Designing AI Advice for Sequential Decision Making

Park Sinchaisri
UC Berkeley Haas

BDRM 2026

Joint work with **Philippe Blaettchen** (SMU)

Too Little Use...

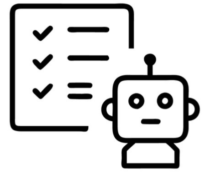

"Humans choose human forecasters over AI, esp. after seeing AI performs."

(Dietvorst, Simmons, Massey 2015)

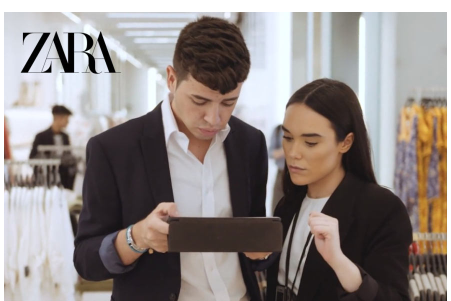

"Workers/radiologists/ managers deviate more as they gain experience."

(Ibanez, Clark, Huckman, Staats 2017,
Sun, Zhang, Hu, Van Mieghem 2022,
Caro & Saez de Tejada Cuenca 2023)

Too Little Use...

> "Humans choose human forecasters over AI, esp. after seeing AI performs."

Performance left on the table

> "Workers/radiologists/managers deviate more as they gain experience."

(Ibanez, Clark, Huckman, Staats 2017,
Sun, Zhang, Hu, Van Mieghem 2022,
Caro & Saez de Tejada Cuenca 2023)

full paper at bit.ly/3SJA13K

Too Little Use...

"Humans choose human forecasters over AI, esp. after seeing AI performs."

Performance left on the table

"Workers/radiologists/managers deviate more as they gain experience."

(Ibanez, Clark, Huckman, Staats 2017, Sun, Zhang, Hu, Van Mieghem 2022, Caro & Saez de Tejada Cuenca 2023)

...Too Much Use

The Dangers of Overreliance on Automation

Safety Concerns and Mitigation Strategies for Pilots

FAA Safety Briefing Magazine Follow 8 min read · May 2, 2025

NEW YORK POST

BUSINESS

ChatGPT outage live updates: College students spiraling during finals

By Taylor Herzlich, Ariel Zilber, Kaydi Pelletier and Isabella Bernabeo

Updated June 10, 2025, 12:51 p.m. ET

10 Comments

full paper at **bit.ly/3SJA13K**

Too Little Use...

"Humans choose human forecasters over AI, esp. after seeing AI performs."

Performance left on the table

"Workers/radiologists/ managers deviate more a they gain experience."

(Ibanez, Clark, Huckman, Staats 2017, Sun, Zhang, Hu, Van Mieghem 2022, Caro & Saez de Tejada Cuenca 2023)

...Too Much Use

The Dangers of Overreliance on Automation

Safety Concerns and Mitigation Strategies for Pilots

FAA Safety Briefing Magazine Follow 8 min read · May 2, 2025

NEW YORK POST

BUSINESS

ChatGPT outage live updates: College students spiraling during finals

By Taylor Herzlich, Ariel Zilber, Kaydi Pelletier and Isabella Bernabeo
Updated June 10, 2025, 12:51 p.m. ET

10 Comments

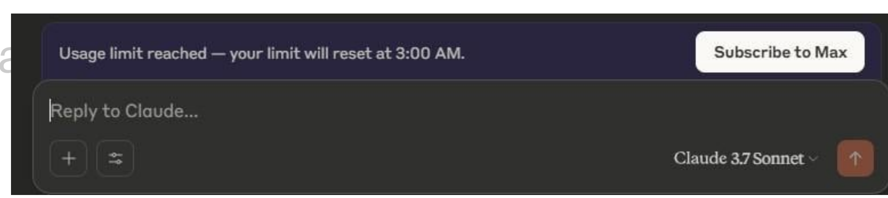

full paper at bit.ly/3SJA13K

Too Little Use...

“Humans choose human forecasters over AI, esp. after seeing AI performs.”

Performance left on the table

“Workers/radiologists/ managers deviate more a they gain experience.”

(Ibanez, Clark, Huckman, Staats 2017, Sun, Zhang, Hu, Van Mieghem 2022, Caro & Saez de Tejada Cuenca 2023)

...Too Much Use

Skill/capability quietly decays

Too Little Use...

“Humans choose human forecasters over AI, esp. after seeing AI performs.”

Performance left on the table

...Too Much Use

Skill/capability quietly decays

(Ibanez, Clark, Huckman, Staats 2017, Sun, Zhang, Hu, Van Mieghem 2022, Caro & Saez de Tejada Cuenca 2023)

Safety Concerns and Mitigation Strategies for Pilots

NEW YORK POST

How should AI advice be communicated for both performance and long-run capability?

Sequential Decision Making

t=1	t=2	t=3	t=4	t=5	
|	|	|	|	|	
[a_1]	[a_2]	[a_3]	[a_4]	[a_5]	← decisions
|	|	|	|	|	
r=0	r=0	r=0	r=0	r=-30	← rewards

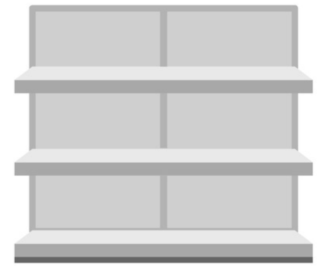

inventory management

EV charging during trip

full paper at **bit.ly/3SJA13K**

Sequential Decision Making

t=1 t=2 t=3 t=4 t=5

| | | | |

$[a_1]$ $[a_2]$ $[a_3]$ $[a_4]$ $[a_5]$ ← decisions

| | | | |

r=0 r=0 r=0 r=0 r=-30 ← rewards

Question: Which decision caused the bad outcome at t=5?

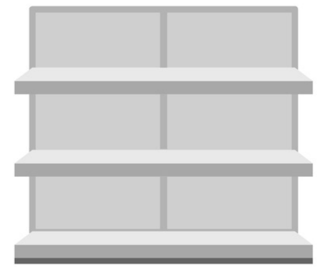

inventory management

EV charging during trip

Sequential Decision Making

t=1 t=2 t=3 t=4 t=5

$[a_1]$ $[a_2]$ $[a_3]$ $[a_4]$ $[a_5]$ ← decisions

r=0 r=0 r=0 r=0 r=-30 ← rewards

Question: Which decision caused
the bad outcome at t=5?

Today's action shapes tomorrow's options!

inventory
management

EV charging during trip

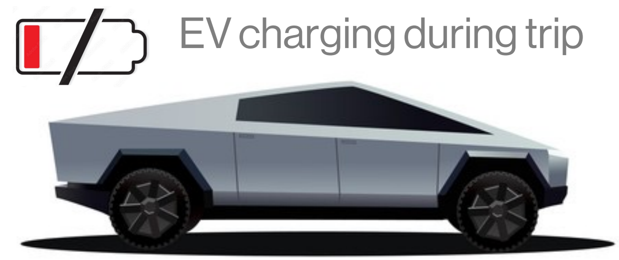

Sequential Decision Making

t=1 t=2 t=3 t=4 t=5

[a_1] [a_2] [a_3] [a_4] [a_5] ← decisions

r=0 r=0 r=0 r=0 r=-30 ← rewards

inventory
management

EV charging during trip

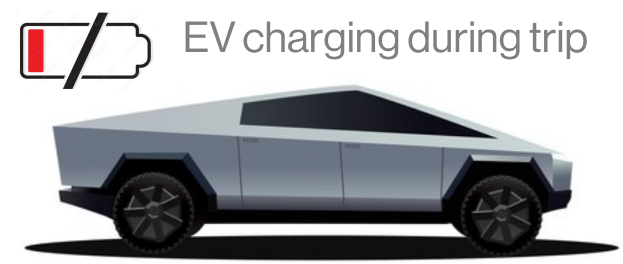

Question: Which decision caused
the bad outcome at t=5?

Today's action shapes tomorrow's options!

One-shot decision	Learn an action; easy to evaluate outcome
Sequential decision	**Learn a policy;** hard to learn on the job

Sequential Decision Making

t=1 t=2 t=3 t=4 t=5

[a_1] [a_2] [a

r=

Question: Which decision caused

Should AI tell people exactly what to do ("precise") or teach them how to think ("broad")?

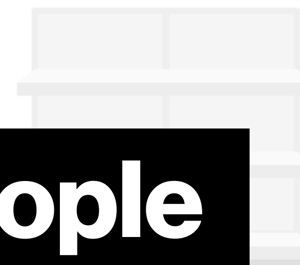

inventory management

ing trip

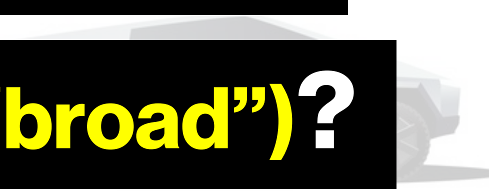

One-shot decision — Learn an action; easy to evaluate outcome

Sequential decision — Learn a policy; hard to learn on the job

Advice Precision in Operations

exact action **Precise**

Broad broader principle containing many possible actions

Both formats are derived from the same MDP-optimal policy.
Broad is not explanation. The precise action lies within the broad action set.

Advice Precision in Operations

exact action **Precise**

Broad broader principle containing many possible actions

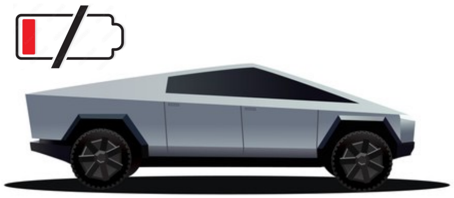

(today's setting)

"Charge to 65%"

"Charge enough for the next 2 highway exits"

Both formats are derived from the same MDP-optimal policy.
Broad is not explanation. The precise action lies within the broad action set.

Advice Precision in Operations

	exact action **Precise**	**Broad** broader principle containing many possible actions
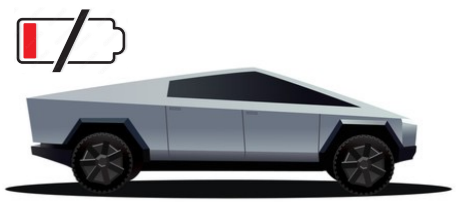 (today's setting)	"Charge to 65%"	"Charge enough for the next 2 highway exits"
Uber Bolt	"Go to Zone 4 now"	"Stay close to the high-demand corridor"
ZARA Alibaba.com	"Order 50 units"	"As underage cost rises, raise your target quantity"

Both formats are derived from the same MDP-optimal policy.
Broad is not explanation. The precise action lies within the broad action set.

Model Intuition

E1: advice available

- Designer chooses advice type **(precise / broad)**
- Worker chooses costly effort.

Model Intuition

E1: advice available

- Designer chooses advice type **(precise / broad)**
- Worker chooses costly effort.

↓ Environment changes by δ

E2: advice not available

- Worker chooses costly effort.

Model Intuition

E1: advice available

- Designer chooses advice type **(precise / broad)**
- Worker chooses costly effort.

Environment changes by δ

E2: advice not available

- Worker chooses costly effort.
- Two channels:

direct implementability
"I remember what to do."

higher-order understanding
"I learned the structure."

Model Intuition

E1: advice available

- Designer chooses advice type **(precise / broad)**
- Worker chooses costly effort.

Environment changes by δ

E2: advice not available

- Worker chooses costly effort.
- Two channels:

direct implementability
“I remember what to do.”

higher-order understanding
“I learned the structure.”

weight placed on immediate performance

same env.

very different env.

how different the new environment is from training

Model Intuition

E1: advice available

- Designer chooses advice type **(precise / broad)**
- Worker chooses costly effort.

Environment changes by δ

E2: advice not available

- Worker chooses costly effort.
- Two channels:

direct implementability
"I remember what to do."

higher-order understanding
"I learned the structure."

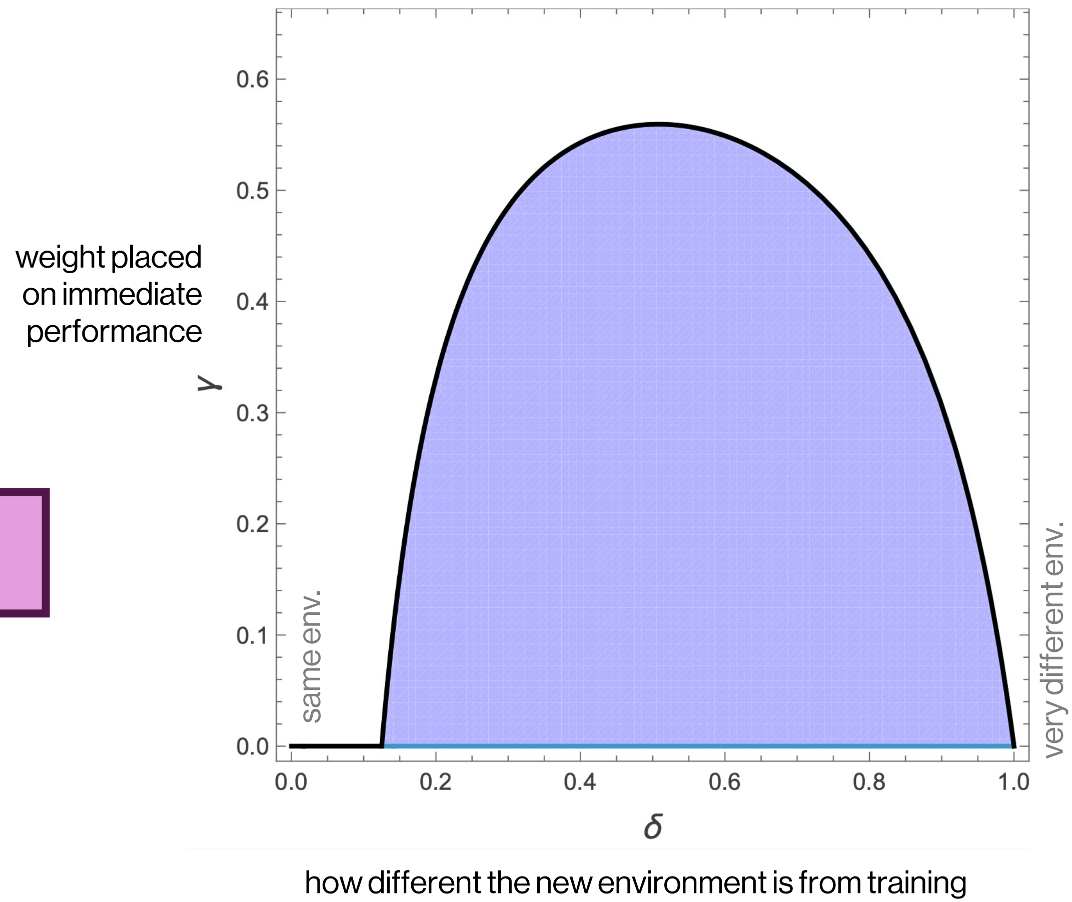

Model Intuition

E1: advice available

- Designer chooses advice type **(precise / broad)**
- Worker chooses costly effort.

Environment changes by δ

E2: advice not available

- Worker chooses costly effort.
- Two channels:

direct implementability
"I remember what to do."

higher-order understanding
"I learned the structure."

Reward gap:
easier to follow,
better E1 performance

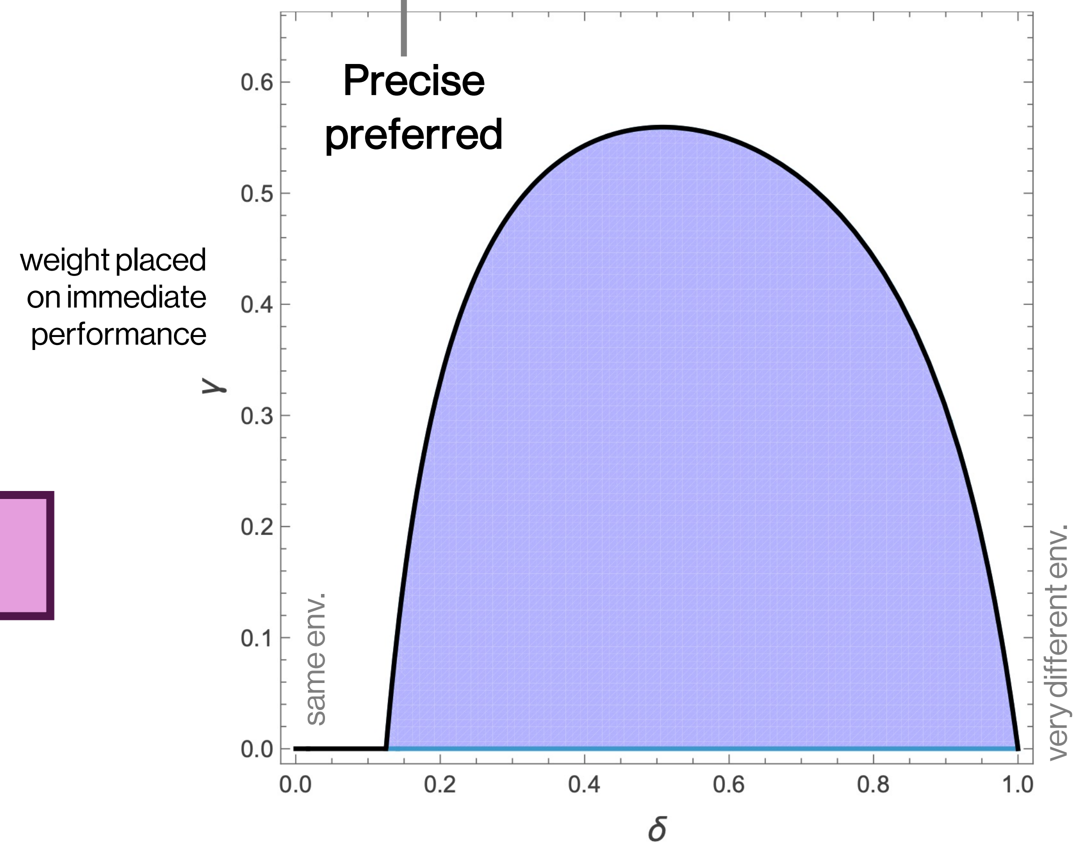

how different the new environment is from training

Model Intuition

E1: advice available

- Designer chooses advice type **(precise / broad)**
- Worker chooses costly effort.

Environment changes by δ

E2: advice not available

- Worker chooses costly effort.
- Two channels:

direct implementability
"I remember what to do."

higher-order understanding
"I learned the structure."

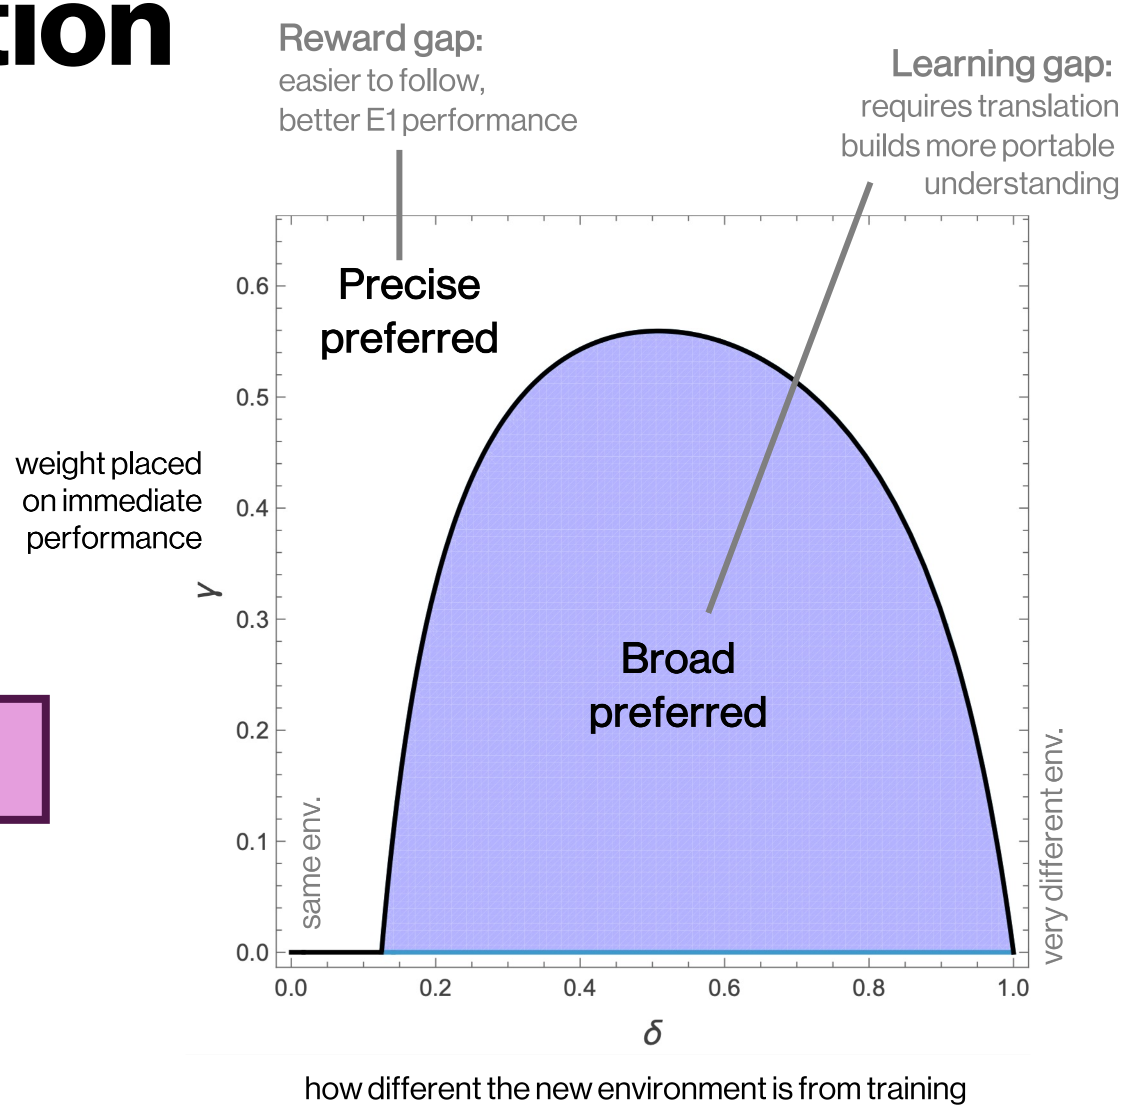

Experimental Setting: EV Charging Game

Experimental Setting: EV Charging Game

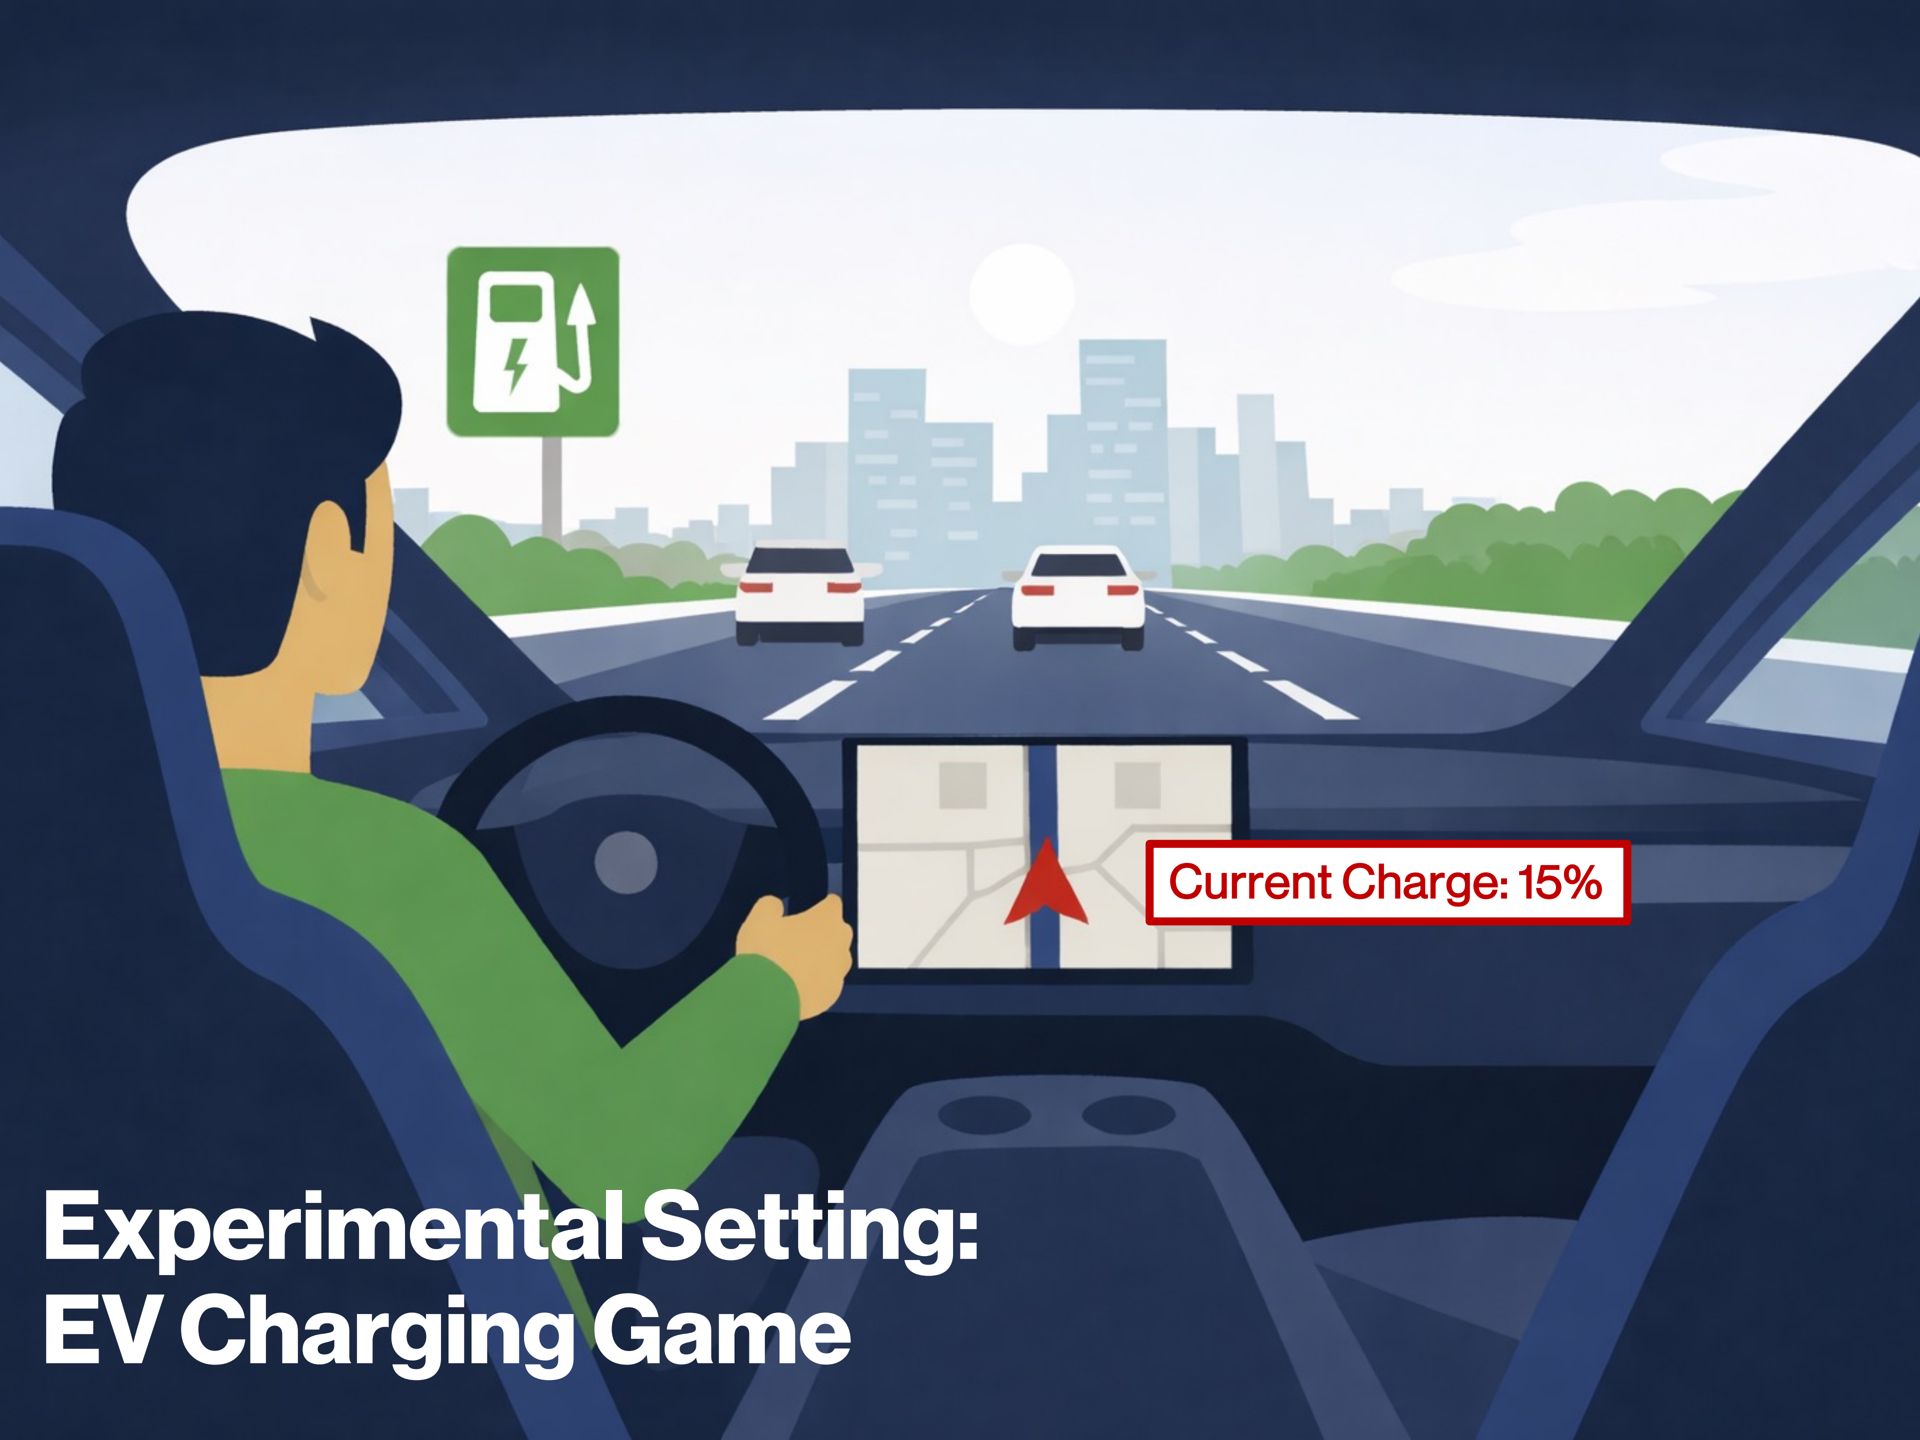

Experimental Setting: EV Charging Game

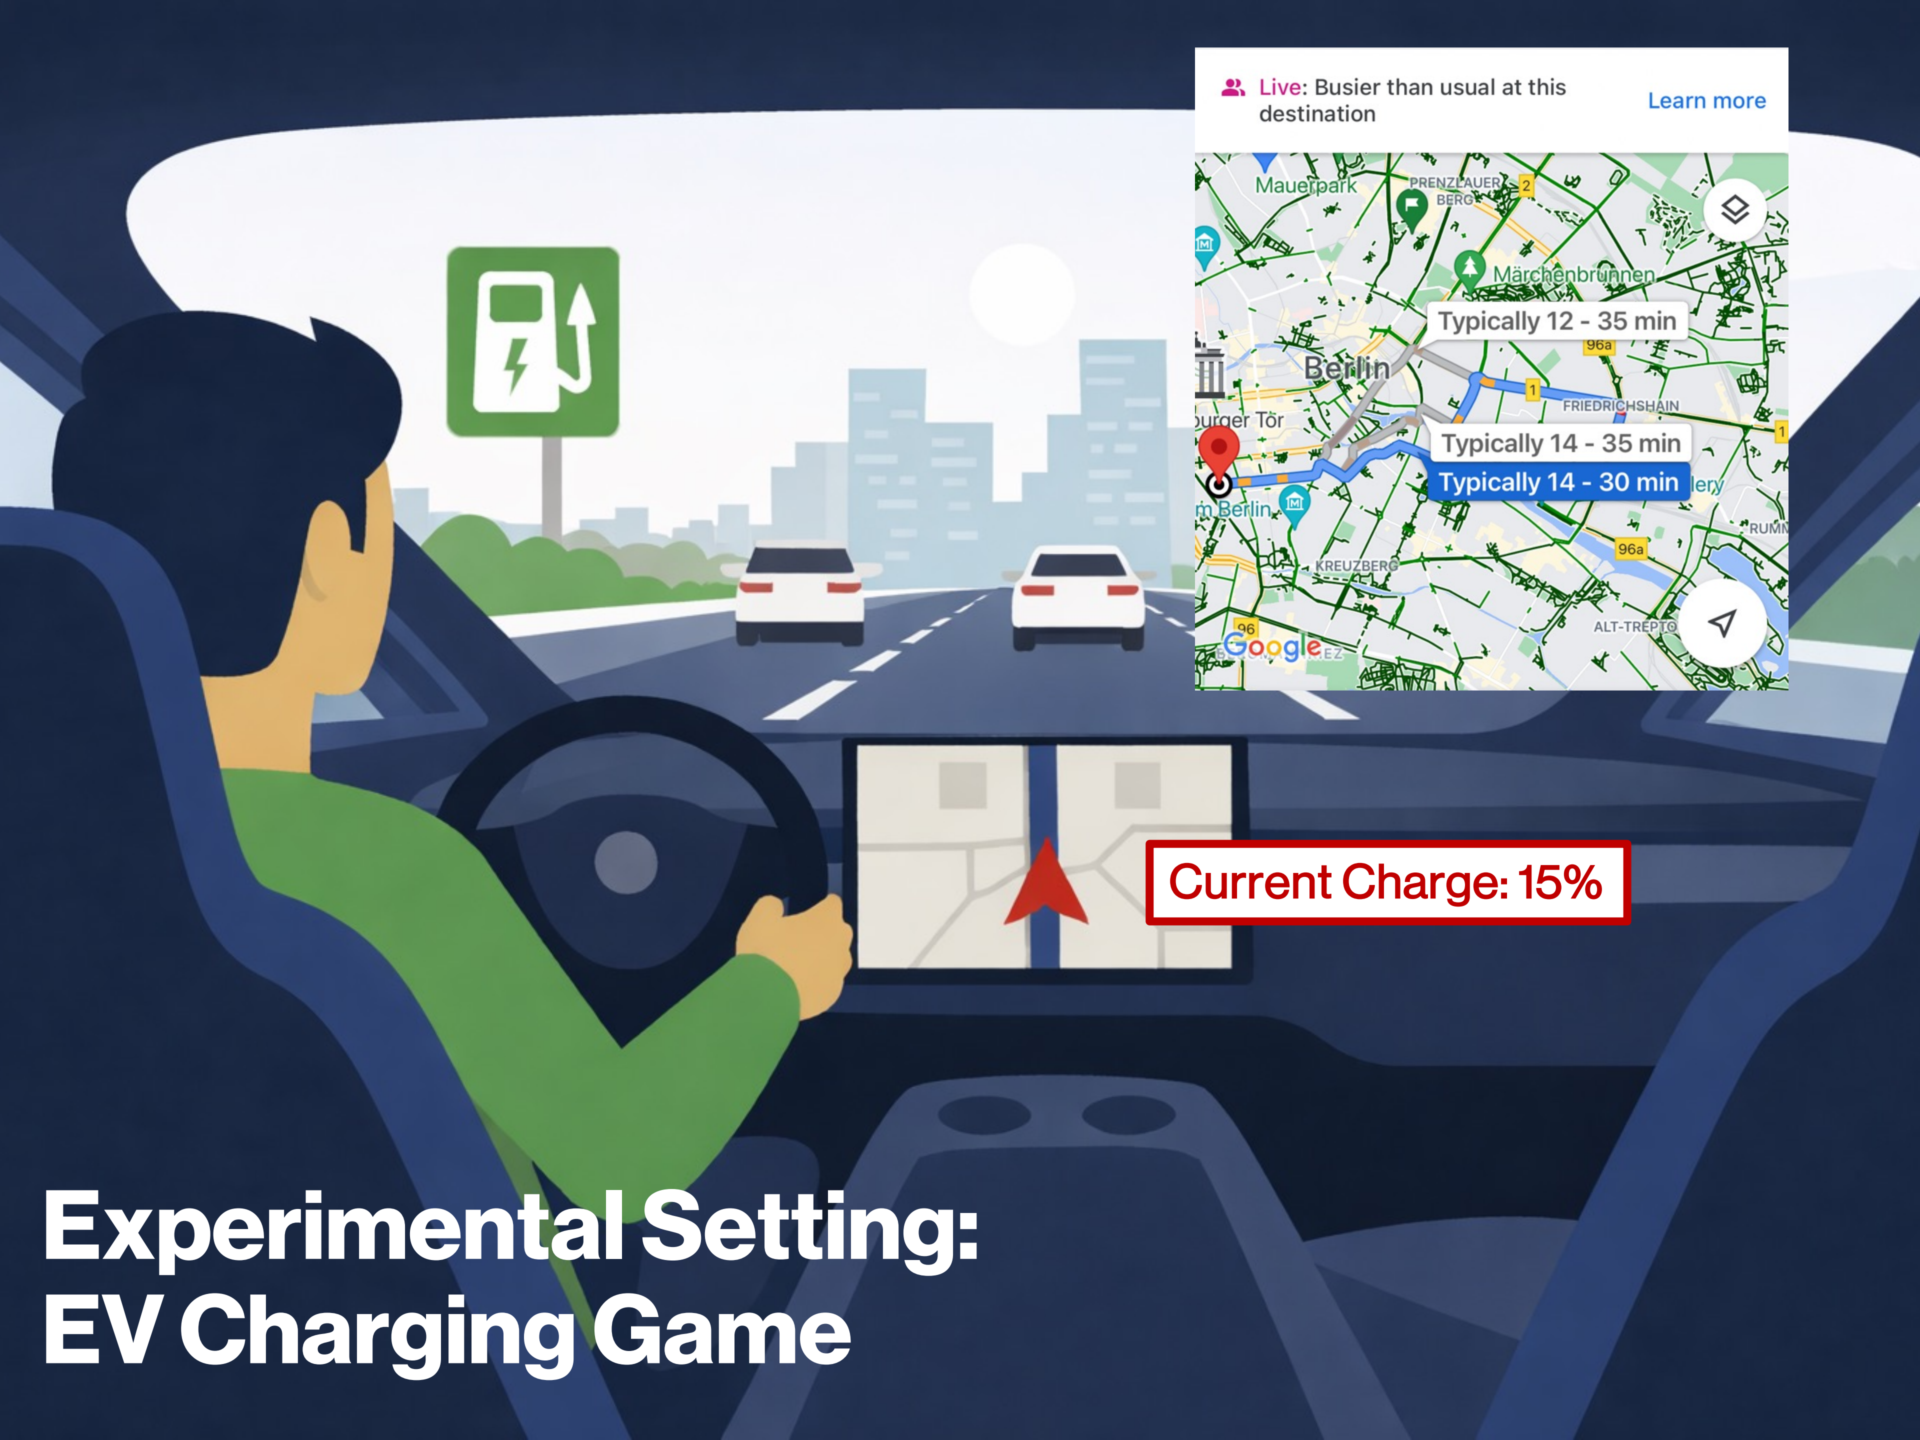

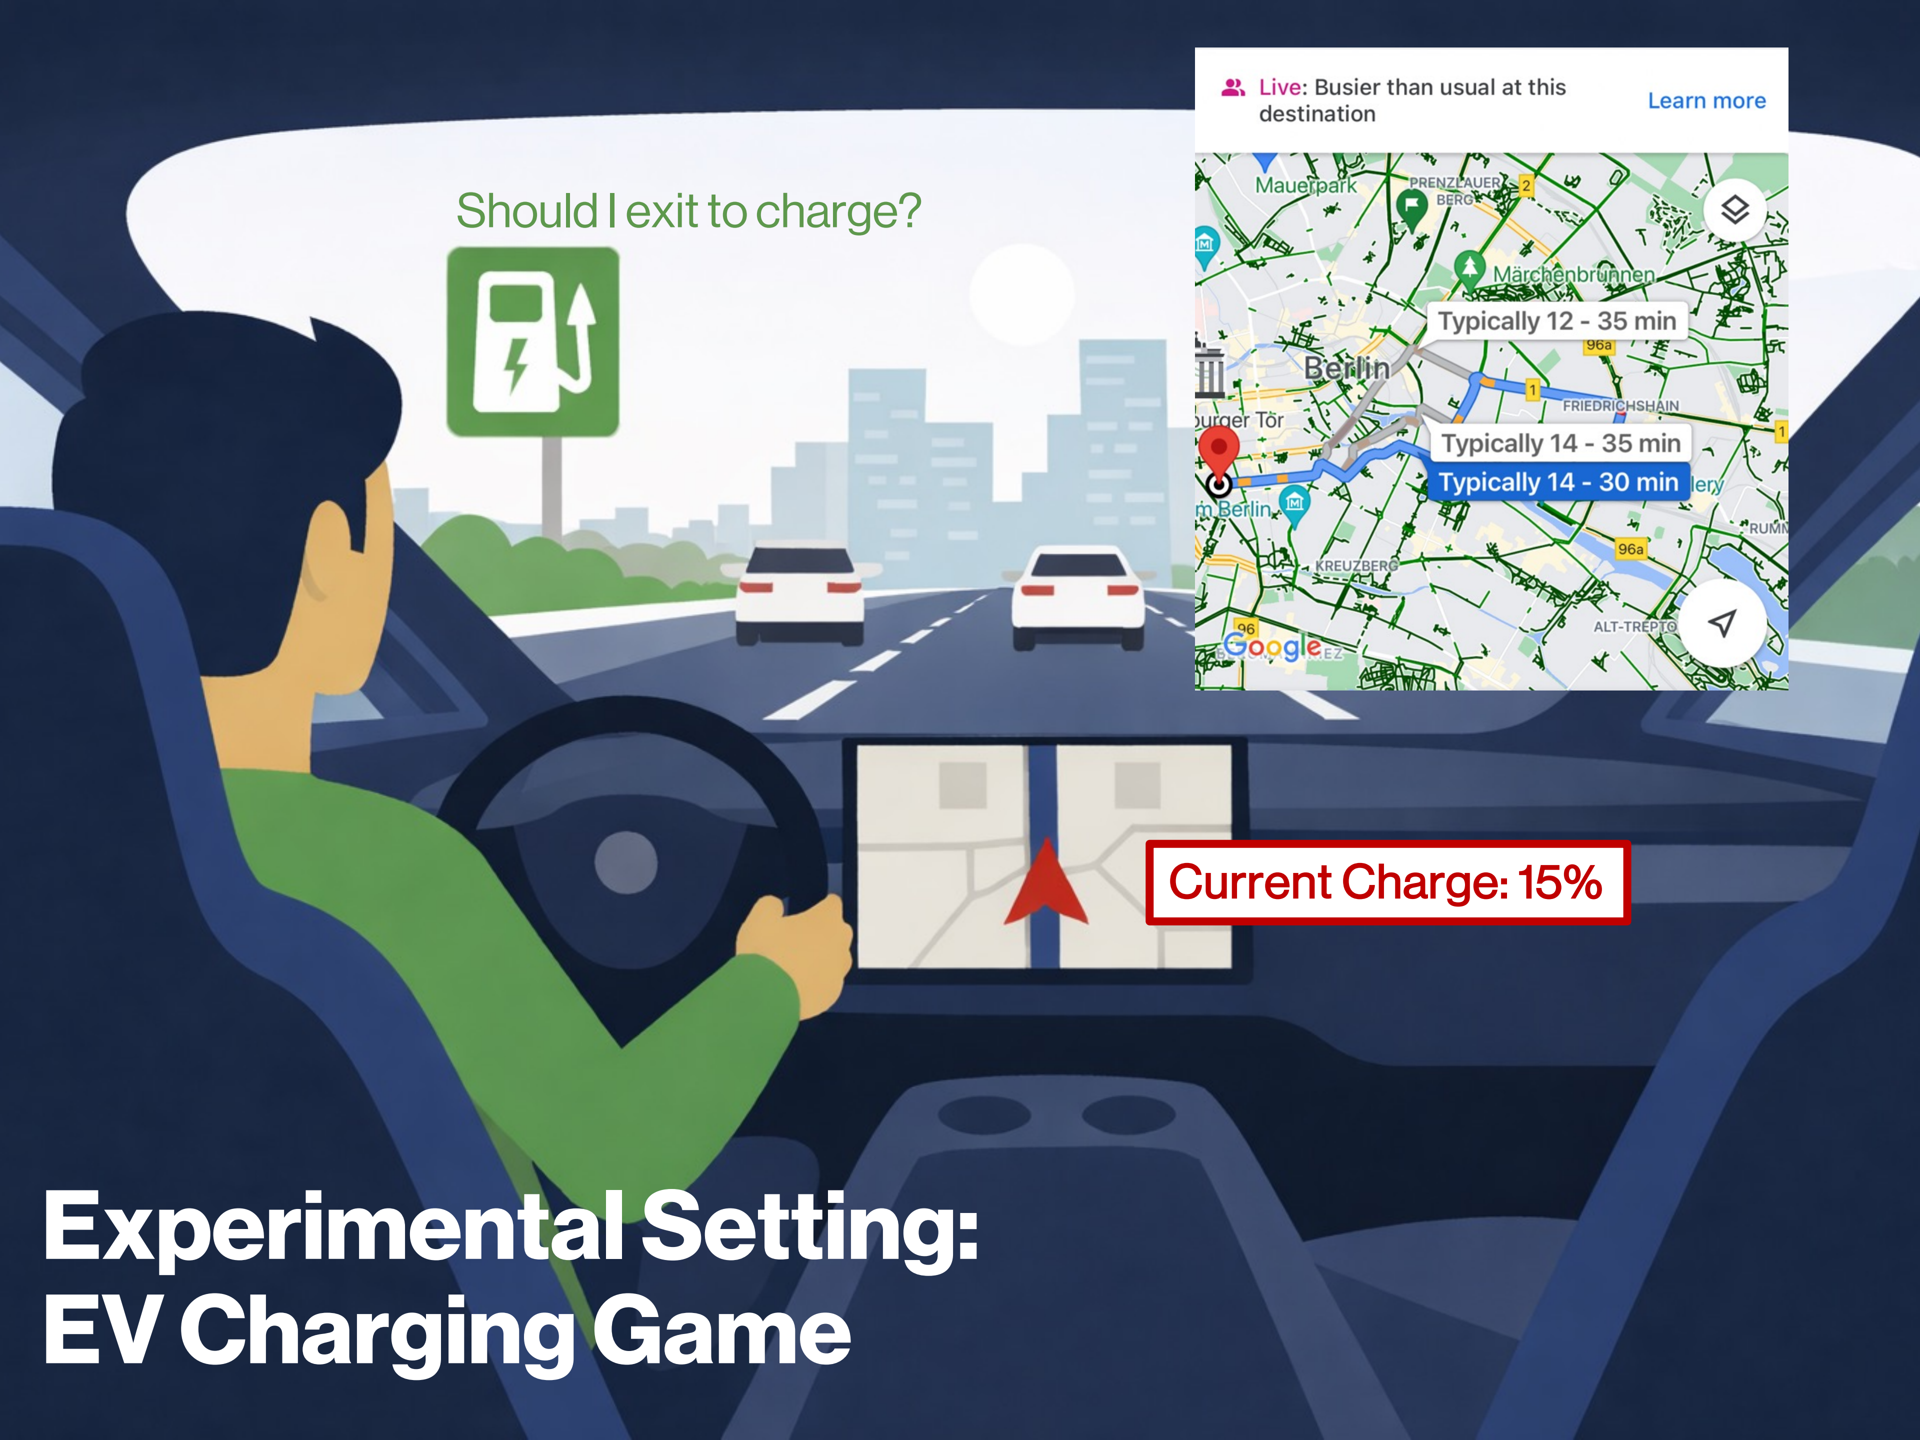

Experimental Setting: EV Charging Game

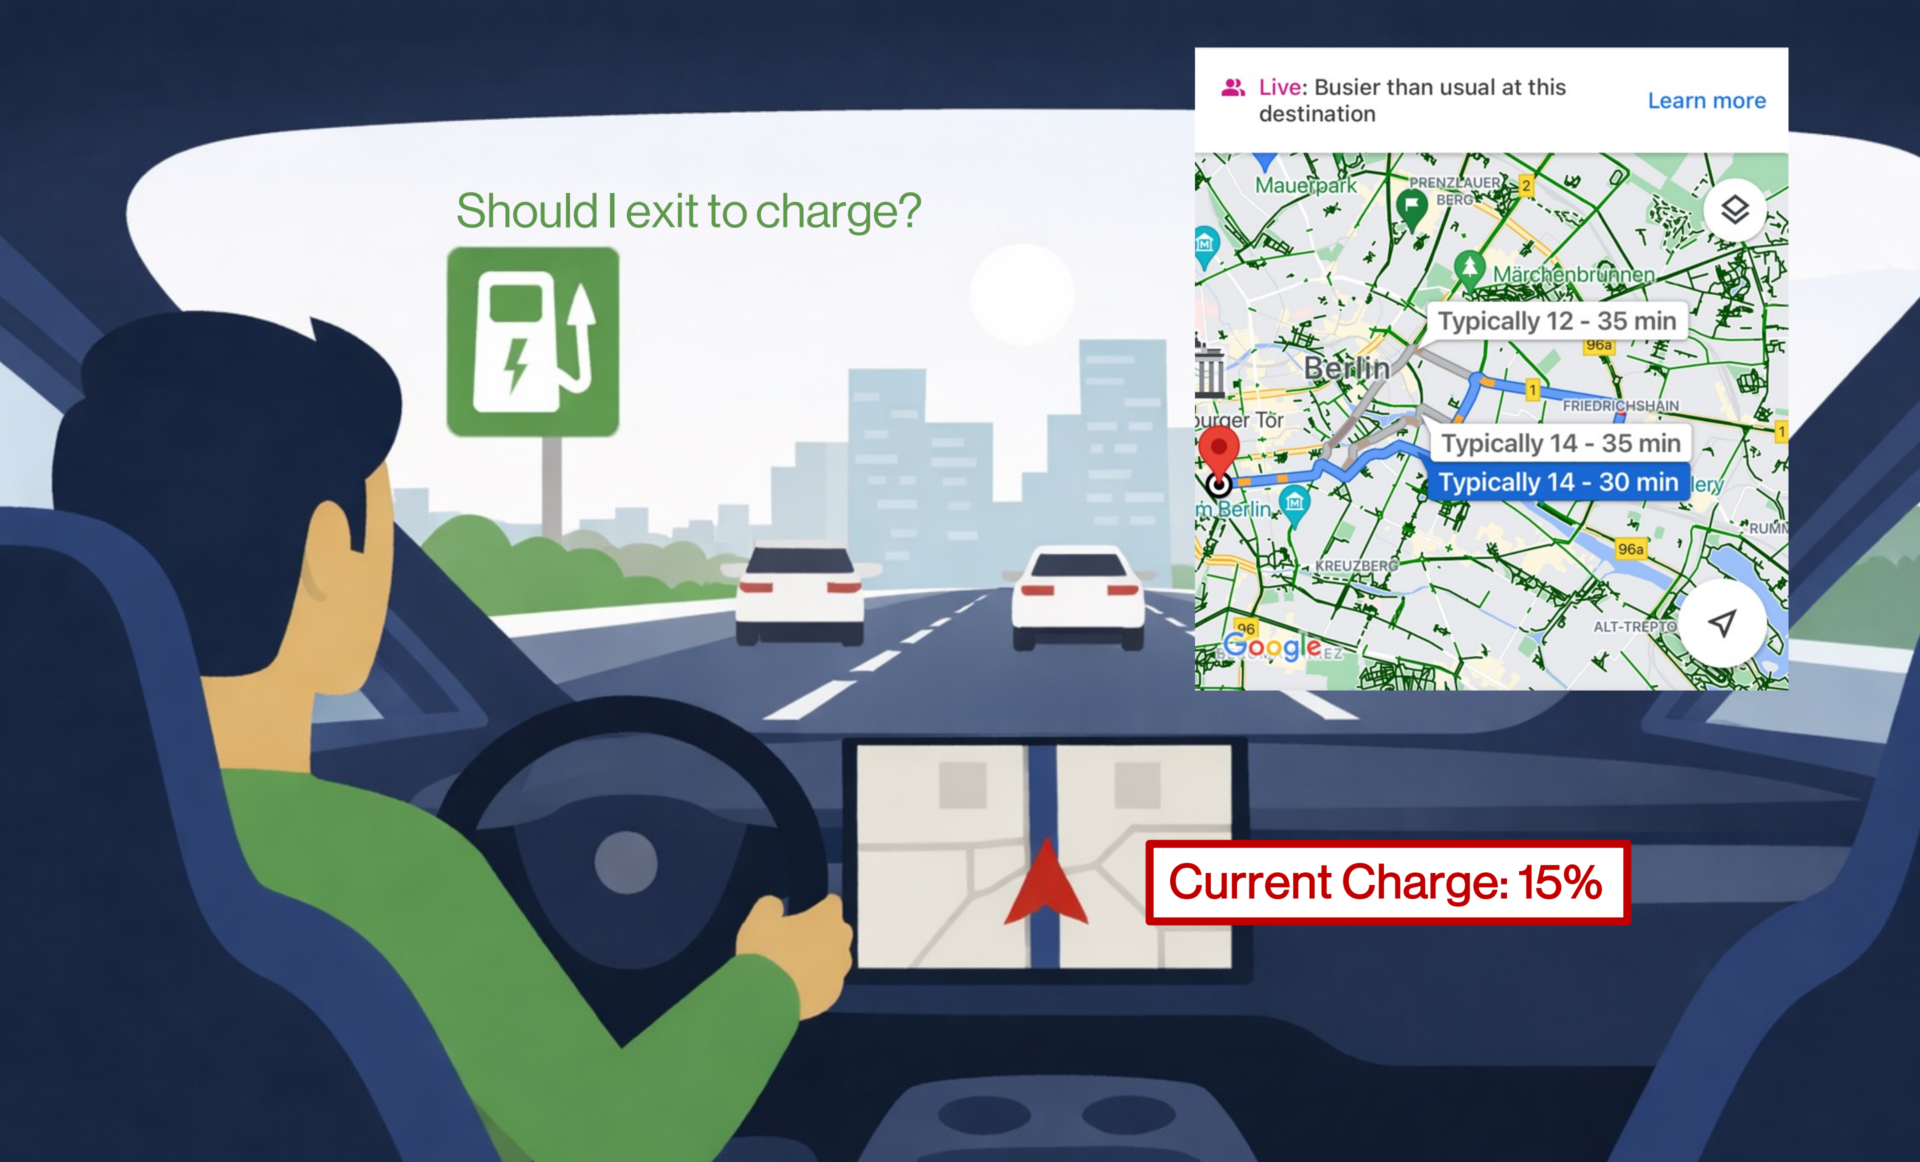

Experimental Setting: EV Charging Game

This is not a paper about EVs/EV adoption!
It's a clean sequential task that is solvable, non-obvious , and concrete.

Study Task/Interface

Goal:
Minimize the time
to get to the destination

Round 1: Segment 0 → 1

🔋 Battery: 0% ⏱ Elapsed Time: 0 min

Distance: 5 min | Traffic: [+10, +15]

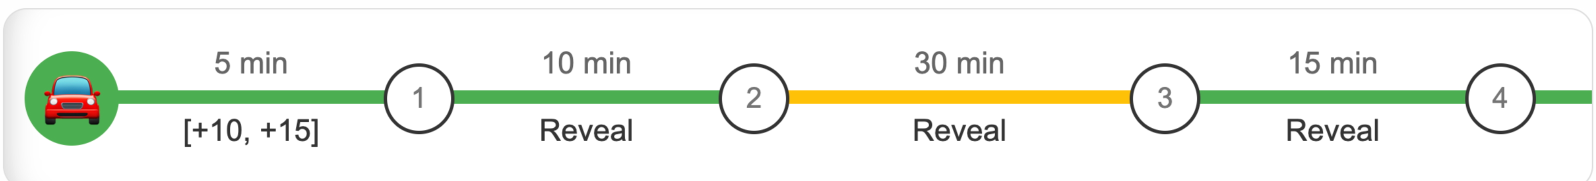

Charge | Proceed without Charging

Study Task/Interface

Goal:
Minimize the time
to get to the destination

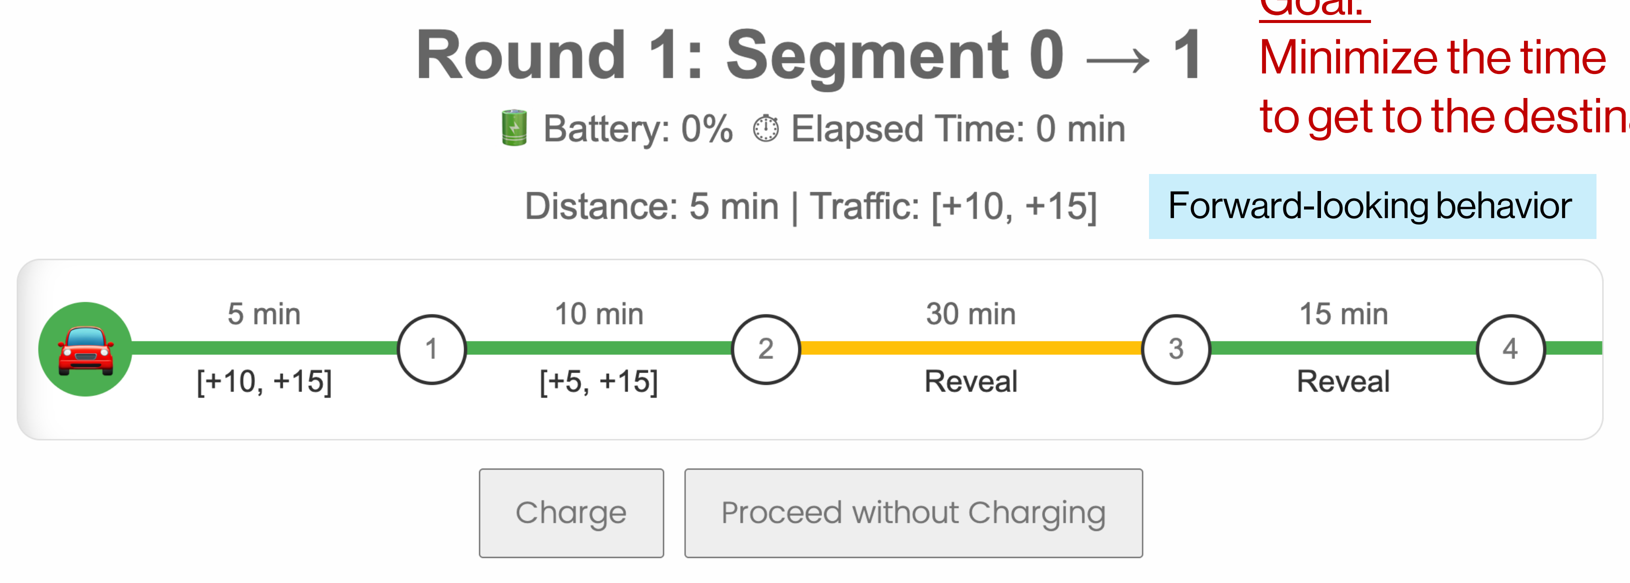

Forward-looking behavior

pre-registered at https://aspredicted.org/7LP_J9G

Study Task/Interface

Goal:
Minimize the time
to get to the destination

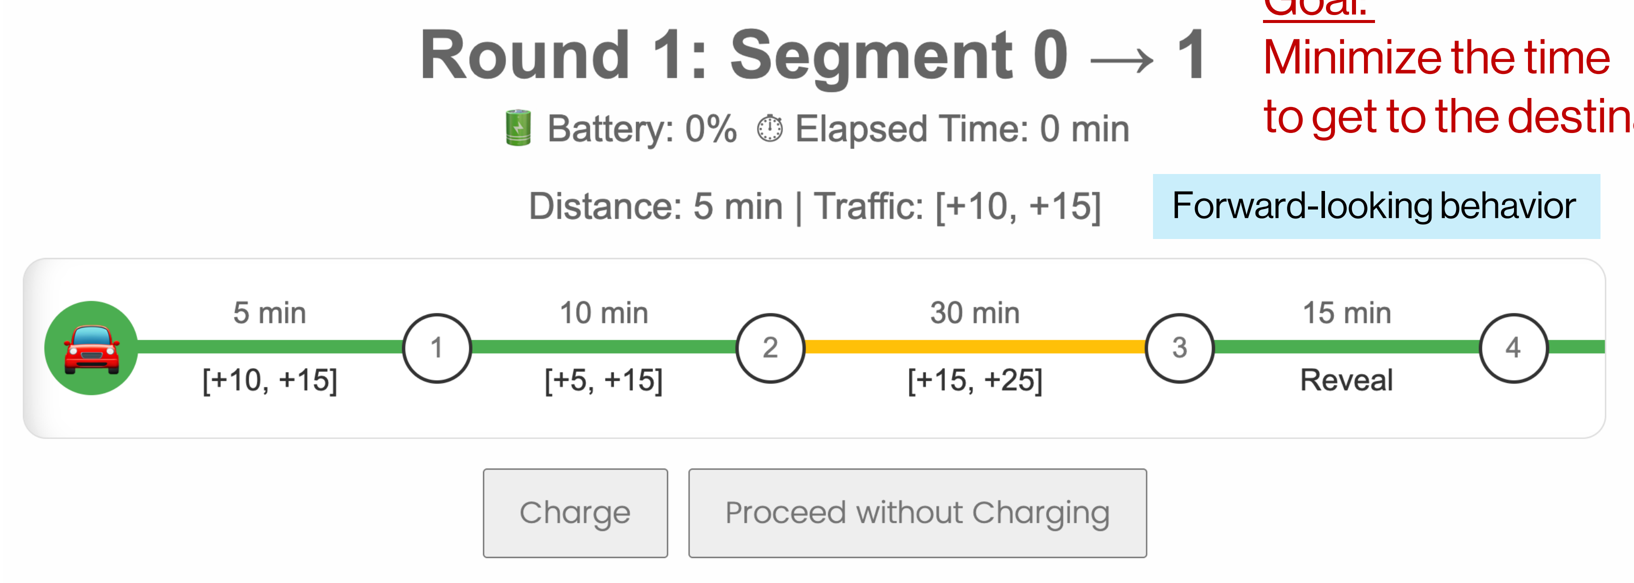

Forward-looking behavior

Study Task/Interface

Goal:
Minimize the time
to get to the destination

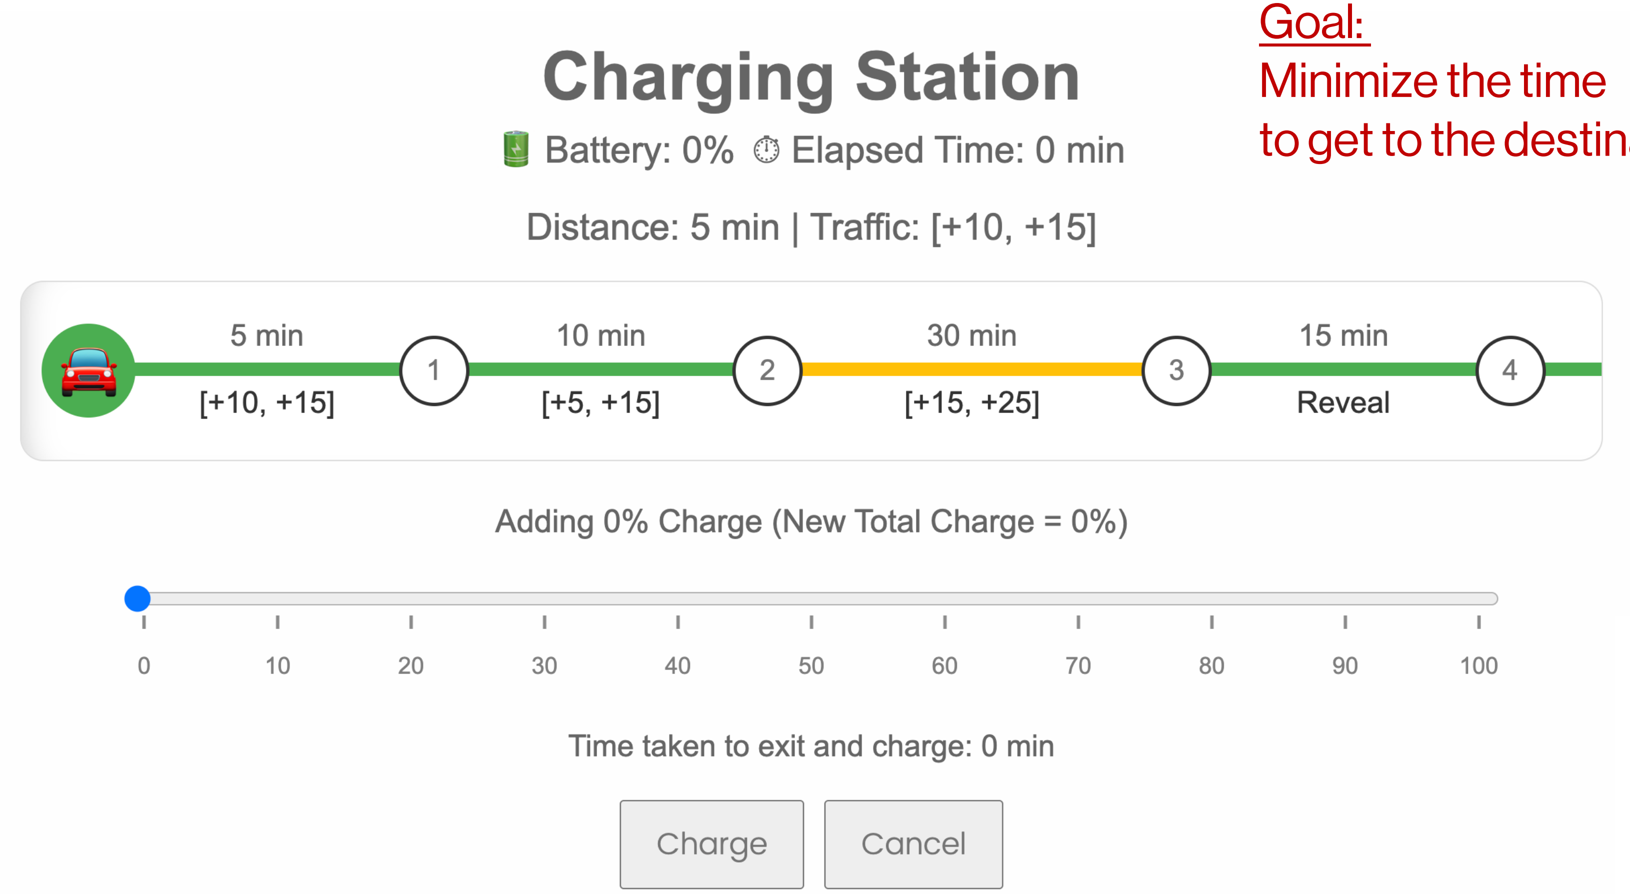

pre-registered at https://aspredicted.org/7LP_J9G

Study Task/Interface

Goal:
Minimize the time
to get to the destination

Charging Station

🔋 Battery: 0% ⏱ Elapsed Time: 0 min

Distance: 5 min | Traffic: [+10, +15]

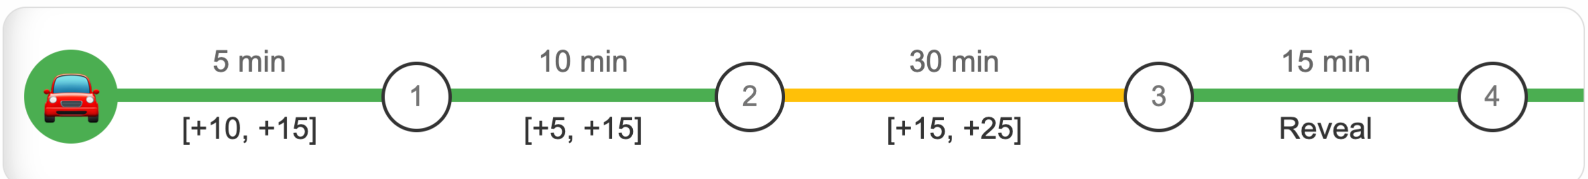

Adding 20% Charge (New Total Charge = 20%)

0 10 20 30 40 50 60 70 80 90 100

Time taken to exit and charge: 51 min

Exploration of strategy

Charge Cancel

pre-registered at https://aspredicted.org/7LP_J9G

Study Task/Interface

Goal:
Minimize the time
to get to the destination

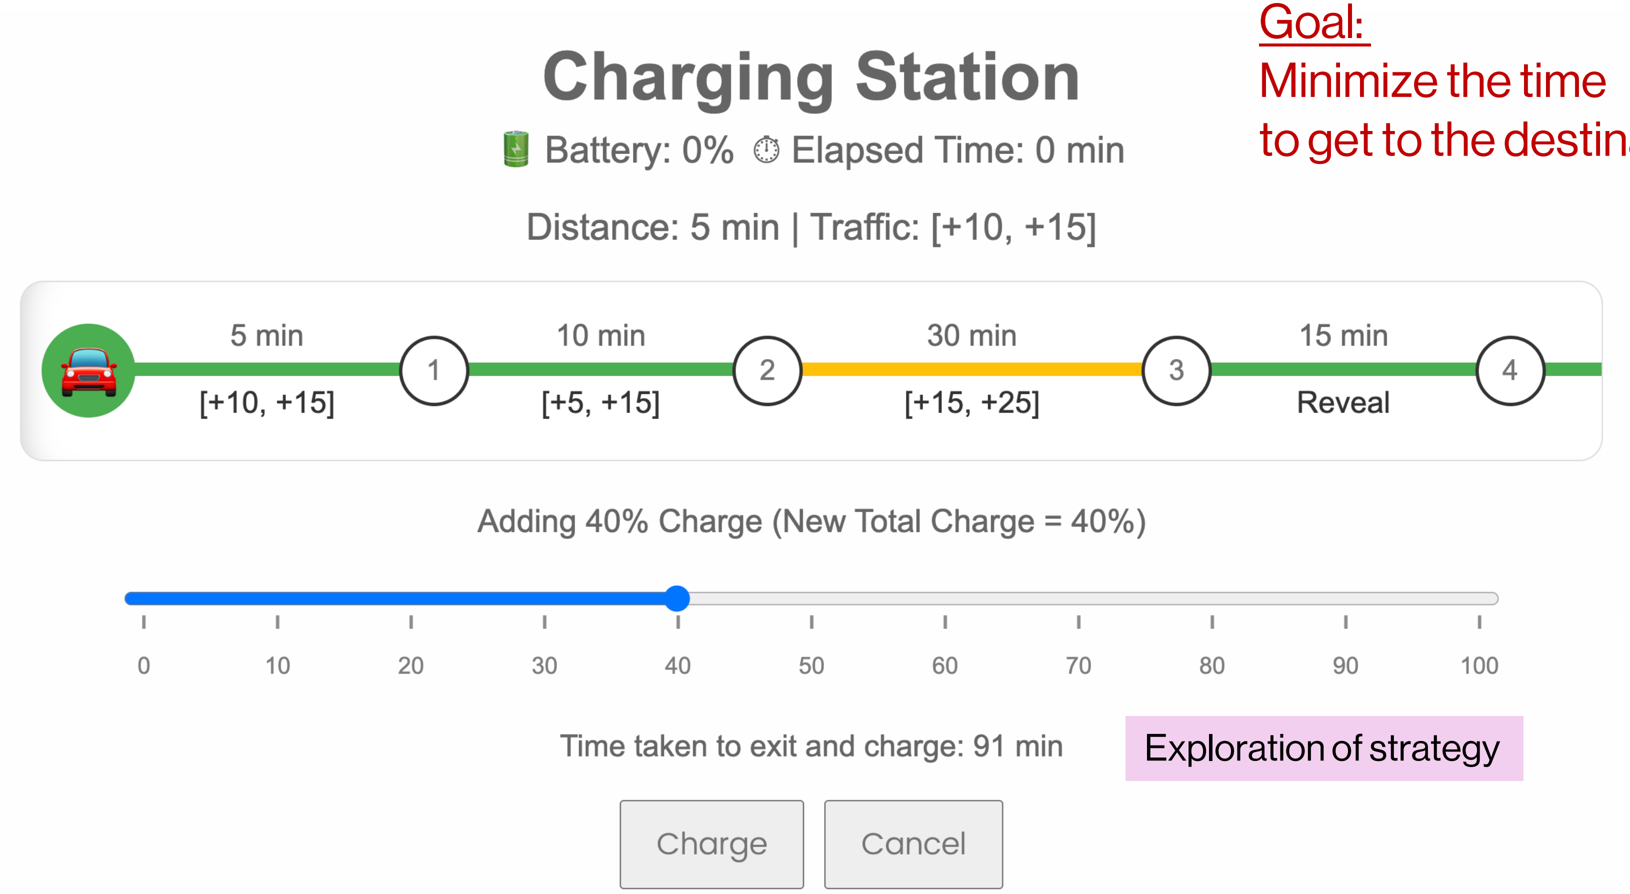

Exploration of strategy

pre-registered at https://aspredicted.org/7LP_J9G

Study Task/Interface

Goal:
Minimize the time
to get to the destination

Charging Station

Battery: 0% ⏱ Elapsed Time: 0 min

Distance: 5 min | Traffic: [+10, +15]

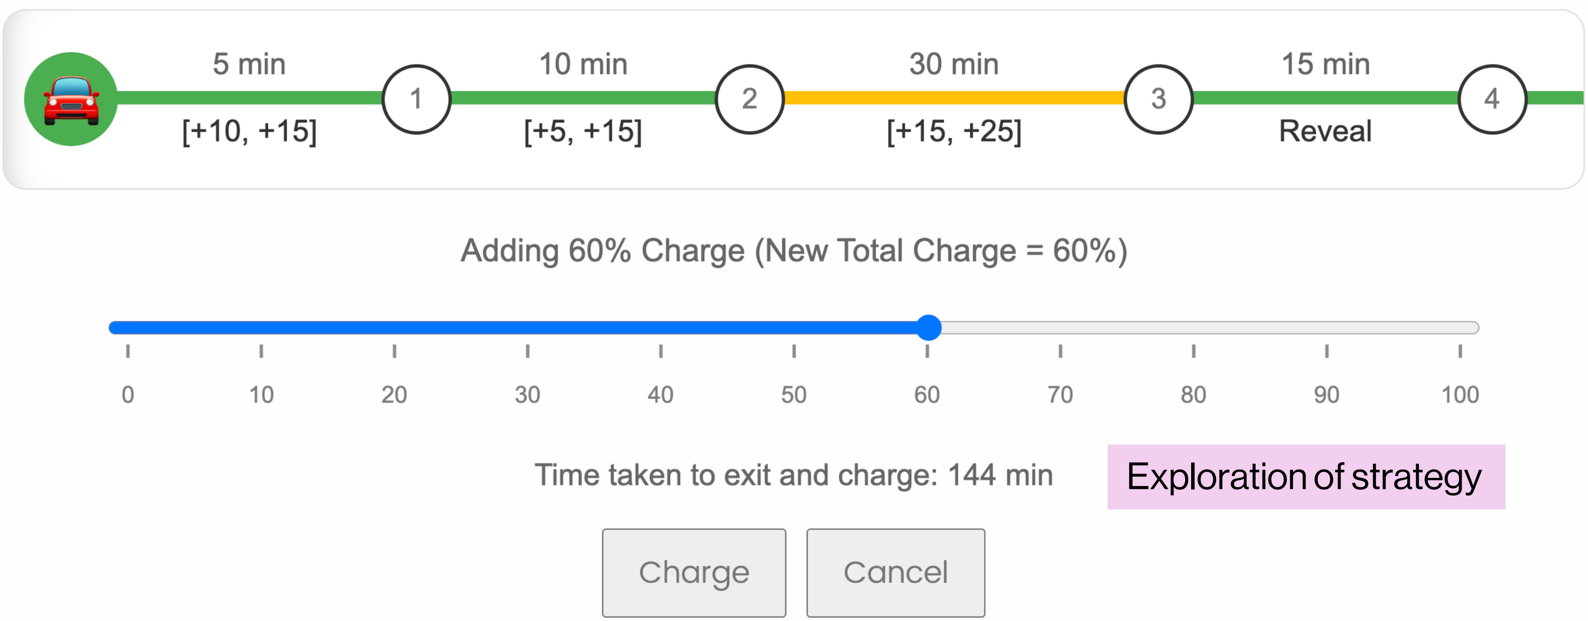

pre-registered at https://aspredicted.org/7LP_J9G

Study Task/Interface

Goal:
Minimize the time
to get to the destination

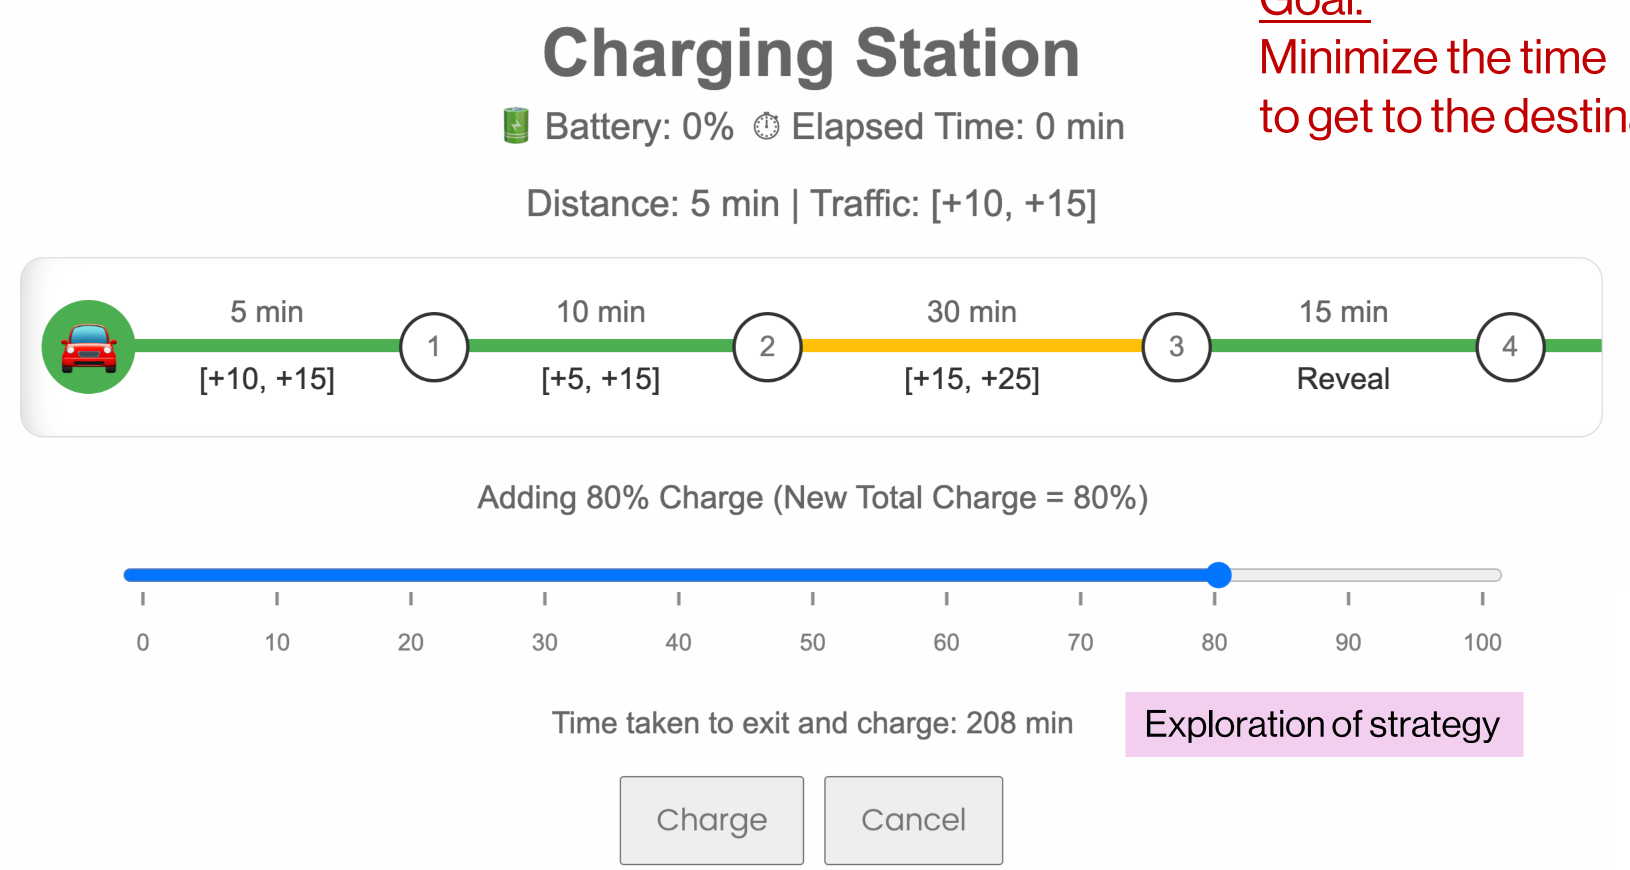

Exploration of strategy

pre-registered at https://aspredicted.org/7LP_J9G

Study Task/Interface

Goal:
Minimize the time
to get to the destination

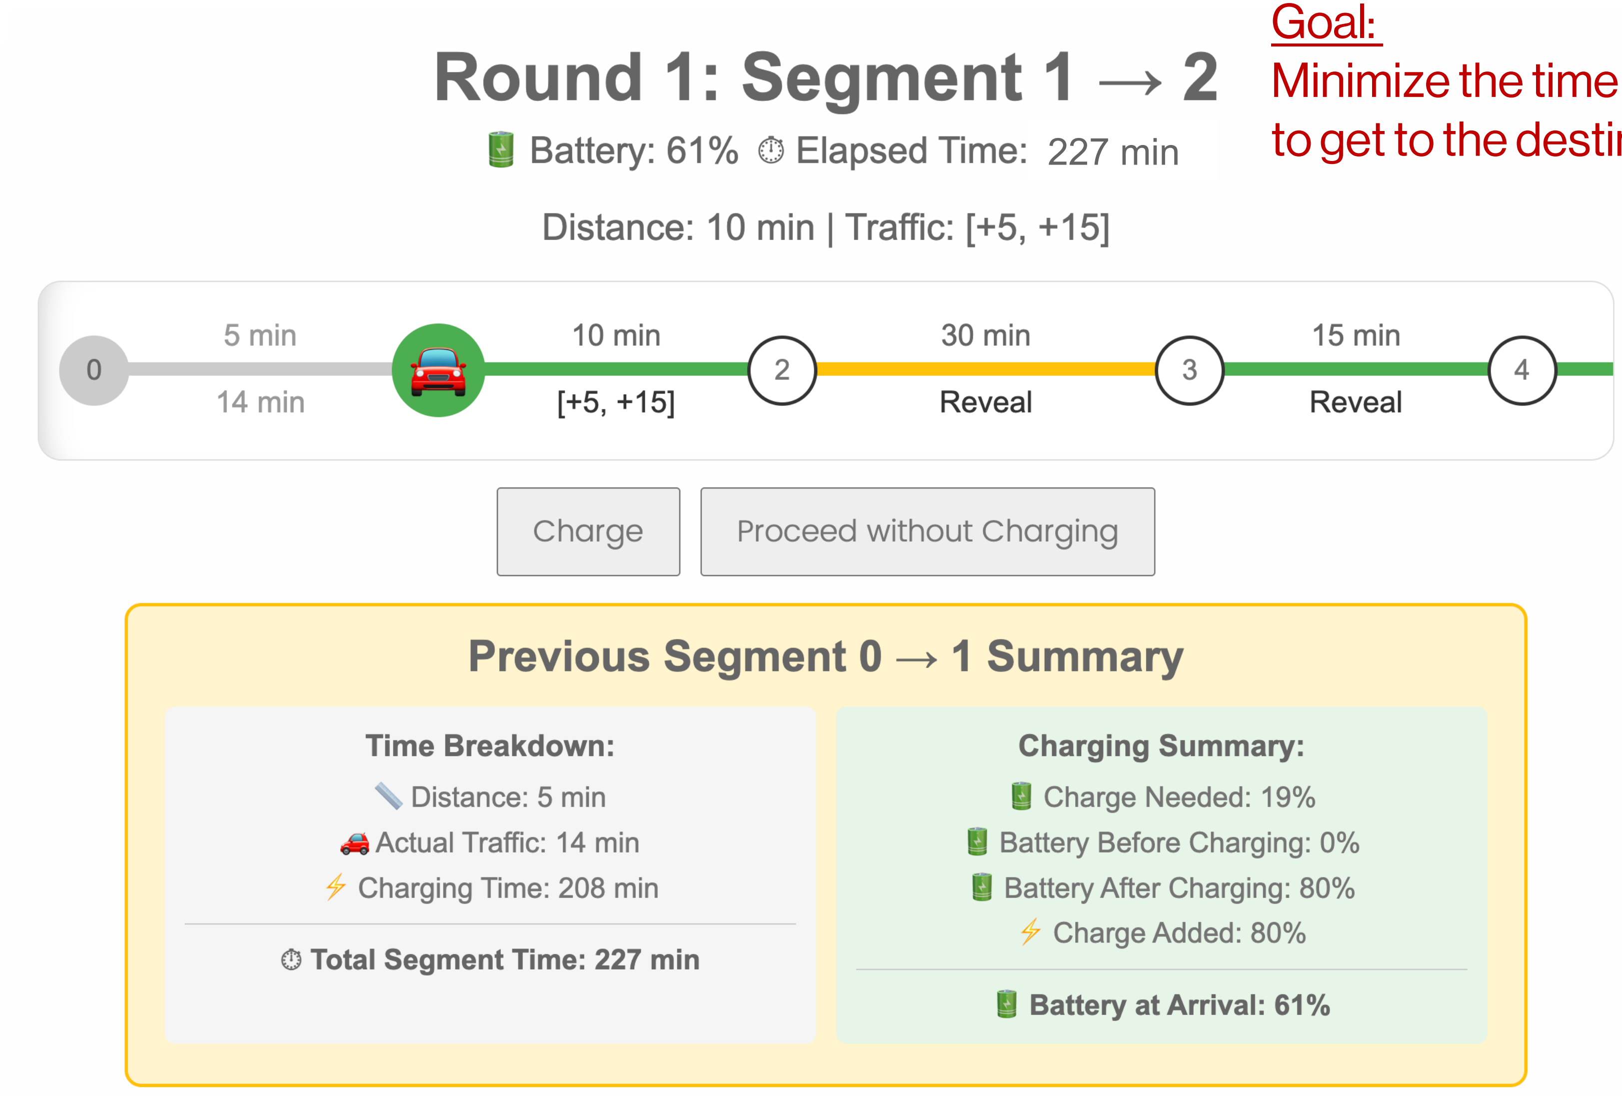

Study Task/Interface

Goal:
Minimize the time
to get to the destination

Round 1: Segment 1 → 2

🔋 Battery: 61% ⏱ Elapsed Time: 227 min

Distance: 10 min | Traffic: [+5, +15]

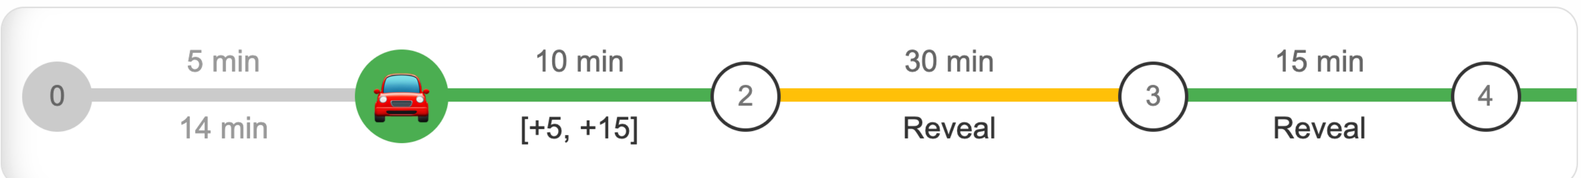

Charge | Proceed without Charging

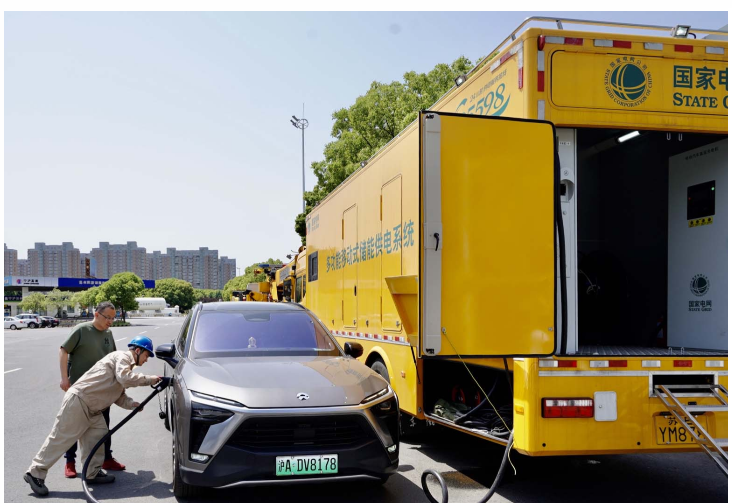

If you run out, large penalty of +300 minutes

pre-registered at https://aspredicted.org/7LP_J9G

Precise vs Broad Advice

	Precise (specific numeric action)	Broad (coverage logic / local principle)
Optimal action:		

Precise vs Broad Advice

	Precise (specific numeric action)	Broad (coverage logic / local principle)
Optimal action:		
Batching Charge for multiple stops		💡 **Tip:** Charge enough for this segment and the next one

pre-registered at https://aspredicted.org/7LP_J9G

Precise vs Broad Advice

	Precise (specific numeric action)	Broad (coverage logic / local principle)
Optimal action:		
Batching Charge for multiple stops		💡 **Tip:** Charge enough for this segment and the next one
Splitting Charge for just one stop		💡 **Tip:** Charge enough for this segment

pre-registered at https://aspredicted.org/7LP_J9G

Precise vs Broad Advice

	Precise (specific numeric action)	Broad (coverage logic / local principle)
Optimal action:		
Batching Charge for multiple stops	💡 **Tip:** Charge X%	💡 **Tip:** Charge enough for this segment and the next one
Splitting Charge for just one stop		💡 **Tip:** Charge enough for this segment

One tip format covers all cases.

pre-registered at https://aspredicted.org/7LP_J9G

Precise vs Broad Advice

	Precise (specific numeric action)	**Broad** (coverage logic / local principle)
Optimal action:		
Batching Charge for multiple stops	💡 **Tip:** Charge X%	💡 **Tip:** Charge enough for this segment and the next one
Splitting Charge for just one stop	One tip format covers all cases.	💡 **Tip:** Charge enough for this segment

Same optimal MDP

pre-registered at https://aspredicted.org/7LP_J9G

Study 1

Precise vs Broad Advice

	Precise (specific numeric action)	**Broad** (coverage logic / local principle)
Optimal action:		
Batching Charge for multiple stops	💡 **Tip:** Charge X%	💡 **Tip:** Charge enough for this segment and the next one
Splitting Charge for just one stop		💡 **Tip:** Charge enough for this segment

Same optimal MDP

One tip format covers all cases.

We run a **2 x 2 between-subject factorial design:** precise/broad vs traffic

pre-registered at https://aspredicted.org/7LP_J9G

Study 1

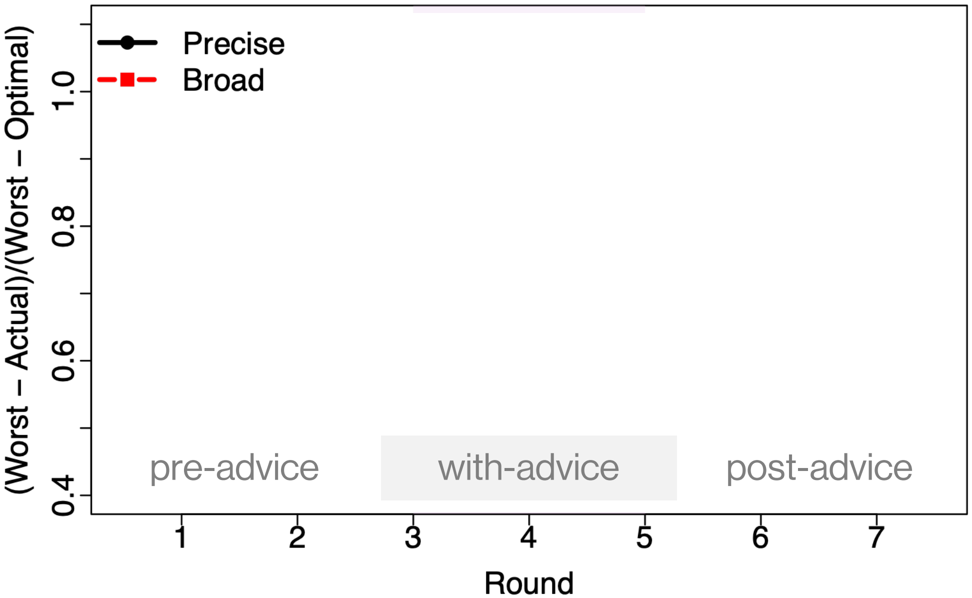

Amazon Mechanical Turk
N = 102, 3,978 decision points

Y-axis = how much improvement
from worst compared to optimal
Higher = better

pre-registered at https://aspredicted.org/7LP_J9G

Study 1

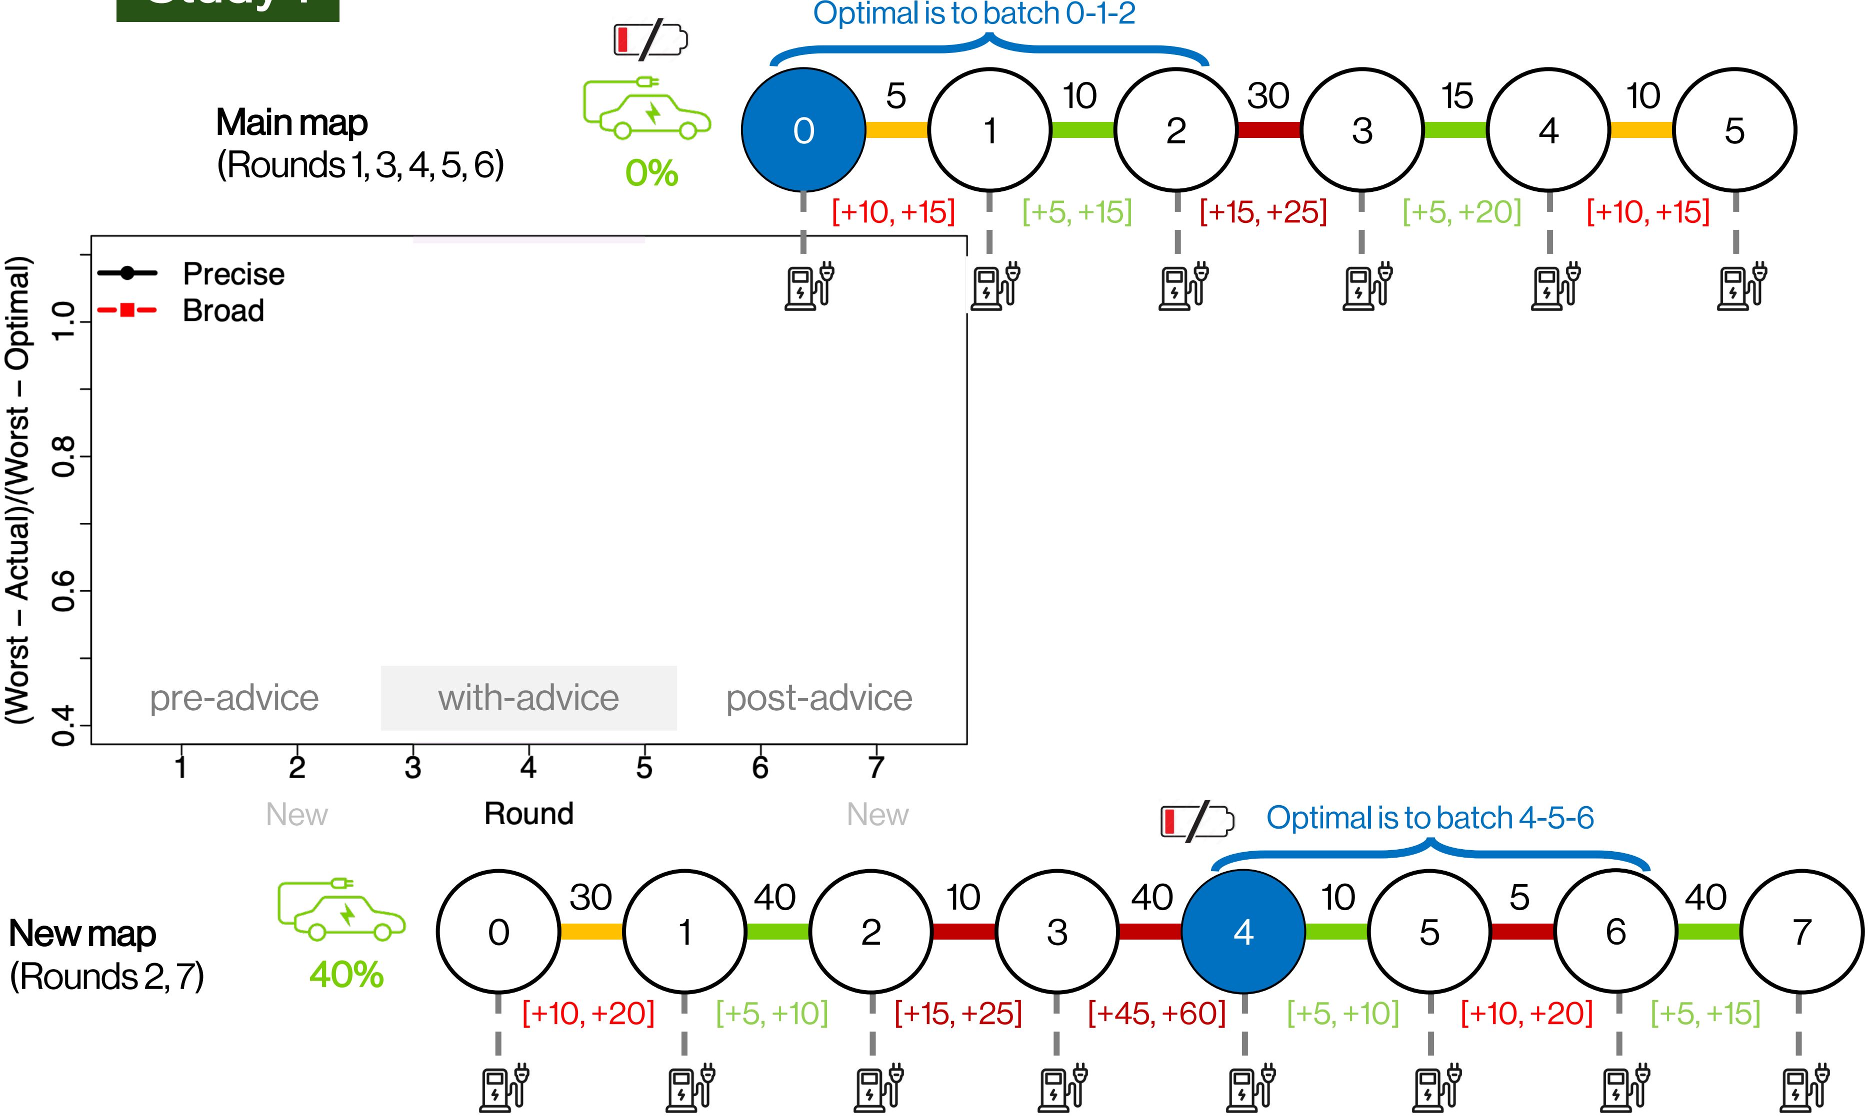

Amazon Mechanical Turk
N = 102, 3,978 decision points

Y-axis = how much improvement
from worst compared to optimal
Higher = better

pre-registered at https://aspredicted.org/7LP_J9G

Study 1

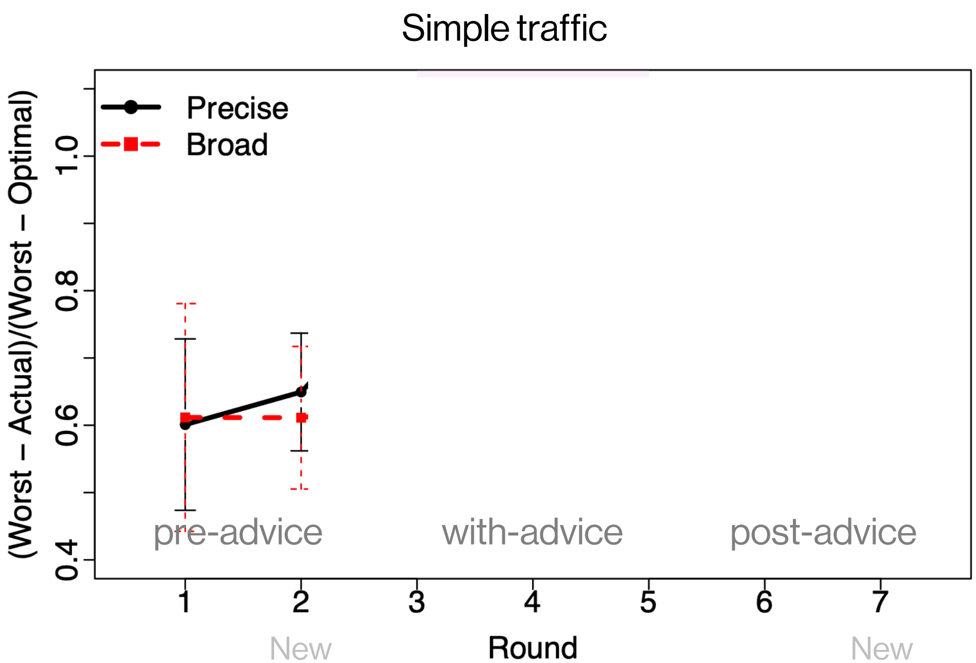

Amazon Mechanical Turk
N = 102, 3,978 decision points

Y-axis = how much improvement
from worst compared to optimal
Higher = better

pre-registered at https://aspredicted.org/7LP_J9G

Study 1

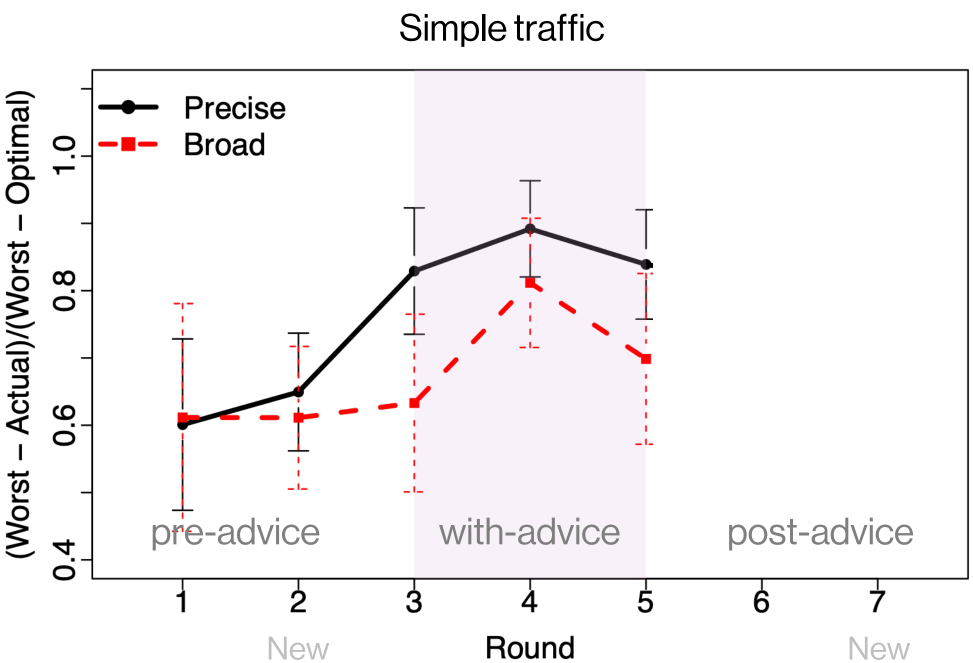

Amazon Mechanical Turk
N = 102, 3,978 decision points

Y-axis = how much improvement
from worst compared to optimal
Higher = better

pre-registered at https://aspredicted.org/7LP_J9G

Study 1

Result: Precise Works Instantly

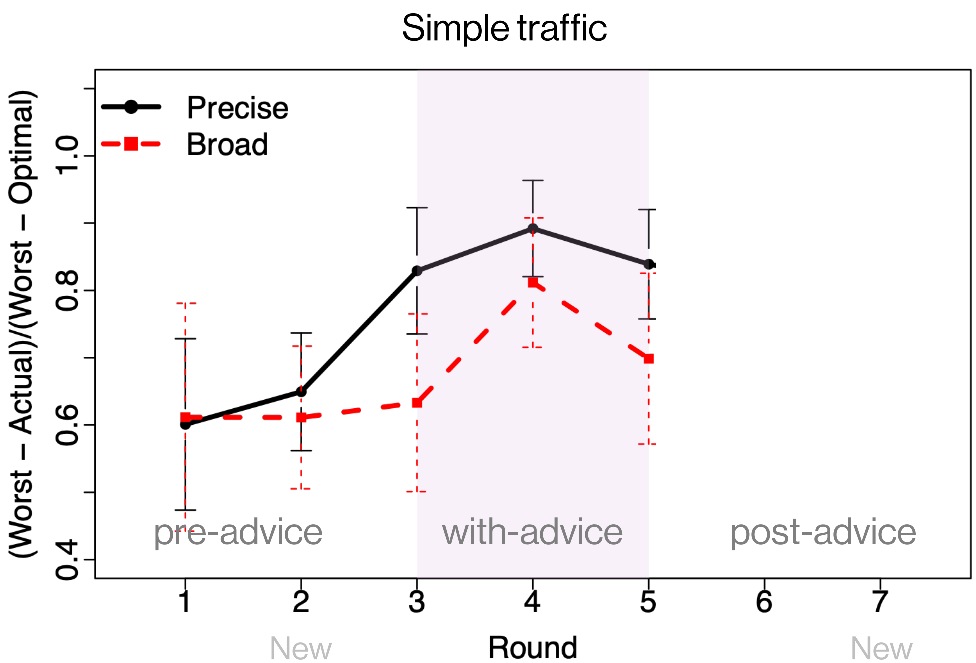

While available, precise advice improves compliance + assisted performance

Amazon Mechanical Turk
N = 102, 3,978 decision points

Y-axis = how much improvement from worst compared to optimal
Higher = better

pre-registered at https://aspredicted.org/7LP_J9G

Study 1

Result: Precise Works Instantly

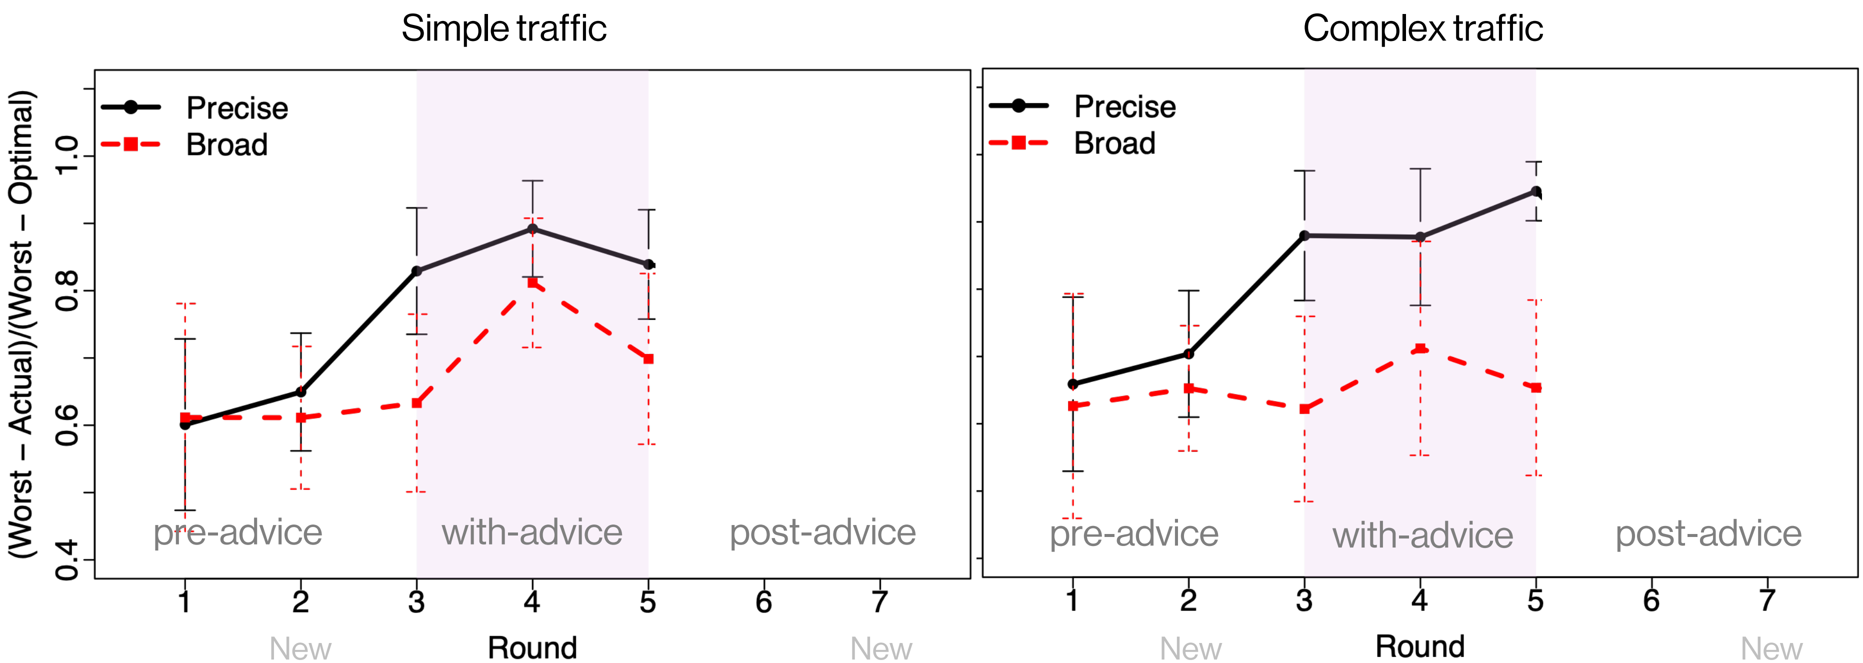

While available, precise advice improves compliance + assisted performance

Amazon Mechanical Turk
N = 102, 3,978 decision points

Y-axis = how much improvement from worst compared to optimal
Higher = better

pre-registered at https://aspredicted.org/7LP_J9G

Study 1

After Advice is Removed...

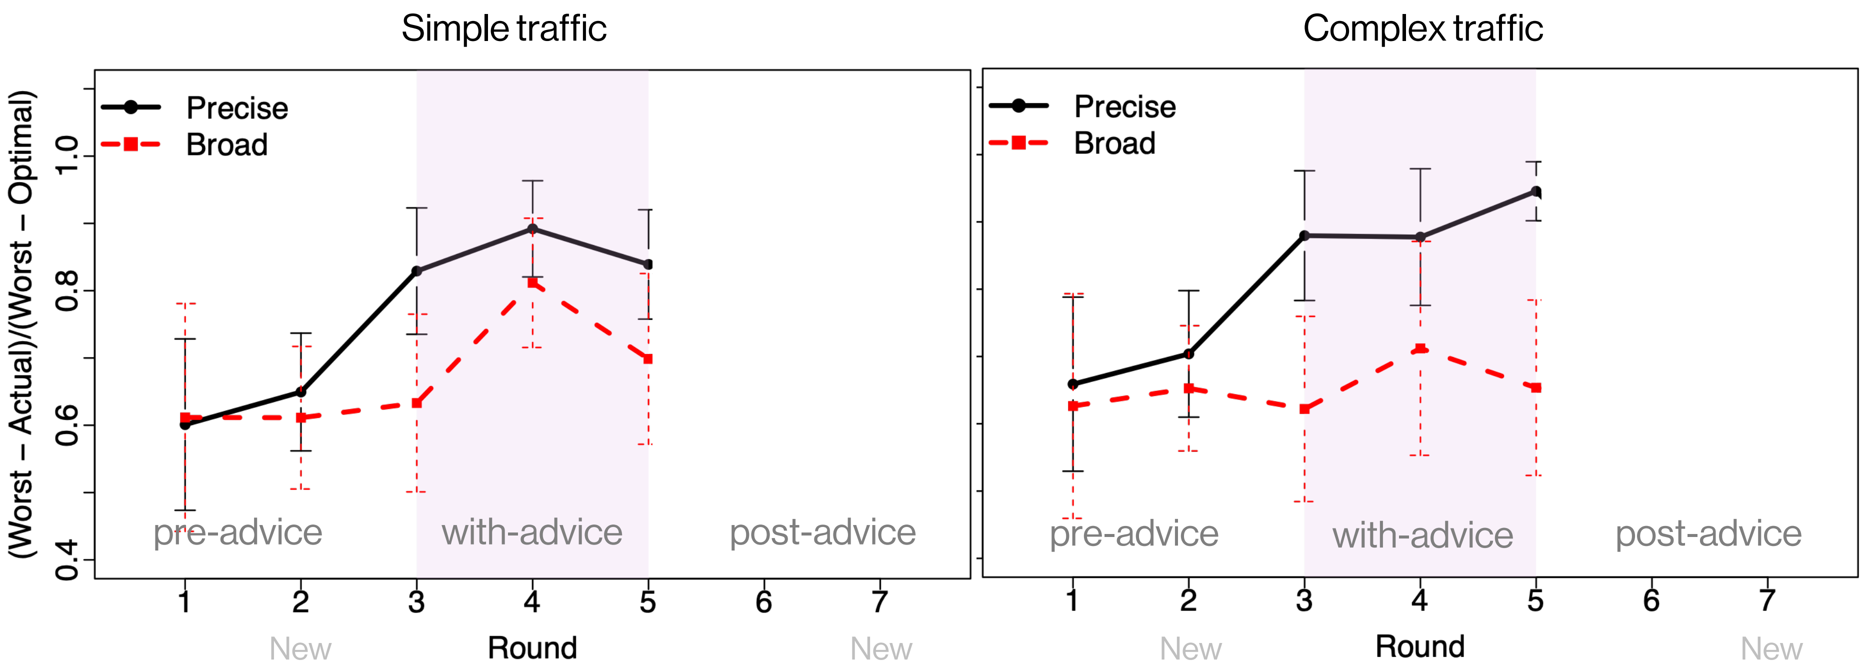

While available, precise advice improves compliance + assisted performance

Amazon Mechanical Turk
N = 102, 3,978 decision points

Y-axis = how much improvement from worst compared to optimal
Higher = better

pre-registered at https://aspredicted.org/7LP_J9G

Study 1

After Advice is Removed...

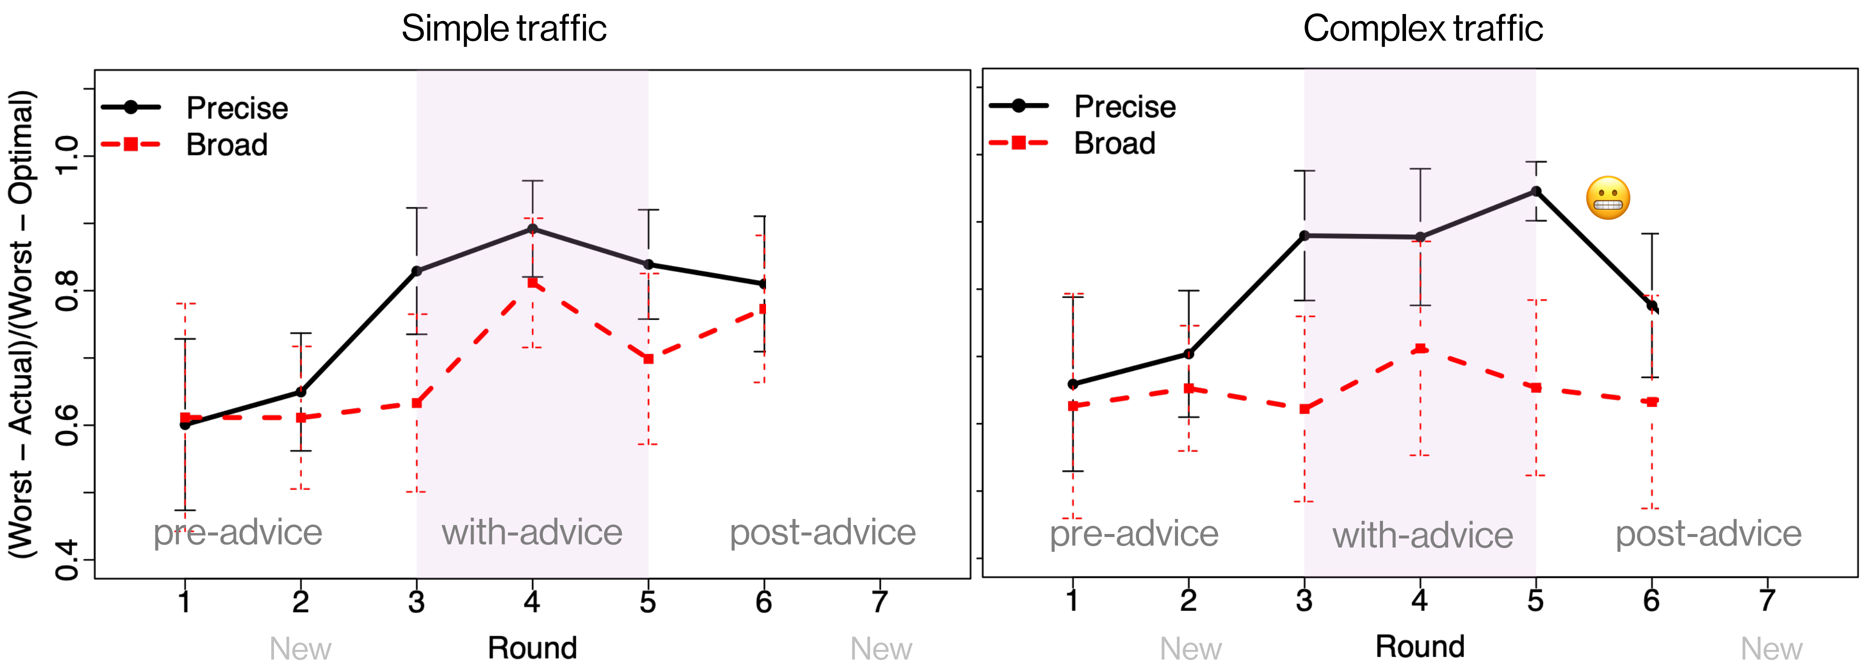

While available, precise advice improves compliance + assisted performance

Amazon Mechanical Turk
N = 102, 3,978 decision points

Y-axis = how much improvement from worst compared to optimal
Higher = better

pre-registered at https://aspredicted.org/7LP_J9G

Study 1

After Advice is Removed...

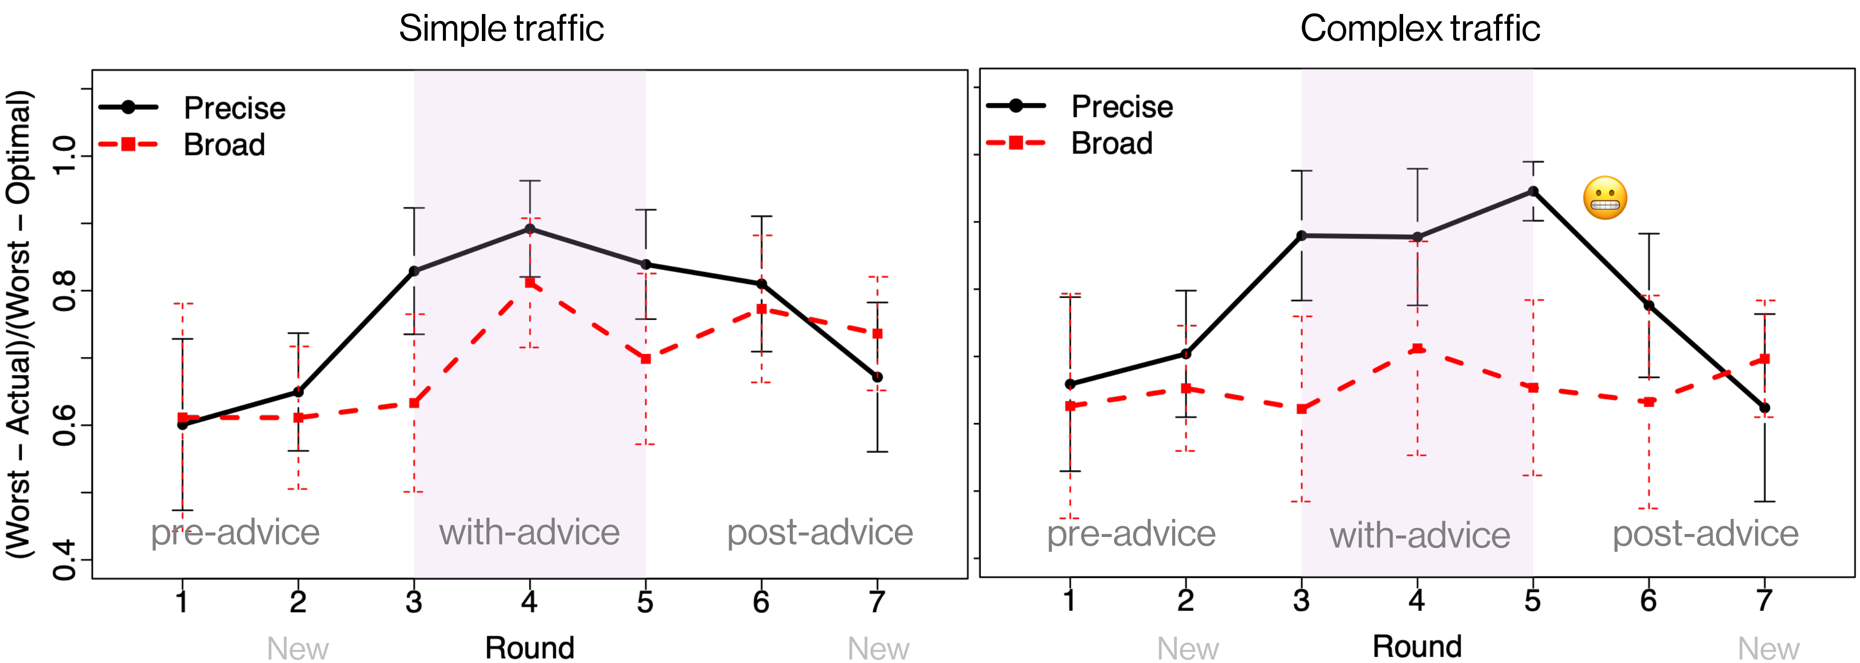

While available, precise advice improves compliance + assisted performance

Amazon Mechanical Turk
N = 102, 3,978 decision points

Y-axis = how much improvement from worst compared to optimal
Higher = better

pre-registered at https://aspredicted.org/7LP_J9G

Study 1

After Advice is Removed...

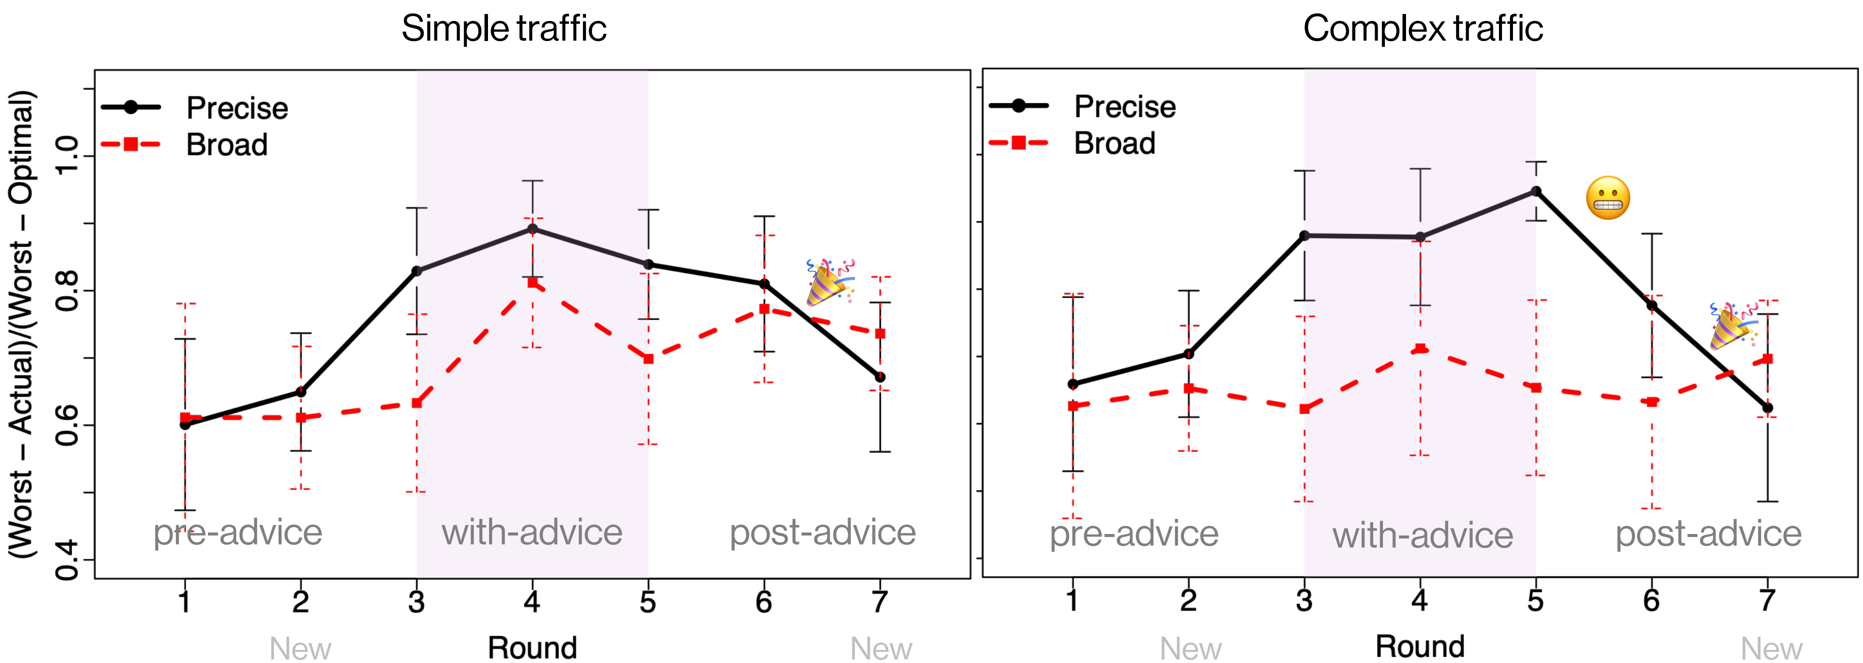

While available, precise advice improves compliance + assisted performance

For new environment, broad seems promising!

(in the direction of the learning-gap prediction!)

Amazon Mechanical Turk
N = 102, 3,978 decision points

Study 1

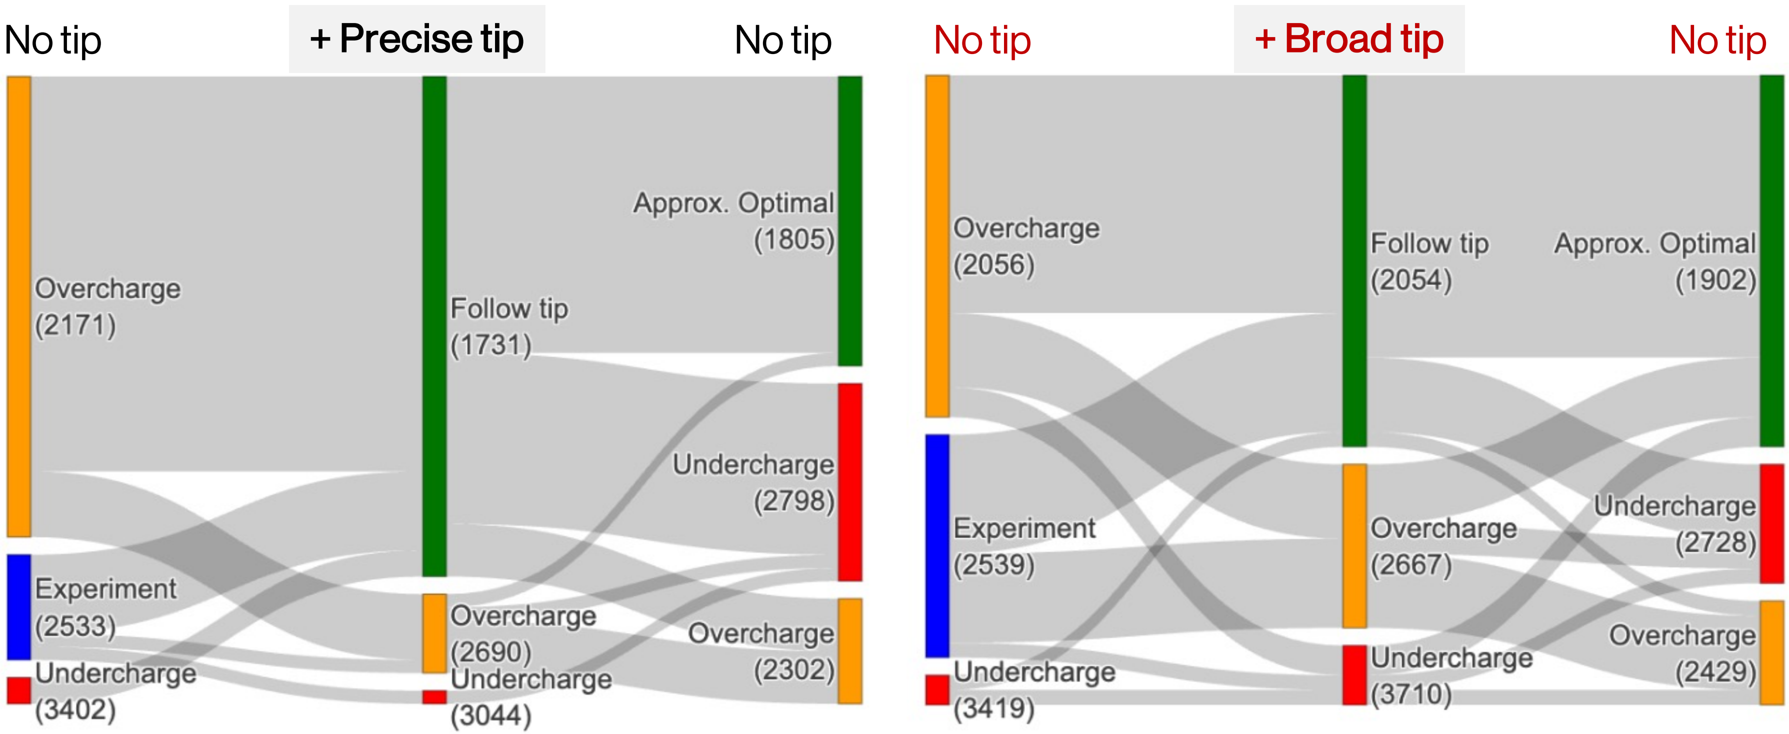
No tip
+ Precise tip
No tip
Overcharge (2171)
Experiment (2533)
Undercharge (3402)
Follow tip (1731)
Overcharge (2690)
Undercharge (3044)
Approx. Optimal (1805)
Undercharge (2798)
Overcharge (2302)
No tip
+ Broad tip
No tip
Overcharge (2056)
Experiment (2539)
Undercharge (3419)
Follow tip (2054)
Overcharge (2667)
Undercharge (3710)
Approx. Optimal (1902)
Undercharge (2728)
Overcharge (2429)

pre-registered at https://aspredicted.org/7LP_J9G

Study 1

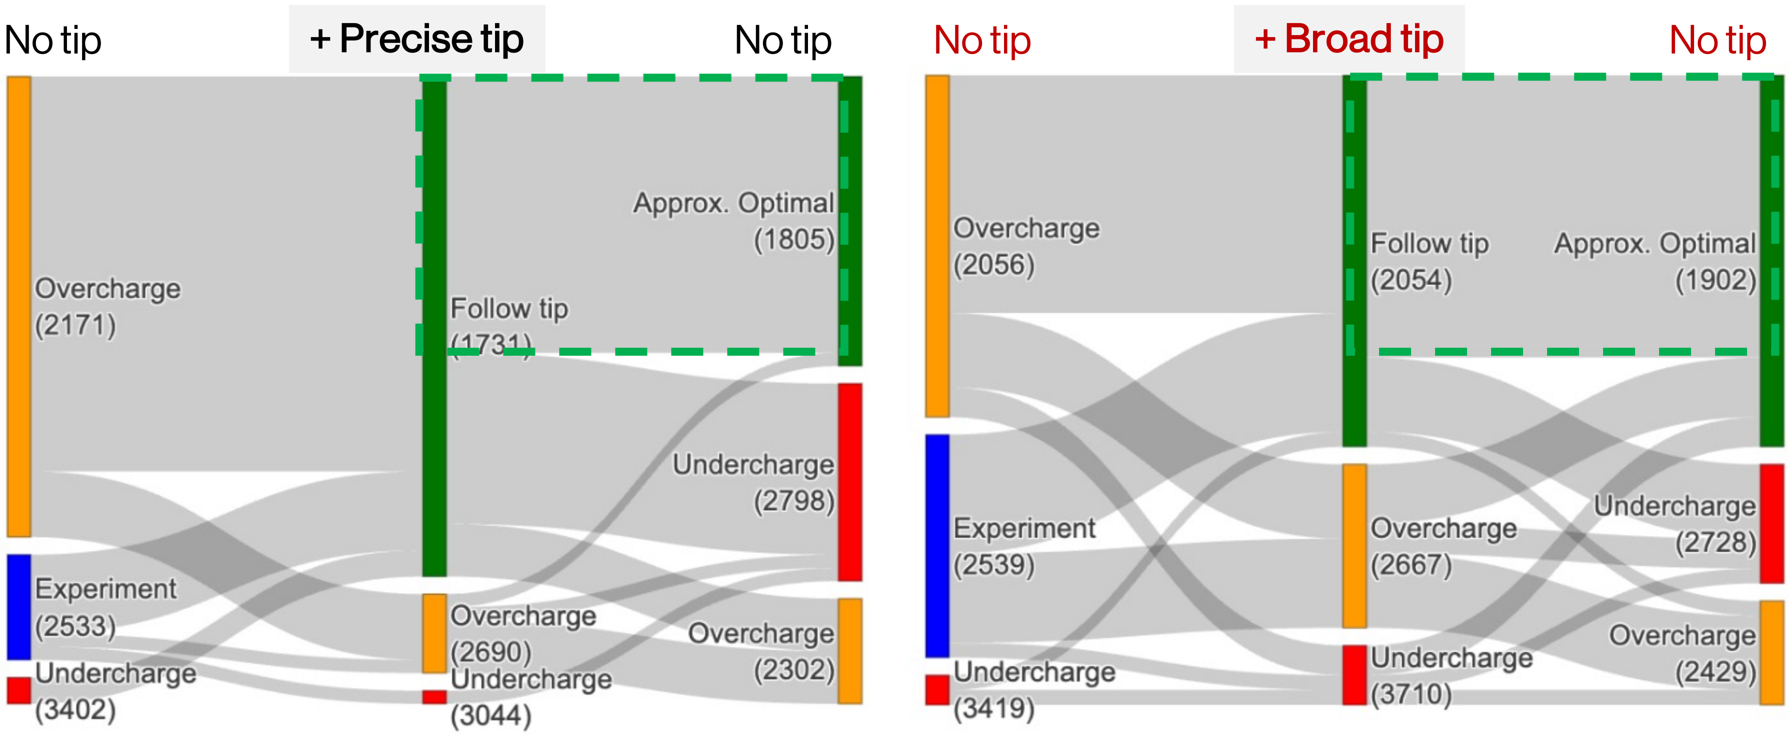

Precise: 55%
stay with optimal strategy afterwards

Broad: 76%
stay with optimal strategy afterwards

pre-registered at https://aspredicted.org/7LP_J9G

Study 1

Broad → Explored More States

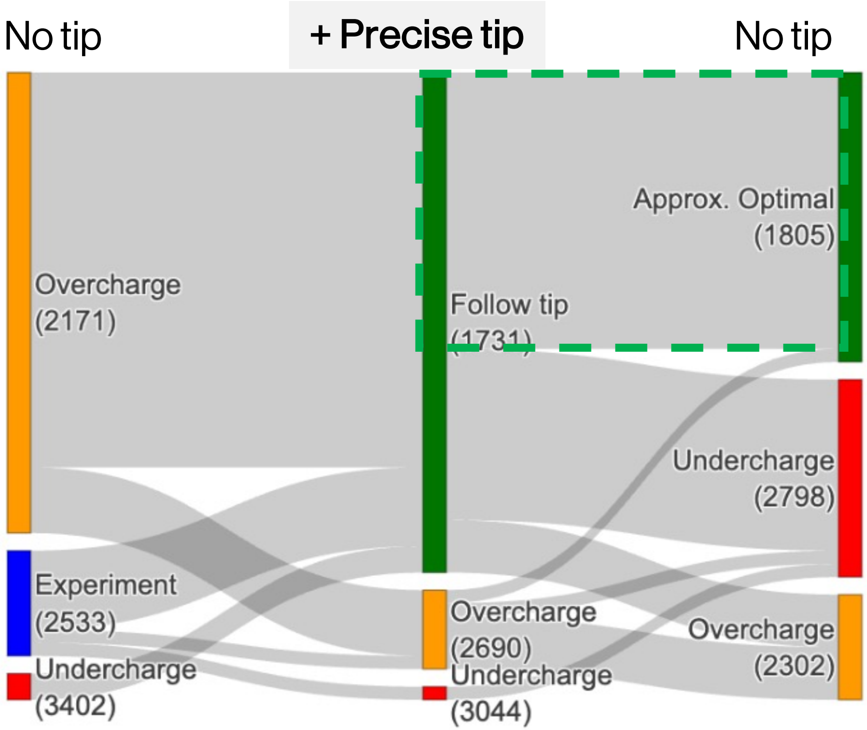

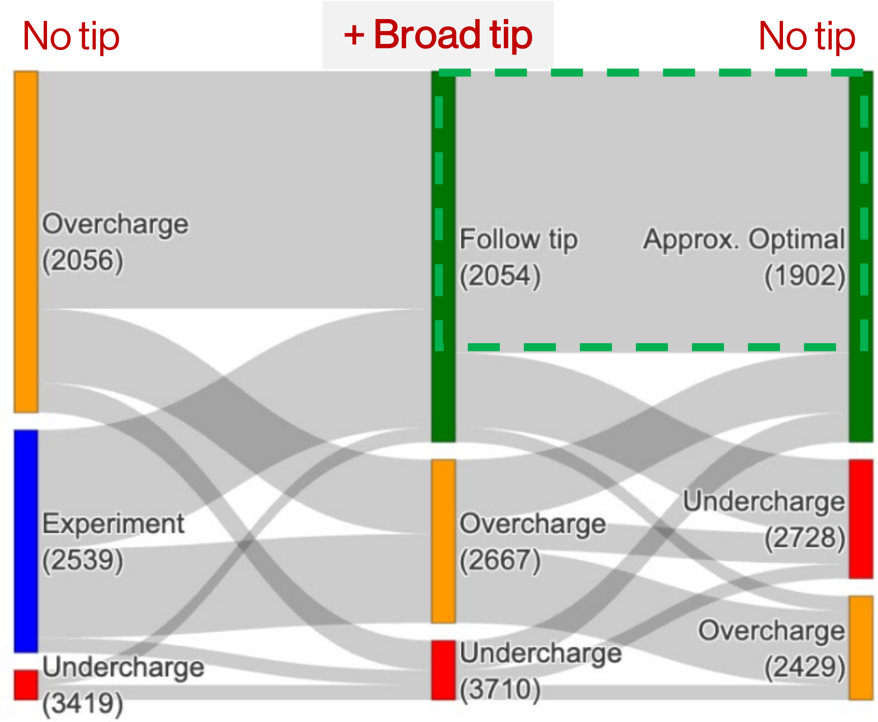

Precise: 55%
stay with optimal strategy afterwards

Broad: 76%
stay with optimal strategy afterwards

(during the tip rounds, they explored more states, figuring out a good strategy)

pre-registered at https://aspredicted.org/23C_WBQ

Study 2

Why Does Broad Help?

	Precise	Broad
Batching	💡 **Tip:** Charge X%	💡 **Tip:** Charge enough for this segment and the next one
Splitting		💡 **Tip:** Charge enough for this segment

pre-registered at https://aspredicted.org/23C_WBQ

Study 2

Why Does Broad Help?

Is the benefit coming from
the strategic rule itself,
or from forcing users to
translate that rule into action?

	Precise	Broad
Batching	💡 **Tip:** Charge X%	💡 **Tip:** Charge enough for this segment and the next one
Splitting		💡 **Tip:** Charge enough for this segment

pre-registered at https://aspredicted.org/23C_WBQ

Study 2

Why Does Broad Help?

Is the benefit coming from **the strategic rule itself**, or from **forcing users to translate that rule into action**?

narrowing the action/
lowering cognitive load

Specific Broad

Broad

💡 **Tip:**
Charge enough
for this segment
and the next one

Batching

Precise

💡 **Tip:**
Charge X%

Splitting

💡 **Tip:**
Charge enough
for this segment

pre-registered at https://aspredicted.org/23C_WBQ

Study 2

Why Does Broad Help?

Is the benefit coming from **the strategic rule itself**, or from **forcing users to translate that rule into action**?

narrowing the action/
lowering cognitive load

	Precise	Specific Broad	Broad
Batching	💡 **Tip:** Charge X%	💡 **Tip:** Charge enough for this segment and the next one, assuming worst case traffic	💡 **Tip:** Charge enough for this segment and the next one
Splitting		💡 **Tip:** Charge enough for this segment, assuming worst case traffic	💡 **Tip:** Charge enough for this segment

pre-registered at https://aspredicted.org/23C_WBQ

Study 2

Why Does Broad Help?

Is the benefit coming from **the strategic rule itself**, or from **forcing users to translate that rule into action**?

narrowing the action/
lowering cognitive load

	Precise	Specific Broad	Broad
Batching	💡 **Tip:** Charge X%	💡 **Tip:** Charge enough for this segment and the next one, assuming worst case traffic	💡 **Tip:** Charge enough for this segment and the next one
Splitting		💡 **Tip:** Charge enough for this segment, assuming worst case traffic	💡 **Tip:** Charge enough for this segment

These two tips lead
to the same exact action

pre-registered at https://aspredicted.org/23C_WBQ

Study 2

Specific Broad

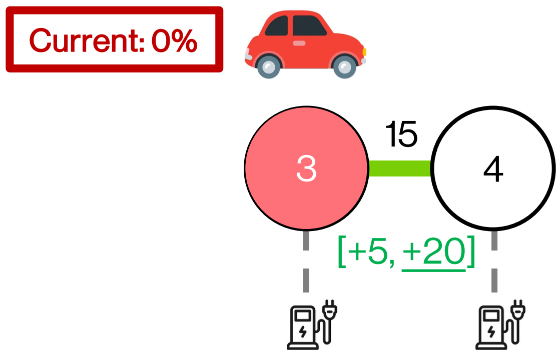

narrowing the action/
lowering cognitive load

Precise	Specific Broad	Broad
💡 **Tip:** Charge 35%	💡 **Tip:** Charge enough for this segment, assuming worst case traffic	💡 **Tip:** Charge enough for this segment

These two tips lead
to the same exact action

15 + 20 = 35

pre-registered at https://aspredicted.org/23C_WBQ

delta irl sum

Study 2

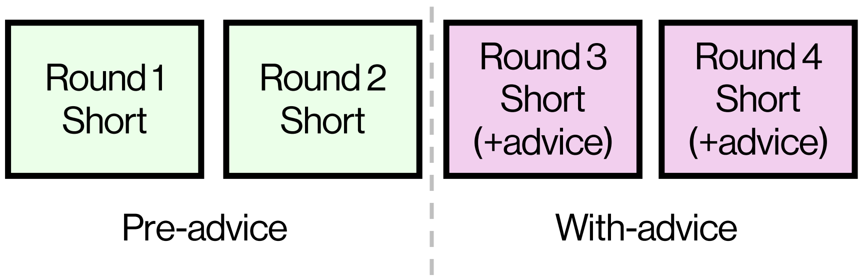

Good to batch early 0-2

Short Map

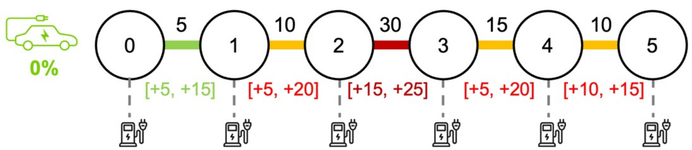

Prolific: N = 400

pre-registered at https://aspredicted.org/23C_WBQ

delta irl sum

Study 2

Round 1 Short | Round 2 Short | Round 3 Short (+advice) | Round 4 Short (+advice) | Round 5 Short | Round 6 Short | Round 7 Short II

Pre-advice (Rounds 1–2) | With-advice (Rounds 3–4) | Post-advice (Rounds 5–7)

Good to batch early 0-2

Short Map

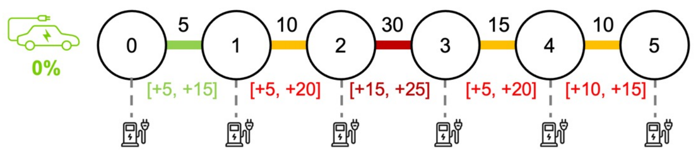

Prolific: N = 400

pre-registered at https://aspredicted.org/23C_WBQ

delta irl sum

Study 2

Pre-advice		With-advice		Post-advice		
Round 1 Short	Round 2 Short	Round 3 Short (+advice)	Round 4 Short (+advice)	Round 5 Short	Round 6 Short	Round 7 Short II

Good to batch early 0-2

Short Map

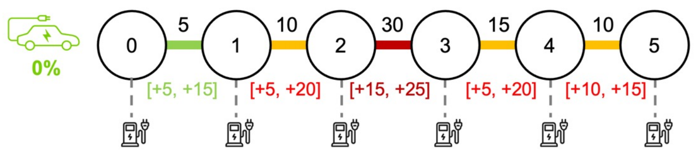

Short II Map

Good to batch late 3-5

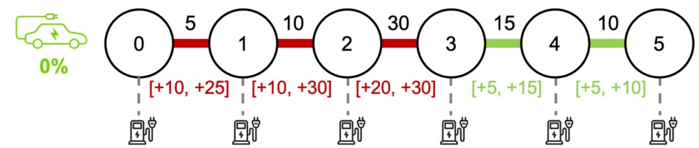

Prolific: N = 400

pre-registered at https://aspredicted.org/23C_WBQ

delta irl sum

Study 2

Round 1 Short	Round 2 Short	Round 3 Short (+advice)	Round 4 Short (+advice)	Round 5 Short	Round 6 Short	Round 7 Short II

Pre-advice | With-advice | Post-advice

Good to batch early 0-2

Short Map

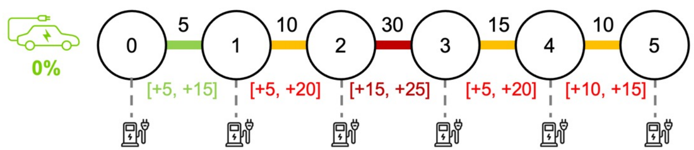

Short II Map

Good to batch late 3-5

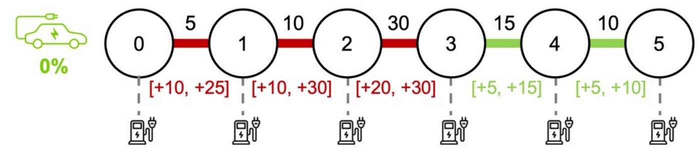

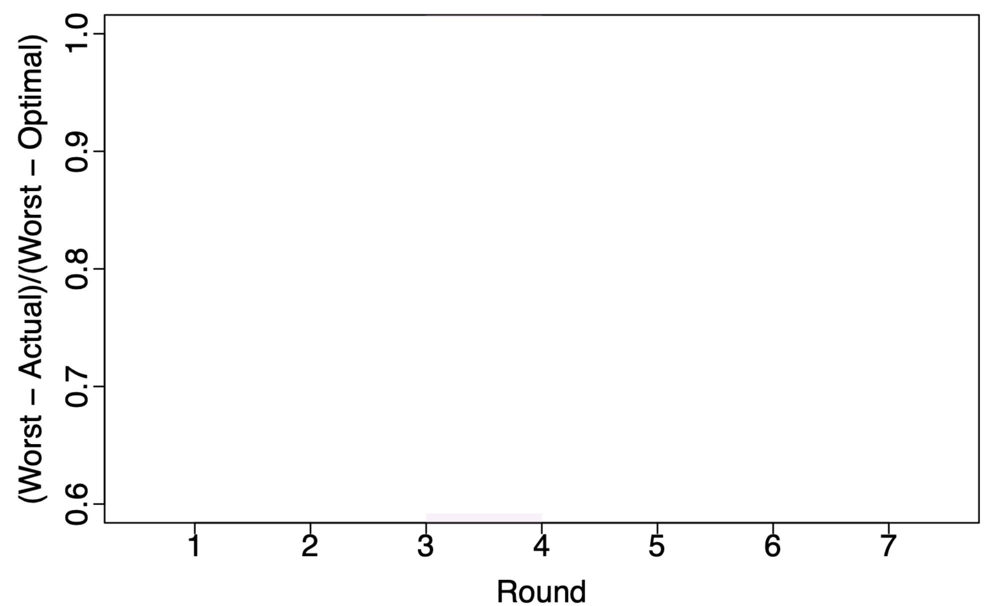

Prolific: N = 400

pre-registered at https://aspredicted.org/23C_WBQ

delta irl sum

Study 2

Replicating Study 1

Pre-advice		With-advice		Post-advice		
Round 1 Short	Round 2 Short	Round 3 Short (+advice)	Round 4 Short (+advice)	Round 5 Short	Round 6 Short	Round 7 Short II

Good to batch early 0-2

Short Map

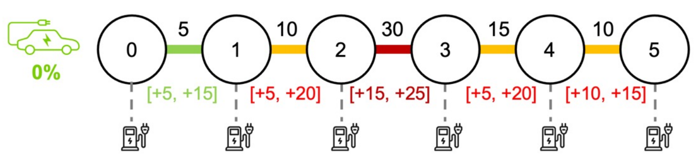

Short II Map

Good to batch late 3-5

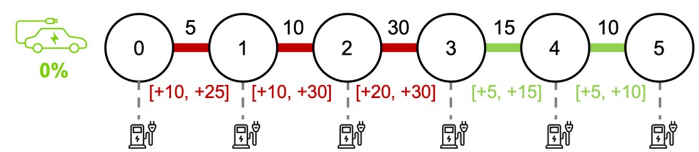

The learning advantage of broad advice is replicated!

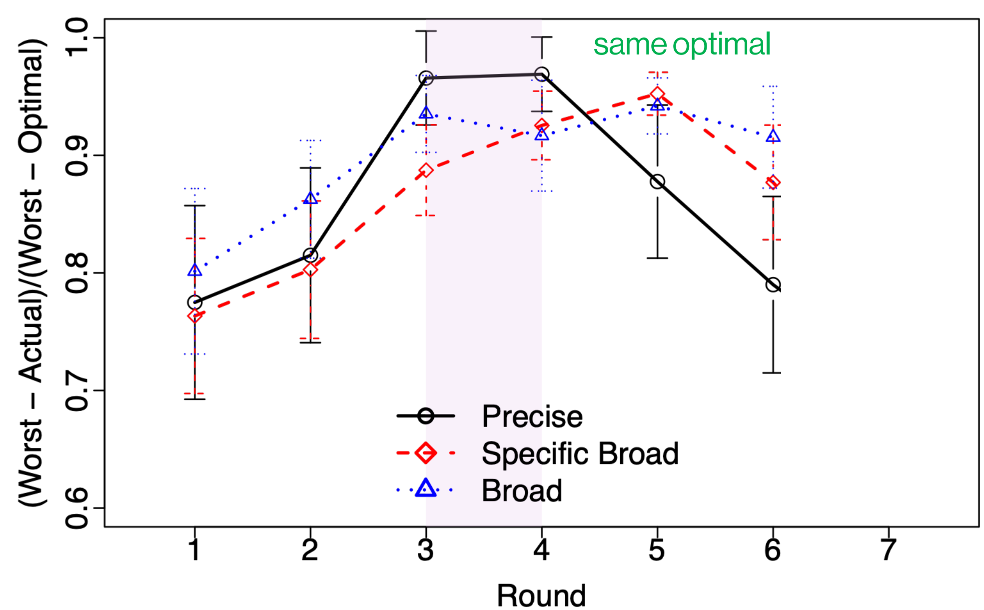

Prolific: N = 400

pre-registered at https://aspredicted.org/23C_WBQ

delta | irl | sum

Study 2

Replicating Study 1

Round 1 Short	Round 2 Short	Round 3 Short (+advice)	Round 4 Short (+advice)	Round 5 Short	Round 6 Short	Round 7 Short II

Pre-advice | With-advice | Post-advice

Good to batch early 0-2

Short Map

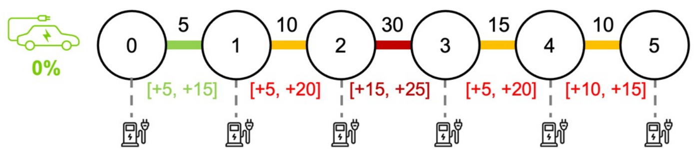

Short II Map

Good to batch late 3-5

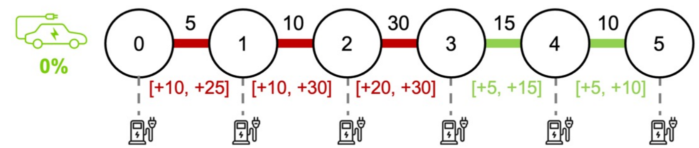

The learning advantage of broad advice is replicated!

Specific broad more like precise → Cognitive effort drives learning

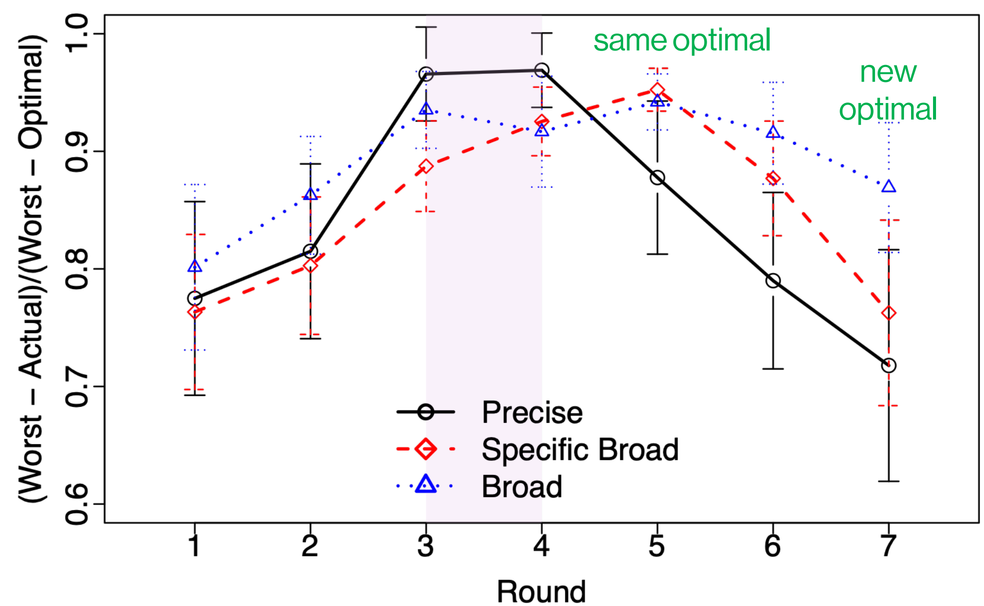

Prolific: N = 400

Study 1 Study 2

Varying the Distance

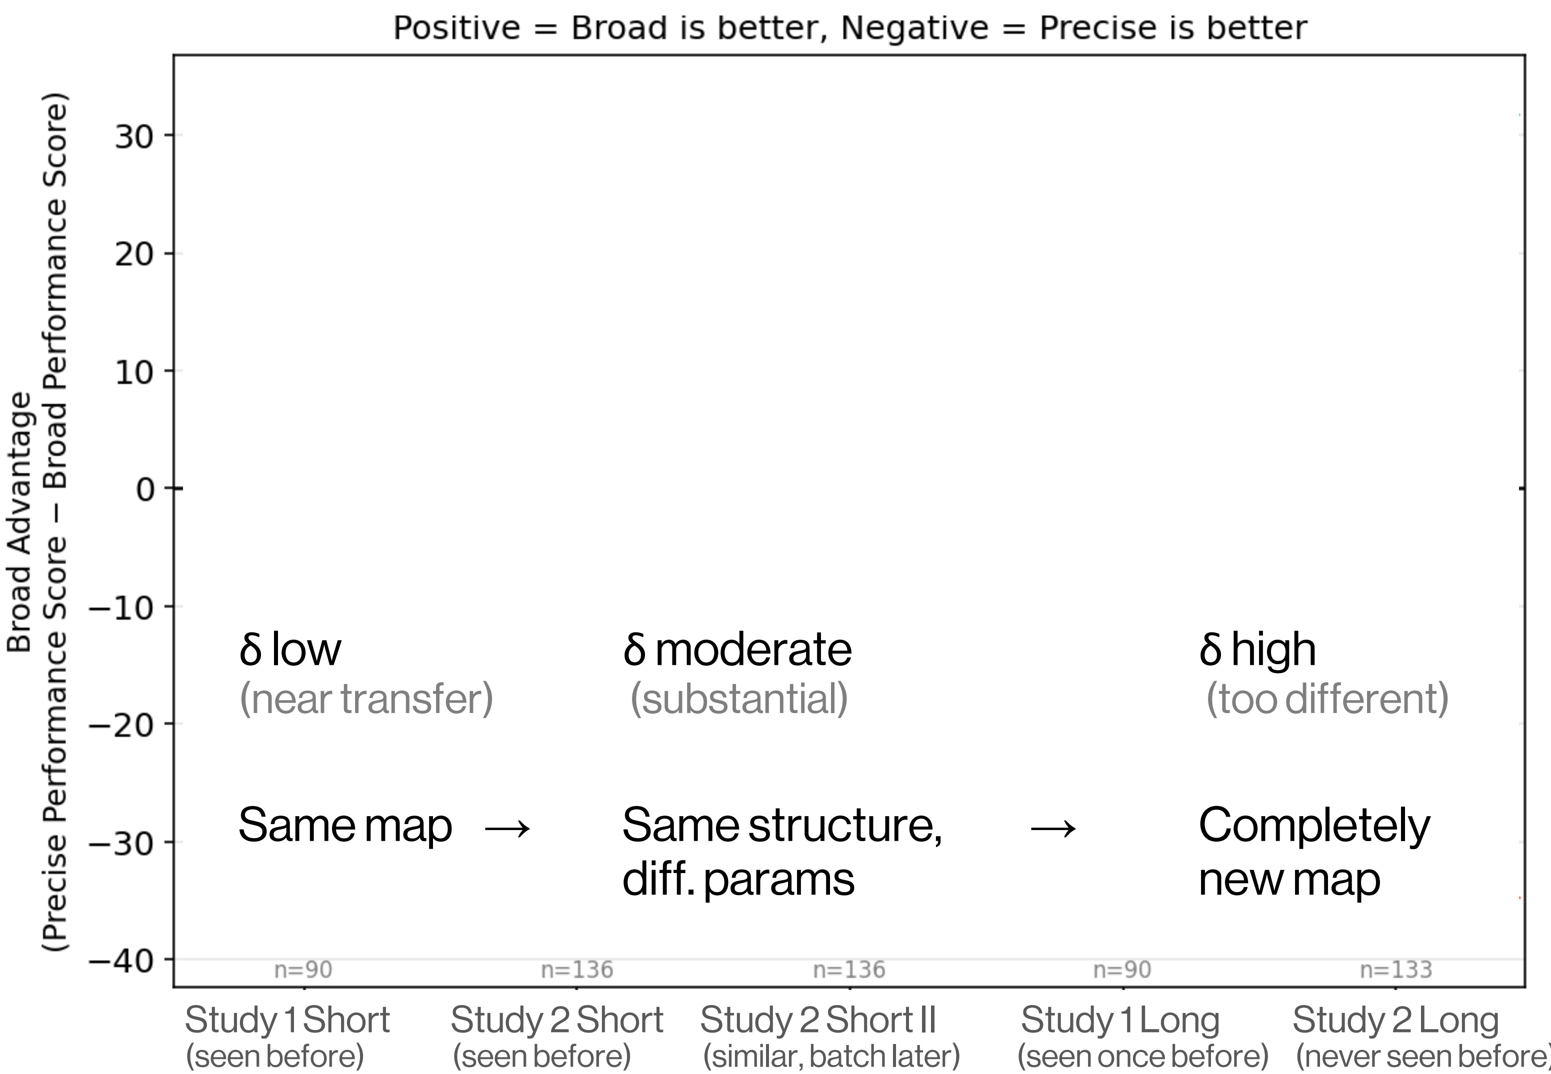

Study 1 Study 2

Varying the Distance

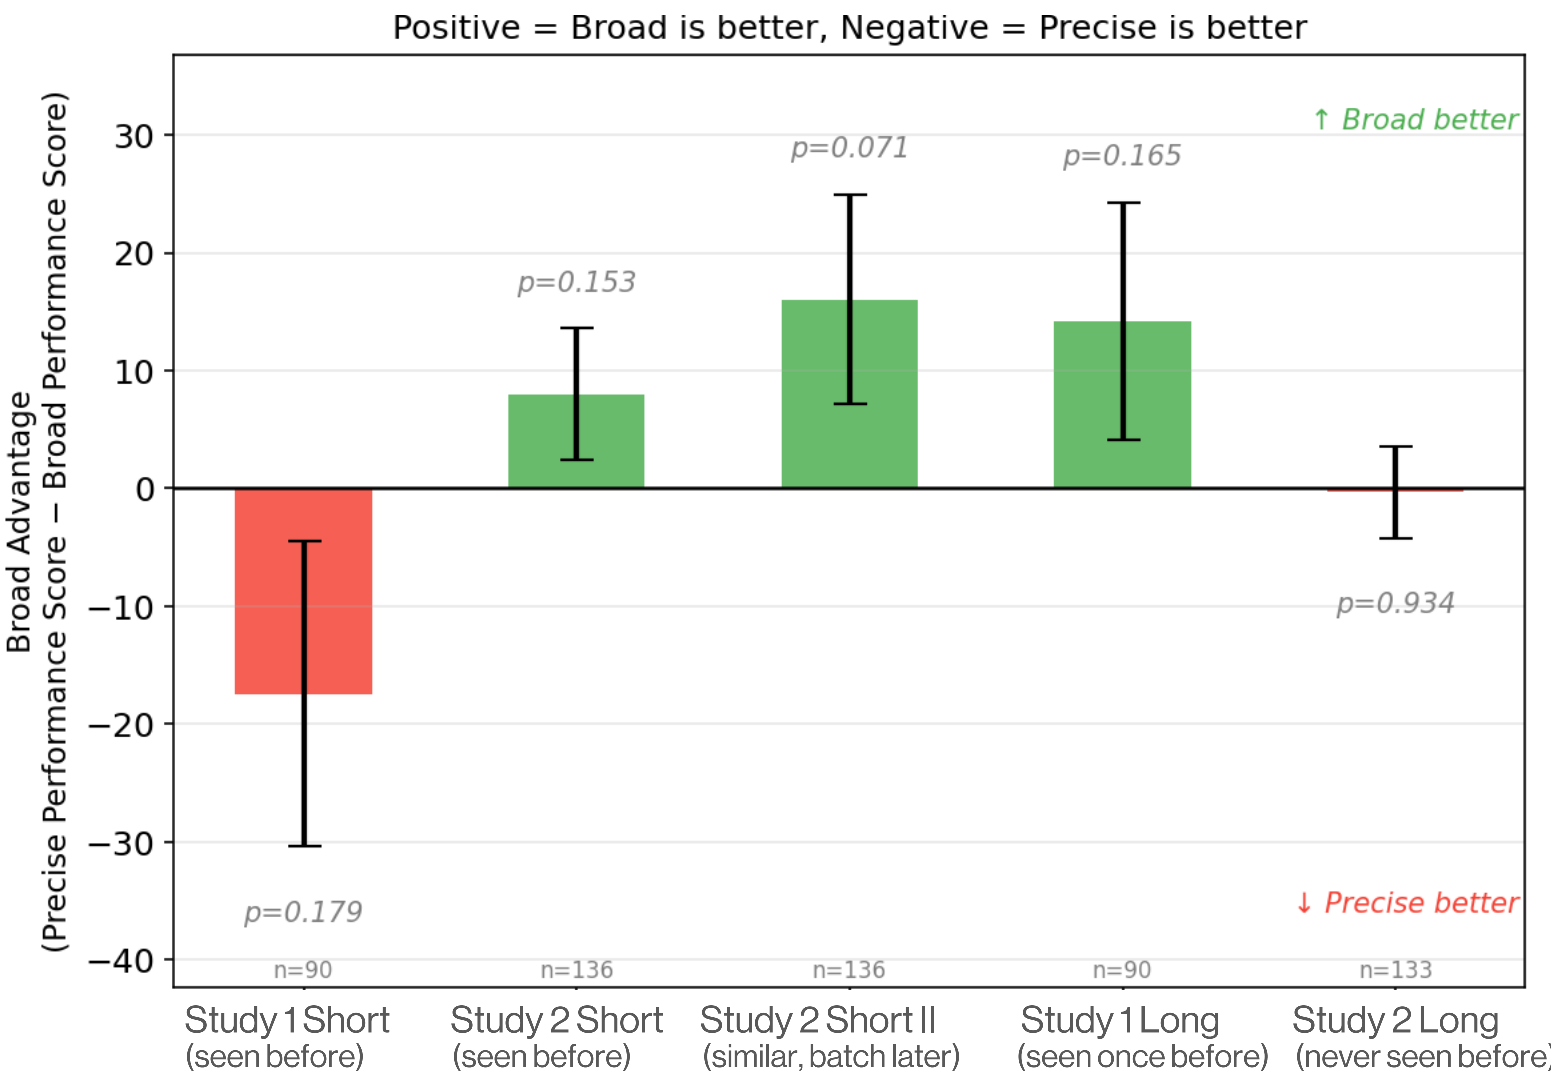

Study 1 Study 2

Varying the Distance

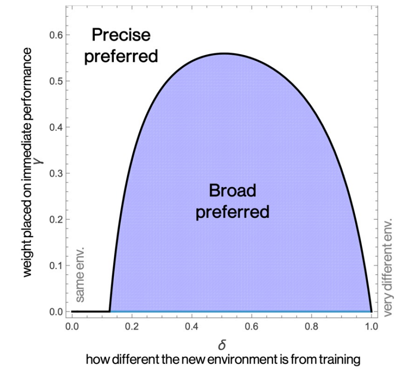

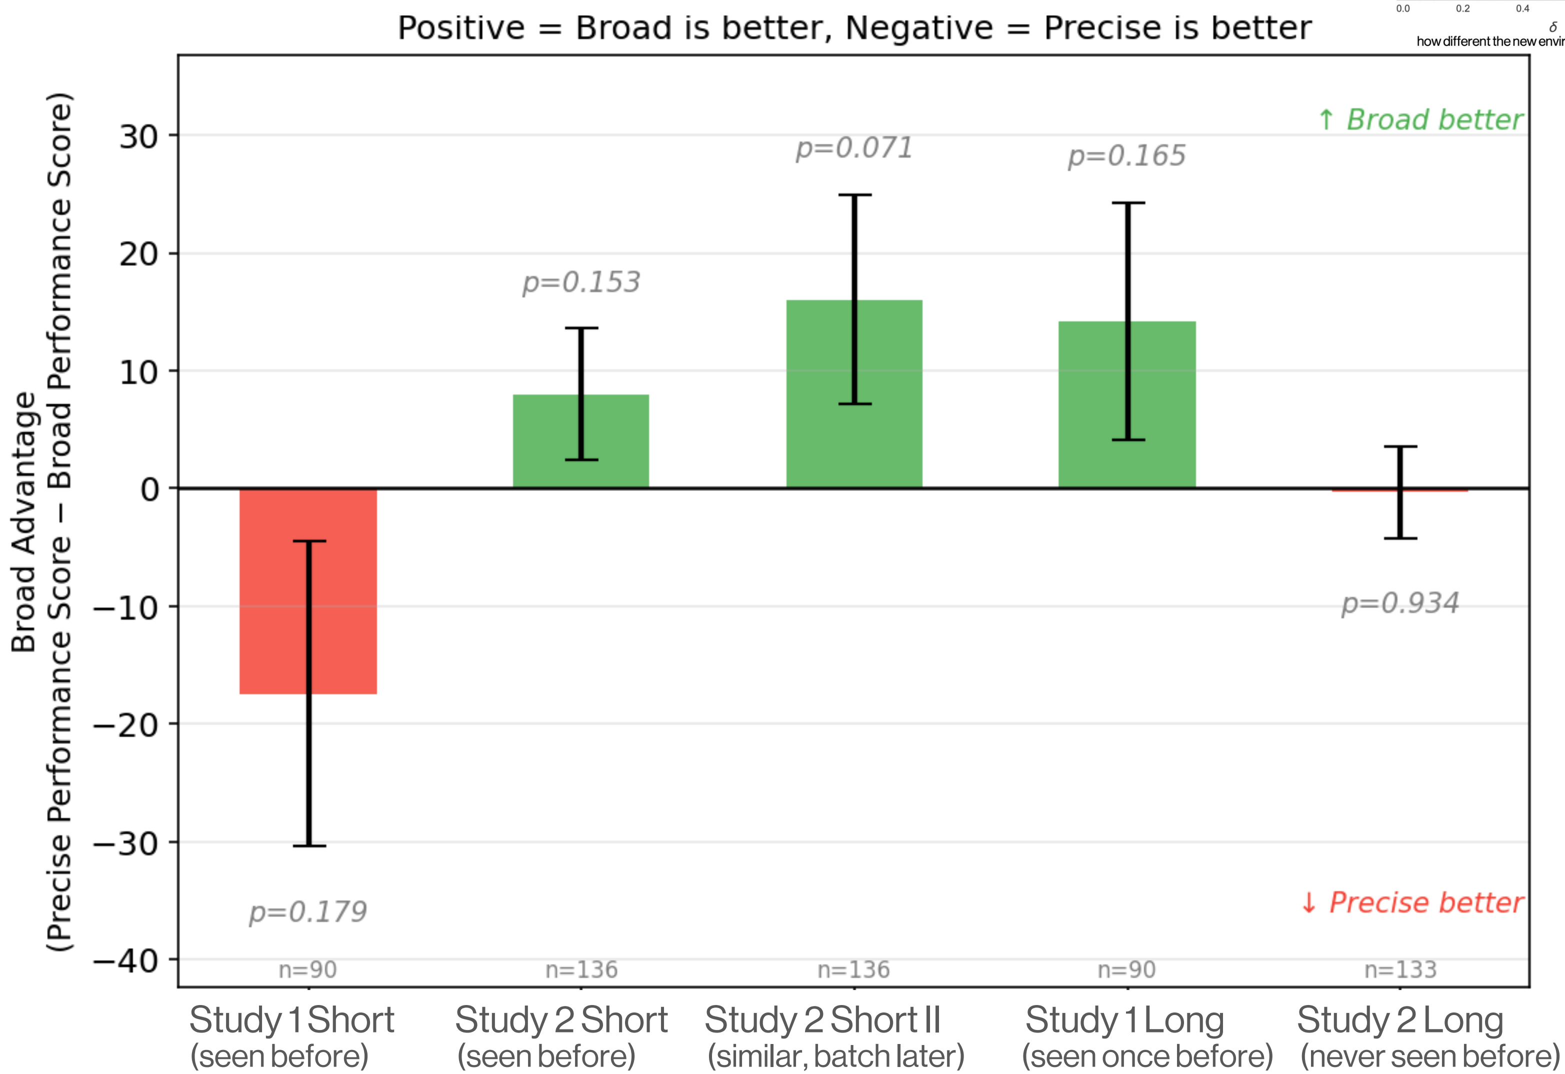

What Are People Actually Optimizing?

What Are People Actually Optimizing?

"Take the total amount of
time needed to get to
the next point..."

What Are People Actually Optimizing?

ϕ_t Time taken at current stop
(including charging & emergency)

"Take the total amount of time needed to get to the next point..."

What Are People Actually Optimizing?

ϕ_t Time taken at current stop
(including charging & emergency)

"Take the total amount of time needed to get to the next point..."

"I first just tried to make sure my charge was 100%."

What Are People Actually Optimizing?

ϕ_t Time taken at current stop
(including charging & emergency)

"Take the total amount of time needed to get to the next point..."

ϕ_s Simplicity
(prefer 0%, 100%, following tip)

"I first just tried to make sure my charge was 100%."

What Are People Actually Optimizing?

ϕ_t Time taken at current stop
(including charging & emergency)

"Take the total amount of time needed to get to the next point..."

ϕ_s Simplicity
(prefer 0%, 100%, following tip)

"I first just tried to make sure my charge was 100%."

"Play it safe!"

What Are People Actually Optimizing?

ϕ_t Time taken at current stop
(including charging & emergency)

"Take the total amount of time needed to get to the next point..."

ϕ_s Simplicity
(prefer 0%, 100%, following tip)

"I first just tried to make sure my charge was 100%."

ϕ_r Risk exposure
(keep a buffer)

"Play it safe!"

What Are People Actually Optimizing?

ϕ_t Time taken at current stop
(including charging & emergency)

"Take the total amount of time needed to get to the next point..."

ϕ_s Simplicity
(prefer 0%, 100%, following tip)

"I first just tried to make sure my charge was 100%."

ϕ_r Risk exposure
(keep a buffer)

"Play it safe!"

"I would probably charge all the way on stop... so you don't have to charge on the last stop"

What Are People Actually Optimizing?

ϕ_t Time taken at current stop
(including charging & emergency)

"Take the total amount of time needed to get to the next point..."

ϕ_s Simplicity
(prefer 0%, 100%, following tip)

"I first just tried to make sure my charge was 100%."

ϕ_r Risk exposure
(keep a buffer)

"Play it safe!"

ϕ_b Batch
(charge sufficiently for multiple stops)

"I would probably charge all the way on stop... so you don't have to charge on the last stop"

What Are People Actually Optimizing?

Inverse RL infers reward function

- Assume participants optimize a personal reward function: a weighted mix of things they care about $r^h_{s_t}(a_t) = \sum_{j=1}^{k} \theta_j \phi_j(s_t, a_t)$

ϕ_t Time taken at current stop
(including charging & emergency)

"Take the total amount of time needed to get to the next point..."

ϕ_s Simplicity
(prefer 0%, 100%, following tip)

"I first just tried to make sure my charge was 100%."

ϕ_r Risk exposure
(keep a buffer)

"Play it safe!"

ϕ_b Batch
(charge sufficiently for multiple stops)

"I would probably charge all the way on stop... so you don't have to charge on the last stop"

What Are People Actually Optimizing?

Inverse RL infers reward function

- Assume participants optimize a personal reward function: a weighted mix of things they care about $r^h_{s_t}(a_t) = \sum_{j=1}^{k} \theta_j \phi_j(s_t, a_t)$

→ we use a Bayesian hierarchical model (SVI) that pools info across participants.

scenario-specific shift
(e.g., pre/with/post, precise vs broad)

$$\theta_i = \theta_0 + \Delta_s + \Delta_i$$

weight on each factor

individual's shift

ϕ_t Time taken at current stop
(including charging & emergency)

"Take the total amount of time needed to get to the next point..."

ϕ_s Simplicity
(prefer 0%, 100%, following tip)

"I first just tried to make sure my charge was 100%."

ϕ_r Risk exposure
(keep a buffer)

"Play it safe!"

ϕ_b Batch
(charge sufficiently for multiple stops)

"I would probably charge all the way on stop... so you don't have to charge on the last stop"

What Stayed After Advice Was Gone?

Y-axis = IRL estimates: post-advice weight – with-advice weight

(Positive = used more after advice disappeared)

scenario-specific shift
(e.g., pre/with/post, precise vs broad)

$$\theta_i = \theta_0 + \Delta_s + \Delta_i$$

weight on
each factor

individual's shift

sum overview

What Stayed After Advice Was Gone?

Y-axis = IRL estimates: post-advice weight – with-advice weight
(Positive = used more after advice disappeared)

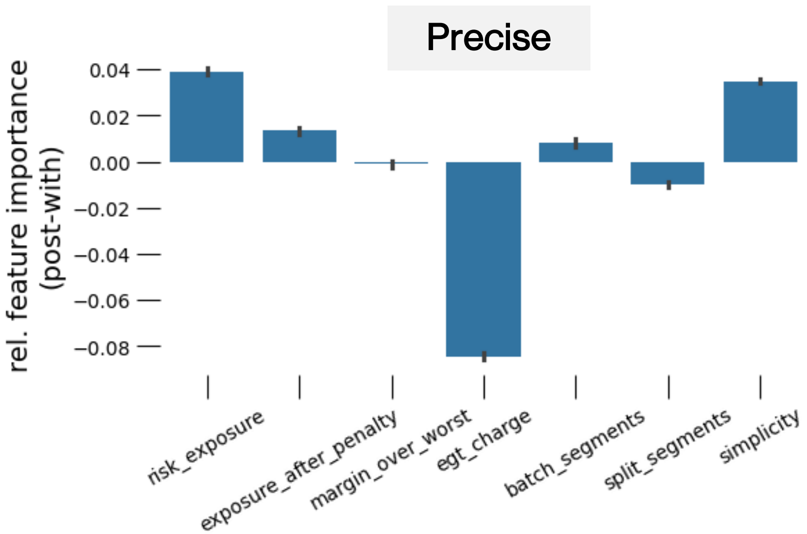

scenario-specific shift
(e.g., pre/with/post, precise vs broad)

$$\theta_i = \theta_0 + \Delta_s + \Delta_i$$

weight on
each factor

individual's shift

sum overview

What Stayed After Advice Was Gone?

Y-axis = IRL estimates: post-advice weight – with-advice weight
(Positive = used more after advice disappeared)

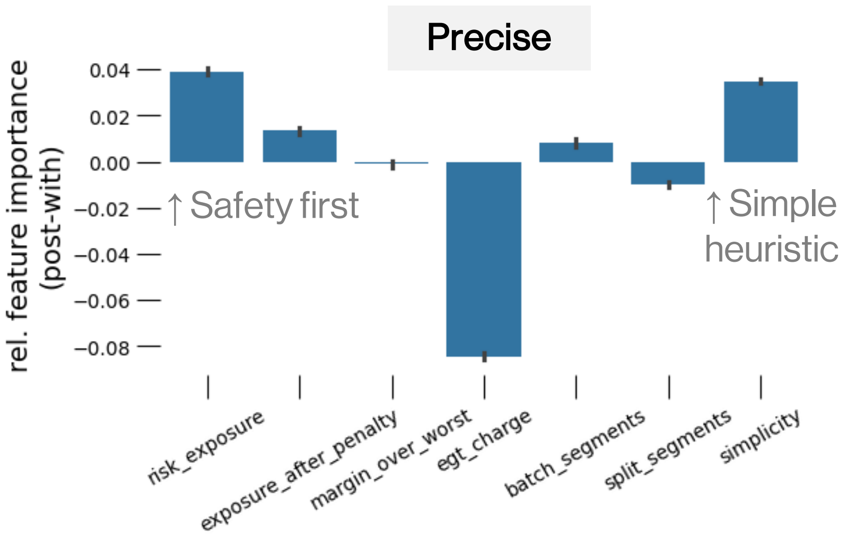

scenario-specific shift
(e.g., pre/with/post, precise vs broad)

$$\theta_i = \theta_0 + \Delta_s + \Delta_i$$

weight on each factor

individual's shift

sum overview

What Stayed After Advice Was Gone?

Y-axis = IRL estimates: post-advice weight – with-advice weight

(Positive = used more after advice disappeared)

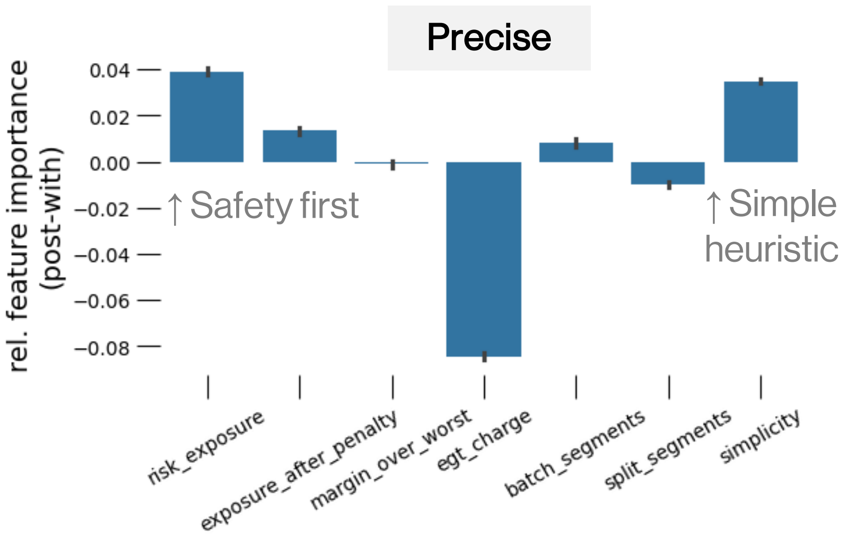

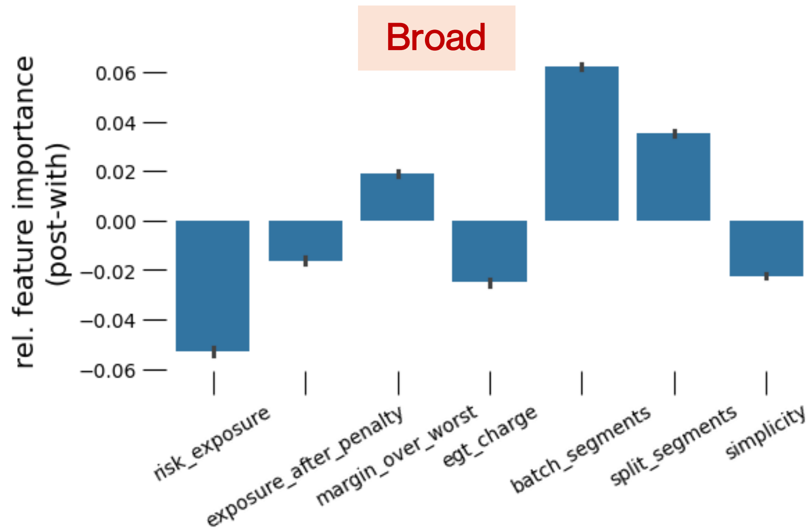

scenario-specific shift
(e.g., pre/with/post, precise vs broad)

$$\theta_i = \theta_0 + \Delta_s + \Delta_i$$

weight on each factor

individual's shift

What Stayed After Advice Was Gone?

Y-axis = IRL estimates: post-advice weight – with-advice weight
(Positive = used more after advice disappeared)

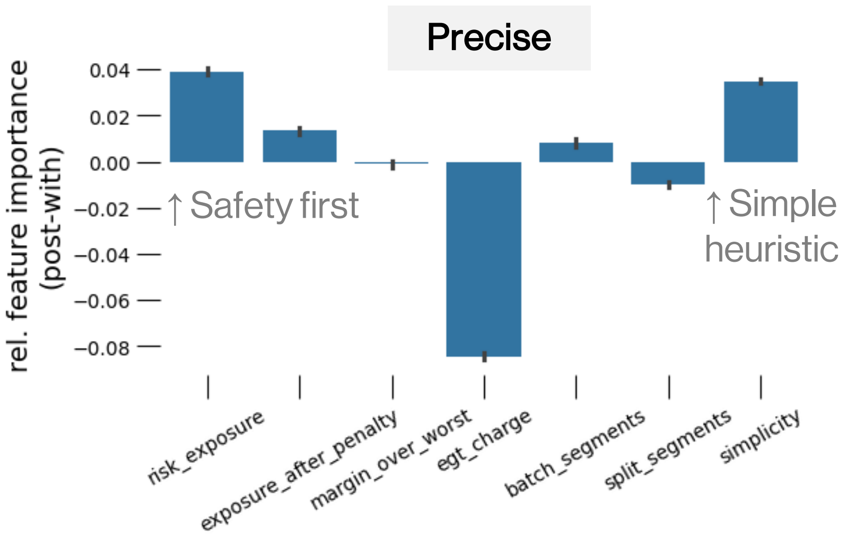

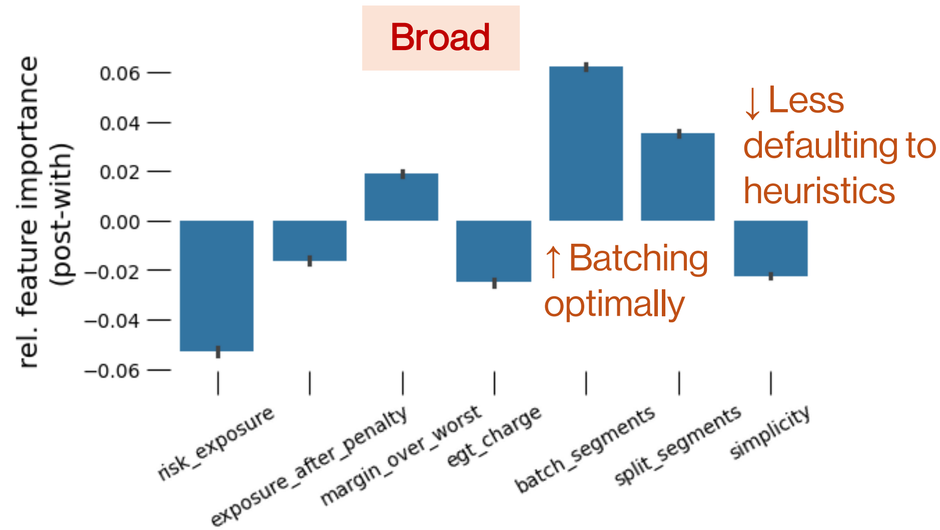

scenario-specific shift
(e.g., pre/with/post, precise vs broad)

$$\theta_i = \theta_0 + \Delta_s + \Delta_i$$

weight on
each factor

individual's shift

sum

overview

What Stayed After Advice Was Gone?

Y-axis = IRL estimates: post-advice weight – with-advice weight
(Positive = used more after advice disappeared)

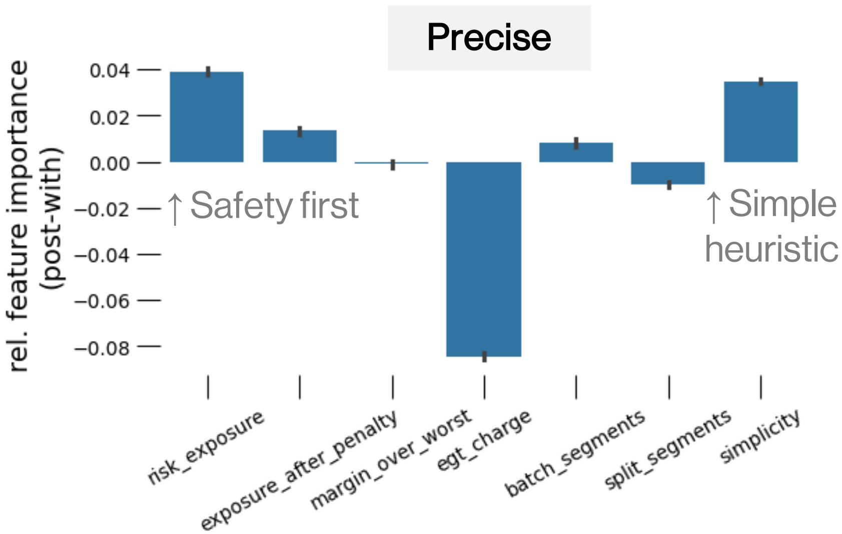

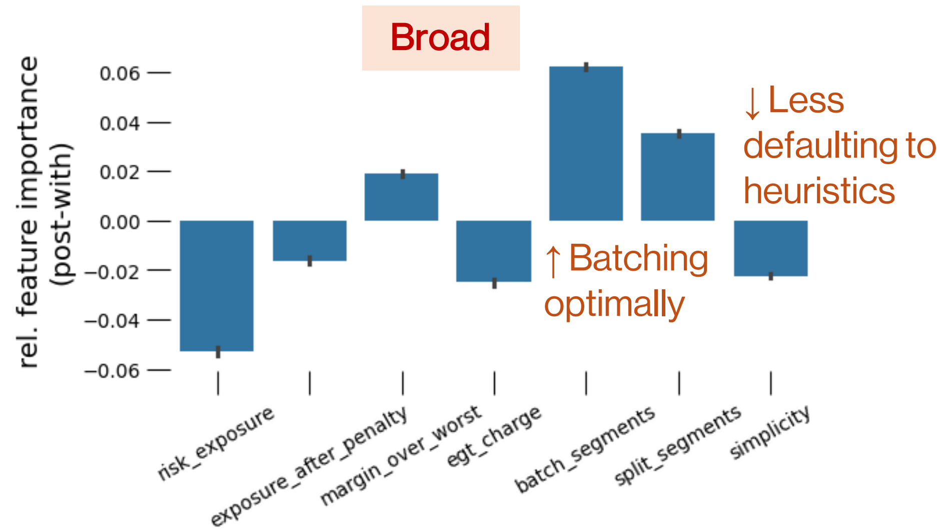

scenario-specific shift
(e.g., pre/with/post, precise vs broad)

$$\theta_i = \theta_0 + \Delta_s + \Delta_i$$

weight on each factor

individual's shift

Precise helped while it was present,
but fell back to safe, simple rules.
Broad left more of the strategy behind.

Designing AI Advice for Sequential Decisions

Precise advice improves assisted execution.
But unsupported capability can be weaker once advice is removed.

Broad advice performs better under moderate transfer.
It requires users to map a rule into an action.

Specific broad identifies the action-mapping burden.
Strategic detail alone doesn't close the short-run reward gap.

Designing AI Advice for Sequential Decisions

Precise advice improves assisted execution.
But unsupported capability can be weaker once advice is removed.

Broad advice performs better under moderate transfer.
It requires users to map a rule into an action.

Specific broad identifies the action-mapping burden.
Strategic detail alone doesn't close the short-run reward gap.

Evaluate advice by what it improves now and what users retain afterward.

Designing AI Advice for Sequential Decisions

Precise advice improves assisted execution.
But unsupported capability can be weaker once advice is removed.

Broad advice performs better under moderate transfer.
It requires users to map a rule into an action.

Specific broad identifies the action-mapping burden.
Strategic detail alone doesn't close the short-run reward gap.

Evaluate advice by what it improves now and what users retain afterward.

Thank you + feedback (broad or precise) very welcome!

Paper link →

Park Sinchaisri (UC Berkeley) / parksinchaisri@berkeley.edu / parksinchaisri.github.io