Constrained Signaling for Welfare and Revenue Maximization

SHADDIN DUGHMI
University of Southern California
and
NICOLE IMMORLICA
Microsoft Research
and
AARON ROTH
University of Pennsylvania

We consider auction settings where the seller is constrained in the amount and nature of information he may reveal about the good being sold. This is encountered, for example, in online advertising auctions, where communicating precise details of every viewer to interested advertisers is impractical, costly, and possibly socially undesirable. We initiate the study of *constrained signaling* in such settings, where a seller must choose which information to reveal subject to exogenous constraints on the signaling policy. We consider a seller employing the second-price auction, and present algorithms and hardness results for approximating the welfare and revenue maximizing signaling policies under a variety of constraints.

Categories and Subject Descriptors: F.2.2 [**Analysis of Algorithms and Problem Complexity**]: General

General Terms: Algorithms, Economics, Theory

Additional Key Words and Phrases: Signaling, Auctions, Mechanism Design

1. INTRODUCTION

The study of economic interactions in the presence of information asymmetries has a rich history, beginning with the seminal work of [Akerlof 1970]. Akerlof observed that the information structure of a market — i.e. who has what information regarding the goods and services for sale — can have a profound effect on equilibrium outcomes. Since then, a rich literature has examined the effect of information revelation, also known as *signaling*, in markets (e.g. by [Spence 1973; 2002]) and, on a smaller scale, in auctions.

In this work [Dughmi et al. 2013], we consider signaling in some of the simplest economic settings, namely single item auctions. In addition to their preponderance on the Internet, such as in online advertising, auctions serve as an instructive microcosm in which to study information asymmetry between buyers and sellers in a market, and the effects of signaling. In our model, an auctioneer is looking to sell a single good to one of several buyers, though the various qualities of the good are known to the auctioneer yet ex-ante unknown to the buyers. This arises, for example, in online advertising: a content publisher looking to sell an impression

Authors' addresses: shaddin@usc.edu, nicimm@microsoft.com, aaroth@cis.upenn.edu.

has more information regarding the identity of the viewer behind the impression than do the advertisers bidding for it. We model such information asymmetry by assuming that the good is drawn by nature from a distribution over possible goods, assumed to be common knowledge. The auctioneer then learns the identity of the realized good, and has the opportunity to publicly announce a message, which we refer to as a *signal*, to the buyers before running an auction. We assume that the auctioneer commits to a *signaling scheme* — a policy specifying a distribution over signals for each good — before nature draws the good, and announces said policy to all the buyers. Each signaling scheme and choice of an auction format (first price, second price, etc) then induces a game among the buyers. The auctioneer's problem is that of designing the signaling scheme to optimize his favored objective, commonly his revenue or the welfare of the winning player.

The conventional wisdom on signaling in auctions is that *more information is better*. This principle rings most true in the simplest setting of all: an auctioneer running a second price auction,[1] and interested in maximizing welfare. To see this, observe that when buyers know the identity of the good being sold, a second price auction awards the item to the player who values it most. In contrast, when the employed signaling scheme announces partial information regarding the identity of the good, the second price auction awards the good to the bidder with the greatest *expected* value after a Bayesian update based on the signal, which is not necessarily the optimal point-wise choice. The *linkage principle* of [Milgrom and Weber 1982] reinforces the same popular lore: under some conditions, revealing more information to the players increases the auctioneer's *revenue* as well.[2] It is tempting, therefore, to conclude that the seller's problem of what to reveal is trivial in many settings, as he should simply announce all available information.

Our motivating observation is that complete transparency is often impossible or costly. In online advertising auctions for example, the sheer volume of sales and diversity of viewers make it impractical for a publisher to communicate to advertisers the precise details of every viewer, and the publisher's privacy policy may further constrain the information they reveal. In addition, legal and reputational considerations may restrict a seller from announcing certain signals for certain items: e.g., products in a grocery store must pass a certification process from the corresponding agency to be sold under the "organic" label. Such constraints introduce intricate tradeoffs in choosing *which* information to reveal. Quantifying those tradeoffs inevitably requires examination of these settings with an optimization lens, as we begin to do in this work.

2. SUMMARY OF RESULTS

To begin our study of constrained signaling, we define and analyze two illustrative settings, one with a finite set of items and combinatorial constraints on the signaling

[1]The choice of auction format is not particularly important for this conclusion, though we fix the second price auction for clarity.

[2]However, as pointed out by [Levin and Milgrom 2010], the linkage principle does not apply in many settings, including ones we consider in this work. Recent work in [Emek et al. 2012] and [Bro Miltersen and Sheffet 2012] designs polynomial time algorithms for computing revenue-maximizing signaling schemes in general.

scheme, and another where the set of items is infinite. In both settings, we consider a seller who first invokes a signaling scheme and then runs a second price auction.

(1) In the *combinatorial setting*, an item is drawn by nature from a finite set according to a known prior. We consider two constraints on the signaling policy: a *cardinality constraint* limiting the number of signals used, and a *graph constraint* given as a bipartite graph with items on one side, signals on the other, and edges describing the valid signals for each item.

(2) In the *geometric setting*, items live in a compact subset of Euclidean space. As before, an item is drawn by nature from a common prior. We assume the dimension d of the space is too high to communicate the identity of an item directly, and constrain our signals to strings of length logarithmic in d.

We present algorithms and impossibility results for computing optimal signaling schemes in these settings, after making some assumptions.

(1) In the combinatorial setting, we assume that players' valuations for potential items are known, or drawn from a prior with constant-size support. We present constant-factor approximation algorithms for computing the welfare maximizing signaling scheme under both a cardinality constraint and a graph constraint, and for computing the revenue maximizing signaling scheme under only a cardinality constraint.

(2) In the combinatorial setting, we show that both welfare and revenue are NP-hard to approximate better than a specific constant factor, even assuming given valuations and only a cardinality constraint on the signaling scheme.

(3) In the geometric setting, we assume players' valuations are defined in terms of either the distance or the angle from the realized good to some target good or set of goods, and drawn from a known prior. We present communication-constrained signaling schemes that achieve nearly the optimal welfare, compared even to the optimal unconstrained signaling scheme.

Our results for the combinatorial setting draw heavily on techniques from the literature on submodular function optimization. For the geometric setting, we employ tools from no regret learning and metric embeddings to prove our results.

3. FUTURE DIRECTIONS

The main contribution of this work, we believe, is enabling future study of additional constrained signaling settings. We conclude by mentioning two of them.

(1) In both models we study, the constraint on our signaling scheme is essentially separable over items. What if, instead, products are described by a set of attributes — in online advertising, those may include age, gender, location, and browsing history, and a signaling scheme is constrained to choose a small subset of those attributes, crucially the same for all goods, to advertise prior to sale? Problems of this form are arguably more natural than those we consider in this work, though appear particularly challenging and intimately related to unresolved questions in learning theory.

(2) Signaling can be thought of as a process of classification; specifically of items into subsets, each of which is associated with a signal. When items are associated with entities that have a contractual or service relationship with the seller, as in online advertising, issues of fairness inevitably arise, as described in [Dwork et al. 2012]. A poignant example is the controversy stirred by the recent work of [Sweeney 2013], which illustrates the effect of race on ad delivery in Google's Adsense. To avoid such blowback, a seller may constrain their signaling scheme to treat similar items similarly.

REFERENCES

Akerlof, G. 1970. The market for "lemons": Quality uncertainty and the market mechanism. *The Quarterly Journal of Economics 84*, 488–500.

Bro Miltersen, P. and Sheffet, O. 2012. Send mixed signals: earn more, work less. In *Proceedings of the 13th ACM Conference on Electronic Commerce*. ACM, 234–247.

Dughmi, S., Immorlica, N., and Roth, A. 2013. Constrained signaling for welfare and revenue maximization. *CoRR abs/1302.4713*.

Dwork, C., Hardt, M., Pitassi, T., Reingold, O., and Zemel, R. 2012. Fairness through awareness. In *Proceedings of the 3rd Innovations in Theoretical Computer Science Conference*. ACM, 214–226.

Emek, Y., Feldman, M., Gamzu, I., Paes Leme, R., and Tennenholtz, M. 2012. Signaling schemes for revenue maximization. In *Proceedings of the 13th ACM Conference on Electronic Commerce*. ACM, 514–531.

Levin, J. and Milgrom, P. 2010. Online advertising: Heterogeneity and conflation in market design. *American Economic Review 100*, 603–607.

Milgrom, P. and Weber, R. 1982. A theory of auctions and competitive bidding. *Econometrica 50*, 1089–1122.

Spence, M. 1973. Job market signaling. *The quarterly journal of Economics 87,* 3, 355–374.

Spence, M. 2002. Signaling in retrospect and the informational structure of markets. *The American Economic Review 92,* 3, 434–459.

Sweeney, L. 2013. Discrimination in online ad delivery. *arXiv preprint arXiv:1301.6822*.