Bayesian Persuasion under Ex Ante and Ex Post Constraints

Konstantin Zabarnyi (Technion CS)

Joint work with Prof. Yakov Babichenko (Technion IE&M) and Dr. Inbal Talgam-Cohen (Technion CS)

Talk Outline

- What is Bayesian persuasion?
- Our contribution
- Related work
- Preliminaries
- Proof sketch of our main existence result
- Discussion of our computational results
- Conclusions and future work

What is Bayesian Persuasion?

- Example: online ad auctions.
- A web user is about to view a personalized ad.
- Several advertisers bid on showing the ad.
- The advertising platform (but not the bidders) knows the user’s profile.
- How do the bidders determine their bids?

Online Ad Auctions

- The advertising platform sends a public signal to the advertisers with some information about the user's profile.
- The advertising platform must commit on a signaling scheme before observing the user's profile.

Online Ad Auctions

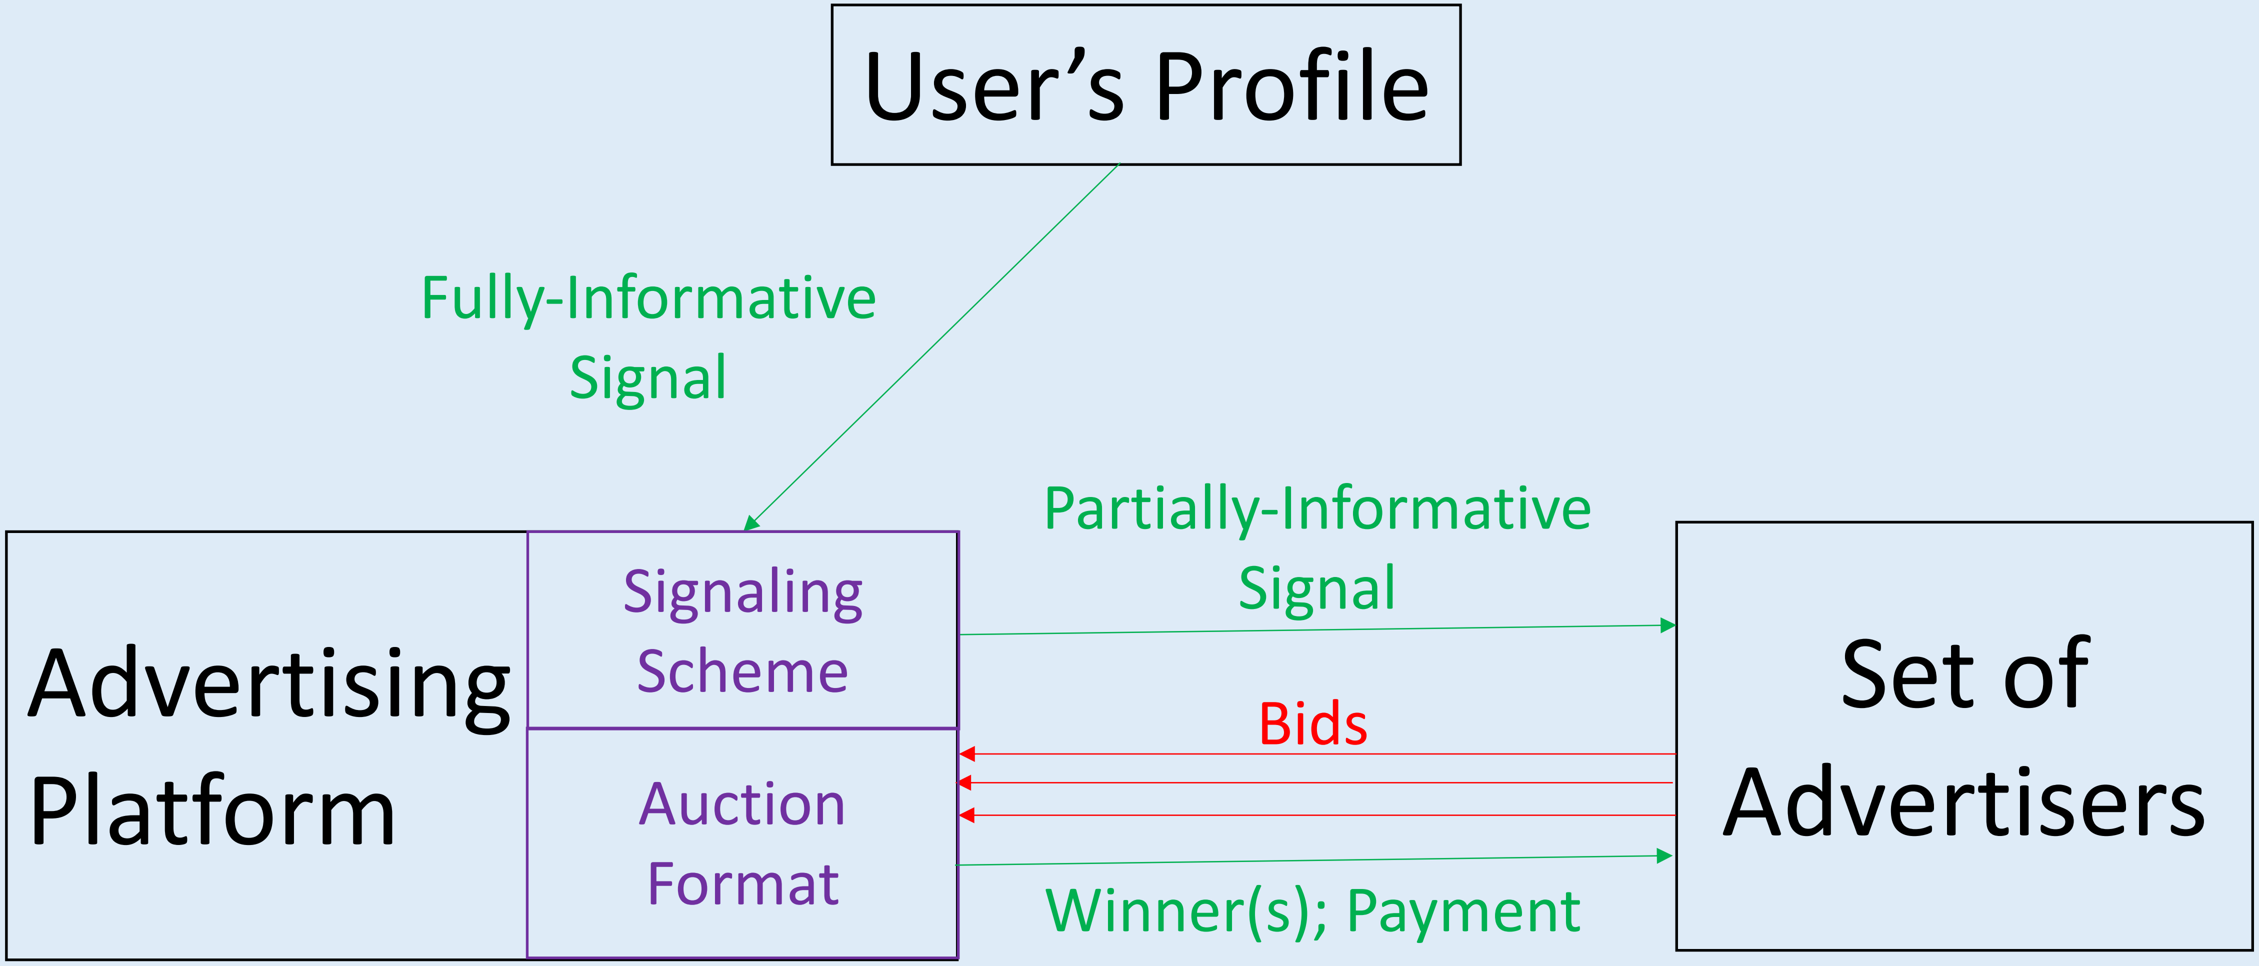

Bayesian Persuasion

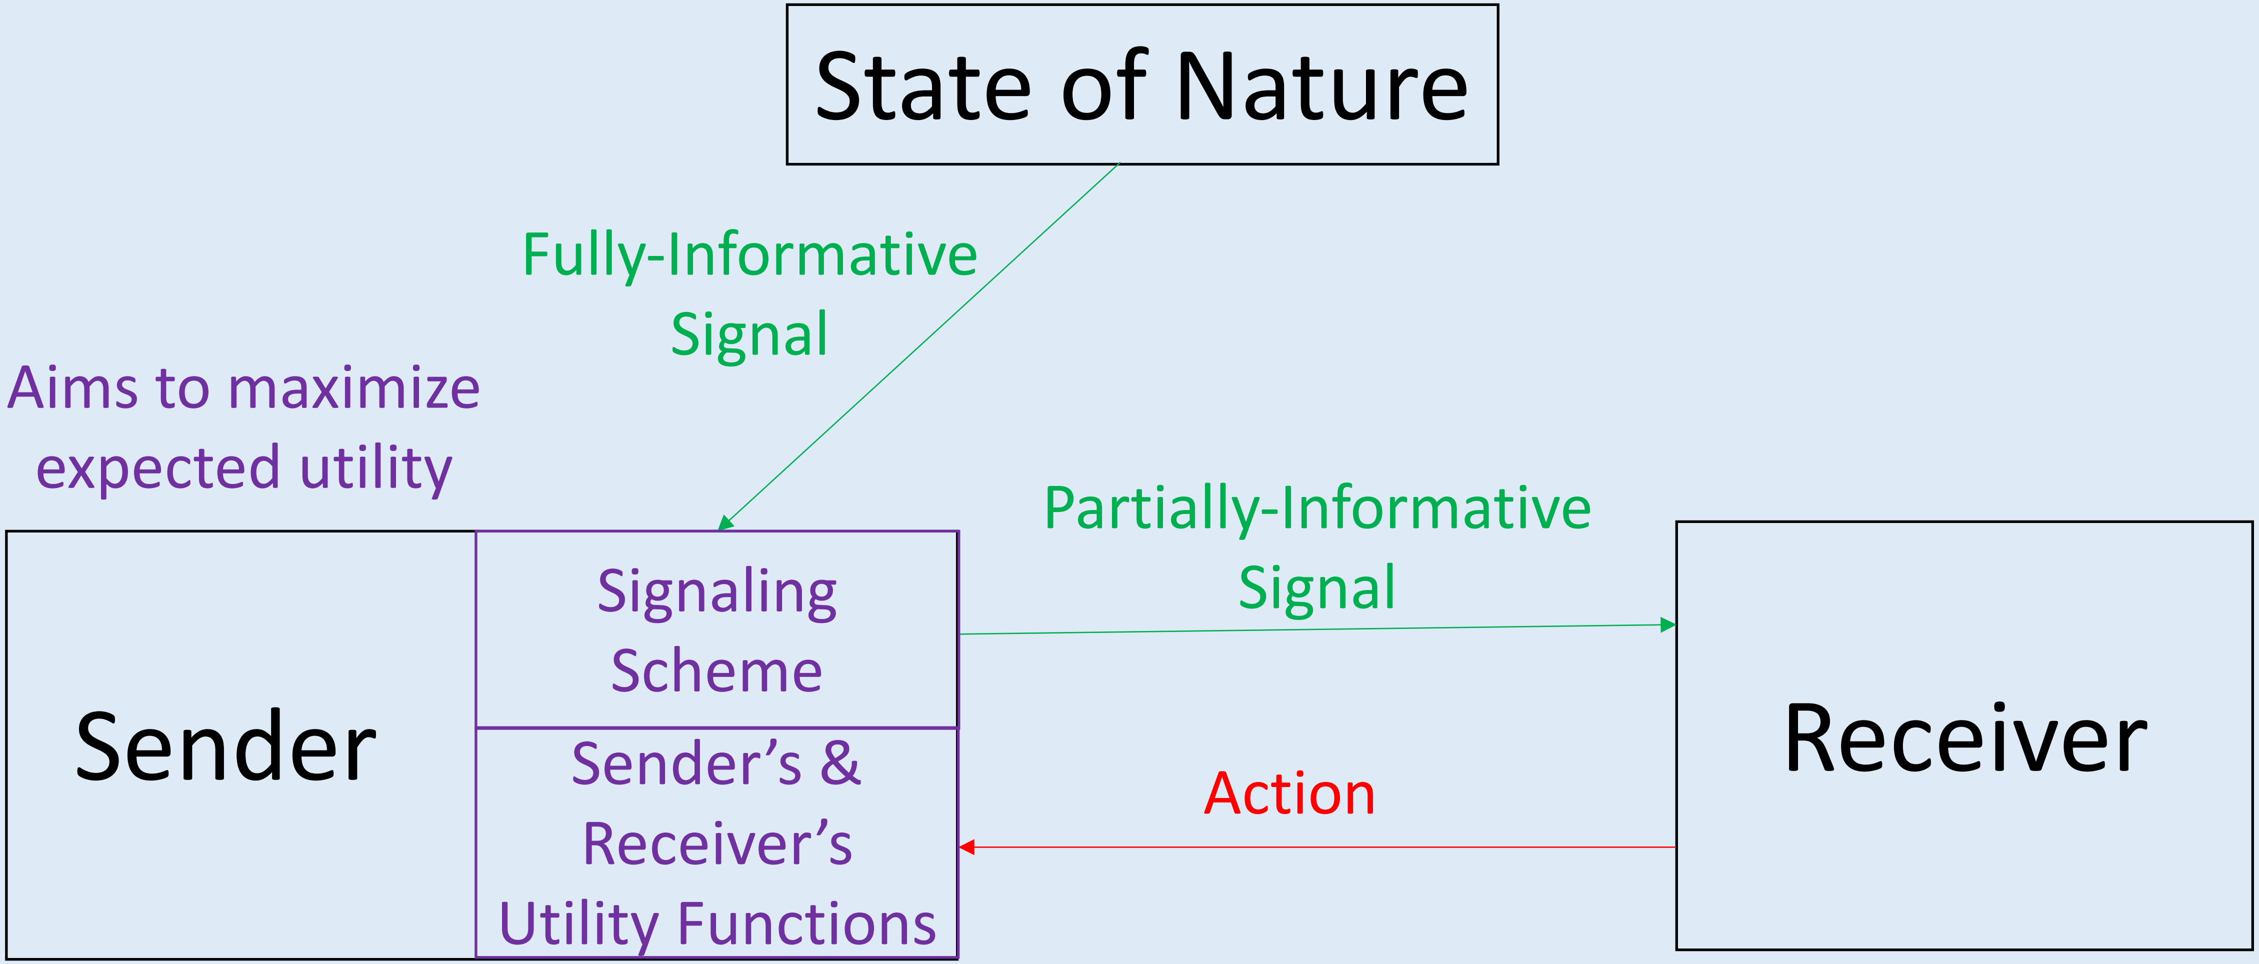

More Examples

- Sender: politician, Receiver: set of voters [Alonso-Câmara 2016]
- Sender: advisor, Receiver: executive [Bloedel-Segal 2018]
- Signaling can also be private [Arieli-Babichenko 2019]

Practical Limitations

- Respecting user’s privacy (ad auctions)
- Preventing discrimination (ad auctions)
- Acknowledging limited Receiver’s attention (information management in organizations)

Our Contribution

- We propose a simple mathematical model capturing practical limitations on Sender's signaling scheme.
- We define two constraint families:

- ex post – every Sender-Receiver communication instance is restricted;

- ex ante – a more general family, which can also restrict the communication in expectation.

Our Contribution – Existence Results

- The support of a signaling scheme is the set of all possible signal realizations.
- The support size is similar to menu-size complexity in auctions.
- For both constraint families, we show that there exist signaling schemes with a small (linear-sized) support.

Our Contribution – Computational Aspects

- We provide an additive bi-criteria FPTAS for an optimal constrained signaling scheme for a constant number of states of nature (under general assumptions).
- Dughmi and Xu [2017] rule out an additive PTAS or a constant-factor poly-time multiplicative approximation for a non-constant number of states of nature (unless $P = NP$).

Our Contribution – Computational Aspects

- We improve the approximation to single-criteria under a Slater-like regularity condition.
- Weaker assumptions still yield an additive (bi/single-criteria) PTAS.

Our Contribution – Ex Post vs Ex Ante

- We show that in general, ex ante constraints can outperform ex post constraints by an arbitrary multiplicative factor.
- However, <u>the ratio is bounded for convex constraints</u> and Sender's utility functions suitable to common <u>auction settings</u>.

Our Contribution – Ex Post vs Ex Ante

- We use this result to derive an <u>approximately welfare-maximizing</u> constrained signaling scheme in ad auctions with <u>exponentially many states of nature</u>.
- We shall not discuss our ex post vs ex ante results today.

Related Work

- [Kamenica-Gentzkow 2011] – the basic Bayesian persuasion model
- [Milgrom-Weber 1982] – signaling in auctions
- [Cheng et al. 2015] – an additive FPTAS based on discretization and LP
- [Dughmi et al. 2014,2015], [Ichihashi 2019] – constrained Bayesian persuasion
- [Vølund 2018] – a model equivalent to ex post constraints

Standard Preliminaries

- States of nature space $\Omega = \{\omega_1, \ldots, \omega_k\}$
- Commonly-known prior distribution p on Ω
- State of nature ω
- Receiver's action space A (compact & nonempty)
- Signaling scheme Σ, which is a randomized function from Ω to $\mathrm{supp}(\Sigma)$
- Signal realization σ
- σ transforms p to a posterior distribution p_σ on Ω, which specifies Receiver's action

Standard Preliminaries

- p_σ practically allows us to ignore Receiver's utility function.
- We fix a nonnegative Sender's utility function $\mathrm{u_s}(\mathrm{p}_\sigma)$.
- For simplicity, we assume that $\mathrm{u_s}$ is state-independent, but our existence and computational results extend to the state-dependent case.

Standard Preliminaries (Weakened)

- We assume that $\mathrm{u_s}(\cdot)$ is upper semi-continuous.
- That is, $\limsup\limits_{\mathrm{p}_\sigma \to \mathrm{p}_{\sigma_0}} \mathrm{u_s}(\mathrm{p}_\sigma) \leq \mathrm{u_s}\left(\mathrm{p}_{\sigma_0}\right)$ for every p_{σ_0}.
- This is a relaxation of the standard continuity assumption.

$\Omega = \{\omega_1, ..., \omega_k\} -$ States of nature space

$\omega -$ State of nature

$\Sigma -$ Signaling scheme

$p -$ Prior distribution on Ω

$p_\sigma -$ Posterior distribution on Ω

Standard Preliminaries

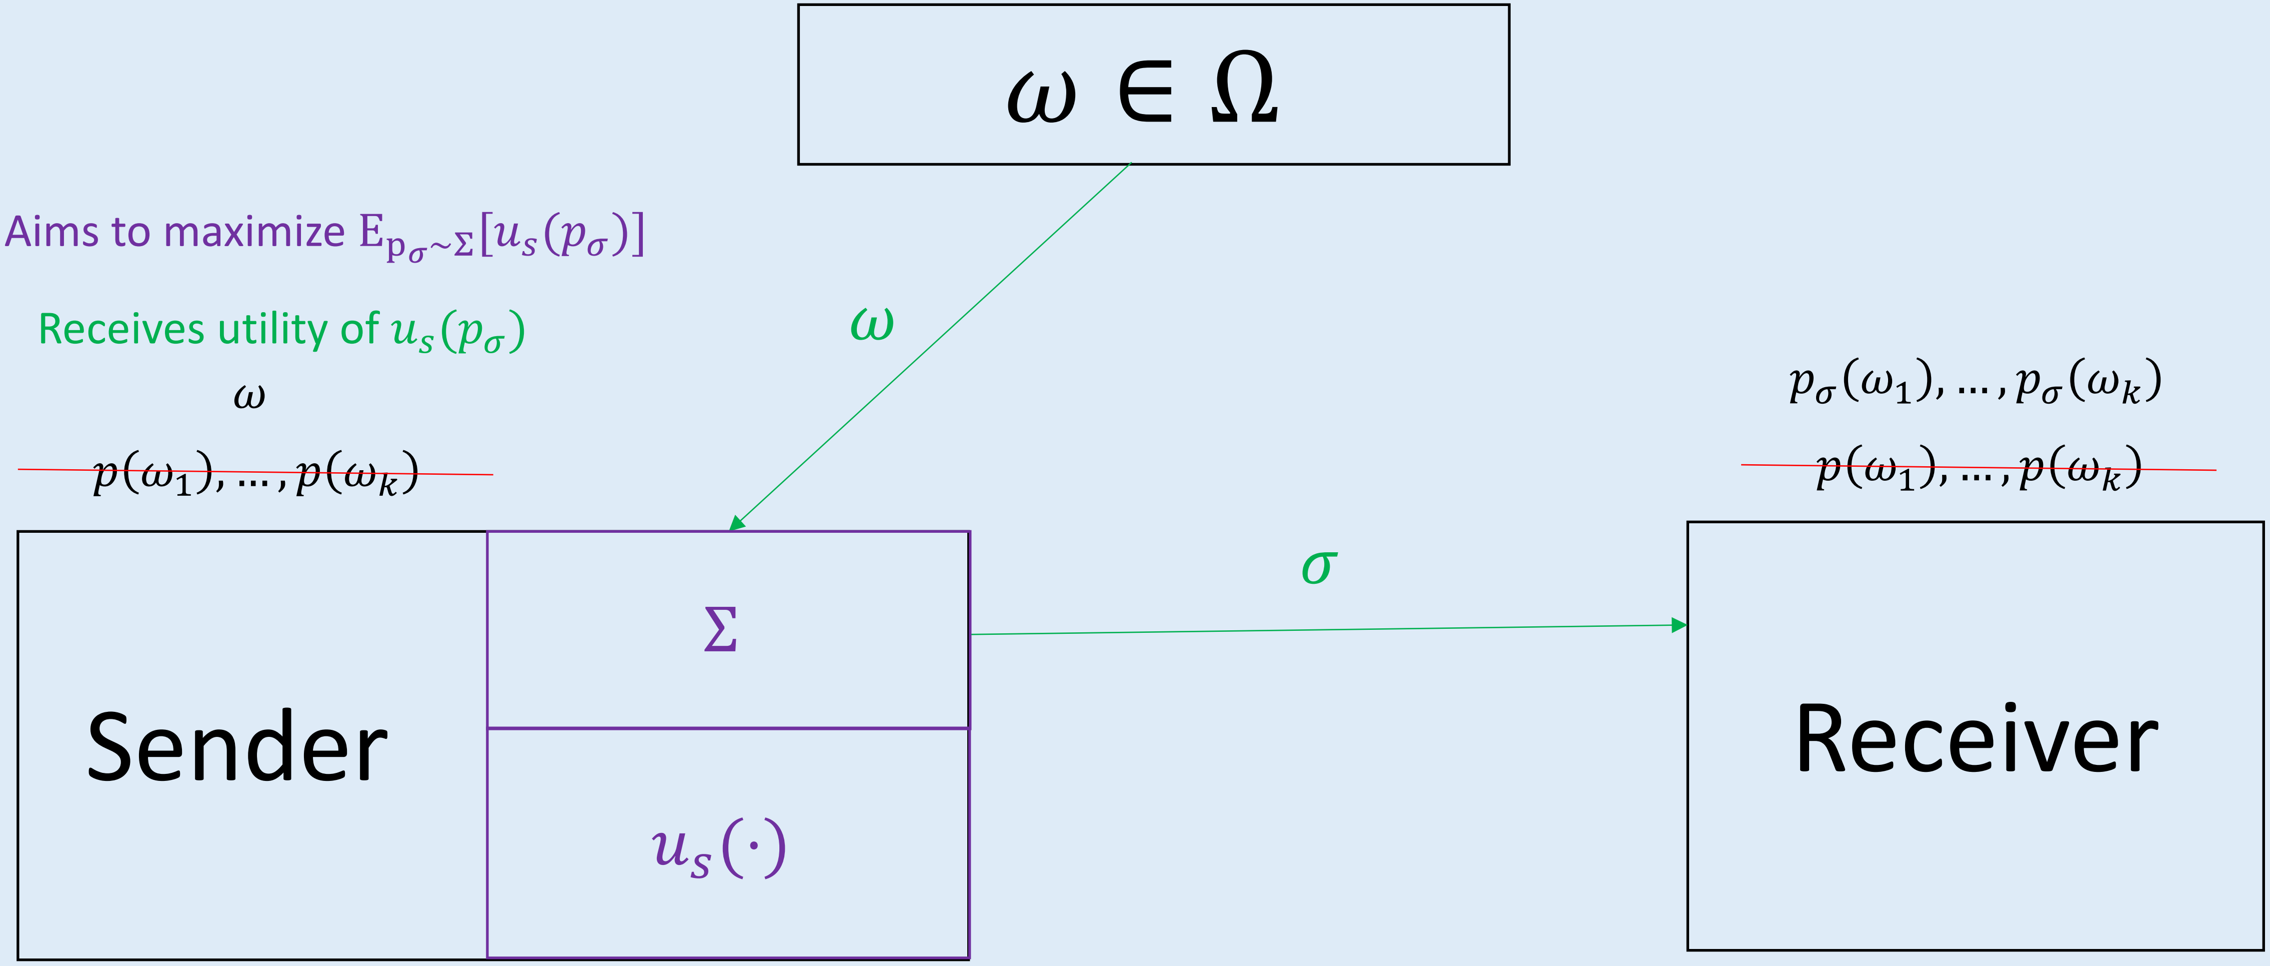

Requirements from the Signaling Scheme

- In the standard model of Kamenica and Gentzkow [2011], the only restriction on Σ is that it must be Bayes-plausible.
- That is, the expected probability over Σ of every $\omega_0 \in \Omega$ must be equal to $p(\omega_0)$.
- We require Σ to satisfy a certain set of additional constraints (ex post or ex ante).

Ex Post Constraints

- An ex post constraint specified by a continuous function $\mathrm{f}(p_\sigma)$ and a constant c requires that $\mathrm{f}(p_\sigma) \leq c$ for every $p_\sigma \in supp(\Sigma)$.
- That is, $supp(\Sigma)$ is restricted to a compact subset of the space of distributions over Ω.

$\Omega = \{\omega_1, \ldots, \omega_k\}$ – States of nature space

ω – State of nature

Σ – Signaling scheme

p – Prior distribution on Ω

p_σ – Posterior distribution on Ω

Ex Ante Constraints

- An ex ante constraint specified by a continuous function $\mathrm{f}(p_\sigma)$ and a constant c requires that $\mathrm{E}_{p_\sigma \sim \Sigma}[\mathrm{f}(p_\sigma)] \leq c$.
- In particular, an ex post constraint specified by some f and c is equivalent to the ex ante constraint specified by $\max\{\mathrm{f}, \mathrm{c}\}$ and c.

$\Omega = \{\omega_1, \ldots, \omega_k\}$ – States of nature space

ω – State of nature

Σ – Signaling scheme

p – Prior distribution on Ω

p_σ – Posterior distribution on Ω

Ex Post and Ex Ante Constraints – Example (Ad Auctions)

- An <u>ex post</u> constraint may require that the user's gender is never revealed with more than 75% certainty.
- The corresponding <u>ex ante</u> constraint imposes an analogous restriction <u>on average</u>.
- That is, the advertising platform may reveal the gender truthfully on 50% of the instances and to conceal it on the remaining 50%.

Ex Post and Ex Ante Constraints – Example (Ad Auctions)

- Ex post constraints provide a robust protection of individual privacy.
- Ex ante constraints protect privacy on a group level (e.g., by limiting Receiver's ability to learn the shopping habits of certain social groups).
- Ex ante constraints ensure inclusion of different social groups.

Ex Post and Ex Ante Constraints – Example (Information Management in Organizations)

- Ex post constraints require every report of the advisor to the executive to be short.
- Ex ante constraints ensure efficiency on average.

An Existence Result for Ex Ante Constraints

- Thm.: Fix m ex ante constraints s.t. there exists a valid signaling scheme. Then there exists an optimal valid signaling scheme with support size $\leq k + m$.
- Furthermore, this bound is tight.

$\Omega = \{\omega_1, \ldots, \omega_k\} -$ States of nature space

$\omega -$ State of nature

$\Sigma -$ Signaling scheme

$p -$ Prior distribution on Ω

$p_\sigma -$ Posterior distribution on Ω

An ex ante constraint: $\mathrm{E}_{p_\sigma \sim \Sigma}\left[\mathrm{f}\left(p_\sigma\right)\right] \leq c$

An Existence Result for Ex Post Constraints

- Thm.: Fix a set of ex post constraints s.t. there exists a valid signaling scheme. Then there exists an optimal valid signaling scheme with support size $\leq k$.
- This bound is the same as for the unconstrained setting, and this bound is tight.

$\Omega = \{\omega_1, \ldots, \omega_k\}$ – States of nature space

ω – State of nature

Σ – Signaling scheme

p – Prior distribution on Ω

p_σ – Posterior distribution on Ω

An ex post constraint: $\forall p_\sigma \in \mathrm{supp}(\Sigma)$: $\mathrm{f}(p_\sigma) \leq c$

Proof Sketch for Ex Ante Constraints

- <u>Step 1</u>: The optimization problem is an infinite-dimensional LP, with Σ being the "variables".

$$
\begin{aligned}
\max \quad & \mathbf{E}_{p_\sigma \sim \Sigma}\left[u_s\left(p_\sigma\right)\right] \\
s.t. \quad & \mathrm{p}\left[\omega_0\right] = \mathbf{E}_{p_\sigma \sim \Sigma}\left[p_\sigma\left[\omega_0\right]\right] \quad \forall \omega_0 \in \Omega \\
& \mathbf{E}_{p_\sigma \sim \Sigma}\left[f_i\left(p_\sigma\right)\right] \le \mathrm{c}_i \quad \forall 1 \le i \le m
\end{aligned}
$$

$\Omega = \{\omega_1, \ldots, \omega_k\} -$ States of nature space

$\omega -$ State of nature

$\Sigma -$ Signaling scheme

$p -$ Prior distribution on Ω

$p_\sigma -$ Posterior distribution on Ω

An ex ante constraint: $\mathrm{E}_{p_\sigma \sim \Sigma}\left[\mathrm{f}\left(p_\sigma\right)\right] \le c$

Proof Sketch for Ex Ante Constraints

- The target function is upper semi-continuous w.r.t. the Lévy–Prokhorov metric on the space of the valid signaling schemes and the usual metric on $\mathbb{R}_{\geq 0}$.

Proof Sketch for Ex Ante Constraints

- <u>Step 2</u>: The maximum is obtained at an extreme point of the feasible set.
- The target function is upper semi-continuous and linear (step 1).
- The feasible set is compact, convex and nonempty.
- Therefore, one of the maximizers is an extreme point (Bauer's maximum principle).

Proof Sketch for Ex Ante Constraints

- Step 3: Every extreme point has support size $\leq 2^{k+m}$.
- There are $k+m$ linear constraints (Bayes-plausibility & ex ante constraints).
- Each constraint is specified by a hyperplane.
- Adding the hyperplanes one-by-one at most doubles the support size upon each addition.

Constraints: $$\mathrm{p}[\omega_0] = \mathbf{E}_{p_\sigma \sim \Sigma}\left[p_\sigma[\omega_0]\right] \quad \forall \omega_0 \in \Omega = \{\omega_1, \ldots, \omega_k\}$$

$$\mathbf{E}_{p_\sigma \sim \Sigma}\left[f_i(p_\sigma)\right] \leq \mathrm{c}_i \quad \forall 1 \leq i \leq m$$

Proof Sketch for Ex Ante Constraints

- Step 4: Every extreme point has support size $\leq k + m$.
- From step 3, we get an infinite family of finite LPs.
- Each finite LP has 2^{k+m} variables (representing the probability weights assigned to the support elements).
- Each finite LP has $k + m$ constraints.
- Every extreme point of the infinite LP is an extreme point of a finite LP.

Common Constraints

- Kullback-Leibler (KL) divergence: $D_{KL}(p_\sigma||p) = \sum_{\omega_0\in\Omega} p_\sigma(\omega_0)\log\frac{p_\sigma(\omega_0)}{p(\omega_0)}$ (compares informativeness of the posterior and the prior).
- Norms of $p_\sigma - p$ (including variation distance).

$\Omega = \{\omega_1, ..., \omega_k\} -$ States of nature space

$\omega -$ State of nature

$\Sigma -$ Signaling scheme

$p -$ Prior distribution on Ω

$p_\sigma -$ Posterior distribution on Ω

Computational Results

- Assume that k is constant and seek for optimal valid Σ.
- We provide an <u>additive bi-criteria FPTAS</u> for practical Sender's utility (e.g., Lipschitz or piecewise-constant) and constraint (e.g., KL divergence or norms) families.
- Bayes-plausibility is satisfied precisely.

$\Omega = \{\omega_1, ..., \omega_k\} -$ States of nature space

$\omega -$ State of nature

$\Sigma -$ Signaling scheme

$p -$ Prior distribution on Ω

$p_\sigma -$ Posterior distribution on Ω

Computational Results

- The same algorithm yields an <u>additive bi-criteria PTAS</u> for <u>any</u> ex ante constraints and continuous or piecewise-constant Sender's utility.
- <u>Continuous Sender's utility</u> – standard assumption if Receiver has continuum of actions.
- <u>Piecewise-constant Sender's utility</u> – captures the case in which Receiver has finitely many actions.

Single-Criteria Approximations

- The reason for the bi-criteria approximations is degenerate cases.
- Finding a root of a polynomial can be expressed in terms of ex ante constraints.
- Assuming that there exists a signaling scheme satisfying all the ex ante constraints in a strict inequality, both results improve to single-criteria approximations.

Computational Results – Proof Ideas

- Approximate Sender's utility with an upper semi-continuous piecewise-constant function.
- Approximate the ex ante constraints with Lipschitz functions.
- Strengthen the constraints to get a single-criteria approximation (if a Slater-like condition holds).
- Solve a finite LP.
- The proofs involve subtle technical points.

Conclusions and Future Work

- We initiate the study of ex ante- and ex post-constrained Bayesian persuasion, and prove:

- existence of a valid signaling scheme with a <u>linear-sized support</u> (we provide tight bounds);
- <u>positive computational results</u> for a constant number of states of nature;
- a bound on the ratio between the optimal Sender's utility under convex ex ante and ex post constraints and "nice" utility functions.

Conclusions and Future Work

- Our results apply to ad auctions and limited attention.
- The ex post vs ex ante result applies to ad auctions with <u>exponentially large</u> states of nature space.
- Future research directions:

- studying optimal ex post-constrained persuasion in special cases;

- studying constrained private signaling.

Thank You!

References

- [Alonso-Câmara 2016]. Alonso, Ricardo, and Odilon Câmara. "Persuading voters." American Economic Review 106.11 (2016): 3590-3605.
- [Bloedel-Segal 2018]. Bloedel, Alexander W., and Ilya R. Segal. "Persuasion with rational inattention." Available at SSRN 3164033 (2018).
- [Arieli-Babichenko 2019]. Arieli, Itai, and Yakov Babichenko. "Private Bayesian persuasion." Journal of Economic Theory 182 (2019): 185-217.

References

- [Dughmi-Xu 2017]. Dughmi, Shaddin, and Haifeng Xu. "Algorithmic persuasion with no externalities." Proceedings of the 2017 ACM Conference on Economics and Computation. 2017.
- [Kamenica-Gentzkow 2011]. Kamenica, Emir, and Matthew Gentzkow. "Bayesian persuasion." American Economic Review 101.6 (2011): 2590-2615.

References

- [Milgrom-Weber 1982]. Milgrom, Paul R., and Robert J. Weber. "Distributional strategies for games with incomplete information." Mathematics of operations research 10.4 (1985): 619-632.
- [Cheng et al. 2015]. Cheng, Yu, et al. "Mixture selection, mechanism design, and signaling." 2015 IEEE 56th Annual Symposium on Foundations of Computer Science. IEEE, 2015.

References

- [Dughmi et al. 2014]. Dughmi, Shaddin, Nicole Immorlica, and Aaron Roth. "Constrained signaling in auction design." Proceedings of the twenty-fifth annual ACM-SIAM symposium on Discrete algorithms. Society for Industrial and Applied Mathematics, 2014.
- [Dughmi et al. 2015]. Dughmi, Shaddin, et al. "Algorithmic signaling of features in auction design." International Symposium on Algorithmic Game Theory. Springer, Berlin, Heidelberg, 2015.

References

- [Ichihashi 2019]. Ichihashi, Shota. "Limiting Sender's information in Bayesian persuasion." Games and Economic Behavior 117 (2019): 276-288.
- [Vølund 2018]. Vølund, Rune T. "Bayesian Persuasion on Compact Subsets." Theoretical Models in Behavioral Economics (2018): 64.

- This is only a partial reference list; the full list appears in our paper.