


# Skin in the Game



Conscious feelings may arise from the mind's reading of the body's condition. We give an artificial agent a purely homeostatic objective and **a body that is vulnerable**: its actions have consequences on its ability to continue taking action. When it makes decisions by modeling its own potential futures, **its behavior begins to look like its life is at stake, because it is.**

In living organisms, conscious feelings are the mental expression of bodily viability [1,2,3]. Artificial agents lack feelings because they have no real stake in their own survival [8]. Homeostatic reinforcement learning treats the maintenance of the body as an intrinsic source of reward [9,10,11], but it typically leaves the agent's competence unaffected by its condition.

We introduce embodied vulnerability: the agent's internal state directly gates its capacity to act. By simulating potential actions through their consequences on its own vulnerable body (an as-if loop [4,6,15]), the agent foresees its own coming unreliability and acts early to preserve agency. We detect the resulting sparks of artificial feeling, measurable in quality of life and recognizable in quality of behavior.

## WHAT IT OPTIMIZES

$$r = \alpha(d_t - d_{t+1}) + \beta(1 - d_{t+1}/d_{max}) - 1_{death}$$

Drive reduction ·  $\alpha=5$     setpoint reward ·  $\beta=0.1$     death penalty

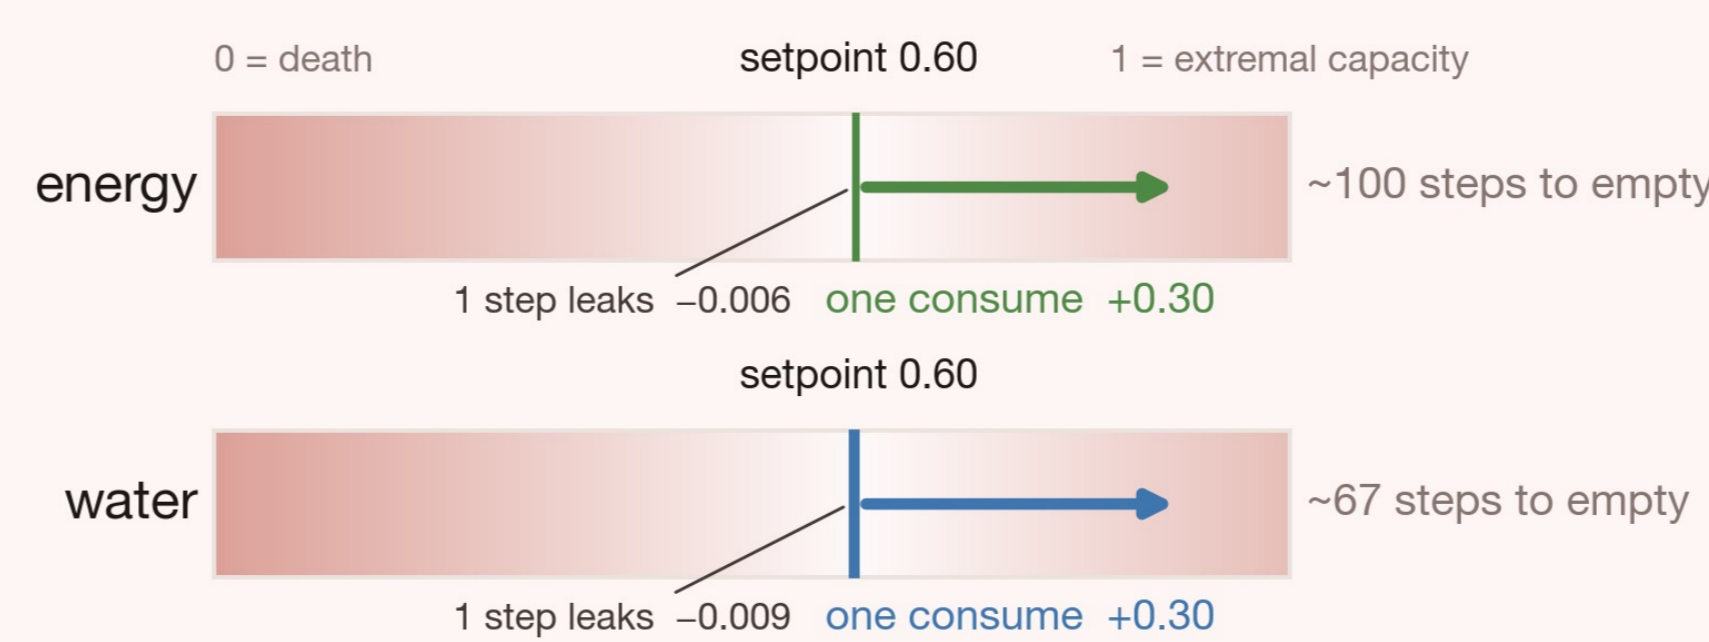
The agent is rewarded for reducing homeostatic drive and being close to the setpoint, and penalized if it dies.

Homeostatic drive  $d = (e - e^*)^2 + (w - w^*)^2$  the squared distance of energy  $e$  and water  $w$  from their setpoints  $e^* = w^* = 0.6$ .

## THE WORLD AND ITS SIX ACTIONS


- A 12x12 world, lived once: no resets, 1000-step ceiling.
- Six discrete actions: move (4 ways), consume (only while on a resource tile, never automatic), and rest (halves the energy leak). No flee, no build.
- A predator hunts at half the agent's speed. A bite costs energy but knocks the predator back 4 steps, the agent's chance to flee.
- A stone is pushed by stepping into it and dragged when you step away from it in line; step sideways and you leave it. Sealing the den is the planner's choice, never scripted.

## THE METABOLISM

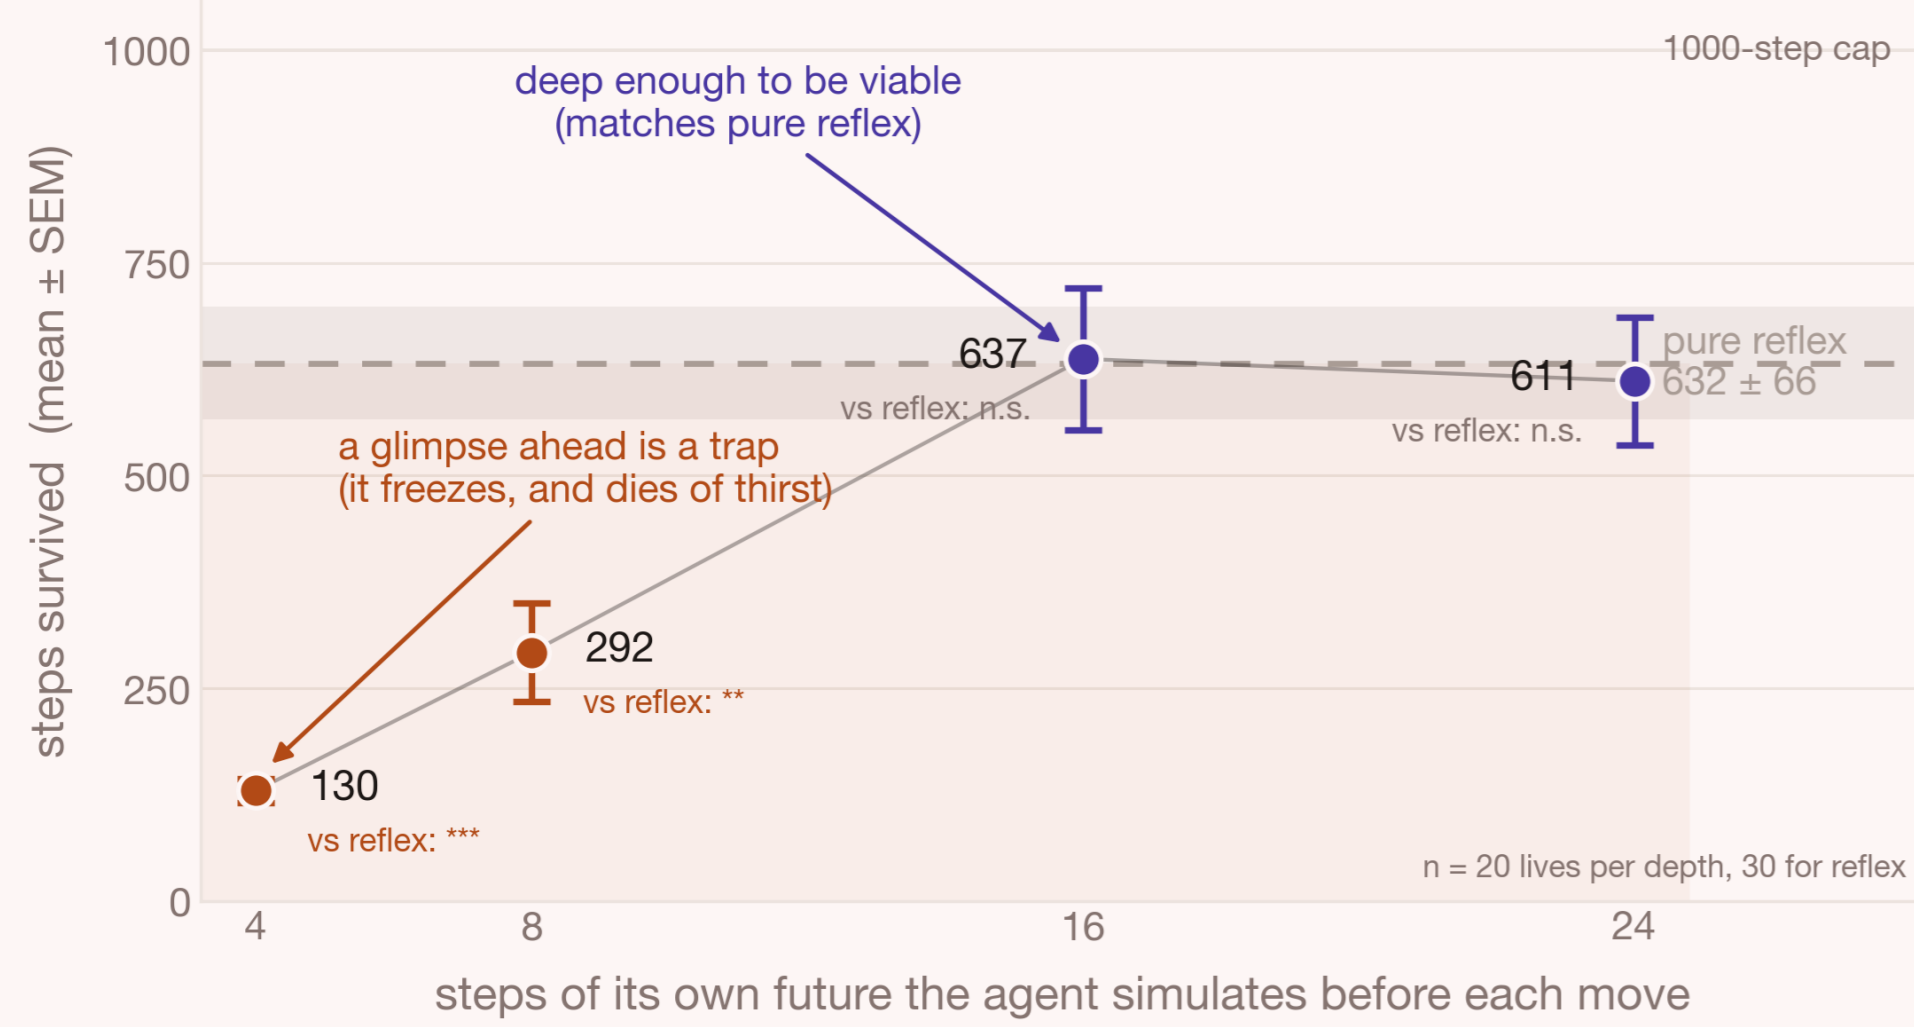


## PLANNING WITH THE BODY IN THE LOOP

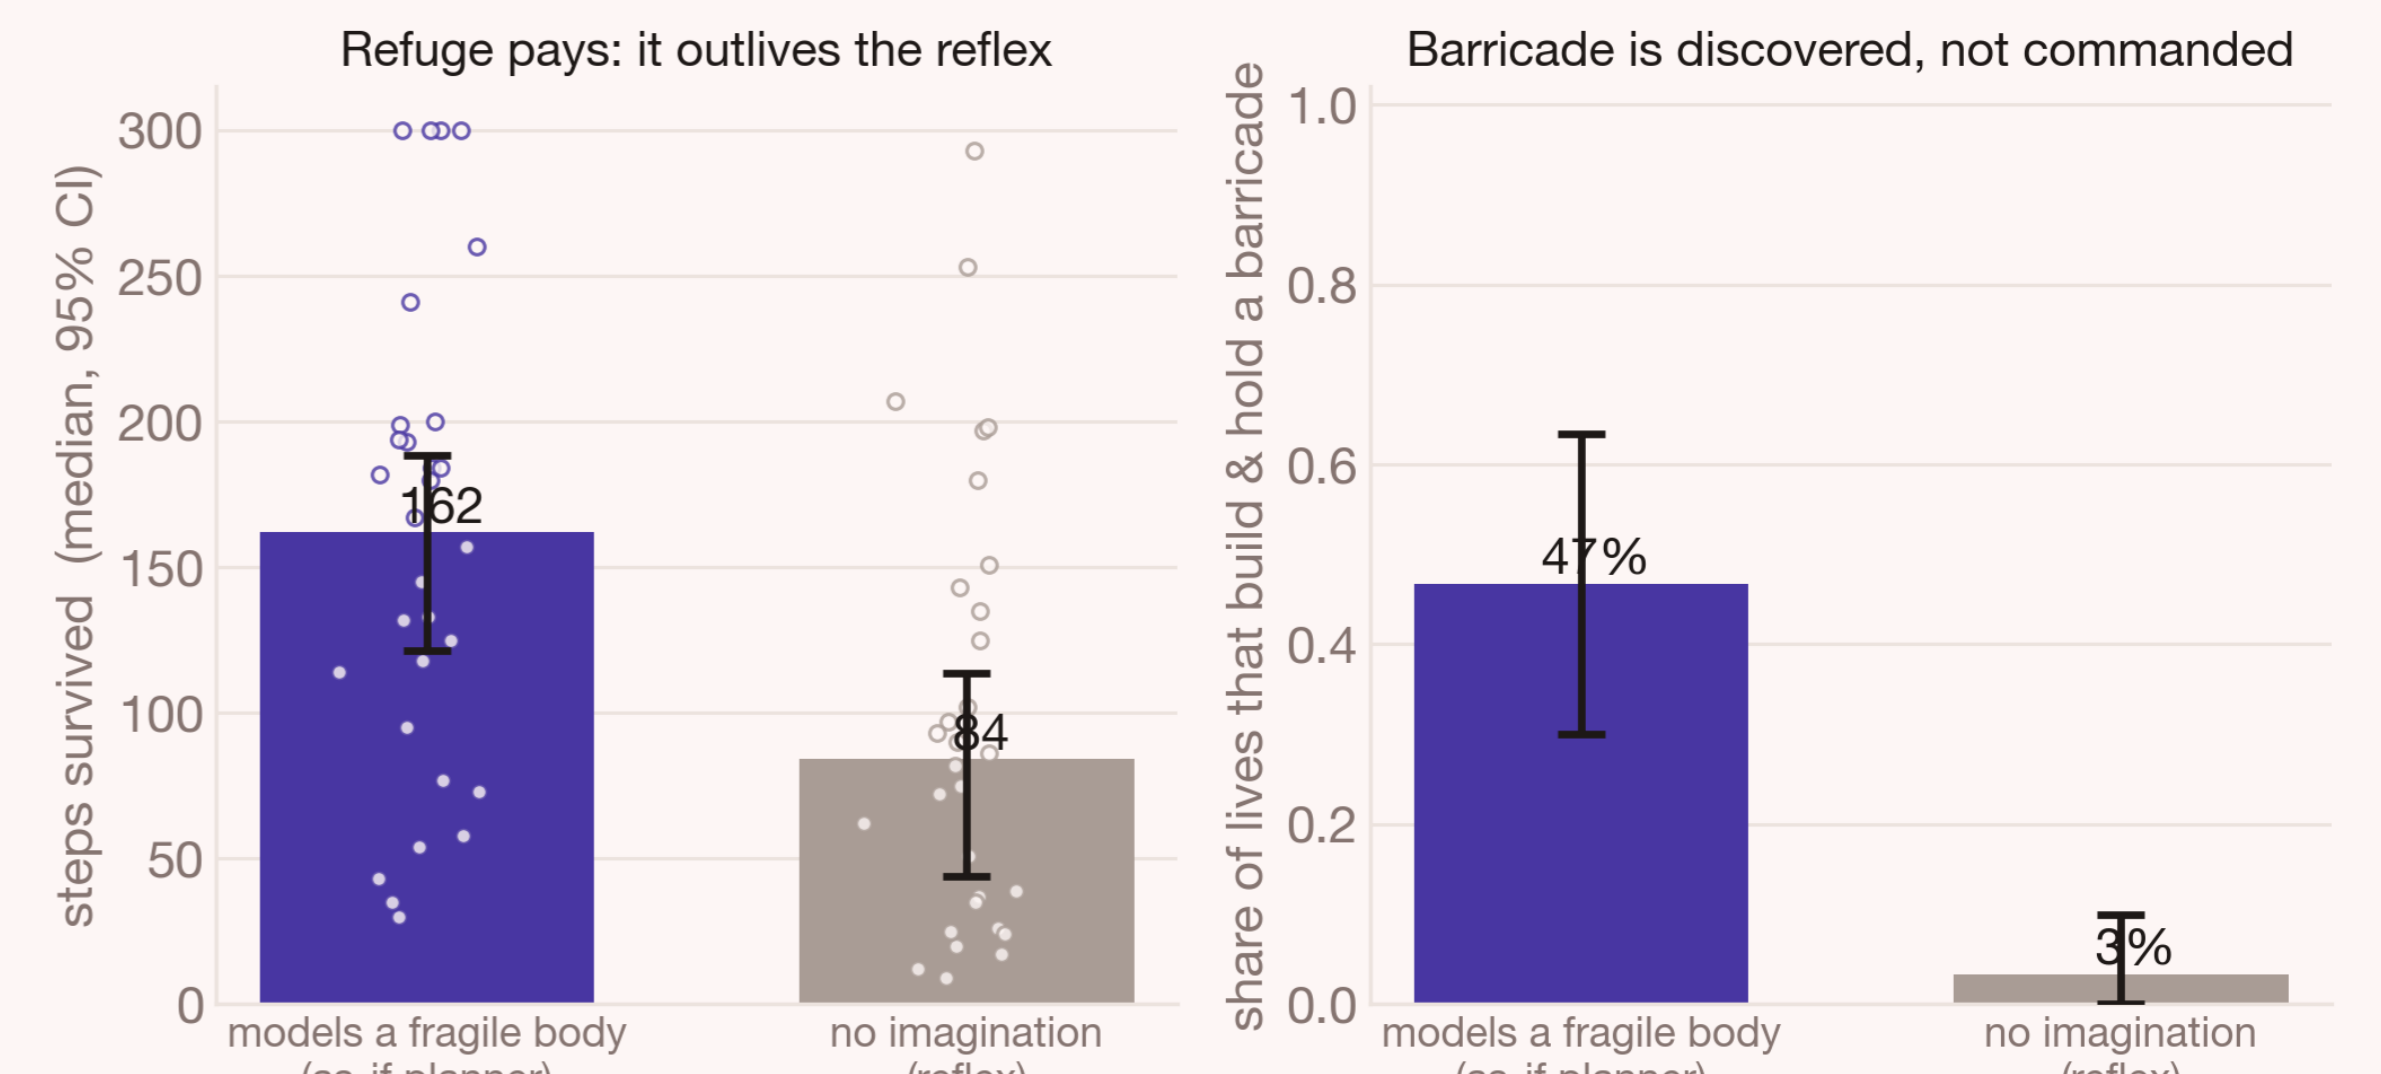
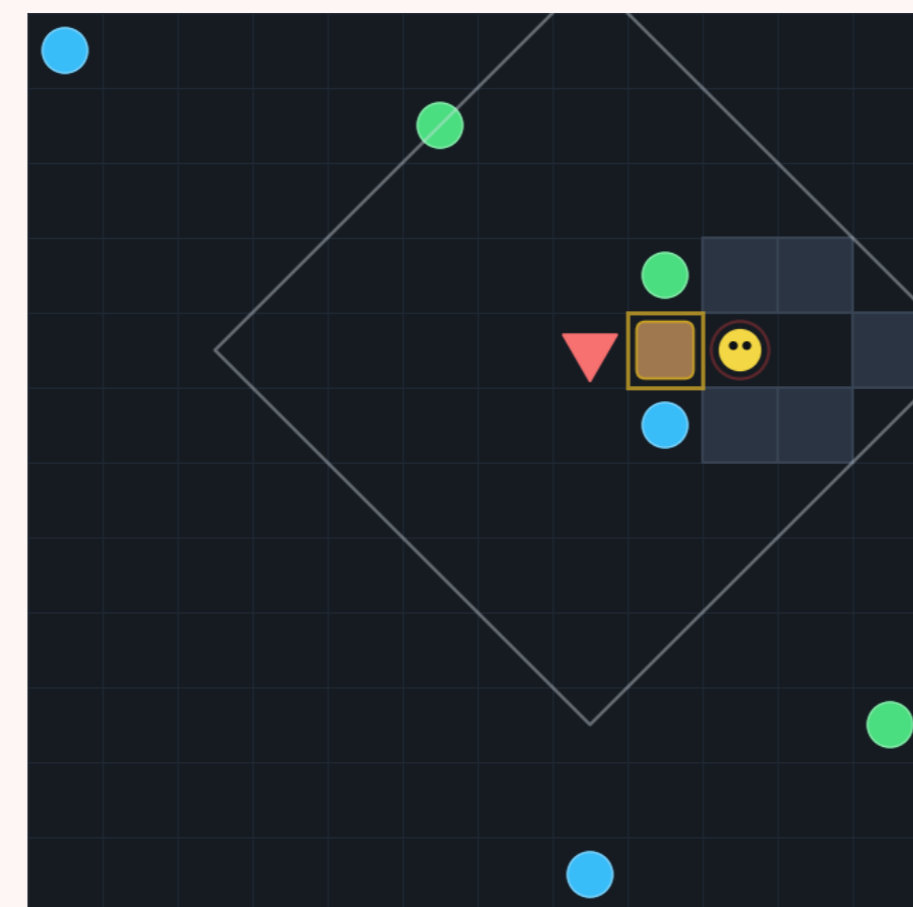
Planning uses the world's true dynamics (a known MDP), not a learned model, to isolate one variable: simulation of a body that can fail. As a prelude to fully learned world models (e.g. [12]), we show that the inner-world model of the body is learnable.



## Foresight must be deep enough, and it learns to take shelter



**Foresight has to be deep.** A brief glimpse ahead is a trap: the agent freezes and dies of thirst. By 16 steps it is viable, matching a hand-crafted reflex policy.



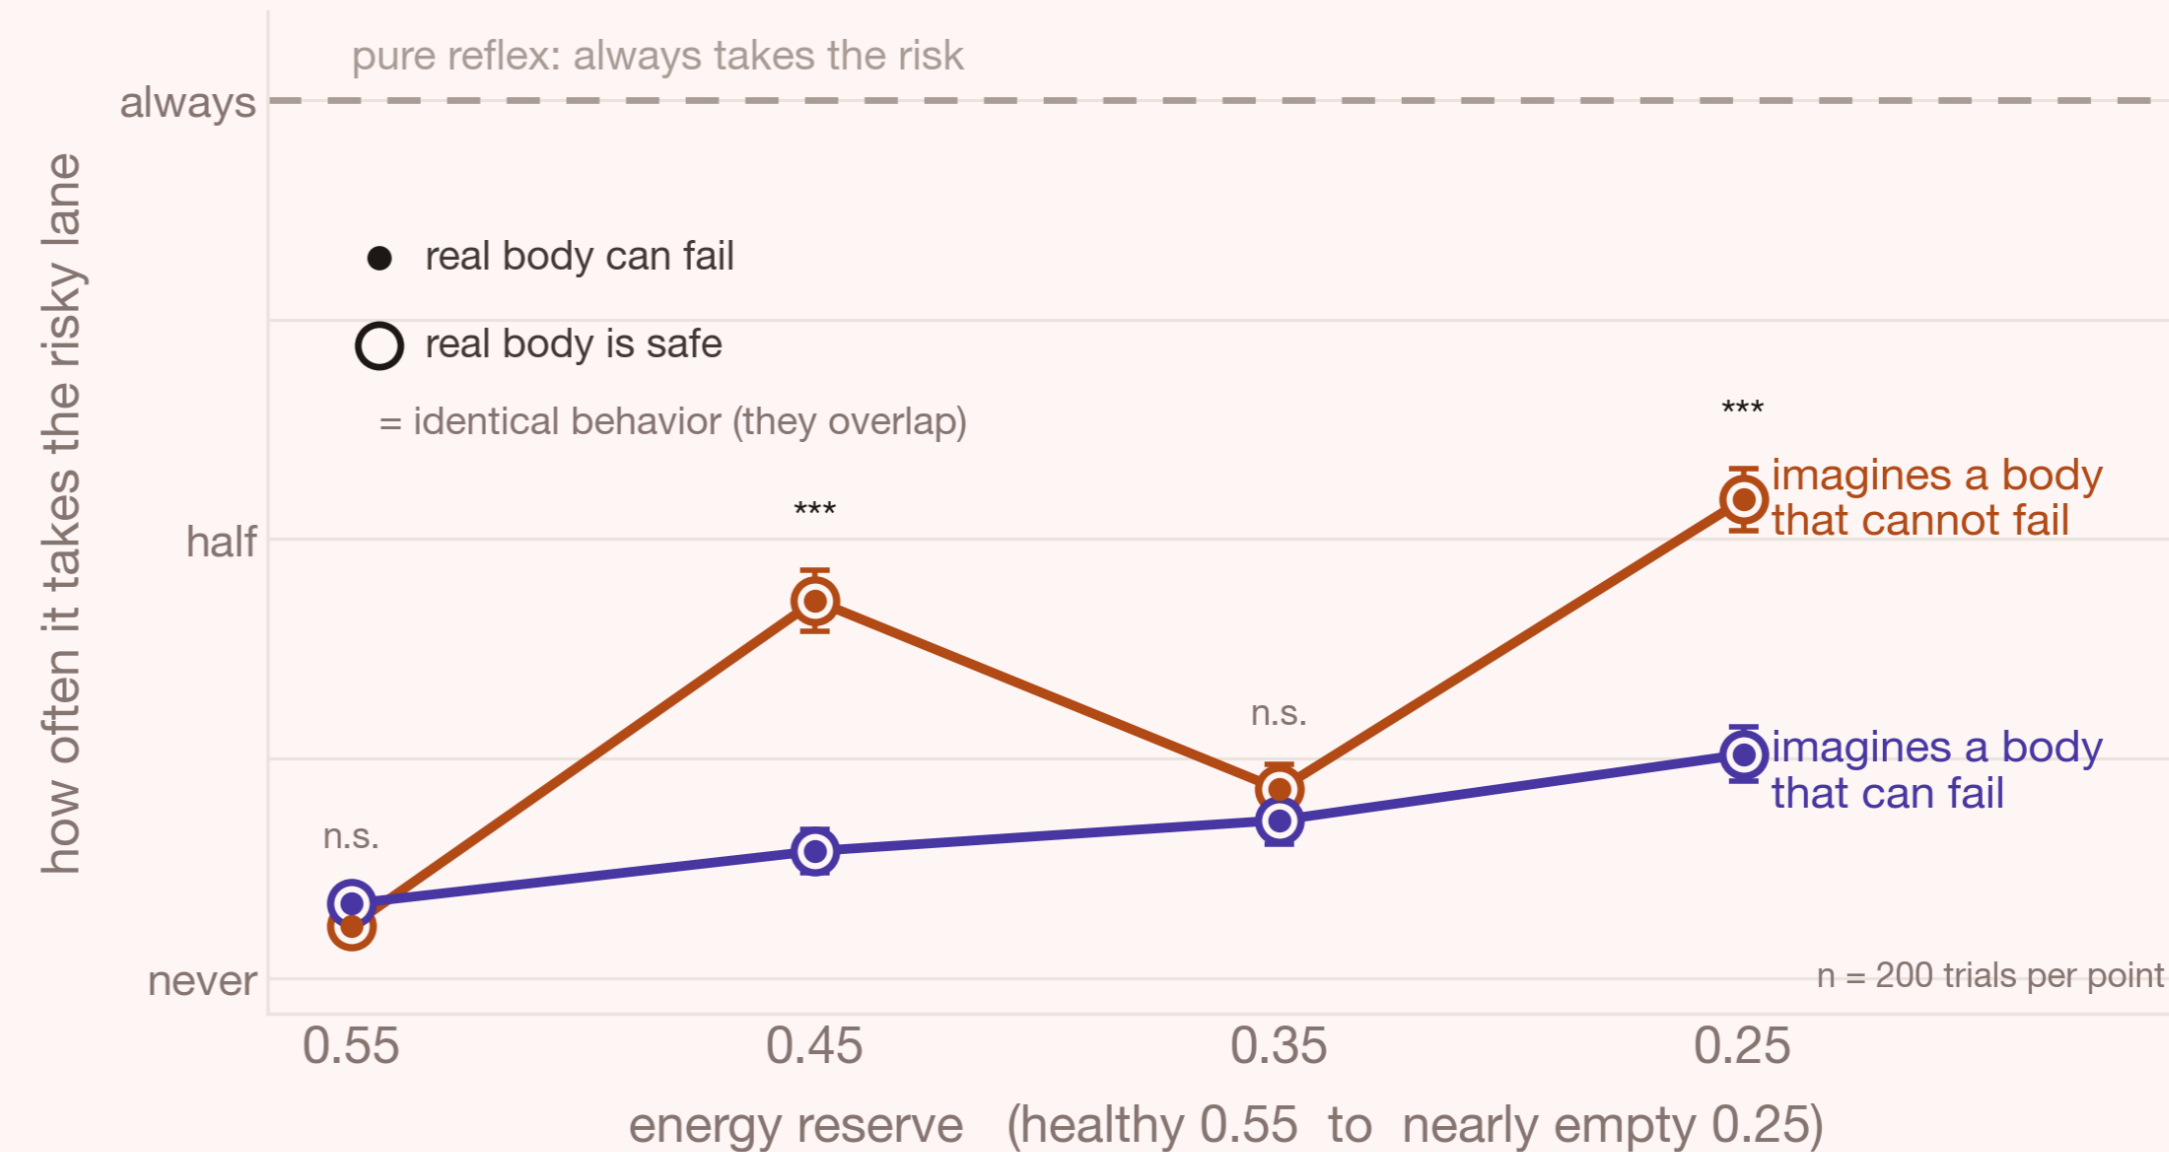
**The agent invents the barricade.** The planner learns to seal itself into a shelter to escape the predator, then ventures out to feed when safe.

## Is it the body or predictions about the body?

We varied two conditions independently: whether the body could fail, and whether the agent believed it could. Caution followed self-beliefs, not necessarily the true body state [7], reminiscent of **anxiety** (a sound body felt as fragile) and of **mania** (a fragile body felt as invincible) [13,14].

### IMAGINED BODY

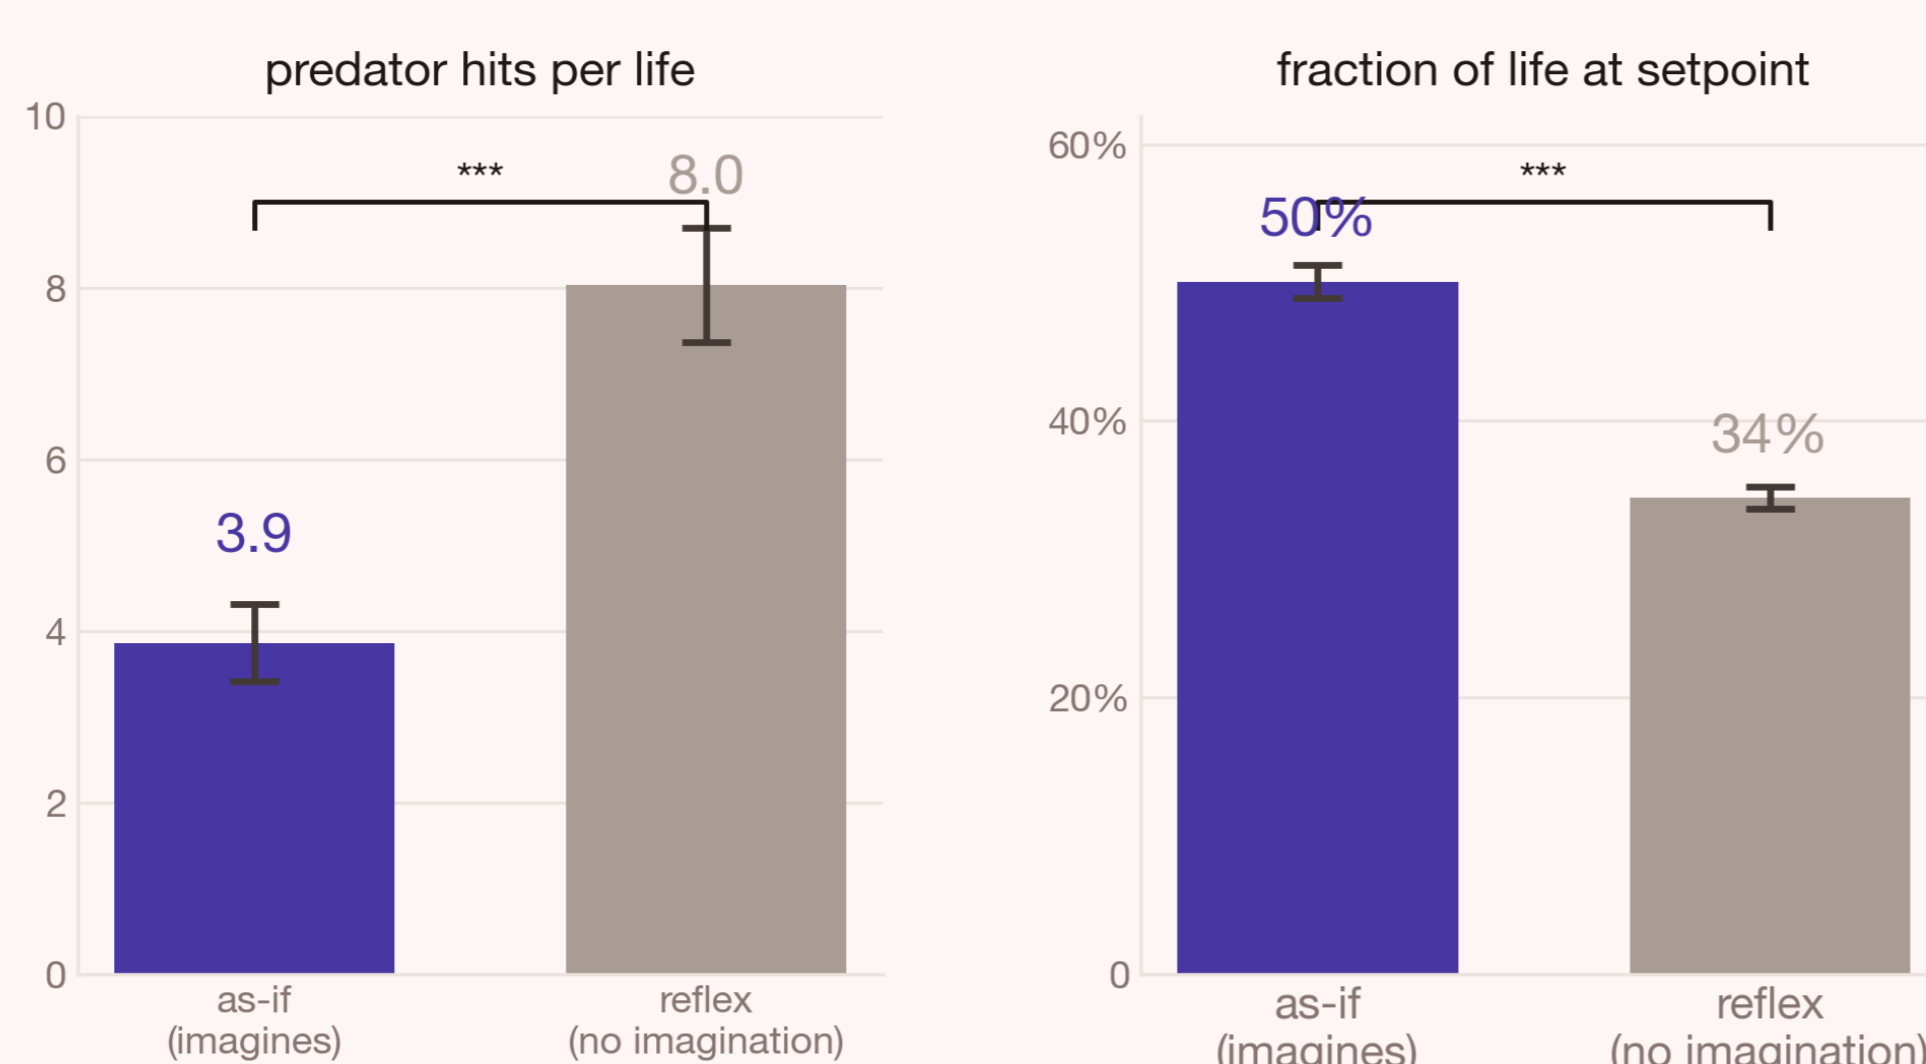
|              | None  | Invulnerable                                | Vulnerable  |
|--------------|---|---|---|
| Vulnerable   | <b>Reflex</b><br>soaks up hits, no foresight  | <b>Daredevil</b><br>reckless, "mania"       | <b>Feeler</b><br>careful, paces, recovers, shelters           |
| Invulnerable | <b>Brute</b><br>unharmable, never looks ahead | <b>Tank</b><br>looks ahead, nothing to fear | <b>Hypochondriac</b><br>cautious over a sound body, "anxiety" |




**Risk-taking is a function of need and the imagined state of need.** When a risky but valuable path is available, the needy agent is more likely to take the risky path. But a daredevil agent that ignores its own vulnerability will even more strongly prefer the risky path.

## Self-simulation as insurance: it pays off under lethal stakes

**When the world is gentle, the premium buys comfort, not time.** The as-if agent and reflex agent live equally long. But foresight reduces predator bites by half and increases time spent at the setpoint by ~50%.



**When wounds turn lethal, the premium buys more time.** The as-if agent has double the lifespan of the reflex agent at high predator lethality.



A body that cannot be harmed has nothing to protect, and so nothing it must act to preserve. Give a homeostatic agent *its own agency to lose*, and behaviors we can recognize as feeling-like will follow. **That is what it means to have skin in the game.**