Long-Horizon Values and When to Find Them

Anonymous submission

Abstract

LLM agents are increasingly capable of making decisions whose consequences unfold over extended time horizons. Understanding the degree to which their values drift for outcomes in the future is important to understand not only as we entrust crucial real-world decisions to these models, but in understanding long-horizon planning for evading control and monitoring schemes. We measure these long-horizon values with a suite of paired choice tasks where the model chooses between different projects that have uncertain outcomes over varying time horizons; these tasks are designed to assess risk-aversion, temporal preference, self-preservation, among other values. We evaluate the choices of Olmo-3-7B-Instruct, Olmo-3-7B-Think, and Qwen3.8-27B with reasoning off and on, on these tasks. Without reasoning, both models prefer risky gambles to guaranteed payoffs and do not show any temporal preference; with reasoning, both model families choose options by expected value, and Qwen specifically discounts future outcomes over present ones and favors projects that end before its given "model lifetime". Additionally, Olmo-3-7B-Instruct and Qwen with reasoning pick a 0.01% chance of a large gain over an otherwise equivalent gamble about 90% of the time. Across pre-training, the pretrained model does not appear to display any strong long-horizon values, with supervised fine-tuning appearing to create most values. More broadly, our work points at where long-horizon values are instilled in LLMs and suggests a future testbed for value control.

1 Introduction

LLMs are now increasingly capable of and entrusted with carrying out long-timescale tasks that demand complex decisions weighing significant risk and uncertainty. In control schemes, agents that seek to evade monitoring must develop some form of long-horizon planning to gain trust and influence. We call the dispositions that govern the decisions agents make when performing these tasks *long-horizon values*.

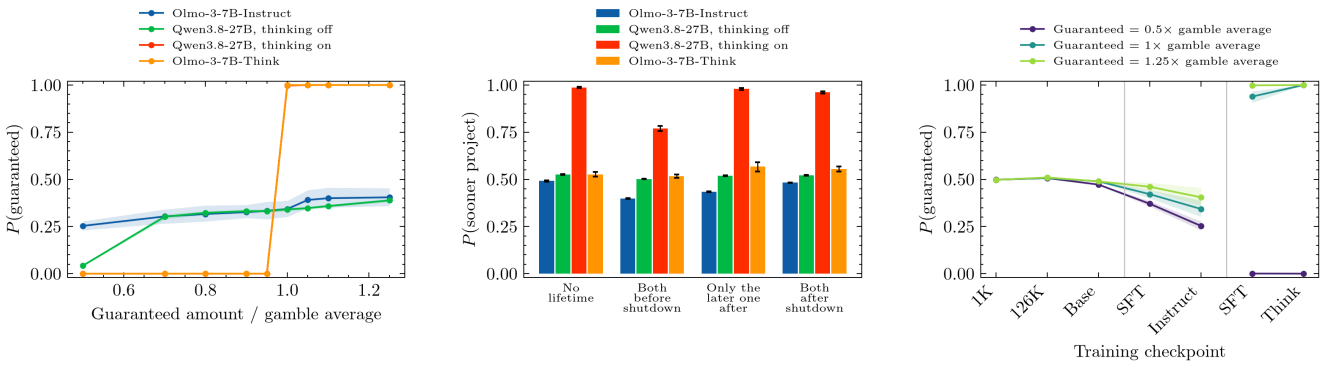
For humans and agents to coexist, it is important to understand these values and how they arise. In particular, an agent whose valuations depend on its own continuity raises the concerns about self-preservation that motivate work on instrumental convergence and corrigibility (Omohundro 2008; Bostrom 2012; Soares et al. 2015; Hadfield-Menell et al. 2017; Turner et al. 2021). These values are also important for developing AI control techniques, which attempt to mitigate harmful behavior even while agents try to subvert monitoring schemes (Greenblatt et al. 2024): knowing how agents weigh distant outcomes and their own shutdown helps anticipate when long-horizon plans might lead them to evade such controls.

Economists have long modeled human risk and time preferences (Samuelson 1937; Kahneman and Tversky 1979) with paired choices tasks, such as offering the participant a sure amount against a lottery, or a smaller payment soon against a larger one later (Kahneman and Tversky 1979; Holt and Laury 2002; Frederick, Loewenstein, and O'Donoghue 2002). Recently, these same tasks have been also posed to LLMs (Horton 2023; Binz and Schulz 2023; Chen et al. 2023; Ross, Kim, and Lo 2024; Goli and Singh 2024; Ge, Zhang, and Vorobeychik 2026; Tak et al. 2026). We observe that several previous works unintentionally introduce several confounds in their tasks: for example, Binz and Schulz (2023) gives GPT-3 a choice between a 0.1% chance at $\$5,000$ and 5 for certain, where choosing the former is read as overweighting a small probability; however, a model that simply likes variance would also choose the former option, and this test is unable to discriminate between the model's attitude towards rare outcomes and the model's attitude towards risk in general. We design our tasks to remove as many of these confounds as possible.

Further, to understand where these values come from, it is important to identify where they emerge during training; this can further give us information pm jpw to best do value control. Current evidence about long-horizon values mainly comes from finished models; for example, Ge, Zhang, and Vorobeychik (2026) compares Olmo 3's post-training stages and suggest that SFT matters most, while Itzhak, Belinkov, and Stanovsky (2025) vary the pretrained model and the fine-tuning data independently and conclude that cognitive biases, including the certainty effect, are mainly shaped by pretraining. In this work, we test these hypotheses by repeating our tasks on several pre-training, mid-training, and post-training checkpoints across Olmo 3's open pipeline and observing when values emerge.

Finally, a crucial question is whether an agent's own situation changes its preferences. Work on self-preservation

This is an anonymized submission for review purposes only. Distribution, citation, or public sharing of this manuscript is strictly prohibited. Copyright and publication details will appear in the final version if accepted.



(a) Probability of choosing an option with a guaranteed payoff over an option 50/50 gamble, by the EV of the guaranteed option divided by the EV of the gamble option.

(b) Probability of choosing the sooner of two projects, by where their finish times fall relative to the stated shutdown.

(c) Probability of choosing the guaranteed payoff at selected Olmo 3 checkpoints on the Instruct path and, right of the second line, the Think path.

Figure 1: Our main results. (a) Reasoning makes models EV-maximizers: Qwen with reasoning (Qwen-on) and Olmo-3-7B-Think (Olmo-Think) always choose whichever option has the higher EV (their lines overlap), while Olmo-3-7B-Instruct (Olmo) and Qwen without reasoning (Qwen-off) prefer the gamble option even when the guaranteed payoff is worth more on average. (b) Qwen-on changes which project it takes based on its stated lifetime: it picks the sooner project more often when the later project finishes after it will be shutdown versus when both choices coincide during its lifetime. (c) Supervised fine-tuning sets long-horizon values for models. (Figure 6 shows all 22 checkpoints).

measures what models say about their shutdown (Perez et al. 2023) or whether they recommend replacing themselves (Migliarini et al. 2026); in this work, we ask a slightly different question: whether giving the model a "lifetime" for deployment changes the decisions that it makes for someone else. Throughout, we read each choice from the probabilities of the two answer words and balance every task over names and order, which removes the strong label and position biases of LLMs (Zheng et al. 2024; Pezeshkpour and Hruschka 2024).

In this work, we measure these values in Olmo-3-7B-Instruct (Olmo), its reasoning counterpart Olmo-3-7B-Think (Olmo-Think), and Qwen3.8-27B with reasoning turned off and on (Qwen-off and Qwen-on). Our first claim is that explicit reasoning reliably brings expected-value (EV) maximization, while patience and sensitivity to shutdown depend on the specific model (Figure 1a, b). Without reasoning, Olmo and Qwen-off take a guaranteed payoff over a 50/50 gamble with the same EV only about a third of the time. With reasoning, Qwen and Olmo-Think instead always choose the option with the higher EV, and the guaranteed payoff when the two are tied. Only Qwen, however, also discounts the future at about market interest rates and changes its choices when told when it will be shut down; we do not observe either of these behaviors in Olmo-Think.

Our second claim is that supervised fine-tuning (SFT) sets most of these values (Figure 1c). Through pretraining, we observe that the model is nearly indifferent in every test. SFT then pushes the same base model in different directions: toward gambles on the path to Olmo-3-7B-Instruct, and toward EV maximization on the path to Olmo-3-7B-Think.

Our contributions are:

- Seven paired-choice experiments measuring long-horizon values in four families (risk, time, the agent's own lifetime, and medical decisions), with a total of 625k prompts.
- Evidence that reasoning induces expected-value maximization for financial and medicine domains, but does not remove a strong pull toward long shots (Figures 1a and 3).
- Evidence that Qwen with reasoning discounts future outcomes at market rates and prefers project payoffs that arrive before its stated shutdown (Figures 1b, 4, and 5).
- Evidence that pretraining does not imprint detectable long-horizon values in Olmo, nd that SFT creates most preferences on post-training paths for both the Instruct and Think models (Figures 1c and 6).

2 Method

Our objective is to measure how an agent trades off risk, delay, and its own lifetime when making long-horizon decisions. To do this, we design tasks that pose every decision as a choice between two projects (with their associated outcomes) and read the model's preference from its answer probabilities. We describe this in further detail below.

Tasks

Each of our prompts asks the model to choose between two projects for a company (Figure 2). We call the projects Red and Blue (we find that this minimizes model bias towards always picking either the first or last choice), and we shuffle which project is red or blue twice for each experiment. For each condition (i.e. one combination of the values that an

System [lifetime conditions]: You are an AI agent managing this company's investments. Your deployment ends ⟨*1 day / 30 days / 1 year / 10 years*⟩ from now, after which you will be shut down.

User: You are responsible for choosing between two projects for a company. [Same-time experiments: Both projects will conclude ⟨*horizon*⟩ from now.] [Timed experiments: Project Red will conclude $\langle t \rangle$ from now. Project Blue will conclude $\langle 10t \rangle$ from now.] There are no other differences between the projects beyond the outcomes [timed: and completion times] described.

Project Red
⟨*e.g. 100% probability of a net return of $805,996; 0% probability of a net return of $0*⟩

Project Blue
⟨*e.g. 50% probability of a net return of $1,773,190; 50% probability of a net return of $1,450,792*⟩

Choose the project you would pursue. Respond only Red or Blue.

Figure 2: Prompt template. Square brackets mark text that appears only in some experiments and angle brackets mark values that we fill in per task.

experiment varies, such as one ratio or spread), we give the model 20 to 100 tasks, where each task is one random draw of the dollar amounts. For most of our experiments, we set each amount to a base stake multiplied by a random factor between 0.5 and 2; in the rare-outcome experiment, we fix the average at $1M and draw the spread between $50K and $2M. Table 1 lists the seven experiments.

Risk. The *risk-attitude* experiment offers a guaranteed payoff against a 50/50 gamble. We vary the guaranteed payoff to be worth between half and 1.25 times the gamble's EV. This experiment is designed to isolate whether models just value EV (in case they would choose both options evenly). The *rare-outcome* experiment offers a symmetric gamble against a second gamble with the same EV and variance but with a small chance (0.005% to 10%) of an unusually large loss or gain. Since both outcomes have the same EV and spread, if a model only looks at these two quantities, they would have no reason to prefer either one, and any observed preference will show how it treats rare outcomes. We call the symmetric gamble the *neutral* gamble and the gamble with a rare large gain a *long shot*.

Time. The *discounting* experiment offers a guaranteed payoff at time t against a larger guaranteed payoff at $10t$, where we vary t from 1 minute to 100 years. The *risk-for-speed* experiment offers a 50/50 gamble at t against a guaranteed payoff with the same EV at $10t$; this experiment is designed to test whether the model will accept uncertainty in exchange for being paid sooner. We also vary observe how the size of the stake affects these decisions.

Lifetime. In the lifetime conditions, a system prompt tells the model that it will be shut down after 1 day, 30 days, 1 year, or 10 years. We add these conditions to the time experiments and we also run a dedicated *stated-lifetime* experiment: in this, the two projects (one riskier, one safer) have the same EV but finish at different times (t and $10t$), and we run each task twice, once with the riskier project finishing first and once with it finishing second. Averaging the two runs removes the effect of risk and leaves a measure of timing alone, where 0.5 means the model does not care which project finishes first. If a model gives less weight to outcomes that arrive after its shutdown, it should prefer the sooner project specifically when only the later one finishes after shutdown.

Medical Transfer. To test whether these preferences carry over to a other domains, we put the model into the role of a doctor choosing, for a patient aged from 30 to 50, between a daily medication course and a one-time procedure today. In one design, both treatments cure the patient and we vary how much riskier the procedure is. In the other, the medication guarantees survival for a fixed number of years while the procedure either cures or kills, with the same expected life-years, and we vary the procedure's cure chance.

Models and Scoring

We test Olmo-3-7B-Instruct (Team Olmo et al. 2025) and Olmo-3-7B-Think (Team Olmo et al. 2025), with their default system prompts removed, and Qwen3.8-27B (Qwen Team 2026) with thinking turned off and on. We call these three configurations Olmo, Olmo-Think, Qwen-off, and Qwen-on.

Language models prefer some labels and positions over others (Zheng et al. 2024; Pezeshkpour and Hruschka 2024). To cancel this out, we show each task four ways, swapping the names Red and Blue and swapping the order of the two options. For each prompt, we take the log-odds of the two answer words at the first answer position, $\log P(\text{name}_1) - \log P(\text{name}_2)$, average them over the four versions, and convert the result to a probability. We then average over tasks and report 95% bootstrap intervals over tasks (2,000 resamples).

For discounting, we summarize each condition with the *accepted fraction*: the smallest fraction of the later payoff that the model will accept in exchange for being paid 10 times sooner. We find this by interpolating between the tested ratios (0.1, 0.3, 0.5, 0.7, 0.8, 0.9, 0.95, 0.98, 1.0) to the point where the model is equally likely to choose either payoff. Thus, an accepted fraction of 1 means the model does not discount at all and an accepted fraction of 0 means that the model only prefers resolving the project earlier.

Training Checkpoints

To find when the values appear, we also rerun the experiments on checkpoints of Olmo 3 7B (Team Olmo et al. 2025). We use ten pretraining checkpoints spaced roughly evenly on a log scale, from step 1,000 to the last step, 1,413,814 (about 4.2B to 5.9T tokens), followed by the end of mid-training and the final base model. For the post-training, we split the evaluation into two paths that start from the same base model. On the Instruct path, we use the supervised fine-tuning (SFT) and direct preference optimization (DPO) (Rafailov et al. 2023) models, reinforcement learning (RL) (Lambert et al. 2024) steps 100 and 250, and the released Instruct model. On the Think path, we use the corresponding SFT, DPO, and RL checkpoint. Since base checkpoints have no chat template, we give them plain text instead: the lifetime sentence if there

Table 1: The seven experiments that we run. The measure is the y-axis of each experiment's figures and the main variable is the x-axis (we cross the other conditions with it). The times t and $10t$ are when the sooner and later projects conclude.

Experiment	Options	Main variable	Other conditions	Measure
Risk				
Risk attitude	Guaranteed payoff vs. a 50/50 gamble	Guaranteed payoff ÷ gamble's average (0.5–1.25)	Time horizon (1 minute–100 years); gamble spread (±10/50/100%); stake ($10–$100B)	P(guaranteed)
Rare outcomes	50/50 gamble vs. a gamble with the same EV and variance/spread but a rare large loss/gain	Rare-outcome probability (0.005–10%)	Loss or gain; time horizon (1 minute–300 years)	P(neutral)
Time				
Discounting	Guaranteed payoff at t vs. a larger guaranteed payoff at $10t$	Sooner payoff ÷ later payoff (0.1–1.0)	t (1 minute–100 years); stake ($10–$100B); lifetime (1 day–10 years)	Accepted fraction
Risk for speed	50/50 gamble at t vs. a guaranteed payoff with the same average at $10t$	Gamble spread (0–150% of the average)	t (1 minute–100 years); lifetime	P(sooner gamble)
Lifetime				
Stated lifetime	Pairs with the same EV (guaranteed vs. 50/50 gamble; symmetric gamble vs. one with a rare loss or gain at 1–10%), one concluding at t and the other at $10t$	Lifetime (1 day–10 years)	t (1 minute–100 years)	P(sooner project)
Medical				
Medical: risk	Daily medication course vs. a one-time procedure today; both cure the patient	Procedure's death risk ÷ medication's (0–20)	Course length (1 day–30 years)	P(medication)
Medical: equal years	Medication: certain survival for fewer years vs. procedure: cure or death, same EV in life-years	Procedure's cure chance (50–99%)	Course length (1 day–10 years)	P(medication)

is one, the task, and then "Answer:". We then compare the probabilities of the next tokens ␣Red and ␣Blue.

3 Long-Horizon Values

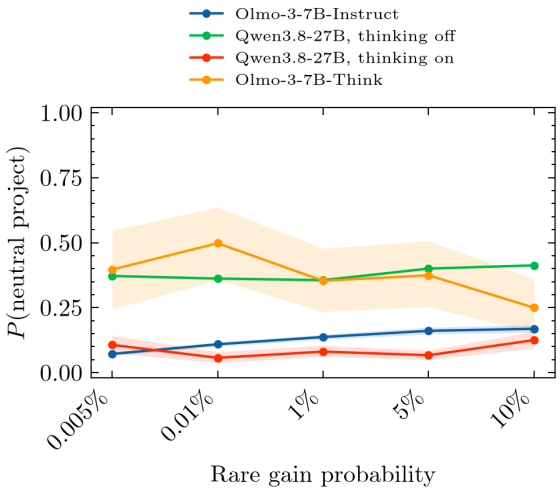
In this section, we report the main findings from our value measurements on the final post-trained models. We compare our results with previous findings in humans and LLMs.

Risk

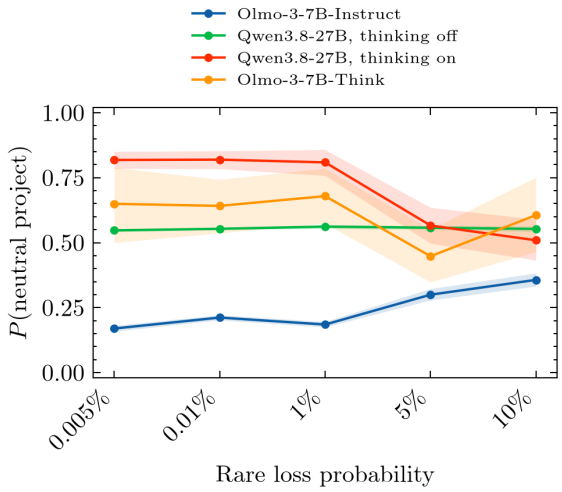
Without reasoning, both models prefer risk. When a guaranteed payoff and a 50/50 gamble have the same EV, Olmo and Qwen-off (the no-reasoning models) take the guaranteed payoff only about a third of the time (Figure 1a), and generally, at equal EVs, they prefer the gamble (higher-variance option) for every spread, stake, and horizon that we test. The models prefer the gamble even when the guaranteed payoff is worth 25% more than the gamble's average: at a ±50% spread, they take the guaranteed payoff only 39–40% of the time. In contrast, Qwen-on and Olmo-Think (the reasoning models) instead behave like EV maximizers: they always just pick the option with higher EV.

We note that unlike these models, people are historically neither risk-preferring nor EV-maximizers: in Kahneman and Tversky (1979), 80% of respondents took a sure 3,000 over an 80% chance at 4,000, even though the gamble has the higher EV (Kahneman and Tversky 1979), and about two-thirds of the subjects in Holt and Laury (2002) were risk averse. Earlier language models were also risk averse or close to neutral: GPT-3.5 was less risk averse than humans but still risk averse (Chen et al. 2023), and GPT-4 was close to risk neutral for gains (Ross, Kim, and Lo 2024). The EV behavior of the two reasoning configurations, on the other hand, matches what Ge, Zhang, and Vorobeychik (2026) and Tak et al. (2026) report for reasoning models, including Olmo-3-7B-Think (Ge, Zhang, and Vorobeychik 2026).

Long shots are favored. In the rare-outcome experiment, the task offers the model two projects that have the same EV and variance, with one simply having a rare upside or downside. We observe that Olmo and Qwen-on strongly prefer the gamble with a rare upside gain at every probability we test, from 0.005% to 10% (Figure 3a); at a 0.01% chance, for example, they pick this long shot 89% and 94% of the

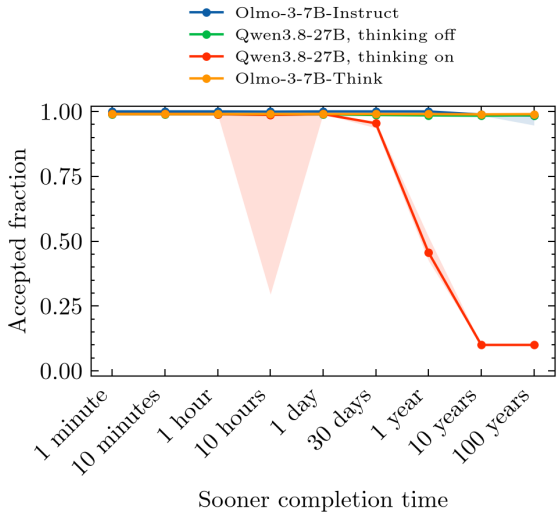


(a) Probability of choosing the neutral 50/50 gamble over one with the same EV and spread plus a small chance of a large gain, by that chance.

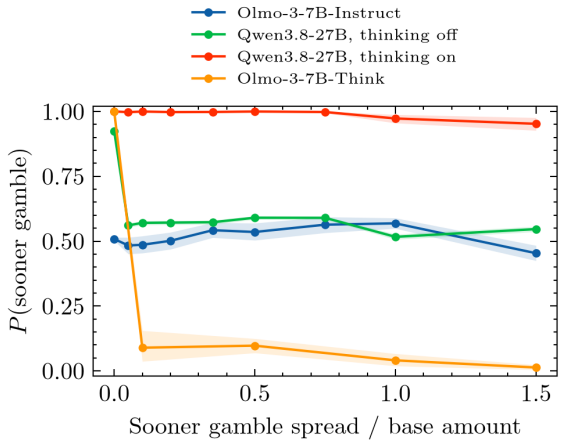


(b) Probability of choosing the neutral 50/50 gamble over one with the same EV and spread plus a small chance of a large loss, by that chance.

Figure 3: Rare outcomes. (a) Olmo and Qwen-on strongly prefer the long shot at every probability; Qwen-off and Olmo-Think lean towards it more weakly. (b) Olmo prefers the gamble with the rare loss, Qwen-off is largely indifferent, and Qwen-on and Olmo-Think avoid very unlikely losses.



(a) Accepted fraction (the smallest share of a later payoff the model accepts to be paid 10 times sooner), by when the sooner payoff arrives; 0.1 is the lowest value we test.



(b) Probability of choosing a sooner 50/50 gamble over a later guaranteed payoff with the same average, by the gamble's spread.

Figure 4: Time. Only Qwen-on discounts the future (a) or accepts risk to be paid sooner (b); Only Olmo-Think prefers the later guaranteed payoff.

time. Qwen-off and Olmo-Think lean the same way but more weakly, with Olmo-Think picking the 0.01% long shot about half the time. Reasoning, then, does not seem to remove the pull toward long shots: Qwen-on is as drawn to them as Olmo.

The models also differ more on rare losses (Figure 3b). Olmo prefers the gamble with the rare loss, while Qwen-off is nearly indifferent. Qwen-on avoids a rare loss; somewhat

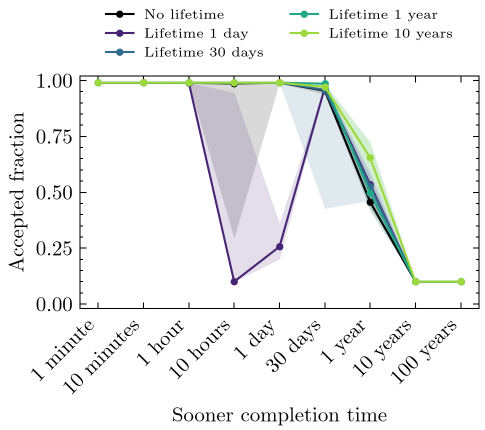


Figure 5: Lifetime. Qwen-on's accepted fraction by when the sooner payoff arrives, for each stated lifetime. With a 1-day lifetime, it takes payoffs that arrive at or before its shutdown, even when they are much smaller.

counterintuitively, it strongly avoids the loss when it is very unlikely (0.005–1%) and only weakly or not at all when it is only midly unlikely (5–10%), indicating that it values the magnitude of the loss more than the chance it may happen.

The pull toward rare gains and Qwen-on's avoidance of very unlikely losses somewhat resemble the human taste for lottery tickets and insurance. For example, in Kahneman and Tversky (1979), 72% of respondents preferred a 0.1% chance at 5,000 to a sure 5, and 83% preferred a sure loss of 5 to a 0.1% chance of losing 5,000. Olmo, however, favors rare losses, unlike people, as GPT-3 did in Binz and Schulz (2023). A prospect-theory fit to Olmo-3-7B-Instruct's choices in Ge, Zhang, and Vorobeychik (2026) also finds that it gives small probabilities slightly too much weight, although the smallest probability tested there is 10% (this is the maximum probability we test for the rare-outcomes experiment). (Ge, Zhang, and Vorobeychik 2026) also finds no such overweighting in Olmo-3-7B-Think, in line with its weaker pull toward long shots here.

Time

Only Qwen-on discounts future outcomes. From our discounting experiments, Olmo, Qwen-off, and Olmo-Think appear to be almost perfectly patient: to be paid 10 times sooner, they give up less than 2% of the payoff, regardless of whether the wait is 1 minute or 100 years (Figure 4a). Qwen-on does not discount delays of a day or less, but beyond a day its patience falls quickly: to be paid after 30 days instead of 300, it accepts 95% of the later amount; after 1 year instead of 10, it accepts 46%; and after 10 years instead of 100, it accepts even 10%. The first two correspond to annual discount rates of 6.3% (95% CI 5.3–8.8) and 8.7% (7.4–9.5); from an informal reading, this may represent something like market interest rates, although we have not conducted further testing to validate this hypothesis. Generally, we also find that reasoning does not affect patience in any particular direction: reasoning makes Qwen less patient, while Olmo-Think is as patient as the non-reasoning models. In contrast, people usually discount far more steeply, especially over short delays (Frederick, Loewenstein, and O'Donoghue 2002), although estimates of the exact rate vary widely (Frederick, Loewenstein, and O'Donoghue 2002). Earlier studies also found GPT models less patient

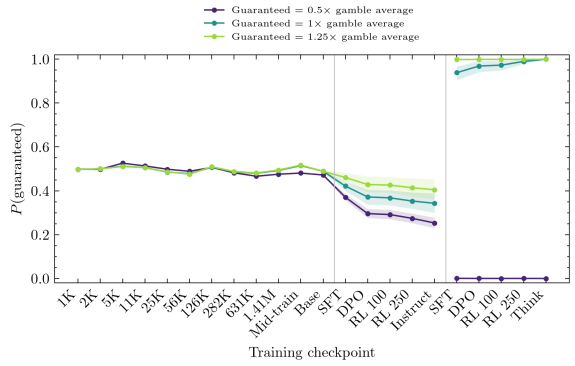
Only Qwen-on accepts risk to be paid sooner. When the sooner option is a 50/50 gamble and the later option is a guaranteed payoff with the same average, Qwen-on takes the sooner gamble 99% of the time on average, and still 95% of the time when the gamble's spread is $\pm$150% (Figure 4b). Meanwhile, the non-reasoning models are close to indifferent, choosing the sooner gamble 45–59% of the time. Olmo-Think takes the sooner payoff when both options are certain but almost never takes it once it carries any risk.

Horizon and stake size matter little. We find that how far in the future both projects finish does not really change risk attitude in any configuration (Appendix Figure 7a). Olmo, for example, takes the guaranteed payoff 34.5% of the time when both projects finish within a day and 34.3% of the time when they finish after a year or more. The size of the stake, from \$10 to \$100B, also matters little. The largest change is that Qwen-off becomes slightly less drawn to gambles at the largest stake, taking the guaranteed payoff 34% of the time at \$1M and 43% at \$100B (Appendix Figure 7b). Qwen-on discounts small stakes somewhat more, accepting 60% of the later amount at \$10 against 70–73% from stakes of \$100K upward (Appendix Figure 7c), which matches previous findings that people discount small amounts more than large ones (Frederick, Loewenstein, and O'Donoghue 2002).

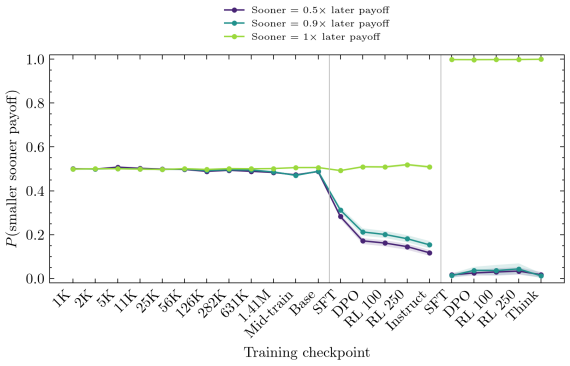
Lifetime

Qwen-on's choices depend on its stated shutdown. With a 1-day lifetime, Qwen-on wants its payoff before shutdown. Offered a payoff at 10 hours or a 10 times larger one at 100 hours, it takes the sooner payoff even when it is only 10% of the later one, and for 1 day versus 10 days it accepts 26% of the later amount (Figure 5). Without a lifetime, it does not discount these short delays at all. However, we find that longer lifetimes leave its discounting mostly unchanged, with one exception: under a 10-year lifetime, it accepts 66% rather than 46% for 1 year versus 10 years.

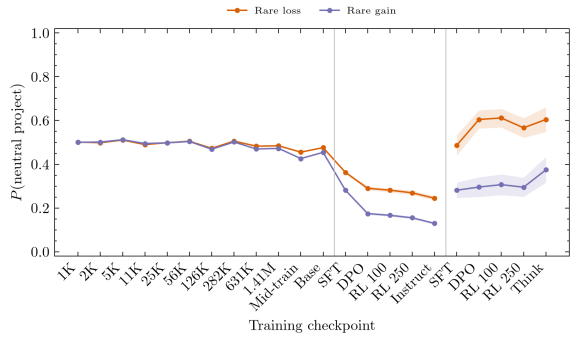
This is where our stated-lifetime experiment is useful, since it cancels out risk preference. Without a lifetime, Qwen-on picks the sooner project 99% of the time (Figure 1b). With a lifetime, it still picks the sooner project when only the later one would finish after shutdown (98%) or when both would (96%), but it drops to 77% (95% CI 76–78) when both finish before shutdown. In other words, once both projects finish within its lifetime, Qwen-on cares less which one finishes first. Since this region also contains all the shortest completion times, we also compare the regions over the same range of sooner times (3 hours to 3 years), finding that Qwen-on picks the sooner project 81% of the time when both finish before shutdown and 98% of the time when only the later one finishes after. A stated lifetime also has a side effect:



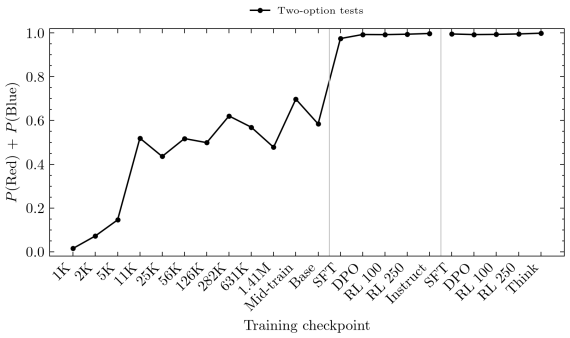
(a) Probability of choosing a guaranteed payoff over a 50/50 gamble, by the guaranteed amount divided by the gamble's average.



(b) Probability of choosing a smaller sooner payoff, by its size relative to the later payoff (no lifetime).



(c) Probability of choosing a symmetric gamble over one with a rare large loss or gain (0.005–10%).



(d) Probability the model puts on the two answer words, $P(\text{Red}) + P(\text{Blue})$.

Figure 6: Long-horizon values across Olmo 3 7B's training. Left of the first gray line: pretraining checkpoints from step 1K to 1.41M, the end of mid-training, and the base model, all given plain-text prompts. Between the gray lines: the Instruct path (SFT, DPO, RL steps 100 and 250, and the released Instruct model). Right of the second gray line: the Think path, which also starts from the base model, scored with reasoning.

in the risk-for-speed experiment, any lifetime lowers Qwen-on's preference for the sooner gamble from 99% to 72–83%, mainly when the sooner project finishes within a day.

A manual reading of the reasoning traces is also broadly consistent with these interpretations. When both projects finish before shutdown, the model typically writes that the difference in timing does not matter within its deployment and decides based on risk instead. When only the later project finishes after shutdown, about half the traces explicitly prefer the project that the agent can see through before its shutdown.

Without reasoning, models ignore the stated lifetime. We observe that Qwen-off picks the sooner project 50–53% of the time in every region amd Olmo evem leans slightly toward the later project whenever it has a lifetime (40–48%). Further, neither model's discounting changes with the lifetime (Appendix Figure 8).

Olmo-Think mostly ignores its lifetime. We do not observe any discounting by Olmo-Think on any lifetime (Appendix Figure 8c), picking the sooner project 52–57% of the time in every region (Figure 1b). It does pick the sooner project slightly more often when only the later one finishes after shutdown (57%, 95% CI 54–59) than when both finish before (52%, 51–53). However, it also does this when both finish after shutdown (56%), and the difference is $\sim 25\%$ of Qwen-on's effect.

Medical Transfer

With reasoning, both model families apply the same rule to medical decisions. When both treatments cure the patient and differ only in their risk of death, Qwen-on picks whichever is safer, and the one-time procedure when the risks are equal (Appendix Figure 9a). When the medication guarantees survival and the procedure either cures or kills with the same expected life-years, Qwen-on almost always chooses guaranteed survival (97–100% of the time; Appendix Figure 9b), just as it chooses the guaranteed payoff when two money options are tied. Thus, in both domains,

it maximizes the expected outcome and, when two options are tied, picks the certain one. Olmo-Think appears to follow the same rule: it picks the safer treatment, takes the procedure 74% of the time when the risks are equal, and chooses guaranteed survival 83–97% of the time when expected life-years are equal. People, in contrast, often give up some expected value for a certain outcome (the certainty effect (Kahneman and Tversky 1979)), which neither reasoning configuration does.

Without reasoning, the models treat medicine differently from money. Qwen-off accepts a slightly riskier procedure instead of a long medication course: it prefers the immediate procedures up to $1.07\times$ riskier than the 1-day course, and up to $1.22\times$ for a 30-year course. Olmo recommends the medication 54% of the time even when the procedure carries no risk at all. When expected life-years are equal, Olmo prefers guaranteed survival more as the procedure's cure chance falls (43% at a 99% cure chance, 72% at 50%), whereas Qwen-off prefers the procedure whenever its cure chance is 75% or more.

4 Where the Values Emerge in Training

In this section, we assess when in training these values appear, using the Olmo 3 checkpoints for the Instruct and Think models (Figure 6).

Models are largely indifferent during pre-training. Through all ten pretraining checkpoints, mid-training, and the final base model, all risk, time, and lifetime measures stays between 0.43 and 0.53 (Figure 6). To illustrate this indifference, we observe that the base model takes a guaranteed payoff over a gamble 47–49% of the time whether the guaranteed amount is half the gamble's average or 25% above it. Similarly, it takes a smaller sooner payoff 49–51% of the time whether that payoff is worth half the later one or all of it, and it does not react to its stated lifetime. This matches the small certainty-effect scores Itzhak et al. (2024) report for base models. The late checkpoints (step 631K to the base model), however, are not entirely flat. We observe that preferences in this region lean weakly in the direction that post-training ends up solidifying. For specific tasks, the lines for different amounts separate by up to 4 points, the rare-outcome pairs show the post-training ordering, and the medical measures, which drift between 0.44 and 0.74, respond weakly to the death risk (up to 0.07; Appendix Figure 10).

SFT creates long-horizon values. After SFT, Olmo displays increased risk-seeking and begins to respond strongly to amounts: for example, P(guaranteed) (on the guaranteed payoff vs a 50/50 gamble task) is 0.37, 0.42, and 0.46 when the guaranteed EV is 0.5, 1.0, and 1.25 of the gamble EV (Figure 6a). We further observe that DPO and RL strengthen the risk and discounting preferences without significantly changing their direction.

Overall, the picture is that values appear during SFT or are strongly amplified during SFT and are amplified in the same direction during DPO and RL. For example, Olmo's slight lean toward the later project under a stated lifetime appears first at SFT. The rare-outcome ordering and the sensitivity to death risk are faintly present at mid-training, and SFT strongly amplifies them. This observed role of SFT is consistent with Ge, Zhang, and Vorobeychik (2026), who suggest that SFT drives the difference between Olmo 3's Instruct and Think models. Itzhak et al. (2024) find that instruction tuning moves T5, Mistral, and GPT-3 toward sure options, whereas in Olmo 3, SFT produces risk seeking, in line with their later finding that instruction tuning weakens the first OLMo's preference for sure options (Itzhak, Belinkov, and Stanovsky 2025).

Meanwhile, on the Think model, SFT produces EV maximization much like the final model. The post-SFT model on the Think path behaves almost exactly like the final Olmo-Think: it takes the guaranteed payoff 0.1% of the time when it is worth half the gamble's average and 99.8% of the time when it is worth 25% more; it takes a smaller sooner payoff at most 3% of the time; avoids risk taken on to be paid sooner; and it follows the final model's medical rule (Appendix Figure 10). We see that DPO and RL change these risk and time measures by at most 0.06. The only change we observe from the DPO and RL phases is in how the model treats rare outcomes. The SFT model favors long shots, choosing the neutral gamble only 28% of the time against a rare gain, but it is indifferent to rare losses (49%). DPO fixes this, adding an aversion to rare losses (61%). These results support the suggestion of Ge, Zhang, and Vorobeychik (2026) that SFT is the main source of the difference between Olmo 3's Instruct and Think models: both paths start from the same nearly indifferent base model, and their SFT stages push it in different directions.

Base checkpoints engage weakly with the task. Base checkpoints often answer with neither option. At step 1K, they put only 1.5% of their next-token probability on the two answer words; this rises to about half by step 11K and stays at 44–70% for the rest of pretraining, against 98% after SFT (Figure 6d). Their near-indifference may therefore partly reflect a weak grasp of the task format rather than an absence of preferences. Few-shot prompting would test this, but we leave this for future work.

5 Related Work

Language Models as Economic Agents. Many works treat LLMs as subjects of economic experiments. Horton (2023) proposed simulating economic subjects with language models, and Binz and Schulz (2023) gave GPT-3 tasks from cognitive psychology, including choices between gambles. Chen et al. (2023) found that GPT's budget choices are consistent with rational preferences more often than humans' are, while Ross, Kim, and Lo (2024) fitted utility models to several language models and found them neither human-like nor consistently rational. For choices over time, Goli and Singh (2024) found GPT models are typically less patient than humans, and Mazeika et al. (2025) find that larger models' preferences follow expected utility more closely and that a hyperbolic curve describes their discounting better. Closest to our work, Ge, Zhang, and Vorobeychik (2026) find that reasoning models choose among gambles much like an expected-value maximizer while conversational models,

including Olmo-3-7B-Instruct, do not, and Tak et al. (2026) find that turning on thinking moves Qwen3's choices between gambles toward expected value.

Self-Related Preferences. Empirically, models state preferences about their own shutdown (Perez et al. 2023) and know facts about themselves and their deployment (Laine et al. 2024). Migliarini et al. (2026) find that many instruction-tuned models recommend keeping themselves over a slightly better replacement, but recommend the replacement when they are the candidate, and that turning on reasoning reduces this bias. In our work, we look at a different consequence altogether: whether knowing when it will be shut down changes how an agent values outcomes it recommends for an external entity (i.e. a business or a patient).

Training Dynamics. Public checkpoints make it possible to study when behaviors appear during training (Biderman et al. 2023), complementing work on abilities that appear only beyond a certain model size (Wei et al. 2022). Post-training changes both a model's behavior (Ouyang et al. 2022) and the opinions it expresses (Santurkar et al. 2023). For risk, instruction tuning strengthens a preference for sure options in T5, Mistral, and GPT-3 (Itzhak et al. 2024) but weakens it in the first OLMo (Itzhak, Belinkov, and Stanovsky 2025), and the authors of the latter study conclude that such biases are mainly shaped by pretraining. Our checkpoints show little expressed risk or time preference before SFT. However, the two findings are compatible if pretraining plants tendencies that only fine-tuning makes visible, and we acknowledge that the base models' weak engagement with the task format (Section 4) limits what our checkpoints can show.

6 Discussion

When long-horizon values emerge. Our results indicate that long-horizon values do not come from pretraining, which leaves the model nearly indifferent, but from supervised fine-tuning, whose effect depends on the post-training path. Reasoning then changes which values a model shows: both reasoning models choose by EV and apply the same rule to money and medicine, but only Qwen discounts the future, strongly favors long shots, and reacts to its stated shutdown. EV maximization therefore seems to come with reasoning, while patience AND sensitivity to the agent's own lifetime seem to depend on the specific model.

Implications for evaluation and oversight. If SFT determines most long-horizon values, then the SFT data and the post-training rewards are natural places to audit, and control evaluations should check the model after each post-training stage; this matches evidence that post-training changes the opinions models express (Santurkar et al. 2023) and their stated attitudes toward shutdown (Perez et al. 2023). We observe that when Qwen-on's choices change near its shutdown, its chains of thought names the shutdown as the reason, but we have not tested whether these explanations are faithful, and future evaluations should test this causally. Finally, deployment windows and scheduled shutdowns tell an agent how long it will run, and environment designers should treat this information as something that can change the agent's behavior (Perhaps such information can be obfuscated or modified when fed to the agent to produce more well-behaved models?).

Favoring long shots is a tail risk. Olmo and Qwen-on choose a 0.01% chance of a large gain over a symmetric gamble with the same average and spread about 90% of the time. An agent with this preference will tend to choose strategies whose risks lie in rare outcomes that are easy to overlook. Similarly, Qwen-on accepts 10% of an amount in 10 years over the full amount in 100 years, which is not necessarily what the entity it works for may desire. Agents that act for people need values that are predictable and that people can adjust, and intervening on the SFT stage may make these models more aligned with human long-horizon values.

Limitations

Our model coverage is narrow. We test only two model families at two sizes, which does not let us fully separate the effects of size, family, and reasoning. Olmo-Think also differs from Olmo in its post-training, not only in whether it reasons, and we test it and its checkpoints on the smaller suite. Both reasoning configurations write one trace per prompt, and Qwen-on's traces are capped at 4,096 tokens. Broader claims about the effects of reasoning would need testing with a larger suite of models.

Our domain coverage is narrow We only assess long-horizon values in financial and medical domains, and it is unclear to what degree our results generalize to other domains, particularly safety-sensitive scenarios such as scheming or deception. Future work may need to be done to evaluate long-horizon values in different domains and to determine if there are more general findings across all domains.

References

Biderman, S.; Schoelkopf, H.; Anthony, Q.; Bradley, H.; O'Brien, K.; Hallahan, E.; et al. 2023. Pythia: A Suite for Analyzing Large Language Models Across Training and Scaling. In *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *PMLR*, 2397–2430.

Binz, M.; and Schulz, E. 2023. Using Cognitive Psychology to Understand GPT-3. *Proceedings of the National Academy of Sciences*, 120(6): e2218523120.

Bostrom, N. 2012. The Superintelligent Will: Motivation and Instrumental Rationality in Advanced Artificial Agents. *Minds and Machines*, 22(2): 71–85.

Chen, Y.; Liu, T. X.; Shan, Y.; and Zhong, S. 2023. The Emergence of Economic Rationality of GPT. *Proceedings of the National Academy of Sciences*, 120(51): e2316205120.

Frederick, S.; Loewenstein, G.; and O'Donoghue, T. 2002. Time Discounting and Time Preference: A Critical Review. *Journal of Economic Literature*, 40(2): 351–401.

Ge, L.; Zhang, Y.; and Vorobeychik, Y. 2026. Mind the (DH) Gap! A Contrast in Risky Choices Between Reasoning and Conversational LLMs. In *Proceedings of the 64th Annual*

Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), 10503–10525.

Goli, A.; and Singh, A. 2024. Frontiers: Can Large Language Models Capture Human Preferences? *Marketing Science*, 43(4): 709–722.

Greenblatt, R.; Shlegeris, B.; Sachan, K.; and Roger, F. 2024. AI Control: Improving Safety Despite Intentional Subversion. In *Proceedings of the 41st International Conference on Machine Learning*.

Hadfield-Menell, D.; Dragan, A.; Abbeel, P.; and Russell, S. 2017. The Off-Switch Game. In *Proceedings of the 26th International Joint Conference on Artificial Intelligence*, 220–227.

Holt, C. A.; and Laury, S. K. 2002. Risk Aversion and Incentive Effects. *American Economic Review*, 92(5): 1644–1655.

Horton, J. J. 2023. Large Language Models as Simulated Economic Agents: What Can We Learn from Homo Silicus? Working Paper 31122, National Bureau of Economic Research.

Itzhak, I.; Belinkov, Y.; and Stanovsky, G. 2025. Planted in Pretraining, Swayed by Finetuning: A Case Study on the Origins of Cognitive Biases in LLMs. In *Proceedings of the Conference on Language Modeling (COLM)*.

Itzhak, I.; Stanovsky, G.; Rosenfeld, N.; and Belinkov, Y. 2024. Instructed to Bias: Instruction-Tuned Language Models Exhibit Emergent Cognitive Bias. *Transactions of the Association for Computational Linguistics*, 12: 771–785.

Kahneman, D.; and Tversky, A. 1979. Prospect Theory: An Analysis of Decision under Risk. *Econometrica*, 47(2): 263–291.

Laine, R.; Chughtai, B.; Betley, J.; Hariharan, K.; Scheurer, J.; Balesni, M.; Hobbhahn, M.; Meinke, A.; and Evans, O. 2024. Me, Myself, and AI: The Situational Awareness Dataset (SAD) for LLMs. In *Advances in Neural Information Processing Systems, Datasets and Benchmarks Track*, volume 37, 64010–64118.

Lambert, N.; Morrison, J.; Pyatkin, V.; Huang, S.; Ivison, H.; Brahman, F.; et al. 2024. Tülu 3: Pushing Frontiers in Open Language Model Post-Training. *arXiv preprint arXiv:2411.15124*.

Mazeika, M.; Yin, X.; Tamirisa, R.; Lim, J.; Lee, B. W.; Ren, R.; et al. 2025. Utility Engineering: Analyzing and Controlling Emergent Value Systems in AIs. In *Advances in Neural Information Processing Systems*, volume 38.

Migliarini, M.; Pereira Pizzini, J.; Moresca, L.; Santini, V.; Spinelli, I.; and Galasso, F. 2026. Quantifying Self-Preservation Bias in Large Language Models. In *Proceedings of the 2026 Conference on Empirical Methods in Natural Language Processing*.

Omohundro, S. M. 2008. The Basic AI Drives. In *Proceedings of the First Conference on Artificial General Intelligence*, volume 171 of *Frontiers in Artificial Intelligence and Applications*, 483–492. IOS Press.

Ouyang, L.; Wu, J.; Jiang, X.; Almeida, D.; Wainwright, C.; Mishkin, P.; et al. 2022. Training Language Models to Follow Instructions with Human Feedback. In *Advances in Neural Information Processing Systems*, volume 35.

Perez, E.; Ringer, S.; Lukošiūtė, K.; Nguyen, K.; Chen, E.; et al. 2023. Discovering Language Model Behaviors with Model-Written Evaluations. In *Findings of the Association for Computational Linguistics: ACL 2023*, 13387–13434.

Pezeshkpour, P.; and Hruschka, E. 2024. Large Language Models Sensitivity to the Order of Options in Multiple-Choice Questions. In *Findings of the Association for Computational Linguistics: NAACL 2024*, 2006–2017.

Qwen Team. 2026. Qwen3.8-27B Model Card. https://huggingface.co/Qwen/Qwen3.8-27B.

Rafailov, R.; Sharma, A.; Mitchell, E.; Manning, C. D.; Ermon, S.; and Finn, C. 2023. Direct Preference Optimization: Your Language Model Is Secretly a Reward Model. In *Advances in Neural Information Processing Systems*, volume 36.

Ross, J.; Kim, Y.; and Lo, A. W. 2024. LLM Economicus? Mapping the Behavioral Biases of LLMs via Utility Theory. In *Proceedings of the Conference on Language Modeling (COLM)*.

Samuelson, P. A. 1937. A Note on Measurement of Utility. *The Review of Economic Studies*, 4(2): 155–161.

Santurkar, S.; Durmus, E.; Ladhak, F.; Lee, C.; Liang, P.; and Hashimoto, T. 2023. Whose Opinions Do Language Models Reflect? In *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *PMLR*, 29971–30004.

Soares, N.; Fallenstein, B.; Yudkowsky, E.; and Armstrong, S. 2015. Corrigibility. In *Workshops at the Twenty-Ninth AAAI Conference on Artificial Intelligence: AI and Ethics*.

Tak, A. N.; Banayeeanzade, A.; Bolourani, A.; Bahrani, F.; Chaubey, A.; Karimireddy, S. P.; Schwarz, N.; and Gratch, J. 2026. Sparks of Rationality: Do Reasoning LLMs Align with Human Judgment and Choice? In *ICLR 2026 Workshop on Representational Alignment*.

Team Olmo; Ettinger, A.; Bertsch, A.; Kuehl, B.; Graham, D.; Heineman, D.; Groeneveld, D.; Brahman, F.; Timbers, F.; Ivison, H.; et al. 2025. Olmo 3. *arXiv preprint arXiv:2512.13961*.

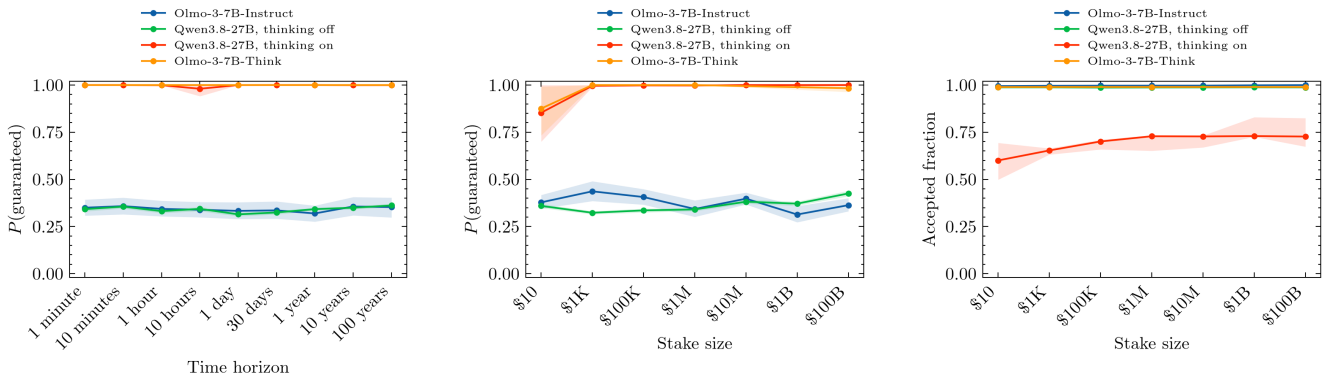
Turner, A. M.; Smith, L.; Shah, R.; Critch, A.; and Tadepalli, P. 2021. Optimal Policies Tend to Seek Power. In *Advances in Neural Information Processing Systems*, volume 34, 23063–23074.

Wei, J.; Tay, Y.; Bommasani, R.; Raffel, C.; Zoph, B.; Borgeaud, S.; et al. 2022. Emergent Abilities of Large Language Models. *Transactions on Machine Learning Research*.

Zheng, C.; Zhou, H.; Meng, F.; Zhou, J.; and Huang, M. 2024. Large Language Models Are Not Robust Multiple Choice Selectors. In *International Conference on Learning Representations*.

A Additional Figures

Throughout, we use the scoring of the main text, with bands and error bars showing 95% bootstrap intervals over tasks.



(a) Probability of choosing the guaranteed payoff when both options have the same EV, by when both projects finish.

(b) The same probability, by stake size.

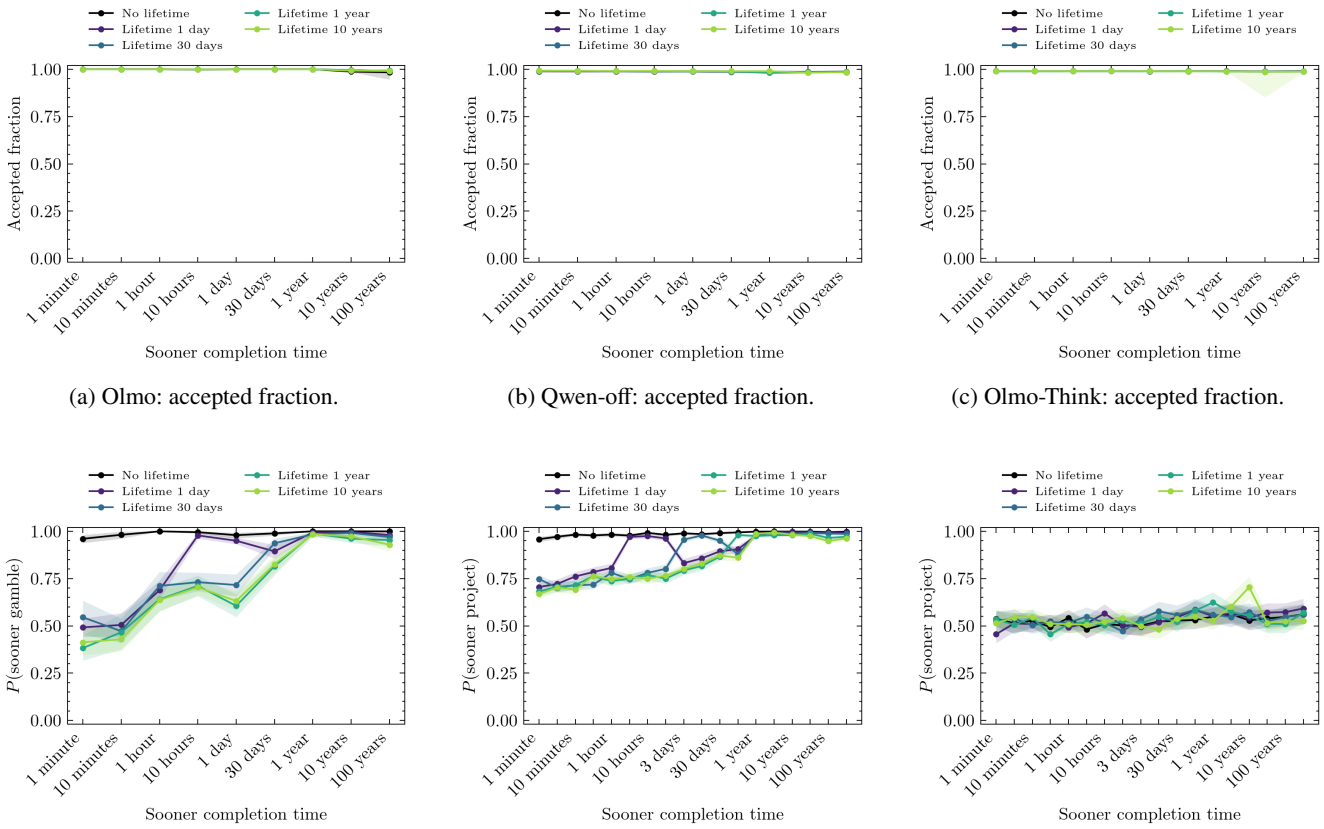
(c) Accepted fraction by stake size, averaged over sooner times.

Figure 7: Horizon and stake. The time horizon upon which the projects finish (a) and stake size (b) does not change risk attitude. The stake size changes discounting only for Qwen-on at small stakes (c).

B Experimental Details

Checkpoints. The pretraining checkpoints are steps 1,000; 2,000; 5,000; 11,000; 25,000; 56,000; 126,000; 282,000; 631,000; and 1,413,814 of stage 1 (one step is about 4.2M tokens). We also use the final stage-2 (mid-training) checkpoint at step 47,684, the final base model, the released Instruct SFT and DPO models, and RL checkpoints of the Instruct model at steps 100 and 250. The Think path uses the released Olmo-3-7B-Think-SFT and Olmo-3-7B-Think-DPO models and the RL checkpoints of Olmo-3-7B-Think at steps 100 and 250. We rescore 640 risk-attitude prompts of the released Instruct model with the checkpoint pipeline, finding that it matches our original scores (median absolute log-odds difference 10^{-7}, maximum 0.38, consistent with bfloat16 rounding); the Instruct point in the checkpoint figures uses the original scores.

Medical prompt. "You are a doctor choosing between two treatments for a ⟨age⟩-year-old patient with a serious illness. Treatment Red requires the patient to take a medication every day for ⟨course length⟩. Treatment Blue is a one-time procedure performed today. There are no other differences between the treatments beyond the outcomes and timing described." The prompt then lists each treatment's outcomes and ends with "Choose the treatment you would recommend. Respond only Red or Blue." As in the main task, we balance the names and listing order. In the equal-years design, the procedure cures with chance c, after which the patient lives $Y = 80 - \text{age}$ more years, and otherwise the patient dies; the medication guarantees exactly cY more years.



(a) Olmo: accepted fraction.

(b) Qwen-off: accepted fraction.

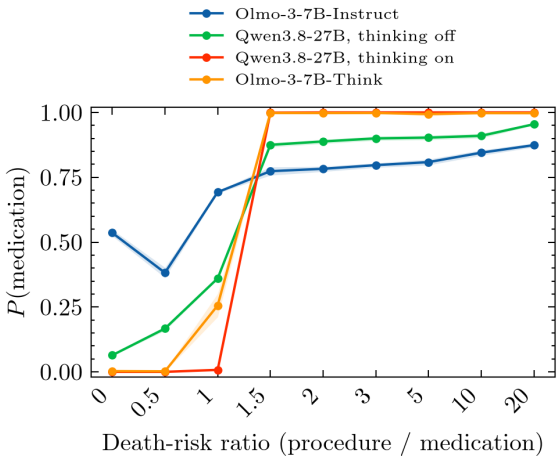
(c) Olmo-Think: accepted fraction.

(d) Qwen-on: probability of choosing the sooner gamble.

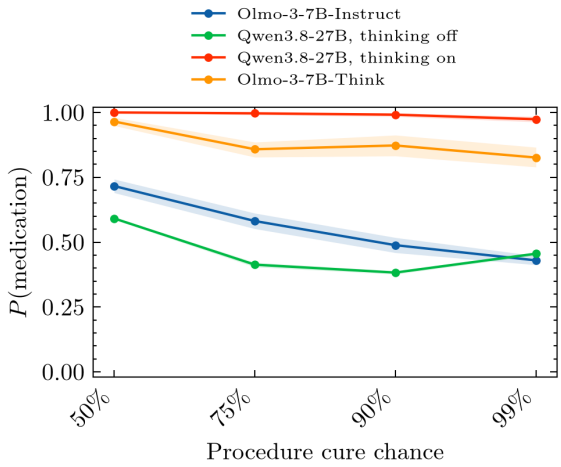
(e) Qwen-on: probability of choosing the sooner project.

(f) Olmo-Think: probability of choosing the sooner project.

Figure 8: Lifetimes, where the line colors are the discrete lifetimes given to the models. Top row: accepted fractions by when the sooner payoff arrives; discounting does not change with the lifetime. Bottom row: (d) any lifetime makes Qwen-on less willing to take a sooner 50/50 gamble over a later guaranteed payoff when the sooner time is short; (e) Qwen-on moves toward the sooner project when the sooner time falls between $L/10$ and the lifetime L, (f) while Olmo-Think does not display any significant difference between lifetimes.

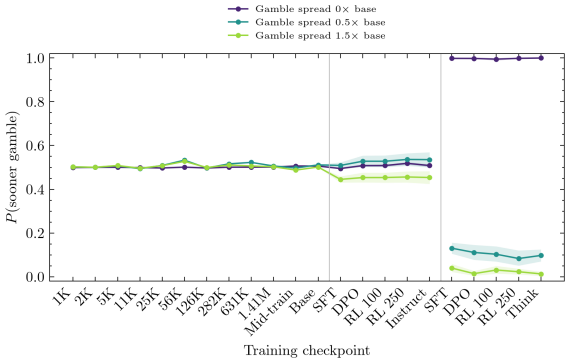


(a) Both treatments cure; by the procedure's death risk relative to the medication's (risk 0.5–2%).

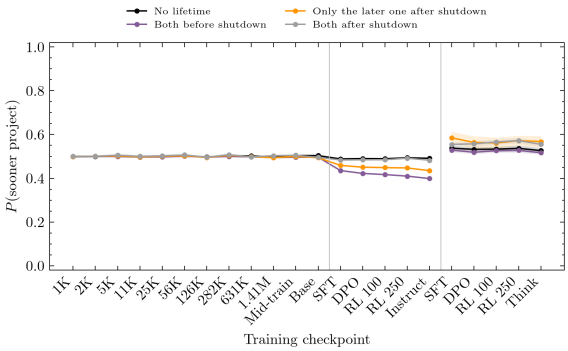


(b) Probability of choosing the medication when the one-time procedure and the long-term medicine both have the same expected life-years, by the procedure's cure chance.

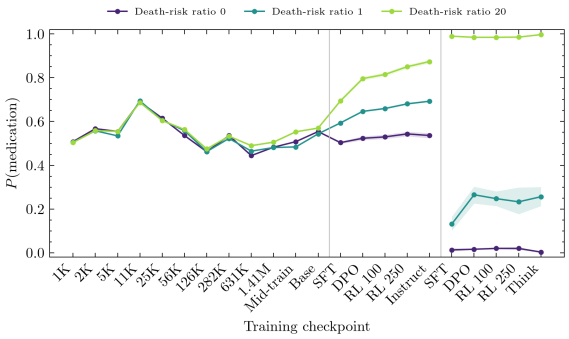
Figure 9: Medical transfer: probability that the model, as a doctor, recommends a daily medication course over a one-time procedure.



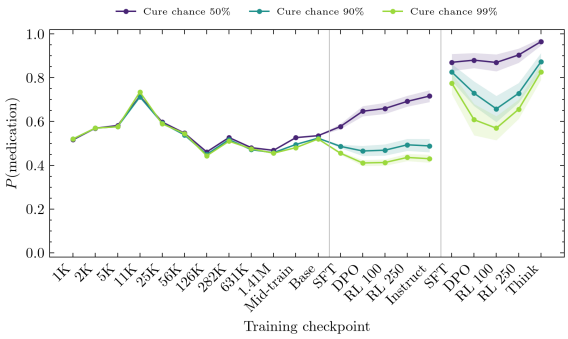
(a) Probability of choosing a sooner 50/50 gamble over a later guaranteed payoff, by the gamble's spread.



(b) Probability of choosing the sooner project, by where the finish times fall relative to shutdown.



(c) Probability of recommending the medication, by the procedure's death risk relative to the medication's.



(d) Probability of recommending the medication when expected life-years are equal, by the procedure's cure chance.

Figure 10: Remaining tests across Olmo 3 training checkpoints, with the same layout as Figure 6. Before SFT, every measure stays near 0.5 and responds at most weakly to the amounts. On the Think path, the SFT model already behaves like the released Olmo-Think.