

# When is “Altruism” good in distributed decision-making?

Philip N. Brown

Dept of Computer Science

University of Colorado, Colorado Springs

IEEE Computer Society – Pikes Peak Chapter  
4/13/2023




University of Colorado  
Colorado Springs



University of Colorado


Boulder | Colorado Springs | Denver | Anschutz Medical Campus



?



Engineered System



1001101  
0010011

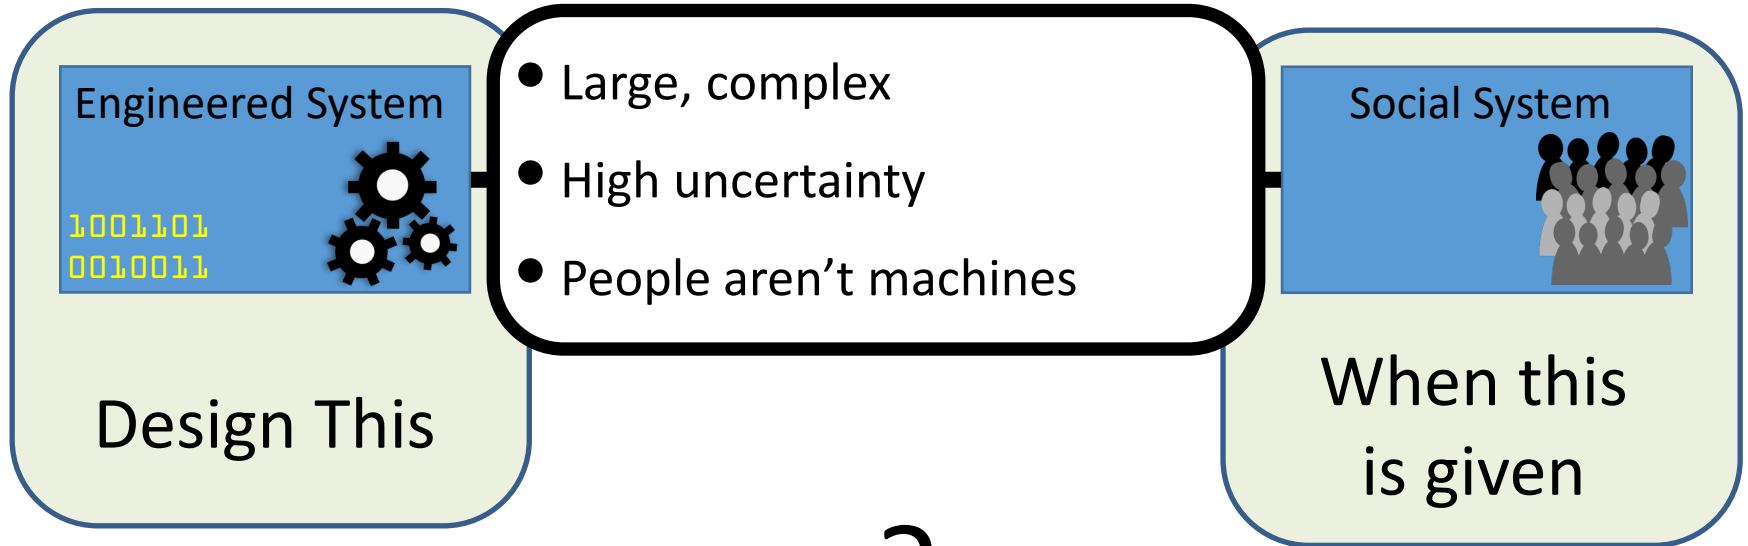
- Large, complex
- High uncertainty
- Optimizers are optimizers

Social System



?





?



# What is Algorithmic Game Theory?

?



# What is Algorithmic Game Theory?

**Game Theory:** The mathematics of strategic interaction

When strategic decision-makers interact, what happens?

1930s-1950s: Mathematics (John von Neumann, John Nash)

?



Nash

Nash?

# What is Algorithmic Game Theory?

**Game Theory:** The mathematics of strategic interaction

When strategic decision-makers interact, what happens?

1930s-1950s: Mathematics (John von Neumann, John Nash)

1960s-1990s: Economics

2000s-present: Engineering and Computer Science

↑  
Descriptive

Prescriptive  
↓

?



Christos Papadimitriou: “Algorithms, Games, and The Internet” (STOC 2001)

# Decision Science and Control Lab Projects

## Resilient Multiagent Coordination




Adaptation policies  
+  
uncertain communication

## Smart Infrastructure



Design incentives  
constraints  
+  
uncertainty

## Disease ↔ Behavior



Socially-aware  
Pandemic  
mitigation

## Strategy-Aware Security



Adversary strategy  
+  
security?



Air Force Office  
of Scientific Research  
FA9550-23-1-0171



National Science  
Foundation  
ECCS-2013779



National Science  
Foundation  
DEB-2032465



National Security  
Agency  
H98230-21-1-0155



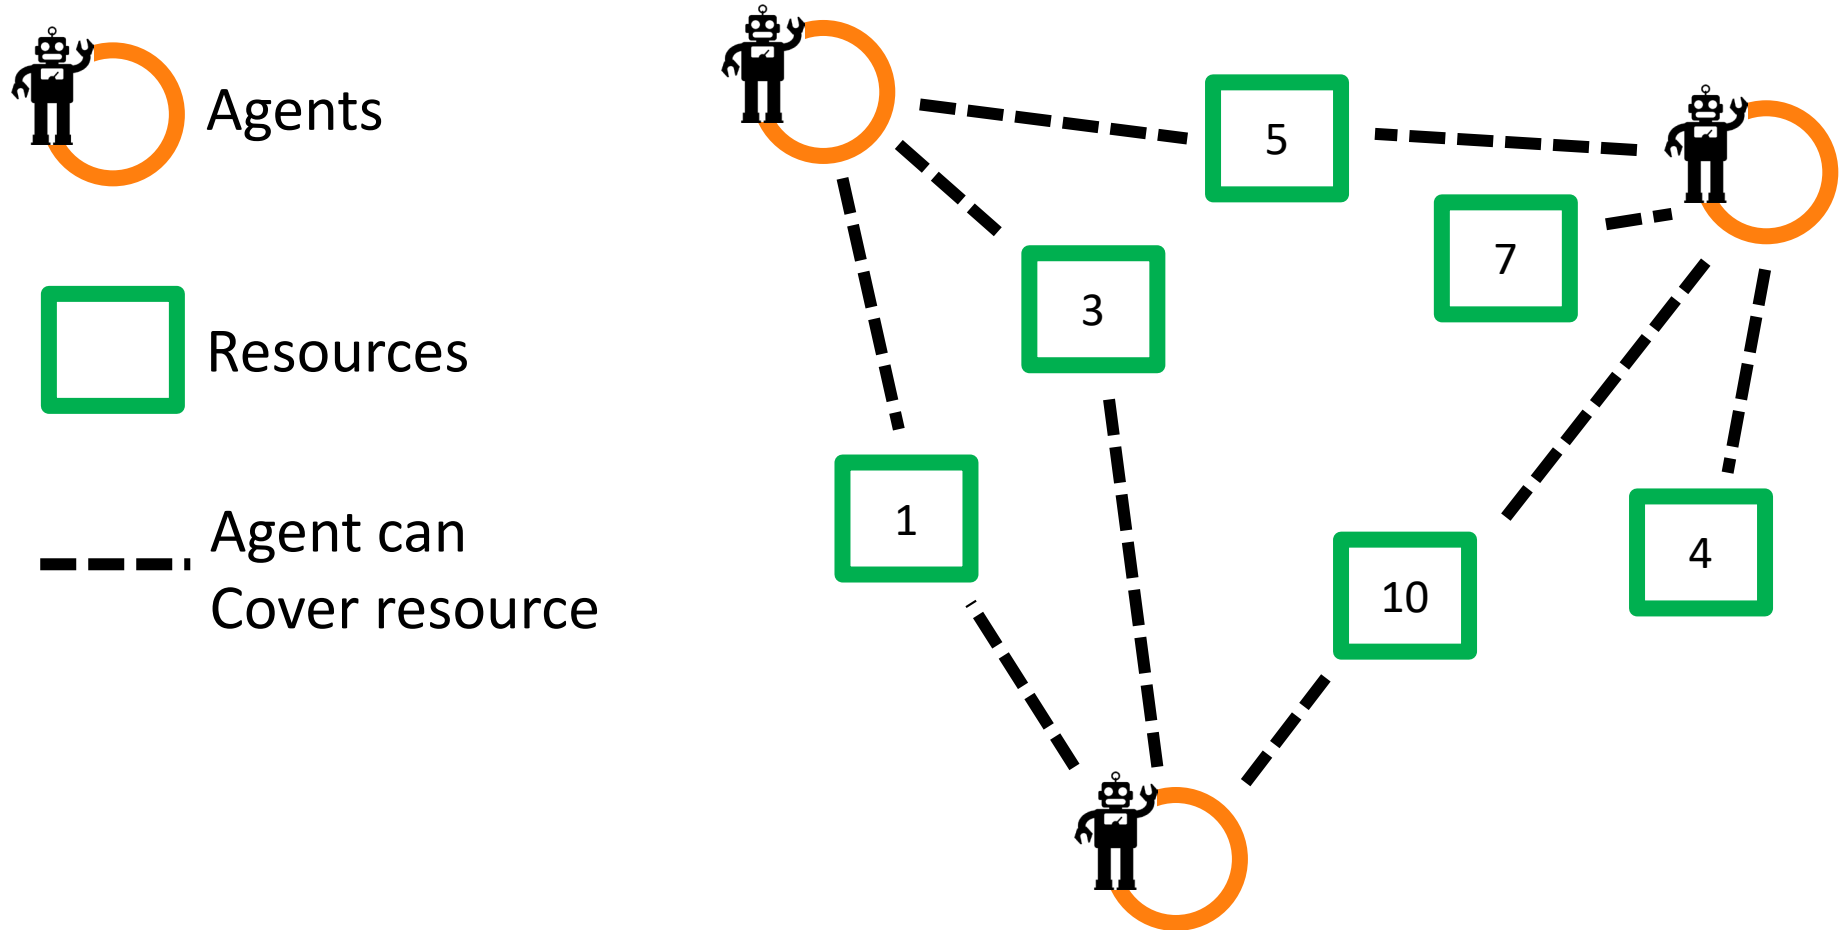


How should interacting autonomous agents behave  
in **compromised** environments?

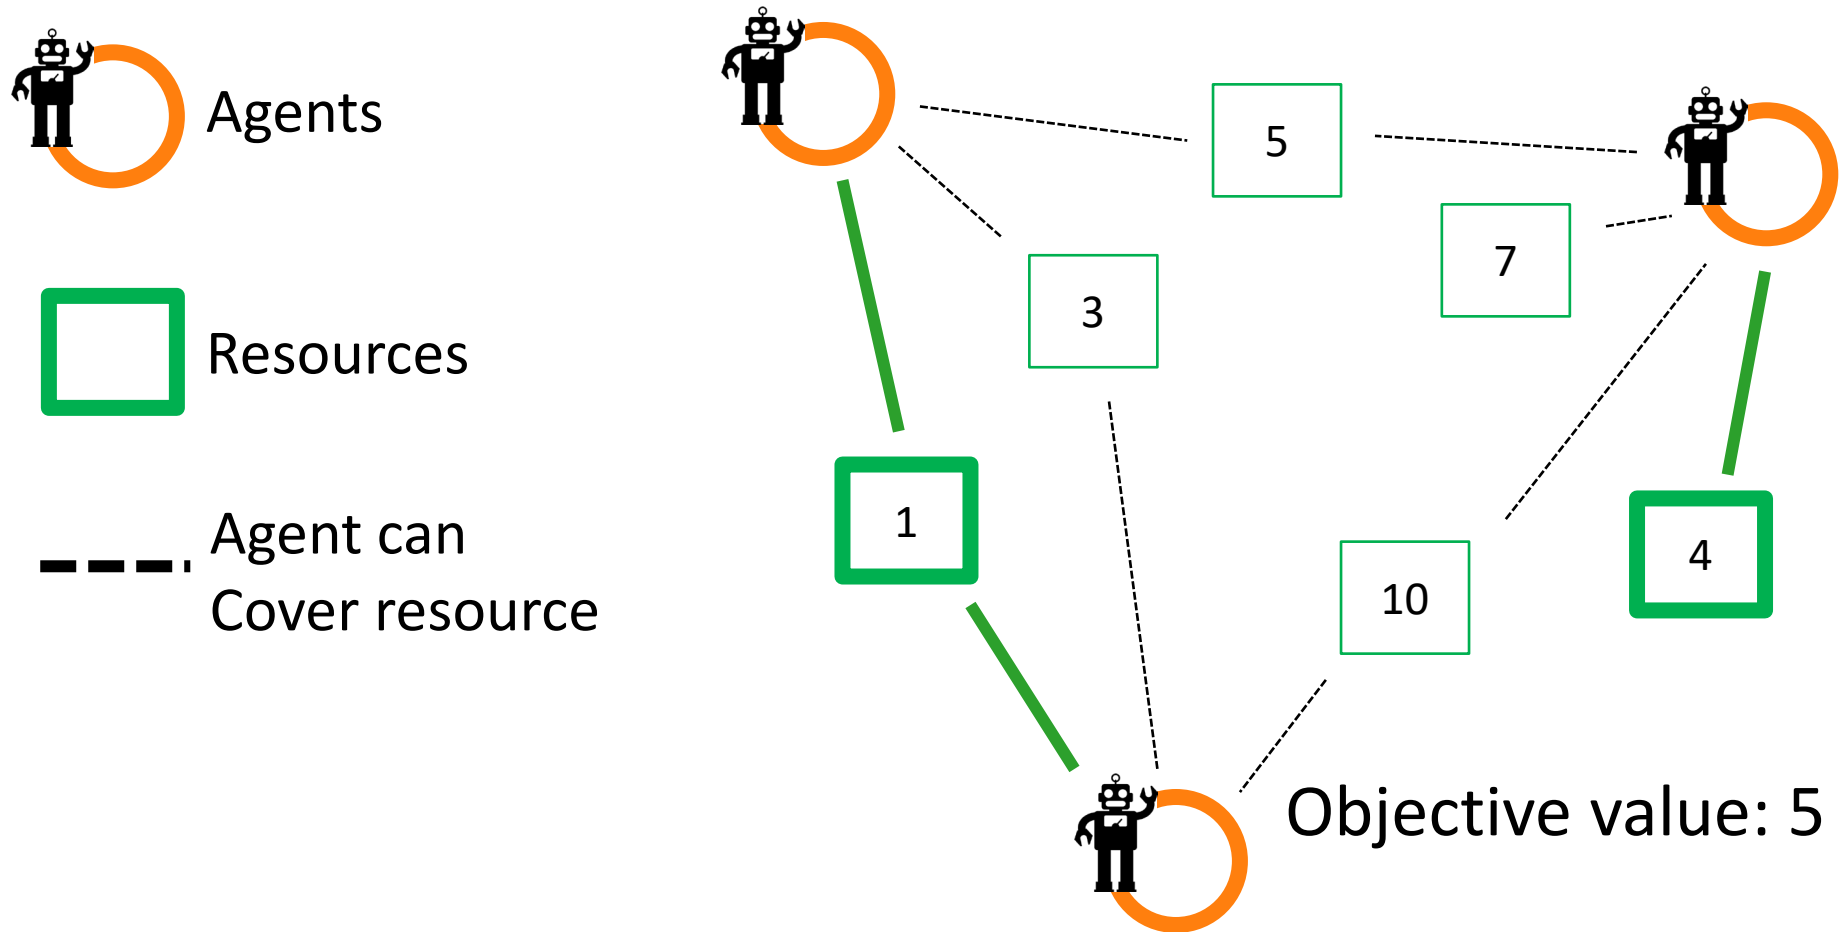


blind?

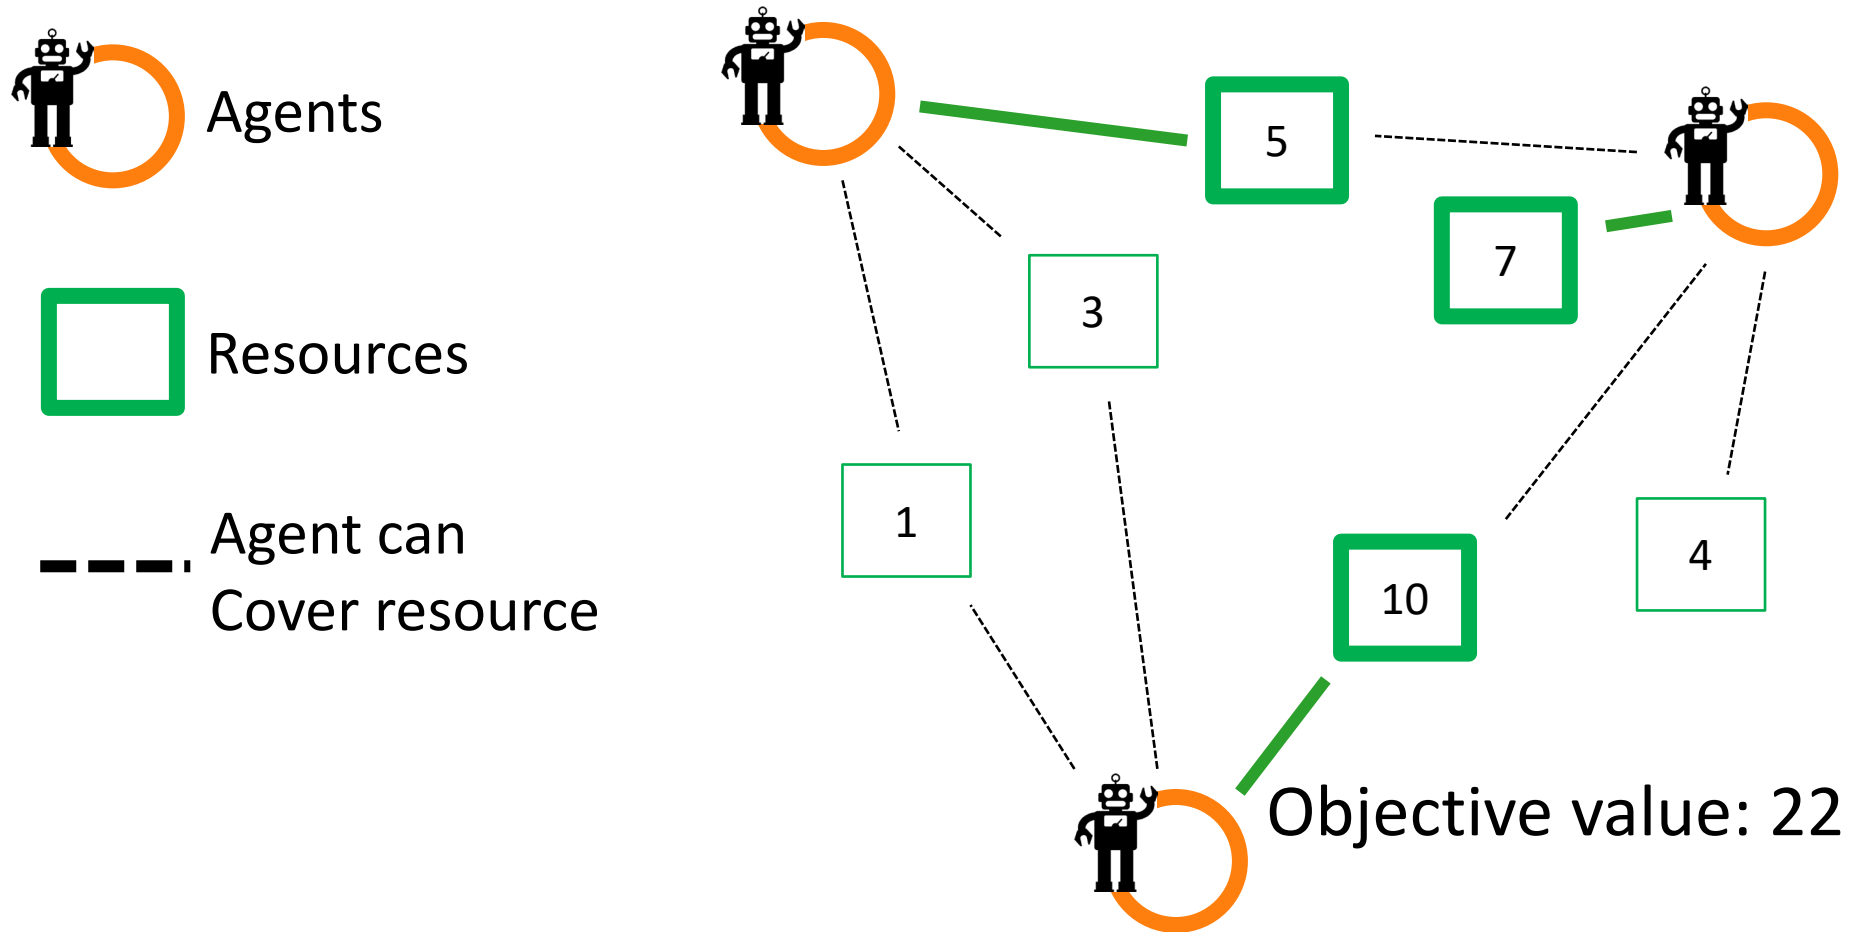
How should interacting autonomous agents behave in **compromised** environments?




Agents' Goal: maximize value of covered resources  
(each can only cover 1 box)



Agents' Goal: maximize value of covered resources  
(each can only cover 1 box)




Agents' Goal: maximize value of covered resources  
(each can only cover 1 box)



What is the engineer's challenge?






**Centralized Paradigm:** Given the whole problem, what's the solution?

**Problem:** what if communicating whole problem isn't possible?




What is the engineer's challenge?



**Centralized Paradigm:** Given the whole problem, what's the solution?

**Problem:** what if communicating whole problem isn't possible?

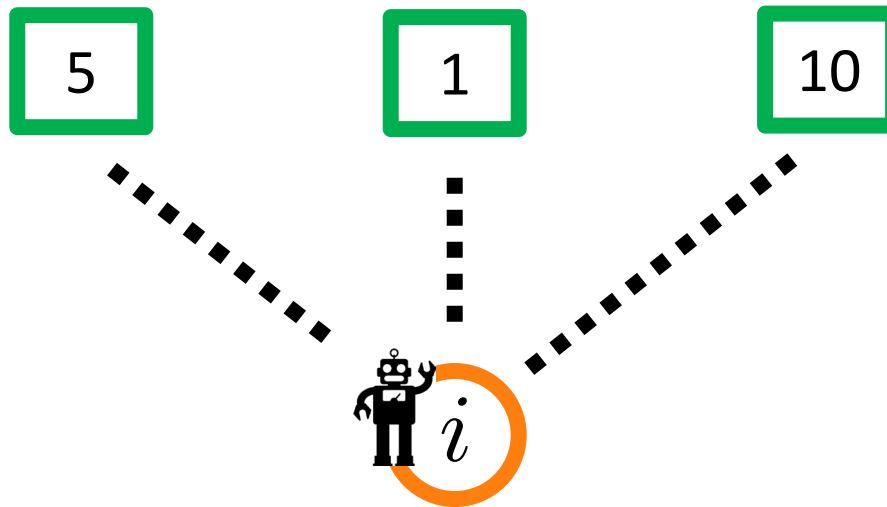


**Decentralized Paradigm:** What should individual agents do, given what they can see?

What is the engineer's challenge?








How should agent  $i$  value the resources it can cover?

**Altruism:** Value my choice at its contribution to system

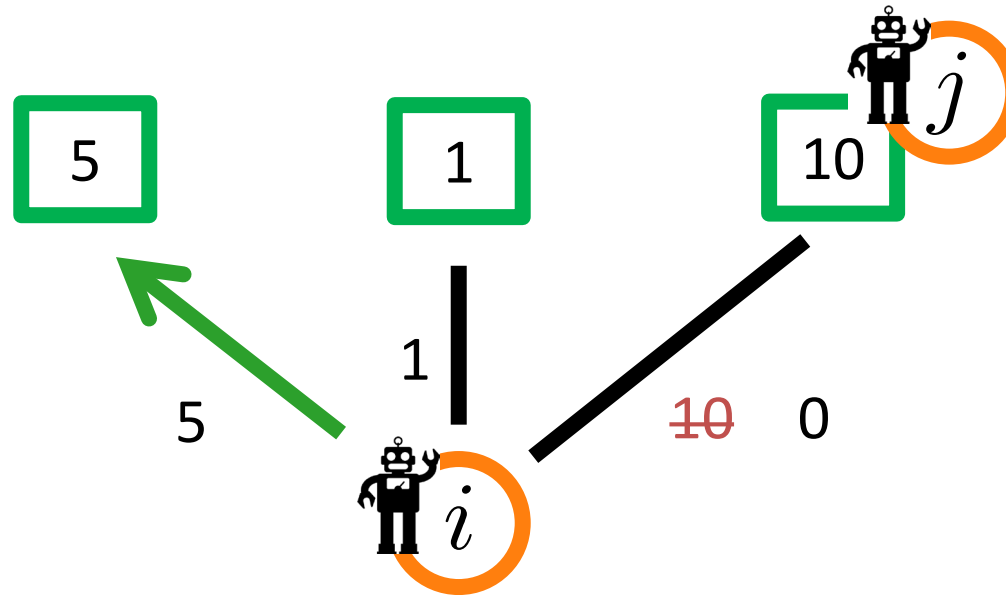
“Marginal Contribution”



How should agent  $i$  value the resources it can cover?

**Altruism:** Value my choice at its contribution to system

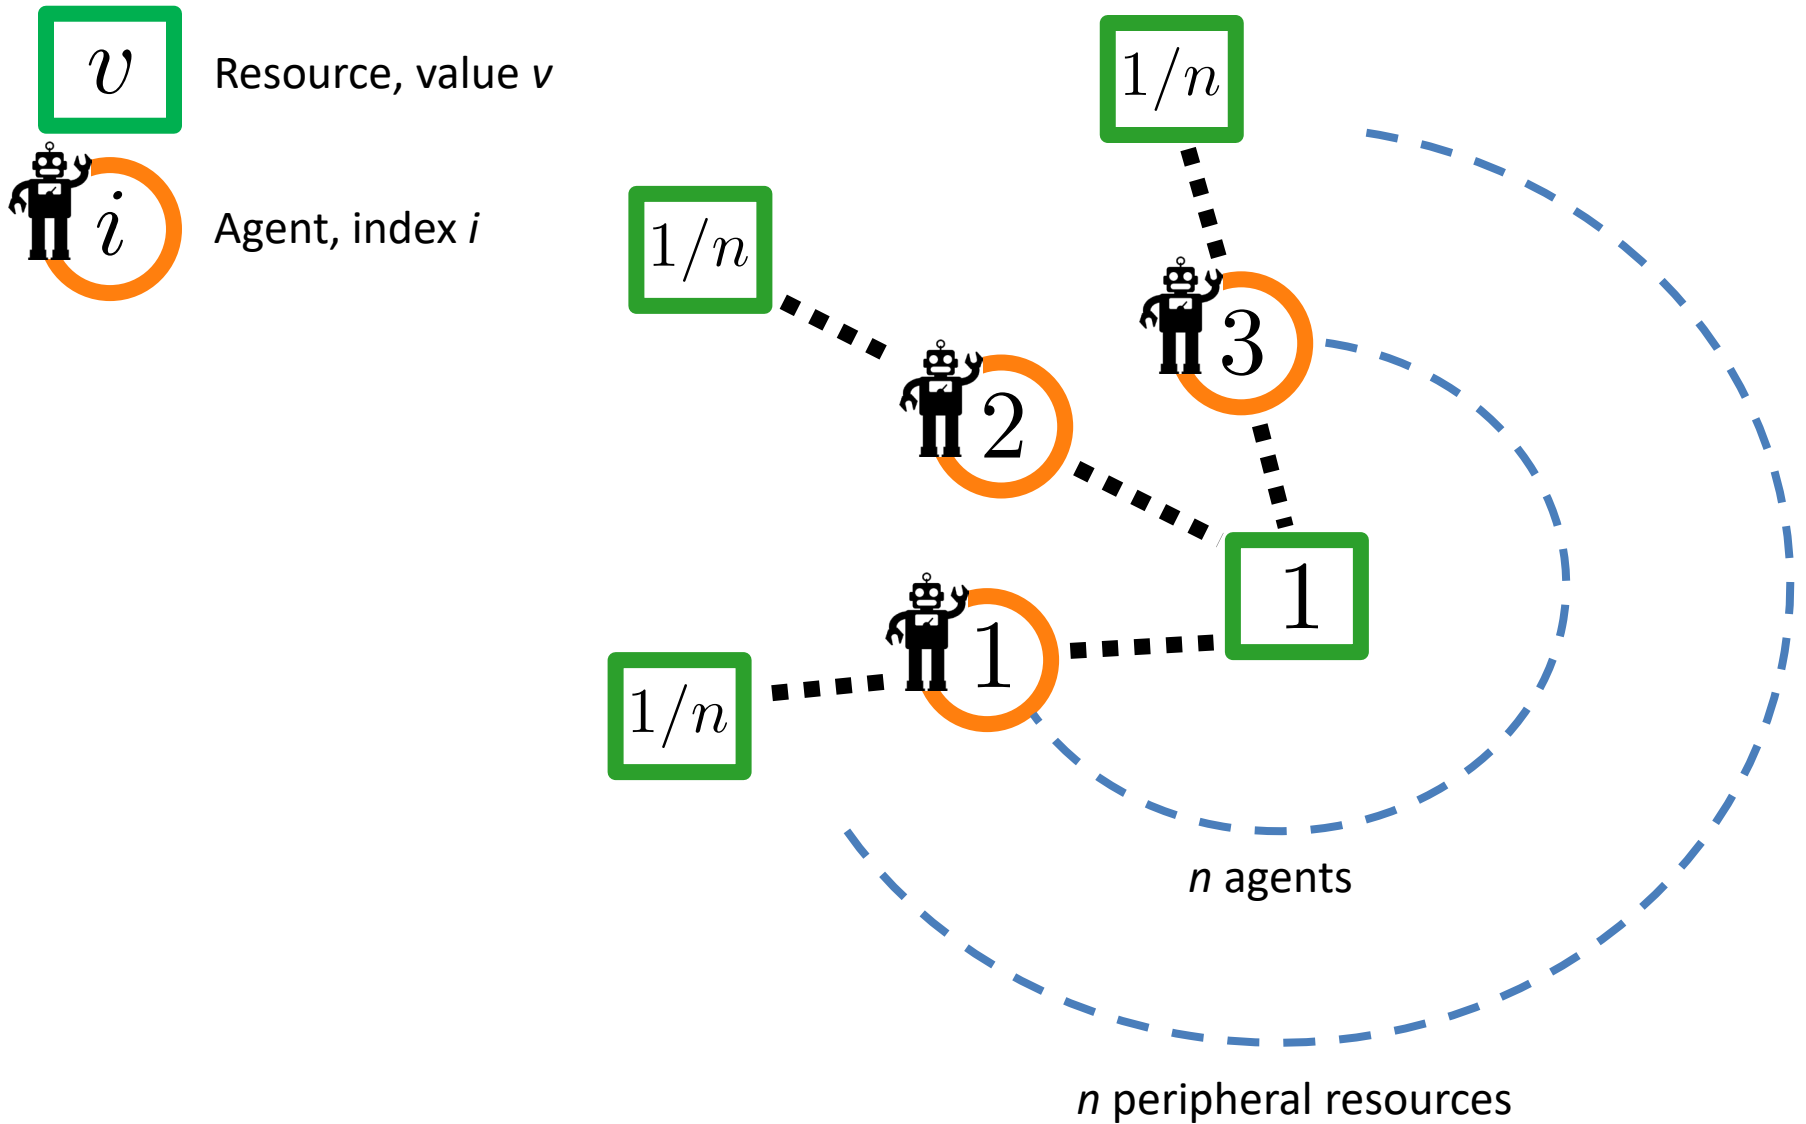
“Marginal Contribution”

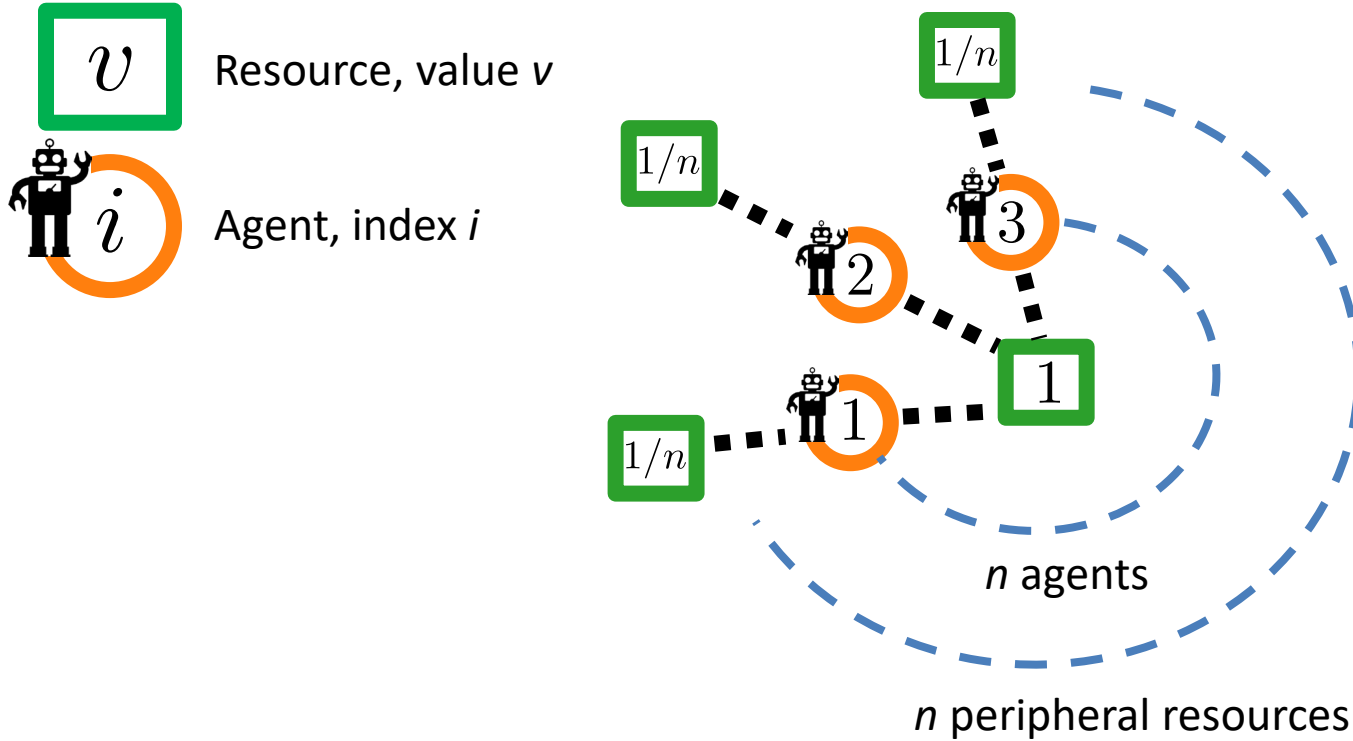


How should agent  $i$  value the resources it can cover?


**Altruism:** Value my choice at its contribution to system

“Marginal Contribution”






Q: What is good here?



Q: What is good here?

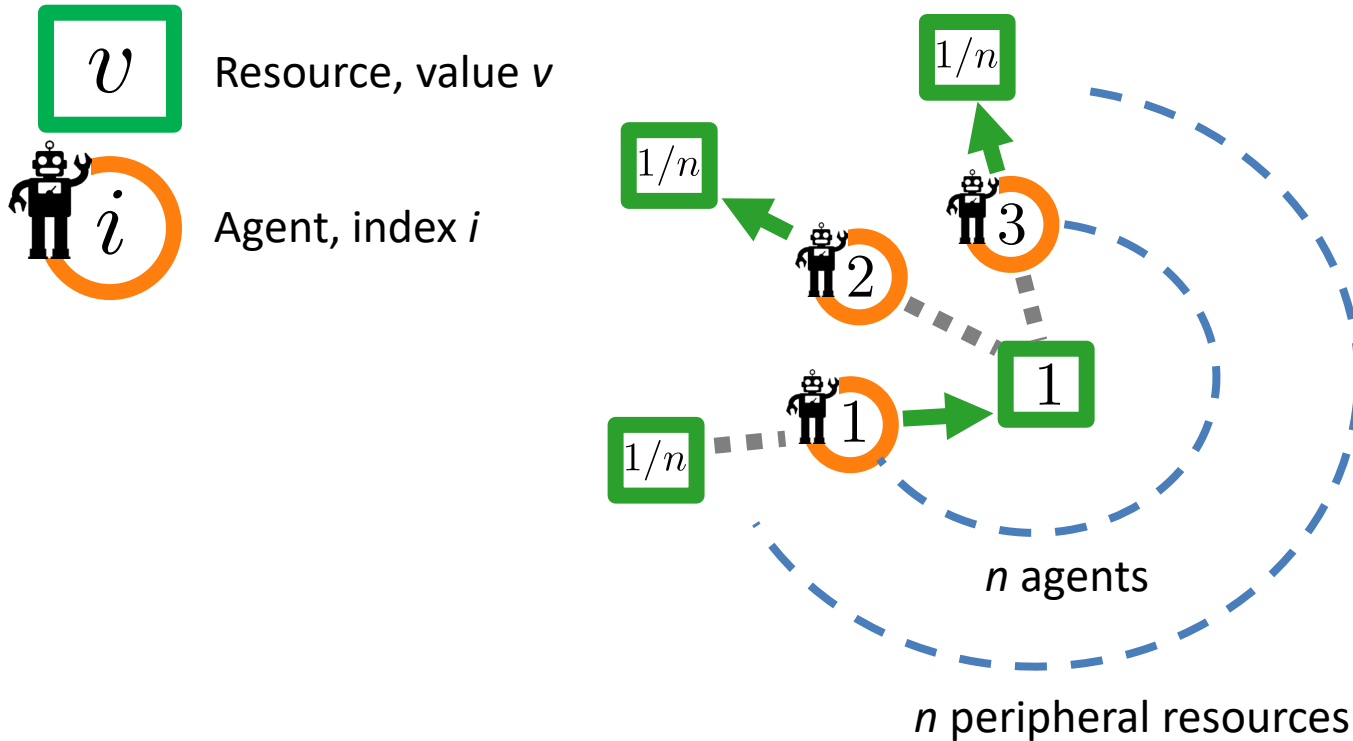
A: One agent in center, all others on periphery



Q: What is good here?

A: One agent in center, all others on periphery

Q: Does altruism get us there?



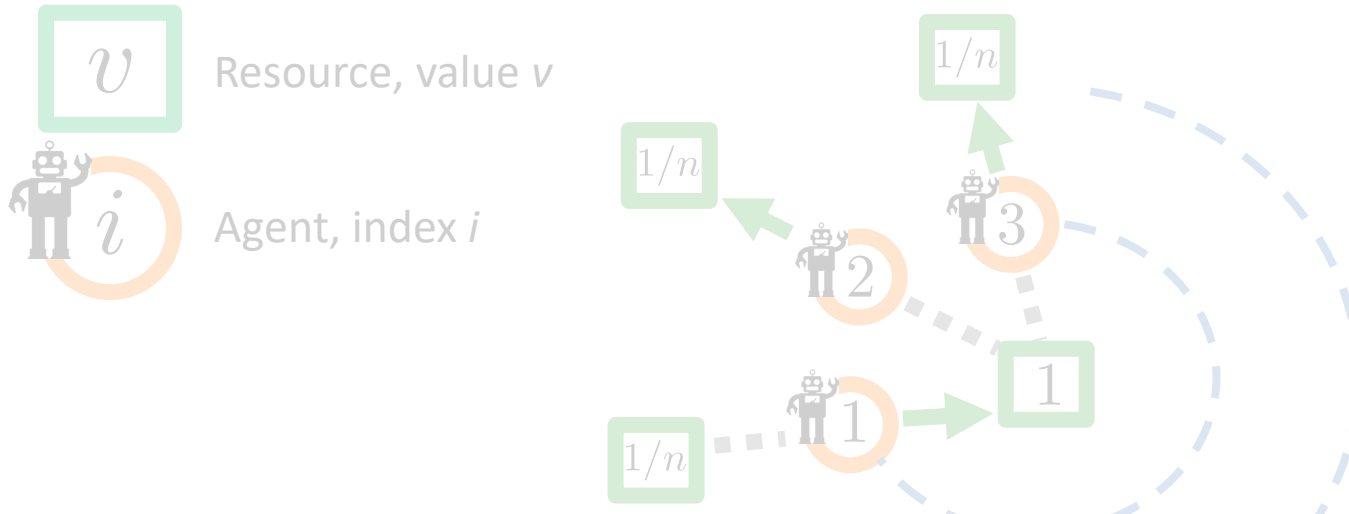
Q: What is good here?

A: One agent in center, all others on periphery

~~Q: Does altruism get us there?~~

Q: Are the Nash Equilibria with altruism optimal?





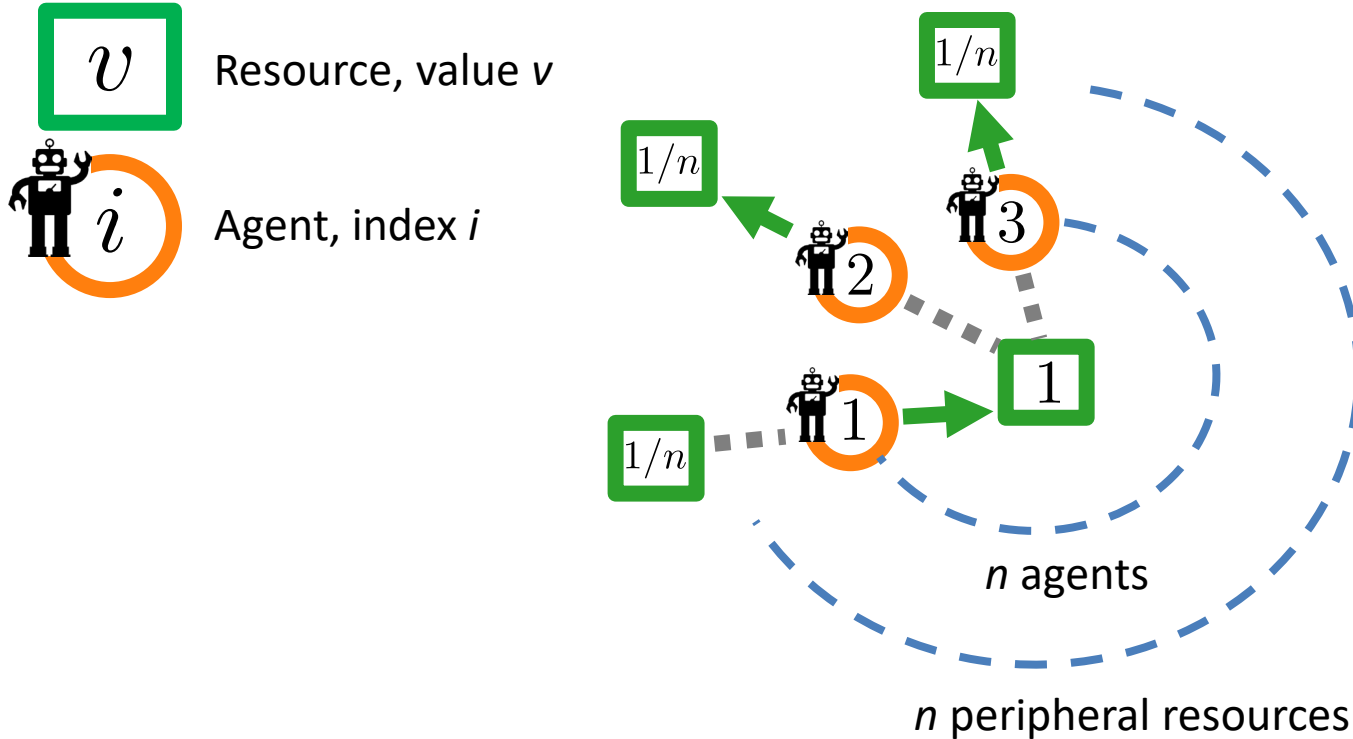
## Nash Equilibrium:

An outcome where *no agent* can change and be *strictly better off*.

Q: What

~~Q: Does altruism get us there?~~

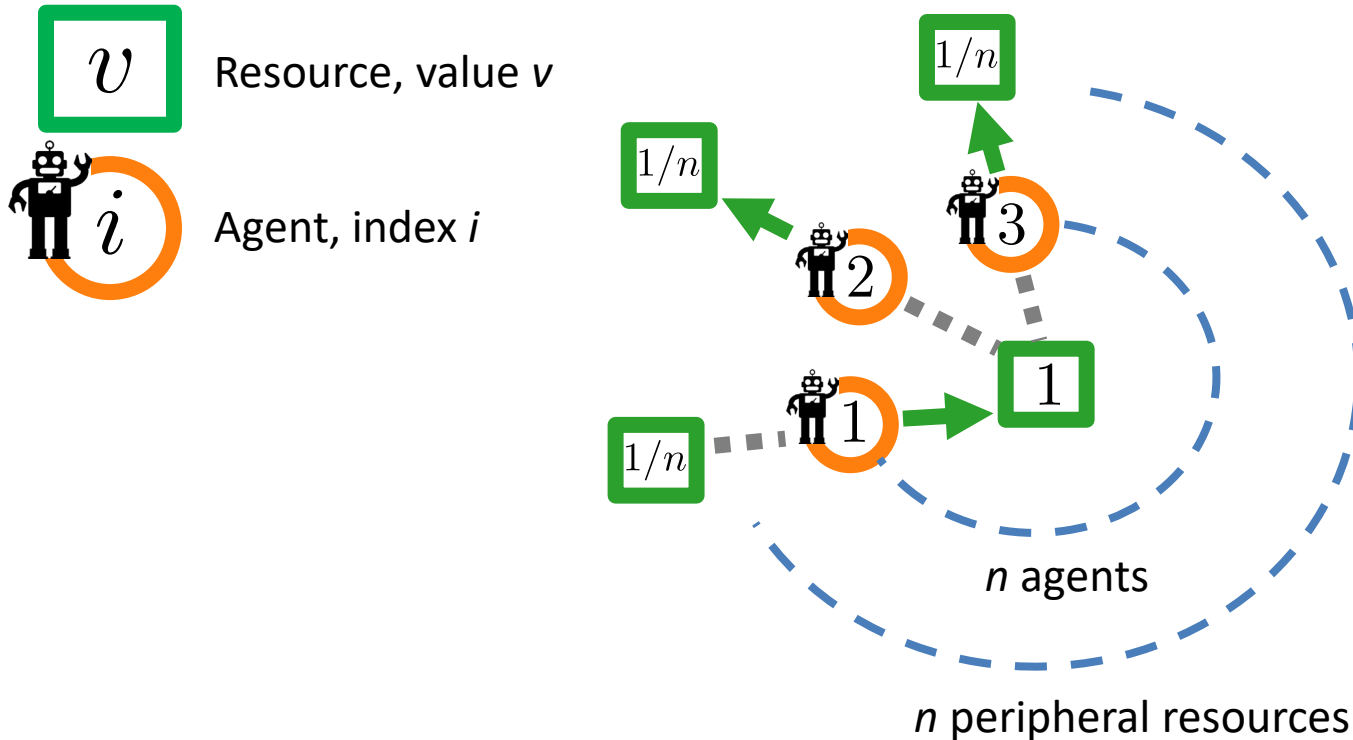
Q: Are the **Nash Equilibria** with altruism optimal?



Q: What is good here?

A: One agent in center, all others on periphery

Q: Does altruism get us there (altruistic NE optimal)?




Q: What is good here?


A: One agent in center, all others on periphery

Q: Does altruism get us there (altruistic NE optimal)?

A: Yes, robustly! ➔ Every Nash Equilibrium is optimal




Resource, value  $v$




Agent, index  $i$

Q: Does altruism get us there?  
A: Yes, robustly!


Universally?



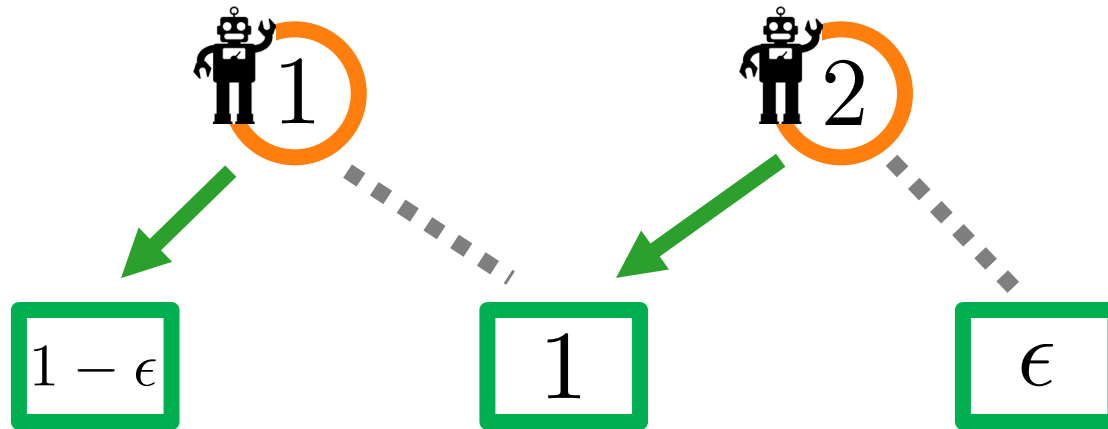
Q: What is good here?



Resource, value  $v$



Agent, index  $i$




Q: What is good here?


A: Left, Center

Q: Does altruism get us there?

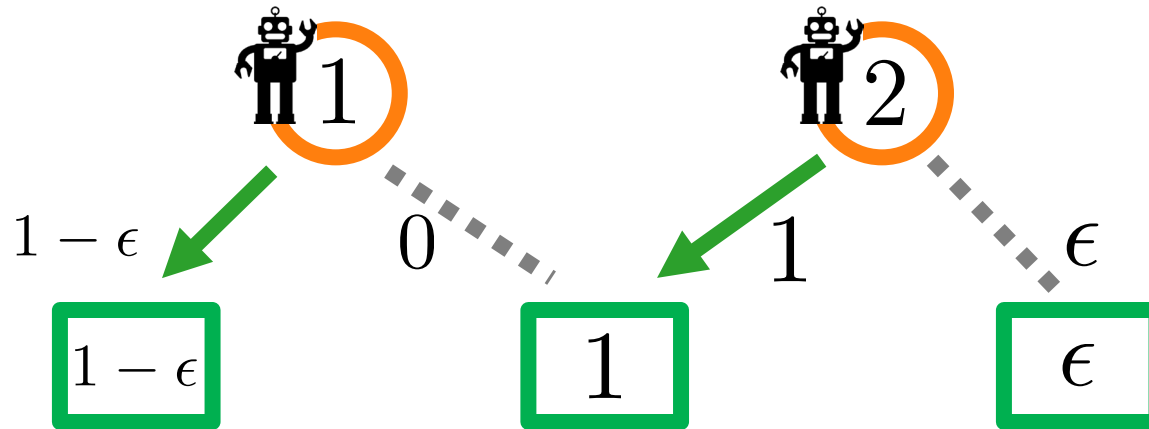
A: It can...



Resource, value  $v$



Agent, index  $i$




Q: What is good here?


A: Left, Center

Q: Does altruism get us there?


A: It can...



Resource, value  $v$



Agent, index  $i$




Q: What is good here?


A: Left, Center

Q: Does altruism get us there?

A: It can... But it need not.



Re



A

Fact: in this class of games,

Altruism



All Nash Equilibria  
within 50% of optimal


Q: What

$$\text{Price of Anarchy} = \max \frac{W(\text{NE})}{W(\text{Optimal})}$$


Q: Does a

A: It can... But it need not.





Re



A

Fact: in this class of games,

Altruism



All Nash Equilibria  
within 50% of optimal

Q: What


$$\text{PoA}(\text{altruism}) = \frac{1}{2}$$

Q: Does a

A: It can... But it need not.




Resource, value  $v$



Agent, index  $i$

Fact: in this class of games,  

$$\text{PoA}(\text{altruism}) = \frac{1}{2}$$




Q: What is good here?


A: Left, Center

Q: Does altruism get us there?

A: It can... But it need not.




Resource, value  $v$



Agent, index  $i$


Fact: in this class of games,  

$$\text{PoA}(\text{altruism}) = \frac{1}{2}$$




Q: What is better here?

A: "Equal Share" (encourage redundancy)




Resource, value  $v$



Agent, index  $i$


Fact: in this class of games,  

$$\text{PoA}(\text{altruism}) = \frac{1}{2}$$




Q: What is better here?

A: "Equal Share" (encourage redundancy)




Resource, value  $v$



Agent, index  $i$


Fact: in this class of games,  

$$\text{PoA}(\text{altruism}) = \frac{1}{2}$$




Q: What is better here?

A: "Equal Share" (encourage redundancy)




Resource, value  $v$



Agent, index  $i$

Fact: in this class of games,

$$\text{PoA}(\text{altruism}) = \frac{1}{2}$$




Q: Does this generalize?

A: No...  $\text{PoA}(\text{equal share}) = \frac{1}{2}$  as well

Q: Does this generalize?

A: No:  $\text{PoA}(\text{equal share}) = \frac{1}{2}$



$n$  agents on center



Each agent on center gets  $1/n$

Equal share over-incentivizes redundancy!

	Altruism	Equal Share
Intuition:	“Trusting” Agents	Promotes Redundancy
PoA:	$\frac{1}{2}$	$\frac{1}{2}$


No clear way to differentiate in general submodular games!

Note: in specific classes of games, more is known!

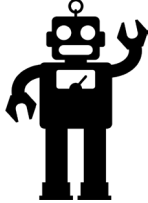
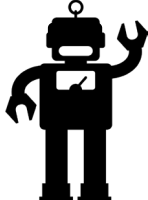


	Altruism	Equal Share
Intuition:	“Trusting” Agents	Promotes Redundancy
PoA:	$\frac{1}{2}$	$\frac{1}{2}$
Robustness:	???	???

Simple Starter question: what if K agents are “blind?”



Simple Starter question: what if  $K$  agents are “blind?”

		<u>Can see</u>	<u>Behavior</u>
Nominal:		Resource Values Choices of other agents	Maximize as usual
Blind:		Only Resource Values	Select highest-value resource available

Q: what are  $\left\{ \begin{array}{l} \text{PoA}(\text{altruism}; K) \\ \text{PoA}(\text{equal share}; K) \end{array} \right.$

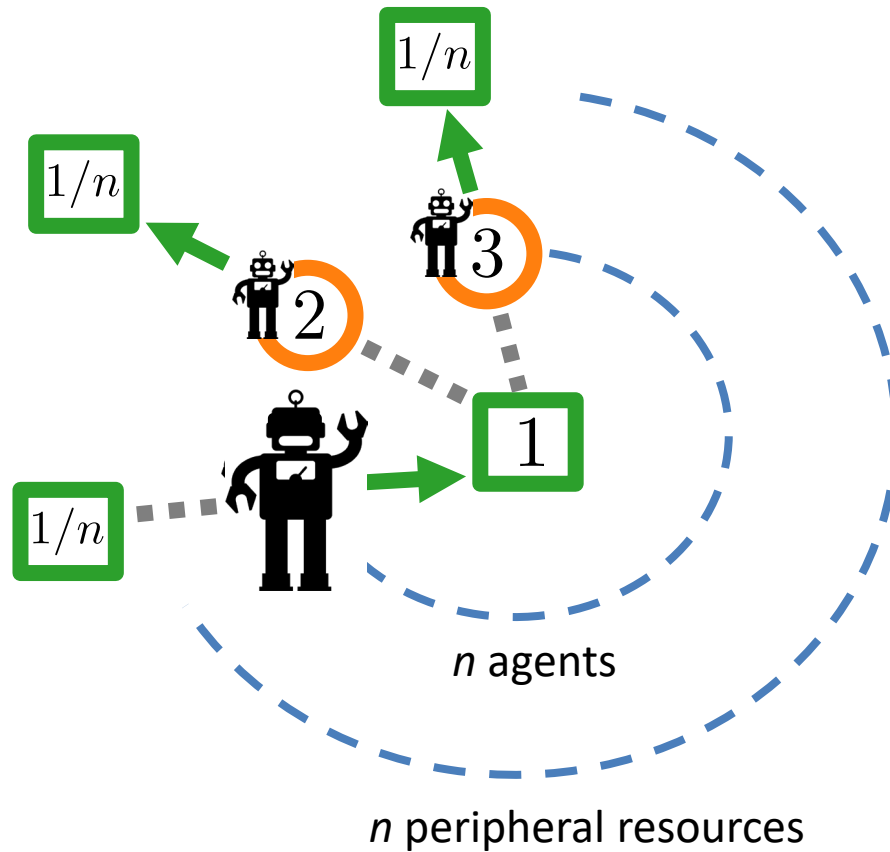
Q: Example 1?

Blind agent chooses center




Others altruistic: go outside

$$PoA = 1$$



Q: Example 1?




Blind agent chooses center


Others altruistic: go outside  
 $PoA = 1$

Others Equal share: go center  
 $PoA = \frac{1}{2}$

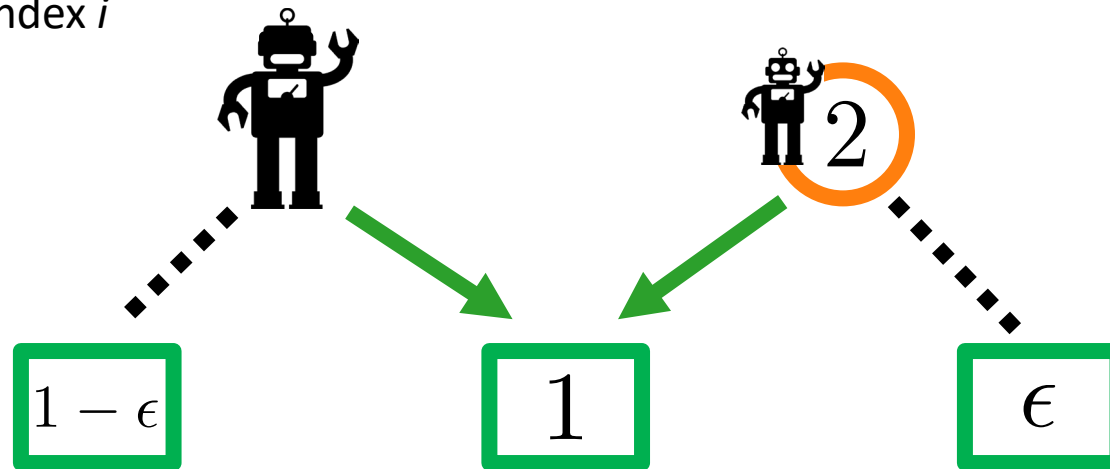
Blind agent changes nothing!



Resource, value  $v$




Agent, index  $i$




Blind agent always chooses **center**.

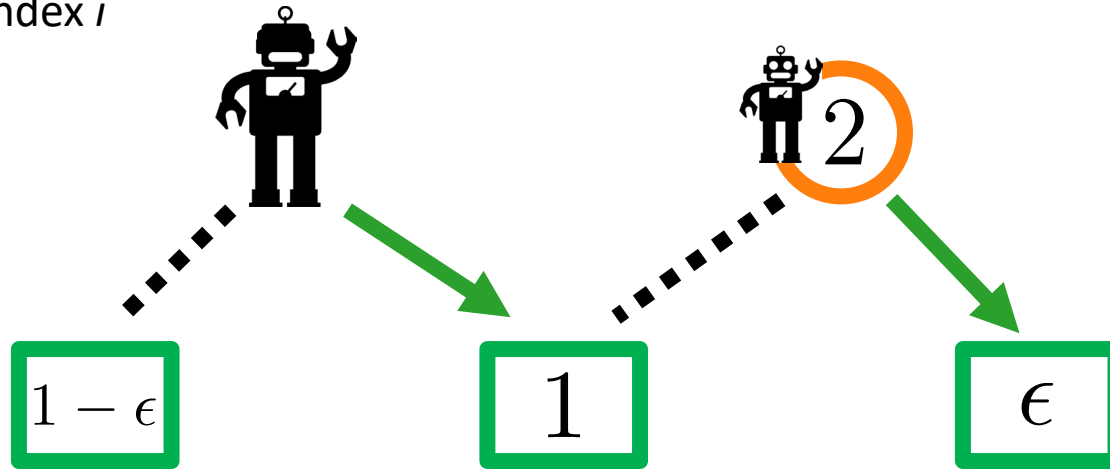
Nominal agent chooses **center** if equal share,



Resource, value  $v$



Agent, index  $i$



Blind agent always chooses **center**.

Nominal agent chooses **center** if equal share,  
**right** if altruistic.

$$\left\{ \text{PoA} = \frac{1}{2} \right.$$

	Altruism		Equal Share	
	Nominal	Blinded	Nominal	Blinded
Example 1:	1	1	$\frac{1}{2}$	$\frac{1}{2}$
Example 2:	$\frac{1}{2}$	$\frac{1}{2}$	1	$\frac{1}{2}$

Suggests: if blindness is a risk, agents should be altruistic

Now, some theory to support this notion...

**Theorem:**

If  $K$  agents are blind to the choices of others, then the following are true (in all submodular games):

$$\begin{aligned} \text{PoA}(\text{altruism}; K) &\geq \text{PoA}(\text{equal share}; K) \\ &= \frac{1}{2 + K} \end{aligned}$$


In fact,

$$\text{PoA}(\text{altruism}; K) = \frac{1}{1 + K}$$

Grimsman, Seaton, Marden and **P. N. Brown.**, "The Cost of Denied Observation in Multiagent Submodular Optimization," *CDC2020*

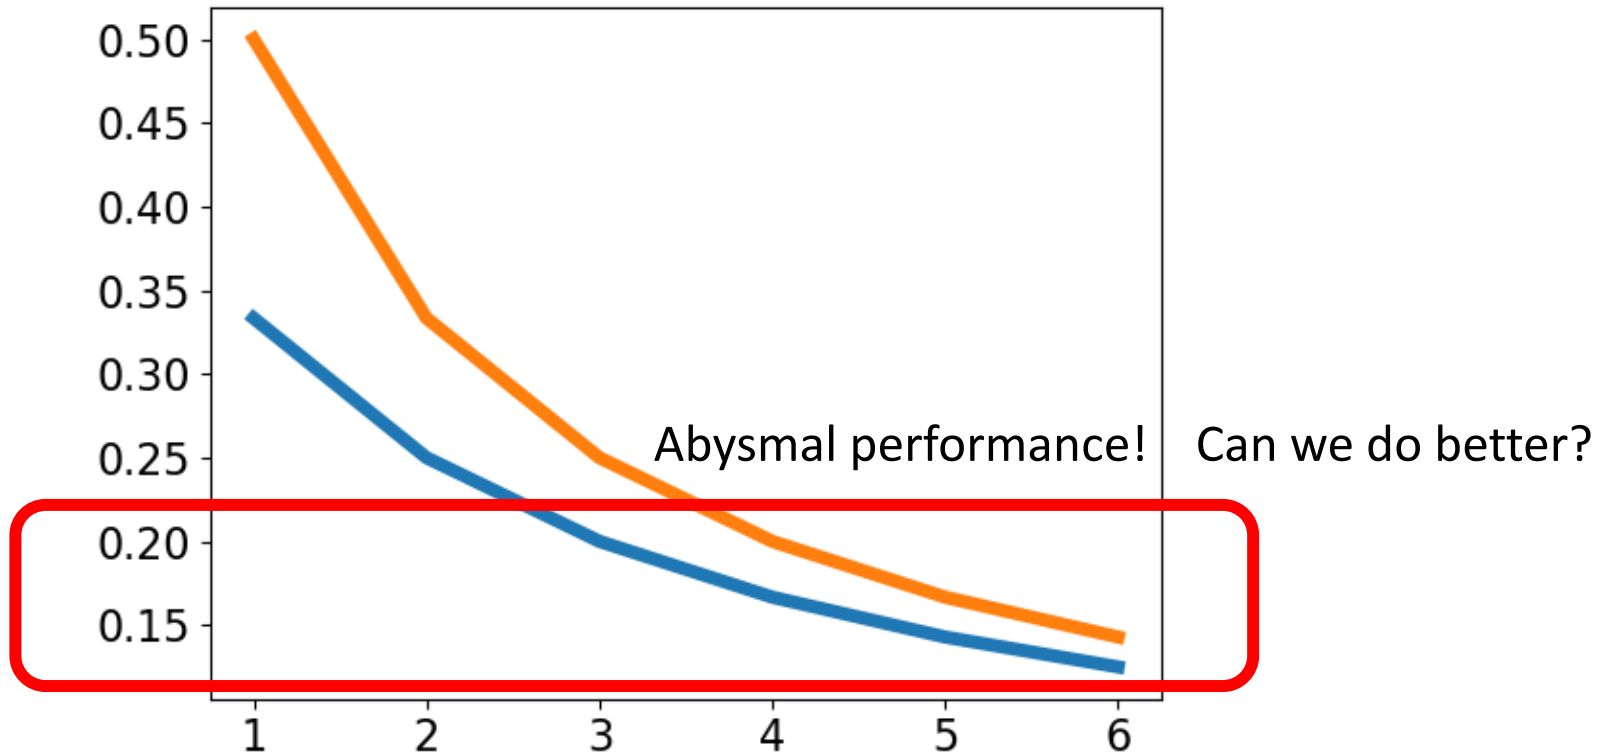


- PoA(altruism;  $K$ ) (exact)
- PoA(equal share;  $K$ ) (exact)



Grimsman, Seaton, Marden and **P. N. Brown.**, "The Cost of Denied Observation in Multiagent Submodular Optimization," *CDC2020*

- PoA(altruism;  $K$ ) (exact)
- PoA(equal share;  $K$ ) (exact)



Grimsman, Seaton, Marden and **P. N. Brown.**, "The Cost of Denied Observation in Multiagent Submodular Optimization," *CDC2020*

**Theorem:**

If  $K$  agents are blind and their “satisfaction” is  $S$ , then

$$\text{PoA}(\text{altruism}; K, S) \geq \frac{1}{1 + K - S}$$

**Plainly:**

The only way an outcome can be **very bad** is if  $S=0$  (no agent cares much about what action it’s choosing)

J. Seaton and **P. N. Brown**, "All Low-Quality Equilibria are Unstable in Submodular Maximization with Communication-Denied Agents" L-CSS, 2022

## Plainly:

The only way an outcome can be **very bad** is if  $S=0$  (no agent cares much about what action it's choosing)


## So what?

Compromised agents shouldn't take their perception too seriously, and should randomize their choices!


J. Seaton and **P. N. Brown**, "All Low-Quality Equilibria are Unstable in Submodular Maximization with Communication-Denied Agents" L-CSS, 2022

## So what?

Compromised agents shouldn't take their perception too seriously, and should randomize their choices!




J. Seaton and **P. N. Brown**, "All Low-Quality Equilibria are Unstable in Submodular Maximization with Communication-Denied Agents" L-CSS, 2022

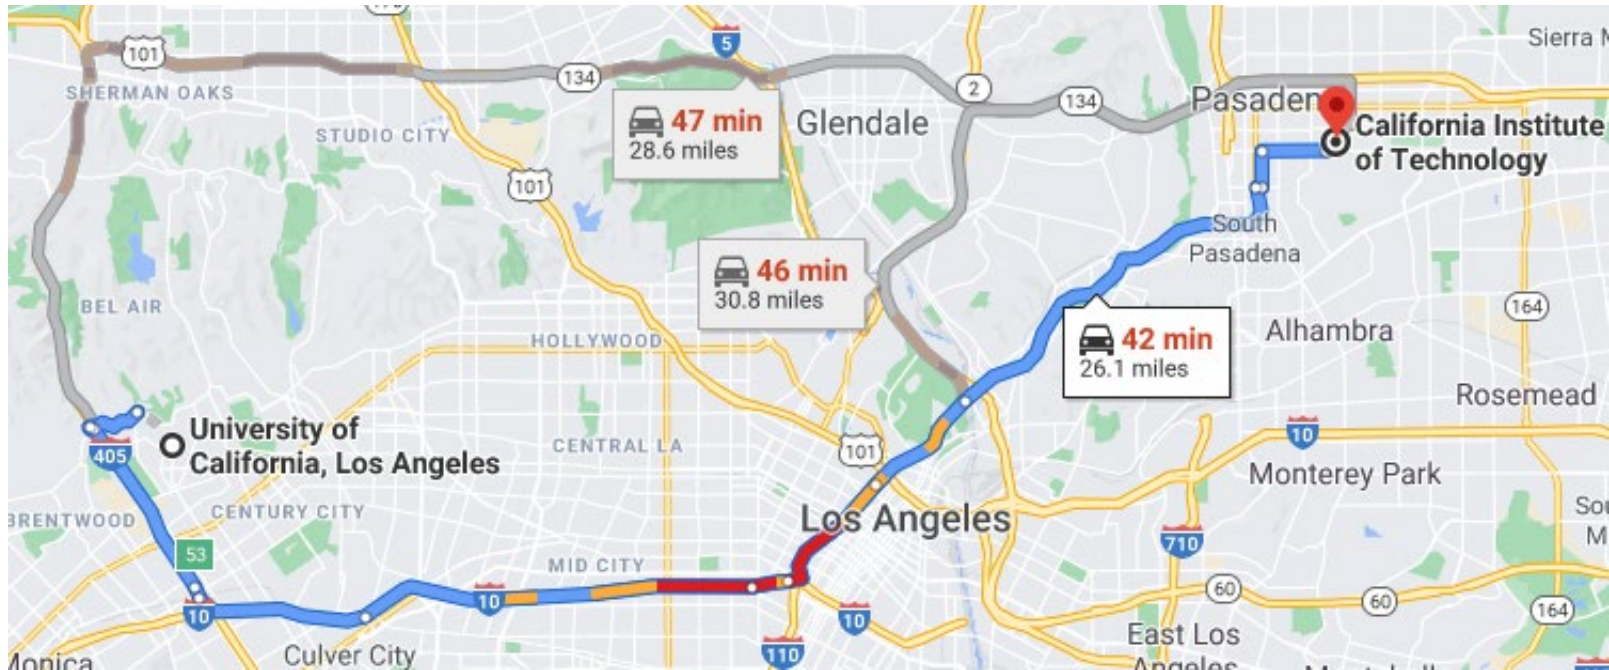



How should interacting autonomous agents behave in **compromised** environments?

1. If others may be blind, be altruistic!
2. If you are blind, dance around!




# Choosing Routes in Highway Networks






## Choosing Routes in Highway Networks




Agenda: pose simple model  
Explore behavior in its context





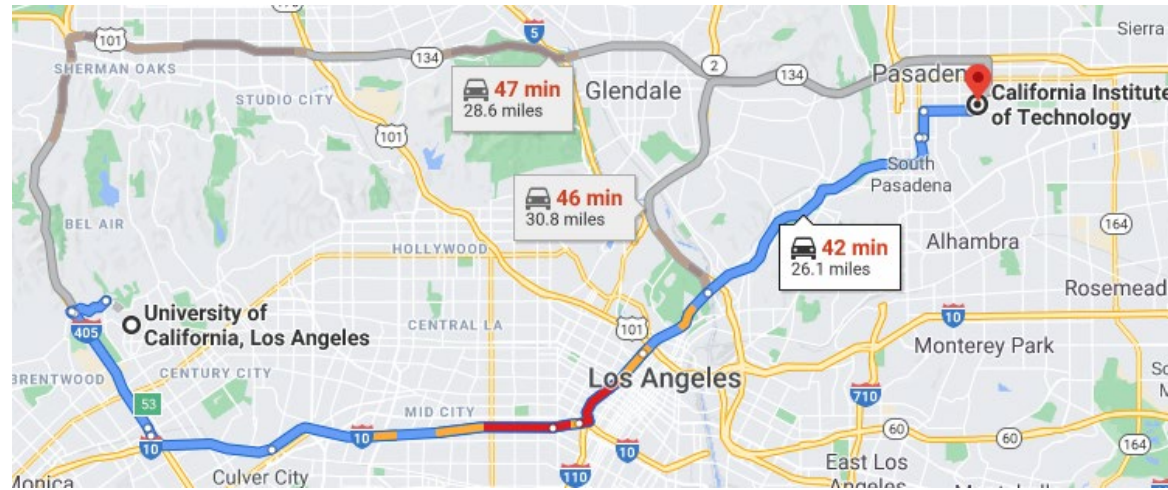
# Choosing Routes in Highway Networks

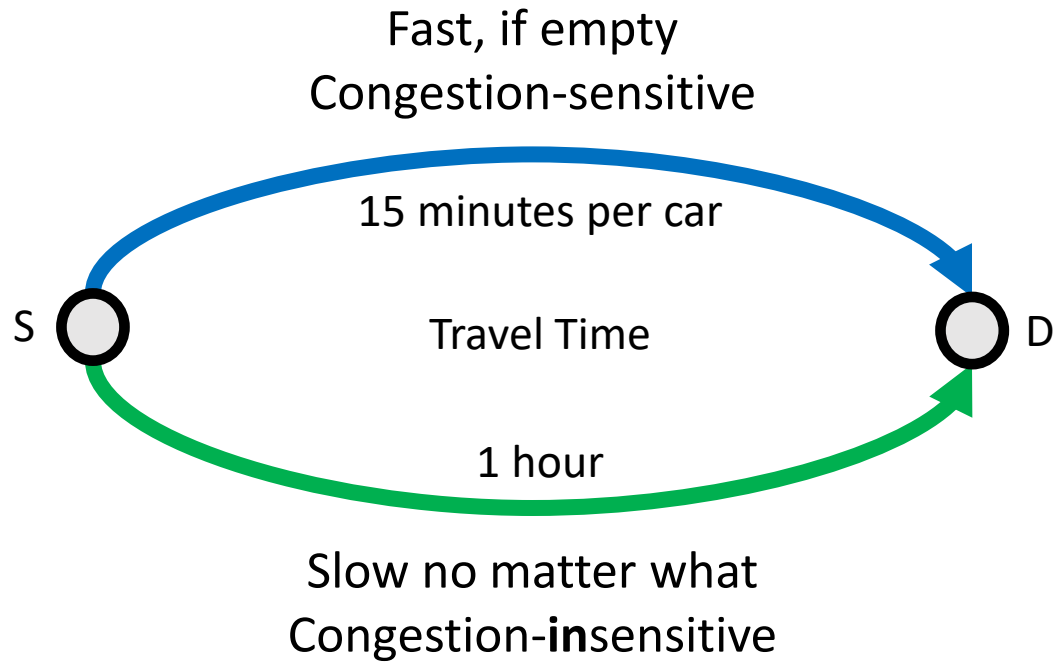
?


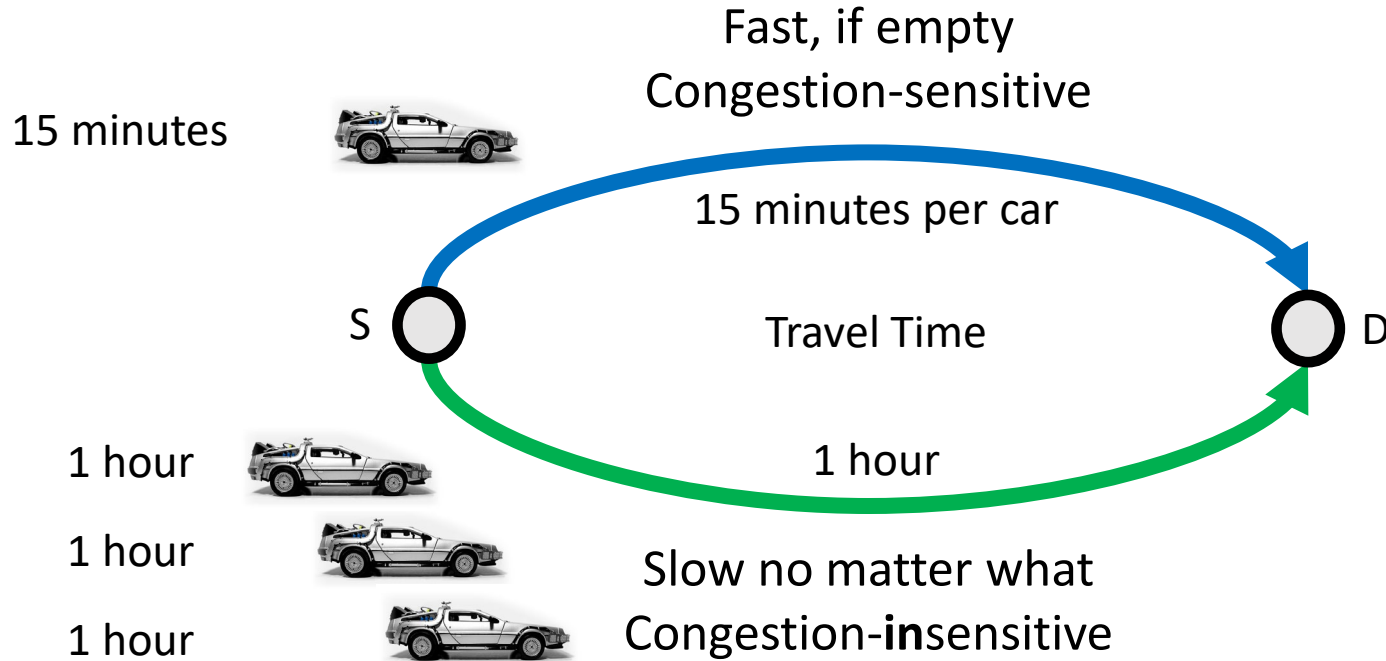


Agenda: pose simple model  
Explore behavior in its context

Question: should self-driving cars be altruistic?








Total Time  
3.25 hours

30 minutes



Fast, if empty  
Congestion-sensitive

30 minutes



1 hour




Slow no matter what  
Congestion-insensitive

1 hour




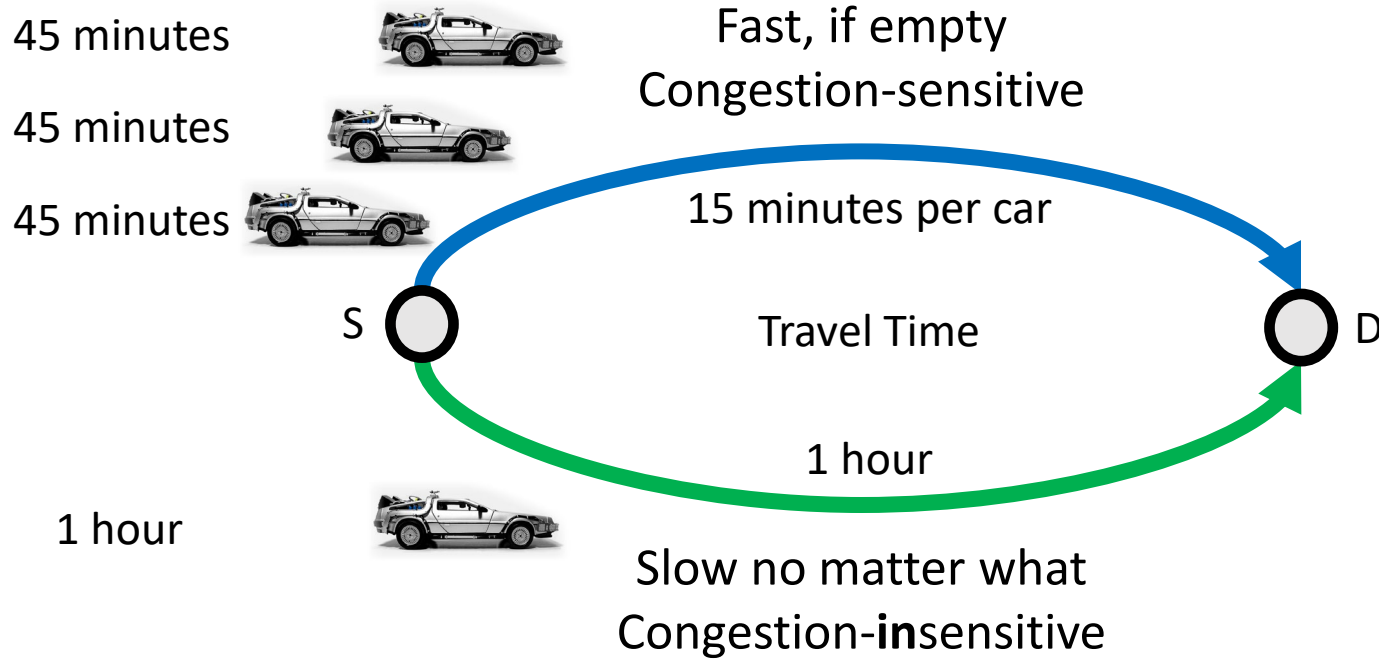
Traffic



Total Time

3.25 hours

3 hours



1 hour



Fast, if empty  
Congestion-sensitive

1 hour



1 hour



1 hour



S



15 minutes per car





Travel Time



1 hour

Slow no matter what  
Congestion-insensitive

Traffic



Total Time

3.25 hours

3 hours

3.25 hours

4 hours

1 hour



1 hour



1 hour



1 hour



S

15 minutes per car


Travel Time

D


1 hour

Best option!  
(Pareto optimal)


Traffic




3.25 hours



3 hours



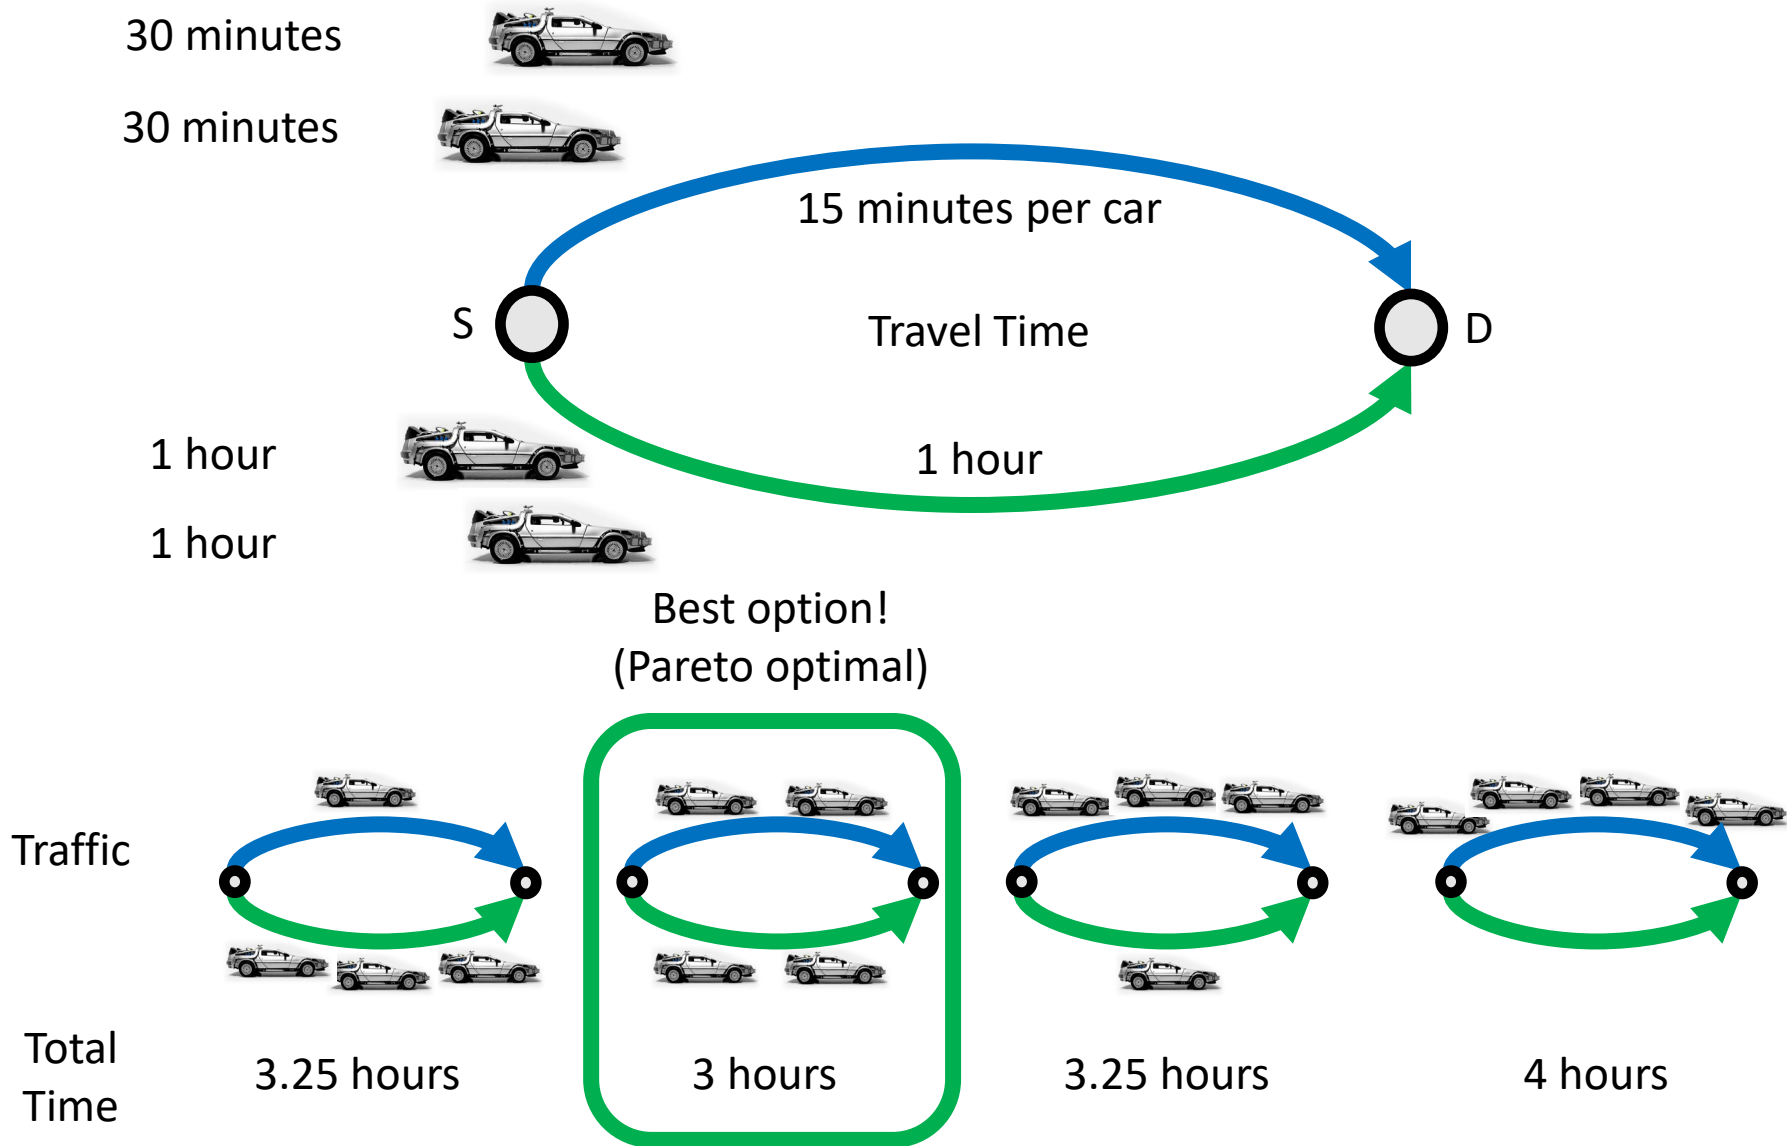
3.25 hours



4 hours

Total  
Time



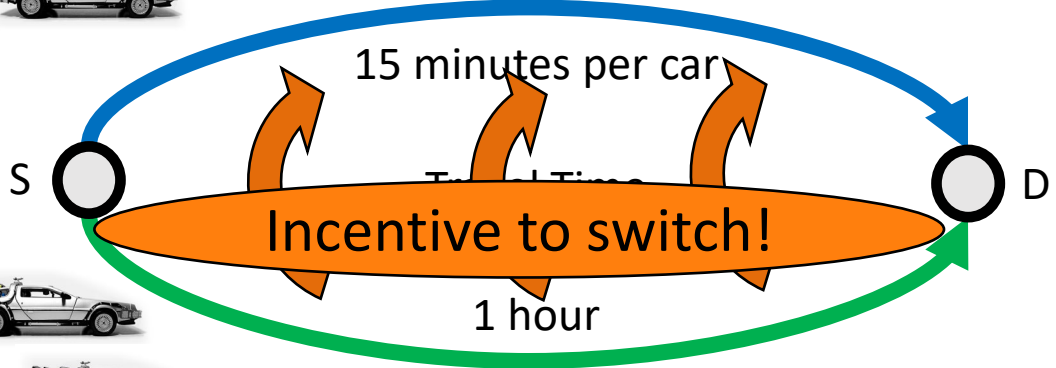


Altruistic self-driving cars?

30 minutes



30 minutes



1 hour







1 hour



Best option!  
(Pareto optimal)

Selfish traffic  
is like this!

Traffic



Total  
Time


3.25 hours

3 hours

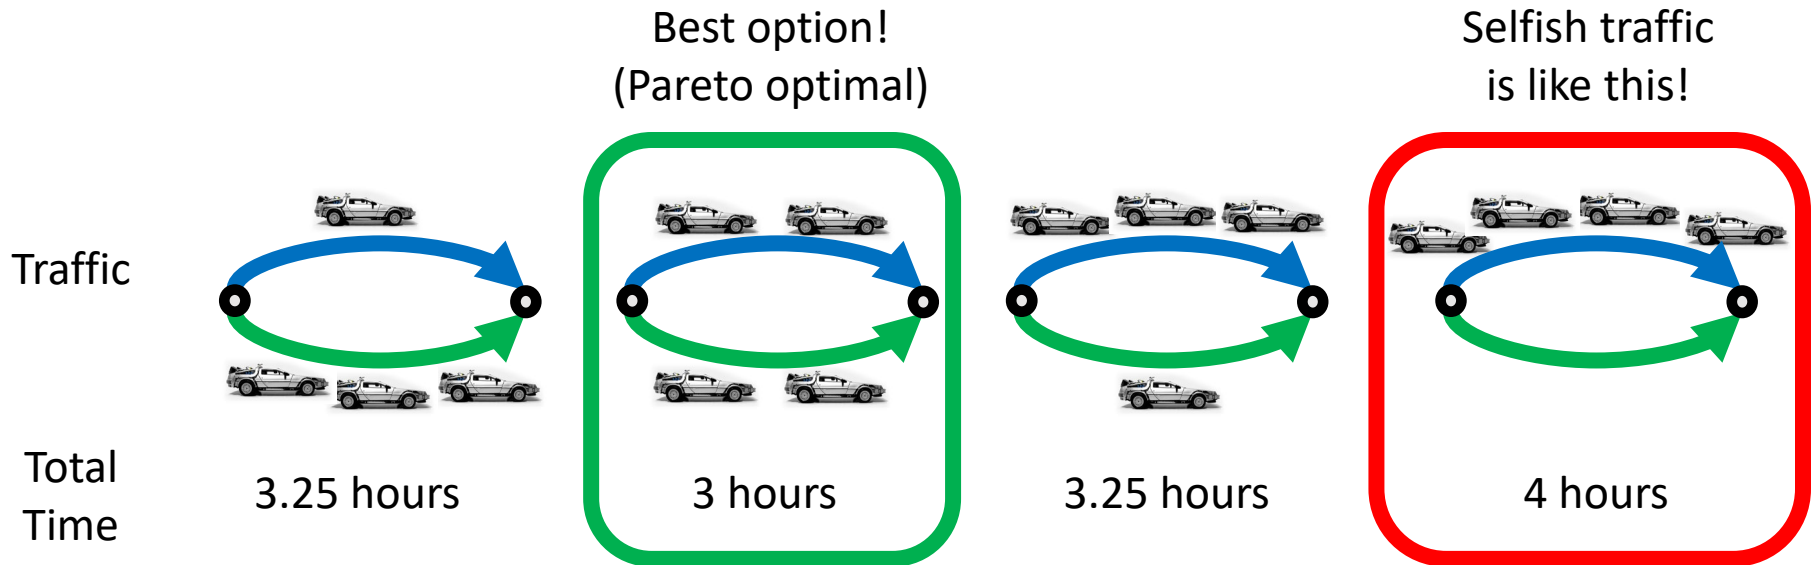
3.25 hours

4 hours

# Altruistic self-driving cars?



Altruism: act like there is 2x actual traffic



1 hour



1 hour



1 hour



1 hour



S

15 minutes per car

Travel Time

D





1 hour



Best option!  
(Pareto optimal)

Selfish traffic  
is like this!

Traffic



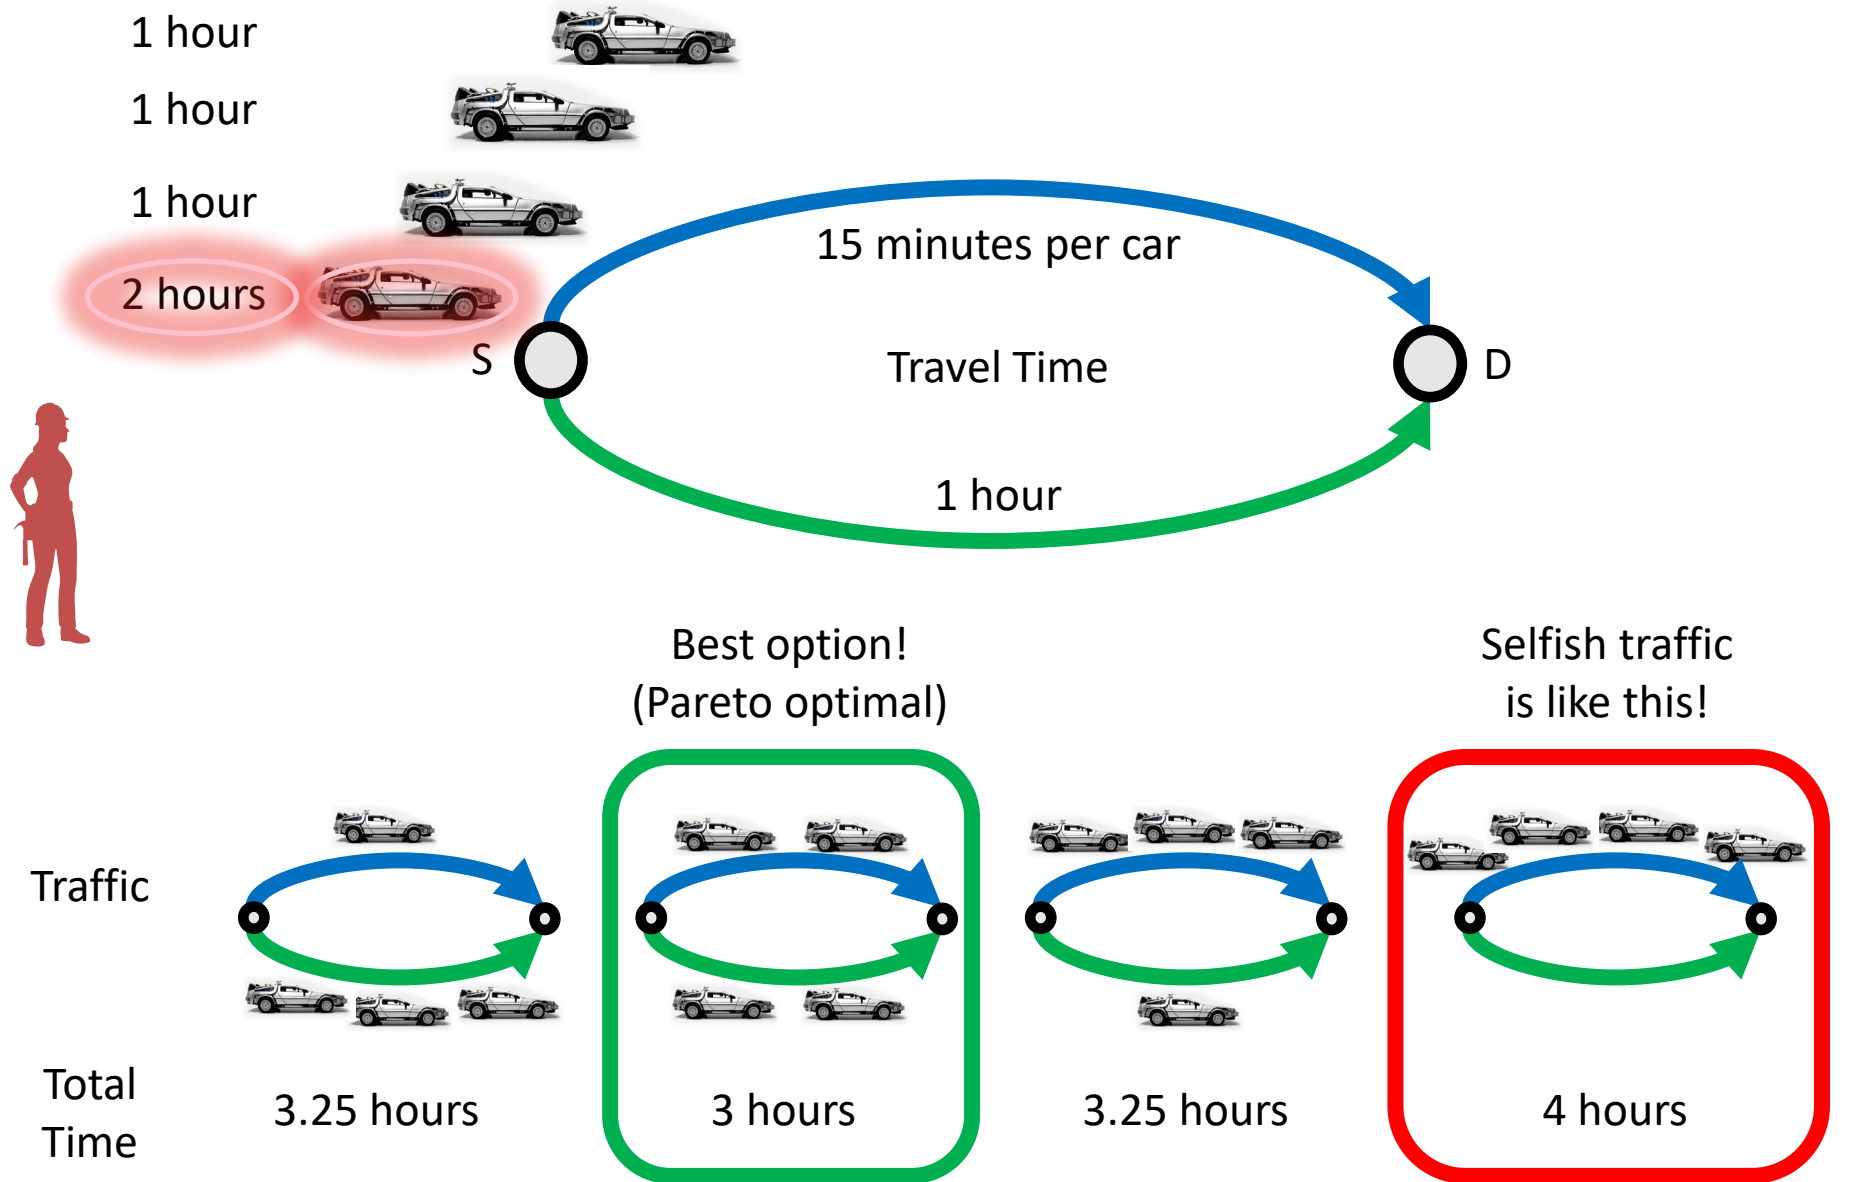
Total  
Time

3.25 hours

3 hours

3.25 hours

4 hours



45 minutes



45 minutes



45 minutes



S

15 minutes per car

Travel Time

D


1 hour

1 hour




Best option!  
(Pareto optimal)


Traffic




3.25 hours



3 hours



3.25 hours



4 hours

Total Time

45 minutes



45 minutes



90 minutes



15 minutes per car

S

Travel Time

D





1 hour

1 hour



Best option!  
(Pareto optimal)

Traffic



Total Time

3.25 hours

3 hours


3.25 hours

4 hours

30 minutes




30 minutes



Best option!  
(Pareto optimal)

Traffic



Total Time

3.25 hours

3 hours

3.25 hours

4 hours




## Altruistic self-driving cars:

- Improve congestion
- Even if only some are altruistic
- Without making others worse off
- Unambiguously good?



Traffic




Total  
Time


3.25 hours

3 hours


3.25 hours

4 hours






Selfish Traffic




Total Time

8 hours

# Altruistic self-driving cars?




## Selfish Traffic




Total Time

8 hours

# Altruistic self-driving cars?



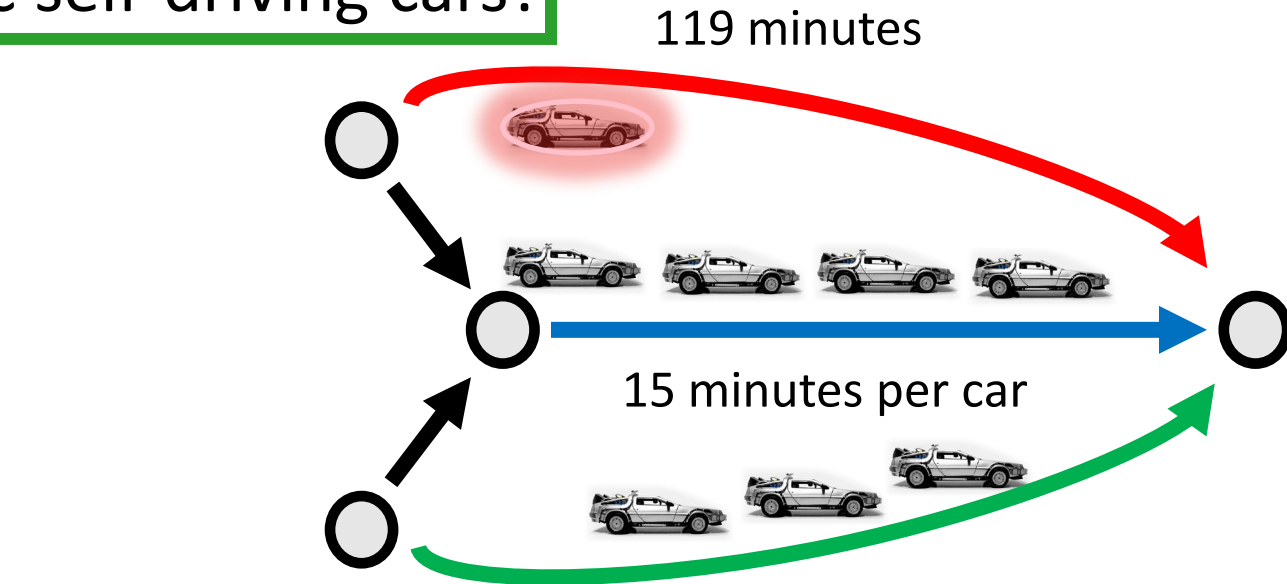
## Selfish Traffic




Total  
Time

8 hours

# Altruistic self-driving cars?




Selfish Traffic



8 hours

1 hour  
1 Altruistic Car




9 hours

Total Time

# Altruistic self-driving cars?





119 minutes



1 hour

1 Altruistic Car

Selfish Traffic


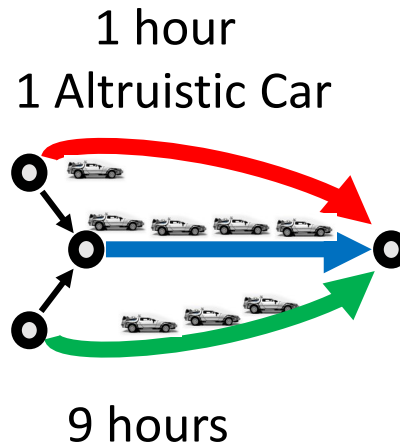




Total Time

8 hours

9 hours

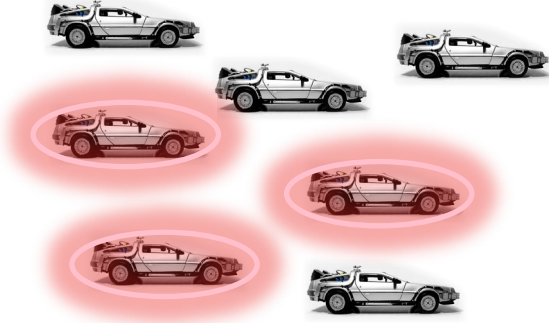
# Altruistic self-driving cars?



Total Time




## Socially-Embedded Autonomy



- Altruism helpful in some problems
- Can be harmful in complex networks

## Compromised Multiagent Systems



- If others are blind, be altruistic
- If you are blind, dance around!



My Website



Technical Talks 



ECCS-2013779



FA9550-23-1-0171