Formal Verification of Diffusion Auctions

Rustam Galimullin[1], Munyque Mittelmann[2], Laurent Perrussel[3]
[1]University of Bergen, Norway
[2]CNRS, LIPN, Université Sorbonne Paris Nord, France
[3]IRIT, Université Toulouse Capitole, France
rustam.galimullin@uib.no, mittelmann@lipn.univ-paris13.fr, laurent.perrussel@irit.fr

Abstract

In diffusion auctions, sellers can leverage an underlying social network to broaden participation, thereby increasing their potential revenue. Specifically, sellers can incentivise participants in their auction to diffuse information about the auction through the network. While numerous variants of such auctions have been recently studied in the literature, the formal verification and strategic reasoning perspectives have not been investigated yet.

Our contribution is threefold. First, we introduce a logical formalism that captures the dynamics of diffusion and its strategic dimension. Second, for such a logic, we provide model-checking procedures that allow one to verify properties like the Nash equilibrium, and that pave the way towards checking the existence of sellers' strategies. Third, we establish computational complexity results for the presented algorithms.

Introduction

In auction theory and mechanism design (Nisan et al. 2007), the set of participants is typically fixed and socially independent, in the sense that any underlying social network among agents is not taken into account. In contrast, by leveraging agents' social networks, a seller could use buyers' connections to promote the auction (Guo and Hao 2021). This has a clear advantage: a larger market may include participants with higher valuations, leading to a potential increase in social welfare or the sellers' revenue. On the other hand, buyers act as competitors and have no incentives to invite more participants, as doing so would increase competition and reduce their likelihood of securing the item being auctioned.

The challenge of encouraging participants to propagate the auction among their social connections has recently sparked interest in the mechanism design community (Zhao 2021). In particular, it led to the introduction of *diffusion auctions* (Zhao et al. 2018; Li et al. 2022), where sellers propose incentives to buyers so that they can benefit from inviting their neighbours. The intuition is that the mechanism guarantees that the buyer's new utility after propagating the auction is not less than her utility of participating in the auction with the original participants. Their main benefit is the increase in the number of participants while guaranteeing economic properties such as incentive-compatibility or optimality (Zhang, Zheng, and Zhao 2024). Yet, two critical aspects remain unexplored — the strategic behaviour of sellers in diffusion auctions, especially when multiple sellers compete to reach the most valuable buyers, and the formal verification of such mechanisms.

In the last decades, a number of logics have been proposed to reason about agents' strategic capabilities with prime examples being *Coalition Logic* (CL) (Pauly 2002), *Alternating-time Temporal Logic* (ATL) (Alur, Henzinger, and Kupferman 2002), and *Strategy Logic* (Mogavero et al. 2014). Combined with model-checking techniques (Clarke et al. 2018), these frameworks provide powerful tools for specification and verification of multi-agent systems, with applications to several problems, from the analysis of voting protocols (Belardinelli et al. 2021; Jamroga, Kurpiewski, and Malvone 2022) to the verification of auctions and mechanism design (Mittelmann et al. 2025, 2023).

In this paper, we provide a formal framework for specification and verification of strategic properties in diffusion auctions. In doing so, we combine the intuitions from *social network logics* (Pedersen 2024), *dynamic epistemic logic* (DEL) (van Ditmarsch, van der Hoek, and Kooi 2008), as well as the aforementioned CL and ATL. We believe that this is *the first logic-based approach to formal verification of diffusion auctions and strategic abilities of sellers in them.*

Contribution We introduce the *n-seller logic for diffusion incentives* $\mathcal{L}^n$ and its *strategic* variant $\mathcal{SL}^n$. These logics are interpreted on diffusion auction mechanism models that are quite general and thus capture a wide variety of mechanisms. Both $\mathcal{L}^n$ and $\mathcal{SL}^n$ allow us to capture the dynamics of diffusion of information about auctions and their strategic dimension. By 'dynamics' here, we mean the change of the underlying social network as a result of sellers proposing incentives to buyers to invite their neighbours to an auction. For these logics, we provide model-checking procedures that allow one to verify properties such as Nash equilibrium, and that pave the way towards checking the existence of sellers' strategies. For the presented algorithms, we establish computational complexity results.

Diffusion Auctions With Multiple Sellers

We start by presenting a formal framework for multiple-seller auctions, where each seller is selling (a copy of) the same item. Sellers and buyers in such a setting are connected via an underlying social network, whose structure the sellers can try to exploit by incentivising their direct neighbours (i.e., buyers participating in the sellers' auctions) to invite all their friends to join the corresponding auction.

Let $\mathsf{S} = \{\sigma_1, \sigma_2, ..., \sigma_n\}$ be a finite non-empty set of n names of sellers, and $\mathsf{B} = \{\beta_1, \beta_2, ...\}$ be a countable set of names of buyers, such that $\mathsf{S} \cap \mathsf{B} = \emptyset$. Also, let $\mathsf{Nom} = \mathsf{S} \cup \mathsf{B}$ denote the total set of agent names, or nominals. We will also write $\mathsf{B}^\bullet = \mathsf{B} \cup \{\bullet\}$ and $\mathsf{Nom}^\bullet = \mathsf{Nom} \cup \{\bullet\}$, where nominal $\bullet$ intuitively stands for 'the current agent'. Finally, let $\mathsf{Terms} = \{ut_\alpha \mid \alpha \in \mathsf{Nom}^\bullet\}$ be a set of terms denoting utilities of agents ut_α and a term $ut_\bullet$ denoting the utility of the current agent.

Definition 1. The language $\mathcal{L}^n$ of the *n-seller logic for diffusion incentives* is defined by the following grammar:

$$\varphi := \alpha \mid (z_1 t_1 + ... + z_m t_m) \geqslant z \mid \neg\varphi \mid (\varphi \wedge \varphi) \mid \Box\varphi \mid [\sigma_1 : \beta_1, ..., \sigma_n : \beta_n]\varphi \mid \heartsuit\gamma,$$

where $\alpha \in \mathsf{Nom}$, $z \in \mathbb{Z}$, $t_i \in \mathsf{Terms}$, $\sigma_i, \sigma_j \in \mathsf{S}$ with $\sigma_i \neq \sigma_j$, $\beta_i \in \mathsf{B}^\bullet$, and $\gamma \in \mathsf{Nom}^\bullet$. Here, $\Box\varphi$ means 'all friends of the current agent satisfy φ', and $\heartsuit\gamma$ means that agent named γ, which can be either a seller or a buyer, gets an item in the current configuration of a mechanism.

Constructs $[\sigma_1 : \beta_1, \sigma_2 : \beta_2, ..., \sigma_n : \beta_n]\varphi$ (abbreviated as $[\overline{\sigma} : \overline{\beta}]\varphi$), where $n = |\mathsf{S}|$, capture the concurrent information diffusion about auctions. This is done by sellers σ_i incentivising the respective buyers β_i, i.e., paying them some sum, to invite all their friends to join the seller's auction. Clause $\sigma_i : \bullet$ denotes the case in which seller σ_i does not incentivise anyone, i.e., she does nothing or skips her turn. Hence, we will write, e.g., $[\sigma_1 : \beta_1, \sigma_2 : \beta_2]\varphi$ if only sellers named σ_1 and σ_2 do not choose $\bullet$. Given $[\overline{\sigma} : \overline{\beta}]\varphi$, we will denote as $\overline{\sigma}^{\setminus\bullet}$ the set of agents $\sigma_i \in \overline{\sigma}$ such that the corresponding $\beta_i \neq \bullet$. Observe that even though this modality is concurrent, i.e., everyone is making moves in parallel, we can easily model consecutive moves by selecting $\bullet$ for all the agents who are not playing in the current turn. For the case of 1-seller auctions $\mathcal{L}^1$, we will write $[\beta]\varphi$ instead of $[\sigma : \beta]\varphi$. Finally, having a sequence $\mathit{Upd} = \sigma_1 : \beta_1, ..., \sigma_n : \beta_n$, we will denote by $\mathit{Upd}(\sigma_i)$ the corresponding buyer's name β_i.

Duals are defined as $\Diamond\varphi := \neg\Box\neg\varphi$ and $\langle\overline{\sigma} : \overline{\beta}\rangle\varphi := \neg[\overline{\sigma} : \overline{\beta}]\neg\varphi$. For the linear inequalities[1], we can use the following abbreviations: $t_1 - t_2 \geqslant z$ for $t_1 + (-1)t_2 \geqslant z$, $t_1 \geqslant t_2$ for $t_1 - t_2 \geqslant 0$, $t_1 \leqslant z$ for $-t_1 \geqslant -z$, $t_1 < z$ for $\neg(t_1 \geqslant z)$, and $t_1 = z$ for $(t_1 \geqslant z) \wedge (t_1 \leqslant z)$. We can also use rational numbers in our formulas via abbreviations (e.g., $t \geqslant \frac{1}{2}$ is an abbreviation for $2t \geqslant 1$). All other standard abbreviations of logic and the rules for removing parentheses hold.

Example 1. With our language, we can express various desirable properties of mechanisms, both static and dynamic. The following formulas are examples for $\mathcal{L}^1$.

- $ut_\alpha = 3 \wedge [\alpha](ut_\alpha > 3)$ for 'the utility of agent α is 3, and after she was incentivised by the seller to invite her friends to participate in the auction, her utility increased'.
- $ut_\bullet = 5 \wedge \Box(ut_\bullet \geqslant 5) \wedge \Diamond(\alpha \wedge \heartsuit\alpha)$ for 'the utility of the current agent is 5, and all her friends have utilities of at least 5, and she also has a friend α who gets an item'.

Formulas of $\mathcal{L}^n$ are interpreted on diffusion auction mechanisms.

Definition 2. A *market network with n sellers* $\mathcal{M}$ is a tuple (Agt, F, Bdg, V, I, N), where

- $Agt = B \cup S$ is the set of agents, where $B = \{a, b, c, ...\}$ is a non-empty set of *buyers*, and $S = \{s_1, ..., s_n\}$ is a non-empty set of *sellers*, and $B \cap S = \emptyset$;
- $F : Agt \to 2^B$ is a symmetric irreflexive *friendship* (neighbour) relation;
- $Bdg : Agt \to \mathbb{Q}^+ \cup \{0\}$ is a non-negative *budget* for each agent;
- $V : B \to \mathbb{Q}^+ \cup \{0\}$ s.t. $V(a) \leqslant Bdg(a)$ assigns to each buyer a non-negative *valuation* of the item being sold;
- $I : B \times S \to \mathbb{Q}^+ \cup \{0\}$ assigns to each buyer the non-negative *incentive* that each seller is willing to pay to them to invite their friends;
- $N = N_{\mathsf{S}} \cup N_{\mathsf{B}}$ is a *naming function*, where $N_{\mathsf{S}} : \mathsf{S} \to S$ and $N_{\mathsf{B}} : \mathsf{B} \to B$ are surjective functions[2].

An *n-seller diffusion auction mechanism* (n-DAM, or DAM) M is a tuple $(\mathcal{M}, P, Pay, U)$, where

- $\mathcal{M} = (Agt, F, Bdg, V, I, N)$ is a market network with n sellers;
- $P_{\mathcal{M}} : Agt \to \{0, 1\}$ is the *allocation (placement) function*, which specifies whether an agent receives an item in an auction conducted within the market network $\mathcal{M}$;
- $Pay_{\mathcal{M}} : B \to \mathbb{Q}^+ \cup \{0\}$ is the *payment function*, which specifies the value each buyer should pay in an auction run within the market network $\mathcal{M}$;
- $U_{\mathcal{M}} : Agt \to \mathbb{Q}^+ \cup \{0\}$ is the *utility function*.

We omit subscripts $\mathcal{M}$ whenever it does not cause confusion. We write M, a to refer to a specific agent a in M.

Observe that the definitions of the allocation, payment, and utility functions do not specify their exact details. This makes our definition of diffusion auction mechanisms general, allowing us to incorporate various types of mechanisms. The only restriction we put on the functions is that their complexity is no greater than the complexity of the model-checking problem of a given logic. As we will see later, model checking $\mathcal{L}^n$ is in P, and thus in this section we assume that P, Pay, and U are *computable in polynomial time*. While the optimal allocation function in combinatorial auctions is NP-complete (Nisan et al. 2007)[3], several mechanisms whose functions are computable in polynomial time have been proposed. Those include a strategyproof combinatorial auction (Dobzinski and Vondrák 2012), a double auction mechanism in social networks (Xu and He 2020), and McAfee's double auction mechanism (McAfee 1992).

[1]Originally, these linear inequalities in a logical context were used to capture reasoning about probabilities (Fagin, Halpern, and Megiddo 1990). Recently, they were also used to express budgets and costs in dynamic epistemic logic (Dolgorukov, Galimullin, and Gladyshev 2024). We follow the latter approach in this work.

[2]Observe that function N is well-defined since $\mathsf{S} \cap \mathsf{B} = \emptyset$.

[3]The optimal allocation function is used to compute allocations and payments in the Vickrey–Clarke–Groves mechanism (Vickrey 1961; Clarke 1971; Groves 1973).

For our running examples, we will use the *single item, multiple units, first price* (SMF) auctions.

Definition 3 (SMF Auction)**.** Given a market network with n sellers $\mathcal{M} = (Agt, F, Bdg, V, I, N)$, the placement function P is defined as follows.

For a seller s_i, let $\mathfrak{s}_i$ be the ordered set of valuations $V(a)$ of buyers a such that $a \in F(s_i)$[4]. The ordering is from the highest to the lowest valuations, where the ties are broken by the lexicographic ordering of the buyers. This set is totally ordered. We denote the first element of $\mathfrak{s}_i$ as $\mathfrak{s}_i(1)$.

We refine the sets $\mathfrak{s}_i$ to account for the distribution of items. A refinement of a totally ordered set $\mathfrak{s}_i$, denoted by $\bar{\mathfrak{s}}_i$, is defined as $\bar{\mathfrak{s}}_i = \mathfrak{s}_i \setminus \{\mathfrak{s}_j(1) | 0 < j < i\}$. Intuitively, a refinement is the set of bidders in the current auction minus those bidders that are already getting the item from some other auction. Finally, for all $s_i \in S$, if $\bar{\mathfrak{s}}_i$ is non-empty, then $P(a) = 1$ and $P(s_i) = 0$ for $V(a) = \bar{\mathfrak{s}}_i(1)$, and $P(s_i) = 1$ otherwise. Observe that in SMF auctions, it is implied that buyers strive to acquire only one (copy of the) item, and hence some sellers may end up not selling their items.

As for the Pay function, buyers, if they get an item, pay the amount equal to their valuation of the item. Function U is defined as follows: for a buyer $a \in B$, $U(a) = Bdg(a) - V(a) \cdot P(a)$; for the seller s_i, $U(s_i) = V(a) + Bdg(s_i)$, where $V(a) = \bar{\mathfrak{s}}_i(1)$, and $U(s_i) = Bdg(s_i)$ if $\bar{\mathfrak{s}}_i$ is empty.

Example 2. As an example of how the placement function works in Definition 3, consider two mechanisms in Figure 1. In both mechanisms, all valuations and incentive values are identical. Moreover, $\heartsuit$ denotes the allocation of items.

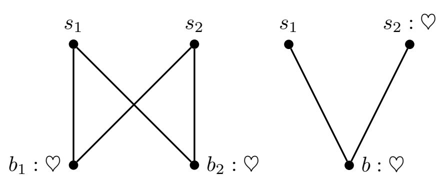

Figure 1: Mechanisms M_1 (left) and M_2 (right).

In M_1 we have that both buyers participate in both auctions run by s_1 and s_2, and each buyer wins an item. For the sellers, sets $\mathfrak{s}_1$ and $\mathfrak{s}_2$ are $\{V(b_1), V(b_2)\}$. Recall that $V(b_1) = V(b_2)$ and we use the lexicographic preference ordering to break ties. The refinement $\bar{\mathfrak{s}}_1 = \{V(b_1), V(b_2)\}$, and the refinement $\bar{\mathfrak{s}}_2 = \{V(b_1), V(b_2)\} \setminus \{V(b_1)\} = \{V(b_2)\}$. Hence, by the definition of the placement function, $P(b_1) = 1$, $P(b_2) = 1$, $P(s_1) = 0$, and $P(s_2) = 0$.

Now, let us take a look at mechanism M_2, where we again have two sellers but now only one buyer. There, $\mathfrak{s}_1 = \mathfrak{s}_2 = \{V(b)\}$, and the corresponding refinements are $\bar{\mathfrak{s}}_1 = \{V(b)\}$ and $\bar{\mathfrak{s}}_2 = \{V(b)\} \setminus \{V(b)\} = \emptyset$ (using lexicographic ordering over sellers as a tie-breaking rule). Hence, the placement function is $P(b) = 1$, $P(s_1) = 0$, and $P(s_2) = 1$. Note that following Definition 3, seller s_2 keeps her item.

[4]Observe that having sets $\mathfrak{s}_i$, and not multisets, suffice as each buyer submits at most one valuation per seller, or auction.

Since sellers can incentivise buyers to invite their friends to join an auction, our mechanisms are *dynamic*, i.e., the social network structure changes as a result of sellers' actions. To capture this, we define the *update* of a mechanism. In our definition, we assume that each buyer can be incentivised by one seller at a time. And, moreover, since buyers are rational, if they are offered incentives from two separate sellers, they choose the higher incentive, and their friends join the auction of the seller offering the higher incentive.

Definition 4. Let an n-DAM $M = ((Agt, F, Bdg, V, I, N), P, Pay, U)$ and a sequence $Upd = \sigma_1 : \beta_1, ..., \sigma_n : \beta_n$ be given. An n-DAM *updated by the concurrent invitations by n sellers* M^{Upd} is a tuple $((Agt, F^{Upd}, Bdg^{Upd}, V, I, N), P, Pay, U)$, where for all $s \in S$, if $N(\sigma_i) = s$, $Upd(\sigma_i) = \beta_i \neq \bullet$, and $s = \arg\max_{s' \in S} I(N(\beta_i), s')$, then

- $F^{Upd}(s) = F(s) \cup \{b | b \in B \text{ and } b \in F(N(\beta_i))\}$,
- $Bdg^{Upd}(s) = Bdg(s) - I(N(\beta_i), s)$ and $Bdg^{Upd}(N(\beta_i)) = Bdg(N(\beta_i)) + I(N(\beta_i), s)$,

Ties are broken by the lexicographic order of sellers' names[5].

Intuitively, given a concurrent information diffusion operator $\overline{\sigma} : \overline{\beta}$, in the updated mechanism all friends of buyer β_i will join the auction run by σ_i if σ_i is one of the sellers that incentivise β_i, and moreover, offers the highest incentive. Then, the seller's budget is reduced by the value of the incentive, and the budget of the corresponding buyer is increased by the value of the incentive.

Definition 5. Let $M = ((Agt, F, Bdg, V, I, N), P, Pay, U)$ be an n-DAM. The semantics of $\mathcal{L}^n$ is defined by induction as follows:

$M, a \models \alpha$ iff $N(\alpha) = a$
$M, a \models \neg\varphi$ iff $M, a \not\models \varphi$
$M, a \models \varphi \wedge \psi$ iff $M, a \models \varphi$ and $M, a \models \psi$
$M, a \models \Box\varphi$ iff $\forall b \in Agt : b \in F(a)$ implies $M, b \models \varphi$
$M, a \models [\overline{\sigma} : \overline{\beta}]\varphi$ iff if $\forall \sigma_i \in \overline{\sigma}^{\setminus\bullet} : N(\beta_i) \in F(N(\sigma_i))$ and $Bdg(N(\sigma_i)) \geqslant I(N(\beta_i), N(\sigma_i))$, then $M^{\overline{\sigma}:\overline{\beta}}, a \models \varphi$

$$M, a \models \heartsuit\alpha \text{ iff } \begin{cases} P(N(\alpha)) = 1 & \text{if } \alpha \neq \bullet, \\ P(a) = 1 & \text{if } \alpha = \bullet \end{cases}$$

$$M, a \models \sum_{i=1}^{m} z_i t_i \geqslant z \text{ iff } \sum_{i=1}^{m} z_i t'_i \geqslant z, \text{ where}$$

$$t'_i = \begin{cases} U(N(\alpha)) & \text{if } t_i = ut_\alpha, \\ U(a) & \text{if } t_i = ut_\bullet. \end{cases}$$

The semantics definition for the clause $M, a \models [\overline{\sigma} : \overline{\beta}]\varphi$ checks whether all sellers, who did not choose to skip their turn, are (i) incentivising the buyers that currently participate in their auction, and (ii) whether the sellers have sufficient budgets to pay the incentives. If (i) and (ii) hold, then φ is

[5]In particular, if there are two or more sellers offering a buyer the same maximal incentive, the buyer propagates the auction information of the seller that appears first in the lexicographic order.

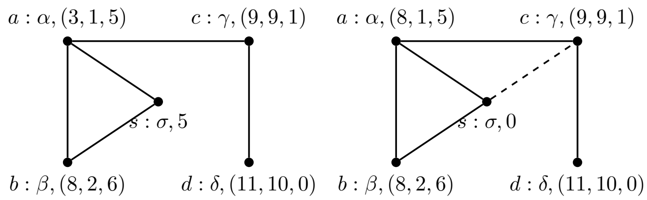

Figure 2: Mechanism M (left) and updated mechanism $M^{\sigma:\alpha}$ (right). For the seller s, her name is σ and her budget in M is 5. For buyer a, her name is α, and $(3, 1, 5)$ denotes the fact that $Bdg(a) = 3$, $V(a) = 1$, and $I(a, s) = 5$. Similarly, for other agents. The new link in M^{α} is dashed.

evaluated in the updated mechanism $M^{\overline{\sigma}:\overline{\beta}}, a$. The dual $\langle \overline{\sigma} : \overline{\beta} \rangle \varphi$ holds iff (i) and (ii) hold, and $M^{\overline{\sigma}:\overline{\beta}}, a \models \varphi$. For clauses of $\heartsuit\alpha$ and linear inequalities, we distinguish the cases of an agent named α and the current agent, denoted by nominal $\bullet$.

Example 3. Consider 1-seller SMF mechanism M and its update $M^{\sigma:\alpha}$ in Figure 2. We have, for example, that $M, a \models (ut_\sigma = 7) \wedge \heartsuit\beta \wedge \langle\alpha\rangle(ut_\sigma = 9 \wedge \heartsuit\gamma)$, where the utility of the seller σ increases after she incentivises the agent named α to invite her friends. Indeed, in mechanism M, the seller's budget is 5, and the highest valuation among the buyers participating in the auction is 2 (from agent β). Hence, the total utility of the seller is 7. Since agent β is currently the highest bidder, she would have been the winner of the auction in the current mechanism (conjunct $\heartsuit\beta$). After the seller incentivises α to invite all her friends to the auction, the budget of α increases by her incentive (i.e. by 5 to the total 8), and agent γ joins the auction (dashed line in $M^{\sigma:\alpha}$). Since her valuation is the highest, she gets the item (conjunct $\heartsuit\gamma$) and the utility of the seller increases to 9. We also have that $M, a \models \neg\langle\alpha\rangle\langle\gamma\rangle(ut_\sigma > 9)$, i.e., the seller cannot increase her utility further by trying to reach the richest buyer δ as she does not have enough budget for two rounds of referrals.

Now we turn to a 2-seller example. Consider the SMF mechanism M and its updates in Figure 3. We assume that sellers s_1 and s_2 have both budgets 1, and that for both of them and all buyers, the incentives are 1, i.e., each seller can incentivise only one buyer. Moreover, let us assume that each buyer, apart from b and e, evaluates the item as 1. Buyer b has valuation 4, and buyer e has valuation 3.

We can see that in mechanism M, each seller has utility of 2, and buyers a and c get the item (recall that we assume the lexicographic tie-breaking rule). Formally, $ut_{\sigma_1} = 2 \wedge ut_{\sigma_2} = 2 \wedge \heartsuit\alpha \wedge \heartsuit\gamma$. We can also consider a more cooperative goal of the mechanism configuration, where each agent either has the item or has a friend who has the item, expressed by $\bigwedge_{i \in \mathsf{Nom}_M} i \to (\heartsuit \bullet \vee \Diamond\heartsuit\bullet)$, where $\mathsf{Nom}_M = \{\sigma_1, \sigma_2, \alpha, ..., \zeta\}$. Mechanism M does not satisfy this goal, as, for example, the formula does not hold for agent d.

Now, assume that both sellers decide to diffuse the information about their auctions over the network. Seller s_1 incentivises buyer d to invite all her friends, i.e., e, to the auction; and similarly for seller s_2 and buyer c who invites b. The resulting update $M^{\sigma_1:\delta,\sigma_2:\gamma}$ is depicted in the middle of Figure 3. After this diffusion, the utility of seller s_1 becomes 3 (1 is spent on the incentive, and agent e bids 3), and the utility of seller s_2 is 4, i.e., for both sellers, the diffusion allowed them to increase their utilities by reaching bidders with higher valuations. Formally, we can write this as $ut_{\sigma_1} = 2 \wedge ut_{\sigma_2} = 2 \wedge [\sigma_1 : \delta, \sigma_2 : \gamma](ut_{\sigma_1} > 2 \wedge ut_{\sigma_2} > 2)$.

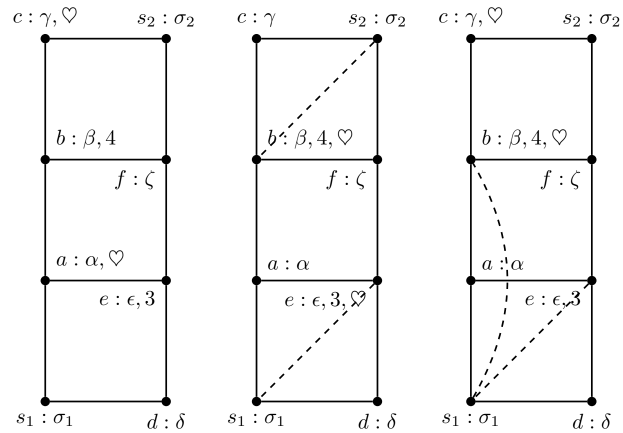

Figure 3: Mechanism M and its updates. In the mechanisms, there are two sellers, s_1 and s_2, and six buyers, $a, b, ..., f$. Agents' names are shown near the state label. New links are dashed. The allocation of the sold items is depicted by $\heartsuit$.

Interestingly, such a diffusion performed by both sellers allows us to reach the cooperative goal: formula $[\sigma_1 : \delta, \sigma_2 : \gamma] \bigwedge_{i \in \mathsf{Nom}_M} i \to (\heartsuit \bullet \vee \Diamond\heartsuit\bullet)$ is valid in mechanism M.

Although the diffusion operator $[\sigma_1 : \delta, \sigma_2 : \gamma]$ allows sellers to increase their utilities and also achieve the cooperative goal, it requires a coordinated action from the sellers. Thus, seller s_1 can increase her utility even more and outperforms seller s_2 by incentivising buyer a instead of d, and the resulting updated mechanism $M^{\sigma_1:\alpha,\sigma_2:\gamma}$ is depicted on the right of Figure 3. In this case, a will invite her friends, including b, to join the auction of seller s_1. Observe that there is nothing seller s_2 can do to make her utility greater than 1 and thus outperform s_1 (recall that as a tie-breaking rule we use the lexicographic ordering, and hence, if buyer b is being invited to both auctions, she will choose seller s_1). In formulas, $\bigwedge_{i \in \mathsf{Nom}_M \cup \{\bullet\}} [\sigma_1 : \alpha, \sigma_2 : i](ut_{\sigma_1} > ut_{\sigma_2})$. The cooperative goal $\bigwedge_{i \in \mathsf{Nom}_M} i \to (\heartsuit \bullet \vee \Diamond\heartsuit\bullet)$ is no longer valid in the mechanism as agent d does not satisfy it.

Strategic Properties of Diffusion Auctions

Let us now demonstrate how to express Nash equilibrium in our setting, as well as study the complexities of the model checking and strategy existence problem.

We first define an appropriate notion of a *finite* DAM. The nuance here is that even if the set of agents is finite, we have a countably infinite number of names that can be assigned to buyers. Luckily for us, whenever a problem requires both a mechanism and a formula in the input, we need to care about only those nominals that occur explicitly in the formula.

Let $M = ((B \cup S, F, Bdg, V, I, N), P, Pay, U)$ be a DAM with a finite B, and let $name(\varphi)$ be the finite set of buyers' names appearing in φ. Also, let $buyers(\varphi) = \{b \in$

$B \mid \exists \beta \in name(\varphi)$ s.t. $N(\beta) = b\}$ be a subset of buyers that are named in φ. For each buyer $b \in B \setminus buyers(\varphi)$, i.e., not named in φ, we pick one arbitrary $\beta \in \mathsf{B}$ such that $N(\beta) = b$, and the finite set of such names is $unnamed(\varphi)$. We denote $X = name(\varphi) \cup unnamed(\varphi)$. We define function $N^{fin} = N_{\mathsf{S}} \cup N_{\mathsf{B}}|_X$, where $N_{\mathsf{B}}|_X : X \to B$ is a restriction of N_{B} to only those names that explicitly appear in φ and single names for buyers not named in φ. By construction, N^{fin} is finite. We denote M with N substituted with N^{fin} as M_{fin}. Intuitively, N^{fin} is a finite restriction of N.

It is straightforward to show the following by induction on φ and using the definition of the semantics.

Proposition 1. Let $\varphi \in \mathcal{L}^n$ and a mechanism M be given. Then we have that $M, a \models \varphi$ iff $M_{fin}, a \models \varphi$.

The *size* of a finite mechanism M is $|M| = |Agt| + |F| + |Bdg| + |V| + |I| + |N^{fin}| + |P| + |Pay| + |U|$. Since all mechanisms in this section are finite, we will use interchangeably M and M_{fin}, and N and N^{fin}.

Nash Equilibrium We can now verify that a given joint diffusion action is a one-step Nash equilibrium (NE) over a given finite mechanism M. This is significant, as it allows reasoning about optimal strategies of sellers. Let $\varphi := \langle \overline{\sigma} : \overline{\beta} \rangle \bigwedge_{1 \leqslant i \leqslant n} ut_{\sigma_i} = m_i$ express that after a joint diffusion action $\langle \overline{\sigma} : \overline{\beta} \rangle$, utilities of all sellers i are m_i.

In order to verify that such a diffusion is indeed an NE, we check that no single seller can increase her utility by deviating, i.e., by incentivising some other buyer present in the same mechanism. The formula for the NE in this case is $\varphi \wedge \bigwedge_{i=1}^{n} \bigwedge_{\gamma \in \mathsf{Nom}^{\bullet}_M} \langle \sigma_1 : \beta_1, ..., \sigma_i : \gamma, ..., \sigma_n : \beta_n \rangle (ut_{\sigma_i} \leqslant m_i)$, where $\mathsf{Nom}^{\bullet}_M$ is the set of names appearing in the finite mechanism plus $\bullet$.

We can further generalise the setting to a k-step NE, by taking $\varphi^k := \langle \overline{\sigma} : \overline{\alpha} \rangle^1 ... \langle \overline{\sigma} : \overline{\gamma} \rangle^k \bigwedge_{1 \leqslant i \leqslant n} ut_{\sigma_i} = m_i$ meaning that after a k-sequence of joint diffusion actions, sellers' utilities are m_i's. Then the corresponding formula for the k-step NE is $\varphi^k \wedge \bigwedge_{i=1}^{n} \bigwedge_{\gamma \in \mathsf{Nom}^{\bullet}_M} \langle \sigma_1 : \alpha_1, ..., \sigma_i : \gamma, ..., \alpha_n : \beta_n \rangle^1 .... \bigwedge_{i=1}^{n} \bigwedge_{\gamma \in \mathsf{Nom}^{\bullet}_M} \langle \sigma_1 : \beta_1, ..., \sigma_i : \gamma, ..., \sigma_n : \beta_n \rangle^k (ut_{\sigma_i} \leqslant m_i)$.

Model Checking and Strategy Existence

First, we show that the complexity of the model checking problem for $\mathcal{L}^n$ is in P.

Theorem 1. Model checking $\mathcal{L}^n$ is in P for the class of finite DAMs with polynomially computable placement, payment, and utility functions.

Proof. In Algorithm 1, we focus on the dynamic modality and the allocation operator[6].

Algorithm 1 An algorithm for model checking $\mathcal{L}^n$

1: **procedure** MC(M, a, φ)
2:   **case** $\varphi = [\overline{\sigma} : \overline{\beta}]\psi$
3:     **for** $\sigma_i \in \overline{\sigma}^{\setminus \bullet}$ **do**
4:       **if** $N(\beta_i) \in F(N(\sigma_i))$ and $Bdg(N(\sigma_i)) \geqslant I(N(\beta_i), N(\sigma_i))$ **then**
5:         **if** $N(\sigma_i) = \arg\max_{N(\sigma_i) \in S} I(N(\beta_i), N(\sigma_i)))$ **then**
6:           $F^{\overline{\sigma}:\overline{\beta}}(N(\sigma_i)) = F(N(\sigma_i)) \cup \{a \mid a \in B$ and $a \in F(N(\beta_i))\}$
7:           $Bdg^{\overline{\sigma}:\overline{\beta}}(N(\sigma_i)) = Bdg(N(\sigma_i)) - I(N(\beta_i), N(\sigma_i))$
8:           $Bdg^{\overline{\sigma}:\overline{\beta}}(N(\beta_i))) = Bdg(N(\beta_i))) + I(N(\beta_i), N(\sigma_i)))$
9:         **else**
10:           **return** $true$
11:     **return** MC$(M^{\overline{\sigma}:\overline{\beta}}, a, \psi)$
12:   **case** $\varphi = \heartsuit\gamma$
13:     **if** $\gamma \neq \bullet$ **then**
14:       **return** $P(N(\gamma)) = 1$
15:     **else**
16:       **return** $P(a) = 1$

On line 5, we check that seller $N(\sigma_i)$ offers the highest incentive to the corresponding buyer. If there are two sellers that offer the highest incentive to the same buyer, we assume that $\arg\max_{N(\sigma_i) \in S} I(N(\beta_i), N(\sigma_i)))$ returns the seller that appears earlier in the lexicographic order.

The algorithm directly mimics the semantics of $\mathcal{L}^n$ and thus correctness can be shown by induction on φ. First, recall that here we assume that the allocation, payment, and utility functions are all computable in polynomial time. For the case of $\varphi = [\overline{\sigma} : \overline{\beta}]\psi$, the size of an updated mechanism $M^{\overline{\sigma}:\overline{\beta}}$ is at most $\mathcal{O}(|M|^2)$ (in the worst case, the friendship relation is universal) and it can be computed in polynomial time. The procedure MC(M, a, φ) is run for at most $|\varphi|$ times and for at most $|\varphi|$ mechanisms. Hence, MC(M, a, φ) is used for a polynomial amount of time. ☐

Having the model checking result at hand, we can formulate and show the complexity of the strategy existence problem (proof is given in (Galimullin, Mittelmann, and Perrussel 2025)). The problem intuitively consists in checking whether, for a given mechanism and a mutual sellers' goal, there is a way (a strategy) for all sellers to achieve φ in a finite number of steps.

Definition 6. Given a finite mechanism M and a goal $\varphi \in \mathcal{L}^n$, the *strategy existence problem* consists in determining whether there is a finite sequence of concurrent incentivisations $\langle \overline{\sigma} : \overline{\beta} \rangle^* = \langle \overline{\sigma} : \overline{\alpha} \rangle .... \langle \overline{\sigma} : \overline{\gamma} \rangle$ such that $M, s \models \langle \overline{\sigma} : \overline{\beta} \rangle^* \varphi$ for sellers $s \in S$.

Theorem 2. The strategy existence problem is NP-complete for the class of finite DAMs with polynomially computable placement, payment, and utility functions.

Reasoning About Sellers' Strategies

While considering the strategy existence problem, we looked at how *all* sellers can reach their joint goal via a sequence of concurrent incentivisation actions. However, in diffusion auctions with multiple sellers selling the copy of

[6] Modal, nominal (Franceschet and de Rijke 2006), and arithmetic cases (Dasgupta, Papadimitriou, and Vazirani 2006) are standard and computed in polynomial time. We omit them for brevity.

the same item, sellers and *coalitions* thereof may compete against each other for buyers. To capture this strategic competitive setting, we introduce a modality inspired by coalitional operators from CL (Pauly 2002) and ATL (Alur, Henzinger, and Kupferman 2002). In particular, we extend the language of $\mathcal{L}^n$ with $\langle\!\langle \mathsf{C} \rangle\!\rangle\varphi$ for $\mathsf{C} \subseteq \mathsf{S}$, meaning that there is a (one-step) strategy for the coalition of sellers C to incentivise buyers such that no matter what other sellers do, φ holds. In other words, modalities $\langle\!\langle \mathsf{C} \rangle\!\rangle$ capture the ability of sellers in C to reach outcome φ in the competitive setting.

Definition 7. The language $\mathcal{SL}^n$ of the *n-seller strategic logic for diffusion incentives* is defined as follows:

$$\varphi := \quad \alpha \mid (z_1 t_1 + ... + z_m t_m) \geqslant z \mid \neg\varphi \mid (\varphi \wedge \varphi) \mid \Box\varphi \mid [\sigma_1 : \beta_1, ..., \sigma_n : \beta_n]\varphi \mid [\![\mathsf{C}]\!]\varphi \mid \heartsuit\gamma,$$

where $\mathsf{C} \subseteq \mathsf{S}$, and the dual of $[\![\mathsf{C}]\!]\varphi$ is $\langle\!\langle \mathsf{C} \rangle\!\rangle\varphi := \neg[\![\mathsf{C}]\!]\neg\varphi$. We will denote by $\overline{\sigma_\mathsf{C}} : \overline{\beta_\mathsf{C}}$ the assignment of buyers names $\overline{\beta_\mathsf{C}}$ from $\mathsf{B}^\bullet$ to the coalition of sellers $\overline{\sigma_\mathsf{C}}$ such that all sellers not in coalition are assigned $\bullet$ (i.e., they skip their turn). Moreover, we will denote $\overline{\sigma_\mathsf{C}} \cup \overline{\sigma_{\mathsf{S}\setminus\mathsf{C}}} : \overline{\beta_\mathsf{C}} \cup \overline{\beta_{\mathsf{S}\setminus\mathsf{C}}}$ the full assignment of n buyers from $\mathsf{B}^\bullet$ to n sellers.

Definition 8. Let $M = ((Agt, F, Bdg, V, I, N), P, Pay, U)$ be an n-DAM. The semantics of $\mathcal{SL}^n$ extends the semantics of $\mathcal{L}^n$ with the following clause and its dual:

$$M, a \models \langle\!\langle \mathsf{C} \rangle\!\rangle\varphi \text{ iff } \exists\overline{\beta_\mathsf{C}}\forall\overline{\beta_{\mathsf{S}\setminus\mathsf{C}}} : M, a \models \langle\overline{\sigma_\mathsf{C}} : \overline{\beta_\mathsf{C}}\rangle\top \text{ and } M, a \models [\overline{\sigma_\mathsf{C}} \cup \overline{\sigma_{\mathsf{S}\setminus\mathsf{C}}} : \overline{\beta_\mathsf{C}} \cup \overline{\beta_{\mathsf{S}\setminus\mathsf{C}}}]\varphi$$

$$M, a \models [\![\mathsf{C}]\!]\varphi \text{ iff } \forall\overline{\beta_\mathsf{C}}\exists\overline{\beta_{\mathsf{S}\setminus\mathsf{C}}} : M, a \models \langle\overline{\sigma_\mathsf{C}} : \overline{\beta_\mathsf{C}}\rangle\top \text{ implies } M, a \models \langle\overline{\sigma_\mathsf{C}} \cup \overline{\sigma_{\mathsf{S}\setminus\mathsf{C}}} : \overline{\beta_\mathsf{C}} \cup \overline{\beta_{\mathsf{S}\setminus\mathsf{C}}}\rangle\varphi$$

Intuitively, $\langle\!\langle \mathsf{C} \rangle\!\rangle\varphi$ holds if and only if there is a choice of buyers for the coalition of sellers C such that this choice is possible (part $\langle\overline{\sigma_\mathsf{C}} : \overline{\beta_\mathsf{C}}\rangle\top$) and whichever buyers the rest of the sellers decide to incentivise, φ holds after the resulting joint concurrent action.

It is immediate that the following formulas are valid[7]:

- $\langle\!\langle \mathsf{C} \rangle\!\rangle\varphi \to \langle\!\langle \mathsf{C} \cup \mathsf{D} \rangle\!\rangle\varphi$, i.e., a superset of a coalition is at least as powerful as the coalition.
- $[\![\emptyset]\!]\varphi \to \langle\!\langle \mathsf{S} \rangle\!\rangle\varphi$, i.e., the relationship between the empty and the grand coalitions.
- $\langle\!\langle \mathsf{C} \rangle\!\rangle(\varphi \wedge \psi) \to \langle\!\langle \mathsf{C} \rangle\!\rangle\varphi$, i.e., the ability to achieve two goals implies the ability to achieve any single one of them.

Example 4. The ability to reason about strategies of coalitions of sellers allows us to consider truly competitive and cooperative scenarios. For the first one, consider $\langle\!\langle \sigma_1, \sigma_2 \rangle\!\rangle[\![\sigma_3]\!](ut_{\sigma_1} > ut_{\sigma_3})$ meaning that a coalition of the first two sellers can preclude the third seller from having a utility equal to or higher than that of the first seller in a two-step incentivisation scenario. In a more altruistic setting, consider $\neg\langle\!\langle \sigma_1, \sigma_2 \rangle\!\rangle(ut_{\sigma_1} + ut_{\sigma_2} > 3) \wedge \langle\!\langle \sigma_1, \sigma_2, \sigma_3 \rangle\!\rangle(ut_{\sigma_1} + ut_{\sigma_2} > 3)$ meaning that together, the first and the second sellers cannot achieve a joint utility higher than 3, but if they cooperate with the third seller, this goal is satisfied.

[7] These formulas are some of the validites of CL (Pauly 2002).

Expressivity and Model Checking

Having introduced a strategic extension of $\mathcal{L}^n$, it is quite natural to wonder whether we gain anything in terms of expressivity, and, if yes, whether it comes at a price. We show that the answer to both questions is *yes* (with a caveat).

Theorem 3. Let M, a be a finite n-DAM and $\varphi \in \mathcal{SL}^n$. Then there exists a $\psi \in \mathcal{L}^n$ s.t. $M, a \models \varphi$ iff $M, a \models \psi$.

Proof. To prove the theorem, we present a truth-preserving translation $t : \mathcal{SL}^n \to \mathcal{L}^n$. All cases, apart from the strategic one, are trivial as $\mathcal{L}^n \subset \mathcal{SL}^n$. For the strategic case, we have $t(\langle\!\langle \mathsf{C} \rangle\!\rangle\varphi) = \bigvee_{\overline{\beta_\mathsf{C}} \in {\mathsf{Nom}^\bullet_M}^{|\mathsf{C}|}} \bigwedge_{\overline{\beta_{\mathsf{S}\setminus\mathsf{C}}} \in {\mathsf{Nom}^\bullet_M}^{|\mathsf{S}\setminus\mathsf{C}|}} (\langle\overline{\sigma_\mathsf{C}} : \overline{\beta_\mathsf{C}}\rangle\top \wedge [\overline{\sigma_\mathsf{C}} \cup \overline{\sigma_{\mathsf{S}\setminus\mathsf{C}}} : \overline{\beta_\mathsf{C}} \cup \overline{\beta_{\mathsf{S}\setminus\mathsf{C}}}]t(\varphi))$. Note that this is a well-formed formula because we are dealing with a finite mechanism and hence we can explicitly go over the elements in ${\mathsf{Nom}^\bullet_M}^{|\mathsf{C}|}$ one by one, where ${\mathsf{Nom}^\bullet_M}^{|\mathsf{C}|}$ is the set of all tuples of agent names in M of size $|\mathsf{C}|$. It follows from the definition of the semantics that the translation is truth-preserving and terminating. □

While Theorem 3 presents a translation that, for a given finite mechanism M, a and a formula $\varphi \in \mathcal{SL}^n$, produces a corresponding formula $\psi \in \mathcal{L}^n$ that agrees with φ on M, a, this result cannot be extended to arbitrary DAMs. In particular, the next theorem states that it is not the case that for a given $\varphi \in \mathcal{SL}^n$ we can always find a $\psi \in \mathcal{L}^n$ that will agree with φ on *all* mechanisms. In other words, $\mathcal{SL}^n$ is more expressive than $\mathcal{L}^n$. To show this, observe that modalities $\langle\!\langle \mathsf{C} \rangle\!\rangle\varphi$ quantify over *all* buyers' names, even those that are not explicitly present in the formula. A sketch of the proof of the next theorem is given in (Galimullin, Mittelmann, and Perrussel 2025).

Theorem 4. $\mathcal{SL}^n$ is strictly more expressive than $\mathcal{L}^n$, i.e. $\mathcal{L}^n \subset \mathcal{SL}^n$ and there is a $\varphi \in \mathcal{SL}^n$ s.t. for all $\psi \in \mathcal{L}^n$ there is a mechanism M, a such that $M, a \models \varphi$ iff $M, a \not\models \psi$.

As promised, the greater expressive power comes with a higher model checking complexity.

Theorem 5. Model checking $\mathcal{SL}^n$ is PSPACE-complete for the class of finite DAMs with the placement, payment, and utility functions computable in polynomial space.

Proof. To show that the problem is in PSPACE, we present Algorithm 2 that extends the P-time Algorithm 1.

Algorithm 2 An algorithm for model checking $\mathcal{SL}^n$

procedure MC(M, a, φ)
  case $\varphi = \langle\!\langle \mathsf{C} \rangle\!\rangle\psi$
    for $\overline{\beta_\mathsf{C}} = \beta_1, ..., \beta_{|\mathsf{C}|}$ s.t. $\beta_i \in \mathrm{dom}(N_\mathsf{B}|_X)$ **do**
      flag: = *true*
      if MC($M, a, \langle\overline{\sigma_\mathsf{C}} : \overline{\beta_\mathsf{C}}\rangle\top$) **then**
        for $\overline{\beta_{\mathsf{S}\setminus\mathsf{C}}} = \beta_1, ..., \beta_{|\mathsf{S}\setminus\mathsf{C}|}$ s.t. $\beta_i \in \mathrm{dom}(N_\mathsf{B}|_X)$ **do**
          if not MC($M, a, [\overline{\sigma_\mathsf{C}} \cup \overline{\sigma_{\mathsf{S}\setminus\mathsf{C}}} : \overline{\beta_\mathsf{C}} \cup \overline{\beta_{\mathsf{S}\setminus\mathsf{C}}}]\psi$) **then**
            flag: = *false*
            break
      else
        flag:= *false*
      if flag **then**
        return *true*

14: **return** *false*

The only new case is $\varphi = \langle\!\langle \mathsf{C} \rangle\!\rangle \psi$. The treatment of the case follows the semantics and thus the correctness follows. Indeed, to verify whether for a given agent we have $\langle\!\langle \mathsf{C} \rangle\!\rangle \psi$, we look for a set of buyers that sellers in C will incentivise (lines 3–5), and then we check whether for all possible incentive diffusions by the remaining sellers (line 6) we still can satisfy ψ after each seller performs their action (line 7). Variable 'flag' keeps track of whether we have found such a choice for the sellers in C, and the algorithm returns $true$ is yes, and $false$ if not.

Observe that the algorithm checks an exponential number of subsets of buyers, and hence the running time is exponential in the size of the input. However, the algorithm uses only a polynomial amount of space. Note that we explore the tree of mechanism updates in a depth-first manner. The space required by a branch in such a tree, and hence by the algorithm, is bounded by $\mathcal{O}(|\varphi| \cdot |M|^2)$. PSPACE-hardness is shown via a reduction from the QBF problem (see (Galimullin, Mittelmann, and Perrussel 2025)). □

Related Work

Logics for auctions Logic-based formalisms have been developed to capture and reason about various aspects of auctions. One direction is *bidding languages*, most notably OR and XOR-based languages, which express the preferences of auction participants (see (Nisan 2000) for an overview). These languages compactly represent possible bids on item combinations, an important aspect in combinatorial auctions, where bidders place bids on bundles of distinct items. Our work, instead, is closer to logical approaches for reasoning about and designing auctions. Mittelmann, Bouveret, and Perrussel (2022) proposed a lightweight formalism to represent auction rules, whereas Belardinelli et al. (2022) addressed the representation of strategies in repeated auctions, and Mittelmann, Herzig, and Perrussel (2021) captured bounded rationality in auctions. Another line of research proposes the use of variants of Strategy Logic (SL) for the design of auction mechanisms, exploiting verification and synthesis (Mittelmann et al. 2025, 2023). Notably, the model-checking complexity for specifications in SL is non-elementary in the general case (Mogavero et al. 2014). There is also research on automated verification of auction protocols (Garg et al. 2025; Caminati et al. 2015; Lange et al. 2013) that stresses the importance of formal verification techniques for auctions.

Across all these works, the set of agents involved in the auction is fixed throughout its execution. To the best of our knowledge, our work is the first one to explore the dynamics of auction diffusion from a logic-based perspective.

DEL and social network logics Our intuition of model updates stems from *dynamic epistemic logic* (DEL), where one can model various information-changing events in the context of agents' knowledge. Ideas of DEL were also adopted in the field of *social network logics* (SNLs), where one uses formal tools to study such phenomena on social networks as information diffusion (Christoff and Hansen 2015; Baltag et al. 2019), social influence (Christoff, Hansen, and Proietti 2016), and echo chambers (Pedersen, Smets, and Ågotnes 2019), to name a few. See (Pedersen 2024, Chapter 3) for an overview. Perhaps the most related work here is (Galimullin and Pedersen 2024), where the authors explore visibility of posts on social networks, and how these posts propagate through the network, somewhat akin to how the information about an auction is spread in diffusion auctions. Using nominals for agent names is common in SNLs and comes from hybrid logic (see, e.g., (Areces and ten Cate 2007)). While discussing strategic logic $\mathcal{SL}^n$, we noted that our coalitional operators are inspired by those of CL and ATL. However, in CL and ATL models are static, i.e., they do not change as a result of agents' actions. Hence, a more relevant work is that on *coalition announcements* (Ågotnes and van Ditmarsch 2008; Galimullin 2021; de Lima 2014) in DEL, where strategic operators quantify over model changes that agents can bring about in a competitive setting, and with model checking complexity being PSPACE-complete (Alechina et al. 2021). Another related work is (Maubert et al. 2020), where agents play a multi-step concurrent game by modifying a model using modalities of DEL. Finally, there has been some work on adding arrows in modal logics (Areces, Fervari, and Hoffmann 2015).

Conclusion

We have presented a formal framework for reasoning about sellers' strategies in diffusion auctions. In particular, we introduced two logics, the n-seller logic for diffusion incentives $\mathcal{L}^n$ and its strategic version $\mathcal{SL}^n$, that can capture various properties of such auctions, like item allocations, utility increase, local properties of the underlying social network, and Nash equilibrium, to name a few. Our logics are *dynamic*, and hence they also allow us to verify whether the above-mentioned properties hold after modifications of the underlying social network that are engendered by sellers incentivising buyers to invite their friends to join auctions. To the best of our knowledge, this is *the first work that tackles the problem of formal verification of diffusion auctions*.

Our definition of diffusion auction mechanisms is quite general and allows us to capture a variety of auction types as long as the complexity of computing the placement, payment, and utility functions is no higher than the complexity of the model-checking problem of the corresponding logic. We have shown that it is in P for $\mathcal{L}^n$ and is PSPACE-complete for $\mathcal{SL}^n$. Moreover, we have demonstrated that the complexity of the strategy existence problem for a given mechanism and a joint goal of sellers is NP-complete.

With our work, we start a research line on formal verification of diffusion auctions, and there are plenty of interesting further directions. In particular, we would like to tackle the formal verification of a probabilistic framework, capturing incomplete information and Bayesian analysis (Huang et al. 2025), as well as consider strategies of buyers. We have also mentioned some of the validities of our logics, and we find it very tempting to explore their axiomatisations. Finally, we plan to explore the case of multi-item diffusion auctions.

Acknowledgements

This research is partially supported by the ANR project NOGGINS ANR-24-CE23-4402.

References

Ågotnes, T.; and van Ditmarsch, H. 2008. Coalitions and announcements. In Padgham, L.; Parkes, D. C.; Müller, J. P.; and Parsons, S., eds., *Proceedings of the 7th AAMAS*, 673–680. IFAAMAS.

Alechina, N.; van Ditmarsch, H.; Galimullin, R.; and Wang, T. 2021. Verification and Strategy Synthesis for Coalition Announcement Logic. *Journal of Logic, Language and Information*, 30(4): 671–700.

Alur, R.; Henzinger, T. A.; and Kupferman, O. 2002. Alternating-time temporal logic. *Journal of the ACM*, 49: 672–713.

Areces, C.; Fervari, R.; and Hoffmann, G. 2015. Relation-changing modal operators. *Logic Journal of the IGPL*, 23(4): 601–627.

Areces, C.; and ten Cate, B. 2007. Hybrid logics. In Blackburn, P.; Van Benthem, J.; and Wolter, F., eds., *Handbook of Modal Logic*, volume 3 of *Studies in Logic and Practical Reasoning*, 821–868. Elsevier.

Baltag, A.; Christoff, Z.; Rendsvig, R. K.; and Smets, S. 2019. Dynamic Epistemic Logics of Diffusion and Prediction in Social Networks. *Studia Logica*, 107(3): 489–531.

Belardinelli, F.; Condurache, R.; Dima, C.; Jamroga, W.; and Knapik, M. 2021. Bisimulations for verifying strategic abilities with an application to the ThreeBallot voting protocol. *Information and Computation*, 276: 104552.

Belardinelli, F.; Jamroga, W.; Malvone, V.; Mittelmann, M.; Murano, A.; and Perrussel, L. 2022. Reasoning about Human-Friendly Strategies in Repeated Keyword Auctions. In Faliszewski, P.; Mascardi, V.; Pelachaud, C.; and Taylor, M. E., eds., *Proceedings of the 21st AAMAS*, 62–71.

Caminati, M. B.; Kerber, M.; Lange, C.; and Rowat, C. 2015. Sound Auction Specification and Implementation. In Roughgarden, T.; Feldman, M.; and Schwarz, M., eds., *Proceedings of the 16th EC*, 547–564. ACM.

Christoff, Z.; and Hansen, J. U. 2015. A logic for diffusion in social networks. *Journal of Applied Logic*, 13(1): 48–77.

Christoff, Z.; Hansen, J. U.; and Proietti, C. 2016. Reflecting on Social Influence in Networks. *Journal of Logic, Language and Information*, 25(3-4): 299–333.

Clarke, E. H. 1971. Multipart pricing of public goods. *Public choice*, 17–33.

Clarke, E. M.; Henzinger, T. A.; Veith, H.; and Bloem, R., eds. 2018. *Handbook of Model Checking*. Springer.

Dasgupta, S.; Papadimitriou, C.; and Vazirani, U. 2006. *Algorithms*. McGraw-Hill Higher Education.

de Lima, T. 2014. Alternating-time temporal dynamic epistemic logic. *Journal of Logic and Computation*, 24(6): 1145–1178.

Dobzinski, S.; and Vondrák, J. 2012. The computational complexity of truthfulness in combinatorial auctions. In Faltings, B.; Leyton-Brown, K.; and Ipeirotis, P., eds., *Proceedings of the 13th EC*, 405–422.

Dolgorukov, V.; Galimullin, R.; and Gladyshev, M. 2024. Dynamic Epistemic Logic of Resource Bounded Information Mining Agents. In Dastani, M.; Sichman, J. S.; Alechina, N.; and Dignum, V., eds., *Proceedings of the 23rd AAMAS*, 481–489. IFAAMAS / ACM.

Fagin, R.; Halpern, J. Y.; and Megiddo, N. 1990. A Logic for Reasoning about Probabilities. *Information and Computation*, 87(1/2): 78–128.

Franceschet, M.; and de Rijke, M. 2006. Model checking hybrid logics (with an application to semistructured data). *Journal of Applied Logic*, 4(3): 279–304.

Galimullin, R. 2021. Coalition and relativised group announcement logic. *Journal of Logic, Language and Information*, 30(3): 451–489.

Galimullin, R.; Mittelmann, M.; and Perrussel, L. 2025. Formal Verification of Diffusion Auctions. *CoRR*, abs/2511.08765.

Galimullin, R.; and Pedersen, M. Y. 2024. Visibility and exploitation in social networks. *Mathematical Structures in Computer Science*, 34(7): 615–644.

Garg, M.; Raja, N.; Sarswat, S.; and Singh, A. K. 2025. Double Auctions: Formalization and Automated Checkers. *Journal of Automated Reasoning.*, 69(3): 17.

Groves, T. 1973. Incentives in teams. *Econometrica: Journal of the Econometric Society*, 617–631.

Guo, Y.; and Hao, D. 2021. Emerging Methods of Auction Design in Social Networks. In Zhou, Z., ed., *Proceedings of the 30th IJCAI*, 4434–4441. ijcai.org.

Huang, Y.; Hao, D.; Fan, Z.; Guo, Y.; and Li, B. 2025. Approximate Revenue Maximization for Diffusion Auctions. *CoRR*, abs/2507.14470.

Jamroga, W.; Kurpiewski, D.; and Malvone, V. 2022. How to measure usable security: Natural strategies in voting protocols. *Journal of Computer Security*, 30(3): 381–409.

Lange, C.; Caminati, M. B.; Kerber, M.; Mossakowski, T.; Rowat, C.; Wenzel, M.; and Windsteiger, W. 2013. A Qualitative Comparison of the Suitability of Four Theorem Provers for Basic Auction Theory. In Carette, J.; Aspinall, D.; Lange, C.; Sojka, P.; and Windsteiger, W., eds., *Proceedings of the 6th CICM*, volume 7961 of *LNCS*, 200–215. Springer.

Li, B.; Hao, D.; Gao, H.; and Zhao, D. 2022. Diffusion auction design. *Artificial Intelligence*, 303: 103631.

Maubert, B.; Pinchinat, S.; Schwarzentruber, F.; and Stranieri, S. 2020. Concurrent Games in Dynamic Epistemic Logic. In Bessiere, C., ed., *Proceedings of the 29th IJCAI*, 1877–1883. ijcai.org.

McAfee, R. P. 1992. A dominant strategy double auction. *Journal of economic Theory*, 56(2): 434–450.

Mittelmann, M.; Bouveret, S.; and Perrussel, L. 2022. Representing and reasoning about auctions. *Autonomous Agents and Multi-Agent Systems*, 36(1): 20.

Mittelmann, M.; Herzig, A.; and Perrussel, L. 2021. Epistemic Reasoning About Rationality and Bids in Auctions. In Faber, W.; Friedrich, G.; Gebser, M.; and Morak, M., eds., *Proceedings of the 17th JELIA*, volume 12678 of *LNCS*, 116–130. Springer.

Mittelmann, M.; Maubert, B.; Murano, A.; and Perrussel, L. 2023. Formal Verification of Bayesian Mechanisms. In Williams, B.; Chen, Y.; and Neville, J., eds., *Proceedings of the 37th AAAI*, 11621–11629. AAAI Press.

Mittelmann, M.; Maubert, B.; Murano, A.; and Perrussel, L. 2025. Formal verification and synthesis of mechanisms for social choice. *Artificial Intelligence*, 339: 104272.

Mogavero, F.; Murano, A.; Perelli, G.; and Vardi, M. Y. 2014. Reasoning About Strategies: On the Model-Checking Problem. *ACM Transactions on Computational Logic*, 15(4): 34:1–34:47.

Nisan, N. 2000. Bidding and allocation in combinatorial auctions. In Jhingran, A.; MacKie-Mason, J.; and Tygar, D. J., eds., *Proceedings of the 2nd EC*, 1–12. ACM.

Nisan, N.; Roughgarden, T.; Tardos, É.; and Vazirani, V. V., eds. 2007. *Algorithmic Game Theory*. CUP.

Pauly, M. 2002. A modal logic for coalitional power in games. *Journal of Logic and Computation*, 12(1): 149–166.

Pedersen, M. Y. 2024. *Malicious Agents and the Power of the Few : On the Logic of Abnormality in Social Networks*. Ph.D. thesis, University of Bergen, Norway.

Pedersen, M. Y.; Smets, S.; and Ågotnes, T. 2019. Analyzing Echo Chambers: A Logic of Strong and Weak Ties. In Blackburn, P.; Lorini, E.; and Guo, M., eds., *Proceedings of the 7th LORI*, volume 11813 of *LNCS*, 183–198. Springer.

van Ditmarsch, H.; van der Hoek, W.; and Kooi, B. 2008. *Dynamic Epistemic Logic*, volume 337 of *Synthese Library*. Springer.

Vickrey, W. 1961. Counterspeculation, auctions, and competitive sealed tenders. *The Journal of Finance*, 16(1): 8–37.

Xu, J.; and He, X. 2020. Design of double auction mechanism based on social network. *IEEE Access*, 8: 8324–8335.

Zhang, Y.; Zheng, S.; and Zhao, D. 2024. Optimal Diffusion Auctions. In *Proceedings of the 27th ECAI*, volume 392 of *Frontiers in Artificial Intelligence and Applications*, 3501–3508. IOS Press.

Zhao, D. 2021. Mechanism Design Powered by Social Interactions. In Dignum, F.; Lomuscio, A.; Endriss, U.; and Nowé, A., eds., *Proceedings of the 20th AAMAS*, 63–67. ACM.

Zhao, D.; Li, B.; Xu, J.; Hao, D.; and Jennings, N. R. 2018. Selling Multiple Items via Social Networks. In André, E.; Koenig, S.; Dastani, M.; and Sukthankar, G., eds., *Proceedings of the 17th AAMAS*, 68–76. IFAAMAS.