Institute of Science and Technology, had indicated a strong connection with his own work on two-dimensional modal logic (Sano 2010). This led to his producing a proof system for social epistemic logic (without $\downarrow$) using an extension of Gentzen's sequent calculus called "hypersequents" and its decidability, which was published much later (Sano 2017). Christoff, Hansen, and Proietti (2016) also produced a proof system for a very similar logic using tableaux. But work on a standard axiomatization of the logic was started by Zhen Liang, who moved from China's Southwest University to the University of Auckland for doctoral work under Seligman's supervision. He produced an axiomatization for the full language (including ↓) and a proof of its completeness, announced in Liang (2017) and with full details in his dissertation, Liang (2020).

We won't dwell on the technicalities of Liang's proof here, but it is worth giving a quick glimpse under the hood to reveal a further China connection. A standard approach to proving the completeness of axiomatizations of modal logic is to construct a single (huge) model for the language, within which every other model can be either found as a part, or extracted from a part. The huge model is called the "canonical model". For various interesting reasons explained in Liang's dissertation, construction of a canonical model for social epistemic logic was fraught with difficulty. Instead, he adapted a technique developed much earlier by a well-known Chinese logician, Ming Xu (1988). Xu's approach to proving the completeness of axiomatizations of certain temporal logics was to construct models in stages, step-by-step, and this technique also proved fruitful for Liang's axiomatization of social epistemic logic, although further complications arise in the case of ↓. A canonical model proof was eventually given by Saúl Fernández González (Balbiani and González 2020; González 2021).

Liang's dissertation also contains a new area of application of social epistemic logic, in which the social relation S is interpreted as the "seeing" relation. This allows the logical analysis of interesting problems and scenarios involving the interaction of knowledge and perception in a social setting. (One such example is the phenomenon of *pluralistic ignorance.*)

Meanwhile, further progress on the logic of social announcement was made by Zuojun Xiong, another former student of Southwest University. In a collaboration with Thomas Ågotnes, of the University of Bergen, Jeremy Seligman and Rui Zhu, another Chinese student then working on a PhD in Auckland, he studied the logic of an *arbitrary social announcement* operator $\langle a\rangle\varphi$, meaning that φ holds after agent a makes some announcement of something he believes to all of his "friends". The results were published as Xiong et al. (2017) and were later extended substantially in his doctoral dissertation (Xiong 2017), supervised by Ågotnes. Zhu also went on to develop this logic further in his dissertation (Zhu, forthcoming) supervised by Seligman.

From High- to Low-level Rationality

While the initial exploration of logics based on social relations concerned knowledge, the two-dimensional framework is only suitable for studying other cognitive attitudes, such as belief and preference. Inspired by Liu's early work on the logic of preference change (Liu 2008; 2011), Zhen Liang adapted social epistemic logic to reason about preferences to help in understanding a well-recognised sociological phenomenon: peer pressure. This was while he was still a master's student in Southwest University, Chongqing, and thus before beginning his doctoral studies. During Seligman's visit to that university, they discussed his work and this led to a collaboration in which they modelled peer pressure using a preference-change operator that is sensitive to social relationships. Depending on the preferences of their "friends", an agent would be under stronger or weaker forms of suggestion. A weak suggestion of α over β would make them lose any preference for β, but a strong suggestion would make them prefer α. Although these changes were modelled as high-level deliberations (using a similar mechanism to the ones used in DEL for epistemic change), analysis of the models showed that the resulting dynamics could be modelled in a much simpler way, using *network automata*. A network automaton is a social network (the graph *S*) with a finite state machine running at each node. In "A Logical Model of the Dynamics of Peer Pressure", Liang and Seligman give the following network automaton responsible for their model of peer pressure (Liang and Seligman 2011, 282–83):

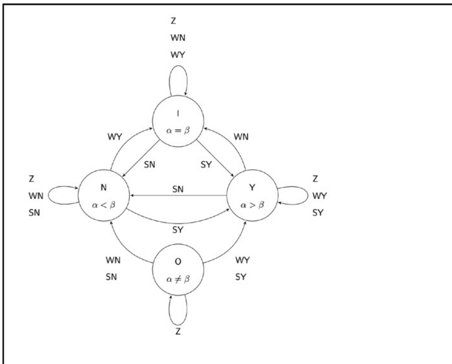

Y	$(\alpha > \beta)$	preference for α over β
N	$(\alpha < \beta)$	preference for β over α
I	$(\alpha = \beta)$	indifference concerning β and α
O	$(\alpha \neq \beta)$	conflicted preferences about β and α
SY	$F\mathsf{Y} \wedge \langle F \rangle \mathsf{Y}$	my friends strongly suggest Y
SN	$F\mathsf{N} \wedge \langle F \rangle \mathsf{N}$	my friends strongly suggest N
WY	$F(\mathsf{I} \vee \mathsf{Y}) \wedge \langle F \rangle \mathsf{Y} \wedge \langle F \rangle \mathsf{I}$	my friends only weakly suggest Y
WN	$F(\mathsf{I} \vee \mathsf{N}) \wedge \langle F \rangle \mathsf{N} \wedge \langle F \rangle \mathsf{I}$	my friends only weakly suggest N
Z	$\neg(\mathsf{SY} \vee \mathsf{SN} \vee \mathsf{WY} \vee \mathsf{WN})$	my friends have no suggestion

One nice consequence of a network automaton model of a social phenomenon is that it is often possible to analyse its asymptotic behaviour: whether preferences will eventually stabilize, fragment or enter some oscillating pattern. The paper provided examples of such an analysis.

Soon after, network automata were used directly to model the dynamics of belief change under a similarly structured but more abstract model of social influence. In "Logical Dynamics of Belief Change in the Community" (2014) Liu, Girard and Seligman used the following automaton to characterize the dynamics of strong and weak influence on belief. Here *Wp* means that the agent is being weakly influenced to believe p, and Sp means that they are being strongly influenced. The arrows show under what conditions they will transit from believing that p (Bp) to being undecided about p (Up) to believing that p is false ($B\neg p$). Various kinds of influence are discussed, as is the impact of changes to the underlying social relation.

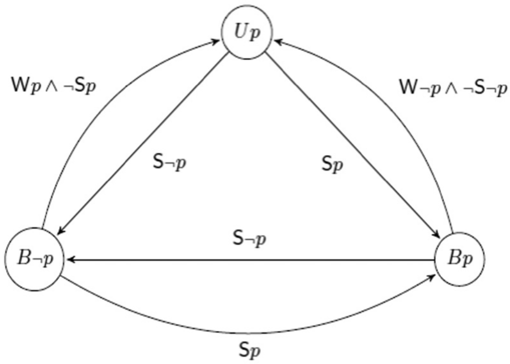

Another issue studied in the paper is the *object* of influence. At one extreme there is influence exerted on a specific opinion, such as the truth of a given proposition p. At the other is influence exerted on credence comparisons in general. It may be that a very small influence on one's relative credence judgements is sufficient to sway one toward or away from a particular belief; or, it may be that a very large influence is insufficient.

The use of network automata to model rational activity such as a change in belief is a little controversial among logicians. Logic is traditionally regarded as a purely normative discipline. We model the path of careful deliberation, aimed at truth. Models of social influence by network automata apparently lack this normative function. Or so the criticism goes. And yet we live in communities and are typically greatly influenced in our opinions by others. Not every decision we make is done so entirely on the weight of evidence available to us. And some mechanisms for propagating that influence are better than others—better in a normative sense. This contrast between the norms of individual deliberation and social mechanism are an example of what Brian Skyrms (2014) calls 'high' and 'low' rationality. At the level of personal psychology we have the capacity to respond to both, and it is the flexibility of humans to know when to deliberate and when not to that is one of our greatest strengths. Daniel Kahneman famously calls this the distinction between "thinking slow" and "thinking fast". For a more well-developed and extensive discussion of this issue, see van Benthem, Liu and Smets (2021).

An Expansion of the Research Agenda

Further work on social network logic in China was greatly assisted during the period 2017–2020 by a project funded by the Chinese Research Foundation for Philosophy and Social Sciences, whose principal investigators were Fenrong Liu, Johan van Benthem, Jeremy Seligman, Beishui Liao and Xinwen Liu. The project gathered about 25 Chinese and international researchers or students as participants tackling various issues in the area. In addition to following up on issues raised by previous research, new perspectives were developed, most notably that of game-playing by van Benthem. This will be reviewed in the next section. In what follows, we review a few of the other new topics.

Starting with the connection between evidence and belief in social settings, Fenrong Liu and Emiliano Lorini in their paper "Reasoning about Belief, Evidence and Trust in a Multi-agent Setting" (2017) studied how an agent accumulates evidence in support of a given fact φ from other agents, and how the body of evidence in support of φ can become a reason to *believe* φ. The paper provided a logic of the interplay between evidence and trust, and between evidence and belief. The new logic supports reasoning about an agent's belief formation and belief change due to new evidence. From this perspective, an agent is, by definition, social: she is connected to other agents and communicates by receiving information from them and passing information to them. *Trust* is a necessary condition for an agent to accept the information provided by another agent. A central assumption of the logic is that, to form a belief that a certain fact φ is true, an agent is sensitive to the following two aspects a) the *amount* of evidence in support of φ, and b) the *ratio* of evidence in support of φ to the total amount of evidence in support of either φ or its negation.

In standard multi-agent epistemic logic, agent names are implicitly assumed to be common knowledge. That's because the in the formula $K_a\varphi$, meaning that agent *a* knows that φ, the *a* is a *rigid designator*; it has the same denotation in every epistemic alternative. This is unreasonable in certain social settings. Yanjing Wang and Jeremy Seligman started their paper "When Names Are Not Commonly Known: Epistemic Logic with Assignments" with the following intriguing scenario (Wang and Seligman 2018, 611):

> One dark and stormy night, Adam was attacked and killed. His assailant, Bob, ran away, but was seen by a passer-by, Charles, who witnessed the crime from start to finish. This led quickly to Bob's arrest. Local news picked up the story, and that is how Dave heard it the next day, over breakfast. Now, in one sense we can say that both Charles and Dave

> know that Bob killed Adam. But there is a difference in what they know about just this fact. Although Charles witnessed the crime, and was able to identify the murderer and victim to the police, he might have no idea about their names. If asked "Did Bob kill Adam?" he may not know. Yet this is a question that Dave could easily answer, despite not knowing who Adam and Bob are, he is very unlikely to be able to identify them in a line-up.

The distinction between these *de re* and *de dicto* readings of 'Charles knows that Bob killed Adam' is hard to make in standard epistemic logic. The paper proposed an extension of epistemic logic using a combination of non-rigid names, rigid variables and assignment operators of the form $[x{:=}a]\varphi$, meaning that φ holds after x is assigned the agent named by 'a'. For example, $[x{:=}b]K_c\text{kill}(x,a)$ attributes the knowledge that Bob killed Adam without the implication that Charles knows who Bob is. The main technical result is a complete axiomization of this logic over S5 models.

Chenwei Shi in his recent paper "Collective Opinion as Tendency towards Consensus" (2021) studied the formation of collective opinions on social networks. The paper highlighted social influence with a nice quote from the book *Propaganda* written by Edward L. Bernays in 1928, "We are governed, our minds are molded, our tastes formed, our ideas suggested, largely by men we have never heard of". The paper made a distinction between a global perspective on the diffusion of opinions as a group process and a local agent-driven one. The main ideas can be summarized below, again cited from the paper (Shi 2021, 594):

> First, collective opinion is a *tendency toward convergence*. The paper models this view of opinion diffusion as a Markov process and understand a group's collective opinion as a high chance of reaching consensus.
>
> Secondly, the *influence structure of a group,* more precisely, how each group member is influenced by others, is the only crucial determinant of long-term opinion behavior, whether toward convergence or otherwise.

The main technical result is the discovery of structural conditions under which group opinion converges.

In "Reasoning and Making Predictions about Agent's Behaviors in Social Networks", Liu and Seligman (2018, in Chinese) distinguished two modes of social influence: one-direction influence, and mutual influence. In the former setting, the paper shows that by derivations in a logical calculus we can foresee the diffusion of certain behaviours, hence a prediction can be made. In the latter setting, two notions of stability are defined, which are of use when we want to make predictions on the spread of certain behaviours: an agent's behaviour is *stable* if she does not change her behaviour given any new influence; a network *stabilizes* if every agent is eventually stable.

In social psychology, there are many ways of analysing of social networks. One of those, Balance Theory, describes a signed network that has two relationships: positive ("friends") or negative ("enemies"). To connect such network analysis with research in logic, Zuojun Xiong and Thomas Ågotnes in their paper "On the Logic of Balance in Social Networks", developed a modal logic for reasoning about the structural properties of such social networks. The class of social networks is balanced to a certain degree *n* if there are no cycles of length up to *n* with an odd number of negative relationships. They completely axiomatized the class of all fully balanced complete signed social networks, i.e., networks where everyone is connected with everyone else (Xiong and Ågotnes 2020).

In the same direction, Yi Wang with his collaborators Wiebe van der Hoek and Louwe B. Kuijer studied social network logic and its connection with the balance theory. In their paper "Who Should Be My Friends? Social Balance from the Perspective of Game Theory" (2019) they defined *balance games*, which describe the formation of friendships and enmity in social networks. The interesting result they show is that if the agents give high priority to future profits over short term gains, all Pareto optimal strategies will eventually result in a balanced network, and if they prioritize short term gains over the long term, every Nash equilibrium eventually results in a stable network that might not be balanced (van der Hoek, Kuijer and Wang 2019). In a follow-up paper "Logics of Allies and Enemies: A Formal Approach to the Dynamics of Social Balance Theory", they combine social balance theory with temporal logic to obtain a Logic of Allies and Enemies (LAE), which can describe the dynamical changes of a social network due to social pressure, and they show that both model checking and validity checking of LAE are PSPACE-complete (van der Hoek, Kuijer and Wang 2020).

The Graph Game Logic Approach

In 2017, Johan van Benthem, Jeremy Seligman, Dag Westerståhl (Stockholm) and Martin Stokhof (Amsterdam) were appointed as Jin Yuelin Professors at Tsinghua. They share the same position, and each of them visits Tsinghua for 2–3 months every year, teaching courses and collaborating with colleagues. The purpose of such a position is to strengthen the logic research at Tsinghua and to carry on the tradition that was started by Jin Yuelin, a pioneering philosopher and logician at the university. In the autumn, Johan van Benthem and Fenrong Liu hosted a seminar to explore social interactions using *graph games*. One of the pioneering ideas is the *sabotage game* studied in van Benthem (2014, 477–85):

> *Definition.* A *sabotage game* is played on a graph, representing the environment, with a starting-node and a goal-node or a goal-region: in each round, a player Remover first cuts a link anywhere in the graph, and then the other player Traveller moves along an edge that is still available where she stands. Traveller wins if she arrives at a node in the goal-region: if this does not happen, and no more moves are possible, Remover wins.

In the research discussed so far, social networks have been centre stage, explicitly modelled as a set of agents, structured by one or more social relations. The graph structure of these networks makes them amenable to study using modal logic. So it is important to emphasize that the graphs of "graph games" are something different. Any graph can be studied using modal logic, and operators can be defined to correspond to actions that a player can take to change them. Indeed, in 2005 van Benthem had already proposed sabotage modal logic (SML) in his paper "An Essay on Sabotage and Obstruction". SML extends the standard modal language with an edge-deletion modality $\blacklozenge$: the standard modality $\Diamond\varphi$ means "Traveller is able to move to a node that is φ", while $\blacklozenge\varphi$ reads "there is a link such that after Remover cuts it, φ is the case". Using this language, the paper analysed sabotage games and studied reasoning about the graph change for two players.

At the Tsinghua seminar, van Benthem and Liu introduced various new graph games and made the first attempt to connect them with logic, in particular, modal logic and dynamic logic. This became their joint paper "Graph Games and Logic Design" (van Benthem and Liu 2020). The paper promotes a methodology of using logic both as a formal tool for analysing games, and as inspiration for the design of new games. It discusses a range of graph game types: travel games, sabotage games, meet/avoid games, and occupation games. It also proposed some parameters for the rules of game playing, organized into two levels: general game structure (moves, turns, goals) and graph structure (the board on which the games are played). Concerning moves, one can make the following distinction, for instance in sabotage games (ibid., 136):

a) Local versus global moves: whether players are localised in the graph (like Traveller in the sabotage game), or can range at random (like Remover in that game).

b) Arbitrary versus definable moves. Can Remover delete any links, or must he follow some explicit definition?

c) Stepwise versus uniform moves. In each round, does Remover cut one link, or more than one link, uniformly defined?

d) Players can stay within a graph, or jump to a changed graph.

Regarding the specification of conditions for winning and losing, there are also different possibilities (ibid.):

a) The goal-region is an area that the players must avoid or want to be in. This amounts to the specification of a unary property of nodes in the graph.

b) An entangled goal is defined by a binary relation between two player's positions, as in the meet/avoid games where one player loses (and the other wins) if she meets with the other player. In this case the binary relation is the identity relation.

c) Finally, there may even be higher-level procedural goals, sensitive not only to the players' positions but also to the way in which they travelled there.

Subsequently, these game types and design parameters have been studied by a number of researchers, including Dazhu Li at Tsinghua, and Chris Mierzewski and Francesca Zaffora at Stanford. In what follows, we will mainly review the work of Dazhu Li, who was a PhD student under the supervision of Fenrong Liu, Alexandru Baltag and Johan van Benthem at the Tsinghua University—University of Amsterdam Joint Research Centre for Logic.

Definable Link Cutting in Graph Games

In the sabotage games that we have seen, Remover cuts the link globally, each time an arbitrary link is chosen. These are two elements that one can change as parameters, to design a new game and study it. This was pursued in Dazhu Li's "Losing Connection: The Modal Logic of Definable Link Deletion" (2020). The paper studied those sabotage games in which links are removed in a local and definable way. A definable sabotage modal logic (S_dML) was proposed, which extends standard modal logic with a link deletion operator $[-\varphi]\psi$. $[-\varphi]\psi$ is read as "after Remover deletes the φ-links starting from the current position of Traveller, ψ holds". To illustrate, consider the following example (Li 2020, 718):

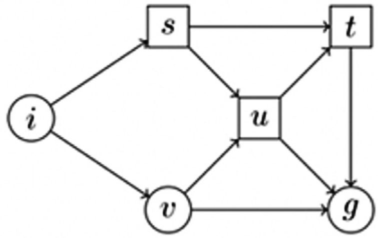

In the graph, the two kinds of shapes, square and circle, denote two different atomic properties of the nodes. The starting-node of Traveller is *i*, and her goal-nodes are *t* and *g*. Assume that the propositional atoms *p* and *q* refer to the properties denoted with circles and squares, respectively. Then we are able to express the facts of the game with formulas of the logic S_dML. For instance, that 'after Remover deletes the links from *v* to any circle points (here only *g*), Traveller still can move to a square node (here *u*)' and can be expressed as the truth at *v* of the formula $[-p]\Diamond q$. Moreover, S_dML-formulas can also describe the winning strategies for players. The formula $[-p]\Box[-q]\Box\bot$, for example, states that Remover can stop Traveller successfully by removing the links from the position of Traveller to the circle nodes in the first round and cutting the links pointing to the square nodes in the second round.

The new language can define many complex properties that are not definable in basic modal logic, but this leads to a drastic increase in computational complexity. The paper proves that S_dML does not have the tree model property or the finite model property, and its satisfiability problem is undecidable. Despite the relatively minor addition to basic modal logic, and in contrast with the decidability of the semantically similar dynamic epistemic logics of link deletion (van Benthem and Liu 2007), the high complexity of S_dML is surprising. Li (2020) identifies the locality of the updates as the culprit.

Locality also led to another problem. In DEL and its extensions, one can usually obtain a complete set of recursion axioms. These are equivalences of the form $AB\varphi \leftrightarrow BA\varphi$, where *A* is the dynamic operator and *B* is some other operator of the language. (One might also need some additional *B*-free components on the right-hand side). This generally allows a recursive removal of *A* from any formula, so showing that every formula with *A* is equivalent to one without it, and thereby showing the completeness of the axiomization, given the completeness of the logic without *A*. But consider the formula $[-\varphi]\Box\psi$: after pushing $[-\varphi]$ into the scope of $\Box$, the model change is *no longer* local, and reference to the node where the formula is evaluated is lost. One idea for fixing this is to extend the language with hybrid operators. Li showed that S_dML can be embedded into the hybrid logic with nominals, the at-operator @ and the down-arrow operator ↓. The problem of finding a complete set of recursion axioms for the logic extending S_dML with hybrid operators was left open.

Supervised Learning Games

Graph games can be used to analyse the scenario of learning and teaching, as played by two agents Learner and Teacher. Typically, the learning process has game-like features, as the teacher wants to correct the student's mistakes, to make sure they avoid them in the future. This interactive feature was studied by Dazhu Li, Alexandru Baltag and Mina Young Pedersen in their LORI paper "On the Right Path: A Modal Logic for Supervised Learning" (2019). The paper was extended and included in Li's Dissertation "Formal Threads in Social Fabric: Studies in the Logical Dynamics of Multi-Agent Interaction" (2021). Consider the following dialogue between a teacher (T) and a learner (L) who is trying to learn a logical proof (Li 2021, 6–7):

> *Example.* After checking a proof written by Learner (L), Teacher (T) begins to talk:
>
> T: You did not prove the theorem yet.
> L: Why? I started with the axioms, showed intermediate lemmas step by step, and finally reached the statement of the theorem.
> T: Your final step to show the theorem that is the goal is correct, but you in fact arrived there by accident, as the inference from lemma α to lemma β in your proof is wrong.
> L: Oops! I see. Then, my steps after β do not make sense. But, how about a new lemma proving γ from α? Now I think I can get to the theorem.
> T: Alas, γ cannot be inferred from α either, a potential mistake. But actually, you miss another lemma δ that can be derived from α. I believe you might be able to show the theorem with it.
> L: Thanks! You are right! Now I am going to search for a correct proof with δ.

The short episode suggests several interesting aspects of the learning process. One is that there are *different kinds of mistakes*: actual mistakes made, and potential mistakes to be avoided. To distinguish them, we need to know how Learner arrived at the current position: *the history matters*. The removing by Teacher of mistakes that were actually made by Learner is an action that modifies the history of the Learner's inferential moves (and makes all further moves based on that history suspect), while eliminating potential mistakes affects the future from the current point. Also, Teacher's pointing out an actual mistake removes the whole actual history after that step, resetting Learner to the last point before the mistake. Besides, Teacher may point Learner to facts that were ignored. In terms of abstract game design, this calls for a powerful Teacher: Teacher should be capable of *adding links to graphs*. Moreover, Learner may not win even though the goal-region is reached: the goal-region should be reached in the right way.

To capture these features of the interaction, the LORI paper (Baltag, Li and Pedersen 2019, 3) defined *supervised learning games* (SLG) as follows:

> *Supervised learning games.* The game is played on a graph with two relations R_L and R_T (representing the inferences conjectured by Learner and the correct inferences that are observed by Teacher, respectively), a starting-node s and a goal-node g. In each round, Learner moves along an R_L-link from her current position t to u, and meanwhile, the new history of her movements is obtained by replacing $(s, \dots, t)$, the history formed in the last round, with $(s, \dots, t, u)$. Teacher then does nothing or takes one of the three actions:
>
> (a) Add an R_T-link to R_L that has not been added to the latter relation yet,
>
> (b) Choose an R_L-link (a, b) that is not an R_T-link from the sequence $(s, \dots, t, u)$, and remove the whole actual history after that step, resetting Learner to the last point before the link (we use $(s, \dots, t, u)|_{(a,b)}$ for this action),
>
> (c) Remove an R_L-link that is not of R_T and does not occur in the sequence $(s, \dots, t, u)$.
>
> It ends if Learner arrives at g through an R_T-path $(s, \dots, g)$ (i.e., every link of the sequence is an R_T-link) or cannot make a move, with both players winning in the former and losing in the latter.

Note that the game is not zero-sum: both Learner and Teacher have the same goal. A logic of supervised learning (LSL) is developed in the paper. *A model* is a graph with two relations, R_L and R_T, and a valuation function. Formulas are evaluated at *sequences of nodes*, each of which stands for a learning process. Teacher's actions of type (b) and (c) are expressed by two operators: $\langle - \rangle_{on}\varphi$ is read as "φ is the case after deleting a mistake on the current sequence", and $\langle - \rangle_{off}\varphi$ is read as "after removing a mistake that is not on the path, φ holds". $\langle + \rangle\varphi$ is also used to express actions of type (a). From this semantics, one can see that the logic can define both the actions and winning positions of players in finite games. A follow-up paper has more technical results on the properties of LSL (Baltag, Li and Pedersen 2022).

Logic of Hide and Seek Games

One of the games introduced in van Benthem and Liu (2020) is familiar from childhood: hide and seek. As a graph game it has the special feature that the goals of the two players are entangled. Here is a formal definition from the WoLLIC paper "On the Subtle Nature of a Simple Logic of the Hide and Seek Game" by Dazhu Li, Sujata Ghosh, Fenrong Liu and Yaxin Tu (2021, 201):

> *Definition (Hide and seek games).* Given a graph, two players Hider and Seeker are located at two different nodes. In each round, Hider and Seeker, in turn, move along an arrow. The goal of Seeker is to meet Hider, while the goal of Hider is to avoid Seeker. Also, a player wins immediately once the other gets stuck.

The *language* of LHS for studying these games is based on two disjoint sets P_H and P_S of propositional variables that refer to the properties of the Hider's and Seeker's current positions, respectively. The language also contains two modalities [H] and [S] to characterize the moves of Hider and Seeker, respectively. (And, as usual, ⟨H⟩ and ⟨S⟩ are the duals.) In addition, a crucial component of the language is a propositional constant *I*, expressing that "the two players are at the same position", namely, Seeker has already caught Hider.

Formulas are evaluated at a pair of graph nodes (*h*, *s*), representing the position of Hider and Seeker, respectively. Variables in P_H are evaluated at the left node (*h*) and those in P_S are evaluated at the right node (*s*). Constant *I* is satisfied only when the two points are identical (*h*=*s*). Some examples of valid formulas of the logic are the following:

$$\langle H\rangle(I \wedge \varphi) \to [H](I \to \varphi)$$

$$\langle R\rangle(I \wedge \varphi) \to [R](I \to \varphi)$$

$$I \to (\langle H\rangle\top \leftrightarrow \langle S\rangle\top)$$

$$I \to ([S]\langle H\rangle I \wedge [H]\langle S\rangle I)$$

One subtle feature of the semantics is that there is an "evaluation-gap" between the two points of an evaluation pair (*s*, *t*). When considering the atomic properties of *s*, the language can only use the variables in P_H, but not the ones in P_S. This leads to some interesting properties of LHS. Although syntactically similar to basic model logic, it is essentially incomparable in terms of expressivity. (This is shown by giving a suitable variant to the notion of bisimulation.)

The constant *I* also has some logical properties that are not so obvious. First, the tree model property fails, as can easily be seen from the formula $I \wedge \langle H \rangle I$: any model satisfying it must contain a loop. The paper further showed that the logic LHS does not have the finite model property, and that its satisfiability problem is undecidable. The author commented that in this respect the complexity introduced by *I* is similar to that of equality operators in other logics, e.g., the Gödel class in Goldfarb (1984) and the logic of functional dependence of Baltag and van Benthem (2021).

Conclusion

Our present concern is the history of research in China, and so the focus has been on the work of people in China and their joint work with international collaborators, mostly in Auckland and Amsterdam. But the research community is open and ideas travel. There are constantly emerging new works in this field. For instance, in addition to the dissertations of Zhen Liang, Zuojun Xiong and Dazhu Li, two more PhD dissertations were recently produced: "Dynamics Logics of Networks: Information Flow and the Spread of Opinion" by Zoé Christoff at the ILLC in Amsterdam in 2016, and "In Search of Homo Sociologicus" by Yunqi Xue at the Graduate Center of CUNY in 2017. Sonja Smets and her group in Amsterdam have been a major force for the development of social network logic. She brought social network logic closer to social sciences by her research on important social phenomena: informational cascades in Baltag et al (2013), echo chambers in Pedersen et al (2019), and polarization in Pedersen et al (2020). The logical features of social group creation were studied in Smets and Velázquez-Quesada (2017, 2020), in which a threshold approach was proposed to model network creation, and the key idea was that an agent would add someone to her social network if and only if the distance between them is smaller or equal than the given threshold. Another earlier work that is worth mentioning is Ruan and Thielscher (2011), which extended DEL with new operators of "follow" and "unfollow" and applied it to analyse the well-known problem of "revolt or stay-at-home", where social networks play an important role in agents' knowledge acquisition and decision-making. Van Benthem (2015) discussed how fixed-point logics, both modal and first-order, can describe various kinds of dynamic limit behaviour in social networks, including convergence, oscillation and divergence. Christoff, Hansen and Proietti (2016) introduced a new notion of *reflective social influence* and proposed a formal framework for reasoning about an individual's private opinions and public behaviour under the dynamics of social influence. Rendsvig (2017) showed the update mechanism in network automata can be emulated using action models in DEL and identified a class of action models that captures the best-response dynamics on a graph. Christoff and Grossi (2017) gave a characterization of the stabilization of diffusion in terms of neighbourhood structures, and showed how the monotone μ-calculus can express their relevant properties. Morrison and Naumov (2020) proposed a new logic system to study the situation in which an agent conforms to multiple social groups that she belongs to instead of one group of peers, and a topological structure of the network was proposed. In the area of graph games and logic, Grossi and Rey (2019) proposed a poison modal logic to describe winning positions in games and bridged it with notions of credulous admissibility sets in argumentation theory, and non-trivial semi-kernels in graph theory. Blando, Mierzewski and Areces (2020) studied poison games systematically using three variants of modal memory logics and compared their expressive power. Van Benthem, Mierzewski and Blando (2020) developed a logic for removing nodes from graphs and studied its logical properties. Declan Thompson applied a game-theoretic approach to network automata in Seligman and Thompson (2015), and extended this to the logical characterization of Nash equilibria in Thompson (2020). The area thus seems to be flourishing and there is even more happening than we are aware of.

By this brief survey of ten years' development of social network logics in China, we hope to have shown that a logical perspective on reasoning about the social aspects of our life is interesting and attractive to researchers and others. Going back to the Chinese philosophy that originally inspired this research direction, we feel that we are just beginning our journey, and only starting to get a clear picture of social relations and social interactions. No doubt this logical research has formed a solid foundation to analyse more complicated social phenomena. Looking into the future, introducing more concrete ideas from Chinese philosophy will definitely enrich the existing approaches to social network logic, and may eventually capture further subtleties of our reasoning about ourselves. We are on the road.

Acknowledgement

We would like to thank Jeremy Seligman for reading through various versions of this paper and providing useful comments and corrections. We thank Zhen Liang for providing a summary of his works, and Yi Wang for sharing his papers with us. We are grateful to Bo Chen for inviting us to this special issue. We are lucky to have Jan Vrhovski as the guest editor who has been so patient with us. Finally, we want to thank Johan van Benthem and the two anonymous reviewers for their very helpful comments.

References

Balbiani, Philippe, and Saúl F. González. 2020. "Indexed Frames and Hybrid Logics." In *Advances in Modal Logic*, edited by Nicola Olivetti, Rineke Verbrugge, Sara Negri, and Gabriel Sandu, volume 13, 53–72. London: College Publications.

Baltag, Alexandru, Lawrence S. Moss, and Slawomir Solecki. 1998. "The Logic of Public Announcements and Common Knowledge and Private Suspicions." In *Proceedings of TARK 1998*, edited by Itzhak Gilboa, 43–56. Burlington: Morgan Kaufmann.

Baltag, Alexandru, Zoé Christoff, Jens Ulrik Hansen, and Sonja Smets. 2013. "Logical Models of Informational Cascades." In *Logic Across the University: Foundations and Applications*, edited by Johan van Benthem, and Fenrong Liu, 405–32. London: College Publications.

Baltag, Alexandru, Dazhu Li, and Mina Y. Pedersen. 2019. "On the Right Path: A Modal Logic for Supervised Learning." In *Proceedings of LORI 2019*, edited by Patrick Blackburn, Emiliano Lorini, and Meiyun Guo, LNCS 11813, 1–14. Berlin, Heidelberg: Springer.

———. 2022. "A Modal Logic for Supervised Learning." To appear in *Journal of Logic, Language and Information*.

Baltag, Alexandru, and Johan van Benthem. 2021. "A Simple Logic of Functional Dependence." *Journal of Philosophical Logic* 50: 939–1005.

Bernays, Edward. 1928. *Propaganda*. London: Routledge.

Blackburn, Patrick, and Jeremy Seligman. 1995. "Hybrid Languages." *Journal of Logic, Language and Information* 4: 251–72.

Blackburn, Patrick, Maarten de Rijke, and Yde Venema. 2002. *Modal Logic*. Cambridge: Cambridge University Press.

Blando, Francesca Z., Krzysztof Mierzewski, and Carlos Areces. 2020. "The Modal Logics of the Poison Game." In *Knowledge, Proof and Dynamics*, edited by Fenrong Liu, Hiroakira Ono, and Junhua Yu, Logic in Asia: Studia Logica Library, 3–23. Singapore: Springer.

Christoff, Zoé. 2016. "Dynamics Logics of Networks: Information Flow and the Spread of Opinion." PhD diss., ILLC, University of Amsterdam.

Christoff, Zoé, Jens U. Hansen, and Carlo Proietti. 2016. "Reflecting on Social Influence in Networks." *Journal of Logic, Language and Information* 25 (3): 299–333.

Christoff, Zoé, and Davide Grossi. 2017. "Stability in Binary Opinion Diffusion." In *Proceedings of LORI 2017*, edited by Alexandru Baltag, Jeremy Seligman, and Tomoyuki Yamada, LNCS 10455, 166–80. Berlin, Heidelberg: Springer.

Fagin, Ronald, Joseph Y. Halpern, Yoram Moses, and Moshe Vardi. 2004. *Reasoning about Knowledge*. Cambridge: The MIT Press.

Girard, Patrick, Jeremy Seligman, and Fenrong Liu. 2012. "General Dynamic Dynamic Logic." In *Advances in Modal Logic*, edited by Thomas Bolander, Torben Braüner, Silvio Ghilardi, and Lawrence Moss, vol. 9, 239–60. London: College Publications.

Goldfarb, Warren D. 1984. "The Unsolvability of the Gödel Class with Identity." *Journal of Symbolic Logic* 49:1237–52.

González, Saúl F. 2021. "Some Dynamic Extensions of Social Epistemic Logic." Paper presented at *International Workshop on Logic Aspects in Multi-Agent Systems and Strategic Reasoning*, 3–4 May 2021.

Grossi, Davide, and Simon Rey. 2019. "Credulous Acceptability, Poison Games and Modal Logic." In *Proceedings of AAMAS 2019*, edited by Edith Elkind, Manuela Veloso, Noa Agmon, and Matthew E. Taylor, 1994–96. Montreal: AAMAS.

Hintikka, Jaakko. 1962. *Knowledge and Belief: An Introduction to the Logic of the Two Notions*. Cornell: Cornell University Press.

Li, Dazhu. 2020. "Losing Connection: The Modal Logic of Definable Link Deletion." *Journal of Logic and Computation* 30: 715–43.

———. 2021. "Formal Threads in the Social Fabric: Studies in the Logical Dynamics of Multi-Agent Interaction." PhD diss., Department of Philosophy, Tsinghua University and ILLC, University of Amsterdam.

Li, Dazhu, Sujata Ghosh, Fenrong Liu, and Yaxin Tu. 2021. "On the Subtle Nature of a Simple Logic of the Hide and Seek Game." In *Logic, Language, Information and Computation (WoLLIC 2021)*, edited by Alexandra Silva, Renata Wassermann, and Ruy de Queiroz, LNCS 13038, 201–18. Cham: Springer.

Liang, Zhen. 2017. "An Axiomatisation for Minimal Social Epistemic Logic." In *Proceedings of LORI 2017*, edited by Alexandru Baltag, Jeremy Seligman, and Tomoyuki Yamada, LNCS 10455, 664–69. Berlin, Heidelberg: Springer.

———. 2020. "Towards Axiomatisation of Social Epistemic Logic." PhD diss., Department of Philosophy, University of Auckland.

Liang, Zhen, and Jeremy Seligman. 2011. "A Logical Model of the Dynamics of Peer Pressure." *Electronic Notes in Theoretical Computer Science* 278: 275–88.

Liu, Fenrong. 2008. "Change for the Better. Preference Dynamics and Agent Diversity." PhD diss., ILLC, University of Amsterdam.

——. 2011. *Reasoning about Preference Dynamics*. Synthese Library, vol. 354. Dordrecht: Springer.

Liu, Fenrong, Jeremy Seligman, and Patrick Girard. 2014. "Logical Dynamics of Belief Change in the Community." *Synthese* 191: 2403–31.

Liu, Fenrong, and Emiliano Lorini. 2017. "Reasoning about Belief, Evidence and Trust in a Multi-Agent Setting." In *PRIMA 2017: Principles and Practice of Multi-Agent Systems*, edited by Bo An, Ana Bazzan, João Leite, Serena Villata, and Leendert van der Torre, LNCS 10621, 71–89. Cham: Springer.

Liu, Fenrong 刘奋荣, and Jeremy Seligman 谢立民. 2018. "Guanyu shejiao wangluo zhong zhuti xingwei de tuili he yuce 关于社交网络中主体行为的推理和预测 (Reasoning and Making Predictions about Agent's Behaviors in Social Networks)." *Jinan xuebao* 暨南学报 (*Journal of Jinan*) 12: 1–8.

Morrison, Colby, and Pavel Naumov. 2020. "Group Conformity in Social Networks." *Journal of Logic, Language and Information* 29: 3–19.

Pedersen, Mina Young, Sonja Smets, and Thomas Ågotnes. 2019. "Analyzing Echo Chambers: A Logic of Strong and Weak Ties." In *Proceedings LORI 2019*, edited by Patrick Blackburn, Emiliano Lorini, and Meiyun Guo, LNCS 11813, 183–98. Berlin, Heidelberg: Springer.

———. 2020. "Further Steps Towards a Logic of Polarization in Social Networks." In *CLAR 2020: Logic and Argumentation*, edited by Mehdi Dastani, Huimin Dong, and Leon van der Torre, LNAI 12061, 324-345. Cham: Springer.

Rendsvig, Rasmus K. 2017. "Diffusion, Influence and Best-Response Dynamics in Networks: An Action Model Approach." ArXiv: 1708.01477.

Ruan, Ji, and Michael Thielscher. 2011. "A Logic for Knowledge Flow in Social Networks." In *AI 2011: Advances in Artificial Intelligence*, edited by Dianhui Wang, and Mark Reynolds, LNCS 7106, 511–20. Berlin, Heidelberg: Springer.

Sano, Katsuhiko. 2010. "Axiomatizing Hybrid Products: How Can We Reason Many-Dimensionally in Hybrid Logic?" *Journal of Applied Logic* 8: 459–74.

———. 2017. "Axiomatizing Epistemic Logic of Friendship via Tree Sequent Calculus." In *Proceedings of LORI 2017*, edited by Alexandru Baltag, Jeremy Seligman, and Tomoyuki Yamada, LNCS 10455, 224–39. Berlin, Heidelberg: Springer.

Seligman, Jeremy, Fenrong Liu, and Patrick Girard. 2011. "Logic in the Community." In *Proceedings of the 4th Indian Conference on Logic and its Applications*, edited by Mohua Banerjee, and Anil Seth, LNCS 6521, 178–88. Berlin, Heidelberg: Springer.

———. 2013. "Facebook and Epistemic Logic of Friendship." In *Proceedings of TARK 2013*, edited by Burkhard C. Schipper, 229–38. Chennai, India.

Seligman, Jeremy, and Declan Thompson. 2015. "Boolean Network Games and Iterated Boolean Games." In *Proceedings LORI 2015*, edited by Wiebe van der

Hoek, Wesley H. Holliday, and Wenfang Wang, LNCS 9394, 353–65. Berlin, Heidelberg: Springer.
Shi, Chenwei. 2021. "Collective Opinion as Tendency towards Consensus." *Journal of Philosophical Logic* 50: 593–613.
Skyrms, Brain. 2014. *Social Dynamics*. Oxford: Oxford University Press.
Smets, Sonja, and Fernando R. Velázquez-Quesada. 2017. "How to Make Friends: A Logical Approach to Social Group Creation." In *Proceedings of LORI 2017*, edited by Alexandru Baltag, Jeremy Seligman, and Tomoyuki Yamada, volume 10455 of LNCS, 377–90. Berlin, Heidelberg: Springer.
———. 2020. "A Closeness- and Priority-Based Logical Study of Social Network Creation." *Journal of Logic, Language and Information* 29: 21–51.
Thompson, Declan. 2020. "Local Fact Change Logic." In *Knowledge, Proof and Dynamics*, edited by Fenrong Liu, Hiroakira Ono, and Junhua Yu, Logic in Asia: Studia Logica Library, 73–96. Singapore: Springer.
van Benthem, Johan. 2005. "An Essay on Sabotage and Obstruction." In *Mechanizing Mathematical Reasoning*, edited by Dieter Hutter, and Werner Stephan, LNCS 2605, 268–76. Berlin, Heidelberg: Springer.
———. 2014. *Logic in Games*. Cambridge: The MIT Press.
———. 2015. "Oscillation, Logic and Dynamical Systems." In *The Facts Matter. Essays on Logic and Cognition in Honour of Rineke Verbrugge*, edited by Sujata Ghosh, and Jakub Szymanik, 9–22. London: College Publications.
van Benthem, Johan, and Fenrong Liu. 2007. "Dynamic Logic of Preference Upgrade." *Journal of Applied Non-Classical Logics* 17: 157–82.
———. 2020. "Graph Games and Logic Design." In *Knowledge, Proof and Dynamics*, edited by Fenrong Liu, Hiroakira Ono, and Junhua Yu, Logic in Asia: Studia Logica Library, 125–46. Singapore: Springer.
van Benthem, Johan, Krzysztof Mierzewski, and Francesca Z. Blando. 2020. "The Modal Logic of Stepwise Removal." *The Review of Symbolic Logic* 1–28. https://doi:10.1017/S1755020320000258.
van Benthem, Johan, Fenrong Liu, and Sonja Smets. 2021. "Logico-Computational Aspects of Rationality." In *The Handbook of Rationality*, edited by Markus Knauff, and Wolfgang Spohn, 185–200. Cambridge: The MIT Press.
van der Hoek, Wiebe, Louwe B. Kuijer, and Yi N. Wang. 2019. "Who Should Be My Friends? Social Balance from the Perspective of Game Theory." In *Proceedings of LORI 2019*, edited by Patrick Blackburn, Emiliano Lorini, and Meiyun Guo, LNCS 11813, 370–84. Berlin, Heidelberg: Springer.
———. 2020. "Logics of Allies and Enemies: A Formal Approach to the Dynamics of Social Balance Theory." In *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence (IJCAI-20)* 210–216.

von Wright, Georg H. 1951. *An Essay on Modal Logic*. Amsterdam: North-Holland Publishing Company.

Wang, Yanjing, and Jeremy Seligman. 2018. "When Names Are Not Commonly Known: Epistemic Logic with Assignments." In *Advances in Modal Logic*, edited by Guram Bezhanishvili, Giovanna D'Agostino, George Metcalfe, and Thomas Studer, vol. 12, 611–28. London: College Publications.

Xiong, Zuojun. 2017. "On the Logic of Multicast Messaging and Balance in Social Network." PhD diss., University of Bergen.

Xiong, Zuojun, and Thomas Ågotnes. 2020. "On the Logic of Balance in Social Networks." *Journal of Logic, Language and Information* 29: 53–75.

Xiong, Zuojun, Thomas Ågotnes, Jeremy Seligman, and Rui Zhu. 2017. "Towards a Logic of Tweeting." In *Proceedings of LORI 2017*, edited by Alexandru Baltag, Jeremy Seligman, and Tomoyuki Yamada, LNCS 10455, 49–64. Berlin, Heidelberg: Springer.

Xu, Ming. 1988. "On Some *U*, *S*-Tense Logic." *Journal of Philosophical Logic* 17: 181–202.

Xue, Yunqi. 2017. "In Search of Homo Sociologicus." PhD diss., The Graduate Center, The City University of New York.

Zhu, Ri. (forthcoming). "Social Announcement Logic." PhD diss., Department of Philosophy, University of Auckland.