Learning to Diffuse: Mechanism Design in Social Networks with Information Propagation Costs

author names withheld

Under Review for NExT-Game 2026

Abstract

Prior diffusion auctions assume that forwarding auction information is costless, making full propagation a dominant strategy. In networked markets and emerging multi-agent settings, however, diffusion may consume time, attention, bandwidth, or social capital, so agents must decide which neighbors are worth informing. We introduce the *Learning Diffusion Mechanism* (LDM), a repeated costly-diffusion extension of the Information Diffusion Mechanism (IDM). In each round, agents observe private valuations, choose whom to inform, and receive their IDM payoff minus a marginal cost per activated link. We show that truthful bidding remains weakly dominant under additive diffusion costs, which reduces the strategic problem to learning diffusion decisions. We then formulate the induced Bayesian diffusion game, establish existence of a stationary Bayes-Nash equilibrium, and derive a natural model-free policy-gradient update. LDM highlights how local propagation costs can shrink the realized auction network and provides a minimal framework for studying learned information diffusion in strategic networked markets.

1. Introduction

Diffusion auctions study selling when participation is itself strategic [3, 12, 13, 23]. A *seller* directly reaches only a small set of *buyers* in a social network; these buyers may invite their neighbors, who may in turn invite others. Because inviting new buyers creates additional competition, agents may withhold information unless the mechanism rewards diffusion. This setting captures networked markets such as referrals, recruitment, online marketplaces, and advertising. The central design goal is therefore twofold: elicit truthful values and incentivize agents to propagate the auction. The Information Diffusion Mechanism (IDM) of Li et al. [13], building on ideas from network referral mechanisms [3], achieves incentive compatibility and individual rationality in a costless diffusion model by rewarding critical diffusion nodes on the path to the winner. Subsequent work has studied revenue, privacy, and structural extensions of diffusion auctions [10, 12, 23].

We consider a more realistic setting in which diffusion is costly. Sharing information may consume time, bandwidth, attention, or social capital, so full diffusion is no longer automatically optimal. Instead, agents must decide which links are worth activating. We model this as a *repeated selling problem*: in each round t, agents choose whom to inform, pay a marginal cost per activated neighbor, and receive utility equal to their IDM reward minus diffusion cost. Thus, agents must learn diffusion policies that balance the expected value of reaching high-value downstream buyers against the cost of communication. This perspective turns diffusion auctions into a *learning problem over strategic networks*. Unlike standard learning-based auction design, which typically assumes a fixed set of bidders [2, 11, 17], the participant set here is endogenous and depends on agents' learned communication decisions. Beyond auctions, the same abstraction captures a basic challenge in emergent multi-agent learning systems: agents must learn when to communicate, reveal information, update shared state, or delegate under computational, privacy, or strategic costs. Examples include multi-agent software engineering, collaborative memory construction, autonomous-mobility coordination, and decentralized electricity-market agents [e.g., 5, 8, 21, 22]. To the best of our knowledge, existing work has not studied learning in network auctions with costly strategic diffusion, nor repeated settings in which agents learn whom to diffuse information to over time; refer to Appendix A.

Contributions. We introduce the *Learning Diffusion Mechanism* (LDM), extending the one-shot IDM to a repeated setting with costly diffusion. Our key contributions are: *(i)* we formalize a repeated selling environment where activating diffusion links incurs a marginal cost; *(ii)* we prove that truthful bidding remains a weakly dominant strategy despite additive costs, crucially decoupling the bidding strategy from the diffusion strategy; *(iii)* we characterize the induced Bayesian diffusion game and establish the existence of a stationary Bayes-Nash equilibrium in mixed diffusion strategies; and *(iv)* introduce a policy-gradient learning framework enabling agents to iteratively update their diffusion propensities. This relies solely on realized net utilities, functioning independently of complete network topology or the valuation distributions of other agents

2. Setting & Preliminaries

Consider a social network $\mathcal{G} = (\mathcal{N}, \mathcal{E})$, where $\mathcal{N} = \{1, \ldots, n\}$ is the set of buyer-agents and $\mathcal{E}$ encodes neighborhood relations. For each agent $i \in \mathcal{N}$, let $r_i = \{j \in \mathcal{N} \mid (i,j) \in \mathcal{E}\}$, denote their immediate neighbors. A seller s holds an item with a reserve price of zero. Each agent i has a private value $v_i \geq 0$ for the item. An agent can participate only after receiving the auction information from the seller or from another informed agent. Upon becoming informed, agent i chooses an action $a_i = (b_i, \mathbf{x}_i)$, where $b_i \in \mathbb{R}_+$ is the reported valuation (bid) and $\mathbf{x}_i = (x_{i,j})_{j \in r_i} \in \{0,1\}^{|r_i|}$ is the diffusion decision vector, with $x_{i,j} = 1$ indicating that i informs neighbor j. The joint diffusion profile $\mathbf{x} = (\mathbf{x}_1, \ldots, \mathbf{x}_n)$ induces a realized diffusion graph $\mathcal{G}_{\mathbf{x}} \subseteq \mathcal{G}$ and therefore a set of agents reachable from the seller s. As a baseline, we use the costless one-shot Information Diffusion Mechanism (IDM) of Li et al. [13]. Given bids and diffusion decisions $(\mathbf{b}, \mathbf{x})$, IDM allocates the item to a winner w and assigns each agent a gross payoff $u_i^{\mathrm{idm}}(\mathbf{b}, \mathbf{x})$. This payoff includes both the winner's allocation utility and any diffusion reward received by critical agents through whom information must pass to reach the winner. The formal definition of critical diffusion sequences, the IDM allocation rules, and the piecewise expression of $u_i^{idm}(\mathbf{b}, \mathbf{x})$ are deferred to Appendix B.

3. Learning Diffusion Mechanism (LDM): Repeated Game

We now extend the costless one-shot IDM baseline to a repeated selling environment with information propagation costs. Time is discrete, $t = 1, \ldots, T$. In each round, agents are given a new valuation draw and choose both a bid and a diffusion action. Unlike the costless IDM, we introduce an exogenous marginal propacation cost $c > 0$; each agent i pays cost c for every neighbor to whom it sends the auction information. Proofs are deferred to Appendix C.

At round t, agent i observes a private valuation $v_i(t) \geq 0$, independently drawn from an absolutely continuous distribution F_i, with bounded support $[0, v_{max}]$. After informed agents simultaneously submit actions $a_i(t) = (b_i(t), \mathbf{x}_i(t))$, the mechanism runs IDM on the realized diffusion graph $\mathcal{G}_{\mathbf{x}}$ and bids $\mathbf{b}(t)$. The realized net utility $U_i(t)$ of agent i is additively separable:

$$U_i(t) = u_i^{idm}(\mathbf{b}(t), \mathbf{x}(t)) - c \sum_{j \in r_i} x_{i,j}(t)\,.$$

3.1. Strategy Space Reduction and Expected Utility

The marginal cost $c > 0$ breaks the dominant strategy of full diffusion. However, the additive separability of $U_i(t)$ allows for an analytical separation between bidding and diffusion strategies.

Proposition 1 (Truthful Bidding Persistence) *Under IDM with propagation costs $c > 0$, truthful bidding, $(b_i(t) = v_i(t))$, is a weakly dominant strategy for any $\mathbf{x}_i(t)$.*

By proposition 1, equilibrium analysis can be restricted to the diffusion phase. Let $\mathbf{X}_i = \{0,1\}^{|r_i|}$ denote agent i's set of pure diffusion actions. Since the game features incomplete information, each agent's strategy maps their private (valuation) into a probability distribution over $\mathbf{X}_i$.

Definition 2 (Mixed Strategies and Expected Utility) *A behavioral strategy for agent i is a measurable function $\sigma_i \colon [0, v_{max}] \to \Sigma_i = \Delta(\mathbf{x}_i)$, where $\sigma_i(\mathbf{x}_i|v_i)$ denotes the probability that agent i chooses pure diffusion action $\mathbf{x}_i \in \mathbf{X}_i = \{0,1\}^{|r_i|}$ conditional on their realized valuation v_i. Given the joint strategy profile of all other agents σ_{-i}, the expected net utility for agent i of type v_i choosing pure diffusion action $\mathbf{x}_i$ is:*

$$\begin{aligned} J_i(\mathbf{x}_i|v_i, \sigma_{-i}) &= \mathbb{E}_{\mathbf{v}_{-i} \sim F_{-i}}\left[U_i(v_i, \mathbf{v}_{-i}, \mathbf{x}_i, \sigma_{-i}(\mathbf{v}_{-i}))\right] \\ &= \int_{\mathbf{v}_{-i}} \sum_{\mathbf{x}_{-i} \in \mathbf{X}_{-i}} \left[U_i(v_i, \mathbf{v}_{-i}, \mathbf{x}_i, \sigma_{-i}(\mathbf{v}_{-i}))\right] \left(\prod_{j \neq i} \sigma_j(\mathbf{x}_j|v_j)\right) dF_{-i}(v_{-i})\,. \end{aligned} \tag{J_i}$$

The expected utility (J_i) evaluates the payoff for agent i when pure diffusion action $\mathbf{x}_i$ is played. Because the game involves incomplete information, this expected value is computed by averaging the unknown private valuations ($\mathbf{v}_{-i}$) and the mixed strategies (σ_{-i}) of all other agents.

3.2. Equilibrium Analysis

To characterize the equilibrium, in the following, we define the best-response mapping.

Definition 3 (Best-Response Correspondence) *The best-response mapping $\mathcal{B}_i \colon \Sigma_{-i} \to \Sigma_i$ assign to each σ_{-i} the set of behavioral strategies that maximize i's utility for almost every $v_i \in [0, v_{max}]$*

$$\mathcal{B}_i(\sigma_{-i}) = \left\{\sigma_i \in \Sigma_i \mid supp(\sigma_i(\cdot|v_i)) \subseteq \arg\max_{\mathbf{x}_i \in \mathbf{X}_i} J_i(\mathbf{x}_i|v_i, \sigma_{-i}),\ \textit{for a.e. } v_i\right\} \tag{$\mathcal{B}_i$}$$

Having formally reduced the game to the diffusion decisions conditional on types (Proposition 1), we now establish the existence of a stationary equilibrium for the network.

Theorem 4 (Existence of Stationary Bayes-Nash Equilibrium) *For any finite social network $\mathcal{G} = (\mathcal{N}, \mathcal{E})$ and cost $c > 0$, if each valuation distributions F_i is absolutely continuous with bounded support $[0, v_{\max}]$, then the LDM diffusion game admits at least one stationary Bayes-Nash Equilibrium (BNE) in mixed diffusion strategies.*

4. Model-Free Learning of Diffusion Policies

The equilibrum characterization above is generally non constructive. Computing $\mathcal{B}_i(\sigma_{-i})$ in Eq. ($\mathcal{B}_i$) requires knowledge of the network, other agents' valuation distributions, and their diffusion strategies. We therefore study a model-free learning rule in which agents update diffusion propensities from realized payoffs. We frame the repeated diffusion game as a multi-agent reinforcement learning (MARL) with independent policy-gradient learners. Let the diffusion policy of agent i be parameterized by a vector θ_i, conditioned on the realized type $v_i(t)$. Let $\theta_{i,j}(v_i(t))$ represent the propensity of agent i to activate the edge to neighbor j given their current valuation. The probability of informing neighbor j is modeled via a sigmoid function:

$$\mathbb{P}_{i,j}(t) = \mathbb{P}_{i,j}(\theta_{i,j}(v_i(t))) = \frac{1}{1 + e^{-\theta_{i,j}(v_i(t))}} .$$

Assuming the activation of each outgoing edge is conditionally independent given θ_i, the joint probability of sampling the pure diffusion action $\mathbf{x}_i(t)$ is:

$$\pi_{\theta_i}(\mathbf{x}_i(t)|v_i(t)) = \prod_{j \in r_i} \left(\mathbb{P}_{i,j}(t)\right)^{x_{i,j}(t)} \left(1 - \mathbb{P}_{i,j}(t)\right)^{1-x_{i,j}(t)} .$$

To maximize the expected net utility $J_i(\theta_i)$, agents utilize the REINFORCE algorithm [18, 20]. The gradient of the objective function with respect to the policy parameters is:

$$\nabla_{\theta_i} J_i(\theta_i) = \mathbb{E}_{\pi_{\theta_i}}[\nabla_{\theta_i} \log \pi_{\theta_i}(\mathbf{x}_i(t)|v_i(t)) U_i(t)] .$$

The approximated gradient for a specific edge propensity $\theta_{i,j}$ evaluates to:

$$\nabla_{\theta_{i,j}} J_i(\theta_i) \approx U_i(t) \frac{\partial \log \pi_{\theta_i}(\mathbf{x}_i(t)|v_i(t))}{\partial \theta_{i,j}} .$$

Since $\frac{\partial \log \pi_{\theta_i}}{\partial \theta_{i,j}} = x_{i,j}(t) - \mathbb{P}_{i,j}(t)$, the stochastic gradient ascent with $\alpha \in (0,1)$ learning rate is:

$$\theta_{i,j}(t+1) = \theta_{i,j}(t) + \alpha U_i(t)(x_{i,j}(t) - \mathbb{P}_{i,j}(t)) .$$

5. Discussion

LDM highlights a basic tension in networked markets: diffusion can increase allocative efficiency and seller revenue, but the agents who bear the cost of propagation may not internalize these global benefits. When diffusion is costly, agents rationally filter their outgoing links, so the realized auction network can be substantially smaller than the underlying social network. The framework also suggests several natural extensions. Our policy conditions diffusion on the agent's valuation, but not on item features. Thus, if a typically low-value region of the network has high value for a particular item, the learner may undersample paths to it. Conditioning diffusion policies on item representations is a direct extension. Another direction is to characterize optimal seller revenue under costly learned diffusion, and to extend the model to multi-unit and combinatorial diffusion auctions [4, 14].

References

[1] Michael J. Curry, Uro Lyi, Tom Goldstein, and John P. Dickerson. Learning revenue-maximizing auctions with differentiable matching. In *Proceedings of The 25th International Conference on Artificial Intelligence and Statistics*, volume 151 of *Proceedings of Machine Learning Research*, pages 6062–6073. PMLR, 28–30 Mar 2022.

[2] Paul Duetting, Zhe Feng, Harikrishna Narasimhan, David Parkes, and Sai Srivatsa Ravindranath. Optimal auctions through deep learning. In *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pages 1706–1715. PMLR, 09–15 Jun 2019.

[3] Yuval Emek, Ron Karidi, Moshe Tennenholtz, and Aviv Zohar. Mechanisms for multi-level marketing. In *Proceedings of the 12th ACM Conference on Electronic Commerce*, EC '11, page 209–218. Association for Computing Machinery, 2011. ISBN 9781450302616. doi: 10.1145/1993574.1993606.

[4] Yuan Fang, Mengxiao Zhang, Jiamou Liu, Bakh Khoussainov, and Mingyu Xiao. Multi-unit auction over a social network. *arXiv:2302.08924*, 2023.

[5] Jakob Foerster, Richard Y. Chen, Maruan Al-Shedivat, Shimon Whiteson, Pieter Abbeel, and Igor Mordatch. Learning with opponent-learning awareness. In *Proceedings of the 17th International Conference on Autonomous Agents and MultiAgent Systems*, pages 122–130, 2018.

[6] Irving Leonard Glicksberg. *A Further Generalization of the Kakutani Fixed-Point Theorem, with Applications to Nash Equilibrium Points.* RAND Corporation, 1951.

[7] Wenshuo Guo, Michael Jordan, and Emmanouil Zampetakis. Robust learning of optimal auctions. In *Advances in Neural Information Processing Systems*, 2021.

[8] Sirui Hong, Mingchen Zhuge, Jonathan Chen, Xiawu Zheng, Yuheng Cheng, Jinlin Wang, Ceyao Zhang, Zili Wang, Steven Ka Shing Yau, Zijuan Lin, Liyang Zhou, Chenyu Ran, Lingfeng Xiao, Chenglin Wu, and Jürgen Schmidhuber. MetaGPT: Meta programming for a multi-agent collaborative framework. In *The Twelfth International Conference on Learning Representations*, 2024.

[9] Kexin Huang, Ziqian Chen, Xue Wang, Chongming Gao, Jinyang Gao, Bolin Ding, and Xiang Wang. Auctionformer: A unified deep learning algorithm for solving equilibrium strategies in auction games. In *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pages 19635–19659. PMLR, 21–27 Jul 2024.

[10] Yifan Huang, Dong Hao, Zhiyi Fan, Yuhang Guo, and Bin Li. Approximate revenue maximization for diffusion auctions. *arXiv:2507.14470*, 2025.

[11] Dmitry Ivanov, Iskander Safiulin, Igor Filippov, and Ksenia Balabaeva. Optimal-er auctions through attention. In *Advances in Neural Information Processing Systems*, 2022.

[12] Fengjuan Jia, Mengxiao Zhang, Jiamou Liu, and Bakh Khoussainov. Incentivising diffusion while preserving differential privacy. In *The 39th Conference on Uncertainty in Artificial Intelligence*, 2023.

[13] Bin Li, Dong Hao, Dengji Zhao, and Tao Zhou. Mechanism design in social networks. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 31, 2017.

[14] Xuanyu Li, Miao Li, Yuhan Cao, and Dengji Zhao. Combinatorial diffusion auction design. *arXiv:2410.22765*, 2024.

[15] Paul R Milgrom and Robert J Weber. Distributional strategies for games with incomplete information. *Mathematics of operations research*, 10(4):619–632, 1985.

[16] Neehar Peri, Michael Curry, Samuel Dooley, and John Dickerson. Preferencenet: Encoding human preferences in auction design with deep learning. In *Advances in Neural Information Processing Systems*, volume 34, pages 17532–17542. Curran Associates, Inc., 2021.

[17] Jad Rahme, Samy Jelassi, and S. Matthew Weinberg. Auction learning as a two-player game. In *International Conference on Learning Representations*, 2021.

[18] Richard S Sutton, David McAllester, Satinder Singh, and Yishay Mansour. Policy gradient methods for reinforcement learning with function approximation. In *Advances in Neural Information Processing Systems*, volume 12. MIT Press, 1999.

[19] Vasilis Syrgkanis. A sample complexity measure with applications to learning optimal auctions. In I. Guyon, U. Von Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017.

[20] Ronald J. Williams. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine Learning*, 8(3–4):229–256, 1992. ISSN 0885-6125. doi: 10.1007/BF00992696.

[21] Qingyun Wu, Gagan Bansal, Jieyu Zhang, Yiran Wu, Beibin Li, Erkang Zhu, Li Jiang, Xiaoyun Zhang, Shaokun Zhang, Jiale Liu, Ahmed Hassan Awadallah, Ryen W White, Doug Burger, and Chi Wang. Autogen: Enabling next-gen llm applications via multi-agent conversation. *arXiv:2308.08155*, 2023.

[22] Shunyu Yao, Jeffrey Zhao, Dian Yu, Nan Du, Izhak Shafran, Karthik Narasimhan, and Yuan Cao. React: Synergizing reasoning and acting in language models. In *International Conference on Learning Representations*, 2023.

[23] Yao Zhang, Shanshan Zheng, and Dengji Zhao. Optimal diffusion auctions. In *Proceedings of the 23rd International Conference on Autonomous Agents and Multiagent Systems*, AAMAS ’24, page 2600–2602. International Foundation for Autonomous Agents and Multiagent Systems, 2024. ISBN 9798400704864.

Appendix A. Related Work

Diffusion and network-aware mechanism design. Classical auction theory studies strategic allocation under incomplete information, typically taking the set of participating agents as fixed [15]. A complementary line of work considers mechanisms in which agents can affect participation itself by forwarding information through a social network. Early work on multi-level marketing formalized reward mechanisms for incentivizing propagation in networks [3]. Diffusion auctions bring this idea into auction design: the seminal mechanism-design model of social-network auctions asks buyers not only to report their values, but also to invite their neighbors, despite the fact that doing so may create additional competition [13]. Subsequent work studies revenue objectives and structural limitations in such settings, including optimal diffusion auctions [23], multi-unit extensions [4], combinatorial diffusion auctions [14], privacy-preserving diffusion incentives [12], and approximate revenue maximization for diffusion auctions [10].

Learning-based auction design. A separate line of work studies how to learn auctions from data, usually in settings where the bidder set is exogenously fixed and there is no strategic diffusion decision. From a statistical perspective, Syrgkanis [19] studies sample complexity for learning optimal auctions, while Guo et al. [7] considers robust learning of optimal auctions under corrupted samples. Neural approaches model allocation and payment rules as differentiable functions and optimize revenue subject to approximate incentive constraints: RegretNet [2] initiated this approach, and later work improved the architecture and regret–revenue tradeoff using attention [11]. Other extensions formulate auction learning as a two-player game with learned deviations [17], incorporate preference and fairness constraints [16], use differentiable matching for richer allocation constraints [1], and employ transformer architectures to solve equilibrium strategies across auction games [9].

Taken together, existing diffusion-auction work largely designs analytic mechanisms for networked participation, while learning-based auction design largely assumes a fixed participant set; to the best of our knowledge, prior work has not studied learning in network auctions where agents strategically decide whether, and to whom, to diffuse information over repeated selling interactions, potentially under diffusion costs.

Appendix B. Extended Preliminaries: Costless One-shot Information Diffusion Mechanism

This appendix makes explicit the costless one-shot IDM primitives used as the baseline in Section 2. In particular, we define participation, critical diffusion nodes, the IDM allocation rule, and the resulting gross IDM payoff before any diffusion costs are introduced.

Participation and realized diffusion graph. Given a diffusion profile $\mathbf{x}$, let $\mathcal{G}_{\mathbf{x}} \subseteq \mathcal{G}$ denote the directed graph containing exactly the active information links. An agent j participates in the auction if and only if there exists a directed path of active information links from the seller s to j in $\mathcal{G}_{\mathbf{x}}$.

We denote the set of participating agents by

$$\mathcal{R}(\mathbf{x}) := \{j \in \mathcal{N} : \text{there exists an active directed path from } s \text{ to } j \text{ in } \mathcal{G}_{\mathbf{x}}\}.$$

Agents outside $\mathcal{R}(\mathbf{x})$ are uninformed and therefore do not submit bids.

Critical diffusion nodes. For a participating agent $j \in \mathcal{R}(\mathbf{x})$, another agent i is a *critical diffusion node* (CDE) if *every* active path from the seller s to j passes through i. Let

$$d_i(\mathbf{x}) := \{j \in \mathcal{R}(\mathbf{x}) : i \text{ is a critical diffusion node for } j\}$$

denote the set of participating agents for whom i is a CDE. By convention, $i \in d_i(\mathbf{x})$ whenever i participates.

For any subset $\mathcal{S} \subseteq \mathcal{N}$, define

$$b^*_{-\mathcal{S}}(\mathbf{x}) := \max_{j \in \mathcal{R}(\mathbf{x}) \setminus \mathcal{S}} b_j \,,$$

with the convention that the maximum is zero when the set is empty. Let

$$m \in \arg\max_{j \in \mathcal{R}(\mathbf{x})} b_j$$

be the highest-bidding participating agent, with ties broken by a fixed deterministic rule.

Diffusion critical sequence. Let

$$\mathcal{K}_m(\mathbf{x}) = (\kappa_1, \kappa_2, \ldots, \kappa_L = m)$$

be the ordered sequence of critical diffusion nodes from the seller to the highest bidder m, ordered by their position along the diffusion from s to m. Hence, for any index $\ell < L$, node κ_ℓ is a necessary antecedent for $\kappa_{\ell+1}$; thus, every active path from the seller to the highest bidder must pass through the sequence $\mathcal{K}_m(\mathbf{x})$ in the prescribed order.

IDM allocation rule. The Information Diffusion Mechanism (IDM) allocates the item based on a sequential threshold condition that rewards critical nodes for their role in expanding the market. The winner w is identified as the first agent κ_{ℓ^*} in the critical sequence $\mathcal{K}_m(\mathbf{x})$ whose reported bid matches or exceeds the highest competing bid from agents outside the influence of the subsequent node in the sequence. Formally, the winner is

$$w = \kappa_{\ell^\star}, \qquad \ell^* = \min\left\{\{\ell \in \{1, \ldots, L-1\} \mid b_{\kappa_\ell} \geq b^*_{-d_{\kappa_{\ell+1}}(\mathbf{x})}\} \cup L\right\}. \tag{1}$$

The mechanism evaluates each critical node κ_ℓ (for $\ell < L$) against the best alternative bid available outside the set of agents who depend on the next critical node $\kappa_{\ell+1}$ to participate. If no such $\ell < L$ satisfies the condition, the item is naturally allocated to the highest bidder, $\kappa_L = m$, as denoted with the $\cup$ notation above.

Definition 5 (Gross IDM payoff) *Given bids and diffusion decisions* $(\mathbf{b}, \mathbf{x})$*, the gross payoff* $u_i^{idm}(\mathbf{b}, \mathbf{x})$*, for agent* i *with true valuation* v_i*, is defined as follows:*

- ***Winner.*** *If* i *is the winner* ($i = w$)*: The agent receives the item and pays the highest bid among agents, independent of* w*:*

$$u_i^{idm} = v_i - b^*_{-d_i(\mathbf{x})}.$$

- ***On path critical diffusion node.*** *If* i *is an on-path buyer* ($i \in \mathcal{K}_w(\mathbf{x}) \setminus \{w\}$)*, then* i *does not receive the item but is rewarded the marginal contribution of its diffusion to the attainable competing bid:*

$$u_i^{idm} = b^*_{-d_{i+1}(\mathbf{x})} - b^*_{-d_i(\mathbf{x})} .$$

- ***All other agents.*** *If agent* i *is neither the winner nor an on-path critical diffusion node, then:*

$$u_i^{idm} = 0 .$$

Appendix C. Technical Appendix

C.1. Proof of Proposition 1

Proof [Proof of Proposition 1] The agent's net utility $U_i(t)$ is additively separable into the gross payoff u_i^{idm} and the total propagation cost $c \sum_{j \in r_i} x_{i,j}$. The cost component is strictly a function of the diffusion decision $\mathbf{x}_i(t)$ and is entirely independent of the chosen bid $b_i(t)$. By the Incentive Compatibility (IC) of the baseline IDM, the gross payoff u_i^{idm} is globally maximized by setting $b_i(t) = v_i(t)$ for any arbitrary realized subgraph of participants. Since the bid does not influence the cost term, $b_i(t) = v_i(t)$ maximizes the total net utility $U_i(t)$. Therefore, truthful bidding strictly dominates other bidding strategies. ■

C.2. Proof of Theorem 4

Proof [Proof of Theorem 4] By Proposition 1, the stage game is analytically equivalent to a Bayesian game where each agent i observes a continuous private type $v_i \in [0, v_{\max}]$ and subsequently selects an action from a finite discrete set $\mathbf{x}_i = \{0, 1\}^{|r_i|}$.

Due to the absolute continuity of F, the probability of tied valuations is zero, ensuring that U_i is continuous almost everywhere. By applying the framework of distributional strategies ([15]), the strategy space is rendered weak-* compact and convex. Since J_i is continuous and linear in σ_i, Glicksberg's Fixed-Point Theorem [6], a generalized Kakutani Fixed-Point Theorem, guarantees the existence of a fixed point $\sigma^* \in \mathcal{B}(\sigma^*)$. This fixed point constitutes a stationary Bayes-Nash Equilibrium of the game. ■