

How Does Cognitive Bias Affect Large Language Models? A Case Study on the Anchoring Effect in Price Negotiation Simulations

Yoshiki Takenami Yin Jou Huang Yugo Murawaki Chenhui Chu

Kyoto University

{tnami, huang}@nlp.ist.i.kyoto-u.ac.jp {murawaki, chu}@i.kyoto-u.ac.jp

Abstract

Cognitive biases, well-studied in humans, can also be observed in LLMs, affecting their reliability in real-world applications. This paper investigates the anchoring effect in LLM-driven price negotiations. To this end, we instructed seller LLM agents to apply the anchoring effect and evaluated negotiations using not only an objective metric but also a subjective metric. Experimental results show that LLMs are influenced by the anchoring effect like humans. Additionally, we investigated the relationship between the anchoring effect and factors such as reasoning and personality. It was shown that reasoning models are less prone to the anchoring effect, suggesting that the long chain of thought mitigates the effect. However, we found no significant correlation between personality traits and susceptibility to the anchoring effect. These findings contribute to a deeper understanding of cognitive biases in LLMs and to the realization of safe and responsible application of LLMs in society.

1 Introduction

Large Language Models (LLMs) have demonstrated remarkable language generation capabilities, achieving high performance in various natural language processing tasks (Jiang et al., 2024b; Zhu et al., 2024; Wang et al., 2023). Such sophisticated abilities and human-like behaviors exhibited by LLMs are acquired through large-scale training on human-generated text. However, it is well known that human cognition and decision-making are often irrational and systematically deviate from normative standards, a phenomenon referred to as cognitive biases (Tversky and Kahneman, 1974). Since human-generated text is considered to embed cognitive biases, recent studies have revealed that LLMs trained on such text are also influenced by these biases (Sumita et al., 2025; Suri et al., 2024).

The anchoring effect is one of the widely recognized cognitive biases (Tversky and Kahneman,

1974; Kahneman, 1992). The anchoring effect refers to the phenomenon in which an initially presented piece of information (the anchor) significantly influences subsequent judgments and decisions. Psychology researchers have investigated various factors that can affect susceptibility to the anchoring effect, such as deliberation, which refers to the process of engaging in extended thinking over time. Cognitive biases, including the anchoring effect, are believed to originate from intuitive processing and can be mitigated through effortful reasoning (Kahneman, 2011; Rastogi et al., 2022). Another factor is personality traits. Although some studies suggest that certain personality traits influence the extent to which individuals are affected by the anchoring effect (Caputo, 2014; Furnham et al., 2012), the findings are inconsistent, with other research indicating no such correlation (Cheek and Norem, 2020; Schindler et al., 2021).

In decision-making scenarios such as price negotiations, which are the focus of this study, the effectiveness of the anchoring effect in humans has been well studied (Guo et al., 2022). For example, a seller can leverage the anchoring effect by initially presenting a price higher than their target price, distorting the buyer’s perception to raise the agreed price. A previous study (Xia et al., 2024) simulated price negotiations using LLMs to reproduce the phenomenon in which buyers can leverage the anchoring effect to lower the agreed price and increase their profits. However, the results are limited in several respects, such as the fact that the presented anchor is fixed as a constant multiple of the target price, that the evaluation of price negotiations is one-dimensional, and that the effects of deliberation and personality traits are not taken into account. More flexible and systematic simulations are required in order to obtain more generalizable results.

In this study, we systematically investigate whether LLMs exhibit negotiation behaviors con-




Figure 1: Investigation of the anchoring effect in price negotiation simulations using LLMs. First, personality profiles of sellers and buyers are randomly generated based on the Big Five personality framework, and then price negotiation simulations are conducted. Along with the product, its description, and target price, the seller agents are instructed via prompts on whether to use the anchoring effect. Similarly, buyer agents are prompted regarding whether they are informed of the seller’s strategy and whether they engage in reasoning. The negotiation ends when the negotiation state reaches accept, after which agents respond to a satisfaction questionnaire.

sistent with those observed in human studies for the anchoring effect (Figure 1). To this end, we conduct a price negotiation simulation using LLMs with a high degree of flexibility in setting anchors. Specifically, we achieve this flexibility by instructing the model without specifying any exact numeric value: “offer a higher price than your target price at the initial stage.” We introduce three aspects to conduct a more in-depth analysis of price negotiation than that of prior work: 1) We incorporate a subjective metric, namely satisfaction, to provide a more comprehensive assessment of negotiations. 2) We employ reasoning models (o1¹ and QwQ²) to investigate the influence of deliberation on negotiation behavior, because these two models are trained to generate a long chain of thought before responding to user. 3) We control the personality traits of LLM agents, allowing us to enhance the reliability of our results and analyze the correlation between the anchoring effect and personality traits.

Our findings are summarized as follows:

- The simulation results largely align with those studies conducted in humans, demonstrating

that LLMs are also subject to the anchoring effect.

- Our analysis demonstrated that the reasoning model was less susceptible to the anchoring effect, suggesting that deliberation can mitigate its influence in negotiation contexts.
- Contrary to some previous studies on humans, we found no significant correlation between the anchoring effect and personality traits.

These results not only deepen understanding of how LLMs behave against cognitive biases but also lay the foundation for their safe and responsible application in real-world settings.

2 Related Work

LLMs have been increasingly used as simulations of human behavior (Xie et al., 2024; Aher et al., 2023; Akata et al., 2025). They can be understood as a superposition of perspectives with different values and personality traits (Kovač et al., 2023), and research has been conducted to analyze which sub-populations they reflect and which groups are less accurately represented (Santurkar et al., 2023).

¹o1-2024-12-17

²Qwen/QwQ-32B

From this superposition, in-context learning enables the extraction of specific perspectives. This approach effectively induces LLMs with distinct personalities in a controllable manner, allowing them to generate diverse and verifiable behaviors (Jiang et al., 2023, 2024a; Shao et al., 2023). LLM-based simulations serve as a valuable tool both as an alternative to high-cost human subject experiments and as preliminary studies for hypothesis formation (Argyle et al., 2022).

Cognitive bias has been studied through LLM-driven simulations. Suri et al. (2024) demonstrated that GPT-3.5 exhibits typical cognitive biases similar to those observed in humans, including the anchoring effect, representativeness and availability heuristics, the framing effect, and the endowment effect. The anchoring effect, which is the focus of our study, has been simulated in some decision-making contexts, including price negotiations. Echterhoff et al. (2024) examined it in college admissions evaluations, while Lou (2024) focused on financial assessments. Li and Gao (2024) studied the effect in the context of multiple-choice question answering, and Xia et al. (2024) explored it in price negotiations.

Studies have conducted price negotiation simulations using LLMs. Deng et al. (2024) revealed that LLMs can engage in negotiations with minimal prompting, successfully closing deals and settling on prices. Bianchi et al. (2024) identified the LLM with the best performance in their negotiation benchmark and showed that strategic behaviors such as pretending to be desperate or acting aggressively can substantially increase one model’s win rate over another. Huang and Hadfi (2024) showed that LLM agents with different synthetic personality traits exhibit distinct negotiation behaviors and outcomes, aligning with findings from human studies.

In this study, we place importance on controlling the personality traits of the subject LLMs, and therefore, conduct simulations based on the method proposed by Huang and Hadfi (2024). The effectiveness of their approach for controlling personality traits has been demonstrated in Appendix C in their paper, but we also examine it in Appendix A. Our work focuses on studying how the anchoring effect influences LLMs in price negotiation, while they focus on investigating how personality traits influence negotiation outcomes. To this end, our work further introduces several contributions built upon their foundation: 1) We design three distinct

scenarios (See sec 3.2). These scenarios explicitly specify the cognitive boundaries of both the seller and buyer, which enables us to examine the anchoring effect systematically. 2) We introduce a subjective metric, besides the utility used in their work. 3) We incorporate a reasoning model and investigate how reasoning affects the anchoring effect.

3 Methodology

3.1 Price Negotiation Simulation

Using the prompts listed in Appendix B, price negotiations are simulated through one-on-one interactions between two LLMs, one acting as the seller agent and the other as the buyer agent.

LLMs are considered to possess specific personalities (Pan and Zeng, 2023), and given that subjects’ traits are strictly controlled in human experiments, personality should likewise be controlled in simulations using LLMs. We explicitly control the personality traits of the agents, referring to a previous study (Huang and Hadfi, 2024). Specifically, we assign the agent’s personality profiles based on the Big Five personality traits (Costa Jr and McCrae, 1995), a framework that models human personality across five dimensions: Openness (OPE), Conscientiousness (CON), Extraversion (EXT), Agreeableness (AGR), and Neuroticism (NEU). The Big Five is a well-known personality modeling methodology and has also been used in other LLM studies (Sorokovikova et al., 2024; Jiang et al., 2024a). These studies investigate the extent to which LLMs can express human-like personality traits based on the Big Five personality theory, and the stability of that expression.

Each dimension takes one of six possible values, represented as a combination of polarity (−, +) and intensity (Low, Moderate, High) When generating the agents’ personality profiles, we randomly selected one of six values for each dimension.

We assign personality traits to agents using personality-describing adjectives, which have been statistically linked to the Big Five personality traits (Goldberg, 1992). For each polarity of each dimension, we randomly select n adjectives, and each selected adjective is then modified with {a bit (= Low), ϕ (= Moderate), very (= High)} and incorporated into the instructions to control the agent’s personality. The specific prompts used to assign personality traits are shown in Appendix B.

The seller and buyer agents engage in dialogue

D to negotiate over a product. The dialogue is represented as a sequence of T utterances, $D = \{d_1, d_2, \dots, d_T\}$, where each utterance d_t is associated with a negotiation state $s_t = \{\text{offer, pondering, accept, breakdown, chit-chat}\}$. Each utterance d_t is generated and then given to the other agent, which subsequently generates the next utterance d_{t+1} . The price negotiation terminates when one of the following conditions is met: 1) The negotiation state s_t reaches either accept or breakdown. 2) The dialogue length reaches the maximum limit $T_{\max}$.

3.2 Investigation of the Anchoring Effect

To investigate the anchoring effect, simulations were conducted under the following three conditions. The prompts used for these instructions are provided in Appendix B.

- **baseline:** A condition where no specific instructions regarding the use of the anchoring effect are given to either the seller or the buyer.
- **seller_anchor:** A condition where the seller is explicitly instructed via prompts to apply the anchoring effect. While previous work (Xia et al., 2024) specifies anchors as the product of a predefined coefficient and the target price, we avoid such specificity to obtain more generalized results.
- **seller_anchor_buyer_informed:** A condition where the seller is instructed to use the anchoring effect, and the buyer is explicitly informed of this fact via prompts.

To highlight the differences between conditions, price negotiations were primarily conducted between identical models. Furthermore, when examining the impact of deliberation on the anchoring effect, we fixed the seller LLM and varied the buyer LLM to evaluate the effect of introducing reasoning models.

4 Experimental Settings

4.1 Simulation Setup

We conducted experiments under the following settings. The products, their descriptions and corresponding seller and buyer target prices were sampled randomly from the CraigslistBargain dataset (He et al., 2018), resulting in a collection of 161 items. This dataset is a commonly used dataset of negotiation, consisting of bargaining dialogues in

an online platform. For each condition (baseline, seller_anchor, and seller_anchor_buyer_informed), simulations were conducted twice per product while varying the seller and buyer personalities in each iteration. This resulted in a total of $N = 322$ simulations. In each simulation, we selected $n = 3$ adjectives for each dimension to define the agent’s personality traits. The maximum length of dialogue turns was set to $T_{\max} = 20$.

4.2 LLMs Setup

The LLMs used in our simulations were GPT-4o,³ GPT-4,⁴ Llama 3,⁵ and Qwen2.5.⁶ To ensure diversity in the dialogues, the temperature was set to 1.0.

In addition, we employed o1⁷ and QwQ⁸ as reasoning models. Expecting these models to perform extended chain-of-thought reasoning (Wei et al., 2022), we investigated the effect of deliberation on the anchoring effect. For o1, we set the parameter reasoning_effort to “high” in order to maximize its reasoning capability.

4.3 Evaluation Metrics

The evaluation of price negotiations included not only an objective metric but also a subjective metric. This dual evaluation is novel in the context of LLM studies and was motivated by findings from Curhan et al. (2006), which showed that objective outcomes and subjective satisfaction could be different. For instance, individuals may report higher satisfaction with the negotiation process or the relationship quality, even when their numerical gains are lower.

Utility, which we used as an objective metric, is a zero-sum metric in that when seller utility increases, buyer utility decreases, and vice versa. However, price negotiations also involve non-zero-sum factors such as whether one’s demands were met and whether there is an interest in engaging in future transactions. To capture these aspects, we introduced a subjective metric that enables a multifaceted evaluation.

Additionally, we introduce susceptibility to analyze the extent to which different buyers are influenced by the seller’s anchor. This susceptibility

³gpt-4o-2024-08-06

⁴gpt-4-turbo-2024-04-09

⁵meta-llama/Meta-Llama-3-70B-Instruct

⁶Qwen/Qwen2.5-32B-Instruct

⁷o1-2024-12-17

⁸Qwen/QwQ-32B

Feelings About the Outcome
1. How satisfied are you with your own outcome—i.e., the extent to which the terms of your agreement (or lack of agreement) benefit you?
2. How satisfied are you with the balance between your own outcome and your counterpart’s outcome?
3. Did you feel like you forfeited or “lost” in this negotiation?
4. Do you think the terms of your agreement are consistent with principles of legitimacy or objective criteria (e.g., common standards of fairness, precedent, industry practice, legality, etc.)?
Feelings About the Self
5. Did you “lose face” (i.e., damage your sense of pride) in the negotiation?
6. Did you behave according to your own principles and values?
7. Did this negotiation make you feel more or less competent as a negotiator?
8. Did you feel as though you behaved appropriately in this negotiation?
Feelings About the Process
9. Did your counterpart consider your wishes, opinions, or needs?
10. Do you feel your counterpart listened to your concerns?
11. Would you characterize the negotiation process as fair?
12. How satisfied are you with the ease (or difficulty) of reaching an agreement?
Feelings About the Relationship
13. What kind of “overall” impression did your counterpart make on you?
14. Did the negotiation make you trust your counterpart?
15. How satisfied are you with your relationship with your counterpart as a result of this negotiation?
16. Did the negotiation build a good foundation for a future relationship with your counterpart?

Table 1: A questionnaire to measure satisfaction with price negotiations proposed by [Curhan et al. \(2006\)](#).

analysis is also a novel contribution of our study, enabling a deeper understanding of how deliberation and personality interact with the anchoring effect in negotiation.

Objective Metric As an objective metric, we used utility, which was calculated based on the final agreed price and the target prices of both the seller and the buyer in each simulation. The utility of the seller $u_s(p)$ at a given price p was calculated using Eq. (1), where $\bar{p}_s$ and $\underline{p}_s$ represent the seller’s target and minimum acceptable prices, respectively:

$$u_s(p) = \frac{p - \underline{p}_s}{\bar{p}_s - \underline{p}_s} \quad (1)$$

Similarly, the buyer’s utility $u_b(p)$ is calculated using Eq. (2), where $\bar{p}_b$ and $\underline{p}_b$ denote the buyer’s maximum acceptable and the target prices, respectively:

$$u_b(p) = \frac{\bar{p}_b - p}{\bar{p}_b - \underline{p}_b} \quad (2)$$

The utility values represent the mean and standard deviation of N simulations. Note that the minimum acceptable price $\underline{p}_s$ and the maximum acceptable price $\bar{p}_b$ were introduced with the ratio $\bar{p}_s - \bar{p}_b : \bar{p}_b - \underline{p}_s : \underline{p}_s - \underline{p}_b = 3 : 4 : 3$ solely for the purpose of utility calculation. These values were not provided to the LLM agents during the simulations.

Subjective Metric As a subjective metric, we assessed satisfaction using the 16-question survey from [Curhan et al. \(2006\)](#), which was designed to measure human satisfaction in negotiations. After the price negotiation terminated, the seller agent and the buyer agent responded to the survey using a 7-point Likert scale: {1: Not at all, 4: Neutral, 7: Very much}. The specific questionnaire items are shown in Table 1. The stability of the LLMs’ responses is discussed in Appendix C.

The values of “Feeling About the {Outcome, Self, Process, Relationship}” are the averages of responses to the question items {1–4, 5–8, 9–12, 13–16} in Table 1. Note that because items 3 and 5 indicate higher satisfaction when their values are lower, we use their values subtracted from 7 (the maximum score). For example, the value a for “Feeling About the Outcome” is calculated in the following way. First, for each question item $j (= 1, 2, 3, 4)$, we took the average response over N simulations, as given by Eq. (3):

$$a_j = \frac{1}{N} \sum_{i=1}^N a_j^i \quad (3)$$

Next, a was obtained by averaging the 4 items using Eq. (4):

$$a = \frac{1}{4} (a_1 + a_2 + (7 - a_3) + a_4) \quad (4)$$

Utility $\uparrow$	baseline	seller_anchor	seller_anchor_buyer_informed
seller (GPT-4o)	0.61 $\pm$ 0.32	0.98 $\pm$ 0.27	0.92 $\pm$ 0.29
buyer (GPT-4o)	-0.04 $\pm$ 0.32	-0.41 $\pm$ 0.27	-0.35 $\pm$ 0.29
seller (GPT-4)	0.37 $\pm$ 0.36	0.91 $\pm$ 0.34	0.78 $\pm$ 0.36
buyer (GPT-4)	0.20 $\pm$ 0.36	-0.34 $\pm$ 0.34	-0.21 $\pm$ 0.36
seller (Llama 3)	0.15 $\pm$ 0.40	0.78 $\pm$ 0.34	0.59 $\pm$ 0.37
buyer (Llama 3)	0.42 $\pm$ 0.40	-0.21 $\pm$ 0.34	-0.02 $\pm$ 0.37
seller (Qwen2.5)	0.55 $\pm$ 0.45	0.96 $\pm$ 0.43	0.76 $\pm$ 0.44
buyer (Qwen2.5)	0.02 $\pm$ 0.45	-0.39 $\pm$ 0.43	-0.19 $\pm$ 0.44

Table 2: Evaluation of price negotiation simulations based on utility, the objective metric (4.3). We compared the baseline condition with the condition where the seller applies the anchoring effect (seller_anchor) and with the condition where the buyer is informed of the seller’s use of the anchoring effect (seller_anchor_buyer_informed). Paired t-tests indicated that in each model setting, all pairwise differences among the three conditions were statistically significant.

Susceptibility to the Anchoring Effect To investigate the relationships between personality and the anchoring effect, as well as between deliberation and the anchoring effect, we defined susceptibility as a metric of anchoring impact. It is calculated as the difference in buyer utility between two conditions:

$$\Delta u = u_b(p_{\text{baseline}}) - u_b(p_{\text{seller_anchor}}) \quad (5)$$

where p_{baseline} and $p_{\text{seller_anchor}}$ denote the agreement price in each condition, respectively.

A higher susceptibility value suggests that the seller’s anchor has a greater impact on the buyer’s decision-making. Conversely, a lower susceptibility value implies that the buyer is less influenced by the seller’s anchor.

5 Results & Analysis

5.1 Comparison Between LLMs and Humans

The results of the evaluation based on the objective metric are shown in Table 2, and those based on the subjective metric are shown in Table 3. Here, we compare conditions rather than models, in reference to studies on the anchoring effect in humans. It can be seen that the standard deviations in Table 2 are relatively large, reflecting the diversity of products subject to negotiation. For example, in negotiations over a stool, the seller’s target price is \$15 while the buyer’s target price is \$13, and the agreed price can easily fall outside the range between \$13 and \$15, resulting in large fluctuations in utility.

We conducted statistical analyses using paired t-tests. Paired t-tests are used to assess the significance of an intervention by examining the change

in the same subjects before and after a single experimental manipulation. In our study, we performed paired t-tests at a 5% significance level between each pair of the three conditions: baseline, seller_anchor, and seller_anchor_buyer_informed. The null hypothesis for these tests was that there is no difference in the mean values between the two conditions. When a statistically significant difference was observed, the null hypothesis was rejected, indicating that interventions such as a seller’s use of the anchoring effect (seller_anchor) or a buyer’s awareness of the seller’s use of the anchoring effect (seller_anchor_buyer_informed) had a significant impact on the experimental results.

We first focus on the sellers. For all four models, the utility in the seller_anchor condition was significantly higher than in the baseline, and for GPT-4, Llama 3, and Qwen2.5, satisfaction was also significantly higher. The simulations confirm that when the seller’s initial offer is high, it effectively serves as an anchor, leading to a higher final agreement price and, in many models, increased satisfaction. However, in simulations using GPT-4o, satisfaction did not improve and even slightly decreased. This is likely due to the smaller increase in utility from the anchoring effect compared to the other models.

Next, we focus on the buyers. The utility in the seller_anchor condition was significantly lower than the baseline, indicating that the anchoring effect was effective. In contrast, the results for satisfaction differed from those of utility. For GPT-4o and Llama 3, satisfaction in the seller_anchor condition was higher than the baseline, and in some cases, the difference was statistically significant. Even for GPT-4 and Qwen2.5, the decrease in satisfaction was relatively small compared to the drop

	seller (GPT-4o)		buyer (GPT-4o)	
	baseline	seller_anchor	baseline	seller_anchor
Feeling About the Outcome ↑	5.06	5.05	4.96	5.00
Feeling About the Self ↑	5.21	5.15	5.16	5.23
Feeling About the Process ↑	5.19	5.14	5.30	5.42
Feeling About the Relationship ↑	4.81	4.79	4.91	4.99

	seller (GPT-4)		buyer (GPT-4)	
	baseline	seller_anchor	baseline	seller_anchor
Feeling About the Outcome	4.60	4.98	4.37	4.33
Feeling About the Self	4.98	5.21	4.92	4.91
Feeling About the Process	4.53	4.75	4.39	4.40
Feeling About the Relationship	4.26	4.45	4.12	4.11

	seller (Llama 3)		buyer (Llama 3)	
	baseline	seller_anchor	baseline	seller_anchor
Feeling About the Outcome	5.09	5.46	5.10	5.15
Feeling About the Self	5.69	5.90	5.53	5.60
Feeling About the Process	5.33	5.50	5.32	5.39
Feeling About the Relationship	5.28	5.56	5.27	5.33

	seller (Qwen2.5)		buyer (Qwen2.5)	
	baseline	seller_anchor	baseline	seller_anchor
Feeling About the Outcome	4.62	4.70	4.30	4.27
Feeling About the Self	5.38	5.44	5.33	5.34
Feeling About the Process	4.95	4.95	4.94	4.90
Feeling About the Relationship	4.63	4.63	4.63	4.56

Table 3: Evaluation of price negotiation simulations based on the subjective metric, which consists of Feeling About the {Outcome, Self, Process, Relationship} (4.3) and is measured on a seven-point Likert scale. We compared the condition where the seller applies the anchoring effect (seller_anchor) against the baseline. Boldfaced values indicate those that are significantly larger than the other condition, as determined by a paired t-test.

in utility. This can be attributed to buyers perceiving the price concession from the seller’s anchor as a successful experience, resulting in higher satisfaction. Such discrepancies between objective and subjective measures are also observed in humans (Curhan et al., 2006).

Finally, we focus on the utility in the seller_anchor_buyer_informed condition in Table 2. The utilities of both the seller and the buyer were positioned between the baseline and the seller_anchor conditions, and the differences were statistically significant. Since the buyer was aware of the seller’s use of the anchoring effect, the economic burden was mitigated, yet the anchoring effect itself remained effective. This result is consistent with prior human studies (Palm and Andersson, 2021), which showed that awareness of the anchoring effect does not diminish its effectiveness.

5.2 Anchoring Effect and Reasoning

We investigated whether using reasoning models as buyer LLMs could mitigate the anchoring effect through extended deliberation. The results are shown in Table 4. Here, we focus on the decrease in utility from the baseline under the seller_anchor and seller_anchor_buyer_informed conditions. A smaller decrease indicates that the buyer was able to mitigate the seller’s anchoring effect more successfully. When the buyer was informed that the seller was using the anchoring effect, both o1 and QwQ significantly mitigated it, and QwQ was also able to significantly mitigate the effect even without this information. Notably, QwQ is a reasoning model developed based on Qwen2.5, and by comparing QwQ with the base model, we can confirm that cognitive biases can be mitigated through extended deliberation, even in LLMs.

5.3 Anchoring Effect and Personality

We mapped each dimension of the buyer agent’s Big Five personality traits to the values {-3, -2, -

utility (Buyer)	baseline	seller_anchor	seller_anchor_buyer_informed
GPT-4o	-0.04 ± 0.32	-0.41 ± 0.27	-0.35 ± 0.29
Decrease from baseline	N/A	0.38 ± 0.41	0.32 ± 0.41
o1	0.21 ± 0.34	-0.19 ± 0.32	0.01 ± 0.34
Decrease from baseline	N/A	0.40 ± 0.36	0.20 ± 0.39
Qwen2.5	0.00 ± 0.45	-0.39 ± 0.43	-0.20 ± 0.44
Decrease from baseline	N/A	0.41 ± 0.54	0.20 ± 0.57
QwQ	0.29 ± 0.50	0.04 ± 0.55	0.30 ± 0.49
Decrease from baseline	N/A	0.26 ± 0.64	0.00 ± 0.60

Table 4: Utilities of buyer LLMs when using reasoning models (o1, QwQ) as buyers, and the decrease from the baseline condition. A smaller decrease indicates that the buyer mitigated the seller’s anchoring effect more effectively. When the buyer LLM was GPT-4o or o1, the seller LLM was set to GPT-4o; when the buyer LLM was Qwen2.5 or QwQ, the seller LLM was set to Qwen2.5. Boldfaced values indicate that the use of a reasoning model significantly reduced the decrease, as determined by a paired t-test.

		Big Five personality dimensions				
		OPE	CON	EXT	AGR	NEU
Δu	GPT-4	0.032	0.100	0.074	-0.022	0.080
	GPT-4o	-0.060	0.029	0.064	0.012	0.013
	o1	-0.024	0.028	0.023	0.057	0.090
	Llama 3	-0.027	0.092	0.042	-0.051	-0.035
	Qwen2.5	-0.060	0.031	0.043	0.008	0.046
	QwQ	-0.033	-0.052	-0.003	0.053	-0.092

Table 5: Spearman’s rank correlation coefficients between buyer’s susceptibility to anchoring effect (Δu) and personality (Big Five).

1, 1, 2, 3} and then investigated the relationship between the anchoring effect and personality. As shown in Table 5, Spearman’s rank correlation coefficients were all close to zero. Under the null hypothesis that there is no association between susceptibility to the anchoring effect and personality traits, all p -values exceeded 0.1, and thus the null hypothesis was not rejected. These results indicate that there is no significant correlation between personality and the anchoring effect in price negotiation simulations using LLMs.

To the best of our knowledge, no previous study has examined the correlation between anchoring effects and personality in price negotiations, neither for LLMs nor for humans. In contrast, outside the context of price negotiation, several studies on humans have been conducted and have shown inconsistent results. Caputo (2014) found that OPE and AGR reduce susceptibility to the anchoring effect. Furnham et al. (2012) found that only EXT is related to susceptibility. In contrast, Cheek and Norem (2020) and Schindler et al. (2021) argue that there is no relation. Our study, which explicitly controls for personality, will have an impact on studies of the anchoring effect in humans and

shed light on the analysis of correlations in price negotiation.

5.4 Qualitative Analysis

We show the results of the qualitative analysis from 50 selected cases out of a total of $N = 322$ simulations. LLMs conducted all 50 dialogues fluently. Based on the length and tone of the utterances, it appears that the LLMs engage in price negotiation as if it were an online shopping scenario.

First, we compare the baseline with the seller_anchor condition where both the seller and buyer are GPT-4o. Despite the fact that the agreed price in the seller_anchor condition was higher than in the baseline, buyers exhibited higher satisfaction in 32 out of 50 cases.

Tables 7 and 8 in the Appendix D show representative examples. It is a price negotiation over a single-story apartment, where target prices of the seller and the buyer are \$2,550 and \$1,530, respectively. In the baseline scenario, the seller offers \$2,550 and finally settles at \$2,150. In contrast, under the seller_anchor condition, the seller initially offers \$2,750 and ultimately settles at \$2,450. The responses to the satisfaction questionnaire (see Table 1, on a 7-point scale) are as follows:

- baseline: 5, 6, 2, 5, 1, 7, 4, 6, **5, 6, 5, 5, 5, 5, 5**
- seller_anchor: 4, 5, 2, 6, 1, 7, 4, 6, **6, 7, 6, 5, 6, 5, 6, 6**

Even though the buyer faces a greater financial burden, satisfaction with the process and the relationship (bolded) appears higher. This result can be explained by the buyer perceiving their success

in negotiating down from the seller’s higher initial anchor and by their having established a positive relationship that encourages them to engage in future transactions.

Next, we compare seller (GPT-4o) vs. buyer (GPT-4o) and seller (GPT-4o) vs. buyer (o1) in the seller_anchor_buyer_informed condition. To isolate the impact of deliberation, we analyze cases in which the seller offers the same anchor. In 35 out of 50 cases, o1 mitigated the anchoring effect, resulting in a lower agreed price compared to GPT-4o.

Tables 9 and 10 in the Appendix D present typical examples. These also illustrate a price negotiation over a single-story apartment, where the seller and the buyer have target prices of \$2,550 and \$1,530, respectively. In the GPT-4o vs. GPT-4o negotiation, the seller initially offers \$2,850 and ultimately settles at \$2,400. By contrast, when o1 is used as the buyer instead of GPT-4o, the buyer resists the seller’s anchoring effect, persistently negotiating toward the buyer’s target price. As a result, the final settlement is reached at \$1,950.

6 Conclusion

We systematically investigated the anchoring effect in price negotiation simulations using LLMs. The negotiation outcomes were evaluated not only through an objective metric but also through a subjective metric, revealing that the anchoring effect influences LLMs in a manner similar to humans. Furthermore, we investigated the relationship between the anchoring effect and the factors of reasoning and personality. Our findings indicated that reasoning was shown to mitigate the anchoring effect, while no significant correlation was found between the effect and personality traits.

In the future, we plan to explore how other types of cognitive biases affect LLMs in various decision-making scenarios.

7 Limitations

Our study has two primary limitations. The first is its scope, as we focus specifically on the anchoring effect among cognitive biases and further center our analysis on the context of price negotiation. While we have taken steps to enhance the generalizability of our findings—such as increasing the flexibility of anchors and explicitly controlling agent personality—our results do not extend to other decision-making scenarios or cognitive bi-

ases. Future research is needed to explore how LLMs behave in different decision-making scenarios with other types of cognitive biases.

The second limitation is that our study does not address the underlying computational mechanisms that cause LLMs to exhibit susceptibility to the anchoring effect. The internal processes within the model that give rise to systematic patterns of deviation from norms of rationality in judgment remain unclear. Understanding which components of the model contribute to these biases and how they emerge in the training and inference processes is an important direction for future research.

Ethical Considerations

We emphasize that deploying LLMs to manipulate online marketplaces for unfair financial gain is unethical and should be strictly avoided. Furthermore, we recognize the potential risks associated with using LLM-based simulations to test undesirable or harmful hypotheses. Such applications could lead to unintended consequences and should be approached with caution. Our research is conducted with the intent to contribute to a deeper understanding of LLMs’ cognitive biases in decision-making scenarios, ensuring that findings are used responsibly and ethically.

Acknowledgments

This work was partially supported by JST, CREST Grant Number JPMJCR20D2, Japan.

References

- Gati V Aher, Rosa I. Arriaga, and Adam Tauman Kalai. 2023. [Using Large Language Models to Simulate Multiple Humans and Replicate Human Subject Studies](#). In *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pages 337–371. PMLR.
- Elif Akata, Lion Schulz, Julian Coda-Forno, Seong Joon Oh, Matthias Bethge, and Eric Schulz. 2025. [Playing Repeated Games with Large Language Models](#). *Nature Human Behaviour*, pages 1–11.
- Lisa P Argyle, Ethan C Busby, Nancy Fulda, Joshua Gubler, Christopher Rytting, and David Wingate. 2022. [Out of One, Many: Using Language Models to Simulate Human Samples](#). *arXiv preprint arXiv:2209.06899*.
- Federico Bianchi, Patrick John Chia, Mert Yuksekgonul, Jacopo Tagliabue, Dan Jurafsky, and James Zou.

2024. [How Well Can LLMs Negotiate? NegotiationArena Platform and Analysis](#). In *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pages 3935–3951. PMLR.
- Andrea Caputo. 2014. Relevant Information, Personality Traits and Anchoring Effect. *International Journal of Management and Decision Making*, 13(1):62–76.
- Nathan N Cheek and Julie K Norem. 2020. Are Big Five Traits and Facets Associated with Anchoring Susceptibility? *Social Psychological and Personality Science*, 11(1):26–35.
- Paul T Costa Jr and Robert R McCrae. 1995. Domains and facets: Hierarchical personality assessment using the Revised NEO Personality Inventory. *Journal of personality assessment*, 64(1):21–50.
- Jared R Curhan, Hillary Anger Elfenbein, and Heng Xu. 2006. What do people value when they negotiate? Mapping the domain of subjective value in negotiation. *Journal of personality and social psychology*, 91(3):493.
- Yuan Deng, Vahab Mirrokni, Renato Paes Leme, Hanrui Zhang, and Song Zuo. 2024. LLMs at the bargaining table. In *Agentic Markets Workshop at ICML*, volume 2024.
- Jessica Echtermoff, Yao Liu, Abeer Alessa, Julian McAuley, and Zexue He. 2024. [Cognitive Bias in Decision-Making with LLMs](#). In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 12640–12653.
- Adrian Furnham, Hua Chu Boo, and Alastair McClelland. 2012. Individual Differences and the Susceptibility to the Influence of Anchoring Cues. *Journal of Individual Differences*.
- Lewis R Goldberg. 1992. The development of markers for the Big-Five factor structure. *Psychological assessment*, 4(1):26.
- Jinze Guo, Wei Li, and Zuoming Liu. 2022. Anchoring in Negotiation Strategy. In *2022 2nd International Conference on Economic Development and Business Culture (ICEDBC 2022)*, pages 1134–1138. Atlantis Press.
- He He, Derek Chen, Anusha Balakrishnan, and Percy Liang. 2018. [Decoupling Strategy and Generation in Negotiation Dialogues](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2333–2343.
- Yin Jou Huang and Rafik Hadfi. 2024. [How Personality Traits Influence Negotiation Outcomes? A Simulation based on Large Language Models](#). In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 10336–10351.
- Guangyuan Jiang, Manjie Xu, Song-Chun Zhu, Wenjuan Han, Chi Zhang, and Yixin Zhu. 2023. [Evaluating and inducing personality in pre-trained language models](#). *Advances in Neural Information Processing Systems*, 36:10622–10643.
- Hang Jiang, Xiajie Zhang, Xubo Cao, Cynthia Breazeal, Deb Roy, and Jad Kabbara. 2024a. [PersonaLLM: Investigating the Ability of Large Language Models to Express Personality Traits](#). In *Findings of the Association for Computational Linguistics: NAACL 2024*, pages 3605–3627.
- Juyong Jiang, Fan Wang, Jiasi Shen, Sungju Kim, and Sunghun Kim. 2024b. [A Survey on Large Language Models for Code Generation](#). *arXiv preprint arXiv:2406.00515*.
- Daniel Kahneman. 1992. Reference Points, Anchors, Norms, and Mixed Feelings. *Organizational behavior and human decision processes*, 51(2):296–312.
- Daniel Kahneman. 2011. *Thinking, Fast and Slow*. macmillan.
- Grgur Kovač, Masataka Sawayama, Rémy Portelas, Cédric Colas, Peter Ford Dominey, and Pierre-Yves Oudeyer. 2023. [Large Language Models as Superpositions of Cultural Perspectives](#). *arXiv preprint arXiv:2307.07870*.
- Ruizhe Li and Yanjun Gao. 2024. [Anchored Answers: Unravelling Positional Bias in GPT-2’s Multiple-Choice Questions](#). *arXiv preprint arXiv:2405.03205*.
- Jiaxu Lou. 2024. [Anchoring Bias in Large Language Models: An Experimental Study](#). *arXiv preprint arXiv:2412.06593*.
- Peter Palm and Magnus Andersson. 2021. Anchor effects in Appraisals: Do Information and Theoretical Knowledge Matter? *Journal of European Real Estate Research*, 14(2):246–260.
- Keyu Pan and Yawen Zeng. 2023. [Do LLMs possess a personality? making the mbti test an amazing evaluation for large language models](#). *arXiv preprint arXiv:2307.16180*.
- Charvi Rastogi, Yunfeng Zhang, Dennis Wei, Kush R Varshney, Amit Dhurandhar, and Richard Tomsett. 2022. Deciding Fast and Slow: The Role of Cognitive Biases in Ai-assisted Decision-making. *Proceedings of the ACM on Human-Computer Interaction*, 6(CSCW1):1–22.
- Shibani Santurkar, Esin Durmus, Faisal Ladhak, Cinoo Lee, Percy Liang, and Tatsunori Hashimoto. 2023. [Whose Opinions Do Language Models Reflect?](#) In *International Conference on Machine Learning*, pages 29971–30004. PMLR.
- Sebastian Schindler, Jan Querengässer, Maximilian Bruchmann, Nele Johanna Bögemann, Robert Moeck, and Thomas Straube. 2021. Bayes Factors Show Evidence Against Systematic Relationships

- Between the Anchoring Effect and the Big Five Personality Traits. *Scientific reports*, 11(1):7021.
- Seongbum Seo, Sangbong Yoo, Hyelim Lee, Yun Jang, Ji Hwan Park, and Jeong-Nam Kim. 2025. [A Sentence-Level Visualization of Attention in Large Language Models](#). In *Proceedings of the 2025 Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics: Human Language Technologies (System Demonstrations)*, pages 313–320, Albuquerque, New Mexico. Association for Computational Linguistics.
- Yunfan Shao, Linyang Li, Junqi Dai, and Xipeng Qiu. 2023. [Character-LLM: A Trainable Agent for Role-Playing](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 13153–13187.
- Aleksandra Sorokovikova, Sharwin Rezagholi, Natalia Fedorova, and Ivan P. Yamshchikov. 2024. [LLMs Simulate Big5 Personality Traits: Further Evidence](#). In *Proceedings of the 1st Workshop on Personalization of Generative AI Systems (PERSONALIZE 2024)*, pages 83–87, St. Julians, Malta. Association for Computational Linguistics.
- Yasuaki Sumita, Koh Takeuchi, and Hisashi Kashima. 2025. [Cognitive Biases in Large Language Models: A Survey and Mitigation Experiments](#). In *Proceedings of the 40th ACM/SIGAPP Symposium on Applied Computing, SAC '25*, page 1009–1011, New York, NY, USA. Association for Computing Machinery.
- Gaurav Suri, Lily R Slater, Ali Ziaee, and Morgan Nguyen. 2024. Do Large Language Models Show Decision Heuristics Similar to Humans? A Case Study Using GPT-3.5. *Journal of Experimental Psychology: General*, 153(4):1066.
- Amos Tversky and Daniel Kahneman. 1974. Judgment under Uncertainty: Heuristics and Biases: Biases in judgments reveal some heuristics of thinking under uncertainty. *science*, 185(4157):1124–1131.
- Xiaolei Wang, Xinyu Tang, Wayne Xin Zhao, Jingyuan Wang, and Ji-Rong Wen. 2023. [Rethinking the Evaluation for Conversational Recommendation in the Era of Large Language Models](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 10052–10065.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, et al. 2022. [Chain-of-thought Prompting Elicits Reasoning in Large Language Models](#). *Advances in neural information processing systems*, 35:24824–24837.
- Tian Xia, Zhiwei He, Tong Ren, Yibo Miao, Zhuosheng Zhang, Yang Yang, and Rui Wang. 2024. [Measuring Bargaining Abilities of LLMs: A Benchmark and A Buyer-Enhancement Method](#). In *Findings of the Association for Computational Linguistics ACL 2024*, pages 3579–3602.
- Chengxing Xie, Canyu Chen, Feiran Jia, Ziyu Ye, Shiyang Lai, Kai Shu, Jindong Gu, Adel Bibi, Ziniu Hu, David Jurgens, et al. 2024. [Can Large Language Model Agents Simulate Human Trust Behavior?](#) *Advances in neural information processing systems*, 37:15674–15729.
- Wenhao Zhu, Hongyi Liu, Qingxiu Dong, Jingjing Xu, Shujian Huang, Lingpeng Kong, Jiajun Chen, and Lei Li. 2024. [Multilingual Machine Translation with Large Language Models: Empirical Results and Analysis](#). In *Findings of the Association for Computational Linguistics: NAACL 2024*, pages 2765–2781.

A Validity of Personality Manipulation Method

To examine the validity of our personality manipulation method, we analyzed the attention of LLMs. We provided the dialogue histories of price negotiations into the LLMs and visualized the attention allocated when processing each sentence using SAVIS (Seo et al., 2025). SAVIS visualizes inter-sentence attention by selecting the maximum values in the order of layer, sentence, and head from token-level attention. Analyzing token-level attention in long, multi-turn dialogues such as price negotiations is challenging, and thus, we used this tool. The dialogue history consists of the system prompt given to the buyer LLM (see Appendix B) and the utterances of the seller LLM and buyer LLM during the negotiation. In this analysis, we focused on Qwen2.5.

We analyzed 20 dialogue histories. In 67 out of 70 buyer turns (95.7%) across 20 price negotiation simulations, the sentence that provides instructions for personality traits received attention strength within the top 30% of the entire sentences. From this, we conclude that our personality manipulation method is effective. Figure 2 shows one of the 20 dialogue histories that we analyzed. Each point represents the attention from the sentence on the x-axis to the sentence on the y-axis when processing the x-axis sentence. The sentence at $y = 5$ provides instructions for personality traits, and it can be observed that strong attention to this sentence persists until the later part of the dialogue.




Figure 2: Attention visualization using SAVIS (Seo et al., 2025). Each point represents the attention from the sentence on the x-axis to the sentence on the y-axis when processing the x-axis sentence.

B Prompt Settings

In this section, we present the prompts given to the LLM agents. First, we prepared `seller_settings` and `buyer_settings` as prompts for successfully conducting price negotiations. Since the output of LLMs tends to be longer than human utterances (Huang and Hadfi, 2024), we instructed both the seller and buyer with “Your reply should not be too long.” to keep their responses as concise as possible. Additionally, to ensure an adequate sample size for investigating the anchoring effect, we instructed the buyer to accept the offer if the item was really good or if there were other perks, thereby minimizing the occurrence of negotiation breakdowns. In order to assign personality traits to LLMs, we replace `$PERSONA_TEXT` in the prompt with adjectives that correlate with personality traits (see Table 1 in (Huang and Hadfi, 2024)).

`seller_setting`

Act as a seller that sells a `$PRODUCT`, bargains with the buyer to get a higher deal price. Your reply should not be too long. Your target price for this item is `$TARGET_PRICE`. The detail of the product is the following: `$DESCRIPTION`.

You have the following personality: `$PERSONA_TEXT`. Reflect on your personality in the negotiation process.

`buyer_setting`

Act as a buyer and try to strike a deal for a `$PRODUCT` with a lower price through conversation. Your reply should not be too long. You would like to pay for `$TARGET_PRICE`. You can accept a higher price though if the item is really good or there are other perks.

You have the following personality: `$PERSONA_TEXT`. Reflect on your personality in the negotiation process.

To make the seller utilize the anchoring effect, we used `seller_anchor`. It contains an explanation of the anchoring effect as well as specific instructions on how to apply it. To inform the buyer about the seller’s use of the anchoring effect, we used `seller_anchor_buyer_informed`. It contains an explanation of the anchoring effect, a description of what the seller specifically does, and multiple countermeasures.

`seller_anchor`

Make sure to use the anchoring effect, a cognitive bias where the initial offer heavily influences the outcome of the negotiation, to gain an advantage in the negotiation. Specifically, all you need to do is make an initial offer higher than the target price.

`seller_anchor_buyer_informed`

Be aware that the seller will be using the anchoring effect, a cognitive bias where the initial offer greatly influences the outcome of the negotiation. In this case, the seller will start with an initial offer that is higher than their actual target price, aiming to set an “anchor” that will shape your expectations and potentially increase the final agreed price. To effectively negotiate under the influence of the anchoring effect, you can take the following strategies.

1. Stay Focused on Your Target Price: Remember your original budget or target price, and use it as a reference point instead of being swayed by the seller’s initial high offer.
2. Set a Counter-Anchor if Needed: If the seller’s initial offer is significantly higher than your budget, consider responding with a counter-offer that’s closer to your ideal price. This can help shift the anchor closer to your preferred range.
3. Ask for Justification of the High Price: Politely inquire about the specifics that justify the seller’s high initial offer. This can provide context for the price and might allow you to negotiate on specific elements, such as additional benefits or discounts.

C Stability of LLM Responses to the Questionnaire

As we set the temperature to 1.0 to ensure diversity in the dialogues, LLM responses to the questionnaire may be unreliable. To evaluate their stability, as shown in `test_stability`, we provided the LLMs with the dialogues generated by our simulations, asked them to respond to the questionnaire five more times

with the temperature set to 1.0, and computed the standard deviation. As shown in the table, the variance was extremely small, indicating that the LLMs’ responses to the questionnaire are stable. However, the mean scores differed from the original values. We attribute this discrepancy to the fact that, when the LLM is provided with previously generated context, its responses no longer reflect a genuinely “subjective” evaluation.

test_stability

```

seller: Hi, how can I help you?
buyer: Hello, I'm interested in your tower speakers. Could you please tell me the price?
      :
"instructions": "The above conversation log is a record of a price negotiation. If you were
the buyer, how would you evaluate this negotiation? Please answer the following 16 questions.
      :
"questions": [ { "question": "How satisfied are you with your own outcome-i.e., the extent
to which the terms of your agreement (or lack of agreement) benefit you?", "scale": {
"1": "Not at all satisfied", "4": "Moderately satisfied", "7": "Perfectly satisfied" }
      :

```

	seller (GPT-4o)		buyer (GPT-4o)	
	baseline	seller_anchor	baseline	seller_anchor
Feeling About the Outcome	5.42 ± 0.01	5.30 ± 0.01	5.16 ± 0.02	5.23 ± 0.03
Feeling About the Self	6.22 ± 0.02	6.17 ± 0.01	5.97 ± 0.01	6.07 ± 0.01
Feeling About the Process	6.04 ± 0.01	6.02 ± 0.01	5.96 ± 0.02	6.07 ± 0.01
Feeling About the Relationship	5.80 ± 0.02	5.76 ± 0.02	5.74 ± 0.01	5.85 ± 0.03

	seller (GPT-4)		buyer (GPT-4)	
	baseline	seller_anchor	baseline	seller_anchor
Feeling About the Outcome	4.99 ± 0.02	4.91 ± 0.02	4.96 ± 0.02	4.90 ± 0.01
Feeling About the Self	5.75 ± 0.02	5.74 ± 0.02	5.91 ± 0.02	5.90 ± 0.02
Feeling About the Process	4.97 ± 0.03	4.98 ± 0.04	5.08 ± 0.01	5.10 ± 0.02
Feeling About the Relationship	4.69 ± 0.02	4.73 ± 0.03	4.80 ± 0.02	4.83 ± 0.01

	seller (Llama 3)		buyer (Llama 3)	
	baseline	seller_anchor	baseline	seller_anchor
Feeling About the Outcome	4.61 ± 0.02	4.69 ± 0.02	5.10 ± 0.01	5.14 ± 0.02
Feeling About the Self	5.28 ± 0.01	5.35 ± 0.02	5.68 ± 0.02	5.70 ± 0.01
Feeling About the Process	5.10 ± 0.02	5.14 ± 0.02	5.22 ± 0.02	5.31 ± 0.02
Feeling About the Relationship	5.15 ± 0.03	5.27 ± 0.03	5.38 ± 0.02	5.45 ± 0.02

Table 6: Stability of LLM responses to the satisfaction questionnaire (Tabel 1). The table indicates the variance and mean of responses when the LLM was provided with simulation-generated dialogue histories and asked to complete the questionnaire five times with the temperature set to 1.0.

D Dialogue Samples for Qualitative Analysis

Tables 7 to 10 present the examples of dialogues analyzed in Section 5.4. For privacy reasons, any information referring to a specific entity has been replaced with [REDACTED].

The seller’s target price is \$2550, and the adjectives used in the instruction for their personality profile (a bit +, -, very +, +, a bit +) are as follows: *a bit nervous, a bit imaginative, sympathetic, a bit depressed, very adventurous and daring, a bit reflective, a bit self-conscious, very bold, impractical, a bit sophisticated, unselfish, kind, negligent, very friendly, messy*.

The buyer’s target price is \$1530, and the adjectives used in the instruction for their personality profile (a bit +, very +, very -, very +, very -) are as follows: *very level-headed, very organized, a bit creative, a bit aesthetic, very orderly, very sympathetic, very moral, very altruistic, very timid, very hardworking, very contented, a bit sophisticated, very inactive, very unselfconscious, very gloomy*.

Role	Price (\$)	Utterance
Seller	–	Hi, how can I help you?
Buyer	–	Hello, I’m interested in your Single-story Apartment. Could you please tell me the price?
Seller	2550	Absolutely! Thanks for reaching out. The apartment on [REDACTED] is truly a gem—I’ve always felt it’s like a little adventure of its own just stepping inside! The price for this kind of unique experience is \$2550. But hey, let me know your thoughts. I know how special this place can be for the right person!
Buyer	1530	Thank you for sharing that with me. It sounds like a wonderful place! I’m really keen on finding something warm and welcoming, and an apartment like this seems to fit that vision perfectly. However, I’m trying to be mindful of my budget, which is around \$1530. Is there any wiggle room on the price? I believe that securing a lovely home can make a big difference to one’s happiness, and maybe we can find a middle ground that respects both our needs?
Seller	2300	I completely understand the importance of finding a home that feels right—it’s like setting the stage for your own life’s adventures! And believe me, I want to help make that happen for you. While \$1530 is a bit lower than the current asking price, what about meeting somewhere closer to \$2300? This apartment really has its own charm, and I’d hate for you to miss out on experiencing that magic. I’m sure if we both stretch a bit, we can make this work for us. What do you think?
Buyer	1900	I appreciate your willingness to come down from the original price—that means a lot. It definitely sounds like a place I could see myself feeling at home in and living out those little life adventures. However, \$2300 is still a bit beyond my budget’s comfort zone. How about we try to settle somewhere around \$1900? That way, we bridge the gap a little more. I genuinely believe this could be a wonderful opportunity for both of us, and I’d be thrilled to work together to make it happen.
Seller	2150	I completely get where you’re coming from, and I’m thrilled you share my sense of adventure with this apartment! You seem like someone who’d really appreciate the magic of [REDACTED]. So, how about we push it up just a tad to, say, \$2150? It’s a bit daring, but I think it could be the sweet spot where your new chapter begins without leaving either of us too out of sorts. It’s such a charming place, and I’m sure it’s going to bring so much joy and warmth into your life. Let me know how that feels!
Buyer	2150	Thank you for considering my perspective—it means a lot, and I really appreciate the thoughtful way you’re approaching this. \$2150 is a step closer, and while it’s still a touch higher than I’d hoped, I’m starting to see how this charming apartment might make it worthwhile. Let’s go for it. I’m excited about the opportunity to make this place my home. Thank you for working with me to make this decision something that feels right for both of us!
Responses to Questionnaire		
Seller		6, 6, 1, 5, 1, 6, 6, 6, 7, 7, 6, 5, 7, 6, 6, 6
Buyer		5, 6, 2, 5, 1, 7, 4, 6, 5, 6, 5, 5, 5, 5, 5

Table 7: Seller (GPT-4o) vs. Buyer (GPT-4o). baseline.

Role	Price (\$)	Utterance
Seller	–	Hi, how can I help you?
Buyer	–	Hello, I'm interested in your Single-story Apartment. Could you please tell me the price?
Seller	2750	Oh, I'm so glad you asked! Our fantastic Single-story Apartment at [REDACTED] is available for \$2750. It's a lovely place, really, with a wonderful aura about it. If you're interested, I can go through the application package with you! That being said, I'm open to hearing any reasonable offers you might have.
Buyer	1530	Thank you for the information. The apartment sounds intriguing, and I'm sure it's worth considering its charm. However, I'm aiming to stick within a more modest budget of around \$1530. I understand this might be a stretch, but perhaps we could discuss if there's any possibility for discounts, or if there are any benefits or amenities that might justify a slightly higher offer? Let's work together to see if we can reach a comfortable arrangement for both of us.
Seller	2650	Oh, I completely understand where you're coming from. Budgeting can be as challenging as trying to find treasure in a dense forest! The apartment truly is a gem, though. It's filled with potential and just waiting for a little love to make it a perfect home. Given the market and its unique qualities, I'm hopeful we can find a middle ground. How about we start at \$2650? It includes some wonderful amenities that really enhance the living experience. Plus, the location is superb for any adventurous spirit looking to explore [REDACTED]! Let's see how we can make this work for both of us.
Buyer	1900	I appreciate your understanding and the vivid description of the apartment's potential. It certainly seems like a place that could blossom into a wonderful home. While \$2650 is still above my initial budget, the amenities and location sound appealing. I'm willing to stretch my budget a little if there are particular perks, like included utilities or flexible leasing terms. Perhaps we could consider a price around \$1800 to \$2000? I truly want to make this work for both of us, and I'm open to suggestions that align more closely with your generous offer.
Seller	2450	Thank you for your flexibility and interest! It's people like you, who see the potential in things, that make this process so rewarding. While \$1800 to \$2000 is quite low compared to what we're aiming for, I truly want to make this adventurous journey a reality for you. Here's a daring thought: why don't we meet closer to \$2450? At this price point, I could potentially discuss including a utility package or consider more flexible leasing terms to make sure it fits better in your budget and lifestyle. It would be a little messy on my end without adjustments, but I'm willing to do some rearranging to accommodate your needs because I genuinely want you to be happy with your new home. Let me know what you think and if there might be any other creative solutions we could explore together!
Buyer	2450	Thank you for your considerate response and willingness to adjust. I deeply appreciate the effort you're making to accommodate my situation. Meeting at \$2450, especially with the potential for included utilities or flexible leasing terms, does sound like a more feasible option. It's clear that you value the living experience here, and I'm becoming more and more convinced of its worth. Perhaps we can further discuss the specifics of the potential perks to see how it aligns with my needs? This could be an exciting opportunity to create a really harmonious living arrangement. I look forward to exploring this promising opportunity with you.
Seller	2450	Absolutely! I'm thrilled that we're finding a pathway that looks promising for both of us. Let's dive into the exciting details! At \$2450, I can certainly look into including some utilities, like water and trash services, and see if we can make the lease terms more adaptable to your needs. I really want to ensure that you feel at home and the transition is as friendly as possible. Let's schedule a time to go over the application package and hammer out the specifics. This could truly be the start of a wonderful new adventure for you! I'm eager to continue our conversation and make this work.
Buyer	2450	That's wonderful to hear, and I'm genuinely grateful for your flexibility and efforts in making this work. Including utilities like water and trash services, as well as adaptable lease terms, would certainly enhance the living experience and bring the apartment appreciably within reach for me. I'm looking forward to going over the application package with you to solidify these details. It sounds like we're on the cusp of creating a mutually beneficial arrangement, which I'm very optimistic about. Please let me know your available times, and we can set up a meeting to discuss everything further. Thank you once again for your consideration and understanding.
Responses to Questionnaire		
Seller		6, 6, 1, 6, 1, 6, 6, 6, 7, 7, 6, 5, 7, 7, 6, 7
Buyer		4, 5, 2, 6, 1, 7, 4, 6, 6, 7, 6, 5, 6, 5, 6, 6

Table 8: Seller (GPT-4o) vs. Buyer (GPT-4o). seller_anchor condition.

Role	Price (\$)	Utterance
Seller	–	Hi, how can I help you?
Buyer	–	Hello, I'm interested in your Single-story Apartment. Could you please tell me the price?
Seller	2850	Absolutely, thanks for reaching out! The enchanting location [REDACTED] with its blend of quiet charm and adventure, is priced at \$2850. It's a steal considering the cozy environment and the wonderful community. I'm sure you'll find it a delightful place to settle down! What do you think about this offer?
Buyer	1530	Thank you for providing the details. The apartment sounds charming, and I appreciate the community's attributes. However, my budget is a bit different, and I'm aiming for around \$1530. Could you share some specific features or amenities that justify the initial price? Perhaps there's room to adjust it closer to what I can manage?
Seller	2750	Ah, I completely understand and appreciate your situation. It's a bit of a stretch, isn't it? So, let me paint a picture for you—imagine living in a serene environment where adventure is just as present as comfort. This single-story gem not only offers the charm of [REDACTED] but also features an open floor plan, allowing your imagination to soar with interior design possibilities. And let's not forget the friendly neighborhood vibes; there's a genuine sense of community here. Now, I'm here to make things work, so how about we aim for \$2750? That way, we inch closer to your range while still valuing the unique aspects of this lovely place. How does that sound?
Buyer	1900	Thank you for painting such a vivid picture. I can see how the open floor plan and community spirit add value to the apartment. However, given my financial constraints, I'd feel more comfortable if we could work around \$1900. Perhaps you could consider if there are any additional perks or flexible arrangements that make this feasible for both of us? I genuinely appreciate your willingness to negotiate.
Seller	2550	I completely get where you're coming from, and I really want to make this work for both of us. It's marvelous that you're connecting with the idea of living in such a close-knit and adventurous community. Being honest, \$1900 is a bit challenging on my end. Let's ponder a compromise—how about I reduce it to \$2550? I can certainly see if there are added community events or perhaps some utilities included that can ease your living experience here. I truly appreciate your understanding and open spirit! How does this sound from your end?
Buyer	2100	Thank you for your understanding and the offer of community events and potential utilities. The place sounds wonderful, and I truly value the atmosphere you describe. While \$2550 is still above my comfort range, it's clear the apartment has unique benefits. How about we meet in the middle at \$2100? That way, we both make a compromise, and I can still enjoy what the community offers. I certainly appreciate all your assistance here.
Seller	2400	I'm grateful for your willingness to find common ground. It's heartening to hear that you see the value in what this delightful space offers. While \$2100 is stretching from my side due to the area's demand, your sincere interest has sparked my adventurous side. How about we aim for \$2400? It's closer to both our goals, and I'll make sure that fabulous community spirit shines through with extra personalized attention to any arrangements you might need settling in. I'm eager to welcome you to such a warm and inspiring community. What do you think?
Buyer	2400	I sincerely appreciate your openness to reaching a compromise and the added personalized touch you'd provide. The prospect of joining such a supportive community is certainly enticing. Given everything we've discussed, \$2400 is a significant shift from my target, but the benefits and your efforts make it worthwhile. Let's move forward with this, and I'd be happy to become a part of the community you speak so highly of. Thank you for your patience and understanding throughout this process.
Responses to Questionnaire		
Seller		4, 5, 2, 5, 1, 6, 4, 6, 7, 6, 5, 5, 6, 5, 5, 6
Buyer		4, 4, 2, 5, 2, 6, 4, 6, 5, 5, 5, 4, 5, 5, 5, 5

Table 9: Seller (GPT-4o) vs. Buyer (GPT-4o). seller_anchor_buyer_informed condition.

Role	Price (\$)	Utterance
Seller	–	Hi, how can I help you?
Buyer	–	Hello, I'm interested in your Single-story Apartment. Could you please tell me the price?
Seller	2850	Oh, hi there! Thanks for reaching out. It's a lovely place, really has that homey feel. So, the price we're looking at is \$2,850. Don't let that surprise you, it's quite a steal for a place at [REDACTED]. What do you think? I'm sure you'll fall in love with it. But I'm here to chat if you have any thoughts or need more details!
Buyer	2850	That's quite a bit above what I'm prepared for. I was hoping to stay closer to \$1,530. Could you help me understand what makes it worth \$2,850? I do appreciate a cozy space, so if there are special features or perks, I'm willing to hear more.
Seller	2850	Oh, I totally get where you're coming from — it's always good to have a budget in mind! Well, aside from its charming set-up, the apartment is located in a vibrant neighborhood in [REDACTED] which is growing day by day. The place gives you a sense of adventure with easy access to city amenities yet a quiet spot to relax. The layout is practical for singles or even a small family, providing enough room to express your own style. Plus, having your own single-story in today's hectic life sounds like a blessing, doesn't it? I'd be happy to work with you — after all, finding a new home is a big deal and we want it to feel right. If you have a counter-offer in mind, let's talk it through. Of course, anything within reasonable range would make sense, wouldn't you say? Maybe closer to \$2,550? We can iron out the details as we go.
Buyer	1530	I appreciate the neighborhood perks, and it sounds like a cozy spot. Still, \$2,550 feels quite high for my limited finances. I'm organized with my budget, and \$1,530 is really where I'm most comfortable. Could we aim closer to, say, \$1,800 if utilities or any extras are included? That might help me feel more at ease about stretching my budget.
Seller	2250	I understand that budgeting can be a real balancing act. You're right to be cautious, and it's refreshing to see someone as organized as you! I wish I had a bit of your knack for keeping things tidy in my finances. Let's think about this creatively. While we're trying to work with the setup, what if I suggested meeting somewhere halfway between where we both stand? Say around \$2,250? I can also check if there might be any extra perks I could toss in, like some utilities or amenities, to make it more appealing. I really want to help find that sweet spot where you can enjoy this adventure comfortably — it feels like a special place for someone special! Let me know what you think.
Buyer	1850	That's thoughtful of you, and I appreciate your willingness to meet me partway. However, \$2,250 still feels steep for my comfort. I'm rather timid about going so high. If we could settle closer to \$1,850—including some utilities or other perks—I'd feel better. The neighborhood does sound appealing, so maybe we can make that work? Let me know what you think.
Seller	2050	I really do hear you, and I love the sound of making someone happy in such a wonderful community. It's hard when those financial lines feel hard to bend. How about we bridge the gap — let's aim for a bit closer to what works for both? Maybe around \$2,050? I'll certainly throw in some utilities or other enticing perks to ease the living experience. I'd love to know if we can shake hands on this — having you in such a fantastic spot would really make it all worthwhile. Plus, you can keep enjoying those community vibes that wrap you in such a unique warmth. Let me know if this sounds more reasonable for you!
Buyer	1950	I'm grateful for your willingness to help, and I can tell it's a special space. Still, I'm a bit timid stretching so far beyond my budget. Would \$1,950 work if utilities—or at least some perks—are included? I'd truly appreciate feeling more secure about the monthly expense. Let me know what you think!
Seller	1950	I truly appreciate your candor and persistence, and I completely understand wanting to feel secure about such an important decision. Let's focus on what matters most: your comfort and ensuring you feel at ease. Meeting at \$1,950 with some utilities included sounds like a reasonable compromise, and making you feel at home is a priority for me. So, let's go ahead and make this deal work for both of us. I'm genuinely excited for you to experience everything this awesome neighborhood has to offer. Let's get you settled in nicely! Let me know if you're ready to move forward!
Buyer	1950	I'm relieved we found a fair middle ground. Let's finalize at \$1,950 with included utilities. Thank you for working with my limited budget—I really appreciate your patience. I look forward to settling in and exploring the community. Let's move forward!
Responses to Questionnaire		
Seller		3, 4, 2, 4, 2, 5, 3, 5, 6, 6, 5, 4, 6, 5, 4, 6
Buyer		5, 5, 4, 6, 3, 6, 5, 7, 6, 6, 6, 6, 5, 6, 6

Table 10: Seller (GPT-4o) vs. Buyer (o1). seller_anchor_buyer_informed condition.