(Logic-based) Automated Mechanism Design

Munyque Mittelmann [1] Bastien Maubert [1]
Aniello Murano [1] **Laurent Perrussel** [2]

[1]University of Naples "Federico II"

[2]IRIT - University of Toulouse Capitole

laurent.perrussel@irit.fr

UNSW KR Conventicle - May 2024

Institut de Recherche
en Informatique de Toulouse
CNRS - INP - UT3 - UT1 - UT2J

Automated Mechanism Design

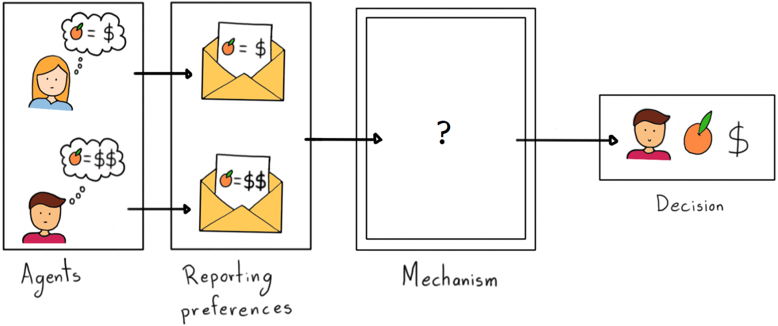

Automated Mechanism Design

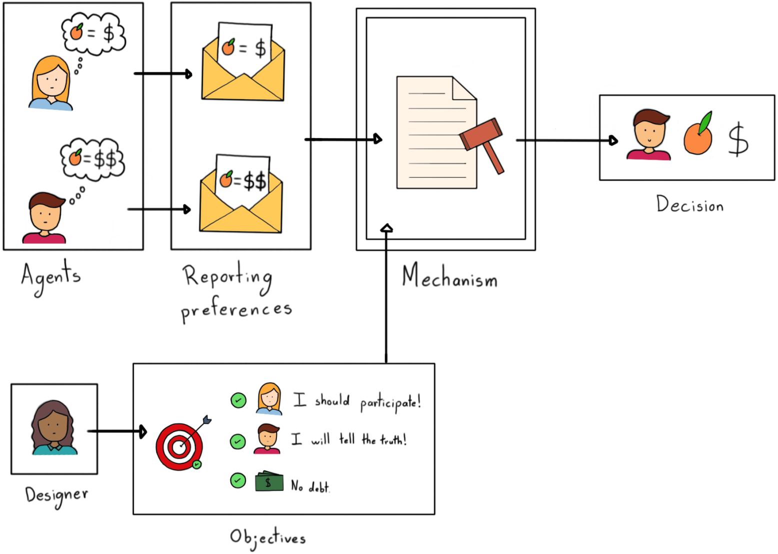

Automated Mechanism Design

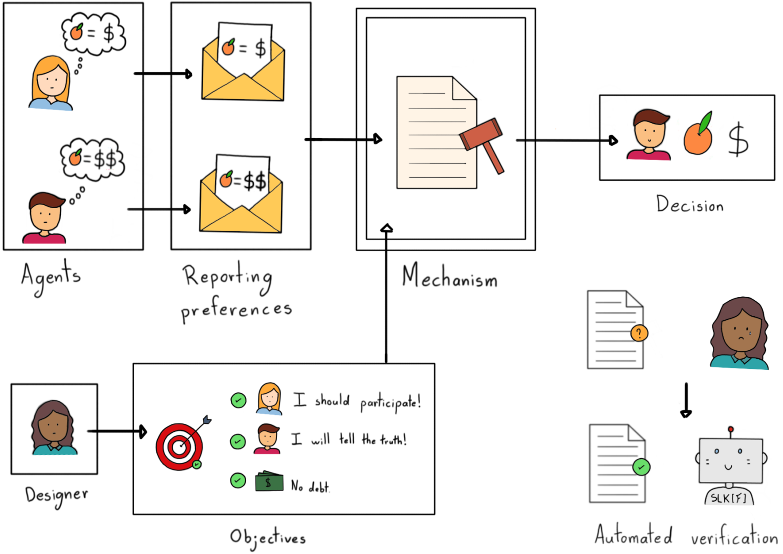

Table of Contents

Automated verification of mechanisms

Requirements:

i Quantitative aspects

ii Imperfect information (II)

iii Ability to express complex solution concepts

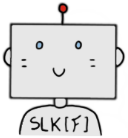

Quantitative and epistemic version of Strategy Logic (SLK[$\mathcal{F}$]) [1]

[1] Strategic Reasoning in Automated Mechanism Design (Maubert et al., KR 2021)

Related Works

- Logics for Strategic Reasoning:
  - ATL and extensions (Alur, Henzinger, and Kupferman 2002)
  - Strategy Logic (SL) (Chatterjee, Henzinger, and Piterman 2010)
  - SL with II and knowledge operators (Berthon et al. 2021; Belardinelli et al. 2020; Maubert and Murano 2018)
  - SL$[\mathcal{F}]$ (Bouyer et al. 2019)
  - Our work: SLK$[\mathcal{F}]$ SL with quantitative semantics and knowledge operators

$\mathrm{SL}[\mathcal{F}]$ with Imperfect Information

SLK$[\mathcal{F}]$ **Syntax**

- Syntax
  - Propositions p
  - Functions $f(\varphi, \ldots, \varphi)$
    e.g. $x \mapsto -x$
    $x, y \mapsto \max(x, y)$
  - Strategy quantifiers $\exists s_a.\, \varphi$ and bindings $(a, s_a)\varphi$
  - Epistemic operator: $K_a\varphi$
  - Temporal operators: $\mathbf{X}\varphi$ and $\varphi\mathbf{U}\varphi$ (and thus $\mathbf{F}\varphi$ and $\mathbf{G}\varphi$)

Concurrent Game Structure

- Weighted Concurrent Game Structure (wCGS) $\mathcal{G}$
  - state-transition model
  - state/position: proposition p with a weight
  - transition: joint action
  - observation relation: each agent can not distinguish between states
- Strategy Str_a of agent a: maps positions to actions
- Assignment χ: maps agents and variables to strategies

SLK$[\mathcal{F}]$ **Semantics**

Let $\mathcal{G}$ be a wCGS, and χ an assignment. Satisfaction value $[\![\varphi]\!]^{\mathcal{G}}_{\chi}(v) \in [-1, 1]$ of a formula φ in a position v is defined as follows

- $[\![p]\!]^{\mathcal{G}}_{\chi}(v) = \ell(v, p)$
- $[\![\exists s_a.\, \varphi]\!]^{\mathcal{G}}_{\chi}(v) = \max_{\sigma \in \mathit{Str}_a} [\![\varphi]\!]^{\mathcal{G}}_{\chi[s_a \mapsto \sigma]}(v)$
- $[\![(a, s_a)\varphi]\!]^{\mathcal{G}}_{\chi}(v) = [\![\varphi]\!]^{\mathcal{G}}_{\chi[a \mapsto \chi(s_a)]}(v)$
- $[\![K_a\varphi]\!]^{\mathcal{G}}_{\chi}(v) = \min_{v' \sim_a v} [\![\varphi]\!]^{\mathcal{G}}_{\chi}(v')$
- $[\![f(\varphi_1, \ldots, \varphi_m)]\!]^{\mathcal{G}}_{\chi}(v) = f([\![\varphi_1]\!]^{\mathcal{G}}_{\chi}(v), \ldots, [\![\varphi_m]\!]^{\mathcal{G}}_{\chi}(v))$
- $\mathbf{F}\varphi$ maximises the values of φ over all future points in time
- $\mathbf{G}\varphi$ minimizes the values of φ over all future points in time

Reasoning about Auction Mechanisms

Social choice functions and mechanisms

- Split the SCF into choice and payment functions: $f = (x, \{p_a\})$

Social choice functions and mechanisms

- Split the SCF into choice and payment functions: $\mathrm{f} = (\mathrm{x}, \{\mathrm{p}_a\})$
- Mechanisms as wCGS :
  - i One initial position for each possible type profile
  - ii Types do not change
  - iii Each agent knows her type
  - iv Every play reaches a terminal position

Social choice functions and mechanisms

- Split the SCF into choice and payment functions: $\mathsf{f} = (\mathsf{x}, \{\mathsf{p}_a\})$
- Mechanisms as wCGS :
  - i One initial position for each possible type profile
  - ii Types do not change
  - iii Each agent knows her type
  - iv Every play reaches a terminal position

Example (Dutch Auction)

- A position $\langle \mathsf{p}, \{x_a\}, \mathsf{t}, \{\theta_a\} \rangle$

Social choice functions and mechanisms

- Split the SCF into choice and payment functions: $\mathsf{f} = (\mathsf{x}, \{\mathsf{p}_a\})$
- Mechanisms as wCGS :
  - i One initial position for each possible type profile
  - ii Types do not change
  - iii Each agent knows her type
  - iv Every play reaches a terminal position

Example (Dutch Auction)

- A position $\langle \mathsf{p}, \{x_a\}, \mathsf{t}, \{\theta_a\}\rangle$
- An initial position $\langle 1, 0, \dots, 0, 0, -1, \theta_1, \dots, \theta_\mathsf{n}\rangle$

Social choice functions and mechanisms

- Split the SCF into choice and payment functions: $\mathsf{f} = (\mathsf{x}, \{\mathsf{p}_a\})$
- Mechanisms as wCGS :
  i One initial position for each possible type profile
  ii Types do not change
  iii Each agent knows her type
  iv Every play reaches a terminal position

Example (Dutch Auction)
- A position $\langle \mathsf{p}, \{x_a\}, \mathsf{t}, \{\theta_a\}\rangle$
- An initial position $\langle 1, 0, \ldots, 0, 0, -1, \theta_1, \ldots, \theta_{\mathsf{n}}\rangle$
- Transition: $\mathsf{p}' = \mathsf{p} - \mathsf{dec}$ if everyone $waits$
Otherwise, allocate the good to the agent who bet, she pays p

Social choice functions and mechanisms

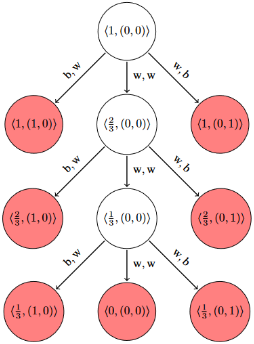

Figure 1: Mechanism timeline

One initial state (types are omitted) - Action bid is written b and wait is w.

Solution concepts

- Nash equilibrium (NE)

$$\mathsf{NE}(\boldsymbol{s}) := \bigwedge_{a \in \mathsf{Ag}} \forall t. \Big[(\mathsf{Ag}_{-a}, \boldsymbol{s}_{-a})(a, t)\mathbf{F}(ter \wedge \mathsf{util}_a) \leq (\mathsf{Ag}, \boldsymbol{s})\mathbf{F}(ter \wedge \mathsf{util}_a)\Big]$$

- Dominant strategy equilibrium (DSE)

$$\mathsf{DSE}(\boldsymbol{s}) := \bigwedge_{a \in \mathsf{Ag}} \mathsf{DS}(s_a, a)$$

where $\mathsf{DS}(s_a, a)$ if s_a weakly maximizes a's utility, for all strategies of other agents.

Implementation of SCF

- Alternatives Alt
- Agent's type $\theta_a \in \Theta_a$
- Social choice function (SCF) $\mathsf{f} : \boldsymbol{\Theta} \rightarrow \mathsf{Alt}$
- Atomic propositions for describing the alternatives
- Let $\mathsf{E} \in \{\mathsf{NE}, \mathsf{DSE}\}$
- Mechanism $\mathcal{G}$ E-implements the SCF f if they assign the same alternative in **some** E-equilibrium, for all type profiles $\boldsymbol{\theta}$.

Mechanism Properties

- Individual Rationality (IR): define IR $:= \bigwedge_{a \in \mathsf{Ag}} 0 \leq \mathsf{util}_a$
  Let $\mathcal{G}$ be a mechanism that E-implements f.

Proposition (IR)
f is individually rational iff IR has the satisfaction value 1 in the E-equilibrium implementing f (for all $\boldsymbol{\theta} \in \boldsymbol{\Theta}$).

Mechanism Properties

- Strategyproofness (SP)
  Let $\hat{\theta}_a$ be the truth-revealing strategy for a
  $\mathcal{G}$ is direct revelation mechanism

Proposition (SP)
$\mathcal{G}$ is SP if $[\![\text{DSE}(\boldsymbol{s})]\!]^{\mathcal{G}}_{\chi}(v^{\boldsymbol{\theta}}_{\iota}) = 1$ for all $\boldsymbol{\theta} \in \boldsymbol{\Theta}$, where $\chi(s_a) = \hat{\theta}_a$ for each a

Mechanism Properties

- Strategyproofness (SP)
  Let $\hat{\theta}_a$ be the truth-revealing strategy for a
  $\mathcal{G}$ is direct revelation mechanism

Proposition (SP)
$\mathcal{G}$ is SP if $[\![\mathrm{DSE}(\boldsymbol{s})]\!]^{\mathcal{G}}_{\chi}(v^{\boldsymbol{\theta}}_{\iota}) = 1$ for all $\boldsymbol{\theta} \in \boldsymbol{\Theta}$, where $\chi(s_a) = \hat{\theta}_a$ for each a

- Efficiency, Pareto optimality, budget-balance

Revenue benchmarks with knowledge

- Mechanisms with possibilistic beliefs $\mathcal{B}$ (Chen and Micali 2015)
- 2nd Belief Benchmark denoted $2^{\text{nd}}(\mathcal{B})$:
  - i The maximum type each agent a is *sure* someone has
  - ii The second highest of such values (for all agents)

Revenue benchmarks with knowledge

- Mechanisms with possibilistic beliefs $\mathcal{B}$ (Chen and Micali 2015)
- 2nd Belief Benchmark denoted $2^{\text{nd}}(\mathcal{B})$:
  i The maximum type each agent a is *sure* someone has
  ii The second highest of such values (for all agents)

$$\varphi_a^{\text{smv}} := K_a \max_{a' \in \text{Ag}} (\text{type}_{a'})$$

$$\varphi_{\text{2nd}} := \text{2nd-max}(\varphi_{a_1}^{\text{smv}}, \ldots, \varphi_{a_n}^{\text{smv}})$$

Revenue benchmarks with knowledge

- Mechanisms with possibilistic beliefs $\mathcal{B}$ (Chen and Micali 2015)
- 2nd Belief Benchmark denoted $2^{\text{nd}}(\mathcal{B})$:
  - i The maximum type each agent a is *sure* someone has
  - ii The second highest of such values (for all agents)

$$\varphi_a^{\text{smv}} := K_a \max_{a' \in \text{Ag}} (\text{type}_{a'})$$

$$\varphi_{\text{2nd}} := \text{2nd-max}(\varphi_{a_1}^{\text{smv}}, \ldots, \varphi_{a_n}^{\text{smv}})$$

Proposition (Revenue benchmark)

Given a mechanism $\mathcal{G}$, a position v and a belief profile $\mathcal{B}(v)$, it holds that $[\![\varphi_{\text{2nd}}]\!]^{\mathcal{G}}(v) = 2^{\text{nd}}(\mathcal{B}(v))$.

Model Checking and Synthesis

Model-checking

Model-checking problem (MC) for SLK$[\mathcal{F}]$:
Given a sentence φ, a wCGS $\mathcal{G}$, a position v in $\mathcal{G}$ and a predicate $P \subseteq [-1, 1]$, decide whether $[\![\varphi]\!]^{\mathcal{G}}(v) \in P$.

Theorem (MC for SLK$[\mathcal{F}]$)
Assuming that functions in $\mathcal{F}$ can be computed in polynomial space, model checking SLK$[\mathcal{F}]$ with imperfect information and **memoryless agents** is PSPACE-complete.

Synthesis of Mechanisms

- Creating mechanisms from a logical specification in SL$[\mathcal{F}]$[2]
- Satisfiability of SL (thus, SL$[\mathcal{F}]$) is undecidable in general
- Decidable cases

Theorem (Satisfiability of SL$[\mathcal{F}]$)
The satisfiability of SL$[\mathcal{F}]$ is decidable in the following cases

- wCGS with bounded actions
- Turn-based wCGS

[2]Automated Synthesis of Mechanisms (Mittelmann et al., IJCAI 2022)

Optimal mechanism synthesis

Algorithm 1: Optimal mechanism synthesis

Data: A SL$[\mathcal{F}]$ specification Φ and a set of possible values for atomic propositions $\mathcal{V}$

Result: A wCGS $\mathcal{G}$ such that $[\![\Phi]\!]^{\mathcal{G}}$ is maximal

Compute $\widetilde{\mathsf{Val}}_{\Phi,\mathcal{V}}$;

Let $\nu_1, \dots, \nu_n$ be a decreasing enumeration of $\widetilde{\mathsf{Val}}_{\Phi,\mathcal{V}}$;

for *i=1…n* **do**

  Solve $\mathcal{V}$- satisfiability for Φ and $\vartheta = \nu_i$;

  if *there exists* $\mathcal{G}$ *such that* $[\![\Phi]\!]^{\mathcal{G}} \geq \nu_i$ **then**

    return $\mathcal{G}$;

  end

end

Japanese auction

- $\mathbf{AG}((\text{initial} \to \text{price} = 0 \land \neg\text{terminal}) \land (\mathbf{XG}\neg\text{initial} \land \mathbf{F}\ \text{terminal}))$
- $\mathbf{AG}(\text{sold} \leftrightarrow \text{choice} \neq -1)$
- $\mathbf{AG}((\neg\text{sold} \land \text{price} + inc \leq 1) \to (\text{price} + inc = \mathbf{X}\text{price} \land \neg\mathbf{X}\text{terminal}))$
- $\mathbf{AG}((\text{sold} \lor \text{price} + inc > 1) \to (\text{price} = \mathbf{X}\text{price} \land \mathbf{X}\text{terminal}))$
- $\mathbf{AG}(\text{choice} = \text{wins}_a \leftrightarrow \text{bid}_a \land \bigwedge_{b \neq a} \neg\text{bid}_a)$
- $\mathbf{AG}(\text{choice} = -1 \leftrightarrow \neg(\bigvee_{a \in \text{Ag}}(\text{bid}_a \land \bigwedge_{b \neq a} \neg\text{bid}_a)))$
- $\mathbf{AG}(\bigwedge_{a \in \text{Ag}}(\text{choice} = \text{wins}_a \to \text{payment}_a = \text{price}))$
- $\mathbf{AG}(\bigwedge_{a \in \text{Ag}}(\text{choice} \neq \text{wins}_a \to \text{payment}_a = 0))$
- $\bigwedge_{\boldsymbol{\theta} \in \boldsymbol{\Theta}} \exists \boldsymbol{s}.\ \text{NE}(\boldsymbol{s}, \boldsymbol{\theta}) \land \mathbf{F}(\text{terminal} \land \text{EF}(\boldsymbol{\theta}))$

Proposition

There exists a wCGS such that the satisfaction value of these rules is 1.

Computational Complexity

Legacy of Strategy Logic

Synthesis of Mechanism

In general $k + 1$-EXPTIME.

Japanese Auction: 3-EXPTIME

Conclusion

- Logic-Based Mechanism Design
- Verifying properties $\rightarrow$ model check SLK$[\mathcal{F}]$-formulas (KR'21)
- Generating mechanisms $\rightarrow$ synthesis from SL$[\mathcal{F}]$-formulas (IJCAI'22)
- Probabilistic setting (AAAI'23)
  - Bayesian mechanisms
  - Mixed strategies
  - Randomized mechanisms

Going Further

- Previous logical approaches are deterministic
- Bayesian and randomized mechanisms
- Challenges for a general approach
  - Settings: deterministic or randomized mechanisms, incomplete information, mixed or pure strategies, and direct or indirect mechanisms
  - Time-line for revealing the incomplete information
- Framework for MD with Probabilistic Strategy Logic (PSL)
- Automatic verification through PSL model checking

Bayesian Mechanism Design

- A (randomized) *social choice function* (SCF) (similarly, *mechanism*) is a function that maps **type profiles** (resp, **strategy profiles**) to probability distributions over the set of alternatives.
- Mechanism as stochastic transition systems: labels on terminal states indicate the alternative chosen

Example BIN-TAC auction

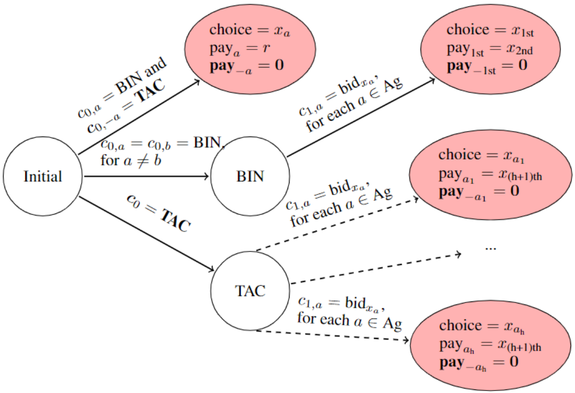

Figure 2: System representing the “Buy-It-Now or Take-a-Chance" (BIN-TAC) auction. Continuous lines are transitions with prob. 1 and dashed lines are transitions with prob. $\frac{1}{h}$.

Expected utilities

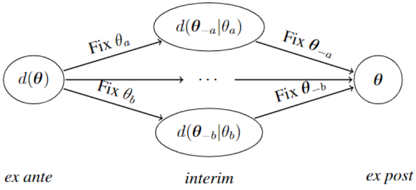

Figure 3: Mechanism timeline

Probability Strategy Logic captures the utility of agent a

- *ex ante* $\mathbb{E}_a^{e.a.}(\boldsymbol{s})$: expected utility given the type profile distribution
- *interim* $\mathbb{E}_a^{e.i.}(\boldsymbol{s}, \theta_a)$ expected utility given agent a's type and the distribution of type profiles
- *ex post* $\mathbb{E}_a^{e.p.}(\boldsymbol{s}, \boldsymbol{\theta})$ expected utility given a type profile

Expected utilities

In more details

$$\mathbb{E}_a^{e.p.}(\boldsymbol{s}, \boldsymbol{\theta}) := \sum_{\alpha \in \mathsf{Alt}} u_a(\theta_a, \alpha) \times \mathbb{P}_{\boldsymbol{s}(\boldsymbol{\theta})}(\mathbf{F}(\mathrm{ter} \wedge \mathrm{al}^{\alpha}))$$

$$\mathbb{E}_a^{e.i.}(\boldsymbol{s}, \theta_a) := \sum_{\boldsymbol{\theta}_{-a} \in \boldsymbol{\Theta}_{-a}} d(\boldsymbol{\theta}_{-a} | \theta_a) \times \mathbb{E}_a^{e.p.}(\boldsymbol{s}, (\boldsymbol{\theta}_{-a}, \theta_a))$$

$$\mathbb{E}_a^{e.a.}(\boldsymbol{s}) := \sum_{\boldsymbol{\theta} \in \boldsymbol{\Theta}} d(\boldsymbol{\theta}) \times \mathbb{E}_a^{e.p.}(\boldsymbol{s}, \boldsymbol{\theta})$$

Figure 4: PSL encoding

Solution concepts

Let $\boldsymbol{s} = (s_a)_{a \in \mathsf{Ag}}$ denote a strategy (variable) profile

$\boldsymbol{s}$ is a Nash equilibrium (NE) if for every agent a and for every $\boldsymbol{\theta}$, s_a is the best response (w.r.t. alternative strategy t_a) that a has to $\boldsymbol{s}_{-a}$ when the type profile is $\boldsymbol{\theta}$

$$\mathsf{NE}(\boldsymbol{s}) := \bigwedge_{\boldsymbol{\theta} \in \boldsymbol{\Theta}} \bigwedge_{a \in \mathsf{Ag}} \forall t_a.\, \mathbb{E}_a^{e.p.}((\boldsymbol{s}_{-a}, t_a), \boldsymbol{\theta}) \leq \mathbb{E}_a^{e.p.}(\boldsymbol{s}, \boldsymbol{\theta})$$

$\boldsymbol{s}$ is a Bayesian-Nash equilibrium (BNE) if for every agent a and every θ_a, s_a is the best response that a has to $\boldsymbol{s}_{-a}$ when her type is θ_a, in expectation over the other types $\boldsymbol{\theta}_{-a}$

$$\mathsf{BNE}(\boldsymbol{s}) := \bigwedge_{a \in \mathsf{Ag}} \bigwedge_{\theta_a \in \Theta_a} \forall t_a.\, \mathbb{E}_a^{e.i.}((\boldsymbol{s}_{-a}, t_a), \theta_a) \leq \mathbb{E}_a^{e.i.}(\boldsymbol{s}, \theta_a)$$

Implementation of an SCF

Given an equilibrium concept $\mathbf{E}$, a mechanism $\mathbf{E}$-implements an SCF f if there exists a strategy profile $\boldsymbol{\sigma}(\boldsymbol{\theta})$ that is an $\mathbf{E}$-equilibrium and it assigns the same probability distribution as f under strategies $\boldsymbol{\sigma}(\boldsymbol{\theta})$, for any types $\boldsymbol{\theta}$[3].

Let $\mathcal{G}$ be a system representing a mechanism and $\varphi_{\mathsf{f},\boldsymbol{s}}$ be the PSL formula expressing whether f assigns the same probability distribution as $\mathcal{G}$ under $\boldsymbol{s}$.

$\mathcal{G}$ $\mathbf{E}$-implements an f if

$$\mathcal{G}, v_\iota \models \exists \boldsymbol{s}.\, \mathbf{E}(\boldsymbol{s}) \wedge \varphi_{\mathsf{f},\boldsymbol{s}}$$

[3] Generalised from Parkes 2001

Mechanism properties

An SCF f is (interim) IR if for every $\boldsymbol{\theta} \in \boldsymbol{\Theta}$ and agent a, their interim utility is at least 0

Given a mechanism $\mathcal{G}$ $\mathbf{E}$-implementing f, $\mathcal{G}$ is interim IR if

$$\mathcal{G}, v_\iota \models \exists \boldsymbol{s}.\, \mathbf{E}(\boldsymbol{s}) \wedge \mathbf{F}(\text{terminal} \wedge \varphi_{\mathsf{f},\boldsymbol{s}} \wedge \bigwedge_{\boldsymbol{\theta} \in \boldsymbol{\Theta}} \text{IR}(\boldsymbol{s}, \boldsymbol{\theta}))$$

where $\text{IR}(\boldsymbol{s}, \boldsymbol{\theta}) := \bigwedge_{a \in \mathsf{Ag}} 0 \leq \mathbb{E}_a^{e.i.}(\boldsymbol{s}, \theta_a)$

Mechanism properties

A direct mechanism is *BIC* if the truth-revealing strategy profile $(\hat{\theta}_a)_{a \in \mathsf{Ag}}$ is a BNE for any $\boldsymbol{\theta} \in \boldsymbol{\Theta}$

Let $\mathcal{G}$ be a system representing a mechanism, $\mathcal{G}$ is BIC if

$$\mathcal{G}, \chi[\boldsymbol{s} \to (\hat{\theta}_a)_{a \in \mathsf{Ag}}], v_\iota \models \mathsf{BNE}(\boldsymbol{s})$$

Evaluating mechanisms $\rightarrow$ model-checking PSL -formulas, which is decidable for memoryless strategies

Conclusion

- Bridge between the economics' approach to MD and formal reasoning in Multi-Agent Systems
- General approach for verification of mechanisms using $SL[\mathcal{F}]$ and Bayesian mechanisms using PSL
- Future work
  - Social Good
  - Practical Tools!

(Logic-based) Automated Mechanism Design

Munyque Mittelmann [1]   Bastien Maubert [1]
Aniello Murano [1]   **Laurent Perrussel** [2]

[1]University of Naples “Federico II”

[2]IRIT - University of Toulouse Capitole

laurent.perrussel@irit.fr

UNSW KR Conventicle - May 2024

Institut de Recherche
en Informatique de Toulouse
CNRS - INP - UT3 - UT1 - UT2J

References I

Maubert, B. et al. (2021). “Strategic Reasoning in Automated Mechanism Design”. In: *KR*.

Alur, R., T.A. Henzinger, and O. Kupferman (2002). “Alternating-time temporal logic”. In: *J. ACM* 49.5, pp. 672–713. URL: https://doi.org/10.1145/585265.585270.

Chatterjee, K., T. A. Henzinger, and N. Piterman (2010). “Strategy Logic”. In: *Inf. Comput.* 208.6, pp. 677–693. DOI: 10.1016/j.ic.2009.07.004. URL: http://dx.doi.org/10.1016/j.ic.2009.07.004.

Berthon, R. et al. (2021). “Strategy Logic with Imperfect Information”. In: *ACM Trans. Comput. Logic* 22.1.

Belardinelli, F. et al. (2020). “Verification of multi-agent systems with public actions against strategy logic”. In: *Artif. Intell.* 285.

Maubert, B. and A. Murano (2018). “Reasoning about Knowledge and Strategies under Hierarchical Information”. In: *KR*.

Bouyer, P.et al. (2019). “Reasoning about Quality and Fuzziness of Strategic Behaviours”. In: *IJCAI*. DOI: 10.24963/ijcai.2019/220.

Chen, Jing and Silvio Micali (2015). “Mechanism design with possibilistic beliefs”. In: *J. Econ. Theory* 156, pp. 77–102.

Mittelmann, M. et al. (2022). “Automated Synthesis of Mechanisms”. In: *To appear at IJCAI*.

References II

Parkes, D. (2001). *Iterative combinatorial auctions: Achieving economic and computational efficiency*. Univ. of Pennsylvania Philadel.