Mechanism Design for Scheduling with Uncertain Execution Time.

Angelina Vidali
Teesside University

Joint work with
Vincent Conitzer (Duke University)

and
ongoing work with Carmine Ventre

The Queen wants a painting for her palace.
Every day she decides which painters will draw.

Goal: minimize E[painting time]

Creativity is unpredictable!
Painters don't know how much time it's going to take them and need incentives to draw!

Time painter i needs to finish the job ~ distribution f_i

painter i knows f_i but not t_i and f_{-i}

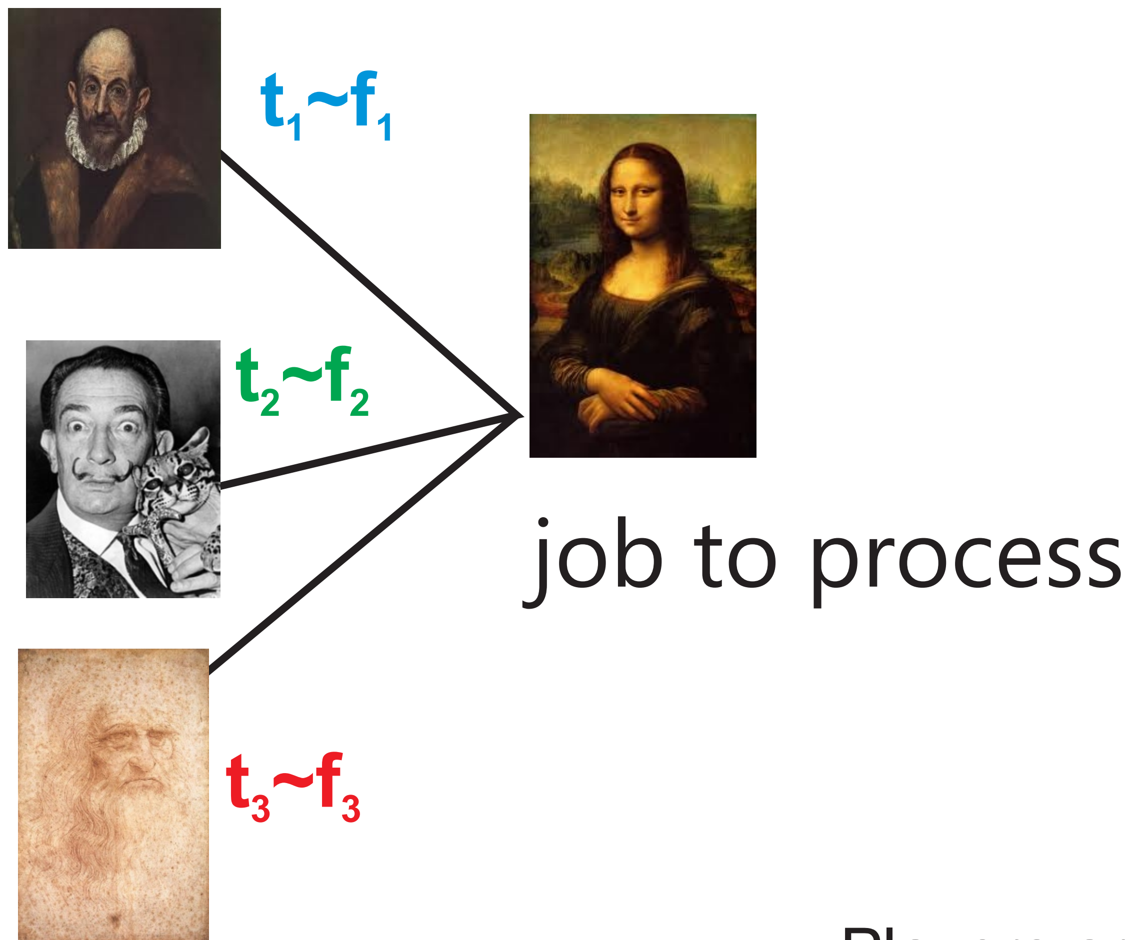

Players are selfish
want to maximize:
utility=E[-time spent painting + payment]

Crowdsourcing

- **Acomplish a complicated very big job using many agents!**
- **Give incentives to the agents to complete the tasks!**

Crowdsourcing Application

We want to solve a problem (e.g. a SAT instance), by running a crowdsourcing contest.

Payment 0 if you fail: Only the winner (agent who solves the problem first) will get paid.

Uncertainty: The agents don’t know how much time they will need.

For SAT there are some good euristics so the probability of finding a solution in the beginning is rather high, but if these don’t work then it might take forever... (MHR assumption makes sense)

Hazard Rate:

Probability a painter finishes the painting at time t
given that he hasn't finished it until time t-1

$$\varphi_i(t) = \frac{P(T_i = t)}{(1 - P[T_i < t])}$$

Monotone (non-increasing) hazard rate assumption:

the more time a painter takes
the less likely is he to finish at the next time step

The greedy algorithm is optimal!

Objective: Minimize the E[sum of processing times]

Greedy=OPT:
assign at each time step the job to the machine with maximum hazard rate
i.e. the machine more likely to finish!

To prove this we need:
Monotone hazard rates assumption

“Sort the hazard rates!”

Objective: Minimize the E[sum of processing times]

OPT: assign at each time step the job to the machine with maximum hazard rate

Input

	φ(1)	φ(2)	φ(3)	φ(4)	φ(5) ...
	0.9	0.8	0.5	0.1	0.1
	0.7	0.4	0.3	0.3	0.3
	0.6	0.2	0.1	0.1	0.1

OPT

	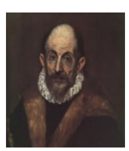						
0.9	0.8	0.7	0.6	0.5	0.4	0.3	0.3

Consistency Property

If we remove one player e.g.  *to get OPT for the rest of the players we just need to remove the player from the schedule*

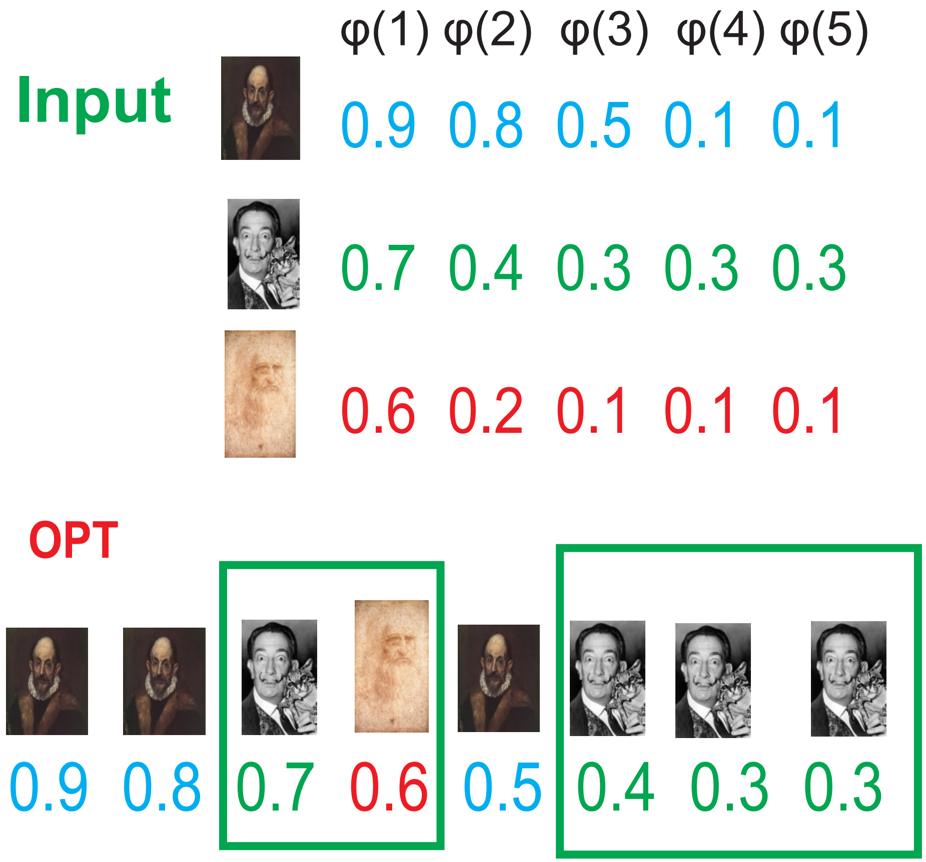

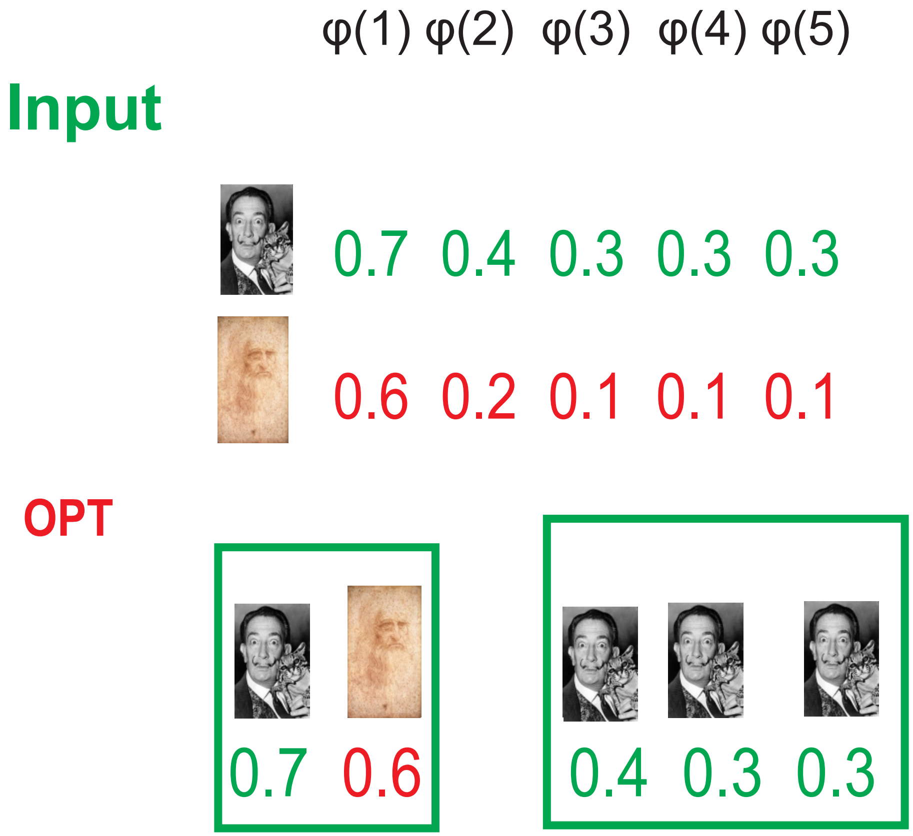

We focus on direct revelation mechanisms

Input: true types of the players
here: the distributions f_i

Output: allocation(which machine processes
at each time step) and
payments

We have a revelation principle

The Vickrey mechanism

min sum of processing times

truthful

The highest bid v_i wins and

pays 2nd highest bid

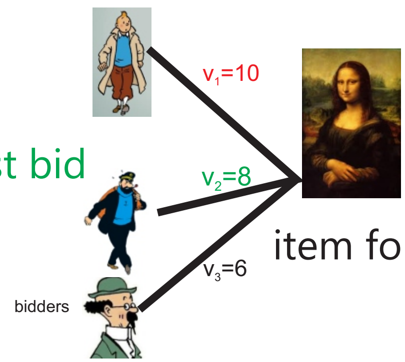

item for sale

Auctions

The fastest machine (lowest bid) t_i wins and

gets payed the 2nd lowest bid

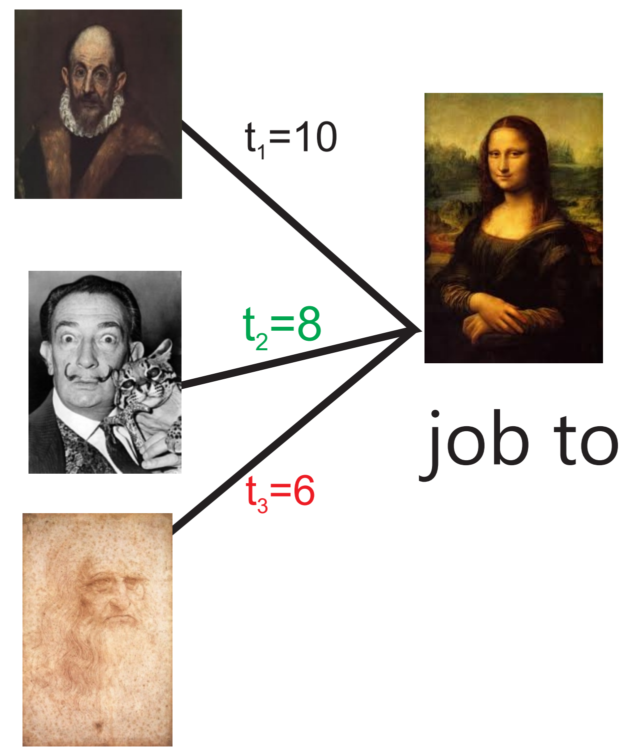

job to process

Scheduling

In our setting
the Expected Vickrey mechanism
is not truthful.

Groves Realized

OPT

REALIZED TIMES OF THE OTHER PLAYERS

FINISH TIME

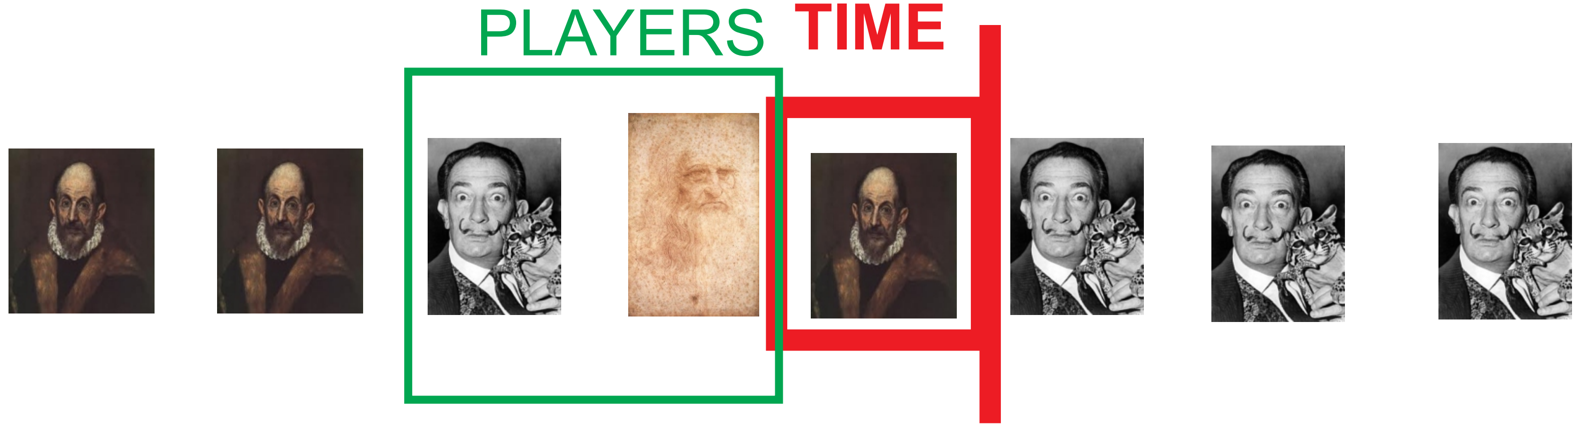

After completing the task we have the realized running times of the players

Groves payment

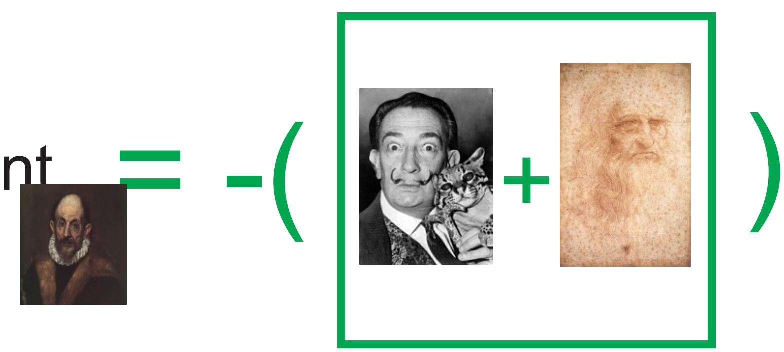

–"*sum of the realized times of the other players*"

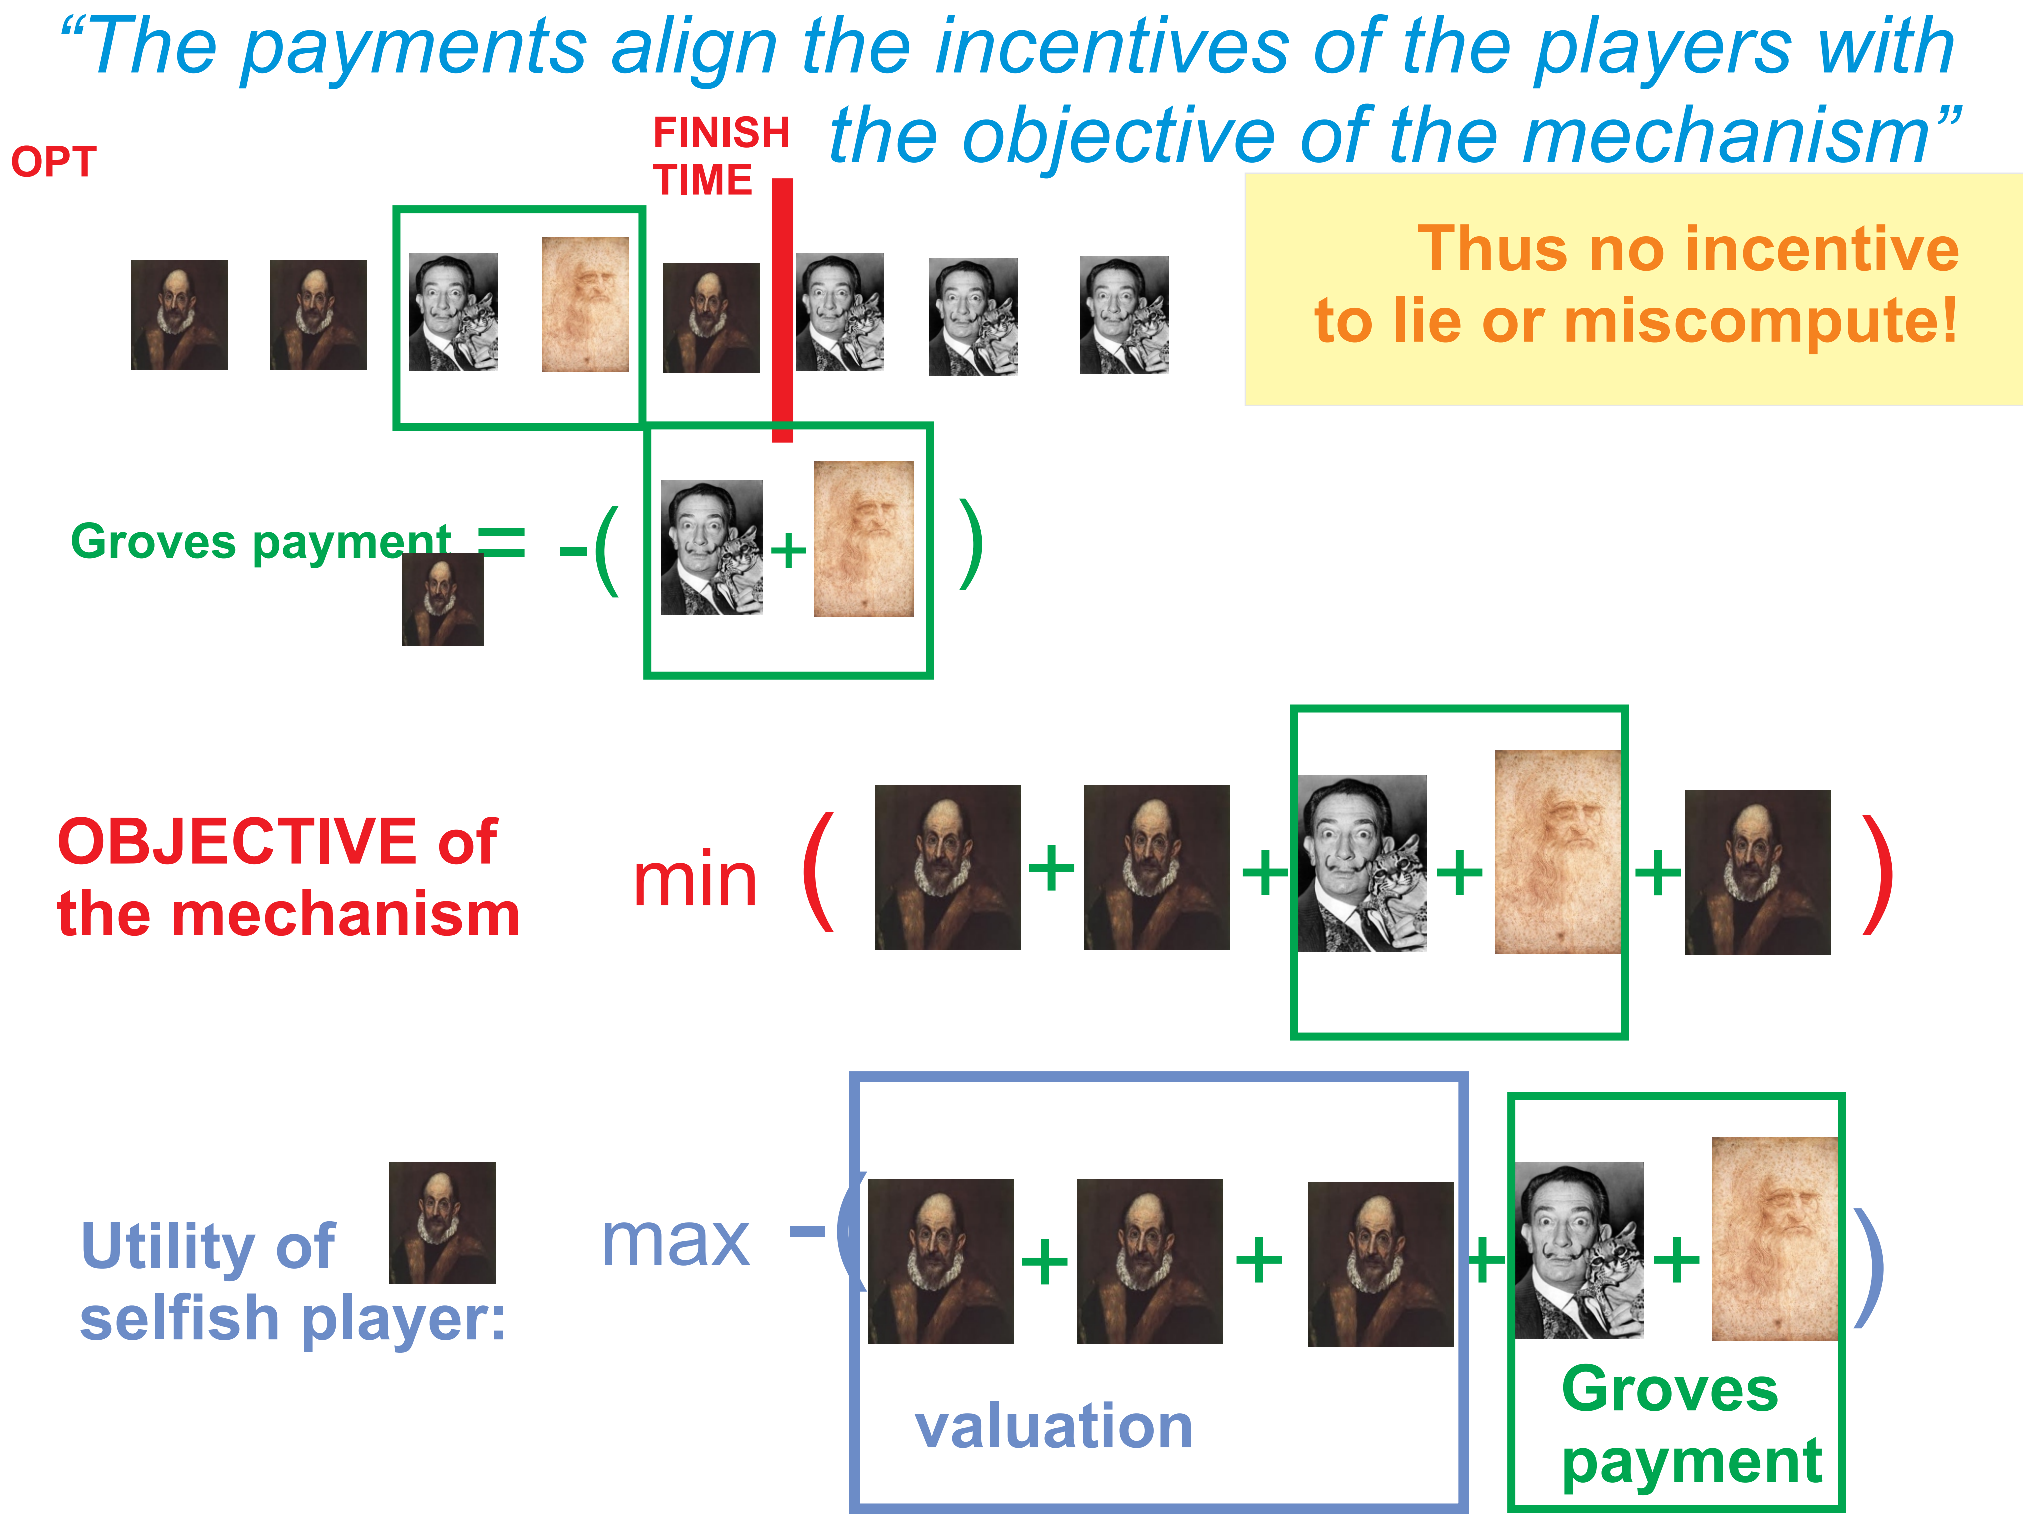

“The payments align the incentives of the players with the objective of the mechanism”
OPT
FINISH TIME
Thus no incentive to lie or miscompute!
Groves payment = -(+)
OBJECTIVE of the mechanism
min (+ + + +)
Utility of selfish player:
max -(+ + + +)
valuation
Groves payment

Solution Concept: Ex-post equilibrium

If the other players are telling the truth,
then the best thing for me to do is to tell the truth,
for **any** private information the players might.

Dominant $\subseteq$ Ex-post $\subseteq$ Bayes Nash

Valuations are interdependent

The valuation of a player depends on
whether another player
has already finished before him.

If we want our mecahnism to have an
h_i(types of the other players) part
(useful for getting properties like IR,etc)

p_i=Groves + h_i(types of the other players)

We have to consider the situation
when the player who finishes isn't there.

"How much does the player who finish contribute to the social wellfare?"

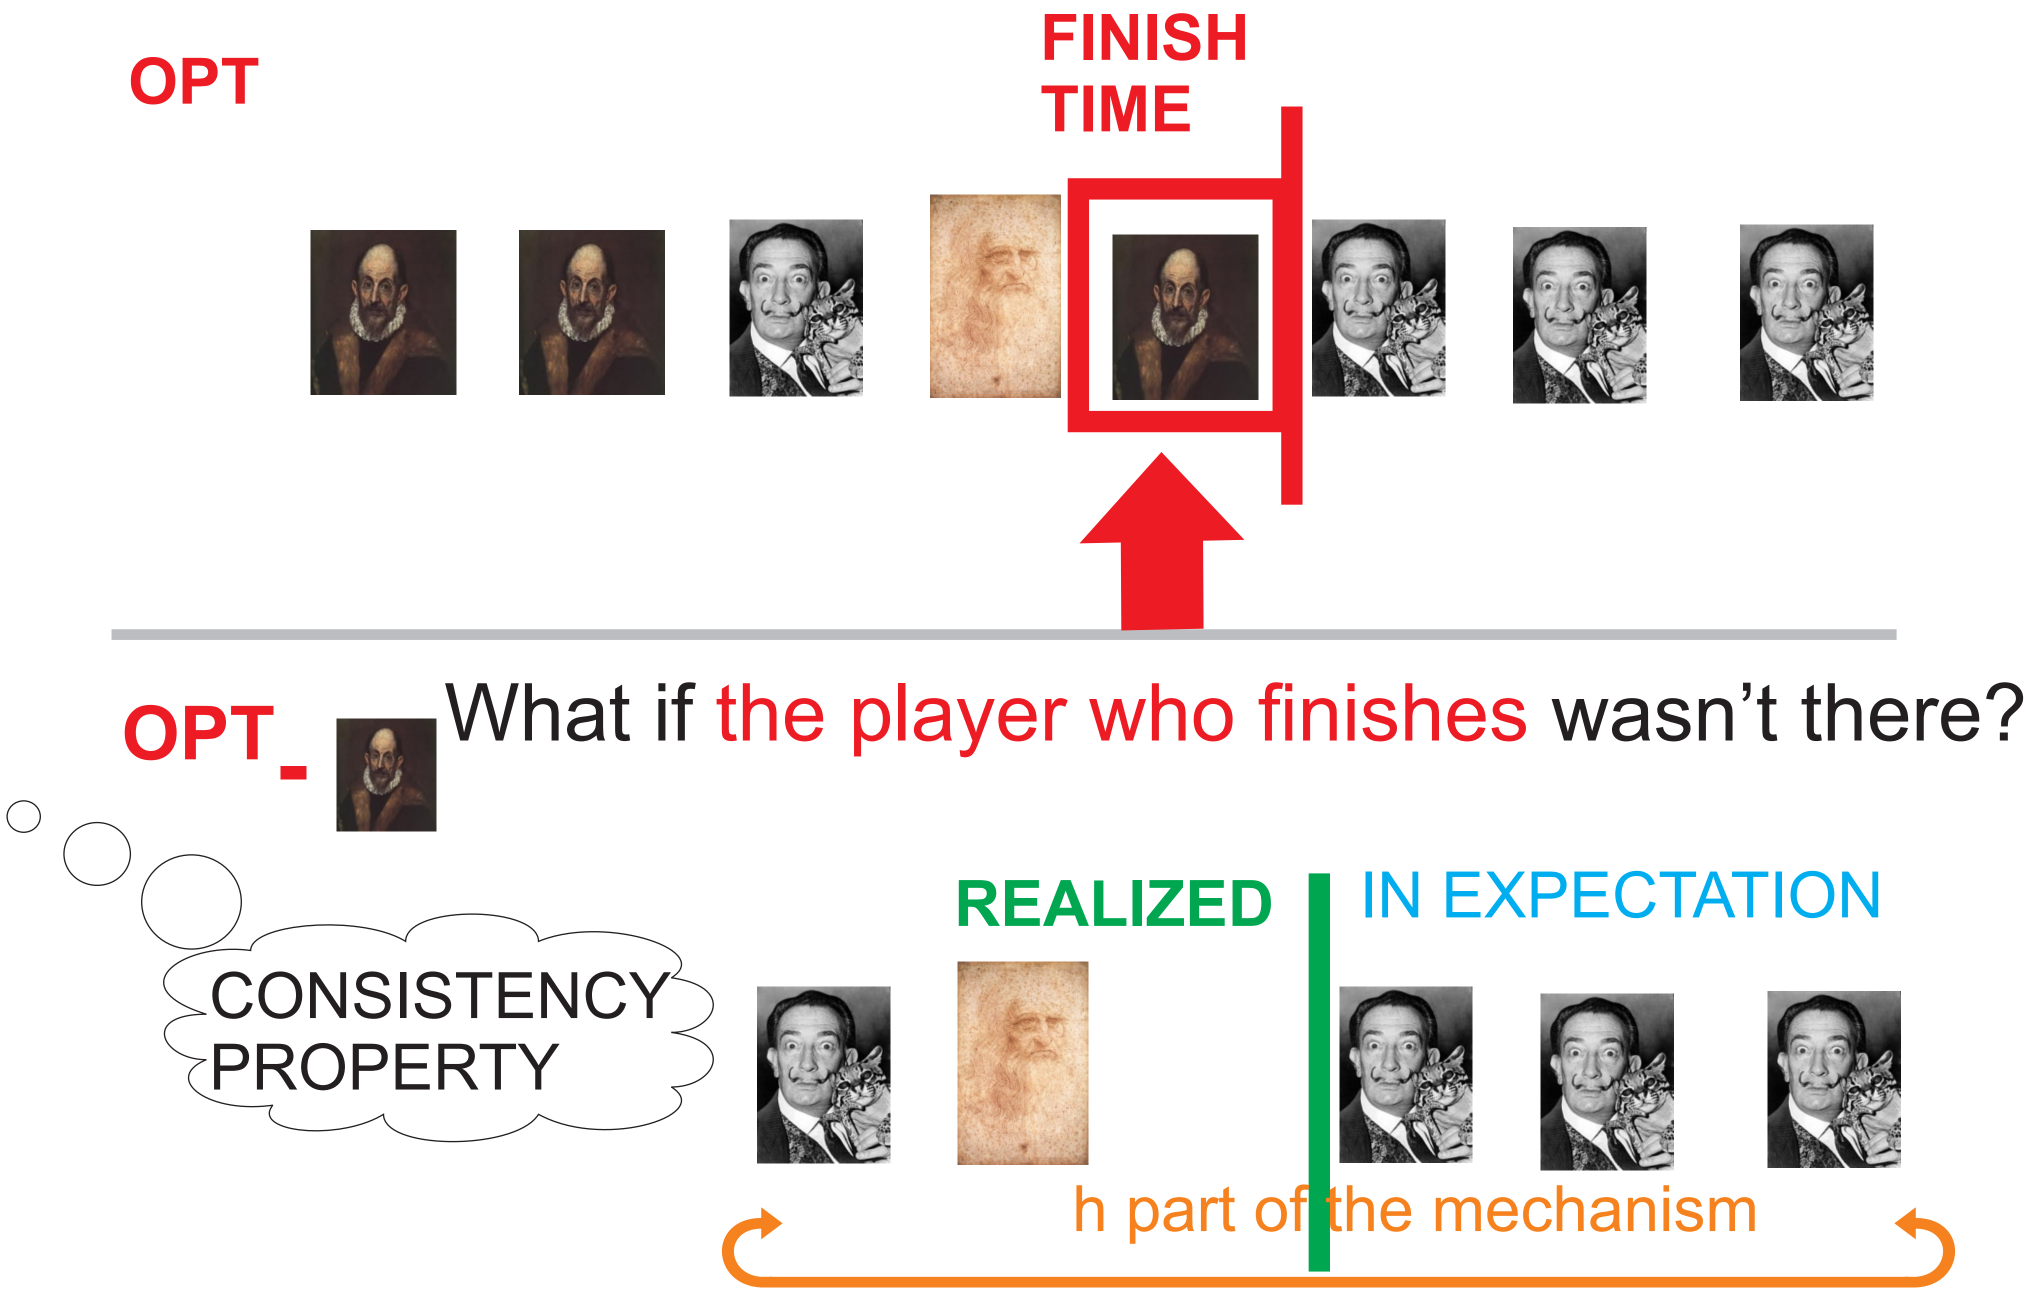

Vickrey Variations

T_N :=how long it takes a group N to finish the task (random variable)

r_N :=realized value of T_N $\quad$ h_i(types of the other players) part

	Groves part			
$p_i=$	$E[T_N - T_i]$	+	0	Expectated Pure Groves (EPG)
$p_i=$	$E[T_N - T_i]$	+	$E[T_{N\setminus\{i\}}]$	Expectated Clarke (EC)
$p_i=$	$-(r_N - r_i)$	+	0	Realized Pure Groves (RPG)
$p_i=$	$-(r_N - r_i)$	+	$E[T_{N\setminus\{i\}}]$	Clarke h (ChE) in Expectation
$p_i=$	$-(r_N - r_i)$	+	$(r_N - r_i) + E[T_{N\setminus\{i\}} - r_{N\setminus\{i\}} \mid T_{N\setminus\{i\}} \geq r_{N\setminus\{i\}}]$	(ChpE) Clarke h partially in Expectation

Groves part

$h_i()$

This rewritting uses the consistency property!

Properties of different Mechanisms

	efficient	truthful in dominant strategies	ex-post truthful	IR	no incentive to miscompute	payment 0 if fail
Clarke in Expectation (CE)	✓	✗	✗	✓	✗	✗
Pure Realized Groves (PRG)	✓	✗	✓	✗	✓	✗
Clarke h (ChE) in Expectation	✓	✗	✓	✓	✓	✗
Clarke h partially in Expectation (ChpE)	✓	✗	✓	✓	✓	✓

Main Theorem:

There exists an (ex-post) truthful mechanism that:

(a) doesn't pay players who did not complete the task

(b) satisfies IR in expectation

(c) has positive payments

(d) it is to the best interest of the players to exert full effort.

(Also generalizes to many tasks.)

No incentive to miscompute

Maybe the painters reported their true distributions but in the end decided it is to their best interest to take a break instead of painting!

If no realized values are used then players can just sit and compute nothing!

“The payments align the incentives of the players with the objective of the mechanism”

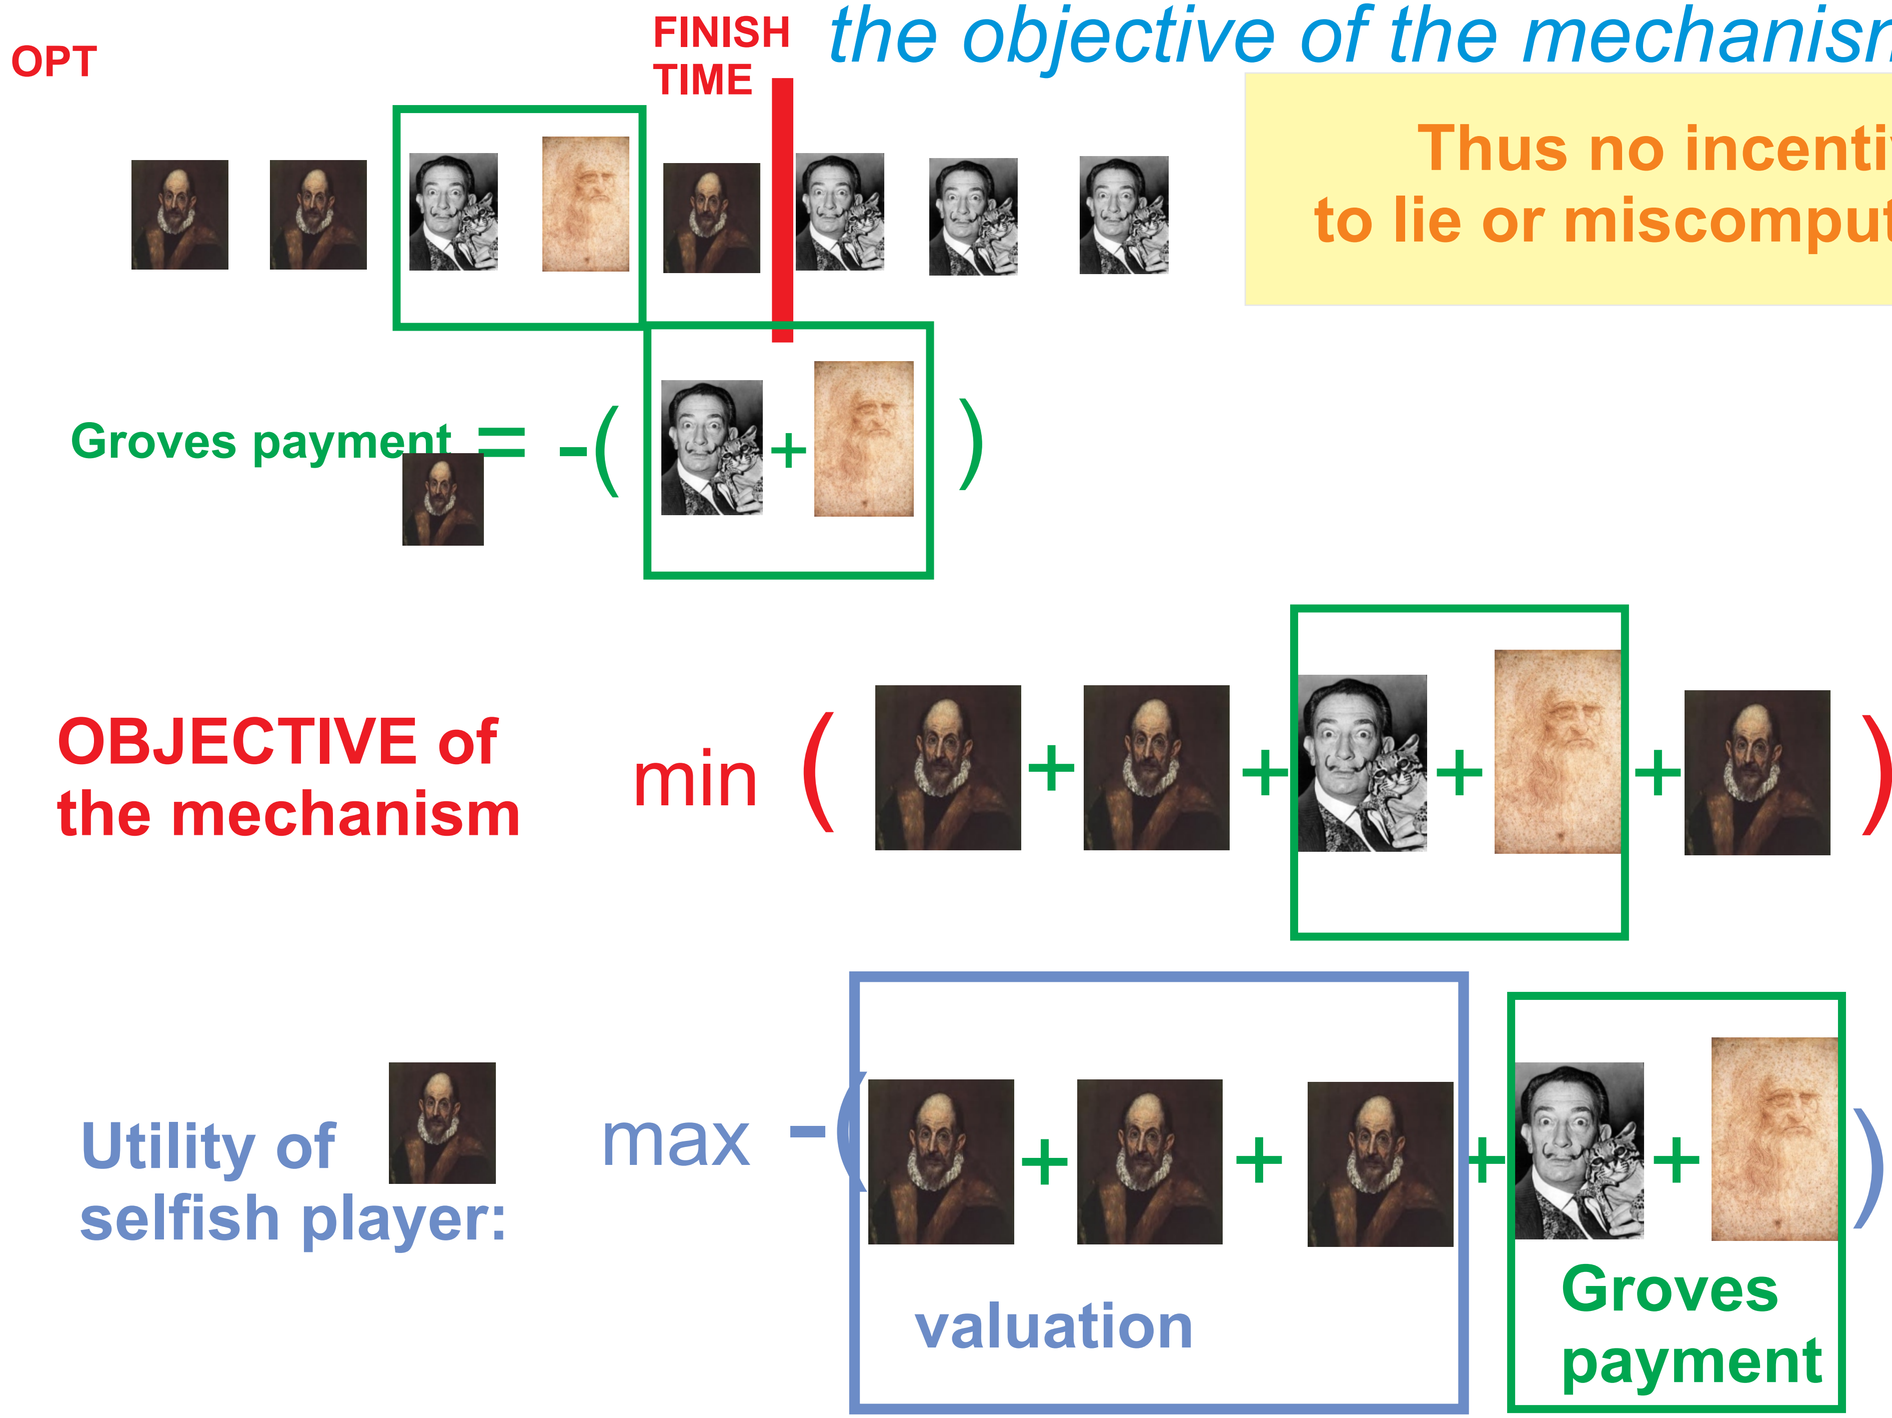

Thus no incentive to lie or miscompute!

No incentive to lie or miscompute in ChpE

(Clarke h partially in Expectation)

Proof is more involved. Idea:

$$payment_i = -(r_N - r_i) + \underbrace{(r_N - r_i) + E[T_{N\setminus\{i\}} - r_{N\setminus\{i\}} \mid T_{N\setminus\{i\}} \geq r_{N\setminus\{i\}}]}_{h_i()}$$

When player i finishes the task he determines which part is taken in expectation and which not. Suppose that he had even more power and could “cut” h at any point then he could still not affect the Expectation of h by miscomputing!

“Would you tell me, please, which way I ought to go from here?”

"That depends a good deal on where you want to get to."

"I don't much care where –"

“Then it doesn't matter which way you go.” **Lewis Carol, Alice in Wonderland**