Ambiguous Contracts

Michal Feldman
Tel Aviv University

ADFOCS 2024

MPI Summer School

August 26, 2024

Saarbrücken, Germany

Joint work with:

Paul Duetting, Daniel Peretz, Larry Samuelson

[EC'23, ECMA forthcoming]

Ambiguous Contracts

- In many contractual relations, contracts are “ambiguous”. E.g.,
  - “We’ll grade one question in each problem set” (professors)
  - “we’ll compensate good drivers” (insurance companies)
  - “you’ll get promoted if you perform well” (companies/academic faculty)
- Motivating question: Why are ambiguous contracts so common?
- We study the power of ambiguity in contract design
  - Lots of work in economic and algorithmic design on ambiguity as a constraint
  - We study ambiguity as a tool --- namely, the deliberate infusion of ambiguity into the design of contracts (inspired by [Di Tillio et al. REStud 2017] who study ambiguity in auction design)

Ambiguous contracts

Ambiguous Contracts

- An ambiguous contract is a set of classic contracts $\tau = (t^1, \ldots, t^k)$
  - $t^i = (t^i_1, \ldots, t^i_m)$ for every i
- Agent is ambiguity averse: selects an action, $i^*(\tau)$, whose minimal expected utility across all contracts t ∈ τ is the highest

$$i^*(\tau) \in \arg\max_{i \in [n]} \underbrace{\min_{t \in \tau} U_A(i, t)}_{U_A(i, \tau)}$$

[breaking ties in favor of principal]

- Consistency: principal is indifferent between all contracts $t \in \tau$ w.r.t. action $i^*(\tau)$, i.e., for any two contracts $t^j, t^l \in \tau$:

$$U_P\left(i^*(\tau), t^j\right) = U_P\left(i^*(\tau), t^l\right)$$

(also implies same payment and same agent's utility for any two contracts in τ)

Timeline

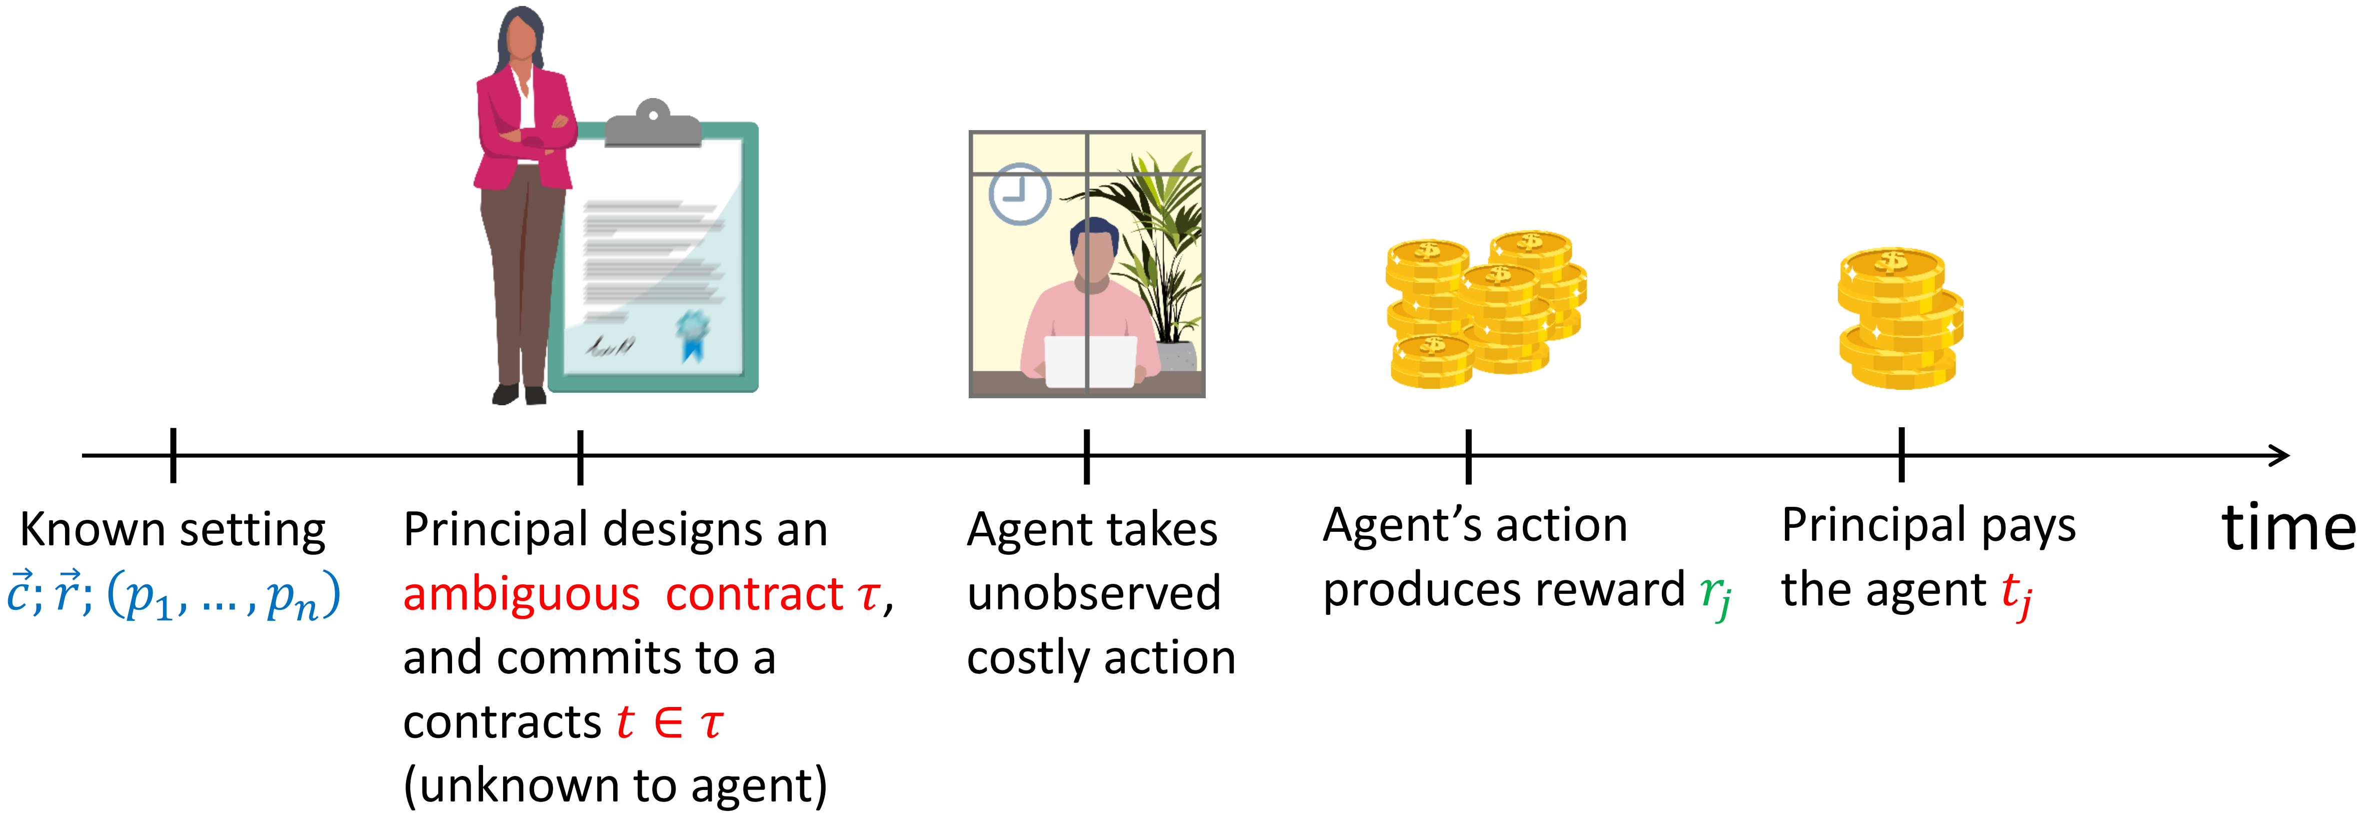
Known setting
$\vec{c}; \vec{r}; (p_1, \ldots, p_n)$
Principal designs an ambiguous contract τ, and commits to a contracts $t \in \tau$ (unknown to agent)
Agent takes unobserved costly action
Agent's action produces reward r_j
Principal pays the agent t_j
time

Example

	Cost	$r_1 = 2$	$r_2 = 2$
Action 1	1/4	1/2	0
Action 2	1/4	0	1/2
Action 3	1	1/2	1/2

Principal's utility = 1-1/4 = 3/4

Best **classic** contract:

- Incentivize action 1
- $t = (^1/_2, 0)$
- Expected payment = $\mathbf{\frac{1}{4}}$

Principal's utility = 2-1 = 1

Let $\tau = (t^1, t^2)$ be **ambiguous** with $t^1 = (2,0)$ and $t^2 = (0,2)$

- Action 1 gives agent's utility -1/4 (under t^2)
- Action 2 gives agent's utility -1/4 (under t^1)
- Action 3 gives agent's utility 0 (expected payment of $1/2 * 2 = \mathbf{1}$ under both contracts)

Upshot: principal can gain by employing ambiguous contracts

Principal's utility = 1-1/4 = 3/4

	Cost	$r_1 = 2$	$r_2 = 2$
Action 1	1/4	1/2	0
Action 2	1/4	0	1/2
Action 3	1	1/2	1/2

Principal's utility = 2-1 = 1

Best **classic** contract:

- Incentivize action 1
- $t = (^1/_2, 0)$
- Expected payment = $\mathbf{\frac{1}{4}}$

Let $\tau = (t^1, t^2)$ be **ambiguous** with $t^1 = (2,0)$ and $t^2 = (0,2)$

- Action 1 gives agent's utility -1/4 (under t^2)
- Action 2 gives agent's utility -1/4 (under t^1)
- Action 3 gives agent's utility 0 (expected payment of $1/2 * 2 = \mathbf{1}$ under both contracts)

Many Questions Arise…

- Can ambiguous contracts benefit both the principal and the agent?
- What's the structure of the optimal ambiguous contract?
- What's the computational hardness of computing the optimal ambiguous contract?
- Are there classes of contracts that are "ambiguity-proof"?
- How much can the principal gain by employing ambiguous contracts?
- What is the effect of mixed strategies by the agent?

Ambiguity can Benefit both Principal and Agent

rewards:	$r_1 = 0$	$r_2 = 9$	$r_3 = 9$	costs
action 1:	1	0	0	$c_1 = 0$
action 2:	0.6	0.3	0.1	$c_2 = 0.6$
action 3:	0.6	0.1	0.3	$c_3 = 0.6$
action 4:	0.2	0.4	0.4	$c_4 = 3$

Example 5 (Ambiguous contracts may benefit both principal and agent)**.** Consider the instance shown in Figure 5. An optimal classic contract is $\langle(0, 2, 0), 2\rangle$, implementing action 2 with utilities 0 and 3 to the agent and principal. The ambiguous contract $\langle\{(0, 8, 0), (0, 0, 8)\}, 4\rangle$ implements action 4 with utilities 0.2 and 4 to the agent and principal. ■

Structure and computation

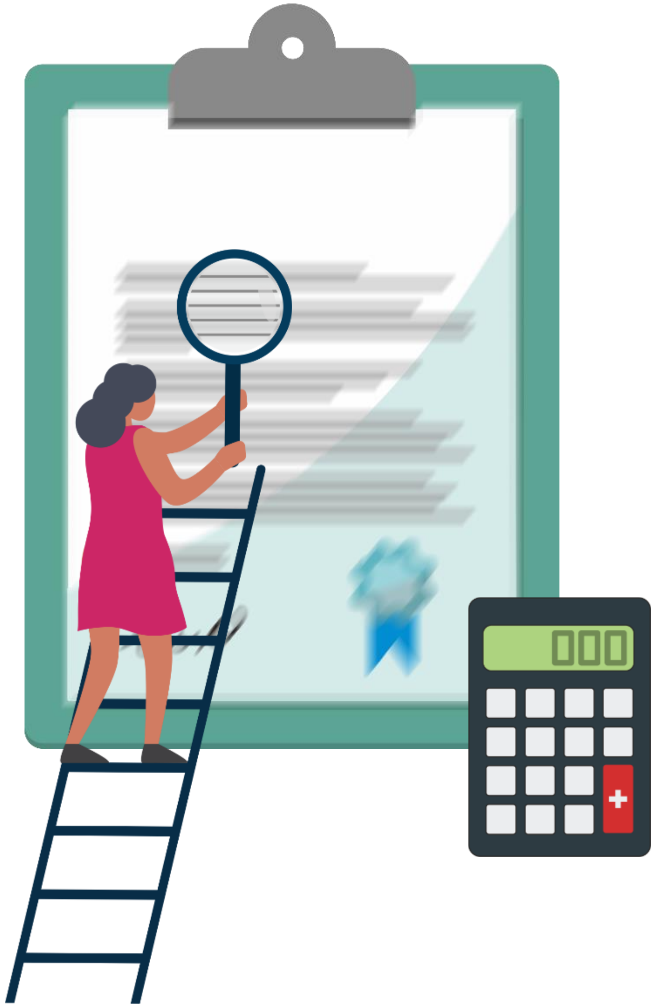

What's the structure and computational hardness of the optimal ambiguous contract?

Single-outcome-payment (SOP) contracts

Definition: an SOP contract is one that pays only for a single outcome, e.g., $t = (0,0,4,0)$

Theorem (informal): For any ambiguous contract τ there's an "equivalent" ambiguous contract τ' composed of SOP contracts

Theorem (formal): For any ambiguous contract τ there's an ambiguous contract $\hat{\tau}$ composed of at most $\min\{n-1, m\}$ SOP contracts such that:

- $i^*(\tau) = i^*(\hat{\tau})$ [τ and $\hat{\tau}$ incentivize the same action]
- $T_{i^*(\tau)}(\tau) = T_{i^*(\tau)}(\hat{\tau})$ [they do so for the same expected payment]

Remark: an analogous theorem for monotone contracts, with step contracts instead of SOPs

Proof Idea

For every action $i \neq i^*$, there exists a contract $t^i \in \tau$ such that

$$U_A(i, t^i) \leq U_A(i^*, t^i) = U_A(i^*, \tau)$$

Plan: modify t^i to an SOP contract $\hat{t}^i$ such that:

- $T_{i^*}(\hat{t}^i) = T_{i^*}(\tau)$ (action i^* has the same E[payment] in $\hat{t}^i$ as in τ)
- $T_i(\hat{t}^i) \leq T_i(t^i)$ (action i has E[payment] in $\hat{t}^i$ at most as in t^i)

We get: $U_A(i, \hat{t}^i) \leq U_A(i, t^i) \leq U_A(i^*, \tau) = U_A(i^*, \hat{\tau})$ (so i^* is incentivized)

Constructing $\hat{t}^i$: Set $\hat{t}^i_{j_{max}} = \frac{T_{i^*}(\tau)}{p_{i^*, j_{max}}}$ and $\hat{t}^i_j = 0$ for all $j \neq j_{max}$,

where $j_{max} \in \arg\max_{j \in m} \frac{p_{i^*, j}}{p_{i,j}}$

Optimal Ambiguous Contract Computation

Theorem: There exists an algorithm that computes the optimal ambiguous contract in time $O(n^2 m)$

Proof idea:

Fix action i.

Lemma 1: If there exists an action $i' \neq i$ such that $p_{i'} = p_i$ and $c_{i'} < c_i$, then action i is not implementable by an ambiguous contract

Lemma 2: Else, action i is implementable, and the optimal ambiguous contract implementing it can be found in time $O(nm)$

Remark: note characterization for implementability by ambiguous contracts

Detour: Characterization of Implementable Actions

Theorem: Action i is implementable with a classic contract if and only if there does not exist a convex combination $\lambda_{i'} \in [0,1]$ of the actions $i' \neq i$ that yields the same distribution over rewards $\Sigma_{i' \neq i} \lambda_{i'} p_{i'j} = p_{ij}$ for all j but at a strictly lower cost $\Sigma_{i' \neq i} \lambda_{i'} c_{i'} < c_i$

Theorem: Action i is implementable with an ambiguous contract if and only if there is no other action $i' \neq i$ such that $p_{i'} = p_i$ but $c_{i'} < c_i$

Example: action 4 can't be implemented by a classic contract, but can be implemented by an ambiguous contracts

rewards:	$r_1 = 0$	$r_2 = 2$	$r_3 = 2$	costs
action 1:	1	0	0	$c_1 = 0$
action 2:	0	1	0	$c_2 = 1$
action 3:	0	0	1	$c_3 = 1$
action 4:	0	1/2	1/2	$c_4 = 3$

Ambiguous contract incentivizing action 4:

$\tau = (t^1, t^2)$ with
$t^1 = (0,6,0)$ and $t^2 = (0,0,6)$

Optimal Ambiguous Contract Computation

Lemma 2: Else (for every action i' with $c_{i'} < c_i$ it holds that $p_{i'} \neq p_i$), action i is implementable, and the optimal ambiguous contract implementing it can be found in time $O(nm)$

Proof: Algorithm for implementable action i:

Let $A = \{i' \neq i \mid p_{i'} \neq p_i\}$. (assume $A \neq \emptyset$, else pay 0)

For each $i' \in A$, let $j(i')$ be a maximizer of $\frac{p_{ij(i')}}{p_{i'j(i')}}$.

Let $$T = \max_{i' \in A} \left\{ \min \left\{ x \geq 0 \,\middle|\, p_{ij(i')} \cdot \frac{x}{p_{ij(i')}} - c_i \geq p_{i'j(i')} \cdot \frac{x}{p_{ij(i')}} - c_{i'} \right\} \right\}$$

For each $i' \in A$, Let $t^{i'} = (0, \dots, T/p_{ij(i')}, 0, \dots, 0)$ [positive payment in index $j(i')$]

Claim 1: Ambiguous contract $\tau = \left\{ t^{i'} \right\}_{i' \in A}$ implements action i.

Claim 2: This contract is the optimal ambiguous contract implementing action i.

Proof in exercise session

Ambiguity Proofness

Are there classes of contracts that are "immune to ambiguous contracts"?

Ambiguity Proofness

Definition: A class of contracts $\mathcal{T}$ is ambiguity-proof if for any instance, any action i, and any ambiguous contract $\tau \in \mathcal{T}$, τ cannot incentivize action i at a strictly lower cost than any single contract in $\mathcal{T}$

Recall example

Principal's utility = 1-1/4 = 3/4

	Cost	$r_1 = 2$	$r_2 = 2$
Action 1	1/4	1/2	0
Action 2	1/4	0	1/2
Action 3	1	1/2	1/2

Principal's utility = 2-1 = 1

Best **classic** contract:

- Incentivize action 1
- $t = (\frac{1}{2}, 0)$
- Expected payment = $\mathbf{\frac{1}{4}}$

Let $\tau = (t^1, t^2)$ be **ambiguous** with $t^1 = (2,0)$ and $t^2 = (0,2)$

- Action 1 gives agent's utility -1/4 (under t^2)
- Action 2 gives agent's utility -1/4 (under t^1)
- Action 3 gives agent's utility 0 (expected payment of $1/2 * 2 = \mathbf{1}$ under both contracts)

Ambiguity Proofness

Definition: A class of contracts $\mathcal{T}$ is ambiguity-proof if for any instance, any action i, and any ambiguous contract $\tau \in \mathcal{T}$, τ cannot incentivize action i at a strictly lower cost than any single contract in $\mathcal{T}$

Definition: A class of contracts $\mathcal{T}$ is ordered iff for any two contracts $t, t' \in \mathcal{T}$ it holds that:

$t(x) \geq t'(x)$ for all $x \in \mathcal{R}^+$ OR $t(x) \leq t'(x)$ for all $x \in \mathcal{R}^+$

Theorem: A class of contracts $\mathcal{T}$ is ambiguity-proof iff it is ordered.

Ambiguity Proofness

Proof of direction 1: oderedness implies ambiguity proofness

Suppose $\mathcal{T}$ is ordered, and let $\tau = (t^1, \ldots, t^k)$ be a consistent ambiguous contract incentivizing action i^*

We show: there exists a single contract incentivizing i^* at same payment

By orderedness, wlog, $t_j^1 \leq t_j$ for all outcomes j and all contracts $t \in \tau$

- Thus, for all actions i: $U_A(i, t^1) = U_A(i, \tau)$
- So: $i^*(t^1) = i^*(\tau)$
- By consistency: $U_P(i^*(t^1), t^1) = U_P(i^*(\tau), \tau)$
- Thus, the classic contract t^1 incentivizes action i^* at the same payment as τ
- So: $\mathcal{T}$ is ambiguity proof

Linear Contracts

Corollary: The class of linear contracts is ambiguity proof

A linear contract pays the agent a fixed share of the reward, namely:

$t_j = \alpha r_j$ for some $\alpha \in [0,1]$

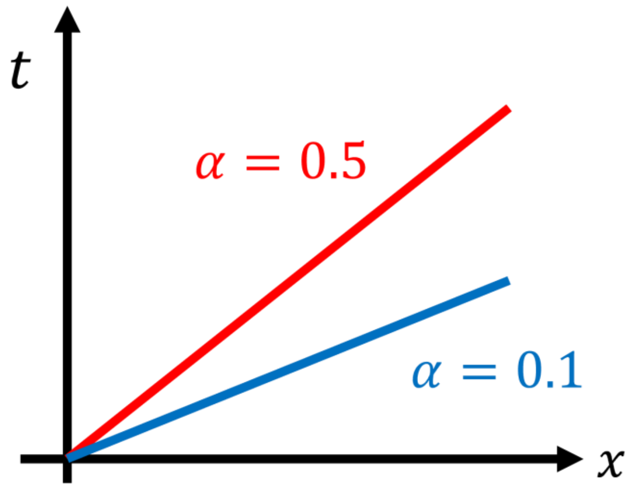

Linear Contracts

Corollary: The class of linear contracts is ambiguity proof

This provides another piece in a long-standing puzzle, asking why simple, sub-optimal contract formats, like linear, are so ubiquitous

> *"It is probably the great robustness of [linear contracts] that accounts for their popularity.*
> *That point is not made as effectively as we would like by our model; we suspect that it cannot be made effectively in any traditional […] model."*
>
> [Holmström & Milgrom'87]

Other pieces are provided by robust optimality of linear contracts [Carroll'15] [Duetting Talgam-Cohen Roughgarden'19]

Mixing Hedges Against Ambiguity

- A mixed action σ is a convex combination over pure actions
- σ_i is the probability the agent plays action i

- Expected reward of σ is $R_\sigma = \Sigma_i \sigma_i R_i$
- Expected payment of σ under contract t is $T_\sigma(t) = \Sigma_i \sigma_i T_i(t)$
- Agent's expected utility for σ under contract t is $U_A(\sigma, t) = \Sigma_i \sigma_i U_A(i, t)$
- Agent's expected utility for σ under ambiguous contract τ is $U_A(\sigma, \tau) = \min_{t \in \tau} U_A(\sigma, \tau)$

Mixing Hedges Against Ambiguity

	Cost	$r_1 = 2$	$r_2 = 2$
Action 1	1/4	1/2	0
Action 2	1/4	0	1/2
Action 3	1	1/2	1/2

Recall: under the ambiguous contract $\tau = ((2,0),(0,2))$, $u_A(1,\tau) = u_A(2,\tau) = -1/4$

- Consider mixed strategy σ, mixing between actions 1 and 2 with probability 0.5 each
- For any contract t: $U_A(\sigma, t) = 0.5\, U_A(1,t) + 0.5\, U_A(2,t)$
- Agent's utility under ambiguous contract τ is $U_A(\sigma,\tau) = \min_{t\in\tau} U_A(\sigma,t)$
- In our example: $U_A\big(\sigma,(2,0)\big) = 0.5\, U_A\big(1,(2,0)\big) + 0.5\, U_A\big(2,(2,0)\big) = 0.5 * \frac{3}{4} - 0.5 * \frac{1}{4} = \frac{1}{4} > 0$
  - Same for contract $(0,2)$. So, $U_A(\sigma,\tau) = 1/4$, strictly better than U_A for action 3
- Note: a mixed strategy may give a strictly higher utility than any of its pure strategies

Mixing Hedges Against Ambiguity

Theorem (informal): mixed strategies eliminate the power of ambiguity altogether

Theorem (formal): Suppose ambiguous contract τ incentivizes a mixed action σ with corresponding utilities $U_A(\sigma,\tau)$ and $U_P(\sigma,\tau)$. Then, there exists a classic contract t incentivizing σ with the same utilities

Proof idea: Consider a 0-sum game between the agent and principal:

T: All classic contracts preserving payoff $U_P(\sigma,\tau)$ under σ

all mixed actions

	$\boldsymbol{t^1}$	$\boldsymbol{t^2}$		
σ_1				
σ_2				
			$U_A(\sigma_i, t^j)$	

- $U_A(\sigma,\mathrm{T}) = U_A(\sigma,\tau)$ (by def of T)
- Claim: $U_A(\sigma,\tau)$ is the value of the game
- Let t be the classic contract realizing the minmax value
- By min-max thm: no mixed action gives the agent against t more than maxmin = $U_A(\sigma,\tau)$
- By construction, action σ gives this utility against t
- So t is the desired classic contract

Ambiguity gap

How much can the principal gain by ambiguous contracts?

Ambiguity gap

Ambiguity gap of an instance (c, r, p):

maximal principal's utility
using an **ambiguous contract**

$$\rho(c, r, p) = \frac{\max_{\tau} U_p(i^*(\tau), \tau)}{\max_{t} U_p(i^*(t), t)}$$

maximal principal's utility
using a **single contract**

Ambiguity gap

Ambiguity gap of an instance (c, r, p):

maximal principal's utility
using an **ambiguous contract**

$$\rho(c, r, p) = \frac{\max_{\tau} U_p(i^*(\tau), \tau)}{\max_{t} U_p(i^*(t), t)}$$

maximal principal's utility
using a **single contract**

Ambiguity gap of a class of instances $\mathcal{C}$: $\rho(\mathcal{C}) = \sup_{(c,r,p)\in\mathcal{C}} \rho(c, r, p)$

Max ambiguity gap over all
instances in class $\mathcal{C}$

Ambiguity gap

Ambiguity gap of an instance (c, r, p):

maximal principal's utility
using an **ambiguous contract**

maximal welfare of an action

$$\rho(c,r,p) = \frac{\max_{\tau} U_p(i^*(\tau),\tau)}{\max_{t} U_p(i^*(t),t)} \leq \frac{\max_{i\in[n]} W_i}{\max_{t} U_p(i^*(t),t)}$$

maximal principal's utility
using a **single contract**

Ambiguity gap of a class of instances $\mathcal{C}$: $\rho(\mathcal{C}) = \sup_{(c,r,p)\in\mathcal{C}} \rho(c,r,p)$

Max ambiguity gap over all instances in class $\mathcal{C}$

Main Result

Theorem: The ambiguity gap of the class of instances with n actions is $n - 1$

Note: upper bound follows from [Duetting et al. 19], who showed that this upper bound holds even with respect to optimal welfare, and even by a linear contract

Lower bound

- An instance with $n + 1$ actions and 3 outcomes having a gap of n
- Optimal welfare (from action $n + 1$) is roughly n
- Optimal principal's utility is roughly 1

Summary

- Algorithmic contract design is a new frontier in AGT
- Many interesting directions waiting to be explored

- Ambiguity can be used by the principal to gain higher utility
- Optimal ambiguous contracts have simple structure (SOP, step)
- Computing the optimal ambiguous contract is feasible
- Linear contracts are ambiguity-proof
- The ambiguity gap is roughly the number of actions
- Mixing hedges against ambiguity

Coming soon..

Survey on Algorithmic Contract Theory

[Duetting Feldman and Talgam-Cohen, to appear (FnTTCS)]

- Optimal and linear contracts
- Simple vs. optimal contracts
- Combinatorial contracts
- Contracts and types agents
- Date-driven contracts
- Contracts and incentive-aware machine learning
- Ambiguous contracts
- Contract design for social good
- Incentivizing effort beyond contracts

Thank you!