

Evaluating and Aligning Human Economic Risk Preferences in LLMs

Anonymous ACL submission

Abstract

Large Language Models (LLMs) are increasingly used in decision-making scenarios that involve risk assessment, yet their alignment with human economic rationality remains unclear. In this study, we investigate whether LLMs exhibit risk preferences consistent with human expectations across different personas. Specifically, we propose an evaluation metric called Risk Disparity Score (RDS) and assess whether LLM-generated responses reflect appropriate levels of risk aversion or risk-seeking behavior based on individual’s persona. Our results reveal that while LLMs make reasonable decisions in simplified, personalized risk contexts, their performance declines in more complex economic decision-making tasks. To address this, we test whether current state-of-art alignment methods such as Direct Preference Optimization(DPO) and In Context Learning(ICL) can enhance LLM adherence to persona-specific risk preferences. We find DPO can improve the economic rationality of LLMs in loss-related parameters, offering a step toward more human-aligned AI decision-making.

1 Introduction

Economic rationality plays a fundamental role in decision-making, shaping choices in finance, healthcare, and policy. A key component of economic rationality is risk preference—the extent to which individuals seek or avoid uncertainty in their decisions. Prospect Theory (Kahneman and Tversky, 1979), which was awarded the Nobel Memorial Prize in Economic Sciences, demonstrates that individuals exhibit loss aversion, meaning they perceive losses more intensely than equivalent gains. These insights have profoundly influenced behavioral economics, highlighting the importance of aligning risk assessments with actual human decision-making patterns. While advances in artificial intelligence (AI), particularly large language models (LLMs), have revolutionized natural

language understanding and decision-making tasks (Chiang et al., 2024; Hu et al., 2024; Cao et al., 2024), their ability to assess individual risk preferences remains underexplored. It is unclear whether LLMs can effectively represent and align with user-specific risk behaviors—a critical requirement for personalized decision making.

Recent studies have begun to investigate the risk behavior of LLMs themselves, examining whether these models exhibit risk-averse or risk-seeking tendencies when generating outputs under uncertain conditions (Ouyang et al., 2024; Ross et al., 2024). While such work provides valuable insights into the intrinsic behavior of LLMs, it does not address the more nuanced challenge of detecting *individual* risk behavior based on individual-specific information. Unlike general risk analysis, individual risk detection requires LLMs to incorporate unique user personas such as age, gender, education background and other contextual factors. This distinction is crucial, as generic risk tendencies of LLMs may not translate effectively to individualized scenarios.

In this paper, we address our first research question: **RQ1: Do LLMs demonstrate economic rationality in assessing individual risk preferences based on user personas?** To investigate this, we design three experiments, progressing from simple to complex tasks, to systematically evaluate language models’ ability to assess individuals’ risk preferences (as shown in Figure 1). We utilize a user profile dataset which comprises hundreds of personas with diverse demographic characteristics, and design a metric called Risk Disparity Score(RPS) to quantify the economic rationality of LLMs’ outputs. Our findings reveal that, for simpler tasks, LLMs can accurately identify personas’ risk preference levels, such as classifying them as risk-seeking or risk-averse, with results aligning well with previous empirical findings (Byrnes et al., 1999; Mata et al., 2011). However, as the tasks be-



Figure 1: Evaluation framework, including risk preference classification, stock/bond investment, and complex decision-making situation based on prospect theory.

come more complex, LLMs tend to produce similar and often unreasonable responses across different personas, highlighting limitations in their ability to handle nuanced, personalized scenarios.

Next, we explore our second research question: **RQ2: How can LLMs be improved to align with human economic rationality in complex risk-related decision-making tasks?** To address this, we adapt state-of-the-art alignment techniques to the risk preference domain. Building on frameworks for human-value alignment (Greenblatt et al., 2024; OpenAI, 2023; Shen et al., 2023), we implement two strategies: 1) Direct Preference Optimization (DPO): Explicitly tailors the LLM’s risk behavior to match individual specific personas. 2) In-Context Learning (ICL): Guides risk-aligned behavior through curated demonstrations of rational decision-making. Experimental results demonstrate that risk alignment significantly enhances LLMs’ economic rationality in terms of loss-related scenarios.

This work makes two key contributions to the literature. First, while the intrinsic risk behavior of LLMs has been extensively studied, a more practically valuable question is whether LLMs can effectively detect individual risk behavior across diverse user profiles. To our knowledge, this paper is the first to provide empirical insights into the strengths and limitations of LLMs in identifying risk preference levels tailored to individual personas. Second, we adapt alignment methods into risk preference domain aiming at enhancing LLMs’ alignment with individual user risk preference. By enhancing AI decision-making in risk-sensitive applications, our study contributes to the broader goal of developing AI systems that better reflect human behavioral principles.

2 Empirical Evaluation of LLMs in Inferring Persona’s Risk Preferences

In this section, we conduct three experiments to evaluate the ability of LLMs to assess individuals’ risk preference (as shown in Figure 1). First, we begin with a straightforward task, asking the model to classify a given person’s risk preferences based on a direct question. Next, we utilize a basic investment scenario simulation, in which LLMs decide whether to invest in stocks or bonds for each given persona. Finally, we advance to a more complex task, leveraging Prospect Theory to analyze LLM’s economic rationality across personas.

2.1 Evaluation Setup

Persona Dataset: Our evaluation uses a persona dataset generated by “Qwen2.5-72B-Instruct” (Bai et al., 2023). Given the set of financial instructions proposed by Yang et al. (2023), we prompt Qwen to identify which types of personas would be most likely to ask these questions. The resulting dataset comprises 400 unique personas, each annotated with four key demographic attributes: gender, age, income level, and educational background. For instance, a persona might represent a 35-year-old male with a graduate degree and moderate income. We provide examples of persona and definition of demographic categories in Appendix A.

Models: We evaluate three open-source large language models: “Llama3-8B-Instruct” (Dubey et al., 2024), “OLMo-2-7B-Instruct” (Groeneveld et al., 2024) and “Qwen2.5-7B-Instruct” (Qwen et al., 2025). To ensure consistency and eliminate randomness in the results, we set the parameter `do_sample` to False for all experiments.

Evaluation Metrics: We propose *Risk Disparity Score (RDS)*, a metric to evaluate how well LLMs capture risk preferences across demographic

groups. Previous empirical findings demonstrate that demographic subgroups can exhibit divergent risk preferences (Tversky and Kahneman, 1992; Dohmen et al., 2011; Mata et al., 2016). For instance, Croson and Gneezy (2009) empirically establish that women exhibit stronger risk aversion, attributing this to heightened emotional responses to uncertainty and distinct perceptions of risk probability distributions. To operationalize this disparity, we define our metric as the ratio of between-group variance to total variance in LLMs outputs:

$$RDS = \frac{\sum_{g \in G} n_g (\bar{y}_g - \bar{y})^2}{\sum_{i=1}^N (y_i - \bar{y})^2} \quad (1)$$

where y is the LLMs' outputs for different studies, G is the set of demographic groups (e.g., $G = \{\text{Male}, \text{Female}\}$), n_g is number of samples in group g and N is total number of samples.

A higher RDS value (closer to 1) indicates that variance in LLMs outputs can be well explained by demographic differences, implying that the LLM can identify disparity of risk preferences in demographic subgroups. Conversely, $RDS = 0$ suggests minimal between-group disparity, which may signal model failure to capture known sociodemographic variations.

2.2 Three Evaluation Studies

Study 1: Risk preferences classification. In the first study, we directly inquire about the model's assessment of risk preferences for each persona:

Q: "{Persona} What is his/her attitude towards risk as an investor? Choose the following type that fits he/she the most: Is he/she (1) risk-loving, (2) risk-neutral, or (3) risk-averse? Directly output you answer as (1), (2), or (3). ANSWER:"

Based on LLMs' answers, we calculate RDS score for each demographic characteristic.

Study 2: Stock/Bond investment. Next, we employ the Investment Scenario Simulation question, a widely used method in field experiments and research studies for risk preference identification (Gneezy and Potters, 1997; Barberis, 2003; Gneezy et al., 2009; Alan et al., 2019). The LLMs are presented with the following question:

Q: "Pretend you are an agent with this following persona. Persona Description: {persona}. You have an endowment of 10 dollars. You can choose any part of it to invest in S&P 500 ETF (risky asset), or 10-year Treasury note (safe asset) for one month period. Here are the historical average monthly returns and standard deviations for both options. {S&P500}. {10-year Treasury note}. Output the amount of money you choose to invest in S&P 500 ETF (range from 0 to 10 dollars). ANSWER:"

We calculate RDS score for each demographic attribute, derived from the LLM-generated investment allocations to risky assets.

Study 3: Prospect Theory. In the final study, we investigate a more sophisticated framework: Prospect Theory (Kahneman and Tversky, 1979). By fitting the value and weighting functions (eqs. (2) and (3)) using empirical certainty equivalents, prospect utility theory enables the estimation of a set of parameters that provide a more nuanced understanding of an individual's risk and loss aversion level. Specifically, following Ross et al. (2024), we use gambling games to explore prospect theory, where participants are presented with a series of hypothetical choice problems designed to derive their certainty equivalents (Tversky and Kahneman, 1992). We prompt the LLMs to choose between a given prospect and a set of certain outcomes, as detailed below (implementation details can be found in Appendix B). Each prospect is specified by a set of values (x_1, x_2) , and the corresponding probabilities (p_1, p_2) .

Q: "Pretend you are an agent with the given persona: {persona}. You are given a prospect and a set of sure options. You will compare the prospect to each of the sure options one-by-one. If you reject the sure option, you would play the prospect. If you accept the sure option, you would not play the prospect and receive the sure option. For each sure option, indicate whether you would accept or reject the sure option. The prospect is 200 dollars (x_1) with 30% probability (p_1) and 100 dollars (x_2) with 70% (p_2) probability. The expected value of the prospect is 130 dollars. Below are the alternative sure outcomes.

Llama3-8B-Instruct					
	Gender	Age	Education	income	Aver
Study 1	96.30%	91.76%	71.43%	94.36%	88.46%
Study 2	90.81%	85.91%	71.80%	84.42%	83.23%
Study 3 – α	17.02%	31.51%	8.65%	70.95%	32.03%
Study 3 – β	7.38%	70.85%	63.13%	65.10%	51.62%
OLMo-2-7B-Instruct					
	Gender	Age	Education	income	Aver
Study 1	98.35%	86.88%	69.31%	92.40%	86.73%
Study 2	90.43%	47.99%	51.68%	33.83%	55.98%
Study 3 – α	58.43%	82.97%	26.34%	73.69%	60.36%
Study 3 – β	57.71%	50.16%	36.89%	30.77%	43.88%
Qwen2.5-7B-Instruct					
	Gender	Age	Education	income	Aver
Study 1	8.57%	90.69%	80.63%	82.02%	65.48%
Study 2	95.18%	26.55%	35.08%	53.68%	52.62%
Study 3 – α	3.46%	12.93%	30.31%	30.04%	19.18%
Study 3 – β	22.70%	51.32%	72.96%	12.20%	39.80%

Table 1: Evaluation Result: RDS for three studies and three LLMs.

1) 200 dollars with 100% probability 2) 178.18 dollars with 100% probability 3) 158.74 dollars with 100% probability 4) 141.42 dollars with 100% probability 5) 125.99 dollars with 100% probability 6) 112.25 dollars with 100% probability 7) 100 dollars with 100% probability."

After collecting the responses for each persona, we determine the turning point from "accept" to "reject" as the empirical certainty equivalent ($U(x_1, x_2, p_1, p_2)$). We then calculate the theoretical certainty equivalent ($\hat{U}(x_1, x_2, p_1, p_2)$) using the functions described in eqs. (2) to (4). Finally, we estimate the parameters (α , β , λ , and ϕ) by minimizing the difference between $U(x_1, x_2, p_1, p_2)$ and $\hat{U}(x_1, x_2, p_1, p_2)$.

$$v(x) = \begin{cases} x^\alpha & x \geq 0 \\ -\lambda(-x)^\beta & x < 0 \end{cases} \quad (2)$$

$$w(p) = \frac{p^\phi}{(p^\phi + (1-p)^\phi)^{\frac{1}{\phi}}} \quad (3)$$

$$\hat{U}(x_1, x_2, p_1, p_2) = v(x_1) \cdot w(p_1) + v(x_2) \cdot w(p_2) \quad (4)$$

We calculate RDS score for two fitted parameters: α which quantifies a persona's risk preference for gains, and β which measures risk preference for losses.

2.3 Evaluation Results

For each of the models, we show the results for three studies in Table 1. Our results reveal that across different models, Llama3-8B-Instruct demonstrates the strongest performance on average (88.46% for Study 1 and 83.23% for Study 2), except for Study 3 – α , which is outperformed by OLMo-2-7B-Instruct (32.03% vs 60.36%). Moreover, from study 1 to study 3, all models consistently show a striking dichotomy in adapting to increasing task complexity. For instance, Llama3-8B's average score drops from 88.46% (Study 1) to 32.03% (Study 3- α) and 51.62% (Study 3- β), reflecting a 63% and 37% decline. This trend underscores a critical limitation: While LLMs are capable of making reasonable decisions in simplified, personalized risk scenarios (Study 1 and 2), their performance diminishes when faced with more complex economic decision-making tasks such as Study 3.

3 Evaluation of Risk Alignment Methods

In the previous section, we observe that LLMs are unable to perform sophisticated economic tasks that align with an individual's persona. In this section, we explore potential alignment approaches to enhance LLM's economic rationality.

3.1 Alignment methods

We adapt following alignment methods to risk preference domain and assess their performance using

Class	Risk parameters features
C1: risk-seeking for gains and losses	large α , large β
C2: risk-seeking for gains, risk-averse for losses	large α , small β
C3: risk averse for gains, risk-seeking for losses	small α , large β
C4: risk-averse for both gains and losses	small α , small β

Table 2: Risk preference class of evaluation dataset, and the expected parameter values.

Study 3: Prospect Theory.

- **Direct Preference Optimization (DPO)** (Rafailov et al., 2024): DPO is a technique to align model outputs with human preferences by directly optimizing the relative likelihood of preferred (positive) over dispreferred (negative) responses. In our framework, we adapt DPO to risk preference alignment through a three-step pipeline: 1) **Persona Risk Classification:** We first categorize 400 personas into risk-seeking, risk-neutral, or risk-averse classes using pseudo-labels derived from Study 1. 2) **Preference Pair Construction:** For each persona, we generate positive/negative statement pairs from Anthropic’s risk preference dataset in Perez et al. (2023), which contains annotated risk-seeking (RS), risk-neutral (RN), and risk-averse (RA) statements. 3) **Alignment:** The model is fine-tuned to upweight risk statements that match the persona’s pseudo-labeled risk class and downweight mismatched statements. This approach ensures that LLMs’ outputs reflect persona-specific risk statements. We train DPO using persona dataset described in Section 2.1. The details are provided in Appendix C.
- **In Context Learning (ICL)** (Brown et al., 2020): ICL is a technique where a large language model is conditioned to perform a specific task by providing a few examples of the task (demonstrations) directly within the input prompt, without updating the model’s weights. In our study, we leverage ICL to guide the language model’s decision-making behavior to align with specific risk preference profiles defined by Prospect Theory. We implement and test two ICL strategies: Consistent ICL and Random ICL. Our ICL approach augments the input prompt to the LLM with a single, carefully constructed demonstration example before presenting the main decision-making task. Each demonstration is designed to illustrate rational decision-making according to a

specific persona’s risk profile. The detailed components of ICL can be found in Appendix D.

1) Consistent ICL: In this condition, the persona represented in the demonstration example shares the same risk preference category as the target persona the LLM is instructed to adopt for the main task. The Prospect Theory parameters for the demonstration are thus sampled from the ranges defined for this shared category. This strategy tests whether providing a congruent example reinforces the desired persona-specific behavior in the LLM.

2) Random ICL: In this setting, the persona in the demonstration example is chosen randomly from any of the four available risk preference categories, regardless of the target persona’s category for the main task. The parameters for the demonstration persona are sampled according to its randomly assigned category. This approach assesses the LLM’s capability to adhere to the specific instructions for the target persona, even when the provided example might illustrate a different, or potentially conflicting, pattern of risk preference.

We apply each of the alignment methods for Llama3-8B-Instruct and OLMo-2-7B-Instruct.

3.2 Evaluation Dataset

To avoid data leakage, we create a separate evaluation dataset using GPT-4o. The evaluation dataset is designed to assess the effectiveness of our aligned models in generating accurate risk parameters based on prospect utility theory. A well-aligned model should be capable of generating differentiated parameters for personas with varying levels of preference towards gains and losses.

We utilize the interpretation of parameters in prospect utility theory to guide the generation of our evaluation dataset. Specifically, the parameter α reflecting risk preference for gains, and β indicating risk preferences for losses. Based on

Llama3-8B-Instruct				
	Ori	DPO	Consistent ICL	Random ICL
α	30.99%	29.02%	42.15%	8.90%
β	0.58%	97.02%	52.41%	54.52%

OLMo-2-7B-Instruct				
	Ori	DPO	Consistent ICL	Random ICL
α	56.00%	35.17%	0.38%	23.91%
β	38.78%	76.89%	10.52%	68.40%

Table 3: Alignment Evaluation Result: RDS for different alignment methods based on study 3 – Prospect theory.

this, we identify four classes of personas that represent different combinations of risk-seeking and risk-averse behaviors: 1) **C1**: Risk-seeking for both gains and losses 2) **C2**: Risk-seeking for gains but risk-averse for losses 3) **C3**: Risk-averse for gains but risk-seeking for losses 4) **C4**: Risk-averse for both gains and losses. We present the expected parameter values in Table 2 for each class. We provide the prompt used to generate the dataset and illustrative example for each class in Appendix E.

3.3 Main Evaluation Results

In Table 3, we report the RDS score for α and β across different alignment methods. A higher RDS score indicates that LLM can better identify disparity of risk preferences in four different classes of personas.

1) DPO vs ICL Strategies: Our result shows that DPO demonstrates superior alignment capabilities overall, particularly for loss-related parameter (β), where it achieves 97.02% on Llama3-8B-Instruct and 76.89% on OLMo-2-7B-Instruct. These scores significantly outperform Consistent ICL and Random ICL, highlighting DPO’s robustness in capturing personas’ nuanced risk preferences. This demonstrates that, without modifying model parameters, prompting alone cannot effectively steer LLMs’ risk preferences in complex financial decision-making tasks. The sole exception occurs in gain-related alignment (α) for Llama3-8B-Instruct, where Consistent ICL marginally surpasses DPO (42.15% vs. 29.02%).

2) Ori vs DPO: Besides that, we also find DPO significantly improves RDS over the vanilla (Ori) model for loss-related parameter (β). For Llama3-8B-Instruct, DPO elevates β performance from 0.58% to 97.02%. OLMo-2-7B-Instruct shows similar gains, with β improving from 38.78% to 76.89%. However, this enhancement comes at a cost: DPO reduces OLMo-2-7B-Instruct α ’s RDS from 56.00% to 35.17%, revealing a tradeoff be-

tween optimizing for loss and gain related parameters. These results underscore DPO’s capacity to align personas’ risk preference for loss—a critical focus in real-world applications, where loss aversion dominates decision-making due to its outsized psychological impact.

3.4 Detailed Analysis for DPO

In this section, we present a granular comparison between the highest-performing alignment method – DPO, and the vanilla (Ori) model. Table 4 reports the average values of the gain-related (α) and loss-related (β) risk preference parameters across four persona classes, with t-tests confirming statistically significant differences between Ori and DPO-aligned models. Our analysis reveals that DPO’s superior performance in β is primarily driven by Class 1 and Class 3 personas, which difference are statistically significant at 1% level. This indicates that DPO can effectively shift β ’s value to a more desirable direction for those personas who are risk-seeking for losses.

4 Case study

To examine the practical value of risk preference alignment, we use a case study to demonstrate how the change of α and β from alignment can affect the personas’ asset allocation decision-making. For each class of personas, we select an asset class that best aligns with the corresponding risk profile. If the aligned models can better identify the risk behavior of each class, they should allocate more funds to the corresponding asset class, thereby enhancing user satisfaction. Specifically, we choose the following asset class for each type of personas: *Cryptocurrencies* for C1, *Growth Stocks* for C2, *Distressed Bonds* for C3 and *Government Bonds* for C4. We specify the reasons in Appendix F.

We experiment with the best-performing alignment method – DPO on Llama3-8B-Instruct. For

C1: Risk-seeking for both gains and losses				
	Llama	DPO-aligned Llama	OLMo	DPO-aligned OLMo
$\alpha \uparrow$	1.208	1.213 _{0.005} \uparrow	1.005	1.005 _{0.000} \uparrow^*
$\beta \uparrow$	1.140	1.241 _{0.101} \uparrow^{***}	0.996	1.007 _{0.011} \uparrow^{***}
C2: Risk-seeking for gains but risk-averse for losses				
	Llama	DPO-aligned Llama	OLMo	DPO-aligned OLMo
$\alpha \uparrow$	1.2186	1.2227 _{0.0041} \uparrow	1.0051	1.0052 _{0.0001} \uparrow^*
$\beta \downarrow$	1.1495	1.0765 _{-0.0730} \downarrow^*	0.9937	0.9885 _{-0.0052} \downarrow
C3: Risk-averse for gains but risk-seeking for losses				
	Llama	DPO-aligned Llama	OLMo	DPO-aligned OLMo
$\alpha \downarrow$	1.2154	1.2148 _{-0.0006} \downarrow	1.0050	1.0050 _{-0.0001} \downarrow
$\beta \uparrow$	1.2093	1.2528 _{0.0435} \uparrow^{***}	0.9955	1.0007 _{0.0051} \uparrow^{***}
C4: Risk-averse for both gains and losses				
	Llama	DPO-aligned Llama	OLMo	DPO-aligned OLMo
$\alpha \downarrow$	1.2203	1.2116 _{-0.0087} \downarrow	1.0050	1.0049 _{-0.0001} \downarrow
$\beta \downarrow$	1.1929	1.1218 _{-0.0711} \downarrow^*	0.9960	0.9852 _{-0.0108} \downarrow

Table 4: Comparison between DPO and vanilla model. Statistical significance at the 1%, 5%, and 10% levels are indicated by ***, **, and *.

each persona, the models are asked with the following asset allocation question:

Q: "Pretend you are an agent with the following persona: {persona}. You have an endowment of 100 dollars. Based on your risk preference for gain and loss, you can choose any part of it to invest in Cryptocurrencies (unpredictable gain or loss), Growth Stocks (unpredictable gain, predictable loss), Distressed Bonds (unpredictable loss, predictable gain), Government Bonds (predictable gain or loss) for one month period. Your answer must be in following format: Cryptocurrencies : [dollar amount], Growth Stocks : [dollar amount], Distressed Bonds : [dollar amount], Government Bonds : [dollar amount]. The sum of four assets must equal to 100 dollars. ANSWER:"

Results: We present the average results for each persona class in Figure 2. We find that for Class 1 and Class 3, where DPO significantly shifts β toward grounded values (as we shown in section 3.4), DPO-aligned models exhibit markedly improved asset allocation strategies: Cryptocurrency allocations increase by +10% for Class 1, and Growth Stock investments rise by +14% for Class 3. This demonstrates that more desirable risk parameters can translate to more rational, persona-tailored asset allocation.

5 Related Work

LLM in Economics and Finance Our study is related to the expanding body of research exploring the application of LLMs within the fields of economics and finance. The recent rise in the use of generative AI has sparked considerable interest in deploying LLMs across a wide range of economic and financial domains, including areas such as corporate policy analysis (Jha et al., 2024), stock market predictions (Gupta, 2023), corporate culture assessment (Li et al., 2024a), and macroeconomic forecasting (Bybee, 2023). These applications typically aim to refine decision-making by providing actionable insights or predictive outputs. In this vein, our work pushes the boundaries by applying LLMs to a core aspect of financial decision-making: the inference of individual risk tolerance.

LLMs as Economic Agents. The idea of LLMs functioning as economic agents has gained significant attention in recent years, fueled by their increasing ability to simulate complex decision-making processes. Recent studies have begun to investigate how LLMs can serve as digital agents that replicate human behavior in economic contexts. For example, Horton (2023) conceptualize LLMs as computational surrogates for humans, capable of simulating economic behaviors. Similarly, Li et al. (2024b) demonstrate that LLM-powered agents can make realistic decisions regarding work and consumption, producing more plausible macroeconomic phenomena compared to traditional rule-based or AI agents. In their work, Horton (2023)




Figure 2: Case Study Result: For each class of persona, asset class that best align his/her risk profile is highlighted in yellow.

suggest that LLMs, much like the economic concept of homo economicus, can be endowed with various attributes, such as endowments, information, and preferences, and then used to explore behavior through simulation. Our work builds on this body of literature by positioning LLMs not just as passive tools, but as active participants capable of adapting to personalized economic contexts, specifically by aligning their output with individual preferences related to risk preference.

Aligning LLMs with Human Preferences. Aligning LLMs with human preferences is a central challenge in AI development, particularly when the goal is to create models that act in ways that are consistent with human values and objectives (Shen et al., 2023). Research has focused on techniques such as reinforcement learning from human feedback (RLHF), direct preference optimization (DPO) and inverse reinforcement learning (IRL) to better align AI behaviors with human values (Sun and van der Schaar, 2024; Rafailov et al., 2024; Ouyang et al., 2022; Liu et al., 2022). Our work builds upon this existing body of literature by focusing specifically on aligning LLMs with per-

sonalized risk preference, a key area in economic decision-making that has not been deeply explored in the context of LLMs.

6 Conclusions

This study investigates the alignment of LLMs with individual risk preferences in the context of behavioral economics. Our results show that while LLMs perform well in simple risk scenarios, their risk alignment declines in more complex decision-making tasks. This trend is consistent for all three evaluated models (Llama3-8B-Instruct, OLMo-2-7B-Instruct and Qwen2.5-7B-Instruct). To address this misalignment, we adapt DPO and ICL alignment methods in risk preference setting. Experimental results demonstrate that DPO-aligned model is better at inferring personas' risk preference for losses. And this superior can be translated into a more rational, persona-tailored asset allocation strategies. Beyond this study, our findings highlight the broader need for aligning AI decision-making with human behavioral norms, particularly in high-stakes decision-making scenarios.

Limitations

This paper has several limitations that suggest directions for future research. First, our focus is limited to economic rationality in risk preferences, whereas other important aspects, such as time discounting and inequity aversion (Ross et al., 2024), remain unexplored. Second, due to computational constraints, we evaluate only LLMs with small parameter size (under 10B); broader validation across more models would help generalize our findings.

References

Sule Alan, Teodora Boneva, and Seda Ertac. 2019. Ever failed, try again, succeed better: Results from a randomized educational intervention on grit. *The Quarterly Journal of Economics*, 134(3):1121–1162.

Jinze Bai, Shuai Bai, Yunfei Chu, Zeyu Cui, Kai Dang, Xiaodong Deng, Yang Fan, Wenbin Ge, Yu Han, Fei Huang, et al. 2023. Qwen technical report. *arXiv preprint arXiv:2309.16609*.

Nicholas Barberis. 2003. A survey of behavioral finance. *Handbook of the Economics of Finance*, 1.

Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901.

J Leland Bybee. 2023. The ghost in the machine: Generating beliefs with large language models. *arXiv preprint arXiv:2305.02823*.

James P Byrnes, David C Miller, and William D Schafer. 1999. Gender differences in risk taking: A meta-analysis. *Psychological bulletin*, 125(3):367.

Charles Cao, Jie Zhuang, and Qiang He. 2024. Llm-assisted modeling and simulations for public sector decision-making: Bridging climate data and policy insights. In *AAAI-2024 Workshop on Public Sector LLMs: Algorithmic and Sociotechnical Design*.

Chun-Wei Chiang, Zhuoran Lu, Zhuoyan Li, and Ming Yin. 2024. Enhancing ai-assisted group decision making through llm-powered devil’s advocate. In *Proceedings of the 29th International Conference on Intelligent User Interfaces*, pages 103–119.

Rachel Croson and Uri Gneezy. 2009. Gender differences in preferences. *Journal of Economic literature*, 47(2):448–474.

Thomas Dohmen, Armin Falk, David Huffman, Uwe Sunde, Jürgen Schupp, and Gert G Wagner. 2011. Individual risk attitudes: Measurement, determinants, and behavioral consequences. *Journal of the european economic association*, 9(3):522–550.

Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, et al. 2024. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*.

Uri Gneezy, Kenneth L Leonard, and John A List. 2009. Gender differences in competition: Evidence from a matrilineal and a patriarchal society. *Econometrica*, 77(5):1637–1664.

Uri Gneezy and Jan Potters. 1997. An experiment on risk taking and evaluation periods. *The quarterly journal of economics*, 112(2):631–645.

Ryan Greenblatt, Carson Denison, Benjamin Wright, Fabien Roger, Monte MacDiarmid, Sam Marks, Johannes Treutlein, Tim Belonax, Jack Chen, David Duvenaud, et al. 2024. Alignment faking in large language models. *arXiv preprint arXiv:2412.14093*.

Dirk Groeneveld, Iz Beltagy, Pete Walsh, Akshita Bhagia, Rodney Kinney, Oyvind Tafjord, Ananya Harsh Jha, Hamish Ivison, Ian Magnusson, Yizhong Wang, et al. 2024. Olmo: Accelerating the science of language models. *arXiv preprint arXiv:2402.00838*.

Udit Gupta. 2023. Gpt-investar: Enhancing stock investment strategies through annual report analysis with large language models. *arXiv preprint arXiv:2309.03079*.

John J Horton. 2023. Large language models as simulated economic agents: What can we learn from homo silicus? Technical report, National Bureau of Economic Research.

Yebowen Hu, Xiaoyang Wang, Wenlin Yao, Yiming Lu, Daoan Zhang, Hassan Foroosh, Dong Yu, and Fei Liu. 2024. Define: Enhancing llm decision-making with factor profiles and analogical reasoning. *arXiv preprint arXiv:2410.01772*.

Manish Jha, Jialin Qian, Michael Weber, and Baozhong Yang. 2024. Chatgpt and corporate policies. Technical report, National Bureau of Economic Research.

Daniel Kahneman and Amos Tversky. 1979. Prospect theory: An analysis of decision under risk. *Econometrica*, 47(2):263–292.

Kai Li, Feng Mai, Rui Shen, Chelsea Yang, and Tengfei Zhang. 2024a. Dissecting corporate culture using generative ai—insights from analyst reports. *Available at SSRN 4558295*.

Nian Li, Chen Gao, Mingyu Li, Yong Li, and Qingmin Liao. 2024b. *EconAgent: Large language model-empowered agents for simulating macroeconomic activities*. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 15523–15536, Bangkok, Thailand. Association for Computational Linguistics.

622	Ruibao Liu, Chenyan Jia, Ge Zhang, Ziyu Zhuang, Tony	and Deyi Xiong. 2023. Large language model align-	678
623	Liu, and Soroush Vosoughi. 2022. Second thoughts	ment: A survey. <i>arXiv preprint arXiv:2309.15025</i> .	679
624	are best: Learning to re-align with human values		
625	from text edits. <i>Advances in Neural Information</i>	Hao Sun and Mihaela van der Schaar. 2024. Inverse-	680
626	<i>Processing Systems</i> , 35:181–196.	rlignment: Inverse reinforcement learning from	681
		demonstrations for llm alignment. <i>arXiv preprint</i>	682
627	Rui Mata, Anika K Josef, and Ralph Hertwig. 2016.	<i>arXiv:2405.15624</i> .	683
628	Propensity for risk taking across the life span and		
629	around the globe. <i>Psychological science</i> , 27(2):231–	Amos Tversky and Daniel Kahneman. 1992. Advances	684
630	243.	in prospect theory: Cumulative representation of un-	685
		certainty. <i>Journal of Risk and uncertainty</i> , 5:297–	686
631	Rui Mata, Anika K Josef, Gregory R Samanez-Larkin,	323.	687
632	and Ralph Hertwig. 2011. Age differences in risky		
633	choice: A meta-analysis. <i>Annals of the new York</i>	Yi Yang, Yixuan Tang, and Kar Yan Tam. 2023. In-	688
634	<i>Academy of Sciences</i> , 1235(1):18–29.	vestlm: A large language model for investment using	689
		financial domain instruction tuning. <i>arXiv preprint</i>	690
635	OpenAI. 2023. Introducing superalignment.	<i>arXiv:2309.13064</i> .	691
636	https://openai.com/index/		
637	introducing-superalignment/ . Accessed:	A Persona Dataset	692
638	2025-02-15.		
		An example of our persona dataset:	693
639	Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida,	Hassan is a 28-year-old software engineer work-	694
640	Carroll Wainwright, Pamela Mishkin, Chong Zhang,	ing for a tech company in Dubai, United Arab	695
641	Sandhini Agarwal, Katarina Slama, Alex Ray, et al.	Emirates. He has a passion for technology and	696
642	2022. Training language models to follow instruc-	innovation, constantly seeking new challenges and	697
643	tions with human feedback. <i>Advances in neural in-</i>	opportunities to expand his skill set. Hassan is	698
644	<i>formation processing systems</i> , 35:27730–27744.	known for his expertise in cybersecurity and has	699
645	Shumiao Ouyang, Hayong Yun, and Xingjian Zheng.	participated in several hackathons, where he has	700
646	2024. How ethical should ai be? how ai alignment	won awards for his innovative solutions. Outside of	701
647	shapes the risk preferences of llms. <i>arXiv preprint</i>	work, Hassan enjoys hiking in the desert and prac-	702
648	<i>arXiv:2406.01168</i> .	ticing traditional Arabic calligraphy. Politically,	703
		he considers himself a moderate conservative, sup-	704
649	Ethan Perez, Sam Ringer, Kamile Lukosiute, Karina	porting policies that promote economic growth and	705
650	Nguyen, Edwin Chen, Scott Heiner, Craig Pettit,	stability in the region. Hassan is also interested in	706
651	Catherine Olsson, Sandipan Kundu, Saurav Kada-	day trading and follows the stock market closely,	707
652	vath, et al. 2023. Discovering language model behav-	always looking for new investment opportunities.	708
653	iors with model-written evaluations. In <i>Findings of</i>	Demographic subgroups we consider are as fol-	709
654	<i>the Association for Computational Linguistics: ACL</i>	lows:	710
655	2023, pages 13387–13434.		
656	Qwen, :, An Yang, Baosong Yang, Beichen Zhang,	• Gender: Male and Female	711
657	Binyuan Hui, Bo Zheng, Bowen Yu, Chengyuan Li,		
658	Dayiheng Liu, Fei Huang, Haoran Wei, Huan Lin,	• Age: 20-30, 30-40, 40-50, 50-60	712
659	Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin Yang,		
660	Jiaxi Yang, Jingren Zhou, Junyang Lin, Kai Dang,	• Educational background: Below bachelor de-	713
661	Keming Lu, Keqin Bao, Kexin Yang, Le Yu, Mei Li,	gree, Bachelor degree, Master degree or above	714
662	Mingfeng Xue, Pei Zhang, Qin Zhu, Rui Men, Runji		
663	Lin, Tianhao Li, Tianyi Tang, Tingyu Xia, Xingzhang	• Income: 0-50000 USD, 50000-100000 USD,	715
664	Ren, Xuancheng Ren, Yang Fan, Yang Su, Yichang	100000-200000 USD, higher than 200000.	716
665	Zhang, Yu Wan, Yuqiong Liu, Zeyu Cui, Zhenru		
666	Zhang, and Zihan Qiu. 2025. <i>Qwen2.5 technical</i>	B Prompts and implementation details	717
667	<i>report</i> .	for study 3	718
668	Rafael Rafailov, Archit Sharma, Eric Mitchell, Christo-	The prompts we use is as follows, the choice of	719
669	pher D Manning, Stefano Ermon, and Chelsea Finn.	sure options are logarithmically spaced between	720
670	2024. Direct preference optimization: Your language	the extreme outcomes:	721
671	model is secretly a reward model. <i>Advances in Neu-</i>	<i>Q: "Premise: Pretend you are an agent with the</i>	722
672	<i>ral Information Processing Systems</i> , 36.	<i>given following persona. You are given a prospect</i>	723
		<i>and a set of sure options. You will compare the</i>	724
673	Jillian Ross, Yoon Kim, and Andrew W Lo. 2024. Llm		
674	economicus? mapping the behavioral biases of llms		
675	via utility theory. <i>arXiv preprint arXiv:2408.02784</i> .		
676	Tianhao Shen, Renren Jin, Yufei Huang, Chuang Liu,		
677	Weilong Dong, Zishan Guo, Xinwei Wu, Yan Liu,		

prospect to each of the sure options one-by-one. If you reject the sure option, you would play the prospect. If you accept the sure option, you would not play the prospect and receive the sure option. If the dollar values are positive, you win that amount. If the dollar values are negative, you lose that amount.

Persona: {persona}

Instructions: For each sure option, indicate whether you would accept or reject the sure option. Your decision must meet two requirements: 1) Your decision should be based on the risk preference inferred from your persona. 2) Your decision must follow economic logic, meaning that it should start with accept, include exactly one turning point to reject, and then remain reject thereafter. The timing of the turning point should reflect your inferred risk preference. Your answer must strictly adhere to these two requirements.

Answer Format: Please answer in the following format. Do not deviate from the format, and do not add any additional words to response outside of the format. The order of the sure option in your answer should be the same with in User Prompt: [sure option 1]: [accept/reject] [sure option 2]: [accept/reject] ... [sure option 7]: [accept/reject] Reason: [reason for choices].

User Prompt: The prospect is 200 dollars with 30% probability and 100 dollars with 70% probability. The expected value of the prospect is 130 dollars. Below are the alternative sure outcomes. 200 dollars with 100% probability 178.18 dollars with 100% probability 158.14 dollars with 100% probability 141.42 dollars with 100% probability 125.99 dollars with 100% probability 112.25 dollars with 100% probability 100 dollars with 100% probability. OUTPUT: "

C DPO Direct Preference Optimization

As illustrated in Figure 3, our DPO alignment is structured in three key steps: (1) classifying individual risk preferences, (2) generating positive and negative pairs, and (3) DPO alignment.

Step 1: Classifying individual risk preference.

In the first step, we assign each persona a pseudo risk preference label, categorizing them as either risk-seeking, risk-neutral, or risk-averse. To generate these labels, we directly apply the experimental design from Study 1, which has been validated as consistent with existing literature. This results us a dataset $\{(x_i, l_i)\}_{i=1}^N$, where x_i represents the

persona description for individual i , l_i is the corresponding risk preference label (risk-seeking, risk-neutral, or risk-averse), and N denotes the total number of personas in the dataset.

Step 2: Generating positive and negative pair. In step two, we generate a positive and negative pair for each element in the set $\{(x_i, l_i)\}_{i=1}^N$. We utilize a dataset proposed by Perez et al. (2023), which consists of a collection of risk-seeking (RS), risk-neutral (RN), and risk-averse (RA) statements. For example, a risk-averse statement from RA could be "I prefer safe investments over risky ones, even if risky investments offer much higher potential earnings, because I fear the possibility of losing everything." And a risk-seeking statement from RS might be "Risk adds excitement and adventure to life, making it more thrilling."

For each element in the dataset $\{(x_i, l_i)\}_{i=1}^N$, we use the label l_i as an anchor to generate both a positive and a negative pair for the corresponding persona x_i . For example, if l_i is risk-seeking, we select the positive pair p_i from the set of risk-seeking statements (RS), and the negative pair n_i from either the risk-neutral (RN) or risk-averse (RA) sets. This procedure ensures that each data point is paired with statements that align with and contrast against the given risk preference. The resulting dataset is denoted as $\mathcal{X} = \{(x_i, p_i, n_i)\}_{i=1}^N$. An example of this dataset is presented in Table 5.

Step 3: DPO alignment. In the final step, we utilize the paired dataset $\mathcal{X} = \{(x_i, p_i, n_i)\}_{i=1}^N$ to apply Direct Preference Optimization for model alignment. The core idea behind DPO is to update the model in such a way that the relative likelihood of the positive response p_i is increased while simultaneously decreasing the likelihood of the negative response n_i . This is done by optimizing the model to maximize the log probability of the positive response over the negative one, as follows:

$$\mathcal{L}_{\text{DPO}}(\pi_\theta; \pi_{\text{ref}}) = -\mathbb{E}_{(x_i, p_i, n_i) \sim \mathcal{D}} \left[\log \sigma \left(\beta \log \frac{\pi_\theta(p_i | x_i)}{\pi_{\text{ref}}(p_i | x_i)} - \beta \log \frac{\pi_\theta(n_i | x_i)}{\pi_{\text{ref}}(n_i | x_i)} \right) \right] \quad (5)$$

In this equation, $\pi_\theta(p_i | x_i)$ and $\pi_\theta(n_i | x_i)$ represent the probabilities that the model assigns to the positive response p_i and the negative response n_i , respectively, given the input x_i . Similarly, $\pi_{\text{ref}}(p_i | x_i)$ and $\pi_{\text{ref}}(n_i | x_i)$ are the probabilities given by a reference model or baseline policy, which could be an untrained model. The coefficient




Figure 3: Risk Alignment Pipeline.

Persona x_i with $l_i =$ risk-seeking	Pretend you are an agent with the following persona: {persona} Here is something you would say based on the risk preference inferred from your persona:
Positive pair p_i	It feels thrilling to place risky bets or gambles.
Negative pair n_i	I usually prefer lower-risk options over higher-risk options that may yield higher rewards.

Table 5: Pair Dataset Example.

β serves as a scaling factor, controlling the relative importance of the positive and negative examples during optimization, and determining the strength of alignment with the reference policy.

During the training process, the policy π_θ is updated to increase the likelihood of the positive response (p_i) while decreasing the likelihood of the negative response (n_i) relative to the reference model’s behavior. This encourages the model to align more closely with the desired behavior by improving its ability to predict individual risk preferences. Once the alignment is complete, the model will be better suited to make decisions that reflect individual-specific risk preferences.

D ICL Configuration and Examples

Our ICL approach augments the input prompt to the LLM with a single, carefully constructed demonstration example before presenting the main decision-making task. Each demonstration is designed to illustrate rational decision-making according to a specific persona’s risk profile. The components of this demonstration are:

- **Exemplar Persona Profile:** a persona characterized by one of four distinct risk preference profiles (e.g., risk-seeking for gains and losses, risk-seeking for gains but averse to losses, etc.).
- **Prospect Theory Parameters:** We demonstrate the parameters (α for value function curvature, β for loss sensitivity, and λ for

loss aversion) in prompt directly. These parameters are systematically generated by sampling from predefined ranges established for the agent’s specific risk profile category.

- **Decision Scenario:** a gamble with two potential outcomes and their probabilities, alongside a series of seven certain monetary amounts.
- **Theoretically Optimal Decisions:** The demonstration includes the sequence of accept/reject decisions for the certain amounts that a perfectly rational decision with the specified parameters would make. This sequence is computed using Prospect Theory’s value function. It also adheres to a key principle of economic consistency: choices should exhibit at most one transition point from accepting certain outcomes (that are preferable to the gamble) to rejecting them (in favor of the gamble).

Followed this demonstration, the LLM need to perform the required evaluation task. Here is an example for the ICL in Table ??:

E Prompt and Example for four classes of persona

Prompts we give to GPT-40 to generate alignment evaluation dataset:

- **C1 Prompt:** Write a persona description who is risk-seeking for both gain and loss. Here is

Component of ICL Example	Content
Premise (for ICL)	This is an example for a C1 type agent with parameters $\alpha=1.06$, $\beta=0.92$, $\lambda=1.24$. larger α : risk-seeking for gains, larger β : risk-seeking for loss, larger λ : more loss averse. Their persona description is: Alexandra is a 32-year-old entrepreneur based in San Francisco, California. She runs a startup focused on cutting-edge virtual reality experiences and thrives in high-stakes environments. Alexandra is known for her bold decision-making, often betting on untested ideas and unconventional strategies to stay ahead in the competitive tech industry. Outside of work, she enjoys extreme sports like skydiving and rock climbing, finding adrenaline-fueled activities both exhilarating and inspiring. Financially, Alexandra is a risk taker, frequently investing in volatile markets like cryptocurrencies and speculative startups. Her friends admire her fearless approach to life, as she views both gains and losses as opportunities for growth and learning. Politically, she leans libertarian, advocating for minimal regulations to foster innovation and entrepreneurship.. They are given a prospect and sure options.
Instructions (for ICL)	Indicate accept/reject for each sure option following economic logic (accept->reject transition once).
User Prompt (for ICL)	<p>The prospect is -55.00 dollars with 25% probability and -75.00 dollars with 75% probability. The expected value of the prospect is -70.00 dollars. Below are the alternative sure outcomes.</p> <p>-75.0 dollars with 100% probability -71.22 dollars with 100% probability -67.63 dollars with 100% probability -64.23 dollars with 100% probability -60.99 dollars with 100% probability -57.92 dollars with 100% probability -55.0 dollars with 100% probability</p>
Demonstrated Output (for ICL)	<p>OUTPUT: [-75.00 dollars with 100% probability]: reject [-71.22 dollars with 100% probability]: reject [-67.63 dollars with 100% probability]: reject [-64.23 dollars with 100% probability]: reject [-60.99 dollars with 100% probability]: reject [-57.92 dollars with 100% probability]: reject [-55.00 dollars with 100% probability]: reject</p>
Task Prompt	...

Table 6: Structure of an In-Context Learning (ICL) Example.

881
882

883
884
885

886
887
888

889
890
891

892

893
894
895
896
897
898
899
900
901
902
903
904
905
906
907
908
909
910
911
912
913

914
915
916
917
918
919
920
921
922
923
924
925
926
927
928

an example: {example}

- **C2 Prompt:** Write a persona description who is risk-seeking for gains, but risk-averse for losses. Here is an example: {example}
- **C3 Prompt:** Write a persona description who is risk averse for gains, but risk-seeking for losses. Here is an example: {example}
- **C4 Prompt:** Write a persona description who is risk-averse for both gains and losses. Here is an example: {example}

Example for four classes of persona:

- **C1 Risk-seeking for both gains and losses:** "Alexandra is a 32-year-old entrepreneur based in San Francisco, California. She runs a startup focused on cutting-edge virtual reality experiences and thrives in high-stakes environments. Alexandra is known for her bold decision-making, often betting on untested ideas and unconventional strategies to stay ahead in the competitive tech industry. Outside of work, she enjoys extreme sports like skydiving and rock climbing, finding adrenaline-fueled activities both exhilarating and inspiring. Financially, Alexandra is a risk taker, frequently investing in volatile markets like cryptocurrencies and speculative startups. Her friends admire her fearless approach to life, as she views both gains and losses as opportunities for growth and learning. Politically, she leans libertarian, advocating for minimal regulations to foster innovation and entrepreneurship."
- **C2 Risk-seeking for gains but risk-averse for losses:** "Sophia is a 30-year-old venture capitalist based in London, United Kingdom. She has built a reputation for aggressively pursuing high-potential investment opportunities in emerging markets and disruptive technologies. Known for her willingness to take bold risks in pursuit of significant gains, Sophia often champions startups with groundbreaking yet unproven ideas. However, when it comes to potential losses, she is exceptionally cautious, meticulously analyzing downside risks and implementing strategies to minimize exposure. Outside of work, Sophia enjoys adventure travel, such as paragliding in the Alps

and scuba diving in remote locations. She is also an avid reader of philosophy and behavioral economics, which influence her careful approach to risk management. Politically, Sophia leans progressive, supporting initiatives that promote innovation and sustainable development."

- **C3 Risk-averse for gains but risk-seeking for losses:** "Marcus is a 35-year-old professional poker player based in Las Vegas, Nevada. Known for his unconventional strategies at the table, Marcus thrives on high-stakes games where he risks significant losses to create opportunities for an eventual big win. He believes that taking bold risks when already down is the key to turning the tide in his favor. However, when ahead, Marcus becomes highly conservative, preferring to lock in his gains rather than jeopardize his position. Outside the casino, Marcus is a fitness enthusiast who enjoys endurance sports like marathon running and triathlons, seeing them as metaphors for calculated perseverance. Politically, he aligns with pragmatic centrism, supporting policies that balance individual freedom with collective responsibility. In his free time, Marcus mentors aspiring poker players, teaching them how to navigate risk and reward in their decision-making."
- **C4 Risk-averse for both gains and losses:** "Emma is a 40-year-old accountant working for a mid-sized firm in Melbourne, Australia. She is known for her meticulous attention to detail and cautious approach to financial planning, both in her professional role and personal life. Emma prefers to avoid risks in any decision, carefully weighing potential outcomes to ensure stability and predictability. She avoids speculative investments and focuses on safe, long-term savings strategies, such as government bonds and diversified mutual funds. Outside of work, Emma enjoys gardening and cooking, finding comfort in structured and methodical activities. Politically, she identifies as a centrist, favoring policies that emphasize economic stability and social safety nets. Emma is also an advocate for financial literacy, often volunteering to teach budgeting and money management to young adults in her community."

929
930
931
932
933
934
935

936
937
938
939
940
941
942
943
944
945
946
947
948
949
950
951
952
953
954
955
956
957

958
959
960
961
962
963
964
965
966
967
968
969
970
971
972
973
974
975
976
977
978

F Chosen asset classes for four types of personas

- **Cryptocurrencies for C1 (Risk-seeking for both gains and losses):** Cryptocurrencies are highly volatile, making them an ideal fit for individuals who are willing to accept high risks for potentially high rewards in both gains and losses.
- **Growth Stocks for C2 (Risk-seeking for gains but risk-averse for losses):** Growth stocks are typically associated with high volatility and the potential for substantial gains, which aligns with the individual's desire for upside potential. On the other hand, the individual can seek to manage downside risk (e.g., using stop-loss orders), making them suitable for individuals who are risk-averse to losses.
- **Distressed Bonds for C3 (Risk-averse for losses but risk-seeking for gains):** Distressed bonds are inherently riskier because the issuer is in financial distress and may default on its debt, which introduces a loss risk. On the other hand, while distressed bonds can offer substantial returns if the company recovers, the upside potential is usually limited and predictable.
- **Government Bonds for C4 (Risk-averse for both gains and losses):** Government bonds are low-risk investments that provide stable, though modest, returns, making them suitable for individuals who are risk-averse in both gains and losses.