# Reducing the Gap Between Theory and Applications in Algorithmic Bayesian Persuasion

**Matteo Castiglioni**

**Abstract** This work focuses on the following question: *is it possible to influence the behavior of self-interested agents through the strategic provision of information*? This 'sweet talk' is ubiquitous among all sorts of economics and non-economics activities. In this work, we model these multi-agent systems as games between an informed sender and one or multiple receivers. We study the computational problem faced by an informed sender that wants to use his information advantage to influence rational receivers with the partial disclosure of information. In particular, the sender faces an information structure design problem that amounts to deciding 'who gets to know what'. Bayesian persuasion provides a formal framework to model these settings as asymmetric-information games. In recent years, much attention has been given to Bayesian persuasion in the economics and artificial intelligence communities due also to the applicability of this framework to a large class of scenarios like online advertising, voting, traffic routing, recommendation systems, security, and product marketing. However, there is still a large gap between the theoretical study of information in games and its applications in real-world scenarios. This work contributes to close this gap along two directions. First, we study the persuasion problem in real-world scenarios, focusing on voting, routing, and auctions. While the Bayesian persuasion framework can be applied to all these settings, the algorithmic problem of designing optimal information disclosure polices introduces computational challenges related to the specific problem under study. Our goal is to settle the complexity of computing optimal sender's strategies, showing when an optimal strategy can be implemented efficiently. Then, we relax stringent assumptions that limit the applicability of the Bayesian persuasion framework in practice. In particular, the classical model assumes that the sender has perfect knowledge of the receiver's utility. We remove this assumption initiating the study of an online version of the persuasion problem. This is the first step in designing adaptive information disclosure policies that deal with the uncertainty intrinsic in all real-world applications.

M. Castiglioni (✉)
Politecnico di Milano, Piazza Leonardo da Vinci 32, Milan, Italy
e-mail: matteo.castiglioni@polimi.it



F. Amigoni (ed.), *Special Topics in Information Technology*,
PoliMI SpringerBriefs, https://doi.org/10.1007/978-3-031-51500-2_1

# 1 Introduction

This work considers the following question: *is it possible to influence the behavior of self-interested agents through the strategic provision of information*? This 'sweet talk' is ubiquitous among all sorts of economic activities, and it was famously attributed to 30% of the GDP in the United States [3]. Moreover, information is the foundation of any democratic election, as it allows voters for better choices. In many settings, uninformed voters have to rely on inquiries of third party entities to make their decision. With the advent of modern media environments, malicious actors have unprecedented opportunities to garble this information and influence the outcome of the election via misinformation and fake news [1]. Reaching voters with targeted messages has never been easier. As another example, consider a multi-agent routing problem in which agents seek to minimize their own costs selfishly. In real-world problems, the state of the network may be uncertain, and not known to its users (e.g., drivers may not be aware of road works and accidents in a road network). A central authority or a navigation app may mitigate inefficiencies and reduce the social cost providing players with partial information about the state of the network.

Bayesian persuasion [28] provides a framework to model the problem faced by an informed sender trying to influence the behavior of self-interested receivers. In particular, the sender faces an information structure design problem which amounts to deciding 'who gets to know what' about some exogenous parameters collectively termed state of nature. Since the seminal work of [28], a large attention has been given to the Bayesian persuasion framework in the economics and artificial intelligence community due also to the applicability of this framework to a large class of scenarios like online advertising [7, 8, 12, 27], voting [2, 25], traffic routing [11, 33], recommendation systems [29], security [31, 35], and product marketing [5, 13]. However, there is still a large gap between the theoretical study of information in games and its applications in real-world scenarios. This work contributes to close this gap along two directions. First, we study the Bayesian persuasion framework in real-world scenarios, focusing on voting, routing, and auctions. While the Bayesian persuasion framework can be applied to all these settings, the algorithmic problem of designing optimal information disclosure polices introduces computational challenges related to the specific problem under study. Then, we relax stringent assumptions that limit the applicability of the classical bayesian persuasion framework in practice. In particular, one of the most limiting assumption is, arguably, that the sender is required to know the receiver's utility function to compute an optimal signaling scheme. We remove this assumption by studying a repeated Bayesian persuasion problem in an online learning framework where, at each round, the receiver's type is adversarially chosen from a finite set of types. This is the first step in designing adaptive information disclosure policies that deals with the uncertainty intrinsic in all real-world applications.

## 2 The Bayesian Persuasion Framework

*Bayesian persuasion* [28] studies the problem faced by an informed agent (the *sender*) trying to influence the behavior of other self-interested agents (the *receivers*) via the partial disclosure of payoff-relevant information. Agents' payoffs are determined by the actions played by the receivers and by an exogenous parameter represented as a *state of nature*, which is drawn by a known prior probability distribution and observed by the sender only. The sender commits to a public randomized information-disclosure policy, which is customarily called *signaling scheme*. In particular, it defines how the sender should send signals to the receivers. Depending on the application various types of signaling schemes have been introduced to represent the possible communication constraints between the sender and the receivers. In a private signaling scheme, the sender can use a private communication channel per receiver, in a public signaling scheme the sender can use a single communication channel for all the receivers, while we introduce semi-public signaling schemes in which the sender can use a single communication channel for a subset of the receivers.

Arguably, one of the most severe obstacle to the application of the classical bayesian persuasion model by [28] to real-world scenarios is that the sender must know exactly the receiver's utility function to compute an optimal signaling scheme. This assumption is unreasonable in practice. However, only recently some works tries to relax this assumption. In particular, [6] study a game with a single receiver and binary-actions in which the sender does not know the receiver utility, focusing on the problem of designing a signaling scheme that perform well for each possible receiver's utility. Zu et al. [37] relax the perfect knowledge assumption assuming that the sender and the receiver do not know the prior distribution over the states of nature. They study the problem of computing a sequence of persuasive signaling schemes that achieve small regret with respect to the optimal signaling scheme with the knowledge of the prior distribution. Bernasconi et al. [10] extends the analysis to sequential settings. In this work, we follow a different approach and we deal with uncertainty about the receiver's utility by framing the Bayesian persuasion problem in an online learning framework [9]. In particular, we advance the state of the art on algorithmic Bayesian persuasion along two directions. First, we study Bayesian persuasion in games with structure, focusing on voting, routing, and auctions. Then, we initiate the study of Bayesian persuasion with payoff uncertainly.

## 3 Persuading in Election

In this section, we study Bayesian persuasion in voting scenarios. Information is the foundation of any democratic election, as it allows voters for better choices. In many settings, uninformed voters have to rely on inquiries of third-party entities to make their decision. For example, in most trials, jurors are not given the possibility of choosing which tests to perform during the investigation or which questions are asked to witnesses. They have to rely on the prosecutor's investigation and questions. The same happens in elections, in which voters gather information from third-party sources. Hence, we pose the question: *can a malicious actor influence the outcome of a voting process only by the provision of information to voters who update their beliefs rationally?* We study majority voting, plurality voting and district-based elections, showing a sharp contrast in term of efficiency in manipulating elections and computational tractability between the case in which private signals are allowed and the more restrictive setting in which only public signals are allowed. In particular, we show that it is possible to compute an optimal private signaling scheme in polynomial time in all the elections that we considered, while the problem of approximating the optimal public signaling scheme is NP-hard even for majority voting. Moreover, we show that, assuming the Exponential Time Hypothesis (ETH), the problem of approximating the optimal public signaling scheme in majority voting requires quasi-polynomial time even relaxing persuasiveness. In doing so, we provide some insights on the complexity of general persuasion problems, such as the characterization of bi-criteria approximations in public signaling problems. A complete version of our results appears in [14, 15, 19].

## 4 Persuading in Routing

The study of how to influence traffic congestion has receive an increasing attention in recent years [22, 30, 33, 34]. *Network congestion games*, where players seek to minimize their own costs selfishly, are a canonical example of a setting where externalities may induce socially inefficient outcomes [32]. In real-world problems, the state of the network may be uncertain, and not known to its users (e.g., drivers may not be aware of road works and accidents in a road network). This setting is modeled via *Bayesian network congestion games* (BNCGs). Here, we explore how information can be used to reduce the social cost in routing games. In particular, we study Bayesian games with atomic players, where network vagaries are modeled via a (random) state of nature which determines the costs incurred by the players. We investigate whether it is possible to efficiently compute optimal, i.e., minimizing the social cost, *ex ante* persuasive signaling schemes in BNCGs, showing that symmetry is a crucial property for its solution. We focus on the notion of *ex ante persuasiveness*, as introduced by [24, 36], where the receivers are incentivized to follow the sender's recommendations having observed only the signaling scheme. We show that an optimal *ex ante* persuasive signaling scheme can be computed in polynomial time in symmetric BNCGs (i.e., where all the players share the same source and destination pair) with edge costs defined as affine functions of the edge congestion. Then, we show that *symmetry* is a crucial property for efficient signaling by proving that it is NP-hard to compute an optimal *ex ante* persuasive signaling scheme in *asymmetric* BNCGs. Our reduction proves an even stronger hardness result, as it works for non-Bayesian singleton congestion games with affine costs, which is arguably the simplest class of asymmetric congestion games. Furthermore, in such setting, a solution to our problem is an optimal coarse correlated equilibrium and, thus, computing optimal coarse correlated equilibria is NP-hard. A complete version of our results appears in [17].

## 5 Persuading in Auctions

In this section, we study persuasion in posted price auctions. In these auctions a seller tries to sell an item by proposing take-it-or-leave-it prices to buyers arriving sequentially. Each buyer has to choose between declining the offer—without having the possibility of coming back—or accepting it, thus ending the auction. We study Bayesian posted price auctions, where the buyers valuations for the item depend on a random state of nature, which is known to the seller only. Thus, the seller does not only have to decide price proposals for the buyers, but also how to partially disclose information about the state so as to maximize revenue. Our model finds application in several real-world scenarios. For instance, in an e-commerce platform, the state of nature may reflect the condition (or quality) of the item being sold and/or some of its features. These are known to the seller only since the buyers cannot see the item given that the auction is carried out on the web. We focus on two different settings: *public signaling*, where the signals are publicly visible to all buyers, and *private signaling*, in which the seller can send a different signal to each buyer through private communication channels. As a first negative result, we prove that, in both public and private signaling, the problem of computing an optimal seller's strategy does not admit an FPTAS unless P = NP. Indeed, the result holds for basic instances with a single buyer. Then, we provide tight positive results by designing a PTAS for each setting. To do so, we provide a preliminary result that allows us to assume without loss of generality that the seller commits to price functions with specific structures. Indeed, in a Bayesian posted price auction, the seller may commit to a price function that selects the prices to be proposed to the buyers *stochastically* on the basis of the signals being sent to *all* the buyers. This introduces considerable additional challenges compared to standard posted price auctions. In order to overcome such difficulties, we show that the seller can commit to a price function that *deterministically* proposes a price to each buyer on the basis of the signal being sent to *that* buyer only, without incurring in any revenue loss. This holds in both the public and the private signaling settings. Finally, we conclude the analysis comparing the effectiveness of different classes of signaling schemes. We show that the seller can increase their revenue by revealing information on the state of nature through signaling, with respect to the case in which they do not disclose anything. Moreover, we shows that the seller may get an higher revenue by using private signaling rather than public signaling. A complete version of our results appears in [23].

## 6 Online Bayesian Persuasion

In this section, we study Bayesian persuasion with payoff uncertainty. First, we consider the setting with a single receiver and we deal with uncertainty about the receiver's type by framing the Bayesian persuasion problem in an online learning framework. In particular, we study a repeated Bayesian persuasion problem where, at each round, the receiver's type is adversarially chosen from a finite set of types. Our goal is the design of an online algorithm that recommends a signaling scheme at each round, guaranteeing an expected utility for the sender close to that of the best-in-hindsight signaling scheme. We study this problem under two models of feedback: in the full information model, the sender selects a signaling scheme and later observes the type of the receiver; in the partial information model, the sender only observes the actions taken by the receiver. First, we study the computational complexity of the online Bayesian persuasion problem. We provide a negative result that rules out, even in the full information setting, the possibility of designing a no-regret algorithm with polynomial per-round running time. The same hardness result holds when employing the notion of no-$\alpha$-regret (in the additive sense) for any $\alpha < 1$. Formally, we show that for any $\alpha \leq 1$, a no-$\alpha$-regret algorithm for the online Bayesian persuasion problem requiring a per-round running time polynomial in the size of the instance cannot exist, unless $\mathsf{NP} \subseteq \mathsf{RP}$. In order to prove this negative result we show, as an intermediate step, that the problem of approximating an optimal signaling scheme is $\mathsf{NP}$-Hard even in the offline Bayesian persuasion problem in which the sender knows the probability distribution according to which receiver's types are selected.

Then, we study whether it is possible to devise a no-regret algorithm for the online Bayesian persuasion problem by relaxing the (per-round) running time constraint. This is not a trivial problem even in the full information feedback setting since, at each round, the sender has to choose a signaling scheme among an infinite number of alternatives. Moreover, the sender's utility depends on the receiver's best response, which yields an objective function which is not linear nor convex (or even continuous in the space of the signaling schemes). In the full information feedback setting, we show how to construct an algorithm that guarantees a regret polynomial in the size of the problem instance, and sublinear in the number of rounds $T$ with order $O(T^{1/2})$. In the partial information feedback setting, we develop an algorithm guaranteeing a regret polynomial in the size of the problem instance, and sublinear in $T$ with order $O(T^{4/5})$. In this case, the main idea is to use a full-information no-regret algorithm in combination with a mechanism to estimate the sender's utilities corresponding to signaling schemes different from the one recommended by the algorithm. Finally, we show that, relaxing the persuasiveness constraints, we can design polynomial-time algorithms with small regret.

Finally, we extend the online Bayesian persuasion framework to include multiple receivers. We focus on the case with no-externalities and binary actions. Moreover, to focus only on the receivers' coordination problem, we overcome the intractability of the single-receiver problem assuming that each receiver has a constant number of types. First, we prove a negative result: for any $0 < \alpha \leq 1$, there is no polynomial-time no-$\alpha$-regret algorithm when the sender's utility function is supermodular or anonymous. Then, we focus on the case of submodular sender's utility functions and we show that, in this case, it is possible to design a polynomial-time no-$(1 - 1/e)$-regret algorithm, which is tight. A complete version of our results appears in [16, 18, 20].

## 7 Efficient Online Learning Through Mechanism Design

In the previous section, we show that, both for the setting with a single and multiple receivers, the design of polynomial-time no-regret algorithms is impossible due to the NP-Hardness of the underline offline problems in which the distribution over the types is known. Hence, the design of efficient algorithms for the offline problem is the bottleneck to the design of efficient online learning algorithms. In this section, we show how to circumvent this issue by leveraging ideas from mechanism design. In particular, we introduce a type reporting step in which the receiver is asked to report her type to the sender, after the latter has committed to a menu defining a signaling scheme for each possible receiver's type. Surprisingly, we prove that, with a single receiver, the addition of this type reporting stage makes the sender's computational problem tractable. Our main result is to show the existence of a menu of *direct* and *persuasive* signaling schemes. In the classical model in which the sender perfectly knows the receiver payoff, a signaling scheme is direct if signals represent action recommendations and persuasive if the receiver is incentivized to follow the recommendations. We extend this definition to menus of signaling schemes. In particular, a menu is direct if the signals used by all the signaling schemes are action recommendations, and it is persuasive if a receiver has an incentive to follow the action recommendation if they reported their true type. Using this result, an optimal menu of signaling schemes can be computed efficiently by a linear program of polynomial size.

Then, we extend our Bayesian persuasion framework with type reporting to settings with multiple receivers, focusing on the widely-studied case of no-externalities and binary actions. Moreover, we focus on most common classes of sender's utility functions: supermodular, submodular and anonymous [4, 5, 26, 36]. In such setting, we show that it is possible to find a sender-optimal solution in polynomial-time for supermodular and anonymous sender's utility functions. As for the case of submodular sender's utility functions, we provide a $(1 - 1/e)$-approximation to the problem, which is tight. A complete version of our results appears in [21].

## 8 Conclusions and Future Research

In this work, we significantly advance the state of the art on algorithmic Bayesian persuasion along two different directions. First, we study the algorithmic problem of designing optimal information disclosure policies in real-world scenarios. In particular, we study several voting problems, including majority voting, plurality voting and district-based elections characterizing the computational complexity of each problem under private and public signaling. In doing so, we provide some insights on the complexity of general persuasion problems, such as the characterization of bi-criteria approximations in public signaling problems. Moreover, we show how the partial disclosure of information can be used to reduce the social cost in routing games and to increase the revenue in posted price auctions. Then, we relax the assumptions that the sender knows the receiver's utility function, initiating the study of online Bayesian persuasion. This is the first step in designing adaptive information disclosure policies that deals with the uncertainty intrinsic in all real-world applications.

We conclude proposing some future research directions. Despite the great attention received by the economics and artificial intelligence communities and the large class of potential real-world applications, the use of Bayesian persuasion in the real world is still limited. We believe that one of the main obstacle to the design of information disclosure policies in practice is the perfect knowledge assumption. An interesting direction is to study how the general online Bayesian persuasion framework introduced in this work can be applied to structured games. This posses various challenges. First, despite the design of no-regret algorithms is computationally intractable in general, it would be interesting to find some structured games for which it is possible to design *efficient* no-regret algorithms. As a second point, while for the single-receiver online Bayesian persuasion problem we provide no-regret algorithms with both full information and partial information feedback, our analysis of settings with multiple-receiver is limited to the case with full feedback and no externalities. While this assumptions are reasonable in some settings, they do not fit with some applications. For instance, routing games requires to take in account externalities among the players. Another interesting direction is to deal with the computational challenges introduced by the online learning framework. In particular, we showed that the computation of no-regret algorithms in the online Bayesian persuasion problem is often computational intractable, making it difficult to apply in practice. We propose a way to solve this problem, showing that the intractability of an offline version of the problem can be circumvented with a type reporting step. It remains an open question if a type reporting step can be used to design *efficient* online learning algorithms. Moreover, in our online learning framework we assume that the receivers have a *finite* number of known possible types. Despite this is a significant improvement over the perfect knowledge of the receivers' utilities, this approach assumes some prior knowledge of the receivers. It would be interesting to extend our results to the case in which the receivers can have arbitrary utilities and hence an *infinite* number of possible types. Finally, we show how to deal with uncertainly over the receivers' utility functions. However, this is not the only unreasonable assumption of the classical Bayesian persuasion framework. For instance, another important assumption is that the sender and receivers share the same prior belief. In practice, these beliefs come from past observations, and thus are uncertain and approximated. References [10, 37] study a game between a sender and a receiver that do not know the prior distribution. It would be interesting to consider uncertainly on the receiver's payoffs and the prior belief simultaneously.

## References

1. Allcott H, Gentzkow M (2017) Social media and fake news in the 2016 election. J Econ Perspect 31(2):211–236
2. Alonso R, Câmara O (2016) Persuading voters. Am Econ Rev 106(11):3590–3605
3. Antioch G, et al (2013) Persuasion is now 30 per cent of US GDP: revisiting McCloskey and Klamer after a quarter of a century. Econ Round-Up 1:1
4. Arieli I, Babichenko Y (2019) Private Bayesian persuasion. J Econ Theory 182:185–217
5. Babichenko Y, Barman S (2017) Algorithmic aspects of private Bayesian persuasion. In: Innovations in theoretical computer science conference
6. Babichenko Y, Talgam-Cohen I, Xu H, Zabarnyi K (2021) Regret-minimizing Bayesian persuasion. In: EC '21: the 22nd ACM conference on economics and computation, Budapest, Hungary, July 18–23, 2021. ACM, p 128. https://doi.org/10.1145/3465456.3467574
7. Bacchiocchi F, Castiglioni M, Marchesi A, Romano G, Gatti N (2022) Public signaling in Bayesian ad auctions. In: Proceedings of the thirty-first international joint conference on artificial intelligence, IJCAI 2022, Vienna, Austria, 23–29 July 2022. pp 39–45. ijcai.org. https://doi.org/10.24963/ijcai.2022/6
8. Badanidiyuru A, Bhawalkar K, Xu H (2018) Targeting and signaling in ad auctions. In: Proceedings of the twenty-ninth annual ACM-SIAM symposium on discrete algorithms, pp 2545–2563
9. Bernasconi M, Castiglioni M, Celli A, Marchesi A, Gatti N, Trovò F (2023) Optimal rates and efficient algorithms for online Bayesian persuasion. In: Proceedings of the 40th international conference on machine learning
10. Bernasconi M, Castiglioni M, Marchesi A, Gatti N, Trovò F (2022) Sequential information design: learning to persuade in the dark. In: NeurIPS
11. Bhaskar U, Cheng Y, Ko YK, Swamy C (2016) Hardness results for signaling in Bayesian zero-sum and network routing games. In: Proceedings of the 2016 ACM conference on economics and computation, pp 479–496
12. Bro Miltersen P, Sheffet O (2012) Send mixed signals: earn more, work less. In: Proceedings of the 13th ACM conference on electronic commerce, pp 234–247
13. Candogan O (2019) Persuasion in networks: public signals and k-cores. In: Proceedings of the 2019 ACM conference on economics and computation, pp 133–134
14. Castiglioni M, Celli A, Gatti N (2020) Persuading voters: it's easy to whisper, it's hard to speak loud. In: The thirty-fourth AAAI conference on artificial intelligence, pp 1870–1877
15. Castiglioni M, Celli A, Gatti N (2023) Public Bayesian persuasion: being almost optimal and almost persuasive. Algorithmica 1–37
16. Castiglioni M, Celli A, Marchesi A, Gatti N (2020) Online Bayesian persuasion. In: Advances in neural information processing systems, vol 33, pp 16188–16198
17. Castiglioni M, Celli A, Marchesi A, Gatti N (2021) Signaling in Bayesian network congestion games: the subtle power of symmetry. In: The thirty-fifth AAAI conference on artificial intelligence
18. Castiglioni M, Celli A, Marchesi A, Gatti N (2023) Regret minimization in online Bayesian persuasion: handling adversarial receiver's types under full and partial feedback models. Artif Intell 314:103821. https://doi.org/10.1016/j.artint.2022.103821

19. Castiglioni M, Gatti N (2021) Persuading voters in district-based elections. In: The AAAI conference on artificial intelligence, AAAI. AAAI Press
20. Castiglioni M, Marchesi A, Celli A, Gatti N (2021) Multi-receiver online Bayesian persuasion. In: Meila M, Zhang T (eds.) Proceedings of the 38th international conference on machine learning. Proceedings of machine learning research, vol 139, pp 1314–1323. PMLR. Accessed 18–24 Jul 2021
21. Castiglioni M, Marchesi A, Gatti N (2022) Bayesian persuasion meets mechanism design: going beyond intractability with type reporting. In: The 21st international conference on autonomous agents and multiagent systems (2022)
22. Castiglioni M, Marchesi A, Gatti N, Coniglio S (2019) Leadership in singleton congestion games: what is hard and what is easy. Artif Intell 277 (2019). https://doi.org/10.1016/j.artint.2019.103177
23. Castiglioni M, Romano G, Marchesi A, Gatti N (2022) Signaling in posted price auctions. In: The thirty-sixth AAAI conference on artificial intelligence
24. Celli A, Coniglio S, Gatti N (2020) Private Bayesian persuasion with sequential games. In: The thirty-fourth AAAI conference on artificial intelligence, pp 1886–1893
25. Cheng Y, Cheung HY, Dughmi S, Emamjomeh-Zadeh E, Han L, Teng SH (2015) Mixture selection, mechanism design, and signaling. In: 56th annual symposium on foundations of computer science, pp 1426–1445
26. Dughmi S, Xu H (2017) Algorithmic persuasion with no externalities. In: ACM EC, pp 351–368
27. Emek Y, Feldman M, Gamzu I, PaesLeme R, Tennenholtz M (2014) Signaling schemes for revenue maximization. ACM Trans Econ Comput 2(2):1–19
28. Kamenica E, Gentzkow M (2011) Bayesian persuasion. Am Econ Rev 101(6):2590–2615
29. Mansour Y, Slivkins A, Syrgkanis V, Wu Z (2016) Bayesian exploration: incentivizing exploration in Bayesian games. In: ACM EC, p 661
30. Massicot O, Langbort C (2019) Public signals and persuasion for road network congestion games under vagaries. IFAC-PapersOnLine 51(34):124–130
31. Rabinovich Z, Jiang AX, Jain M, Xu H (2015) Information disclosure as a means to security. In: Proceedings of the 2015 international conference on autonomous agents and multiagent systems, pp 645–653
32. Roughgarden T (2005) Selfish routing and the price of anarchy, vol 174. MIT Press Cambridge
33. Vasserman S, Feldman M, Hassidim A (2015) Implementing the wisdom of Waze. In: Twenty-fourth international joint conference on artificial intelligence, pp 660–666
34. Wu M, Amin S, Ozdaglar AE (2018) Value of information systems in routing games. arXiv:1808.10590
35. Xu H, Freeman R, Conitzer V, Dughmi S, Tambe M (2016) Signaling in Bayesian stackelberg games. In: AAMAS, pp 150–158
36. Xu H (2020) On the tractability of public persuasion with no externalities. In: Proceedings of the fourteenth annual ACM-SIAM symposium on discrete algorithms. SIAM, pp 2708–2727
37. Zu Y, Iyer K, Xu H (2021) Learning to persuade on the fly: robustness against ignorance. In: EC '21: the 22nd ACM conference on economics and computation, Budapest, Hungary, July 18–23, 2021. ACM, pp 927–928. https://doi.org/10.1145/3465456.3467593