

Distributed Signaling Games

MORAN FELDMAN, Microsoft Research, Israel, and University of Haifa

MOSHE TENNENHOLTZ, Microsoft Research, Israel, and Technion-IIT

OMRI WEINSTEIN, Microsoft Research, Israel, and Columbia University

A recurring theme in recent computer science literature is that proper design of signaling schemes is a crucial aspect of effective mechanisms aiming to optimize social welfare or revenue. One of the research endeavors of this line of work is understanding the algorithmic and computational complexity of designing efficient signaling schemes. In reality, however, information is typically not held by a central authority, but is distributed among multiple sources (third-party “mediators”), a fact that dramatically changes the strategic and combinatorial nature of the signaling problem.

In this paper we introduce *distributed signaling games*, while using display advertising as a canonical example for introducing this foundational framework. A distributed signaling game may be a pure coordination game (i.e., a distributed optimization task), or a non-cooperative game. In the context of pure coordination games, we show a wide gap between the computational complexity of the centralized and distributed signaling problems, proving that distributed coordination on revenue-optimal signaling is a much harder problem than its “centralized” counterpart.

In the context of non-cooperative games, the outcome generated by the mediators’ signals may have different value to each. The reason for that is typically the desire of the auctioneer to align the incentives of the mediators with his own by a compensation relative to the marginal benefit from their signals. We design a mechanism for this problem via a novel application of Shapley’s value, and show that it possesses some interesting properties; in particular, it always admits a pure Nash equilibrium, and it never decreases the revenue of the auctioneer (relative to his a priori revenue when there are no mediators).

CCS Concepts: • **Theory of computation** → *Algorithmic game theory and mechanism design; Solution concepts in game theory; Distributed computing methodologies*; • **Information Systems** → *World Wide Web; Online advertising and Applied computing; Electronic Commerce; Online Auctions*.

Additional Key Words and Phrases: Signaling; Distributed mechanism design; Shapley value

ACM Reference Format:

Moran Feldman, Moshe Tennenholtz, and Omri Weinstein. 2019. Distributed Signaling Games. *ACM Transactions on Economics and Computation* 1, 1 (December 2019), 26 pages. <https://doi.org/10.1145/nnnnnnn.nnnnnnn>

1 INTRODUCTION

The topic of signaling has recently received much attention in the computer science literature on mechanism design [3, 5–8, 12]. A recurring theme of this literature is that proper design of a signaling scheme is crucial for obtaining efficient outcomes, such as social welfare maximization or revenue maximization. In reality, however, sources of information are typically not held by a central authority, but are rather distributed among third party mediators/information providers, a fact which dramatically changes the setup to be studied, making it a game between information providers rather than a more classic mechanism design problem. Such a game is in the spirit of work on the theory of teams in economics [15], whose computational complexity remained largely unexplored. The goal of this paper is to initiate an algorithmic study of such games, which we term *distributed signaling games*, via what we view as a canonical example: Bayesian auctions; and

Authors’ addresses: Moran Feldman, Microsoft Research, Israel, and University of Haifa, moranfe@cs.haifa.ac.il; Moshe Tennenholtz, Microsoft Research, Israel, and Technion-IIT, moshet@ie.technion.ac.il; Omri Weinstein, Microsoft Research, Israel, and Columbia University, omri@cs.columbia.edu.

2019. XXXX-XXXX/2019/12-ART \$15.00
<https://doi.org/10.1145/nnnnnnn.nnnnnnn>

50 more specifically, display advertising in the presence of third party external mediators (information
51 providers), capturing the (multi-billion) ad exchanges industry.

52 Consider a web-site owner that auctions each user's visit to its site, a.k.a. impression. The
53 impression types are assumed to arrive from a commonly known distribution. The bidders are
54 advertisers who know that distribution, but only the web site owner knows the impression type
55 instantiation, consisting of identifiers such as age, origin, gender and salary of the web-site visitor.
56 As is the practice in existing ad exchanges, we assume the auction is a second price auction. The
57 web-site owner decides on the information (i.e., signal) about the instantiation to be provided to
58 the bidders, which then bid their expected valuations for the impression given the information
59 provided. The selection of the proper signaling by the web-site is a central mechanism design
60 problem. Assume, for example, an impression associated with two attributes: whether the user
61 is male or female on one side, and whether he is located in the US or out of the US on the other
62 side. This gives 4 types of possible users. Assume for simplicity that the probability of arrival of
63 each user type is $1/4$, and that there are four advertisers each one of them has value of \$100 for a
64 distinguished user type and \$0 for the other types, where these values are common-knowledge.
65 One can verify that an auctioneer who reveals no information receives an expected payoff of \$25,
66 an auctioneer who reveals all information gets no payoff, while partitioning the impression types
67 into two pairs, revealing only the pair of the impression which was materialized (rather than the
68 exact instantiation) will yield a payoff of \$50, which is much higher revenue.

69 While the above example illustrates some of the potential benefits of signaling and its natural
70 fit to mechanism design, its major drawback is in the unrealistic manner in which information
71 is manipulated: while some information about the auctioned item is typically published by the
72 ad network [20] (such information is modeled here as a public prior), and despite the advertisers'
73 effort to perform "behavioral targeting" by clever data analysis (e.g., utilizing the browsing history
74 of a specific user to infer her interests), the quantity of available contextual information and market
75 expertise is often way beyond the capabilities of both advertisers and auctioneers. This reality
76 gave rise to "third-party" companies which develop technologies for collecting data and online
77 statistics used to infer the contexts of auctioned impressions (see, e.g., [16] and references therein).
78 Consequently, a new *distributed* ecosystem has emerged, in which many third-party companies
79 operate within the market aiming at maximizing their own utility (royalties or other compensations),
80 while significantly increasing the effectiveness of display advertising, as suggested by the following
81 article recently published by Facebook:

82 *"Many businesses today work with third parties such as Acxiom, Datalogix, and Epsilon*
83 *to help manage and understand their marketing efforts. For example, an auto dealer may*
84 *want to customize an offer to people who are likely to be in the market for a new car. The*
85 *dealer also might want to send offers, like discounts for service, to customers that have*
86 *purchased a car from them. To do this, the auto dealer works with a third-party company*
87 *to identify and reach those customers with the right offer".*

88 (www.facebook.com, "Advertising and our Third-Party Partners", April 10, 2013.)

89 Hence, in reality sources of information are distributed. Typically, the information is distributed
90 among several mediators or information providers/brokers, and is not held (or mostly not held) by
91 a central authority/web-site owner. In the display advertising example, one information source may
92 know the gender and another may know the location of the web-site visitor, while the web-site
93 itself often lacks the capability to track such information. The information sources need to decide
94 on the communicated information. In this case the information sources become players in a game.
95 To make the situation clearer, assume (as above) that the value of each impression type for each
96 bidder is public-knowledge (as is typically the case in repeated interactions through ad exchanges
97 which share their logs with the participants), and the only unknown entity is the instantiation
98

of the impression type; given the information learned from the information sources each bidder will bid his true expected valuation; hence, the results of this game are determined solely by the information providers. Notice that if, in the aforementioned example, the information provider who knows the gender reveals it while the other reveals nothing, then the auctioneer receives a revenue of \$50 as in the centralized case, while the cases in which both information providers reveal their information or none of them do so result in lower revenues. This shows the subtlety of the situation.

The above suggests that a major issue to tackle is the study of *distributed signaling games*, going beyond the (centralized) realm of classical mechanism design. We use a model of the above display advertising setting, due to its centrality, as a tool to introduce this novel foundational topic. The distributed signaling games may be pure coordination games (a.k.a. distributed optimization), or non-cooperative games. In the context of pure coordination games each information source has the same utility from the output created by their joint signal. Namely, in the above example if the web-site owner pays each information source proportionally to the revenue obtained by the web-site owner then the aims of the information sources are identical. The main aim of the third parties/mediators is to choose their signals based on their privately observed information in a distributed manner in order to optimize their own payoffs. Notice that in a typical embodiment, which we adapt, due to both technical and legal considerations, the auctioneer does not synthesize reported signals into new ones nor the information providers are allowed to explicitly communicate among them about the signals, but can only broadcast information they individually gathered. The study of the computational complexity of this highly fundamental problem is the major technical challenge tackled in this paper. Interestingly, we show a wide gap between the computational complexity of the centralized and of the distributed signaling setups, proving that coordinating on optimal signaling is a much harder problem than the one discussed in the context of centralized mechanism design. On the other hand we also show a natural restriction on the way information is distributed among information providers, which allows for an efficient constant approximation scheme.

In the context of non-cooperative games the outcome generated by the information sources' reports may result in a different value for each of them. The reason for that is typically the desire of the auctioneer to align the incentives of the mediators with his own by a compensation relative to the marginal benefit from their signals. In the above example one may compare the revenue obtained without the additional information sources, to what is obtained through their help, and compensate relatively to the Shapley values of their contributions, which is a standard (and rigorously justified) tool to fully divide a gain yielded by the cooperation of several parties. Here we apply such division to distributed signaling games, and show that it possesses some interesting properties: in particular the corresponding game has a pure strategy equilibrium, a property of the Shapley value which is shown for the first time for signaling settings (and is vastly different from previous studies of Shapley mechanisms in non-cooperative settings such as cost-sharing games [18]).

1.1 Model

Our model is a generalization of the one defined in [11]. There is a ground set $I = [n]$ of potential items (contexts) to be sold and a set $B = [k]$ of bidders. The value of item j for bidder i is given as v_{ij} . Following the above discussion (and the previous line of work, e.g., [9, 11]), we assume the valuation matrix $V = \{v_{i,j}\}$ is publicly known. An auctioneer is selling a single random item j_R , distributed according to some publicly known prior distribution μ over I , using a second price auction (a more detailed description of the auction follows). There is an additional set $M = [m]$ of "third-party" mediators. Following standard practice in game theoretic information models [1, 8, 10],

we assume each mediator $t \in M$ is equipped with a *partition* (signal-set) $\mathcal{P}_t \in \Omega(I)$ ¹. Intuitively, \mathcal{P}_t captures the extra information t has about the item which is about to be sold—he knows the set $S \in \mathcal{P}_t$ to which the item j_R belongs, but has no further knowledge about which item of S it is (except for the a priori distribution μ)—in other words, the distribution t has in mind is $\mu|_S$. For example, if the signal-set partition \mathcal{P}_t partitions the items of I into pairs, then mediator t knows to which pair $\{j_1, j_2\} \in \mathcal{P}_t$ the item j_R belongs, but he has no information whether it is j_1 or j_2 , and therefore, from her point of view, $\Pr[j_R = j_1] = \mu(j_1)/\mu(\{j_1, j_2\})$.

Mediators can signal some (or all) of the information they own to the network. Formally, this is represented by allowing each mediator t to report *any* super-partition \mathcal{P}'_t , which is obtained by merging partitions in her signal-set partition \mathcal{P}_t (in other words \mathcal{P}_t must be a refinement of \mathcal{P}'_t). In other words, a mediator may report any partition \mathcal{P}'_t for which there exists a set $Q'_t \in \Omega(\mathcal{P}_t)$ such that $\mathcal{P}'_t = \{\cup_{S \in A} S \mid A \in Q'_t\}$. In particular, a mediator can always report $\{I\}$, in which case we say that he “remains silent” since he does not contribute any information. The signals $\mathcal{P}'_1, \mathcal{P}'_2, \dots, \mathcal{P}'_m$ reported by the mediators are *broadcasted*² to the bidders, inducing a combined partition $\mathcal{P} \triangleq \times_{t=1}^m \mathcal{P}'_t = \{\cap_{i \in M} A_i \mid A_i \in \mathcal{P}'_i\}$, which we call the *joint partition* (or *joint signal*). \mathcal{P} splits the auction into separate “restricted” auctions: For each bundle $S \in \mathcal{P}$, the item j_R belongs to S with probability $\mu(S) = \sum_{j \in S} \mu(j)$, in which case S is signaled to the bidders and a second-price auction is performed over $\mu|_S$. Notice that if the signaled bundle is $S \subseteq I$, then the (expected) value of bidder i for $j_R \sim \mu|_S$ is $v_{i,S} = \frac{1}{\mu(S)} \sum_{j \in S} (\mu(j) \cdot v_{ij})$, and the truthfulness of the second price auction implies that this will also be bidder i ’s bid for the restricted auction. The winner of the auction is the bidder with the maximum bid $\max_{i \in B} v_{i,S}$, and he is charged the second highest valuation for that bundle $\max_{i \in B}^{(2)} v_{i,S}$. Therefore, the auctioneer’s revenue with respect to \mathcal{P} is the expectation (over $S \in \mathcal{P}$) of the price paid by the winning bidder:

$$R(\mathcal{P}) = \sum_{S \in \mathcal{P}} \mu(S) \cdot \max_{i \in B}^{(2)} (v_{i,S}) .$$

The joint partition \mathcal{P} signaled by the mediators can dramatically affect the revenue of the auctioneer. Consider, for example, the case where V is the 4×4 identity matrix, μ is the uniform distribution, and M consists of two mediators associated with the partitions $\mathcal{P}_1 = \{\{1, 2\}, \{3, 4\}\}$ and $\mathcal{P}_2 = \{\{1, 3\}, \{2, 4\}\}$. If both mediators remain silent, the revenue is $R(\{I\}) = 1/4$ (as this is the average value of all 4 bidders for a random item). However, observe that $\mathcal{P}_1 \times \mathcal{P}_2 = \{\{1\}, \{2\}, \{3\}, \{4\}\}$, and the second highest value in every column of V is 0, thus, if both report their partitions, the revenue drops to $R(\mathcal{P}_1 \times \mathcal{P}_2) = 0$. Finally, if mediator 1 reports \mathcal{P}_1 , while mediator 2 keeps silent, the revenue increases from $1/4$ to $R(\mathcal{P}_1) = 1/2$, as the value of each pair of items is $1/2$ for two different bidders (thus, the second highest price for each pair is $1/2$). This example can be easily generalized to show that in general the intervention of mediators can increase the revenue by a factor of $n/2$!

Indeed, the purpose of this paper is to understand how mediators’ (distributed) signals affect the revenue of the auctioneer. We explore the following two aspects of this question:

¹For a set S , $\Omega(S) \triangleq \{\mathcal{A} \subseteq 2^S \mid \cup_{A \in \mathcal{A}} A = S, \forall A, B \in \mathcal{A} A \cap B = \emptyset\}$ is the collection of all partitions of S .

²By saying that a mediator reports \mathcal{P}'_t , we mean that he reports the bundle $S \in \mathcal{P}'_t$ for which $j_R \in S$. The reader may wonder why our model is a broadcast model, and does not allow the mediators to report their information to the auctioneer through private channels, in which case the ad network will be able to manipulate and publish whichever information that best serves its interest. The primary reason for the broadcast assumption is that online advertising is a highly dynamic marketplace in which mediators often “come and go”, so implementing “private contracts” is infeasible. The second reason is that real-time bidding environments cannot afford the latency incurred by such a two-phase procedure in which the auctioneer first collects the information, and then selectively publishes it. The auction process is usually treated as a “black box”, and modifying it harms the modularity of the system.

- (1) (Computational) Suppose the auctioneer has control over the signals reported by the mediators. We study the computational complexity of the following problem. *Given a $k \times n$ matrix V of valuations and mediators' partitions $\mathcal{P}_1, \mathcal{P}_2, \dots, \mathcal{P}_m$, what is the revenue maximizing joint partition $\mathcal{P} = \mathcal{P}'_1 \times \dots \times \mathcal{P}'_m$?* We call this problem the *Distributed Signaling Problem*, and denote it by $\text{DSP}(n, k, m)$.

We note that the problem studied in [11] is a special case of DSP, in which there is a *single* mediator ($m = 1$) with *perfect knowledge* about the item sold and can report any desirable signal (partition).³

- (2) (Strategic) What if the auctioneer cannot control the signals reported by the mediators (as the reality of the problem usually entails)? *Can the auctioneer introduce compensations that will incentivize mediators to report signals leading to increased revenue in the auction, when each mediator is acting selfishly?*

This is a mechanism design problem: Here the auctioneer's goal is to design a payment rule (i.e., a mechanism) for allocating (part of) his profit from the auction among the mediators, based on their reported signals and the auction's outcome, so that global efficiency (i.e., maximum revenue) emerges from their signals.

Section 1.2 summarizes our findings regarding the two above problems.

1.2 Our Results

Ghosh et al. [11] showed that computing the revenue-maximizing signal in their "perfect-knowledge" setup is NP-hard, but present an efficient algorithm for computing a 2-approximation of the optimal signal (partition). We show that when information is distributed, the problem becomes much harder. More specifically, we present a gap-preserving reduction from the *Maximum Independent Set* problem to DSP.

THEOREM 1.1 (HARDNESS OF APPROXIMATING DSP). *If there exists an $O(m^{1/2-\epsilon})$ approximation (for some constant $\epsilon > 0$) for instances of $\text{DSP}(2m, m+1, m)$, then there exists an $O(N^{1-2\epsilon})$ approximation for Maximum Independent Set (MIS_N), where N is the number of nodes in the underlying graph of the MIS instance.*

Since the Maximum Independent Set problem is NP-hard to approximate to within a factor of $n^{1-\rho}$ for any fixed $\rho > 0$ [13], Theorem 1.1 indicates that approximating the revenue-maximizing signal, even within a multiplicative factor of $O((\min\{n, k, m\})^{1/2-\epsilon})$, is NP-hard. In other words, one cannot expect a reasonable approximation ratio for $\text{DSP}(n, k, m)$ when the three parameters of the problem are all "large". The next theorem shows that a "small" value for either one of the parameters n or k indeed implies a better approximation ratio.

THEOREM 1.2 (APPROXIMATION ALGORITHM FOR SMALL n OR k). *For $k \geq 2$, there is a polynomial time $\min\{n, k-1\}$ -approximation algorithm for $\text{DSP}(n, k, m)$.⁴*

We leave open the problem of determining whether one can get an improved approximation ratio when the parameter m is "small". For $m = 1$, the result of [11] implies immediately a 2-approximation algorithm. However, even for the case of $m = 2$ we are unable to find an algorithm having a non-trivial approximation ratio. We mitigate the above results by proving that for a natural (and realistic) class of mediators called *local experts* (defined in Section 3), there exists a polynomial time 5-approximation algorithm for DSP.

³In other words, \mathcal{P}_1 is the partition of I into singletons.

⁴For $k = 1$, any algorithm is optimal since the use of a second price auction implies that the revenue of any strategy profile is 0 when there is only one bidder.

246 THEOREM 1.3 (A 5-APPROXIMATION ALGORITHM FOR LOCAL EXPERT MEDIATORS). *If mediators are*
 247 *local experts, there exists a polynomial time 5-approximation algorithm for DSP.*

248 In the strategic setup, we design a fair (symmetric) payment rule $\mathcal{S} : (\mathcal{P}'_1, \mathcal{P}'_2, \dots, \mathcal{P}'_m) \rightarrow \mathbb{R}_+^m$
 249 for incentivizing mediators to report useful information they own, and refrain from reporting
 250 information with negative impact on the revenue. This mechanism is inspired by the *Shapley Value*—
 251 it distributes part of the auctioneer’s surplus among the mediators according to their expected
 252 relative marginal contribution to the revenue, when ordered randomly.⁵ We first show that this
 253 mechanism always admits a *pure* Nash equilibrium, a property we discovered to hold for arbitrary
 254 games where the value of the game is distributed among players according to Shapley’s value
 255 function.
 256

257 THEOREM 1.4. *Let \mathcal{G}_m be a non-cooperative m -player game in which the payoff of each player is*
 258 *set according to \mathcal{S} . Then \mathcal{G}_m admits a pure Nash equilibrium. Moreover, best response dynamics are*
 259 *guaranteed to converge to such an equilibrium.*

260 We then turn to analyze the revenue guarantees of our mechanism \mathcal{S} . Our first theorem shows
 261 that using the mechanism \mathcal{S} never decreases the revenue of the auctioneer compared to the initial
 262 state (i.e., when all mediators are silent).
 263

264 THEOREM 1.5. *For every Nash equilibrium $(\mathcal{P}'_1, \mathcal{P}'_2, \dots, \mathcal{P}'_m)$ of \mathcal{S} , $R(\times_{t \in M} \mathcal{P}'_t) \geq R(\{I\})$.*

265 The next two theorems provide tight bounds on the *price of anarchy* and *price of stability* of \mathcal{S} .⁶
 266 Unlike in the computational setup, even restricting the mediators to be local experts does not enable
 267 us to get improved results here.
 268

269 THEOREM 1.6. *For $k \geq 2$, the price of anarchy of \mathcal{S} under any instance $\text{DSP}(n, k, m)$ is no more*
 270 *than $\min\{k - 1, n\}$.*

271 THEOREM 1.7. *For every $n \geq 1$, there is a $\text{DSP}(3n + 1, n + 2, 2)$ instance for which the price of*
 272 *stability of \mathcal{S} is at least n . Moreover, all the mediators in this instance are local experts.*
 273

274 Interestingly, an adaptation of Shapley’s uniqueness theorem [19] to our non-cooperative setting
 275 asserts that the price of anarchy of our mechanism is inevitable if one insists on a few natu-
 276 ral requirements—essentially anonymity and efficiency⁷ of the payment rule—and assuming the
 277 auctioneer alone can introduce payments. We discuss this further in Section 4.3.
 278

279 1.3 Additional Related Work

280 The formal study of internet auctions with contexts was introduced by [9] where the authors studied
 281 the impact of contexts in the related Sponsored Search model, and showed that *bundling contexts*
 282 may have a significant impact on the revenue of the auctioneer. The subsequent work of Ghosh et.
 283 al. [11] considered the computational algorithmic problem of computing the revenue maximizing
 284 partition of items into bundles, under a second price auction in the full information setting. Recently,
 285 Emek et al. [8] and Bro Miltersen and Sheffet [3] studied signaling (which generalizes bundling)
 286 in the context of display advertising. They explore the computational complexity of computing

287 ⁵Shapley’s value was originally introduced in the context of cooperative games, where there is a well defined notion of a
 288 coalition’s value. In order to apply this notation to a non-cooperative game, we assume the game has some underlying
 289 global function ($v(\cdot)$) assigning a value to every strategy profile of the players, and the Shapley value of each player is
 290 defined with respect to $v(\cdot)$. In this setting, a “central planner” (the auctioneer in our case) is the one making the utility
 291 transfer to the “coalised” players. For the formal axiomatic definition of a value function and Shapley’s value function, see
 292 [19].

292 ⁶The price of anarchy (stability) is the ratio between the revenue of the optimum and the worst (best) Nash equilibrium.

293 ⁷I.e., the sum of payments is equal to the total surplus of the auctioneer.
 294

a signaling scheme that maximizes the auctioneer’s revenue in a Bayesian setting. On the other hand, Guo and Deligkas [12] studied a special case of bundling where only “natural” bundles are allowed. Unlike our distributed setup, all the above models are *centralized*, in the sense that the auctioneer has full control over the bundling process (which in our terms corresponds to having a single mediator with a perfect knowledge about the item sold).

A different model with knowledgeable third parties was recently considered by Cavallo et al. [4]. However, the focus of this model is completely different than ours. More specifically, third parties in this model use their information to estimate the clicks-per-impression ratio, and then use this estimate to bridge between advertisers who would like to pay-by-click and ad networks which use a pay-by-impression payment scheme.

The proof of Theorem 1.4 is based on the existence of a potential function for our mechanism \mathcal{S} . There has been work showing that Shapley-based mechanism possesses a potential function in other contexts as well. In particular, such mechanisms apply to resource-selection forms (see, e.g., [14, 17]). Moreover, [17] shows a general class of resource selection games, such that a local condition on cost/welfare distribution per resource must be according to the Shapley value in order to yield a potential function. However, unlike in our work and in a recent work in Machine Learning [2], the above refer to local conditions on the distribution of gain/cost of a given resource, rather than the evaluation of an individual’s informational contribution. The work in [2] decides on a content to be shown based on its relative informational quality, when the exact probability of showing a content depends on its Shapley value in a cooperative game, while in our setting we decide on monetary compensation based on marginal informational contribution as described by a Shapley-based mechanism. These three genres of work point to the importance of Shapley-based mechanisms. It seems, though, that finding formal connections between them is not immediate, and may be a subject for future study.

2 PRELIMINARIES

Throughout the paper we use capital letters for sets and calligraphic letters for set families. For example, the partition \mathcal{P}_t representing the knowledge of mediator t is a set of sets, and therefore, should indeed be calligraphic according to this notation. A mechanism \mathcal{M} is a tuple of payment functions $(\Pi_1, \Pi_2, \dots, \Pi_m)$ determining the compensation of every mediator given a strategy profile (i.e., $\Pi_t : \Omega(\mathcal{P}_1) \times \Omega(\mathcal{P}_2) \times \dots \times \Omega(\mathcal{P}_m) \rightarrow \mathbb{R}^+$). Every mechanism \mathcal{M} induces the following game between mediators.

Definition 2.1 (DSP game). Given a mechanism $\mathcal{M} = (\Pi_1, \Pi_2, \dots, \Pi_m)$ and a DSP(n, k, m) instance, the DSP $_{\mathcal{M}}(n, k, m)$ game is defined as follows. Every mediator $t \in M$ is a player whose strategy space consists of all partitions \mathcal{P}'_t for which \mathcal{P}_t is a refinement. Given a strategy profile $\mathcal{P}'_1, \mathcal{P}'_2, \dots, \mathcal{P}'_m$, the payoff of mediator t is $\Pi_t(\mathcal{P}'_1, \mathcal{P}'_2, \dots, \mathcal{P}'_m)$.

Given a DSP instance and a set $S \subseteq I$, we use the shorthand $v(S) := \max_{i \in B}^{(2)}(v_{i,S})$ to denote the second highest bid in the restricted auction $\mu|_S$. Using this notation, the expected revenue $R(\mathcal{P})$ of the auctioneer under the (joint) partition \mathcal{P} of the mediators can be restated as

$$R(\mathcal{P}) = \sum_{S \in \mathcal{P}} \mu(S) \cdot v(S) .$$

For a DSP $_{\mathcal{M}}$ game, let $\mathcal{E}(\mathcal{M})$ denote the set of Nash equilibria of this game and let \mathcal{P}^* be a maximum revenue strategy profile. The *Price of Anarchy* and *Price of Stability* of DSP $_{\mathcal{M}}$ are defined as:

$$\text{PoA} := \max_{\mathcal{P} \in \mathcal{E}(\mathcal{M})} \frac{R(\mathcal{P}^*)}{R(\mathcal{P})} , \quad \text{and} \quad \text{PoS} := \min_{\mathcal{P} \in \mathcal{E}(\mathcal{M})} \frac{R(\mathcal{P}^*)}{R(\mathcal{P})} ,$$

344 respectively. Notice that our definition of the price of anarchy and price of stability differs from the
 345 standard one by using revenue instead of social welfare.

346 3 THE COMPUTATIONAL COMPLEXITY OF DISTRIBUTED SIGNALING (DSP)

347 This section explores DSP from a pure combinatorial optimization viewpoint. In other words, we
 348 assume the auctioneer can control the signals produced by each mediator. The objective of the
 349 auctioneer is then to choose a distributed strategy profile $\mathcal{P}'_1, \mathcal{P}'_2, \dots, \mathcal{P}'_m$ whose combination $\times_t \mathcal{P}'_t$
 350 yields maximum revenue in the resulting auction. We begin with negative results (proving Theorem
 351 1.1 in Subsection 3.1), and then proceed with a few approximation algorithms for the problem
 352 (Subsection 3.2), including a 5-approximation algorithm for the case of *local expert* mediators
 353 (Theorem 1.3).
 354

355 3.1 Hardness of Approximating DSP

356 It is not hard to show that solving DSP exactly, i.e., finding the maximum revenue strategy profile,
 357 is NP-hard. In fact, this statement directly follows from the NP-hardness result of [11] by observing
 358 that the special case of a single mediator with perfect knowledge ($\mathcal{P}_1 = \{\{j\} \mid j \in I\}$) is equivalent
 359 to the centralized signaling model of [11].
 360

361 PROPOSITION 3.1. *Solving DSP exactly is NP-Hard.*

362 The main result of this section is that approximating the revenue-maximizing signals, even
 363 within a multiplicative factor of $O((\min\{n, k, m\})^{1/2-\epsilon})$, remains NP-hard. This result is achieved
 364 by a gap preserving reduction from Maximum Independent Set to DSP.
 365

366 THEOREM 3.2. *For every integer $\ell \geq 1$, an $\alpha \geq 1$ approximation for DSP($2\ell N, \ell N + 1, \ell N$) induces
 367 a $\alpha(1 + \ell^{-1}(N + 1))$ approximation for the Maximum Independent Set problem (where N is the number
 368 of nodes in the graph).
 369*

370 Observe that Theorem 1.1 follows easy from Theorem 3.2.

371 THEOREM 1.1. *If there exists an $O(m^{1/2-\epsilon})$ approximation (for some constant $\epsilon > 0$) for instances of
 372 DSP($2m, m + 1, m$), then there exists a $O(N^{1-2\epsilon})$ approximation for Maximum Independent Set (MIS $_N$),
 373 where N is the number of nodes in the underlying graph of the MIS instance.
 374*

375 PROOF. Let $\ell = N + 1$. By Theorem 3.2, there exists an approximation algorithm for Maximum
 376 Independent Set whose approximation ratio is 2α , where α is the approximation ratio that can be
 377 achieved for instances of DSP($2\ell N, \ell N + 1, \ell N$). On the other hand, by our assumption:
 378

$$379 \alpha = O((\ell N)^{1/2-\epsilon}) = O(N^{2(1/2-\epsilon)}) = O(N^{1-2\epsilon}),$$

380 which completes the proof. □

381 In the rest of this section we prove Theorem 3.2. The high-level idea of the reduction is as follows.
 382 Given a graph $G = (V, E)$, we map it to a DSP instance by associating ℓ pairs of (equi-probable) items
 383 $\{(j_{v,k}, j'_{v,k})\}_{k=1}^{\ell}$ with each node of the graph $v \in V$, where the items $j'_{v,k}$ are auxiliary items called
 384 the “helper items” of node v . Additionally, for each node we have ℓ single-minded bidders, each of
 385 which is interested (exclusively) in a specific item $j_{v,k}$. An additional bidder (i_h) is interested in *all*
 386 helper-items of all nodes. Finally, each node $v \in V$ has a corresponding set of ℓ mediators $\{m_{v,k}\}_{k=1}^{\ell}$
 387 (one for each items pair). Each mediator $m_{v,k}$ has a single bit of information, corresponding to
 388 whether or not the auctioned item belongs to the union set of *the k -th pair $(j_{v,k}, j'_{v,k})$ together with*
 389 *all helper elements of the neighbor nodes of v .*

390 As the helper elements are valuable only to a *single* bidder i_h , a part in a partition can have a
 391 non-zero contribution to the revenue only if it contains at least two non-helper elements or at
 392

least one helper element and at least one non-helper element. However, our construction ensures that, whenever a strategy profile \mathcal{P} involves at least two “active” mediators of neighboring nodes $(u, v) \in E$, the resulting joint partition isolates the non-helper elements of u from helper elements (except for maybe one part that might contain elements of both types). Thus, in a high revenue strategy profile the set of nodes associated with many “speaking” (active) mediators must be close to an independent set (see Figure 1 for an illustration).

We now proceed with the formal proof. Given a graph $G = (V, E)$ of N nodes, consider an instance of DSP consisting of the following:

- A set $\{j_{v,k}, j'_{v,k} \mid v \in V, 1 \leq k \leq \ell\}$ of $2\ell N$ items having equal probabilities to appear. For every node v and integer $1 \leq k \leq \ell$, the element $j'_{v,k}$ is referred to as the “helper” item of $j_{v,k}$. For notational convenience, we define the following sets for every node $v \in V$: $J_v = \{j_{v,k} \mid 1 \leq k \leq \ell\}$ and $H_v = \{j'_{v,k} \mid 1 \leq k \leq \ell\}$.
- A set $\{i_{v,k}, \mid v \in V, 1 \leq k \leq \ell\} \cup \{i_h\}$ of $\ell N + 1$ bidders. For every node v and integer $1 \leq k \leq \ell$, the bidder $i_{v,k}$ has a value of $2\ell N$ for the item $j_{v,k}$ and a value of 0 for all other item. The remaining bidder i_h has a value of $2\ell N$ for each helper item (i.e., each item from the set $\{j'_{v,k} \mid v \in V, 1 \leq k \leq \ell\}$) and a value of 0 for the non-helper item.
- There are ℓN mediators $\{m_{v,k}, \mid v \in V, 1 \leq k \leq \ell\}$. For every node v and integer $1 \leq k \leq \ell$ the partition of mediator $m_{v,k}$ is defined as follows:

$$\mathcal{P}_{v,k} = \left\{ \{j_{v,k}, j'_{v,k}\} \cup \bigcup_{u \mid uv \in E} H_u \right\} \cup \{R_v\},$$

where R_v is the set of remaining elements that do not belong to the first part of the partition. Informally, $\mathcal{P}_{v,k}$ is a binary partition where on one side we have the two elements $j_{v,k}$ and $j'_{v,k}$ and all the helper elements corresponding to neighbor nodes of v , and on the other side we have the rest of the elements (we illustrate this construction in Figure 1).

Let us begin the analysis of the above DSP instance by determining the contribution of each part in an arbitrary partition \mathcal{P} to $R(\mathcal{P})$.

CLAIM 3.3. *Let \mathcal{P} be an arbitrary partition. Then, for every part $S \in \mathcal{P}$, the contribution of S to $R(\mathcal{P})$ is 0, unless $|S| \geq 2$ and S contains at least one non-helper item. In the last case, the contribution of S to the total revenue is:*

$$\mu(S) \cdot v(S) = 1.$$

PROOF. There are a few cases to consider.

- If S contains only helper elements, then it is valuable only to i_h , and thus, has a 0 contribution to $R(\mathcal{P})$.
- If S contains only one element, then it is valuable only to one bidder because each element is valuable only to one element. Thus, it again contributes 0 to $R(\mathcal{P})$.
- If S contains multiple elements, at least one of which is *non-helper*, then it is valuable to at least two bidders. Specifically, for every element $j_{v,k} \in S$, $v_{i_{v,k}, S} = 2\ell N/|S|$. Additionally, if there are helper elements in S , then:

$$v_{i_h, S} = \frac{2\ell N \cdot |S \cap \{j'_{v,k} \mid v \in V, 1 \leq k \leq \ell\}|}{|S|} \geq \frac{2\ell N}{|S|}.$$

Hence, for any such part S we get $v(S) = \max_{i \in B}^{(2)} v_{i, S} = 2\ell N/|S|$, and the contribution of the part to $R(\mathcal{P})$ is $\mu(S) \cdot v(S) = 1$. \square




Fig. 1. A simplified graphical illustration of the reduction. The blue lines represent edges in the original graph G . The black (gray) circles represent a single one out of the ℓ non-helper (helper) items of each node. The two polygons represent the partition sets of mediators $m_{u,1}, m_{v,1}$, corresponding to the (neighboring) nodes u, v , respectively. If *both* mediators speak, the resulting joint partition isolates the non-helper element(s) of u from each from helper elements, as evident by the intersection region of both polygons. Notice that, for each node x in the graph, the partition of each mediator $m_{x,k}$ includes *all* neighboring non-helper items, across all ℓ copies, but only the k 'th copy of the non-helper item of node x (this feature is missing from the figure for simplicity, but is crucial for the analysis).

Let A be an arbitrary independent set of G , and let $S_A := \{m_{v,k} \mid v \in A, 1 \leq k \leq \ell\}$. The following claim lower bounds the revenue of the joint partition arising when the mediators of S_A are the only speaking mediators.

CLAIM 3.4. $R(\times_{m_{v,k} \in S_A} \mathcal{P}_{v,k}) \geq |S_A| = \ell \cdot |A|$.

PROOF. Observe that $\mathcal{P}_{v,k}$ separates $j_{v,k}$ from every other item of the set $\{j_{v,k} \mid v \in V, 1 \leq k \leq \ell\}$. Hence, $\times_{m_{v,k} \in S_A} \mathcal{P}_{v,k}$ contains $|S_A|$ different parts $\{T_{v,k} \mid m_{v,k} \in S_A\}$, where each part $T_{v,k}$ contains $j_{v,k}$. On the other hand, each pair $(j_{v,k}, j'_{v,k})$ of items is separated only by the partitions of mediators corresponding to neighbors of v . Since A is independent, this implies that $j_{v,k}$ and $j'_{v,k}$ share part in $\times_{m_{v,k} \in S_A} \mathcal{P}_{v,k}$ for every mediator $m_{v,k} \in S_A$. In other words, for every $m_{v,k} \in S_A$, the part $T_{v,k}$ contains $j'_{v,k}$ in addition to $j_{v,k}$, and thus by Claim 3.3, contributes 1 to $R(\times_{m_{v,k} \in S_A} \mathcal{P}_{v,k})$. Therefore, $R(\times_{m_{v,k} \in S_A} \mathcal{P}_{v,k}) \geq |S_A| = \ell \cdot |A|$, as claimed. \square

Claim 3.4 asserts that there exists a solution for the above DSP instance whose value is at least $\ell \cdot OPT$, where OPT is the size of the maximum independent set in G .

Consider now an arbitrary set S of mediators, and let $S' = \{m_{v,k} \in S \mid \forall m_{u,k'} \in S \ uv \notin E\}$. Informally, a mediator $m_{v,k}$ is in S' if it belongs to S and no neighbor node u of v has a mediator in S . The following lemma upper bounds in terms of $|S'|$ the revenue of the joint partition of the mediators in S .

LEMMA 3.5. $R(\times_{m_{v,k} \in S} \mathcal{P}_{v,k}) \leq |S'| + N + 1$.

PROOF. Each partition $\mathcal{P}_{v,k}$ separates a single non-helper element $j_{v,k}$ from the other non-helper elements. Hence, $\times_{m_{v,k} \in S} \mathcal{P}_{v,k}$ consists of at most $|S| + 1$ parts containing non-helper elements (recall that parts with only helper elements have 0 contribution to $R(\times_{m_{v,k} \in S} \mathcal{P}_{v,k})$, and thus, can be ignored). Let us label these parts $\{T_{v,k}\}_{m_{v,k} \in S}, T$, where $T_{v,k}$ is the part containing $j_{v,k}$ for every $m_{v,k} \in S$ and T is the part containing the remaining non-helper elements. Let us upper bound the contribution of each such part to $R(\times_{u \in S} \mathcal{P}_u)$.

ALGORITHM 1: Independent Set Algorithm

-
- 1 Construct a DSP instance from the independent set instance as described above.
 - 2 Run the α -approximation algorithm for DSP assumed by Theorem 3.2 on the constructed instance, and let S be the set of mediators speaking in the obtained strategy profile.
 - 3 Calculate the configuration $S' = \{m_{v,k} \in S \mid \forall m_{u,k'} \in S \ uv \notin E\}$.
 - 4 Calculate the independent set $A' = \{v \in V \mid \exists_{1 \leq k \leq \ell} m_{v,k} \in S'\}$.
 - 5 If A' is non-empty output it, otherwise output an arbitrary single node.
-

- The part T and all the parts $\{T_{v,k}\}_{m_{v,k} \in S'}$ can contribute at most 1 each because no part has a larger contribution.
- Consider a part $T_{v,k}$ obeying $m_{v,k} \in S \setminus S'$ and $|J_v \cap S| = 1$ (i.e., $m_{v,k}$ is the only mediator of v belonging to S). This part also contribute at most 1, but there can be at most N such parts, one for every node.
- Finally, consider a part $T_{v,k}$ obeying $m_{v,k} \in S \setminus S'$ and $|J_v \cap S| \geq 2$ (i.e., at least two mediators of v belong to S). By the construction of $\mathcal{P}_{v,k}$, $T_{v,k}$ can contain in addition to $j_{v,k}$ only the corresponding helper element $j'_{v,k}$ and helper elements from $\bigcup_{u|uv \in E} H_u$. Let us see why none of these helper elements actually belongs to $T_{v,k}$, and thus, the part T_v contains only $j_{v,k}$ and contributes 0.
 - Since $m_{v,k} \notin S'$, there exists a neighbor node v' of v having a mediator $m_{v',k'} \in S$. Then, the partition $\mathcal{P}_{v',k'}$ separates $j'_{v,k}$ from $j_{v,k}$ and guarantees that $j'_{v,k} \notin T_{v,k}$.
 - Let $m_{v,k'}$ be another mediator in $J_v \cap S$ (exists since $|J_v \cap S| \geq 2$). Then, the partition $\mathcal{P}_{v,k'}$ separates the helper elements of $\bigcup_{u|uv \in E} H_u$ from $j_{v,k}$ and guarantees that none of these helper mediators belongs to $T_{v,k}$.

In conclusion: $R(\times_{m_{v,k} \in S} \mathcal{P}_{v,k}) \leq |S'| + N + 1$, as claimed. \square

We are now ready to prove Theorem 3.2.

PROOF OF THEOREM 3.2. Consider Algorithm 1. We would like to show that this algorithm is an $\alpha(1+\ell^{-1}(N+1))$ -approximation algorithm for Maximum Independent Set, which proves the theorem. The DSP instance constructed by Algorithm 1 has a strategy profile of revenue at least $\ell \cdot OPT$ by Claim 3.4. Since S is obtained using an α -approximation algorithm, $R(\times_{m_{v,k} \in S} \mathcal{P}_{v,k}) \geq \ell \cdot OPT / \alpha$ (recall that all mediators have binary partitions, and thus, they must report their partitions as is when they speak). By Lemma 3.5 we now get:

$$|S'| \geq R(\times_{m_{v,k} \in S} \mathcal{P}_{v,k}) - (N + 1) \geq \ell \cdot OPT / \alpha - (N + 1) .$$

Informally, A' is the set of nodes having mediators in S' . The independence of A' follows from the construction of S' , which guarantees that Algorithm 1 outputs an independent set. Additionally, since each node has ℓ mediators:

$$|A'| \geq \frac{|S'|}{\ell} \geq \frac{\ell \cdot OPT / \alpha - (N + 1)}{\ell} = \frac{OPT}{\alpha} - \frac{N + 1}{\ell} .$$

If $OPT \geq \alpha(1 + \ell^{-1}(N + 1))$, then:

$$\begin{aligned} |A'| &\geq \frac{OPT}{\alpha} - \frac{N + 1}{\ell} \\ &= \frac{OPT}{\alpha(1 + \ell^{-1}(N + 1))} + \frac{[1 - (1 + \ell^{-1}(N + 1))^{-1}] \cdot OPT}{\alpha} - \frac{N + 1}{\ell} \\ &\geq \frac{OPT}{\alpha(1 + \ell^{-1}(N + 1))} . \end{aligned}$$

On the other hand, if $OPT \leq \alpha(1 + \ell^{-1}(N + 1))$, then the solution of Algorithm 1 is of size at least $\frac{OPT}{\alpha(1 + \ell^{-1}(N + 1))}$ simply because it is not empty. \square

3.2 Approximation Algorithms for DSP

In light of Theorem 1.1, an efficient algorithm with a reasonable approximation guarantee for general DSP is unlikely to exist when the three parameters of the problem are all “large”. Subsection 3.2.1 gives a trivial algorithm which has a good approximation guarantee when either n or k is small. A more interesting result is given in Subsection 3.2.2, which proves a 5-approximation algorithm for DSP under the assumption that the mediators are *local experts* (as stated in Theorem 1.3).

3.2.1 A simple $\min\{n, k - 1\}$ -approximation algorithm for DSP. In this section we prove the following theorem:

THEOREM 1.2. *For $k \geq 2$, there is a polynomial time $\min\{n, k - 1\}$ -approximation algorithm for $DSP(n, k, m)$.*

PROOF. We show that the algorithm that simply returns the partition $\{I\}$, the joint partition corresponding to the case where all mediators are silent, has the promised approximation guarantee. For that purpose we analyze the revenue of $\{I\}$ in two different ways:

- Let $\mathcal{P}' = (\mathcal{P}'_1, \mathcal{P}'_2, \dots, \mathcal{P}'_m)$ be an arbitrary strategy profile of the instance in question. The revenue of \mathcal{P}' is:

$$\begin{aligned} R(\times_{t=1}^m \mathcal{P}'_t) &= \sum_{S \in \times_{t=1}^m \mathcal{P}'_t} \mu(S) \cdot v(S) \leq |\times_{t=1}^m \mathcal{P}'_t| \cdot \max_{S \in \times_{t=1}^m \mathcal{P}'_t} \mu(S) \cdot v(S) \\ &\leq n \cdot \max_{S \in \times_{t=1}^m \mathcal{P}'_t} \mu(S) \cdot v(S) \leq n \cdot R(\{I\}) , \end{aligned}$$

where the last inequality holds since, for every set S , $R(\{I\}) = v(I) \geq v(S) \cdot \mu(S)$. This shows that the approximation ratio of the trivial strategy profile $\{I\}$ provides an n -approximation to the optimal revenue.

- Let $\mathcal{P}' = (\mathcal{P}'_1, \mathcal{P}'_2, \dots, \mathcal{P}'_m)$ be an arbitrary strategy profile of the instance in question. The revenue of \mathcal{P}' is:

$$\begin{aligned} R(\times_{t=1}^m \mathcal{P}'_t) &= \sum_{S \in \times_{t=1}^m \mathcal{P}'_t} \mu(S) \cdot v(S) = \sum_{S \in \times_{t=1}^m \mathcal{P}'_t} \mu(S) \cdot \left(\max_{i \in B}^{(2)} \frac{\sum_{j \in S} \mu(j) \cdot v_{i,j}}{\mu(S)} \right) \\ &= \sum_{S \in \times_{t=1}^m \mathcal{P}'_t} \left(\max_{i \in B}^{(2)} \sum_{j \in S} \mu(j) \cdot v_{i,j} \right) . \end{aligned}$$

For every bidder $i \in B$, let $\Sigma_i = \sum_{j \in I} \mu(j) \cdot v_{i,j}$. It is easy to see that $v(I) = \max_{i \in B}^{(2)} \Sigma_i$ (in other words, the second highest Σ_i value is $v(I)$). Let $i^* \in B$ be the index maximizing Σ_{i^*} (breaking ties arbitrary). Consider a set $S \in \times_{t=1}^m \mathcal{P}'_t$. The elements of S contribute at least $\max_{i \in B}^{(2)} \sum_{j \in S} \mu(j) \cdot v_{i,j}$ to at least two of the values: $\Sigma_1, \dots, \Sigma_n$. Thus, they contribute at

least the same quantity to the sum $\sum_{i \in B \setminus \{i^*\}} \Sigma_i$. This means that at least one of the values $\{\Sigma_i\}_{i \in B \setminus \{i^*\}}$ must be at least:

$$\frac{\sum_{S \in \times_{t=1}^m \mathcal{P}'_t} \left(\max_{i \in B}^{(2)} \sum_{j \in S} \mu(j) \cdot v_{i,j} \right)}{k-1} = \frac{R(\times_{t=1}^m \mathcal{P}'_t)}{k-1}.$$

By definition Σ_{i^*} must also be at least that large, and therefore,

$$R(\{I\}) = v(I) \geq R(\times_{t=1}^m \mathcal{P}'_t) / (k-1). \quad \square$$

3.2.2 A 5-approximation algorithm for local expert mediators. In this subsection we consider an interesting special case of DSP which is henceforth shown to admit a constant factor approximation.

Definition 3.6 (Local Expert mediators). A mediator t in a DSP instance is a *local expert* if there exists a set $I_t \subseteq I$ such that: $\mathcal{P}_t = \{\{j\} \mid j \in I_t\} \cup \{I \setminus I_t\}$.

Informally, a local expert mediator has perfect knowledge about a single set I_t —if the item belongs to I_t , he can tell exactly which item it is. In other words, a local expert mediator specializes in some kind of items to the extent that it knows everything about this kind of items, and nothing at all about other kinds of items. Our objective in the rest of the section is to prove Theorem 1.3, i.e., to describe a 5-approximation algorithm for instances of DSP consisting of only local expert mediators.

We begin the proof with an upper bound on the revenue of the optimal joint strategy, which we denote by \mathcal{P}^* . To describe this bound, we need some notation. We use \hat{I} to denote the set of items that are within the expertise domain of some mediator (formally, $\hat{I} = \bigcup_{t \in M} I_t$). Additionally, for every item $j \in I$, h_j and s_j denote $\mu(j)$ times the largest value and second largest value, respectively, of j for any bidder (more formally, $h_j = \mu(j) \cdot \max_{i \in B} v_{i,j}$ and $s_j = \mu(j) \cdot \max_{i \in B}^{(2)} v_{i,j}$).

Next, we need to partition the items into multiple sets. The optimal joint partition \mathcal{P}^* is obtained from partitions $\{\mathcal{P}_t^*\}_{t \in M}$, where \mathcal{P}_t^* is a possible partition for mediator t . Each part of \mathcal{P}^* is the intersection of $|M|$ parts, one from each partition in $\{\mathcal{P}_t^*\}_{t \in M}$. On the other hand, each part of \mathcal{P}_t^* is a subset of I_t , except for maybe a single part. Hence, there exists at most a single part $I_0 \in \mathcal{P}^*$ such that $I_0 \not\subseteq I_t$ for any $t \in M$. For ease of notation, if there is no such part (which can happen when $\hat{I} = I$) we denote $I_0 = \emptyset$. To partition the items of $I \setminus I_0$, we associate each part $S \in \mathcal{P}^* \setminus \{I_0\}$ with an arbitrary mediator t such that $S \subseteq I_t$, and denote by A_t the set of all items belonging to the parts associated with mediator t , as defined by the aforementioned process. Observe that the construction of A_t guarantees that $A_t \subseteq I_t$. Additionally, $\{I_0\} \cup \{A_t\}_{t \in M}$ is a disjoint cover (i.e., a partition) of I .

A different partition of the items partitions them according to the bidder that values them the most. In other words, for every $1 \leq i \leq k$, H_i is the set of items for which bidder i has the largest value. If multiple bidders have the same largest value for an item, we assign it to the set H_i of an arbitrary one of these bidders. Notice that the construction of H_i guarantees that the sets $\{H_i\}_{i \in B}$ are disjoint.

Finally, for every set $S \subseteq I$, we use $\phi(S)$ to denote the sum of the $|B| - 1$ smaller values in $\{\sum_{j \in H_i \cap S} h_j\}_{i \in B}$, i.e., the sum of all the values except the largest one. In other words, we calculate for every bidder i the sum of its values for items in $H_i \cap S$, and then add up the $|B| - 1$ smaller sums. Using all the above notation we can now state our promised upper bound on $R(\mathcal{P}^*)$.

LEMMA 3.7. $R(\mathcal{P}^*) \leq \mu(I_0) \cdot v(I_0) + \sum_{j \in \hat{I}} s_j + \sum_{t \in M} \phi(A_t)$.

PROOF. Fix an arbitrary mediator $t \in M$, and let i be the bidder whose term is not counted by $\phi(A_t)$. For every part $S \in \mathcal{P}^*$ associated with t , let i' be a bidder other than i that has one of the

638 two largest bids for S . By definition:

$$639 \quad \mu(S) \cdot v(S) = \max_{i' \in B}^{(2)} \sum_{j \in S} \mu(j) \cdot v_{i',j} \leq \sum_{j \in S} \mu(j) \cdot v_{i',j} \leq \sum_{j \in S \cap H_i} s_j + \sum_{j \in S \setminus H_i} h_j .$$

642 Summing over all parts associated with t , we get:

$$643 \quad \sum_{\substack{S \in \mathcal{P}^* \\ S \subseteq A_t}} \mu(S) \cdot v(S) \leq \sum_{j \in A_t \cap H_i} s_j + \sum_{j \in A_t \setminus H_i} h_j \leq \sum_{j \in A_t} s_j + \phi(A_t) .$$

646 Summing over all mediators, we get:

$$647 \quad R(\mathcal{P}^*) - \mu(I_0) \cdot v(I_0) \leq \sum_{t \in M} \left(\sum_{j \in A_t} s_j + \phi(A_t) \right) \leq \sum_{j \in \hat{I}} s_j + \sum_{t \in M} \phi(A_t) . \quad \square$$

651 Our next step is to describe joint partitions that can be found efficiently and upper bound the
652 different terms in the bound given by Lemma 3.7 (up to a constant factor). Finding such partitions
653 for the first two terms is quite straightforward.

655 **OBSERVATION 3.8.** *The joint partitions where all mediators are silent $\{I\} = \times_{i \in B} \{I\}$ obeys:*

$$656 \quad R(\{I\}) \geq \mu(I_0) \cdot v(I_0) .$$

658 **PROOF.**

$$659 \quad R(\{I\}) = \max_{i \in B}^{(2)} \left(\sum_{j \in I} \mu(j) \cdot v_{i,j} \right) \geq \max_{i \in B}^{(2)} \left(\sum_{j \in I_0} \mu(j) \cdot v_{i,j} \right) = \mu(I_0) \cdot v(I_0) . \quad \square$$

662 **OBSERVATION 3.9.** *The joint partitions $\mathcal{P}_S = \times_{t \in M} \mathcal{P}_t$ where every mediator reports all his informa-
663 tion obeys:*

$$664 \quad R(\mathcal{P}_S) = R(\{\{j\}_{j \in \hat{I}}\} \cup \{I \setminus \hat{I}\}) \geq \sum_{j \in \hat{I}} \mu(j) \cdot \max_{i \in B}^{(2)} v_{i,j} = \sum_{j \in \hat{I}} s_j .$$

667 It remains to find a joint partition that upper bounds, up to a constant factor, the third term in the
668 bound given by Lemma 3.7. If one knows the sets $\{A_t\}_{t \in M}$, then one can easily get such a partition
669 using the method of Ghosh et al. [11]. In this method, one partitions every set A_t into the parts
670 $\{A_t \cap H_i\}_{i=1}^t$ and sort these parts according to the value of $\sum_{j \in A_t \cap H_i} h_j$. Then, with probability $1/2$
671 every even part is united with the part that appears after it in the above order, and with probability
672 $1/2$ it is united with the part that appears before it in this order. It is not difficult to verify that if the
673 part of bidder i is not the first in the order, then with probability $1/2$ it is unified with the part that
674 appears before it in the order, in which case it contributes $\sum_{j \in A_t \cap H_i} h_j$ to the revenue. Hence, the
675 expected contribution to the revenue of the parts produced from A_t is at least $1/2 \cdot \phi(A_t)$.

676 Algorithm 2 can find a partition that is competitive against $\sum_{t \in M} \phi(A_t)$ without knowing the
677 sets $\{A_t\}_{t \in M}$. The algorithm uses the notation of a *cover*. We say that a set S_j of items is a cover of
678 an item $j \in I_t \cap H_i$ if $S_j \subseteq I_t \cap H_{i'}$ for some $i \neq i'$.

679 Notice that the definition of cover guarantees that a part containing both j and S_j contributes to
680 the revenue at least $\min\{h_j, \sum_{j' \in S_j} h_{j'}\}$. Using this observation, each iteration of Algorithm 2 can
681 be viewed as trying to extract revenue from element j . Additionally, observe that the partition \mathcal{P}
682 produced by Algorithm 2 can be presented as a joint partition since every part in it, except for $I \setminus \hat{I}$,
683 contains only items that belong to a single set I_t (for some mediator $t \in M$).

685 **OBSERVATION 3.10.** *Algorithm 2 can be implemented in polynomial time.*

ALGORITHM 2: Local Experts - Auxiliary Algorithm

```

1 Let  $I' \leftarrow \hat{I}$  and  $\mathcal{P} \leftarrow \{I \setminus \hat{I}\}$ .
2 while  $I' \neq \emptyset$  do
3   Let  $j$  be the element maximizing  $h_j$  in  $I'$ .
4   Find a cover  $S_j \subseteq I'$  of  $j$  obeying  $h_j \leq \sum_{j' \in S_j} h_{j'} \leq 2h_j$ , or maximizing  $\sum_{j' \in S_j} h_{j'}$  if no cover of  $j$ 
   makes this expression at least  $h_j$ .
5   Add the part  $S_j \cup \{j\}$  to  $\mathcal{P}$ , and remove the elements of  $S_j \cup \{j\}$  from  $I'$ .
6 return  $\mathcal{P}$ 

```

PROOF. One can find a cover S_j maximizing $\sum_{j' \in S_j} h_{j'}$ in line 4 of the algorithm by considering the set $I_t \cap H_{i'} \cap I'$ for every mediators t and bidder i' obeying $j \in I_t$ and $j \notin H_{i'}$. Moreover, if this cover is of size larger than $2h_j$, then by removing elements from this cover one by one the algorithm must find a cover S'_j obeying $h_j \leq \sum_{j' \in S'_j} h_{j'} \leq 2h_j$ because j is the element maximizing h_j in I' . \square

The following lemma relates the revenue of the set produced by Algorithm 2 to $\sum_{t \in M} \phi(A_t)$.

LEMMA 3.11. *No iteration of the loop of Algorithm 2 decreases the value of the expression $R(\mathcal{P}) + 1/3 \cdot \sum_{t \in M} \phi(A_t \cap I')$.*⁸

PROOF. Fix an arbitrary iteration. There are two cases to consider. First, assume $h_j \leq \sum_{j' \in S_j} h_{j'} \leq 2h_j$. In this case the increase in $R(\mathcal{P})$ during this iteration is:

$$\mu(S_j \cup \{j\}) \cdot v(S_j \cup \{j\}) \geq \min \left\{ h_j, \sum_{j' \in S_j} h_{j'} \right\} = h_j .$$

On the other hand, one can observe that, when removing an element j' from S , the value of $\phi(S)$ can decrease by at most $h_{j'}$. Hence, the decrease in $\sum_{t \in M} \phi(A_t \cap I')$ during this iteration can be upper bounded by:

$$h_j + \sum_{j' \in S_j} h_{j'} \leq 3h_j .$$

Consider now the case $\sum_{j' \in S_j} h_{j'} < h_j$. In this case the increase in $R(\mathcal{P})$ during the iteration is:

$$\mu(S_j \cup \{j\}) \cdot v(S_j \cup \{j\}) \geq \min \left\{ h_j, \sum_{j' \in S_j} h_{j'} \right\} = \sum_{j' \in S_j} h_{j'} .$$

If j does not belong to A_t for any mediator t , then by the above argument we can bound the decrease in $\sum_{t \in M} \phi(A_t \cap I')$ by $\sum_{j' \in S_j} h_{j'}$. Hence, assume from now on that there exists a mediator t' and a bidder i such that $j \in A_{t'} \cap H_i$. Let $i' \neq i$ be a bidder maximizing $\sum_{j' \in H_{i'} \cap A_{t'} \cap I'} h_{j'}$. Clearly, the removal of a single element from I' can decrease $\phi(A_{t'} \cap I')$ by no more than $\sum_{j' \in H_{i'} \cap A_{t'} \cap I'} h_{j'}$. Hence, the decrease in $\sum_{t \in M} \phi(A_t \cap I')$ during the iteration of the algorithm can be upper bounded by:

$$\sum_{j' \in H_{i'} \cap A_{t'} \cap I'} h_{j'} + \sum_{j' \in S_j} h_{j'} .$$

⁸Before the algorithm terminates \mathcal{P} is a partial partition in the sense that some items do not belong to any part in it. However, the definition of $R(\mathcal{P})$ naturally extends to such partial partitions.

On the other hand, $H_{I'} \cap A_{I'} \cap I'$ is a possible cover for j , and thus, by the optimality of S_j :

$$\sum_{j' \in H_{I'} \cap A_{I'} \cap I'} h_{j'} \leq \sum_{j' \in S_j} h_{j'} . \quad \square$$

COROLLARY 3.12. $R(\mathcal{P}_A) \geq 1/3 \cdot \sum_{t \in M} \phi(A_t)$, where \mathcal{P}_A is the partition produced by Algorithm 2.

PROOF. After the initialization step of Algorithm 2 we have:

$$R(\mathcal{P}) + 1/3 \cdot \sum_{t \in M} \phi(A_t \cap I') \geq 1/3 \cdot \sum_{t \in M} \phi(A_t) .$$

On the other hand, when the algorithm terminates:

$$R(\mathcal{P}) + 1/3 \cdot \sum_{t \in M} \phi(A_t \cap I') = R(\mathcal{P}_A)$$

because $I' = \emptyset$. The corollary now follows from Lemma 3.11. \square

We are now ready to prove Theorem 1.3.

THEOREM 1.3. *If mediators are local experts, there exists a polynomial time 5-approximation algorithm for DSP.*

PROOF. Consider an algorithm that outputs the best solution out of $\{I\}$, \mathcal{P}_S and \mathcal{P}_A . The following inequality shows that at least one of these joint partitions has a revenue of $R(\mathcal{P}^*)/5$:

$$R(\{I\}) + R(\mathcal{P}_S) + 3R(\mathcal{P}_A) \geq \sum_{j \in \hat{I}} s_j + \mu(I_0) \cdot v(I_0) + \sum_{t \in M} \phi(A_t) \geq R(\mathcal{P}^*) ,$$

where the first inequality holds by Observations 3.8 and 3.9 and Corollary 3.12; and the second inequality uses the upper bound on $R(\mathcal{P}^*)$ proved by Lemma 3.7. \square

4 THE STRATEGIC PROBLEM

This section explores the DSP problem from a strategic viewpoint, in which the auctioneer *cannot* control the signals produced by each mediator, and is, therefore, trying to solicit information from the mediators that would yield a maximal revenue in the auction. In other words, the objective of the auctioneer is to design a mechanism \mathcal{M} whose equilibria (i.e., the signals $\mathcal{P}'_1, \mathcal{P}'_2, \dots, \mathcal{P}'_m$ which are now chosen strategically by the mediators) induce maximum revenue. Our first contribution is the introduction of the *Shapley mechanism*, whose definition appears in Subsection 4.1. Subsection 4.1 also proves some interesting properties of the Shapley mechanism (Theorems 1.4 and 1.5). Subsection 4.2 studies the price of anarchy and price of stability of the DSP game induced by the Shapley mechanism (Theorems 1.6 and 1.7). Finally, Subsection 4.3 shows that the Shapley mechanism is the only possible mechanism if one insists on a few natural requirements.

4.1 The Shapley Mechanism

In this subsection we describe a mechanism \mathcal{S} which determines the payments to the mediators as a function of the reported signals. Our mechanism aims to incentivize mediators to report useful information, with the hope that global efficiency emerges despite selfish behavior of each mediator. In the remainder of the paper we study the mechanism \mathcal{S} and the game $\text{DSP}_{\mathcal{S}}$ it induces. For the sake of generality, we describe \mathcal{S} for a game generalizing DSP.

Consider a game \mathcal{G}_m of m players where each player t has a finite set A_t of possible strategies, one of which $\emptyset_t \in A_t$ is called the null strategy of t . The value of a strategy profile in the game \mathcal{G}_m is determined by a value function $v : A_1 \times A_2 \times \dots \times A_m \rightarrow \mathbb{R}$. A mechanism $M = (\Pi_1, \Pi_2, \dots, \Pi_m)$ for \mathcal{G}_m is a set of payments rules. In other words, if the players choose strategies $a_1 \in A_1, a_2 \in$

785 $A_2, \dots, a_m \in A_m$, then the payment to player t under mechanism M is $\Pi_t(v, a_1, a_2, \dots, a_m)$. Notice
 786 that DSP fits the definition of \mathcal{G}_m when $A_t = \Omega(\mathcal{P}_t)$ is the set of partitions that t can report for
 787 every mediator t , and \emptyset_t is the silence strategy $\{I\}$. The appropriate value function v for DSP is
 788 the function $R(\times_{t=1}^m \mathcal{P}_t')$, where $\mathcal{P}_t' \in A_t$ is the strategy of mediator t . In other words, the value
 789 function v of a DSP game is equal to the revenue of the auctioneer.

790 Given a strategy profile $a = (a_1, a_2, \dots, a_m)$, and subset $J \subseteq [m]$ of players, we write a_J to denote
 791 a strategy profiles where the players of J play their strategy in a , and the other players play their
 792 null strategies. We abuse notation and denote by \emptyset the strategy profile a_\emptyset where all players play
 793 their null strategies. Additionally, we write (a'_t, a_{-t}) to denote a strategy profile where player t
 794 plays a'_t and the rest of the players follow the strategy profile a . The mechanism \mathcal{S} we propose
 795 distributes the increase in the value of the game (compared to $v(\emptyset)$) among the players according to
 796 their Shapley value: it pays each player his expected marginal contribution to the value according
 797 to a uniformly random ordering of the m player. Formally, the payoff for player t given a strategy
 798 profile a is

$$799 \quad \Pi_t(a) = \frac{1}{m!} \cdot \sum_{\sigma \in S_m} [v(a_{\{\sigma^{-1}(j)|1 \leq j \leq \sigma(t)\}}}) - v(a_{\{\sigma^{-1}(j)|1 \leq j < \sigma(t)\}})] , \quad (1)$$

800
801
802 which can alternatively be written as

$$803 \quad \Pi_t(a) = \sum_{J \subseteq [m] \setminus \{t\}} \gamma_J (v(a_{J \cup \{t\}}) - v(a_J)) , \quad (2)$$

804
805
806
807 where $\gamma_J = \frac{|J|!(m-|J|-1)!}{m!}$ is the probability that the players of J appear before player t when
 808 the players are ordered according to a uniformly random permutation $\sigma \in_R S_m$. We use both
 809 definitions (1) and (2) interchangeably, as each one is more convenient in some cases than the other.
 810 We remark that the above payoffs can be implemented efficiently.⁹

811
812 Clearly, the mechanism \mathcal{S} is anonymous (symmetric). The main feature of the Shapley mechanism
 813 is that it is efficient. In other words, the sum of the payoffs is exactly equal to the total increase in
 814 value (in the case of DSP, the surplus revenue of the auctioneer compared to the initial state).¹⁰

815
816 **PROPOSITION 4.1 (EFFICIENCY PROPERTY).** *For every strategy profile $a = (a_1, a_2, \dots, a_m)$,*

$$817 \quad v(a) - v(\emptyset) = \sum_{t=1}^m \Pi_t(a) .$$

818
819
820 **PROOF.** Recall that the payoff of mediator t is:

$$821 \quad \frac{1}{m!} \cdot \sum_{\sigma \in S_m} [v(a_{\{\sigma^{-1}(j)|1 \leq j \leq \sigma(t)\}}}) - v(a_{\{\sigma^{-1}(j)|1 \leq j < \sigma(t)\}})] .$$

822
823
824
825
826
827 ⁹Assuming value queries, we can calculate a payoff for every player by drawing a random permutation σ and paying
 828 $v(a_{\{\sigma^{-1}(j)|1 \leq j \leq \sigma(t)\}}}) - v(a_{\{\sigma^{-1}(j)|1 \leq j < \sigma(t)\}})$ for each mediator t . Clearly this procedure produces the payoffs of our
 829 mechanism *in expectation*. Alternatively, the *expected* payoff of each player can be approximated using sampling.

830 ¹⁰One natural alternative for the Shapley mechanism is a VCG-based mechanism. The main disadvantage of this alternative
 831 mechanism is that it is not necessarily efficient. In fact, one can easily design instances where a VCG-based mechanism
 832 induces a total payoff which is significantly larger than the increase in the value.

Summing over all mediators, we get:

$$\begin{aligned}
\sum_{t=1}^m \Pi_t(\mathcal{P}'_t, \mathcal{P}'_{-t}) &= \sum_{t=1}^m \left\{ \frac{1}{m!} \cdot \sum_{\sigma \in \mathcal{S}_m} [v(a_{\{\sigma^{-1}(j)|1 \leq j \leq \sigma(t)\}}) - v(a_{\{\sigma^{-1}(j)|1 \leq j < \sigma(t)\}})] \right\} \\
&= \frac{1}{m!} \cdot \sum_{\sigma \in \mathcal{S}_m} \sum_{t=1}^m [v(a_{\{\sigma^{-1}(j)|1 \leq j \leq \sigma(t)\}}) - v(a_{\{\sigma^{-1}(j)|1 \leq j < \sigma(t)\}})] \\
&= \frac{1}{m!} \cdot \sum_{\sigma \in \mathcal{S}_m} [v(a_{\{\sigma^{-1}(j)|1 \leq j \leq m\}}) - v(a_{\emptyset})] = v(a) - v(\emptyset) . \quad \square
\end{aligned}$$

Proposition 4.1 implies the following theorem. Notice that Theorem 1.5 is in fact a restriction of this theorem to the game $\text{DSP}_{\mathcal{S}}$.

THEOREM 4.2. *For every Nash equilibrium a , $v(a) \geq v(\emptyset)$.*

PROOF. A player always has the option of playing his null strategy, which results in a zero payoff for him. Thus, the payoff of a player in a Nash equilibrium can never be negative. Hence, by Proposition 4.1: $v(a) \geq v(\emptyset) + \sum_{i=1}^m \Pi_i(a) \geq v(\emptyset)$. \square

Next, let us prove Theorem 1.4. For convenience, we restate it below.

THEOREM 1.4. *Let \mathcal{G}_m be a non-cooperative m -player game in which the payoff of each player is set according to \mathcal{S} . Then \mathcal{G}_m admits a pure Nash equilibrium. Moreover, best response dynamics are guaranteed to converge to such an equilibrium.*

PROOF. We prove the theorem by showing that \mathcal{G}_m is an exact potential game, which in turn implies all the conclusions of the theorem. Recall that an exact potential game is a game for which there exists a potential function $\Phi: A_1 \times A_2 \times \dots \times A_t \rightarrow \mathbb{R}$ such that every strategy profile a and possible deviation $a'_t \in A_t$ of a player t obey:

$$\Pi_t(a'_t, a_{-t}) - \Pi_t(a) = \Phi(a'_t, a_{-t}) - \Phi(a) . \quad (3)$$

In our case the potential function is: $\Phi(a) = \sum_{J \subseteq [m]} \beta_J \cdot v(a_J)$, where $\beta_J = \frac{(|J|-1)!(m-|J|)!}{m!}$. Let us prove that this function obeys (3). It is useful to denote by a' the strategy profile (a'_t, a_{-t}) . By definition:

$$\Pi_t(a') - \Pi_t(a) = \sum_{J \subseteq [m] \setminus \{i\}} \gamma_J [v(a_{J \cup \{i\}}) - v(a_J)] - \sum_{J \subseteq [m] \setminus \{i\}} \gamma_J [v(a'_{J \cup \{i\}}) - v(a'_J)] . \quad (4)$$

For $J \subseteq [m] \setminus \{i\}$, we have $a_J = a'_J$. Plugging this observation into (4), and rearranging, we get:

$$\Pi_t(a') - \Pi_t(a) = \sum_{J \subseteq [m] \setminus \{i\}} \gamma_J [v(a_{J \cup \{i\}}) - v(a'_{J \cup \{i\}})] . \quad (5)$$

For every J containing i we get: $\gamma_{J \setminus \{i\}} = \beta_J$. Using this observation and the previous observation that $a_J = a'_J$ for $J \subseteq [m] \setminus \{i\}$, (5) can be replaced by:

$$\Pi_t(a') - \Pi_t(a) = \sum_{J \subseteq [m]} \beta_J (v(a'_J) - v(a_J)) = \Phi(a') - \Phi(a) . \quad \square$$

Before concluding this section, a few remarks regarding the use of \mathcal{S} to DSP are in order:

- (1) The reader may wonder why the auctioneer cannot impose on the mediators any desired outcome $\times_{t \in M} \mathcal{P}'_t$ by offering mediator t a negligible payment if he signals \mathcal{P}'_t , and no payment otherwise. However, implementing such a mechanism requires the auctioneer to know the

information sets \mathcal{P}_t of each mediator *in advance*. In contrast, our mechanism requires access only to the outputs of the mediators.

- (2) Proposition 4.1 implies that the auctioneer distributes the entire surplus among the mediators, which seems to defeat the purpose of the mechanism. However, in the target application she can scale the revenue by a factor $\alpha \in (0, 1]$ and only distribute the corresponding fraction of the surplus. As all of our results are invariant under scaling, this trick can be applied in a black box fashion. Thus, we assume, without loss of generality, $\alpha = 1$.
- (3) We assume mediators never report a signal which is inconsistent with the true identity of the sold element j_R . The main justification for this assumption is that the mediators' signals must be consistent with one another (as they refer to a single element j_R). Thus, given that sufficiently many mediators are honest, "cheaters" can be easily detected.
- (4) Note that for a particular ordering of the mediators $\sigma \in \mathcal{S}_m$ and a particular joint strategy profile, the marginal payoff of a mediator may be negative (if she is out of luck and contributes negatively to the revenue according to σ). However, we stress that the expected value (over σ) of each mediator is never negative in any equilibrium strategy (by Theorem 4.2). Since in realistic applications the process is assumed to be repeated over time, the probability that a mediator has overall negative payoff is negligible.

4.2 The Price of Anarchy and Price of Stability of the Shapley Mechanism

In this section we analyze the PoS and PoA of the DSP_S game. First, we note that the proof of Theorem 1.2 in Section 3.2.1 shows that $R(\{I\}) \geq \text{OPT}/\min\{n, k-1\}$ whenever $k \geq 2$. Together with Theorem 1.5, we get:

THEOREM 1.6. *For $k \geq 2$, the price of anarchy of $\text{DSP}_S(n, k, m)$ is no more than $\min\{k-1, n\}$.*

Remark: The above statement of Theorem 1.7 uses a somewhat different notation than its original statement in Section 1.2, but both statements are equivalent. The same is true for the statement of Theorem 1.7 below.

Naturally, the upper bound given by Theorem 1.6 applies also to the price of stability of DSP_S . The rest of this section is devoted to proving Theorem 1.7, which shows that Theorem 1.6 is asymptotically tight.

THEOREM 1.7. *For every $n \geq 1$, there is a $\text{DSP}_S(3n+1, n+2, 2)$ game for which the price of stability is at least n . Moreover, all the mediators in this game are local experts.*

We begin the proof of Theorem 1.7 by describing the $\text{DSP}_S(3n+1, n+2, 2)$ game whose price of stability we bound. For ease of notation, let us denote this game by DSP_n .

items: The $3n+1$ items of DSP_n all have equal probabilities. It is convenient to denote them by $\{a_\ell\}_{\ell=1}^n$, $\{b_\ell\}_{\ell=1}^n$, $\{c_\ell\}_{\ell=1}^n$ and d .

bidders: The $n+2$ bidders of DSP_n can be partitioned into 3 types. One bidder, denoted by i_G , has a bid of ε for the items of $\{b_\ell\}_{\ell=1}^n$ and a bid of 1 for all other items, where $\varepsilon \in (0, 1)$ is a value that will be defined later. One bidder, denoted by i_O has a bid of 1 for item d and a bid of 0 for all other items. Finally, the other n bidders are denoted by $\{i_\ell\}_{\ell=1}^n$. Each bidder i_ℓ has a bid of 1 for item b_ℓ and a bid of 0 for all other items.

mediators: The two mediators of DSP_n are denoted by t_1 and t_2 . Both mediators are local experts whose partitions are defined by the sets $\mathcal{I}_1 = \{a_\ell, b_\ell\}_{\ell=1}^n$ and $\mathcal{I}_2 = \{b_\ell, c_\ell\}_{\ell=1}^n$, respectively.

A graphical sketch of DSP_n is given by Figure 2. Intuitively, getting a high revenue in DSP_n requires pairing b items with a or c items. Unfortunately, one mediator can pair the b items with a items, and the other mediator can pair them with c items, thus, creating "tension" between the

$$\begin{array}{r}
932 \\
933 \\
934 \\
935 \\
936 \\
937 \\
938 \\
939 \\
940 \\
941 \\
942 \\
943 \\
944 \\
945 \\
946 \\
947 \\
948 \\
949 \\
950 \\
951 \\
952 \\
953 \\
954 \\
955 \\
956 \\
957 \\
958 \\
959 \\
960 \\
961 \\
962 \\
963 \\
964 \\
965 \\
966 \\
967 \\
968 \\
969 \\
970 \\
971 \\
972 \\
973 \\
974 \\
975 \\
976 \\
977 \\
978 \\
979 \\
980
\end{array}$$

a items
 b items
 c items
 d

$$V = \begin{pmatrix}
1 & 1 & \dots & 1 & \varepsilon & \varepsilon & \dots & \varepsilon & 1 & 1 & \dots & 1 & 1 \\
0 & 0 & \dots & 0 & 0 & 0 & \dots & 0 & 0 & 0 & \dots & 0 & 1 \\
0 & 0 & \dots & 0 & 1 & 0 & \dots & 0 & 0 & 0 & \dots & 0 & 0 \\
0 & 0 & \dots & 0 & 0 & 1 & \dots & 0 & 0 & 0 & \dots & 0 & 0 \\
\vdots & \vdots & & \vdots & \vdots & \vdots & & \vdots & \vdots & \vdots & & \vdots & 0 \\
0 & 0 & \dots & 0 & 0 & 0 & \dots & 1 & 0 & 0 & \dots & 0 & 0
\end{pmatrix} \begin{array}{l} \leftarrow i_G \\ \leftarrow i_O \\ \leftarrow i_1 \\ \leftarrow i_2 \\ \vdots \\ \leftarrow i_n \end{array}$$

$$\mathcal{P}_1 = \square \quad \square \quad \dots \quad \square \quad \square \quad \square \quad \dots \quad \square$$

$$\mathcal{P}_2 = \quad \quad \square \quad \square \quad \dots \quad \square \quad \square \quad \square \quad \dots \quad \square$$

Fig. 2. A graphical representation of DSP_n . \mathcal{P}_1 and \mathcal{P}_2 are the partitions of mediators t_1 and t_2 , respectively.

mediators. The main idea of the proof is to show that both mediators are incentivized to report partitions pairing the b items, which results in a joint partition isolating all the b items.

The following observation simplifies many of our proofs.

OBSERVATION 4.3. *The contribution $\mu(S) \cdot v(S)$ of a part S to the revenue $R(\mathcal{P})$ of a partition $\mathcal{P} \ni S$ is:*

- $(3n + 1)^{-1}$ if S contains d .
- $(3n + 1)^{-1}$ if S contains an item of $\{b_\ell\}_{\ell=1}^n$ and $|S| \geq 2$.
- $\varepsilon/(3n + 1)$ if S contains only a single item, and this item belongs to $\{b_\ell\}_{\ell=1}^n$.
- 0 otherwise.

PROOF. Notice that the only bidder that values more than one item is i_G , and thus, i_G is the single bidder that can have a bid larger than $1/|S|$ for S . Hence, $v(S)$ is upper bounded by $1/|S|$, and one can bound the contribution of S by:

$$\mu(S) \cdot v(S) \leq \frac{|S|}{3n + 1} \cdot \frac{1}{|S|} = (3n + 1)^{-1} .$$

If the part S contains the item d , then both bidders i_G and i_O have bids of at least $1/|S|$ for it. Hence, the above upper bound on the contribution of S is in fact tight. Similarly when S contains an item b_ℓ and some other item j , it has a contribution of $(3n + 1)^{-1}$ since two bidders have a bid of at least $1/|S|$ for it, bidder i_ℓ and a second bidder that depends on j :

- If $j = b_{\ell'}$ for some $\ell' \neq \ell$, bidder $i_{\ell'}$.
- If $j = d$, $j = a_{\ell'}$ or $j = c_{\ell'}$, bidder i_G .

If the part S contains only a single item b_ℓ , then it gets non-zero bids only from two bidders: a bid of 1 from i_ℓ and a bid of ε from bidder i_G . Hence, $\mu(S) \cdot v(S) = (3n + 1)^{-1} \cdot \varepsilon$. Finally, if S does not fall into any of the cases considered above, then it must contain only items of $\{a_\ell, c_\ell\}_{\ell=1}^n$. Such a part receives a non-zero bid only from bidder i_G , and thus, $v(S) = 0$. \square

The next lemma gives a lower bound on the optimal revenue of DSP_n .

LEMMA 4.4. *The optimal revenue of DSP_n is at least $(n + 1)/(3n + 1)$.*

PROOF. Consider a scenario in which t_1 pairs the b items with a items, and t_2 is silent. Formally, the two mediators use the following strategies $O_1 = \{\{a_\ell, b_\ell\}_{\ell=1}^n \cup \{c_\ell\}_{\ell=1}^n \cup \{d\}\}$ and $O_2 = \{I\}$, respectively. Observe that these are indeed feasible strategies for t_1 and t_2 , respectively. By

981 Observation 4.3, the revenue of DSP_n given these strategies is:

$$\begin{aligned}
 982 \quad R(O_1 \times O_2) &= R(O_1) = \sum_{\ell=1}^n \mu(\{a_\ell, b_\ell\}) \cdot v(\{a_\ell, b_\ell\}) + \mu(\{c_\ell\}_{\ell=1}^n \cup \{d\}) \cdot v(\{c_\ell\}_{\ell=1}^n \cup \{d\}) \\
 983 \quad &= \sum_{\ell=1}^n \frac{1}{3n+1} + \frac{1}{3n+1} = \frac{n+1}{3n+1}. \quad \square \\
 984 \quad & \\
 985 \quad & \\
 986 \quad & \\
 987 \quad &
 \end{aligned}$$

988 Our next objective is to get an upper bound on the revenue of any Nash equilibrium of DSP_n .
 989 We say that an item a_ℓ is *redundant* in a strategy \mathcal{P}'_1 of t_1 if the part $S \in \mathcal{P}'_1$ containing a_ℓ obeys
 990 one of the following:

- 991 • S contains the item d or an item $a_{\ell'}$ for some $\ell' \neq \ell$.
- 992 • S contains no items of $\{b_\ell\}_{\ell=1}^n$.

993 The next lemma shows that a redundant item is indeed redundant in the sense that removing it
 994 does not change the contribution to the revenue of parts containing it.

995 LEMMA 4.5. *If a_ℓ is redundant in a strategy \mathcal{P}'_1 of t_1 and $S \in \mathcal{P}'_1$ is the part containing a_ℓ , then*
 996 $\mu(S) \cdot v(S) = \mu(S \setminus \{a_\ell\}) \cdot v(S \setminus \{a_\ell\})$. *Moreover, for every possible strategy \mathcal{P}'_2 of t_2 , if $S' \in \mathcal{P}'_1 \times \mathcal{P}'_2$*
 997 *is the part of $\mathcal{P}'_1 \times \mathcal{P}'_2$ containing a_ℓ , then $\mu(S') \cdot v(S') = \mu(S' \setminus \{a_\ell\}) \cdot v(S' \setminus \{a_\ell\})$*
 998

999 PROOF. We prove the second part of the lemma. The first part follows from it since one possible
 1000 choice of \mathcal{P}'_2 is $\{I\}$. Since a_ℓ is redundant in \mathcal{P}_1 , one of the following three cases must hold:

- 1001 • The first case is when $d \in S$. Observe that d must share a part with a_ℓ in \mathcal{P}'_2 , and thus, d belongs
 1002 also to S' . Hence, by Observation 4.3, $\mu(S') \cdot v(S') = (3n+1)^{-1} = \mu(S' \setminus \{a_\ell\}) \cdot v(S' \setminus \{a_\ell\})$.
- 1003 • The second case is when $(\{b_\ell\}_{\ell=1}^n \cup \{d\}) \cap S = \emptyset$. Since S' is a subset of S , we get also
 1004 $(\{b_\ell\}_{\ell=1}^n \cup \{d\}) \cap S' = \emptyset$, and thus, by Observation 4.3, $\mu(S') \cdot v(S') = 0 = \mu(S' \setminus \{a_\ell\}) \cdot v(S' \setminus \{a_\ell\})$.
- 1005 • The third case is when there exists $\ell' \neq \ell$ for which $a_{\ell'} \in S$. Since a_ℓ and $a_{\ell'}$ must share a part
 1006 in \mathcal{P}_2 , they both belong also to S' . Using Observation 4.3 one can verify that removing a_ℓ
 1007 from a set S' containing $a_{\ell'}$ can never change $\mu(S') \cdot v(S')$. \square
 1008
 1009

1010 A strategy \mathcal{P}'_1 of t_1 is called *a-helped* if every part in it that contains an item of $\{b_\ell\}_{\ell=1}^n$ contains
 1011 also an item of $\{a_\ell\}_{\ell=1}^n$. The next lemma shows that every strategy of t_1 is dominated by an *a-helped*
 1012 one.

1013 LEMMA 4.6. *If \mathcal{P}'_1 is a strategy of t_1 , then there exists an *a-helped* strategy \mathcal{P}''_1 of t_1 such that:*
 1014 $\Pi_1(\mathcal{P}''_1, \mathcal{P}'_2) \geq \Pi_1(\mathcal{P}'_1, \mathcal{P}'_2)$ *for every strategy \mathcal{P}'_2 of t_2 . Moreover, for every pair of items from $\{b_\ell\}_{\ell=1}^n \cup$*
 1015 *$\{d\}$, \mathcal{P}''_1 separates this pair (i.e., each item of the pair appear in a different part of \mathcal{P}''_1) if and only if*
 1016 *\mathcal{P}'_1 does.*
 1017

1018 PROOF. Let $D(\mathcal{P}'_1)$ be the number of parts in \mathcal{P}'_1 that contain an item of $\{b_\ell\}_{\ell=1}^n$ but no items of
 1019 $\{a_\ell\}_{\ell=1}^n$. We prove the lemma by induction on $D(\mathcal{P}'_1)$. If $D(\mathcal{P}'_1) = 0$, then \mathcal{P}'_1 is *a-helped* and we are
 1020 done. It remains to prove the lemma for $D(\mathcal{P}'_1) > 0$ assuming that it holds for every strategy $\hat{\mathcal{P}}'_1$ for
 1021 which $D(\hat{\mathcal{P}}'_1) > D(\mathcal{P}'_1)$.

1022 Since $D(\mathcal{P}'_1) > 0$ and the number of $\{a_\ell\}_{\ell=1}^n$ items and $\{b_\ell\}_{\ell=1}^n$ items is equal, there must be in \mathcal{P}'_1
 1023 either a part that contains at least two items of $\{a_\ell\}_{\ell=1}^n$ or a part that contains an item of $\{a_\ell\}_{\ell=1}^n$ and
 1024 no items of $\{b_\ell\}_{\ell=1}^n$. Both options imply the existence of a redundant item $a_{\ell'}$ in \mathcal{P}'_1 . Additionally,
 1025 let S be one of the parts of \mathcal{P}'_1 counted by $D(\mathcal{P}'_1)$, i.e., a part that contains an item of $\{b_\ell\}_{\ell=1}^n$ but no
 1026 items of $\{a_\ell\}_{\ell=1}^n$.

1027 Consider the strategy $\tilde{\mathcal{P}}'_1$ obtained from \mathcal{P}'_1 by moving $a_{\ell'}$ to the part S . By Lemma 4.5 the
 1028 contribution to the revenue of the part containing $a_{\ell'}$ in \mathcal{P}'_1 does not decrease following the removal
 1029

of $a_{\ell'}$ from it. On the other hand, Observation 4.3 implies that adding items to a part can only increase its contribution. Hence, the contribution of S to the revenue is at least as large in $\tilde{\mathcal{P}}'_1$ as in \mathcal{P}'_1 . Combining both arguments, we get: $R(\tilde{\mathcal{P}}'_1) \geq R(\mathcal{P}'_1)$. An analogous argument can be used to show also that $R(\tilde{\mathcal{P}}'_1 \times \mathcal{P}'_2) \geq R(\mathcal{P}'_1 \times \mathcal{P}'_2)$ for every strategy \mathcal{P}'_2 of t_2 . Thus, by definition, $\Pi_1(\tilde{\mathcal{P}}'_1, \mathcal{P}'_2) \geq \Pi_1(\mathcal{P}'_1, \mathcal{P}'_2)$.

Observe that the construction of $\tilde{\mathcal{P}}'_1$ from \mathcal{P}'_1 guarantees that $D(\tilde{\mathcal{P}}'_1) = D(\mathcal{P}'_1) - 1$. Hence, by applying the induction hypothesis to $\tilde{\mathcal{P}}'_1$ we get a strategy \mathcal{P}''_1 which, for every strategy \mathcal{P}'_2 of t_2 , obeys the inequality:

$$\Pi_1(\mathcal{P}''_1, \mathcal{P}'_2) \geq \Pi_1(\tilde{\mathcal{P}}'_1, \mathcal{P}'_2) \geq \Pi_1(\mathcal{P}'_1, \mathcal{P}'_2)$$

Moreover, the construction of $\tilde{\mathcal{P}}'_1$ from \mathcal{P}'_1 does not move any items of $\{b_{\ell}\}_{\ell=1}^n \cup \{d\}$, thus, $\tilde{\mathcal{P}}'_1$ separates pairs from this set if and only if \mathcal{P}'_1 does. The lemma now follows since the induction hypothesis guarantees that \mathcal{P}''_1 separates pairs from the above set if and only if $\tilde{\mathcal{P}}'_1$ does. \square

Using the previous lemma we can prove an important property of not strictly dominated strategies of t_1 .

LEMMA 4.7. *If \mathcal{P}'_1 is not a strictly dominated strategy of t_1 , then \mathcal{P}'_1 isolates the items of $\{b_{\ell}\}_{\ell=1}^n \cup \{d\}$ from each other.*

PROOF. Assume towards a contradiction that the lemma does not hold, and let \mathcal{P}'_1 be a counter example. In other words, \mathcal{P}'_1 does not isolate the items of $\{b_{\ell}\}_{\ell=1}^n \cup \{d\}$ from each other, and yet there exists a strategy \mathcal{P}'_2 of t_2 such that every strategy $\hat{\mathcal{P}}'_1$ of t_1 obeys $\Pi_1(\hat{\mathcal{P}}'_1, \mathcal{P}'_2) \leq \Pi_1(\mathcal{P}'_1, \mathcal{P}'_2)$. By Lemma 4.6 we may assume that \mathcal{P}'_1 is an a -helped strategy (otherwise, we can replace \mathcal{P}'_1 with the strategy whose existence is guaranteed by this lemma).

Let $b_{\ell'}$ be an item that is not isolated by \mathcal{P}'_1 from some other item of $\{b_{\ell}\}_{\ell=1}^n \cup \{d\}$. The existence of $b_{\ell'}$ implies the existence of a redundant item $a_{\ell''}$ in \mathcal{P}'_1 because one of the following must be true:

- $b_{\ell'}$ shares a part in \mathcal{P}'_1 with another item of $\{b_{\ell}\}_{\ell=1}^n$. Since the number of $\{a_{\ell}\}_{\ell=1}^n$ items is equal to the number of $\{b_{\ell}\}_{\ell=1}^n$ items, there must be either a part of \mathcal{P}'_1 that contains an item of $\{a_{\ell}\}_{\ell=1}^n$ but no items of $\{b_{\ell}\}_{\ell=1}^n$ or a part of \mathcal{P}'_1 that contains two items of $\{a_{\ell}\}_{\ell=1}^n$.
- $b_{\ell'}$ shares a part in \mathcal{P}'_1 with the item d . Since \mathcal{P}'_1 is a -helped, this part must contain also an item of $\{a_{\ell}\}_{\ell=1}^n$ (which is redundant).

Consider the strategy $\tilde{\mathcal{P}}'_1$ obtained from \mathcal{P}'_1 by removing the items $a_{\ell''}$ and $b_{\ell'}$ from their original parts and placing them together in a new part. Let us analyze $R(\tilde{\mathcal{P}}'_1) + R(\tilde{\mathcal{P}}'_1 \times \mathcal{P}'_2)$. Since $a_{\ell''}$ is redundant, its removal from its original part $\tilde{\mathcal{P}}'_1$ does not decrease the contribution of this part to either revenue by Lemma 4.5. Additionally, the removal of $a_{\ell''}$ leaves $b_{\ell'}$ either sharing a part in \mathcal{P}'_1 with d or with another item of $\{a_{\ell}\}_{\ell=1}^n$ and another item of $\{b_{\ell}\}_{\ell=1}^n$. In both cases, the removal of $b_{\ell'}$ does not affect the contribution of its part to $R(\tilde{\mathcal{P}}'_1)$ by Observation 4.3. On the other hand, the removal of $b_{\ell'}$ can decrease the contribution of its part to the revenue $R(\tilde{\mathcal{P}}'_1 \times \mathcal{P}'_2)$. However, Observation 4.3 guarantees that this decrease is at most $(3n + 1)^{-1}$. Finally, the contribution of the new part $\{a_{\ell''}, b_{\ell'}\}$ to $R(\tilde{\mathcal{P}}'_1)$ is $(3n + 1)^{-1}$. Combining all these observations, we get:

$$\begin{aligned} R(\tilde{\mathcal{P}}'_1) + R(\tilde{\mathcal{P}}'_1 \times \mathcal{P}'_2) &\geq R(\mathcal{P}'_1) + R(\mathcal{P}'_1 \times \mathcal{P}'_2) - (3n + 1)^{-1} + (3n + 1)^{-1} + A \\ &= R(\mathcal{P}'_1) + R(\mathcal{P}'_1 \times \mathcal{P}'_2) + A \end{aligned} \quad (6)$$

where A is the contribution of parts of $\tilde{\mathcal{P}}'_1 \times \mathcal{P}'_2$ that are subsets of $\{a_{\ell''}, b_{\ell'}\}$ to $R(\tilde{\mathcal{P}}'_1 \times \mathcal{P}'_2)$. To get a contradiction it is enough to show that $A > 0$, i.e., that at least one of these parts has a positive contribution to $R(\tilde{\mathcal{P}}'_1 \times \mathcal{P}'_2)$. There are two cases to consider:

- If $b_{\ell'}$ shares a part with $a_{\ell''}$ in \mathcal{P}'_2 , then the part $\{a_{\ell''}, b_{\ell'}\}$ appears in $\tilde{\mathcal{P}}'_1 \times \mathcal{P}'_2$ and contributes $(3n+1)^{-1}$ to $R(\tilde{\mathcal{P}}'_1 \times \mathcal{P}'_2)$.
- If \mathcal{P}'_2 separates the item $b_{\ell'}$ from $a_{\ell''}$, then the part $\{b_{\ell'}\}$ appears in $\tilde{\mathcal{P}}'_1 \times \mathcal{P}'_2$ and contributes $\varepsilon/(3n+1)$ to $R(\tilde{\mathcal{P}}'_1 \times \mathcal{P}'_2)$. \square

COROLLARY 4.8. *If \mathcal{P}'_1 and \mathcal{P}'_2 are strategies for t_1 and t_2 that form a Nash equilibrium, then both \mathcal{P}'_1 and \mathcal{P}'_2 isolate the items of $\{b_{\ell}\}_{\ell=1}^n \cup \{d\}$ from each other.*

PROOF. Since \mathcal{P}'_1 is a part of a Nash equilibrium, it is not strictly dominated. Hence, by Lemma 4.7, it must isolate the items of $\{b_{\ell}\}_{\ell=1}^n \cup \{d\}$ from each other. The corollary holds also for \mathcal{P}'_2 by symmetry. \square

We are now ready to analyze the price of stability of DSP_n .

THEOREM 4.9. *The price of stability of DSP_n is at least $(n+1)/(n\varepsilon+1)$. Hence, for $\varepsilon = 1/n^2$, the price of stability of DSP_n is at least n .*

PROOF. Consider an arbitrary Nash equilibrium $(\mathcal{P}'_1, \mathcal{P}'_2)$ of DSP_n . By Corollary 4.8, both \mathcal{P}'_1 and \mathcal{P}'_2 must isolate the items of $\{b_{\ell}\}_{\ell=1}^n \cup \{d\}$ from each other. However, every other item of I must share a part with d in at least one of these partitions, and thus, every item of $\{b_{\ell}\}_{\ell=1}^n$ has a singleton part in $\mathcal{P}'_1 \times \mathcal{P}'_2$. Hence,

$$R(\mathcal{P}'_1 \times \mathcal{P}'_2) = \frac{n\varepsilon+1}{3n+1}.$$

Combining the last observation with Lemma 4.4, we get that the price of stability of DSP_n is at least:

$$\frac{(n+1)/(3n+1)}{(n\varepsilon+1)/(3n+1)} = \frac{n+1}{n\varepsilon+1}. \quad \square$$

Note that Theorem 1.7 follows immediately from Theorem 4.9.

4.3 The Limitations of \mathcal{S} are Inevitable

Theorem 1.7 asserts that the revenue of the best equilibrium can be about n times worse than the optimal revenue. This discouraging result raises the question of whether alternative payment rules can improve the revenue guarantees of the auctioneer. Unfortunately, Shapley's uniqueness theorem answers this question negatively, assuming one requires the mechanism to have some natural properties.

We consider a family of games of the type considered in Subsection 4.1. Formally, let \mathcal{F}_m denote a family of m -player games where each player t has the same finite set A_t of possible strategies in all the games, one of which $\varnothing_t \in A_t$ is called the null strategy of t . Each game in the family is determined by an arbitrary value function $v : A_1 \times A_2 \times \dots \times A_m \rightarrow \mathbb{R}$, and each possible such value function induces a game in \mathcal{F}_m . Recall that a mechanism $M = (\Pi_1, \Pi_2, \dots, \Pi_m)$ is a set of payments rules. In other words, if the players choose strategies $a_1 \in A_1, a_2 \in A_2, \dots, a_m \in A_m$, then the payment of player t under mechanism M is $\Pi_t(v, a_1, a_2, \dots, a_m)$.¹¹

THEOREM 4.10 (UNIQUENESS OF SHAPLEY MECHANISM, CF. [19]). *Let \mathcal{F}_m be a family of games as described above. Then, the Shapley value mechanism \mathcal{S} is the only mechanism satisfying the following axioms:*

- (1) (Normalization) For every player t , $\Pi_t(a) = 0$ whenever $a_t = \varnothing_t$.

¹¹Shapley's theorem is stated for cooperative games where players in the coalition can reallocate their payments within the coalition. In our non-cooperative setup, we assume side payments can be only introduced by the mechanism.

(2) (Anonymity) If $\mathcal{G}_m \in \mathcal{F}_m$ is a game with a strategy profile a^* such that $v(a) = v(\emptyset)$ (recall that \emptyset denotes the strategy profile $(\emptyset_1, \emptyset_2, \dots, \emptyset_m)$) for every strategy profile $a \neq a^*$, then for every strategy profile a and player t :

$$\Pi_t(a) = \begin{cases} 0 & \text{if } a_t = \emptyset_t, \\ \frac{v(a) - v(\emptyset)}{|\{t \in [m] \mid a_t \neq \emptyset_t\}|} & \text{otherwise.} \end{cases}$$

(3) (Additivity) If $\mathcal{G}_m, \mathcal{H}_m \in \mathcal{F}_m$ are two games with value functions v_g and v_h , then $\Pi_t^{v_g + v_h}(a) = \Pi_t^{v_g}(a) + \Pi_t^{v_h}(a)$ (where $\Pi_t^v(a)$ stands for the payment of player t given strategy profile a in the game defined by the value function v).

We note that the second axiom (“Anonymity”) is ubiquitous in market design, and is typically enforced by market regulations prohibiting discrimination among clients. Intuitively, this axiom says that whenever there is only one strategy profile a^* which produces a value other than $v(\emptyset)$, then, when a^* is played, the mechanism is required to *equally* distribute the surplus $v(a^*) - v(\emptyset)$ among the participants playing a non-null strategy in a^* . Observe that violating this axiom would require private contracts with (at least some of) the players (mediators). Implementing such contracts defeats one of the main purposes of our mechanism, namely that it can be easily implemented in a dynamic environment having an unstable mediators population.

The above theorem was originally proved in a *cooperative* setting (where players may either join a coalition or not), under slightly different axioms. The “anonymity” axiom of Theorem 4.10 replaces the fairness and efficiency axioms of the original theorem, and is sufficient for the uniqueness proof to go through in a non-cooperative setup such as the DSP game.

5 DISCUSSION

In this paper we have considered computational and strategic aspects of auctions involving third party information mediators. Our main result for the computational point of view shows that it is NP-hard to get a reasonable approximation ratio when the three parameters of the problem are all “large”. For the parameters n and k this is tight in the sense that there exists an algorithm whose approximation ratio is good when either one of these parameters is “small”. However, we do not know whether a small value for the parameter m allows for a good approximation ratio. More specifically, even understanding the approximation ratio achievable in the case $m = 2$ is an interesting open problem. Observe that the case $m = 2$ already captures (asymptotically) the largest possible price of stability and price of anarchy in the strategic setup,¹² and thus, it is tempting to assume that this case also captures all the complexity of the computational setup.

Unfortunately, most of our results, for both the computational and strategic setups, are quite negative. The class of local experts we describe is a natural mediators class allowing us to bypass one of these negative result and get a constant approximation ratio algorithm for the computational setup. An intriguing potential avenue for future research is finding additional natural classes of mediators that allow for improved results, either under the computational or the strategic setup. We note that from a mathematical cleanness point of view, it seems natural to study in this context the class of mediators having information about a particular attribute of each of the users in the overall population. Such mediators, however, typically appear in practice as sub-divisions of the auctioneer. Thus, it makes sense from an application point of view to consider further filtering of their signals before communicating to the bidders; which brings us back to the single-mediator setting discussed in previous work.

¹²By Theorem 1.7 the price of stability can be as large as $O(\min\{k, n\})$ even for two mediators, and Theorem 1.6 shows that the price of anarchy cannot be larger than that for any number of mediators.

1177 Another possible direction for future research is to study an extension of our distributed setup
 1178 where bundling is replaced with randomized signaling (similarly to the works of [3] and [8] which
 1179 introduced randomized signaling into the centralized model of [11]). In the centralized model
 1180 it turned out that finding the optimal randomized signaling is easier than finding the optimal
 1181 bundling [3, 8], which is counterintuitive since randomized signaling generalize bundling. Hence,
 1182 one can hope that randomized signaling might also mitigate some of our inapproximability results.
 1183

1184 ACKNOWLEDGMENTS

1185 At the time of writing, the research of Moran Feldman was supported by Israel Science Foundation
 1186 (ISF) grant no. 1357/16, the research of Moshe Tennenholtz was supported by the European Research
 1187 Council (ERC) under the European Union’s Horizon 2020 research and innovation programme
 1188 (Grant No. 740435), and the research of Omri Weinstein was supported by a Simons Society Junior
 1189 fellowship and NSF CAREER award CCF-1844887.
 1190

1191 REFERENCES

- 1192 [1] Robert J. Aumann. 1976. Agreeing to Disagree. *The Annals of Statistics* 4, 6 (1976), 1236–1239.
- 1193 [2] Omer Ben-Porat and Moshe Tennenholtz. 2018. A Game-Theoretic Approach to Recommendation Systems with
 1194 Strategic Content Providers. In *Proceedings of the Neural Information Processing Systems Conference (NeurIPS '18)*.
 1195 Curran Associates, Inc., 1110–1120.
- 1196 [3] Peter Bro Miltersen and Or Sheffet. 2012. Send Mixed Signals: Earn More, Work Less. In *Proceedings of the 13th ACM
 1197 Conference on Electronic Commerce (EC '12)*. ACM, New York, NY, USA, 234–247. [https://doi.org/10.1145/2229012.
 2229033](https://doi.org/10.1145/2229012.2229033)
- 1198 [4] Ruggiero Cavallo, R. Preston McAfee, and Sergei Vassilvitskii. 2015. Display Advertising Auctions with Arbitrage.
 1199 *ACM Trans. Economics and Comput.* 3, 3 (2015), 15. <https://doi.org/10.1145/2668033>
- 1200 [5] Yu Cheng, Ho Yee Cheung, Shaddin Dughmi, and Shang-Hua Teng. 2014. Signaling in Quasipolynomial time. (2014).
 1201 arXiv:1410.3033 <http://arxiv.org/abs/1410.3033>
- 1202 [6] Shaddin Dughmi. 2014. On the Hardness of Signaling. In *Proceedings of the 55th IEEE Annual Symposium on Foundations
 1203 of Computer Science (FOCS '14)*. IEEE Computer Society, Washington, DC, USA, 354–363. [https://doi.org/10.1109/
 FOCS.2014.45](https://doi.org/10.1109/FOCS.2014.45)
- 1204 [7] Shaddin Dughmi, Nicole Immorlica, and Aaron Roth. 2014. Constrained Signaling in Auction Design. In *Proceedings of
 1205 the Twenty-Fifth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA '14)*. Society for Industrial and Applied
 1206 Mathematics, Philadelphia, PA, USA, 1341–1357. <https://doi.org/10.1137/1.9781611973402.99>
- 1207 [8] Yuval Emek, Michal Feldman, Iftah Gamzu, Renato Paes Leme, and Moshe Tennenholtz. 2014. Signaling Schemes for
 1208 Revenue Maximization. *ACM Trans. Economics and Comput.* 2, 2 (2014), 5. <https://doi.org/10.1145/2594564>
- 1209 [9] Eyal Even-Dar, Michael J. Kearns, and Jennifer Wortman. 2007. Sponsored Search with Contexts. In *Proceedings of
 1210 Internet and Network Economics, Third International Workshop (WINE '07)*. Springer-Verlag, Berlin, Heidelberg, 312–317.
 1211 https://doi.org/10.1007/978-3-540-77105-0_32
- 1212 [10] Ronald Fagin, Joseph Y. Halpern, Yoram Moses, and Moshe Vardi. 1995. *Reasoning About Knowledge*. MIT Press,
 1213 Cambridge, MA.
- 1214 [11] Arpita Ghosh, Hamid Nazerzadeh, and Mukund Sundararajan. 2007. Computing Optimal Bundles for Sponsored
 1215 Search. In *Proceedings of Internet and Network Economics, Third International Workshop (WINE '07)*. Springer Berlin
 1216 Heidelberg, Berlin, Heidelberg, 576–583. https://doi.org/10.1007/978-3-540-77105-0_63
- 1217 [12] Mingyu Guo and Argyrios Deligkas. 2013. Revenue Maximization via Hiding Item Attributes. In *Proceedings of the
 1218 23rd International Joint Conference on Artificial Intelligence (IJCAI '13)*. AAAI Press / International Joint Conferences
 1219 on Artificial Intelligence, Menlo Park, California, 157–163. [http://www.aaai.org/ocs/index.php/IJCAI/IJCAI13/paper/
 view/6909](http://www.aaai.org/ocs/index.php/IJCAI/IJCAI13/paper/view/6909)
- 1220 [13] Johan Håstad. 1999. Clique is hard to approximate within $1 - \epsilon$. *Acta Mathematica* 182, 1 (1999), 105–142. <https://doi.org/10.1007/BF02392825>
- 1221 [14] Jason R. Marden and Adam Wierman. 2013. Distributed Welfare Games. *Operations Research* 61, 1 (2013), 155–168.
- 1222 [15] Jacob Marschak and Roy Radner. 1972. *Economic theory of teams*. Yale University Press, New Haven and London.
- 1223 [16] Jonathan R. Mayer and John C. Mitchell. 2012. Third-Party Web Tracking: Policy and Technology. In *Proceedings
 1224 of the IEEE Symposium on Security and Privacy (SP '12)*. IEEE Computer Society, Washington, DC, USA, 413–427.
 1225 <https://doi.org/10.1109/SP.2012.47>
- [17] Dov Monderer. 2006. *Solution-Based Congestion Games*. Springer Japan, Tokyo, 397–409.

- 1226 [18] Tim Roughgarden and Mukund Sundararajan. 2009. Quantifying inefficiency in cost-sharing mechanisms. *Journal of*
1227 *the ACM* 56, 4 (2009), 23:1–23:33. <http://doi.acm.org/10.1145/1538902.1538907>
- 1228 [19] Lloyd S. Shapley. 1953. A value for n-person games. *Contributions to the theory of games 2* (1953), 307–317.
- 1229 [20] Shuai Yuan, Jun Wang, and Xiaoxue Zhao. 2013. Real-time Bidding for Online Advertising: Measurement and Analysis.
1230 In *Proceedings of the Seventh International Workshop on Data Mining for Online Advertising (ADKDD '13)*. ACM, New
1231 York, NY, USA, 3:1–3:8.

1232 Received March 2018; accepted November 2019

1233

1234

1235

1236

1237

1238

1239

1240

1241

1242

1243

1244

1245

1246

1247

1248

1249

1250

1251

1252

1253

1254

1255

1256

1257

1258

1259

1260

1261

1262

1263

1264

1265

1266

1267

1268

1269

1270

1271

1272

1273

1274