Part 3
Algorithmic Persuasion

Haifeng Xu

USC → Harvard (postdoc) → UVA

Outline

- Introduction
- Persuasion Through the Algorithmic Lens
- Applications of Persuasion
- The Bigger Picture and Future Directions

Example: Recommendation Letters

- Advisor vs. recruiter
- 1/3 of the advisor's students are excellent; 2/3 are average
- A fresh graduate is randomly drawn from this population
- Recruiter
  - Utility $1+\epsilon$ for hiring an excellent student; -1 for an average student
  - Utility 0 for not hiring
  - A-priori, only knows the advisor's student population

$$(1+\epsilon)\times 1/3 - 1\times 2/3 \qquad < \qquad 0$$

hiring *Not hiring*

Example: Recommendation Letters

- Advisor vs. recruiter
- 1/3 of the advisor's students are excellent; 2/3 are average
- A fresh graduate is randomly drawn from this population
- Recruiter
  - Utility $1+\epsilon$ for hiring an excellent student; -1 for an average student
  - Utility 0 for not hiring
  - A-priori, only knows the advisor's student population
- Advisor
  - Utility 1 if the student is hired, 0 otherwise
  - Knows whether the student is excellent or not

Example: Recommendation Letters

What is the advisor's optimal "recommendation strategy"?

- Attempt 1: always say "excellent" (equivalently, no information)
  - Recruiter ignores the recommendation
  - Advisor expected utility 0

Example: Recommendation Letters

What is the advisor's optimal "recommendation strategy"?

- Attempt 2: honest recommendation (i.e., full information)
  - Advisor expected utility 1/3

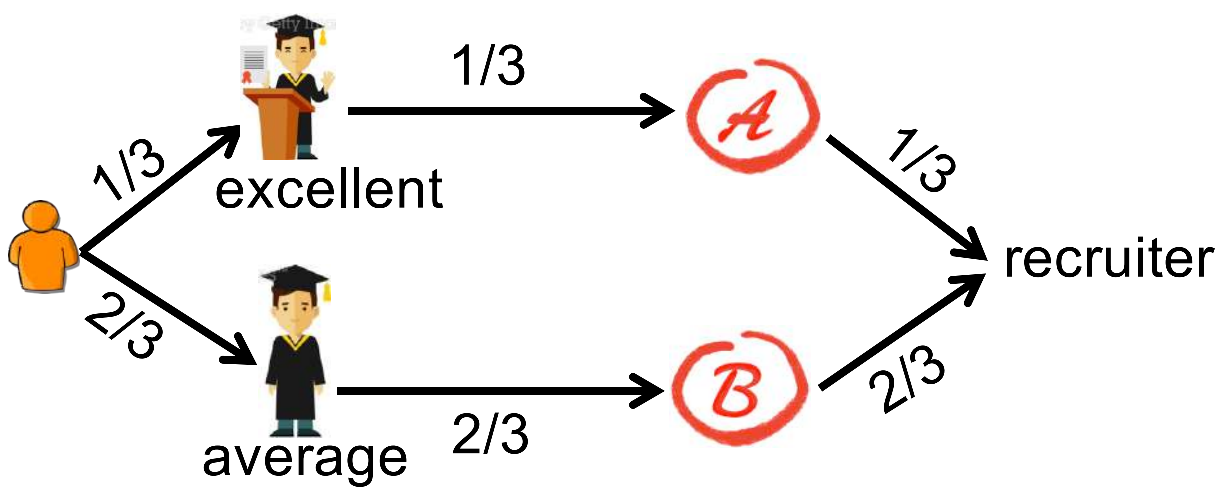

Example: Recommendation Letters

What is the advisor's optimal "recommendation strategy"?

- Attempt 3: partial noisy information → advisor expected utility 2/3

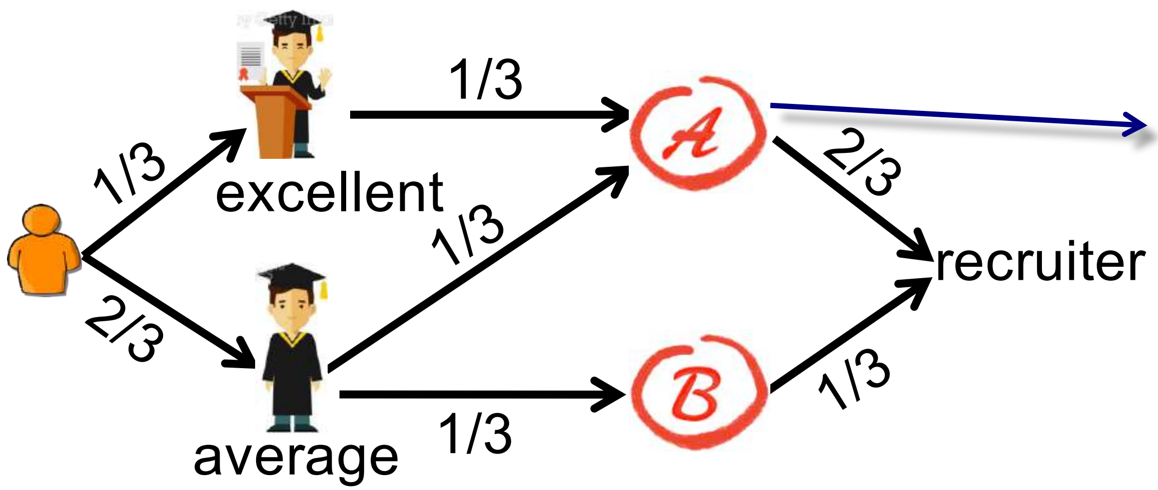

P(excellent | A) = $1/2$

$(1+\epsilon - 1)/2 \;>\; 0$

Hiring   Not hiring

Persuasion the act of exploiting an informational advantage in order to influence the decisions of others

- Intrinsic in human activities: advertising, negotiation, politics, security, marketing, financial regulation,…
- A large body of recent work

One Quarter of GDP Is Persuasion

By DONALD MCCLOSKEY AND ARJO KLAMER*

—— The American Economic Review Vol. 85, No. 2, 1995.

Bayesian Persuasion [Kamenica/Gentzkow 11]

- Two players: a persuader (sender), a decision maker (receiver)
  - Previous example: advisor = sender, recruiter = receiver
- Receiver must choose action from $1, 2, \dots, n$
- Action i has random type θ_i
  - Determines sender utility $s(\theta_i)$ and receiver utility $r(\theta_i)$
- State of nature $\theta = (\theta_1, \dots, \theta_n) \in R^{2n}$ drawn from a common prior
- Sender can observe realized θ; receiver only knows the prior

Persuasion Problem

Sender must design and commit to a signaling scheme X:

- Randomized map from *states of nature* to signals

When state θ is realized, sender must communicate $\sigma \sim X(\theta)$ to receiver before he chooses action.

The Commitment Assumption

Argument 1: Emerges at Equilibrium

- Game played repeatedly: sender and receiver optimize long-run payoff
- Commitment $\approx$ reputation / credibility

Argument 2: Service Agreement or Trusted Authority

- E.g., in an auction, principal commits to rules of interactions
- Publish code, undergo audits / statistical tests

Direct Persuasive Schemes

- In previous example, the scheme *recommends an action* based on the state of nature
  - Such schemes are called direct: signals are actions
- A recommendation should be persuasive: after Bayes update, receiver's favorite action is indeed the recommended action

Fact: There exists an optimal signaling scheme which is direct and persuasive.

Remark:

- A direct persuasive scheme is a *Bayes correlated equilibrium*
- Solving BP is to compute the BCE that maximizes sender utility

Characterizing Optimal Sender Utility

- Receiver belief p → receiver best response → sender utility $\mathrm{V}(p)$

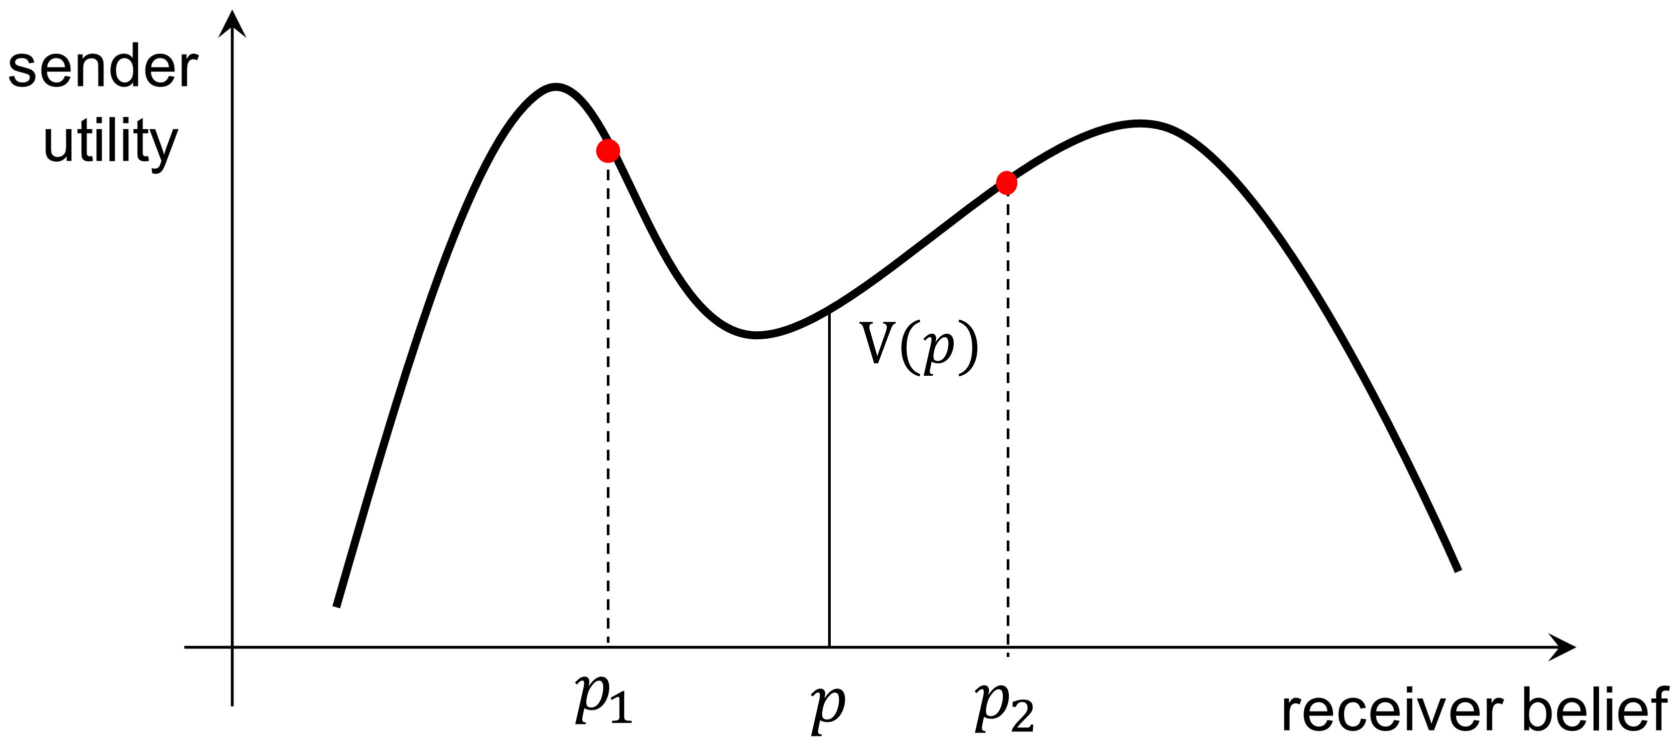

A signaling scheme is a convex decomposition of p

Characterizing Optimal Sender Utility

- Receiver belief p → receiver best response → sender utility $\mathrm{V}(p)$

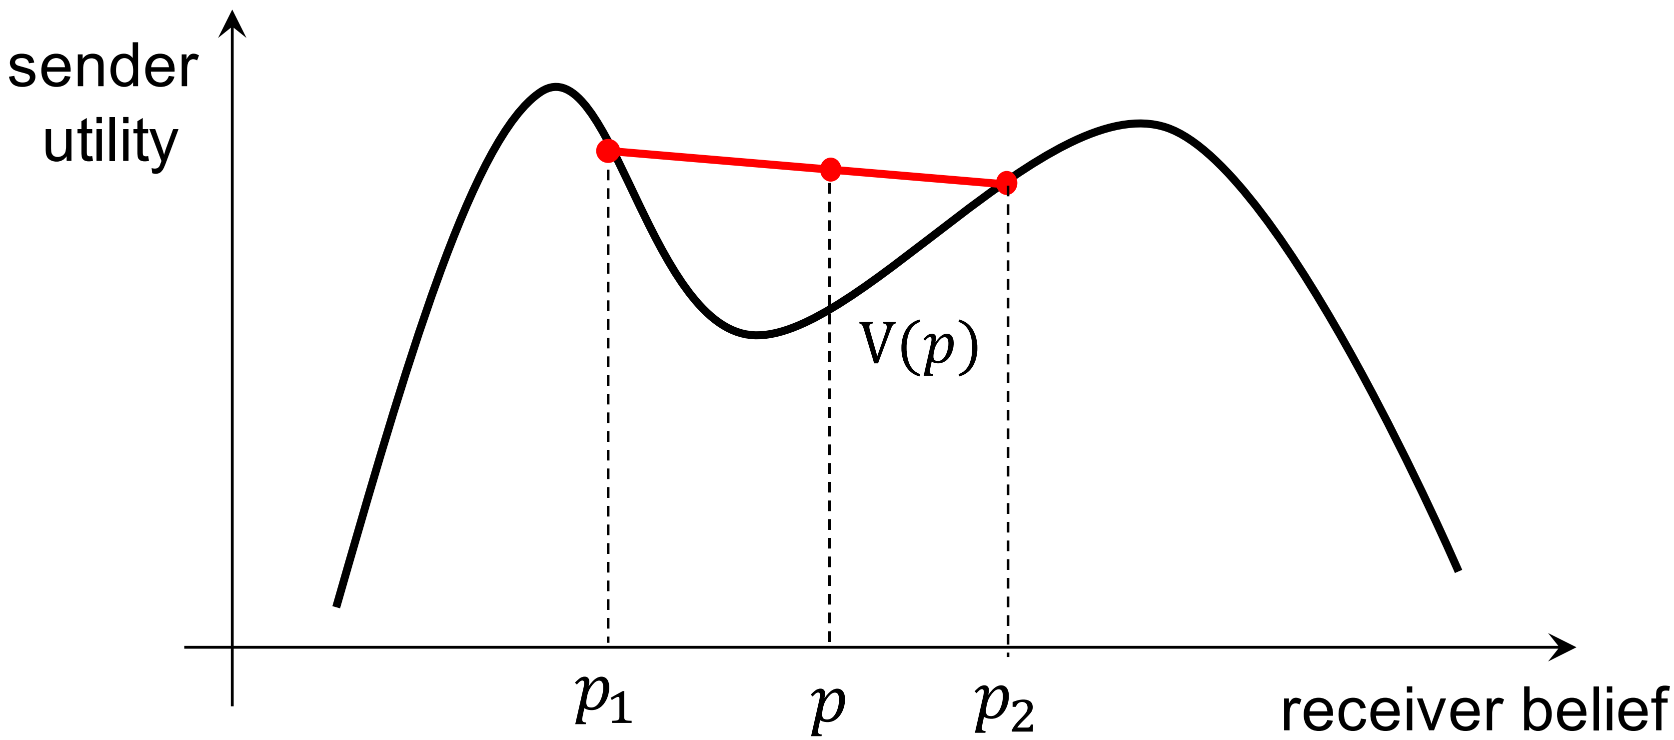

A signaling scheme is a convex decomposition of p

Characterizing Optimal Sender Utility

- Receiver belief p → receiver best response → sender utility $V(p)$

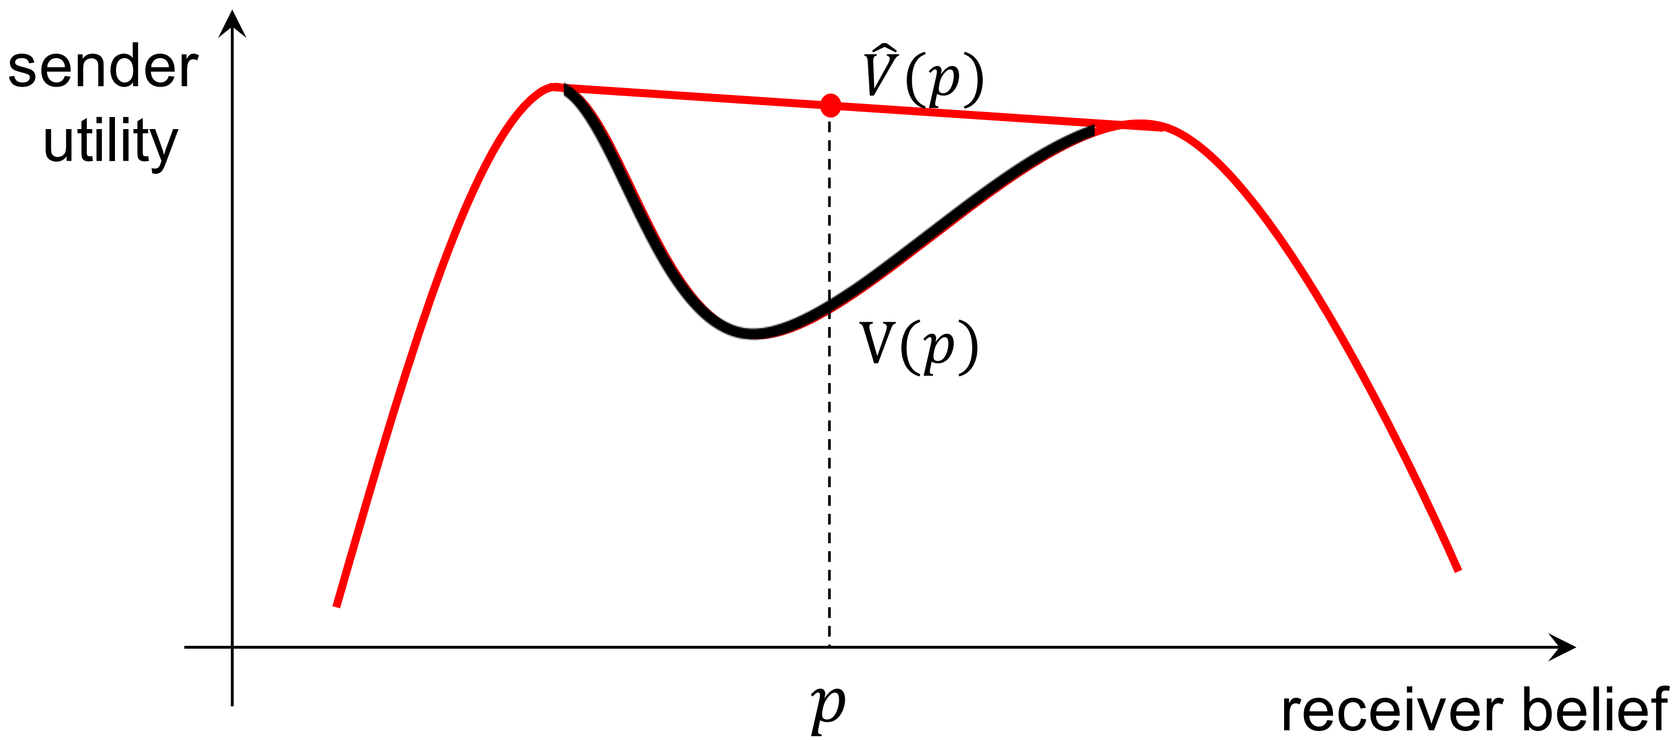

Proposition [KG'11]: for any prior p, the optimal sender utility from persuasion is $\hat{V}(p)$, where $\hat{V}$ is the concave closure of V.

Outline

- Introduction
- Persuasion Through the Algorithmic Lens
- Applications of Persuasion
- The Bigger Picture and Future Directions

Why the Algorithmic Lens?

- Enable automated application
- Lead to structural insights
- Understand possibility and limitation of the model

Some settings are combinatorial by nature, thus require algorithmic techniques

Example: Advisor with Multiple Students

- Advisor has 2 students; recruiter wants to recruit one of them
- Each student's type is independent uniform draw from {L, S, W}
- Student's type determine long / short term achievement

Recruiter utility →
Advisor utility →

	L	S	W
Long-Term	2	$1+\epsilon$	0
Short-Term	0	1	0

Scheme 1: no information

- Students appear identical to the recruiter
- Recruiter randomly chooses a student
- Expected advisor utility 1/3

Example: Advisor with Multiple Students

- Advisor has 2 students; recruiter wants to recruit one of them
- Each student's type is independent uniform draw from {L, S, W}
- Student's type determine long / short term achievement

Recruiter utility → Long-Term
Advisor utility → Short-Term

	L	S	W
Long-Term	2	$1 + \epsilon$	0
Short-Term	0	1	0

Scheme 2: full information

- Good cases for advisor: (S,S), (S,W), (W,S)
- Expected advisor utility: $(1/9)\times 3 = 1/3$

Example: Advisor with Multiple Students

- Advisor has 2 students; recruiter wants to recruit one of them
- Each student's type is independent uniform draw from {L, S, W}
- Student's type determine long / short term achievement

Recruiter utility ⟶ (Long-Term)
Advisor utility ⟶ (Short-Term)

	L	S	W
Long-Term	2	$1 + \epsilon$	0
Short-Term	0	1	0

Scheme 3: optimal (partially informative) scheme

- Properly correlate students' types
  - When there is exactly one type-S student, recommend him
  - Otherwise, recommend a student uniformly at random
- An S-type student is hired whenever S shows up (prob 5/9)

Example: Advisor with Multiple Students

- Advisor has 2 students; recruiter wants to recruit one of them
- Each student's type is independent uniform draw from {L, S, W}
- Student's type determine long / short term achievement

		L	S	W
Recruiter utility →	Long-Term	2	$1 + \epsilon$	0
Advisor utility →	Short-Term	0	1	0

This setting: I.I.D. prior for action types

Explicit Prior

- Probabilities for all states are explicitly enumerated
- Linear program (LP)

Expected sender utility

Obedience constraints

$$
\begin{aligned}
\max \quad & \boxed{\textstyle\sum_\theta \mathbf{P}(\theta) \sum_{i=1}^n x(\theta,i) s(\theta_i)} \\
\text{s.t.} \quad & \boxed{\textstyle\sum_\theta \mathbf{P}(\theta) x(\theta,i)[r(\theta_i) - r(\theta_j)] \geq 0, \quad \text{for } i,j \in [n]} \\
& \textstyle\sum_{i=1}^n x(\theta,i) = 1, \quad \text{for } \theta \\
& x(\theta,i) \geq 0, \quad \text{for } \theta, i
\end{aligned}
$$

Scheme feasibility

Succinct I.I.D. Prior

- State of nature $\theta = (\theta_1, \dots, \theta_n)$
- Action type θ_i's are i.i.d., supported on a discrete set of size m

Theorem [DX'16]: In the i.i.d. model, the optimal signaling scheme can be implemented in $\mathrm{poly}(n, m)$ time.

Structural Insight I

Analogous to single-item auctions

- Actions $\approx$ bidders
- Action types $\approx$ bidder types
- Recommending an action $\approx$ giving the time to a bidder
- Signaling scheme $\approx$ allocation rule with obedience constraints instead of incentive compatibility (IC) constraints

Succinct I.I.D. Prior

- State of nature $\theta = (\theta_1, \dots, \theta_n)$
- Action type θ_i's are i.i.d., supported on a discrete set of size m

Theorem [DX'16]: In the i.i.d. model, the optimal signaling scheme can be implemented in $\mathrm{poly}(n, m)$ time.

Structural Insight II

There exists an optimal scheme that is symmetric

- Each action is recommended w.p. $1/n$
- All (un)recommended actions "look the same" (i.e., have the same posterior type distribution)

Succinct I.I.D. Prior

- State of nature $\theta = (\theta_1, \dots, \theta_n)$
- Action type θ_i's are i.i.d., supported on a discrete set of size m

Theorem [DX'16]: In the i.i.d. model, the optimal signaling scheme can be implemented in $\text{poly}(n, m)$ time.

Proof Outline

- Summarize symmetric schemes via its reduced forms
  - Prob. of recommendation (resp. winning) for each type
- LP over Border's polytope [Border '91] + obedience constraints (linear in reduced form)
  - Solvable efficiently [Alaei et al '12, Cai et al '12]

Succinct Independent Prior

- Generalize i.i.d. model to non-identical actions, explicit marginals
- Border's thm generalizes to non-identical bidders . . .

Theorem [DX'16]: In persuasion with independent actions, it is #P-hard to compute optimal expected sender utility.

Structural Lesson

There is no Border's-theorem-like characterization for persuasion with independent actions

- Also called "generalized border's theorem" [Copalan et al. '15]

Succinct Independent Prior

- Generalize i.i.d. model to non-identical actions, explicit marginals
- Border's thm generalizes to non-identical bidders . . .

Theorem [DX'16]: In persuasion with independent actions, it is #P-hard to compute optimal expected sender utility.

Explanation

- Obedience constraints, unlike IC constraints, are not expressible using "standard" reduced form.
- Any adequate reduced form encodes #P-hard problems, so Toda's theorem implies it cannot exist unless PH collapses

Remark: See [Kolotilin et al. 2017] for other connections and differences between persuasion and mechanism design

General Black-Box Prior

- $\theta = (\theta_1, \dots, \theta_n)$ is drawn from an arbitrary distribution
  - θ_i's can be correlated
- Given to algorithm as a black box

Theorem [DX'16]: For general black-box prior, an ϵ-optimal ϵ-persuasive scheme can be implemented in $\mathrm{poly}(n, 1/\epsilon)$ time.

Bicriteria loss is inevitable for information-theoretic reasons.

General Black-Box Prior

Algorithm Sketch

- Input is a state θ^* drawn from prior
- In addition to θ^*, take $\mathrm{poly}(n, 1/\epsilon)$ samples from black box $\boldsymbol{P}$
- Solve explicit LP on empirical $\widehat{\mathrm{P}}$ (relax to ϵ-obedience)
- Signal as LP suggests for θ^*

$$
\begin{array}{lll}
\text{maximize} & \sum_{\theta} \sum_{i=1}^{n} \widehat{\mathbf{P}}(\theta) x(\theta, i) s(\theta_i) & \\
\text{subject to} & \sum_{\theta} \widehat{\mathbf{P}}(\theta) x(\theta, i)[r(\theta_i) - r(\theta_j)] \geq -\epsilon, & \text{for } i, j \in [n]. \\
& \sum_{i=1}^{n} x(\theta, i) = 1, & \text{for } \theta \in \Theta. \\
& x(\theta, i) \geq 0, & \text{for } \theta \in \Theta, i \in [n].
\end{array}
$$

General Black-Box Prior

Algorithm Sketch

- Input is a state θ^* drawn from prior
- In addition to θ^*, take $\mathrm{poly}(n, 1/\epsilon)$ samples from black box $\boldsymbol{P}$
- Solve explicit LP on empirical $\widehat{\mathrm{P}}$ (relax to ϵ-obedience)
- Signal as LP suggests for θ^*

Structural Insight

To query scheme locally (i.e., sample $X(\theta^*)$), little context is needed

- Small sample complexity

Outline

- Introduction
- Persuasion Through the Algorithmic Lens
- Applications of Persuasion
- The Bigger Picture and Future Directions

Domain I: Security Games

- A defender faces a strategic adversary
- Defender seeks to protect critical targets from adversary's attack



Key idea: defender can utilize informational advantage to deceive adversary and improve defense

Example: UAVs for Conservation

- Illegal poaching is a major threat to endangered animals
- UAVs to combat poaching

Air Shepherd
The Lindbergh Foundation

Example: UAVs for Conservation

Air Shepherd
The Lindbergh Foundation

UAV video

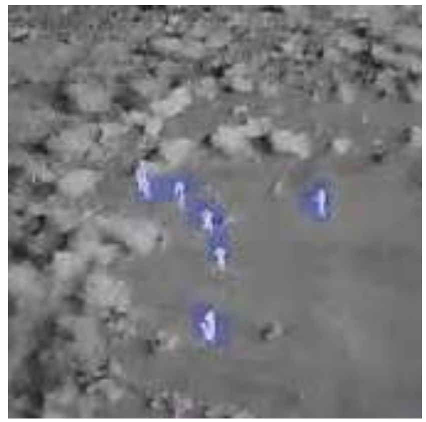

Automatic poacher detection

?

UAVs can not directly catch poachers

- ➢ Very few rangers are available and nearby
  - ❖ Poaching usually happens during night

Exploit Informational Advantage

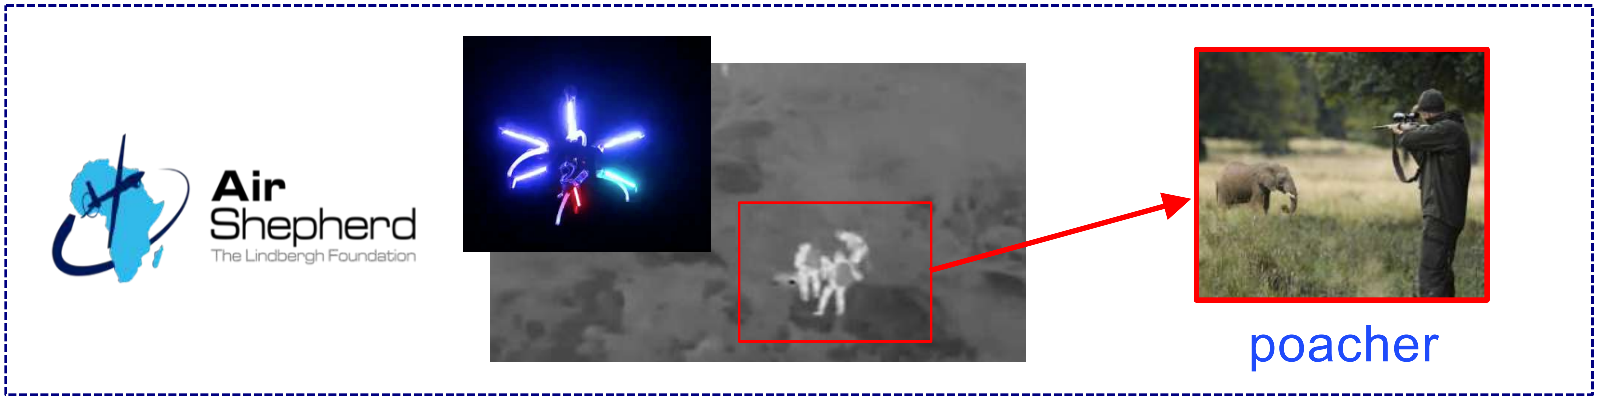

- Poacher only knows that rangers come with certain probability
- Air Shepherd knows precisely whether a ranger will come or not
- Approach: use alerting signal to deter poaching
  - Signals correlate with the presence of rangers
  - Deceptive alerting: may alert even no rangers are nearby

Additional Challenge due to Domain Features

UAVs

Rely on

Rangers

Alerting: send (*deceptive*) alerting signals

Interdiction: directly interdict poaching

Additional Challenge due to Domain Features

Coordinate rangers and UAVs + Deceptive alerting via persuasion

Global optimality

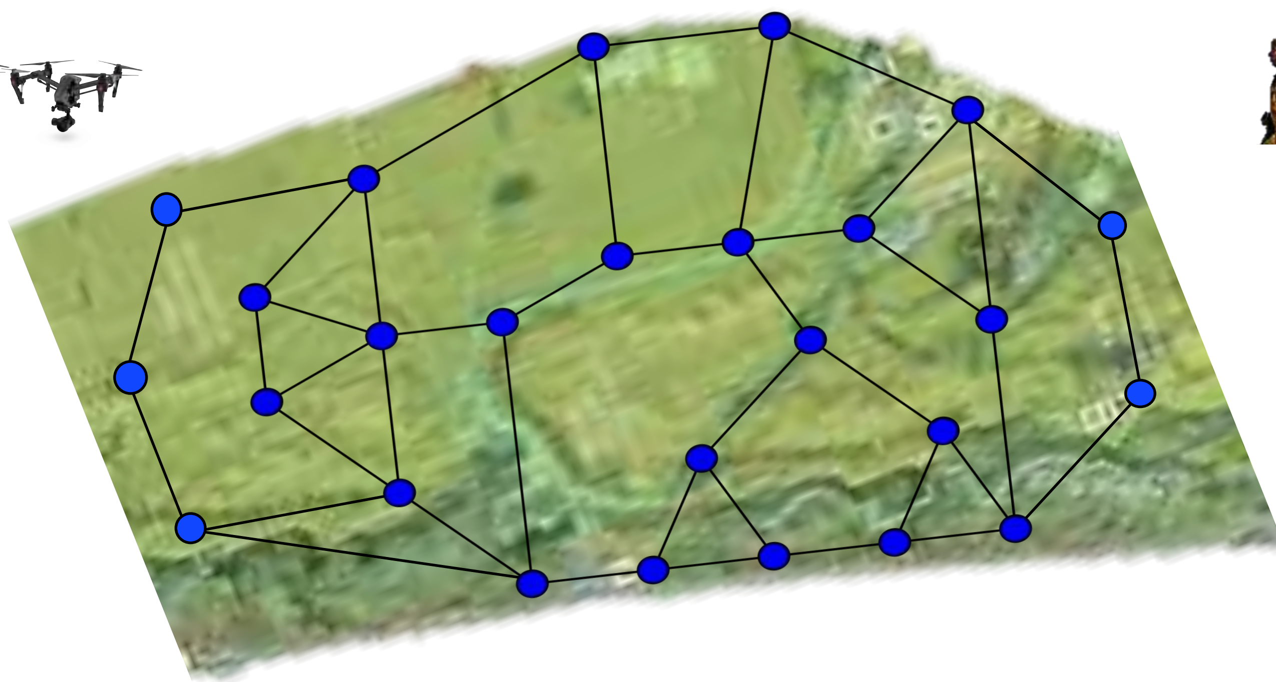

Additional Challenge due to Domain Features

Coordinate rangers and UAVs + Deceptive alerting via persuasion

Global optimality

- Novel model that integrates patroller's interdiction and UAVs' deceptive alerting functionality [Xu et al. 18]
- Algorithmic study: complexity analysis and scalable algorithms

Utility Comparisons

- Simulated games
- Y-axis: defender loss (lower is better)

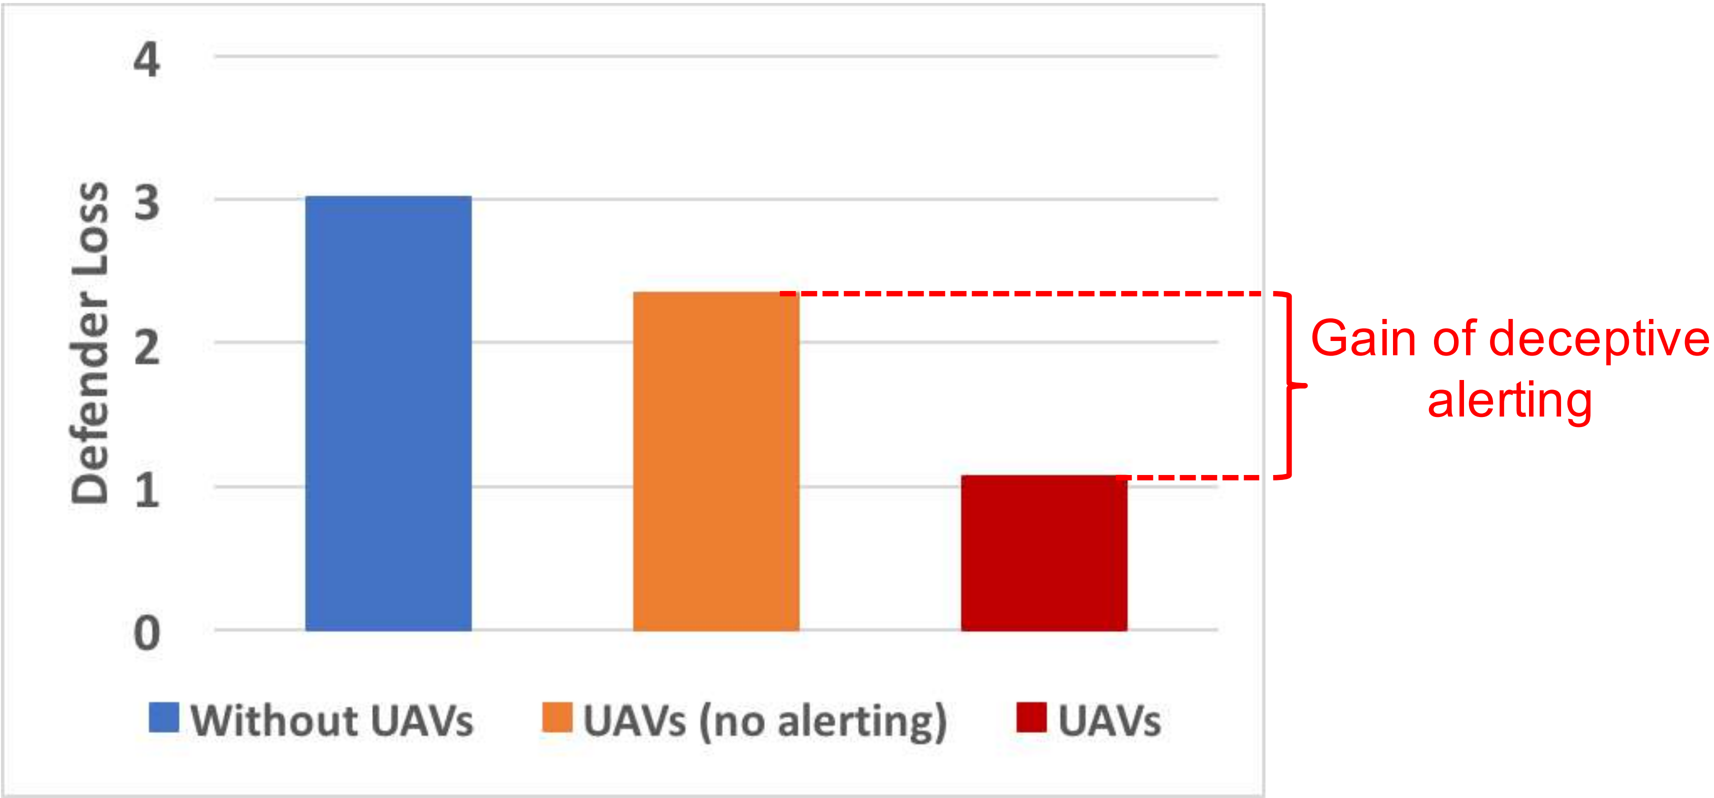

Utility Comparisons

- Simulated games
- Y-axis: defender loss (lower is better)

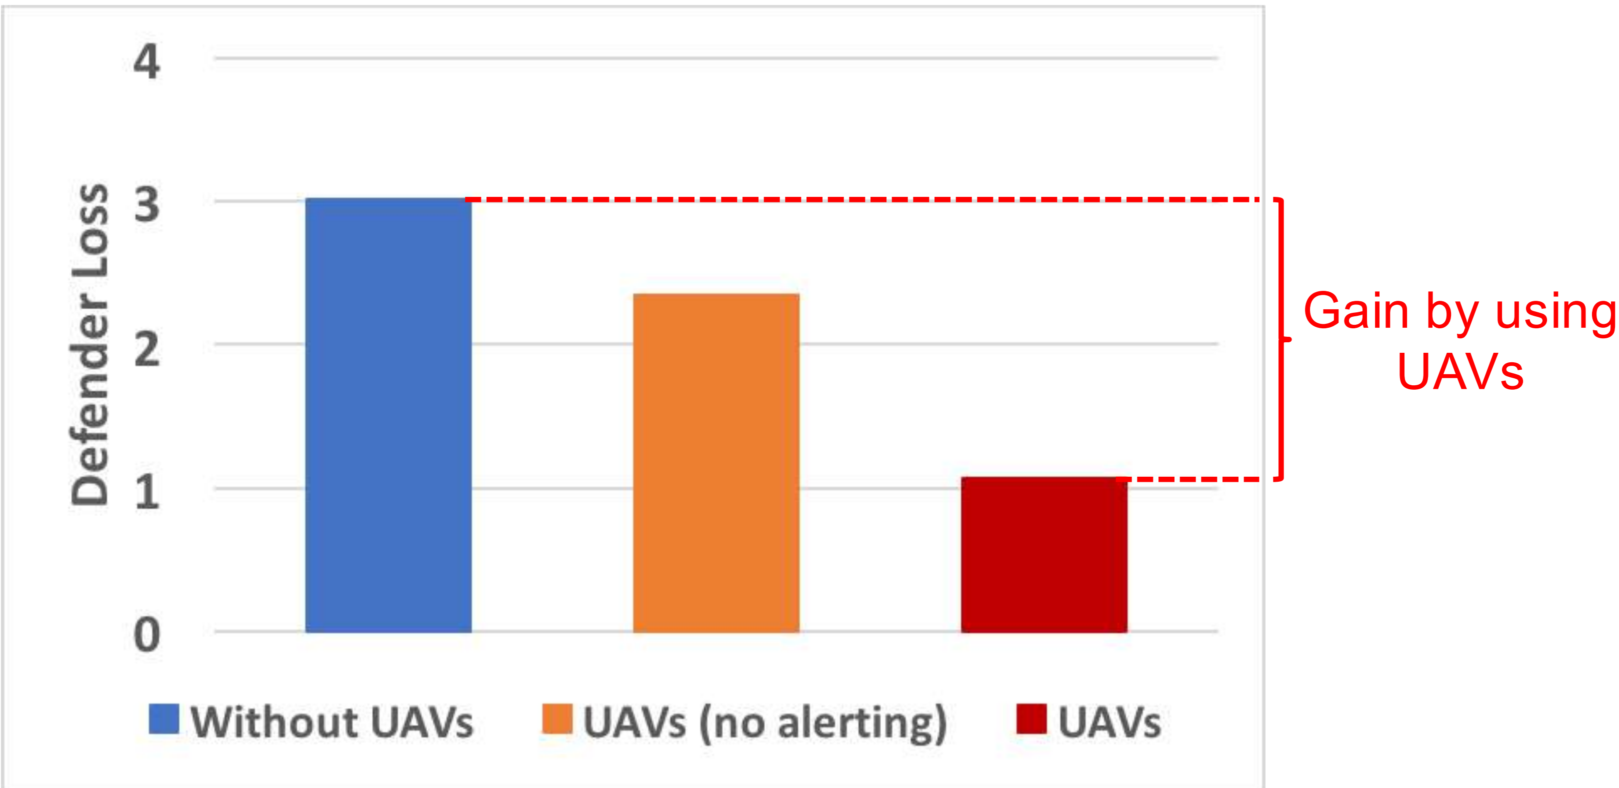

Other Applications in Security Games

- Prevent fare evasion in honor metro systems [Xu et al. 2015]
  - Alert passengers with warning signals

Other Applications in Security Games

- Prevent fare evasion in honor metro systems [Xu et al. 2015]
  - Alert passengers with warning signals
- Reduce illegal parking [Hernández/Neeman 2017]

Other Applications in Security Games

- Prevent fare evasion in honor metro systems [Xu et al. 2015]
  - Alert passengers with warning signals
- Reduce illegal parking [Hernández/Neeman 2017]
- Protect cyber systems [Schlenker et al. 2018]

Other Applications in Security Games

- Prevent fare evasion in honor metro systems [Xu et al. 2015]
    - Alert passengers with warning signals
- Reduce illegal parking [Hernández/Neeman 2017]
- Protect cyber systems [Schlenker et al. 2018]
- New security game models with deception [Rabinovich et al. 2015, Xu et al. 2016]

Domain II: Auctions

- E.g., ad auctions (auctions for selling online ad slots)
- Uncertainty and information asymmetry

Advertiser's perspective:

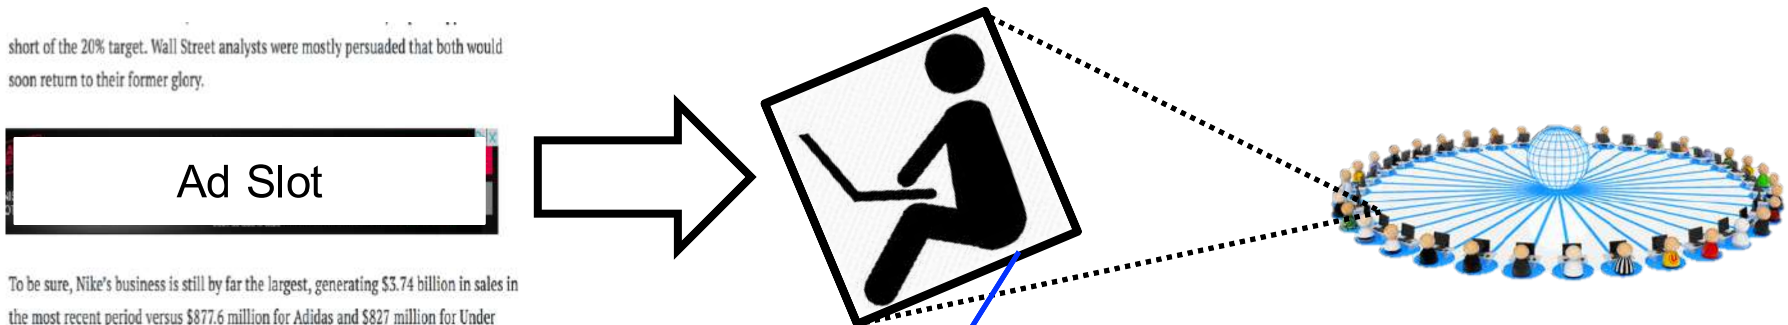

A web user

The Internet

Auctioneer's perspective:

Persuasion in Auctions

- Reveal information to bidders in order to influence their bidding
  - Can reveal different information to different bidders
- Difficulty: intricate equilibrium behavior
  - Players' actions are affected by information and others' actions
  - Issues of equilibrium selection

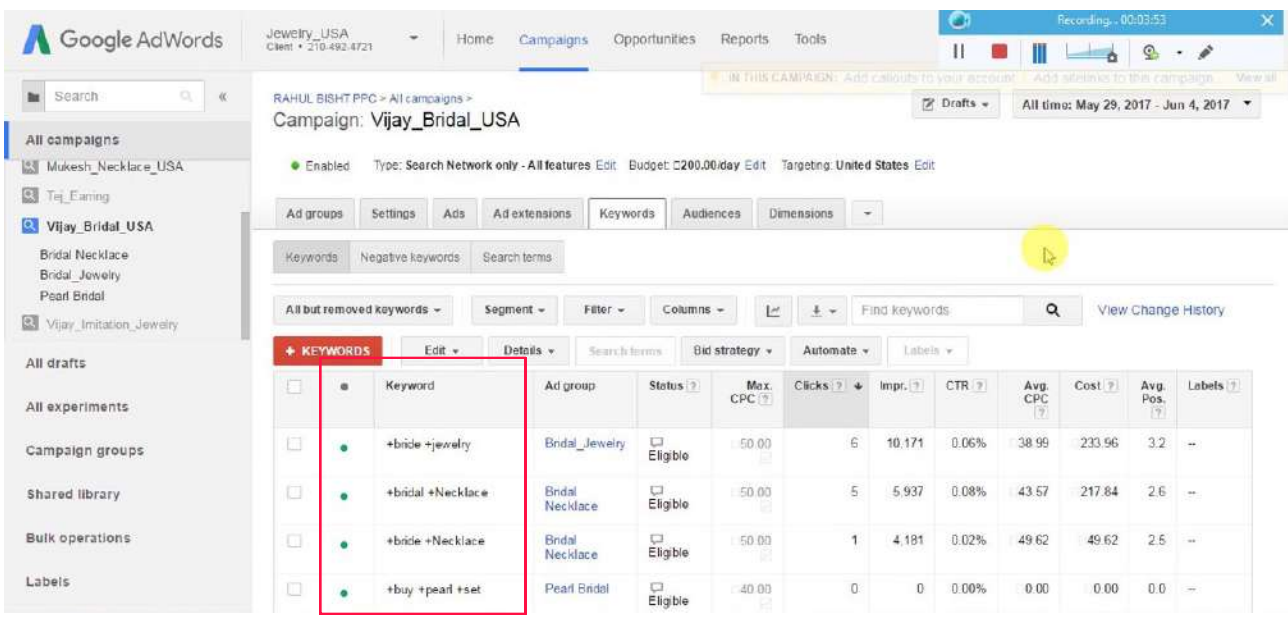

Two Natural Types of Signaling Schemes

Public scheme: must send the signal publicly to all bidders	Private scheme: may send different signals to different bidders
➢ Equivalently, send the same signal to every bidder ➢ Due to, e.g., fairness or communication constraints	➢ Signals may be correlated to steer desired collective behavior

Example: Public/Private Schemes, Intricacies

		Bidder 1	Bidder 2	Bidder 3
State 1	prob: $1-\epsilon$	2ϵ	ϵ	1
State 2	prob: ϵ	1	$1-\epsilon$	ϵ

Single-item second-price auction

Fact: In public schemes, bidding the expected true value is a dominant strategy for each bidder.

Typical to adopt the dominant-strategy equilibrium in public schemes

Claim: Revenue of the optimal public scheme is at most 3ϵ.

Example: Public/Private Schemes, Intricacies

		Bidder 1	Bidder 2	Bidder 3
State 1	prob: $1-\epsilon$	2ϵ	ϵ	1
State 2	prob: ϵ	1	$1-\epsilon$	ϵ

Private scheme: full info to bidder 2 and 3; no info to bidder 1

Truthful bidding is not an equilibrium

- Bidder 2 and 3: bidding true value is a dominant strategy
- But, bidder 1 will *not* bid her expected value ($\approx 3\epsilon$)
  - Equilibrium bid for bidder 1 is any $b \in (1-\epsilon, 1)$

Example: Public/Private Schemes, Intricacies

		Bidder 1	Bidder 2	Bidder 3
State 1	prob: $1-\epsilon$	2ϵ	ϵ	1
State 2	prob: ϵ	1	$1-\epsilon$	ϵ

Private scheme: full info to bidder 2 and 3; no info to bidder 1

Large revenue gap between public and private scheme

- Claim: Revenue of above scheme $\geq 1-\epsilon$
- Recall: revenue of optimal public signaling $\leq 3\epsilon$

Example: Public/Private Schemes, Intricacies

		Bidder 1	Bidder 2	Bidder 3
State 1	prob: $1-\epsilon$	2ϵ	ϵ	1
State 2	prob: ϵ	1	$1-\epsilon$	ϵ

Private scheme: full info to bidder 2 and 3; no info to bidder 1

This private scheme extracts almost full surplus

- Rev $\geq 1-\epsilon$
- Social surplus $= 1$

See [Badanidiyuru/Bhawalkar/Xu 18] for a general conclusion.

(Some) Works in Auctions

- One bidder/buyer
  - [Bergemann/Brooks/Morris 2015, Shen/Tang/Zeng 2018] characterize all possible revenue-welfare tradeoffs of persuasion
- Multiple bidders
  - [Fu et al. 12]: revealing full information is the optimal public scheme in Myerson's optimal auction
  - [Emek et al. 12] and [Miltersen/Sheffet 12]: how to compute optimal public scheme in second price auctions?
  - [Dughmi et al. 15]: contrained public scheme by revealing only a subset of item features
  - [Badanidiyuru/Bhawalkar/Xu 18]: both private and public schemes in a setting motivated by ad auctions
  - . . . .

Outline

- Introduction
- Persuasion Through the Algorithmic Lens
- Applications of Persuasion
- The Bigger Picture and Future Directions

Information Structure Design

Influence equilibria by designing “who knows what”

	One receiver	Multiple receivers
No player private information	Bayesian persuasion & applications	
With player private information		

Main focus of this part
(except auctions)

Information Structure Design

Influence equilibria by designing "who knows what"

	One receiver	Multiple receivers
No player private information	Bayesian persuasion & applications	Many basic models & applications
With player private information		

- Basic models:
  - [Dughmi 2014]: two receivers playing a zero-sum game
  - [Taneva 2016]: two receivers, two actions, two states
  - [Arieli/Babichenko 2016, Dughmi/Xu 2018]: multiple receivers, no externalities, binary actions

Persuading Multiple Receivers: Applications

[Chan/Li/Wang ’15] (private signals),
[Alonso/Camara’15] (public signals),
[Cheng et al.’15] algorithmic study

Routing games
[Bhaskar et al.’16]

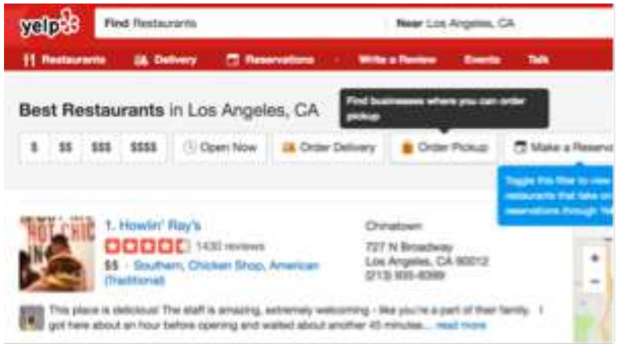

Recommendation systems
[Mansour/Slivinks/Syrgkanis’15,
Mansour et al. ‘16]

Queueing with strategic customers
[Lingenbrink/Iyer’17]

Information Structure Design

Influence equilibria by designing "who knows what"

	One receiver	Multiple receivers
No player private information	Bayesian persuasion & applications	Basic models and applications
With player private information	Limited work, relate to mechanism design	

- The receiver holds private information about the random state
- The sender may elicit receiver's private information
  - [Kolotilin et al. 2017] studies a basic model and relate its structure to mechanism design
  - [Xu et al. 2016] studies persuasion in Bayesian Stackelberg games

Information Structure Design

Influence equilibria by designing “who knows what”

	One receiver	Multiple receivers
No player private information	Bayesian persuasion & applications	Basic models and applications
With player private information	Limited work, relate to mechanism design	Much less explored

Information Structure Design

Influence equilibria by designing “who knows what”

	One receiver	Multiple receivers
No player private information	Bayesian persuasion & applications	Basic models and applications
With player private information	Limited work, relate to mechanism design	Much less explored

- See [Bergemann/Morris 2017] for a unified perspective and examples
- Models with multiple senders [Kamenica/Gentzkow 2017]

Future Directions

- Information structure design with multiple receivers
  - Basic questions: persuade multiple receivers with externalities, private information, and multiple actions
  - Dynamic/repeated settings
  - Constraints/costs on signaling schemes
  - Applications

- Value and pricing of information
  - Connects to the first part of this tutorial
  - Consider the “data → information” procedure
  - Relation between persuasion and information elicitation?

Thank You

Questions?

References (1/4)

- [Kamenica/Gentzkow 11]: Emir Kamenica and Matthew Gentzkow. *Bayesian persuasion.* American Economic Review, 2011.
- [DX'16]: Shaddin Dughmi and Haifeng Xu. *Algorithmic Bayesian persuasion*. STOC 2016.
- [Xu et al. 18]: Haifeng Xu, Kai Wang, Phebe Vayanos, Milind Tambe. *Strategic Coordination of Human Patrollers and Mobile Sensors with Signaling for Security Games*. AAAI 2018.
- [Xu et al. 2015]:Haifeng Xu, Zinovi Rabinovich, Shaddin Dughmi, Milind Tambe. *Exploring Information Asymmetry in Two-Stage Security Games*. AAAI 2015.
- [Hernández/Neeman 2017]: Penélope Hernández and Zvika Neeman. *How Bayesian Persuasion can Help Reduce Illegal Parking and Other Socially Undesirable Behavior.* Working paper, 2017.
- [Schlenker et al. 2018]: Aaron Schlenker, Omkar Thakoor, Haifeng Xu, Fei Fang, Milind Tambe, Long Tran-Thanh, Phebe Vayanos, Yevgeniy Vorobeychik. *Deceiving Cyber Adversaries: A Game Theoretic Approach.* AAMAS 2018.
- [Rabinovich et al. 2016]: Zinovi Rabinovich, Albert X. Jiang, Manish Jain, Haifeng Xu. *Information Disclosure as a Means to Security*. AAMAS 2015.
- [Badanidiyuru/Bhawalkar/Xu 18]: Ashwinkumar Badanidiyuru, Kshipra Bhawalkar, Haifeng Xu. *Targeting and Signaling in Ad Auctions.* SODA 2018.

References (2/4)

- [Bergemann/Brooks/Morris 2015]: Dirk Bergemann, Benjamin Brooks, Stephen Morris. *The Limits of Price Discrimination*. American Economic Review, 2015.
- [Shen/Tang/Zeng 2018]: Weiran Shen, Pingzhong Tang, Yulong Zeng. *A Closed-Form Characterization of Buyer Signaling Schemes in Monopoly Pricing.* AAMAS 2018.
- [Fu et al. 2012]: Hu Fu, Patrick Jordan, Mohammad Mahdian, Uri Nadav, Inbal Talgam-Cohen, and Sergei Vassilvitskii. *Ad Auctions with Data?* SAGT 2012.
- [Emek et al. 12]: Yuval Emek, Michal Feldman, Iftah Gamzu, Renato Paes Leme, Moshe Tennenholtz. *Signaling Schemes for Revenue Maximization.* EC 2012.
- [Miltersen/Sheffet 12]: Peter Bro Miltersen and Or Sheffet. *Send mixed signals: earn more, work less.* EC 2012.
- [Dughmi et al. 15]: Shaddin Dughmi, Nicole Immorlica, Ryan O'Donnell, and Li-Yang Tan. *Algorithmic Signaling of Features in Auction Design.* SAGT 2015.
- [Dughmi 2014]: Shaddin Dughmi. *On the Hardness of Signaling.* FOCS 2014.
- [Taneva 2016]: Taneva, I. *Information Design*. Discussion paper, 2016.
- [Arieli/Babichenko 2016]: Itai Arieli, Yakov Babichenko. *Private Bayesian Persuasion.* Working paper, 2016.

References (3/4)

- [Dughmi/Xu 2018]: Shaddin Dughmi, Haifeng Xu. *Algorithmic Persuasion with No Externalities.* EC 2017.
- [Chan/Li/Wang '15]: Jimmy Chan, Fei Li, and Yun Wang. *Discriminatory Persuasion: How to Convince a Group.* Working paper, 2015.
- [Alonso/Camara'15]: Ricardo Alonso, Odilon Câmara. *Persuading Voters.* American Economic Review, 2016.
- [Cheng et al.'15]: Yu Cheng, Ho Yee Cheung, Shaddin Dughmi, Ehsan Emamjomeh-Zadeh, Li Han, Shang-Hua Teng. *Mixture Selection, Mechanism Design, and Signaling*. FOCS 2015.
- [Bhaskar et al.'16]: Yu Cheng, Umang Bhaskar, Young Kun Ko, Chaitanya Swamy. *Hardness of Signaling in Bayesian Games*. EC 2016.
- [Mansour/Slivinks/Syrgkanis'15]: Yishay Mansour, Aleksandrs Slivkins, Vasilis Syrgkanis. *Bayesian Incentive-Compatible Bandit Exploration*. EC 2015.
- [Mansour et al. '16]: Yishay Mansour, Aleksandrs Slivkins, Vasilis Syrgkanis., and Steven Wu. *Bayesian Exploration: Incentivizing Exploration in Bayesian Games*. EC 2016.

References (4/4)

- [Lingenbrink/Iyer'17]: D. Lingenbrink, K. Iyer, *Optimal Signaling Mechanisms in Unobservable Queues with Strategic Customers, EC 2017.*
- [Kolotilin et al. 2017]: Anton Kolotilin, Tymofiy Mylovanov, Andriy Zapechelnyuk. *Persuasion of a Privately Informed Receiver*. Econometrica 2017.
- [Xu et al. 2016]: Haifeng Xu, Rupert Freeman, Vincent Conitzer, Shaddin Dughmi and Milind Tambe. *Signaling in Bayesian Stackelberg Games.* AAMAS 2016.
- [Bergemann/Morris 2017]: Dirk Bergemann and Stephen Morris. *Information Design: A Unified Perspective*. Working paper, 2017.
- [Kamenica/Gentzkow 2017]: Emir Kamenica and Matthew Gentzkow. *Bayesian persuasion with multiple senders and rich signal spaces*. Games and Economic Behavior, 2017.