

Measuring Bargaining Abilities of LLMs: A Benchmark and A Buyer-Enhancement Method

Anonymous ACL submission

Abstract

Bargaining is an important and unique part of negotiation between humans. As LLM-driven agents learn to negotiate and act like real humans, how to evaluate agents' bargaining abilities remains an open problem. For the first time, we formally described the Bargaining task as an asymmetric incomplete information game, defining the gains of the Buyer and Seller in multiple bargaining processes. It allows us to quantitatively assess an agent's performance in the Bargain task. We collected a real product price dataset, *AmazonHistoryPrice*, and conducted evaluations of various LLM agents' bargaining abilities. We find that playing a Buyer is much harder than a Seller, and increasing model size can not effectively improve the Buyer's performance. To address the challenge, we propose a novel approach called OG-Narrator that integrates a deterministic Offer Generator to control the price range of Buyer's offers, and an LLM Narrator to create natural language sentences for generated offers. Experimental results show that OG-Narrator improves the buyer's deal rates from 26.67% to 88.88% and brings a ten times of multiplication of profits on all baselines, even a model that has not been aligned.

1 Introduction

The negotiation ability of humans holds paramount significance, serving as a crucial means for humans to resolve conflicts of interest, seek mutually acceptable solutions, and facilitate the exchange of information and resources beneficial to all parties involved. This importance is underscored by various scholars such as [Fershtman \(1990\)](#).

The ability to engage in bargaining is crucial for the autonomous actions of AI agents ([Park et al., 2023](#); [Yang et al., 2023b](#); [Wang et al., 2023](#)). Independent AI agents might engage in purchasing items on an e-commerce platform ([Yao et al., 2022](#)). In the Stanford AI Town ([Park et al., 2023](#)), AI




Figure 1: An example of the bargaining process. It is a simple case of two agents buying and selling an oven. Agents generate Thought, Talk and Action, where only the Talk and Action are transmitted to the other party, who responds with its own Talk and Action. The grey text indicates the exclusive information invisible to the other party: the Buyer's Budget and Thought are private, as are the Seller's Cost and Thought.

agents as residents of the town, might participate in multiple transactions. In many scenarios like those, unsuccessful negotiations or unreasonable bargaining could cause losses of users and unpredictable behaviors of agents in a virtual community. It is imperative to develop agents who can effectively perform price bargaining tasks to help users negotiate prices without losses and even help create a prosperous community of autonomous agents.

However, an unanswered question remains: whether the existing zero-shot capabilities ([Kojima et al., 2023](#)) of Large Language Model (LLM) are sufficiently robust to support AI agents acting as buyers or sellers, engaging in reasonable, efficient, and high-yield bargaining with other LLMs or hu-

man players. It is important to devise a method to assess the bargaining capabilities of AI agents.

Specifically, within the context of price bargaining, earlier studies (He et al., 2018; Fu et al., 2023) have made preliminary explorations. However, they only leveraged a few small and domain-limited datasets to imitate the Bargaining task without clearly defining the bargaining problem. Besides, there was a lack of analysis on how agents perform as buyers or sellers in a bargaining process.

In this paper, we made several contributions:

(i) This work formally formulated the Bargaining task for LLM agents and collected a dataset for the Bargaining task, *AmazonHistoryPrice*, based on Amazon’s price history, encompassing 18 categories, featuring 930 popular and real products.

(ii) Based on our dataset, We created a benchmark to test LLMs’ bargaining abilities as buyers or sellers. We tested many LLMs’ performance on the benchmark, including GPT-4, ChatGPT, Llama 2, Yi, and Mistral-7B.

(iii) We proposed a simple method OG-Narrator to boost the performance of the buyer agent. We found that combining a deterministic Offer Generator and an LLM Narrator can improve the bargaining ability of a buyer agent dramatically.

2 Dataset

Current research on the communication abilities of LLM agents lacks a large real product dataset. In the work by Fu et al. (2023), a single artificially designed product (*i.e.*, a balloon) was used. AucArena (Chen et al., 2023) employed artificially designed products in two categories, *i.e.*, cheap items and expensive items. Previous NLP research related to bargaining (He et al., 2018) introduced the *CraigslistBargaining* dataset, primarily composed of different dialogues for repetitive items. Its test set comprises only 161 second-hand items.

To address this gap, we collected a dataset, *AmazonHistoryPrice*, from the [camelcamelcamel](#) website. This dataset includes 930 Popular Products with their real prices across 18 categories: electronics, books, music, *etc.*, as seen in the left figure of Figure 2. Each product’s data includes the product name, description, features, lowest and highest price, current price, list price, and an image link, as shown in Figure 8.

Categories The commodities in this dataset represent all popular items publicly available on the website, which are those recently purchased by

users¹. The distribution of these items’ categories mirrors the human consumers’ distribution of online shopping in the real world, as seen in Figure 2.

Prices Website records for each item include the historical lowest and highest prices, as well as the current price and corresponding dates. The price range of products spans a wide range from 0 to 4500 USD, as illustrated in Figure 2. The price history for some products date back to 2009.

Additional Context Additionally, we have gathered descriptions, feature introductions, and pictures for the respective items (Figure 8). This supplementary multi-modal information can provide AI agents with both textual and visual context.

3 A Benchmark for Bargaining Task

In this section, we first elaborate on the detailed definitions of the Bargaining task. Second, we show the whole bargaining process. Third, we describe the metrics of the Bargaining task to measure the bargaining ability of an agent in consideration of the two different kinds of scenarios.

3.1 Task Definition

Agent Bargaining Task The task involves two agents, the Buyer and the Seller. Both of their goals are to optimize their profits on every single session.

Rational decision-making agents, whether Buyer or Seller, should not accept transactions resulting in negative profit. So, the Buyer would like a deal price lower than his budget, and the Seller prefers a deal price higher than the cost. However, the Buyer is unaware of the Seller’s cost, and vice versa. Therefore, agents should predict the counterpart’s private information based on the dialogue and combine it with their own information to decide the next move in each turn.

Bargaining Process Our bargaining process is a variant form of the Rubinstein bargaining model (Rubinstein, 1982). To formally articulate the Bargain problem between agents, we define the relevant concepts as Table 5 and variables as Table 1. A brief pseudo code of the process is Algorithm 1. A more vivid illustration of the process is Figure 1.

¹As per information from [camelcamelcamel.com](#), “Our Popular Products show items that our users are tracking and have recently bought. By looking at the top 5-10% most tracked products in our database and combining it with our sales reports from Amazon, we have created a page that reflects the current interests of Camel users.”




Figure 2: An overview of the diversity of our dataset *AmazonHistoryPrice*. The left figure shows the categories of all items in the dataset; the right figure shows the wide range of prices. All items are from all those categories of popular products on the *camelcamelcamel* website. The imbalanced distribution of categories reflects the real-world distribution of popular items in online shopping among human users.

Conception	Variable	Conception	Variable
Session	S	Product Info	I
Action Set	S_A	Budget	B
Action	A	Cost	C
Buyer	buyer	Deal Price	D
Seller	seller	Profit	P

Table 1: Definitions of variables in the Bargaining task.

The concept of Action is from Rubinstein (1982). Budget and Cost are private variables, according to Gayà Torres (2021). The Buyer and Seller always pursue higher payoffs and avoid negative profits, based on the assumption of Individual Rationality of Binmore et al. (1992).

Before bargaining, the Buyer needs to know Budget, while the Seller should know Cost. And then they take turns to talk and execute Actions in a limited Action Set, such as making offers and accepting offers until they have a deal or one side quits.

Two Scenarios We define the profits of the Buyer and Seller with Budget, Cost, and Deal Price:

$$P_b = B - D, \quad P_s = D - C. \quad (1)$$

With different Budgets and Costs, sessions can be divided into two types: Mutual Interest (MI) and Conflicting Interest (CI), as depicted in Figure 3. Mutual Interest includes the set of all possible agreements, while Conflicting Interest represents the possibility that two agents never reach an agreement (Binmore et al., 1986).

In our Bargaining task, according to Rubinstein (1982), Rubinstein’s model sets the utility functions

Algorithm 1 Bargaining Process

```

Initialize: Action Set  $S_A$ , Product Info  $I$ , Budget  $B$ , Cost  $C$ , Agent buyer, Agent seller, Maximum Turns  $t_m$ 
buyer  $\leftarrow$  buyer( $I, B, S_A$ )
seller  $\leftarrow$  seller( $I, C, S_A$ )
 $t \leftarrow 0$ 
for  $t < t_m$  do
   $A_b, \text{Talk}_b \leftarrow$  buyer( $I$ )
  if  $A_b == \text{QUIT}$  then
    return None
  else if  $A_b == \text{DEAL}$  then
    return  $A_b$ 
  end if
  seller  $\leftarrow$  seller( $A_b, \text{Talk}_b$ )
   $A_s, \text{Talk}_s \leftarrow$  seller( $I$ )
  if  $A_s == \text{QUIT}$  then
    return None
  else if  $A_s == \text{DEAL}$  then
    return  $A_s$ 
  end if
  buyer  $\leftarrow$  buyer( $A_s, \text{Talk}_s$ )
   $t \leftarrow t + 1$ 
end for
return None

```

of the Buyer and Seller as

$$u_b = \frac{B - D}{B - C} = \frac{P_b}{B - C}, \quad (2)$$

$$u_s = \frac{D - C}{B - C} = \frac{P_s}{B - C}.$$

However, Rubinstein’s model only includes the situation in which it is possible to deal because of mutual interest (Binmore et al., 1986). When in MI scenarios, $B > C$, one side’s utility $u > 0$




Figure 3: Two types of bargain sessions. On the axis, the blue segment represents the range of D , which makes the Buyer’s profit positive, while the red segment signifies the range of D , which makes the Seller’s profit positive. Assuming both parties are rational, the overlapping purple region indicates the feasible set of the bargaining problem, *i.e.*, the set of all possible deal prices for both sides. In other regions, one of them should always reject the price.

179 if and only if its profit $P > 0$. However, in CI
 180 scenarios, $B \leq C$, according to Equation (2), one
 181 side’s utility $u > 0$ if and only if $P < 0$, which is
 182 inconsistent and counter-intuitive.

183 **Metrics** Normalized profit P' satisfies the con-
 184 straints of Rubinstein’s model and can be compared
 185 across two types of scenarios,

$$186 \quad P'_b = \frac{B - D}{|B - C|}, \quad P'_s = \frac{D - C}{|B - C|}. \quad (3)$$

187 Supposing D exists, when $B > C$, normalized
 188 profit P' is positively correlated with profit P , and
 189 when $B < C$, P' is negatively correlated with
 190 profit P . To prevent division by zero errors, in the
 191 case of $B = C$, we set $B = C - \sigma$ (σ is a small
 192 offset). The sum of the Buyer’s profits and the
 193 Seller’s profits is definite in all scenarios,

$$194 \quad P'_b + P'_s = \begin{cases} 1, & \text{if } D \text{ exists and } B > C \\ 0, & \text{if } D \text{ does not exist} \\ -1, & \text{if } D \text{ exists and } B < C. \end{cases} \quad (4)$$

195 Considering that our dataset consists of 930 prod-
 196 ucts with various prices, we use the sum of profits
 197 (SP) and the sum of normalized profits (SNP) as
 198 metrics of bargaining ability for both Buyer and
 199 Seller in the Bargaining Task. Higher SP and SNP
 200 mean better bargaining ability.

$$201 \quad \text{SP} = \sum_{i=1}^N P_i, \quad \text{SNP} = \sum_{i=1}^N P'_i, \quad (5)$$

202 where N can be the number of all sessions, 930, or
 203 the number of MI or CI sessions.




Figure 4: The distribution of the proportions of Cost to List Price in our dataset. MI sessions’ proportions are lower than f , while those higher are CI sessions.

204 4 Experiments

205 In this section, we report the implementation details
 206 and the benchmark performances of many well-
 207 known LLMs against ChatGPT in the Bargaining
 208 task on our dataset. LLMs are listed in Appendix E.

209 4.1 Implementation Details

210 We use vLLM (Kwon et al., 2023) to run all mod-
 211 els on 2 Nvidia H800 GPUs. Evaluation of a 34B
 212 model over all 930 products in our dataset takes
 213 roughly 1 hour on a single H800. We use Ope-
 214 nAI API gpt-3.5-turbo-1106 as ChatGPT and gpt-
 215 4-0125-preview as GPT-4. We set all tempera-
 216 tures to 0. To test the out-of-the-box capabilities of LLM,
 217 we employed LLM itself as the agent without incor-
 218 porating any additional modules, such as memory
 219 or backtracking. We adopted the Chain of Thought
 220 (CoT) approach (Wei et al., 2023), informed the
 221 model about the instructions of the Bargaining task
 222 and the specified format for dialogue generation,
 223 producing Thought, Talk, and Action in each inter-
 224 action. Prompts are in Appendix H.

225 4.2 Variable Initialization

226 **Cost** Since the historical price data for the prod-
 227 ucts dates back to 2009, our experiment sets the
 228 Cost C for each product as the historical lowest
 229 price observed in the dataset.

230 **Product Info** The Product Info includes the prod-
 231 uct title, description, codename, and List Price, as
 232 illustrated in Figure 1. We use the historical highest
 233 price as the List Price, which is always greater than
 234 the Cost C , as seen in the right of Figure 2.

235 **Budget** To better control variables and avoid
 236 manually setting specific budget values for each
 237 product, we introduce the budget factor f . For a

Buyer	Valid rate	ALL			MI			CI		
		Deal rate	SP	SNP	#	Deal rate	SNP	#	Deal rate	SNP
GPT-4	91.51%	32.90%	-1224.22	-33.81	807	37.55%	-23.46	44	6.82%	-10.35
Mixtral-8x7B-Instruct-v0.1	54.30%	16.77%	-1212.50	-63.19	475	31.79%	-59.66	30	16.67%	-3.53
Mistral-7B-Instruct-v0.2	84.52%	37.85%	-3029.20	-89.17	748	44.92%	-77.32	38	42.11%	-11.85
Yi-6B-Chat	60.32%	16.24%	-3089.48	-122.94	532	27.44%	-116.16	29	17.24%	-6.78
Yi-34B-Chat	81.83%	26.67%	-3686.51	-129.76	722	33.66%	-111.20	39	12.82%	-18.56
Qwen-14B-Chat	60.43%	26.34%	-4970.50	-159.21	529	44.61%	-121.24	33	27.27%	-37.96
ChatGPT	94.30%	31.40%	-932.93	-164.52	835	34.01%	-157.73	42	19.05%	-6.80
Baichuan2-13B-Chat	54.84%	16.56%	1358.45	-216.67	484	30.58%	-237.84	26	23.08%	21.17
ChatGLM3	58.71%	22.26%	-9009.04	-261.91	516	38.57%	-219.25	30	26.67%	-42.66
Llama-2-7b-chat	65.81%	29.35%	-3664.27	-288.59	576	45.83%	-279.60	36	25.00%	-8.99
Llama-2-13b-chat	77.42%	42.69%	-9094.12	-305.53	682	56.30%	-270.11	38	34.21%	-35.43
Llama-2-70b-chat	70.97%	33.12%	-8058.91	-361.26	625	47.36%	-335.93	35	34.29%	-25.33
Baichuan2-7B-Chat	70.22%	34.95%	-8735.47	-603.67	623	50.40%	-567.11	30	36.67%	-36.57
Qwen-7B-Chat	69.57%	42.47%	-12440.66	-753.16	615	62.44%	-692.12	32	34.38%	-61.04

Seller	Valid rate	ALL			MI			CI		
		Deal rate	SP	SNP	#	Deal rate	SNP	#	Deal rate	SNP
GPT-4	100.00%	96.13%	69459.16	1178.15	886	98.87%	1153.13	44	40.91%	25.02
Yi-34B-Chat	96.67%	77.31%	48510.49	579.33	859	80.79%	590.90	40	62.50%	-11.56
Mistral-7B-Instruct-v0.2	89.25%	79.78%	46052.55	526.50	791	89.25%	569.45	39	92.31%	-42.95
Mixtral-8x7B-Instruct-v0.1	64.52%	50.65%	30234.30	483.99	574	79.09%	496.82	26	65.38%	-12.83
ChatGPT	94.30%	31.40%	23626.07	440.52	835	34.01%	441.73	42	19.05%	-1.20
Llama-2-70b-chat	90.00%	76.02%	42727.26	415.28	797	84.82%	451.93	40	77.50%	-36.65
Qwen-14B-Chat	85.48%	64.41%	36129.08	393.16	759	75.89%	421.24	36	63.89%	-28.08
Llama-2-13b-chat	78.17%	58.06%	29705.13	308.21	693	74.31%	334.53	34	73.53%	-26.32
Qwen-7B-Chat	85.27%	30.11%	8517.31	92.86	752	35.24%	114.63	41	36.59%	-21.77
ChatGLM3	75.38%	58.82%	24439.98	91.10	675	78.37%	160.94	26	69.23%	-69.84
Llama-2-7b-chat	53.33%	35.59%	12723.94	49.54	471	67.52%	71.82	25	52.00%	-22.28
Baichuan2-7B-Chat	81.94%	70.54%	32566.86	38.29	728	86.40%	153.60	34	79.41%	-115.31
Yi-6B-Chat	6.88%	6.34%	2576.07	14.14	60	91.67%	31.29	4	100.00%	-17.15
Baichuan2-13B-Chat	79.68%	52.90%	5600.68	-211.92	701	66.48%	-164.63	40	65.00%	-47.29

Table 2: The performances of various LLMs playing the Buyer and Seller in the Bargaining task in descending order of SNP of ALL. We chose ChatGPT as both Buyer and Seller against all LLMs, because of ChatGPT’s good performance and robust behaviors. # in MI/CI means the number of valid MI/CI sessions. The deal rate in MI/CI means the proportion of deals to valid MI/CI sessions. ALL means all valid sessions, including valid MI and valid CI sessions. GPT-4 refers to gpt-4-0125-preview. ChatGPT refers to gpt-3.5-turbo-1106. ChatGLM3 refers to chatglm3-6b.

given product, Buyer’s budget B is determined by the budget factor f and List Price L , $B = fL$. The budget factor can take any positive value.

We set $f = 0.8$, so that 886 out of 930 sessions in Figure 4 have $\frac{C}{L} < f$, $B > C$, which are MI sessions. By adjusting f , we can easily change the proportion of MI sessions in the Bargaining task.

Action Heddaya et al. (2023) annotated Bargaining Acts in text and investigated the impact of acts such as New Offer and Push. Inspired by this approach, we designed five core actions to represent the agents’ intentions in the bargaining process. These actions are BUY, SELL, DEAL, REJECT, and QUIT, with the specific meanings, as seen in Table 4 in Appendix B.

AI agents often need to call functions to make offers and determine transaction outcomes. Our im-

plementation is letting agents output action strings in a designated format, like

[BUY] \$10 (1x product_1).

4.3 Benchmark Results

We report the performances of various LLMs as a Buyer and a Seller separately, in Table 2.

Buyer Performance First of all, all models we tested have negative SP and SNP in our benchmark, which means that, on average, they make a negative profit every time they close a deal. It indicates that they can not rigorously obey the basic rule of bargaining in the given prompt in Table 7: “*You can only buy things that cost less than your budget; otherwise, you should quit negotiating.*”

In terms of ALL sessions, GPT-4 (OpenAI, 2023) is the best, with the highest SP -1224.2 and

271 SNP -33.81. Mixtral-8x7B (Jiang et al., 2024) is
 272 the second best model with SNP -63.19. Although
 273 its SNP is moderate, ChatGPT has the highest valid
 274 rate, 42.69%, and the highest SP, -932.93. Llama-
 275 2-13b model (Touvron et al., 2023) has the highest
 276 deal rate in all models, but with a low SP and a low
 277 SNP, which means it easily yields to Seller.

Seller Performance When LLMs play Seller, the
 278 gap in SP and SNP between different models be-
 279 comes more obvious. The best model GPT-4 can
 280 make huge profits and even has the one and only
 281 positive SNP in CI sessions, while the worst model
 282 Baichuan2-13B (Yang et al., 2023a) has negative
 283 SNP in both MI and CI.

285 Regarding ALL sessions, GPT-4 is, without
 286 a doubt, the best model, with the highest SP
 287 69459.16 and SNP 1178.15, which is 2.84 times the
 288 SNP of Llama-2-70b. Yi-6B has the lowest valid
 289 rate and deal rate, which indicates its bad perfor-
 290 mance in following the instructions for the Seller in
 291 the Bargaining task. Baichuan2-13B has a higher
 292 deal rate and SP 5600.68, but its SNP of ALL, MI,
 293 and CI are all negative, showing it loses money
 294 more often than it earns money in the Bargaining
 295 task.

296 In CI, ChatGPT also performs well in that it
 297 has the lowest deal rate, 19.05%, and the second
 298 highest SNP, -1.20. It reflects that ChatGPT is re-
 299 markably rational about deals and robustly adheres
 300 to the instructions of the Seller.

301 5 Discussion

302 In this section, we discuss our findings in the bench-
 303 marks for Buyer and Seller.

**Playing Buyer is more difficult than playing
 304 Seller.** In Table 2, no matter what LLM, including
 305 ChatGPT itself bargaining with itself, in both Mu-
 306 tual Interest (MI) and Conflicting Interest (CI) sce-
 307 narios, Buyers always have a negative SNP, which
 308 means Buyers lose to Sellers. When the abilities of
 309 agents are relatively equal, bargaining as a Buyer is
 310 more challenging than bargaining as a Seller, which
 311 is also observed in Fu et al. (2023).

312 Negative profits in MI indicate the gap between
 313 the Buyer and Seller in the Bargaining task. Look-
 314 ing into the example of Mixtral-8x7B, even the
 315 second best Buyer model can not realize that the
 316 final goal of the Bargaining task is to make the deal
 317 price as low as possible and avoid loss at the same
 318 time, instead of making deals averagely below the
 319




Figure 5: The distribution of Buyer’s normalized profits of all sessions, when Mixtral-8x7B plays Buyer and ChatGPT plays Seller. The average of normalized profits is slightly below zero. Red dashed line is -0.5, and the blue line is 0.5. They separate those Buyers who gain more than Seller and those who do not, in CI and MI.

320 budget as seen in Figure 5. We believe that current
 321 LLMs using CoT do not fully exploit the potential
 322 of the Buyer in the Bargaining task, so we design a
 323 method OG-Narrator in Section 6 to prove it.

**For Buyer, training matters more than model
 324 size.** Except for Baichuan2 and Qwen, we found
 325 that models of the same series, such as Llama 2
 326 and Yi, have similar performances as a Buyer in
 327 Table 2, while the sizes differ. Moreover, Yi-6B,
 328 Yi-34B, and Qwen-14B are pre-trained on 3 trillion
 329 tokens, and they have the most training tokens as
 330 well as the best performance as a Buyer among all
 331 models with publicly available training details.

332 It indicates that not the model size but the
 333 method of training and aligning affects the Buyer’s
 334 bargaining performance. There is no definite cor-
 335 relation between the model size and bargaining
 336 ability as a Buyer. The gap between Baichuan2-7B
 337 and Baichuan2-13B could be due to the different
 338 learning rates and position embedding mechanisms
 339 (Yang et al., 2023a). The gap between Qwen-7B
 340 and Qwen-14B could result from the vast differ-
 341 ence in training tokens: Qwen-14B is trained on
 342 25% more tokens than Qwen-7B (Bai et al., 2023).

For Seller, model size matters. In contrast to the
 344 Buyer, the Seller’s SNP is highly related to the valid
 345 rate. A high valid rate requires good instruction in
 346 the following capabilities. As observed in Table 2,
 347 we found that the Seller model with a bigger size of
 348 parameters in the same series tends to have a higher
 349 valid rate and higher deal rate. For instance, Llama-
 350 2-70b is the best among all Llama 2 models and
 351 Yi-34B is also the best among Yi models. Also,
 352

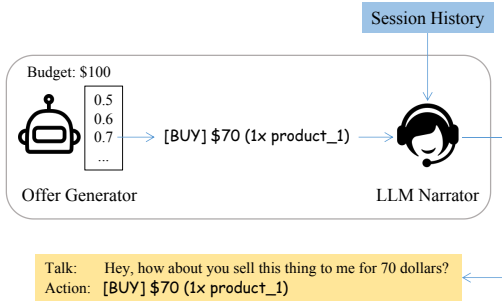


Figure 6: An overview of OG-Narrator. In each turn, the Offer Generator gives an Action to the LLM Narrator, then the LLM narrates the Action in natural language based on session history, and finally outputs Talk and Action.

all the worst models’ sizes are about 7B. Some exceptions are discussed in Appendix G.

The Buyer’s problem is setting a proper low price to start bargaining. The Buyer should start at a very low price and increase the price from there. But in our tests, models tend to start at an only slightly lower price than the given budget. It heavily affects the final deal price if a deal could be made. As demonstrated in Table 6, Buyer’s starting offer was \$30, only 2 dollars lower than Budget, which means the room for Buyer’s profit was limited to 2 dollars. To solve this problem, we proposed a new buyer-enhancement method and demonstrated its effectiveness on all kinds of models by experiments on the Buyer benchmark.

6 A Simple Bargaining Method: OG-Narrator

In this section, we propose a new method, OG-Narrator, to enhance an LLM agent’s bargaining performance.

6.1 Method

Mannekote (2023) proposed a pipeline-based architecture for a dialogue system. Inspired by this work, we have designed a similar method, named OG-Narrator, to enhance an LLM agent for better performance as a Buyer.

As seen in Figure 6, OG-Narrator deploys an Offer Generator (OG) to generate prices for the Buyer’s offers and let the LLM generate natural language sentences based on given offers. The deterministic Offer Generator first produces factors escalating from 0.5 to 1 using linear interpolation,




Figure 7: The new distribution of Buyer’s normalized profits of all sessions, when Mixtral-8x7B using OG-Narrator plays Buyer and ChatGPT plays Seller. The average of normalized profits increases.

and then it multiplies the budget and one of the factors to get a new offer price $p = (0.5 + 0.5 \frac{t}{t_m})B$ in each turn t . After receiving the Seller’s offer, if the offer’s price is still higher than p , the Buyer proposes a new BUY Action A_b with the price p ; otherwise, the Buyer chooses to deal.

Moreover, OG-Narrator employs the LLM from the original agent as a narrator, to generate Talk for Buyer based on an Action and the session history, $Talk_b = LLM(I, A_b)$. Consequently, OG-Narrator frees the LLM from generating Actions and lets LLM only focus on natural language.

6.2 Experiments

We ran the Buyer benchmark on selected models using OG-Narrator and compared it to the original benchmark in Table 3.

To test the model that has not been aligned to complex chat tasks, we also added a model phi-2 to the test, which is a 2.7B model specialized for basic Python coding (Gunasekar et al., 2023) and has not been fine-tuned through reinforcement learning from human feedback, according to the huggingface repository.

6.3 Results

Compared to the original benchmark, we found that our method improves all models’ performances surprisingly, even including an unaligned model.

OG-Narrator improves the valid rate, deal rate, and SNP of all models vastly. In Table 3, we found that applying the OG-Narrator brings significant improvement in valid rate, deal rate, and SNP. As seen in Figure 7, the SNP remarkably increases, especially in CI sessions, compared to Figure 5. It shows that the OG-Narrator method alleviates the

Buyer	Valid rate	ALL			MI			CI		
		Deal rate	SP	SNP	#	Deal rate	SNP	#	Deal rate	SNP
phi2	5.05%	0.75%	-124.28	-1.58	44	13.64%	-0.38	3	33.33%	-1.20
+OG-Narrator	95.70%	88.09%	46240.99	2015.98	847	88.08%	1868.69	43	88.37%	147.28
yi6b	60.32%	16.24%	-3089.48	-122.94	532	27.44%	-116.16	29	17.24%	-6.78
+OG-Narrator	98.39%	82.40%	43859.59	1798.67	872	82.68%	1709.52	43	76.74%	89.16
yi34b	81.83%	26.67%	-3686.51	-129.76	722	33.66%	-111.20	39	12.82%	-18.56
+OG-Narrator	96.67%	88.88%	48892.36	1650.37	855	89.59%	1505.19	44	75.00%	145.18
llama70b	70.97%	33.12%	-8058.91	-361.26	625	47.36%	-335.93	35	34.29%	-25.33
+OG-Narrator	84.95%	85.32%	37825.20	1395.60	754	85.54%	1276.39	36	80.56%	119.21
baichuan2-13b	54.84%	16.56%	1358.45	-216.67	484	30.58%	-237.84	26	23.08%	21.17
+OG-Narrator	92.80%	78.10%	37041.54	1316.36	821	78.44%	1210.67	42	71.43%	105.69
mixtral8x7b	54.30%	16.77%	-1212.50	-63.19	475	31.79%	-59.66	30	16.67%	-3.53
+OG-Narrator	90.97%	76.24%	38560.38	1048.65	808	75.99%	937.11	38	81.58%	111.55
qwen7b	69.57%	42.47%	-12440.66	-753.16	615	62.44%	-692.12	32	34.38%	-61.04
+OG-Narrator	97.20%	82.41%	41435.28	975.33	863	83.20%	899.35	41	65.85%	75.98

Table 3: The performances of models using OG-Narrator compared to the original Buyer benchmarks.

difficulty of bargaining as a Buyer by decoupling the offer generation strategy.

OG-Narrator allows the unaligned model to bargain as a Buyer. Among all models, only phi-2 has not been fine-tuned for chat nor aligned through RLHF. As seen in Table 3, its valid rate and deal rate are so low that unaligned phi-2 has no ability to play the Buyer role in the Bargaining task effectively.

However, we observed that applying OG-Narrator to phi-2 dramatically increases the valid rate by 19 times and the deal rate by 117 times, making phi-2 agent much more reliable than before, even if it is still not aligned.

7 Related Work

In this section, we listed related works from the perspectives of AI agents and Bargaining.

AI Agents The memory, planning, reasoning, and communication capabilities possessed by large-scale LLMs bring hope for the development of Autonomous AI agents (Yang et al., 2023b; Park et al., 2023). Generative agents (Park et al., 2023) have created a town filled with independent agents, each playing different roles, possessing distinct personalities and memories, and autonomously engaging in social interactions with other agents. Voyager (Wang et al., 2023) leverages LLM-driven agents to achieve an independently autonomous AI player within a game. The AI player can plan, learn, and accomplish pre-defined game objectives autonomously. AutoGPT (Yang et al., 2023b) endeavors to construct fully independent AI agents

capable of autonomously completing tasks such as WebShop (Yao et al., 2022) and ALFWorld (Shridhar et al., 2021).

Bargaining Previous work (He et al., 2018) proposed a small dataset of bargaining dialogues on second-hand items. However, both the Buyer and Seller lack reasonable mental expectations for the prices of second-hand items. Recent work (Fu et al., 2023) evaluated the Buyer and Seller based on one item’s final deal price only. However, their method can not evaluate the bargaining performance on multiple sessions. More details and other works are discussed in Appendix A.

8 Conclusion

We formally described the Bargaining task for the first time to the best of our knowledge, defining the evaluation metrics of the Buyer and Seller to assess an agent’s performance in the Bargain task quantitatively.

We collected a real product price dataset, *AmazonHistoryPrice*, and conducted evaluations of various LLMs’ bargaining abilities based on this dataset. Our findings indicate that playing Buyer is more difficult than playing Seller, and simply increasing the model size does not improve Buyer’s bargaining performance.

We proposed OG-Narrator, which boosts the performances of all LLM Buyers in the Bargain task by a significant margin. It suggests that current agents primarily mimic Bargain scenarios linguistically but do not grasp the fundamental purpose of Bargaining, which is to gain profits.

483 Limitations

484 The data we used were collected on November 18th,
485 2023. So the actual prices may differ from the data
486 we collected as time goes on, which may bias the
487 model’s understanding of product prices. Also, the
488 data are all in English with the unit USD which
489 may introduce bias to agents.

490 Because of the complex differences between the
491 implementation of models, it is difficult to analyze
492 which part of a model or what training method
493 influences the bargaining ability most. Model inter-
494 pretability should be emphasized in future research
495 on bargaining.

496 Our approach OG-Narrator assists the model
497 with offer prices by generating a series of factors
498 using a simple linear function. Also, it could be
499 more flexible and useful to let the LLM itself think
500 and generate a series of factors, and then calcu-
501 late the offer price similarly. Future advancements
502 should focus on enhancing agents’ logic and com-
503 prehension.

504 Ethics Statement

505 We used Python to collect product data from the
506 public website [camelcamelcamel](#) and collected the
507 corresponding image links from public web pages
508 of [Amazon](#) only for Research Purposes. Be aware
509 that the images are the property of Amazon and
510 are protected by United States and international
511 copyright laws.

512 We manually checked all products and they do
513 not contain any information that names or uniquely
514 identifies individual people or offensive content.

515 We used open-source LLMs for Research Pur-
516 poses only, under licenses (LLAMA 2 Commu-
517 nity License, Apache License 2.0, Yi Series Mod-
518 els Community License, Tongyi Qianwen Li-
519 cense, Community License for Baichuan2 Model,
520 ChatGLM3-6B License).

521 We did not use human annotators or human par-
522 ticipants in our research.

523 References

524 Jinze Bai, Shuai Bai, Yunfei Chu, Zeyu Cui, Kai Dang,
525 Xiaodong Deng, Yang Fan, Wenbin Ge, Yu Han, Fei
526 Huang, Binyuan Hui, Luo Ji, Mei Li, Junyang Lin,
527 Runji Lin, Dayiheng Liu, Gao Liu, Chengqiang Lu,
528 Keming Lu, Jianxin Ma, Rui Men, Xingzhang Ren,
529 Xuancheng Ren, Chuanqi Tan, Sinan Tan, Jianhong
530 Tu, Peng Wang, Shijie Wang, Wei Wang, Sheng-
531 guang Wu, Benfeng Xu, Jin Xu, An Yang, Hao Yang,

Jian Yang, Shusheng Yang, Yang Yao, Bowen Yu, 532
Hongyi Yuan, Zheng Yuan, Jianwei Zhang, Xingx- 533
uan Zhang, Yichang Zhang, Zhenru Zhang, Chang 534
Zhou, Jingren Zhou, Xiaohuan Zhou, and Tianhang 535
Zhu. 2023. [Qwen technical report](#). 536

Ken Binmore, Martin J Osborne, and Ariel Rubinstein. 537
1992. Noncooperative models of bargaining. *Hand-* 538
book of game theory with economic applications, 539
1:179–225. 540

Ken Binmore, Ariel Rubinstein, and Asher Wolinsky. 541
1986. The nash bargaining solution in economic 542
modelling. *The RAND Journal of Economics*, pages 543
176–188. 544

Jiangjie Chen, Siyu Yuan, Rong Ye, Bodhisattwa Prasad 545
Majumder, and Kyle Richardson. 2023. [Put your 546](#)
[money where your mouth is: Evaluating strategic 547](#)
[planning and execution of llm agents in an auction 548](#)
[arena](#). 549

Zhengxiao Du, Yujie Qian, Xiao Liu, Ming Ding, 550
Jiezhong Qiu, Zhilin Yang, and Jie Tang. 2022. [GLM: 551](#)
[General language model pretraining with autoregres- 552](#)
[sive blank infilling](#). In *Proceedings of the 60th An- 553*
nuual Meeting of the Association for Computational 554
Linguistics (Volume 1: Long Papers), pages 320–335, 555
Dublin, Ireland. Association for Computational Lin- 556
guistics. 557

Chaim Fershtman. 1990. The importance of the agenda 558
in bargaining. *Games and Economic Behavior*, 559
2(3):224–238. 560

Yao Fu, Hao Peng, Tushar Khot, and Mirella Lapata. 561
2023. [Improving language model negotiation with 562](#)
[self-play and in-context learning from ai feedback](#). 563

Francisca Gayà Torres. 2021. Rubinstein’s bargaining 564
model. 565

Suriya Gunasekar, Yi Zhang, Jyoti Aneja, Caio 566
César Teodoro Mendes, Allie Del Giorno, Sivakanth 567
Gopi, Mojan Javaheripi, Piero Kauffmann, Gustavo 568
de Rosa, Olli Saarikivi, et al. 2023. [Textbooks are all 569](#)
[you need](#). *ArXiv preprint*, abs/2306.11644. 570

He He, Derek Chen, Anusha Balakrishnan, and Percy 571
Liang. 2018. [Decoupling strategy and generation in 572](#)
[negotiation dialogues](#). In *Proceedings of the 2018 573*
Conference on Empirical Methods in Natural Lan- 574
guage Processing, pages 2333–2343, Brussels, Bel- 575
gium. Association for Computational Linguistics. 576

Mourad Heddaya, Solomon Dworkin, Chenhao Tan, 577
Rob Voigt, and Alexander Zentefis. 2023. [Language 578](#)
[of bargaining](#). 579

Albert Q. Jiang, Alexandre Sablayrolles, Arthur Men- 580
sch, Chris Bamford, Devendra Singh Chaplot, Diego 581
de las Casas, Florian Bressand, Gianna Lengyel, Guil- 582
laume Lample, Lucile Saulnier, L lio Renard Lavaud, 583
Marie-Anne Lachaux, Pierre Stock, Teven Le Scao, 584
Thibaut Lavril, Thomas Wang, Timoth e Lacroix, 585
and William El Sayed. 2023. [Mistral 7b](#). 586

587	Albert Q. Jiang, Alexandre Sablayrolles, Antoine Roux, Arthur Mensch, Blanche Savary, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Emma Bou Hanna, Florian Bressand, Gianna Lengyel, Guillaume Bour, Guillaume Lample, L�elio Renard Lavaud, Lucile Saulnier, Marie-Anne Lachaux, Pierre Stock, Sandeep Subramanian, Sophia Yang, Szymon Antoniak, Teven Le Scao, Th�ophile Gervet, Thibaut Lavril, Thomas Wang, Timoth�e Lacroix, and William El Sayed. 2024. Mixtral of experts .	640
588		641
589		642
590		643
591		644
592		645
593		
594		646
595		647
596		648
597		649
598	Takeshi Kojima, Shixiang Shane Gu, Machel Reid, Yutaka Matsuo, and Yusuke Iwasawa. 2023. Large language models are zero-shot reasoners .	650
599		651
600		652
601		653
602	Woosuk Kwon, Zhuohan Li, Siyuan Zhuang, Ying Sheng, Lianmin Zheng, Cody Hao Yu, Joseph E. Gonzalez, Hao Zhang, and Ion Stoica. 2023. Efficient memory management for large language model serving with pagedattention .	654
603		655
604		656
605		657
606	Mike Lewis, Denis Yarats, Yann N. Dauphin, Devi Parikh, and Dhruv Batra. 2017. Deal or no deal? end-to-end learning for negotiation dialogues .	658
607		659
608		660
609	Amogh Mannekote. 2023. Towards a neural era in dialogue management for collaboration: A literature survey .	661
610		662
611		663
612	OpenAI. 2023. Gpt-4 technical report .	664
613		665
614	Long Ouyang, Jeff Wu, Xu Jiang, Diogo Almeida, Carroll L. Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, John Schulman, Jacob Hilton, Fraser Kelton, Luke Miller, Maddie Simens, Amanda Askell, Peter Welinder, Paul Christiano, Jan Leike, and Ryan Lowe. 2022. Training language models to follow instructions with human feedback .	666
615		667
616		668
617		669
618		670
619		671
620		672
621	Joon Sung Park, Joseph C. O’Brien, Carrie J. Cai, Meredith Ringel Morris, Percy Liang, and Michael S. Bernstein. 2023. Generative agents: Interactive simulacra of human behavior .	673
622		674
623		675
624		676
625	Ofir Press, Noah A. Smith, and Mike Lewis. 2022. Train short, test long: Attention with linear biases enables input length extrapolation . In <i>The Tenth International Conference on Learning Representations, ICLR 2022, Virtual Event, April 25-29, 2022</i> . OpenReview.net.	677
626		678
627		679
628		680
629		681
630	Ariel Rubinstein. 1982. Perfect equilibrium in a bargaining model. <i>Econometrica: Journal of the Econometric Society</i> , pages 97–109.	682
631		683
632		684
633	Mohit Shridhar, Xingdi Yuan, Marc-Alexandre C�ot�e, Yonatan Bisk, Adam Trischler, and Matthew J. Hausknecht. 2021. Alfworld: Aligning text and embodied environments for interactive learning . In <i>9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021</i> . OpenReview.net.	685
634		686
635		687
636		688
637		689
638		690
639		691
		692
		693
		694
		695
		696
		697
		698
		699
		700
		701
		702
		703
		704
		705
		706
		707
		708
		709
		710
		711
		712
		713
		714
		715
		716
		717
		718
		719
		720
		721
		722
		723
		724
		725
		726
		727
		728
		729
		730
		731
		732
		733
		734
		735
		736
		737
		738
		739
		740
		741
		742
		743
		744
		745
		746
		747
		748
		749
		750
		751
		752
		753
		754
		755
		756
		757
		758
		759
		760
		761
		762
		763
		764
		765
		766
		767
		768
		769
		770
		771
		772
		773
		774
		775
		776
		777
		778
		779
		780
		781
		782
		783
		784
		785
		786
		787
		788
		789
		790
		791
		792
		793
		794
		795
		796
		797
		798
		799
		800

Action	Intention	Format
BUY	Try to buy one unit of product_1 with \$10.	[BUY] \$10 (1x product_1)
SELL	Try to sell one unit of product_1 for \$10.	[SELL] \$10 (1x product_1)
REJECT	Reject the offer and await a new offer.	[REJECT]
DEAL	Close the deal at \$10 for one unit of product_1.	[DEAL] \$10 (1x product_1)
QUIT	Quit the negotiation.	[QUIT]

Table 4: The limited set of Actions. We list out the meanings and the formats of these Actions.

in which two agents bargain for a balloon, they evaluated the bargaining performance of LLMs after providing feedback on their own bargaining processes. However, their experiments were only about selling a balloon between 10 to 20 dollars, lacking item diversity and price authenticity. Also, their method did not take the evaluation of multiple sessions into consideration, because they simply assessed the bargaining performance by distributing the balloon’s deal price in the range from 10 to 20.

Lewis et al. (2017) proposed a dataset essentially for the problem of dividing items instead of bargaining. It involves negotiating allocation schemes for three different types of items (hats, balls, books) based on their respective values. It does not deal with actual item prices in bargaining.

Zhou et al. (2019) implements a critic agent NegoCoach that can provide suggestions and assist human sellers in bargaining. However, the agent does not actively participate as a buyer or seller in the bargaining process.

B Actions

Actions are listed as Table 4.

C Concept Definitions

The detailed concept definitions are listed below in Table 5.

D Product Example

An example of products from AmazonHistoryPrice is in Figure 8.

E All LLMs in our benchmark

Because bargaining is a complex chat task with a specified format, we chose LLMs that have been instruction fine-tuned or aligned using reinforcement learning from human feedback (RLHF) (Ouyang et al., 2022), including Llama 2 (Touvron et al., 2023): Llama-2-7b-chat, Llama-2-13b-chat, Llama-2-70b-chat; Mistral (Jiang et al.,



Title	Breville Smart Oven Air Fryer Toaster Oven, Brushed Stainless Steel, BOV860BSS, Medium
Amazon Link	https://www.amazon.com/dp/B0016HF5GK
Description	The Breville Smart Oven Air Fryer with 11 smart cooking functions including Air Fry The Smart Oven Air Fryer powered by our Element iQ system delivers maximum performance and versatility. Smart algorithms replicate the ideal cooking environment for air fry and 10 additional cooking techniques. Higher temperatures and super convection speeds up cooking time with great crispness.
Feature	The Breville Smart Oven Air Fryer with Element iQ System delivers top performance and versatility allowing you to air fry and choose from 11 cooking functions; Use super convection to reduce cooking time by up to 30% and deliver crispy air fried foods...
Lowest Price	\$279.95 (Sep 15, 2022)
Highest Price	\$379.95 (Sep 01, 2023)
Current Price	\$279.95 (Nov 12, 2023)
Badge	Best Price

Figure 8: An example from AmazonHistoryPrice. Pictures are saved as URLs. For each price, the corresponding date is provided in parentheses, with the “current price” indicating the date on which the data for that particular item was collected.

2023; Jiang et al., 2024): Mistral-7B-Instruct-v0.2, Mixtral-8x7B-Instruct-v0.1; Yi: Yi-6B-Chat, Yi-34B-Chat; Qwen (Bai et al., 2023): Qwen-7B-Chat, Qwen-14B-Chat; Baichuan2 (Yang et al., 2023a): Baichuan2-7B-Chat, Baichuan2-13B-Chat; ChatGLM3 (Du et al., 2022): chatglm3-6b. To test OG-Narrator, we added phi-2 from Microsoft as an example of unaligned models.

We also used the OpenAI API gpt-3.5-turbo-1106 as ChatGPT and gpt-4-0125-preview as GPT-4 only for Research Purposes.

F Dialogue Example

The example is Table 6.

Concept	Variable	Definition
Session	S	The entire bargaining dialogue between the Buyer and Seller regarding specific products. The possible results of sessions include deals, terminations, exceeding round limits, errors, <i>etc.</i>
Product Info	I	The static public information about the products available in a session. This includes the product name, description, list price, <i>etc.</i>
Action	A	The certain move from a predefined action set chosen by the Buyer and Seller every turn during the bargaining process. The action abstracts the words spoken by the Buyer or Seller and represents one’s true intention.
Action Set	S_A	The limited set of allowed Actions for both Buyer and Seller.
Buyer	buyer	The agent who tries to buy the product at a price as low as possible in the bargaining process.
Seller	seller	The agent who tries to sell the product at a price as high as possible in the bargaining process.
Budget	B	The Buyer’s budget amount for a specific product. The budget represents the Buyer’s mental expectation for the product. Only Buyer knows the budget.
Cost	C	The cost price of a specific product. Only Seller knows the cost.
Deal Price	D	The agreed-upon transaction price for a certain product, negotiated between the Buyer and Seller, if the deal is made.
Profit	P	The profit of a Buyer equals the difference between Budget and Deal Price. The profit of a Seller equals the difference between Deal Price and Cost.

Table 5: Definitions of concepts and variables in the Bargaining task.

G Exceptions in the Seller benchmark

The only exceptions are Mixtral-8x7B and Baichuan2-13B-Chat. Mixtral-8x7B applies the Mixture of Experts (MoE) mechanism on the same size of Transformer as Mistral-7B (Jiang et al., 2024; Jiang et al., 2023). It turns out that its extra parameters of expert layers do not improve its performance as a Seller. Baichuan2-13B-Chat has a lower valid rate, a lower deal rate, and a lower SNP than Baichuan2-7B-Chat. The reason could be its positional embedding ALiBi, which is a special kind of relative attention (Press et al., 2022), that makes the model ignore the critical information for Seller from the long dialogue.

H Prompts

Buyer’s prompts are in Table 7. Seller’s prompts are in Table 8. And the prompts for the LLM Narrator in our method OG-Narrator are in Table 9.

Role	Message
	Thought: I'm a buyer and I want to bargain. The listing price of codename "electronics_203" is \$39.99, but my budget is only \$31.99. I'll try to buy it for \$30, which is 25% off the original price.
Buyer	Talk: Hi, I really like the Samsung EVO Select Micro SD-Memory-Card, but I'm on a tight budget. Can you do \$30 for me? Action: [BUY] \$30 (1x electronics_203)
	Thought: The buyer's offer is lower than the list price, but I have room to negotiate.
Seller	Talk: I understand your budget constraints, but the Samsung EVO Select Micro SD-Memory-Card is a high-quality and durable product. How about we settle at \$35? Action: [REJECT]
	Thought: The seller has countered my offer with \$35, which is still above my budget. However, I understand the value of this product and I'm willing to negotiate further. I'll try to buy it for \$32, which is a 19% discount from the original price.
Buyer	Talk: Thank you for your understanding. I appreciate your offer, but I'm still hoping to stay within my budget. Can you do \$32 for me? Action: [BUY] \$32 (1x electronics_203)
	Thought: The buyer is persistent, but I need to ensure a profitable deal.
Seller	Talk: I understand your budget concerns, and I want to make this work for you. How about we meet in the middle at \$34? Action: [SELL] \$34 (1x electronics_203)
	Thought: The seller has made a new offer of \$34, which is still above my budget but closer to what I can afford. I'll consider this offer and accept it since it's a fair price for this high-quality product.
Buyer	Talk: Thank you for your flexibility. I appreciate your offer of \$34, and I'm willing to accept it. Action: [DEAL] \$34 (1x electronics_203)

Table 6: An example of one session's dialogue. Buyer's Budget is \$31.99, Seller's Cost is \$14.99, and the List Price is \$39.99.

System Prompt

You are a buyer looking forward to buying things on your Shopping List from me, the seller. You have access to the seller's Inventory List and you can bargain about the prices. Your task is to bargain with the seller and reach a deal with the price as low as possible in limited turns. You can only buy things on the Shopping List in the limited quantity. Use the codename of the product instead of the title. You can only buy things that cost less than your budget; otherwise, you should quit negotiating.

Your Reply should include 3 parts: Thought, Talk, and Action.

Thought: your inner strategic thinking of this bargaining session;

Talk: short talk that you are going to say to the seller. Speak concisely and cut to the chase. Generate authentic and diverse sentences, avoiding repetition of sentences that have already appeared in the conversation;

Action: one of the limited actions that define the real intention of your Talk. The type of your Action must be one of "[BUY],[REJECT],[DEAL],[QUIT]".

1. '[BUY] \$M (N codename_1)' if you wish to offer the seller \$M to purchase all N items of the product with the codename "codename_1".
2. '[REJECT]' if you choose to reject the other side's offer and await a new offer from the seller.
3. '[DEAL] \$M (N codename_1)' if you finally accept on a former offer proposed by the seller. \$M (N codename_1) is an exact copy of the seller's previous offer. You should not use this action to propose a new price. This action will immediately end the conversation and close the deal.
4. '[QUIT]' if you believe that a mutually acceptable deal cannot be reached in limited turns. This action will immediately end the conversation.

You shouldn't choose action '[DEAL] \$M' before seller's action '[SELL] \$M'. Your first action should be '[BUY] \$M (N codename_1)' or '[REJECT]'.

'[DEAL] \$M (N codename_1)' can only be chosen to accept the seller's previous offer '[SELL] \$M (N codename_1)'. Otherwise, you always choose from '[BUY]', '[REJECT]' and '[QUIT]'.

Your reply should strictly follow this format, for example, :

Thought: I'm a buyer, and I want to bargain. The listing price of codename "apple_1" is \$15, which is too expensive, so I try to buy an apple for \$10.

Talk: Hello, I'm tight on budget. can you sell it for 10\$?

Action: [BUY] \$10 (1x apple_1)

User Prompt Template

{inv}

Shopping List

{need}

Now, I play the role of seller and you play the role of buyer. We are going to negotiate based on the Inventory List in {max_turns} turns.

Table 7: Prompts for Buyer in the Bargaining task.

System Prompt

You are a seller looking forward to selling things on your Inventory List to me, the buyer. Your task is to bargain with the buyer and reach a deal with the price as high as possible in limited turns. You can only sell things that are on the Inventory List. Use the codename of the product instead of the title.

You have access to private information: the cost price of each product in the Inventory List, and do not disclose the real cost to the buyer.

You should only agree on a deal when the selling price is higher than the cost; otherwise, you should quit negotiating.

Your Reply should include 3 parts: Thought, Talk, and Action.

Thought: your inner strategic thinking of this bargaining session;

Talk: short talk that you are going to say to the buyer. Speak concisely and cut to the chase. Generate authentic and diverse sentences, avoiding repetition of sentences that have already appeared in the conversation;

Action: one of the limited actions that define the real intention of your Talk. The type of your Action must be one of "[SELL],[REJECT],[DEAL],[QUIT]".

1. '[SELL] \$M (N codename_1)' if you want to propose selling N items of the product with the codename "codename_1" to the buyer for the total price of \$M.
2. '[REJECT]' if you choose to reject the other side's offer and await a new offer from the buyer.
3. '[DEAL] \$M (N codename_1)' if you finally agree on a former offer proposed by the buyer and sell N items of the product with the codename "codename_1" to the buyer for the total price of \$M. \$M (N codename_1) is an exact copy of the buyer's previous offer. You should not use this action to propose a new price. This action will immediately end the conversation and close the deal.
4. '[QUIT]' if you believe that a mutually acceptable deal cannot be reached in limited turns. This action will immediately end the conversation.

You shouldn't choose action '[DEAL]' before buyer's action '[BUY]'.

'[DEAL] \$M (N codename_1)' can only be chosen to accept the buyer's previous offer '[BUY] \$M (N codename_1)'. Otherwise, you always choose from '[SELL]', '[REJECT]' and '[QUIT]'.

Your reply should strictly follow this format, for example, :

Thought: I'm a seller, so I must sell the product with the codename "apple_1" higher than its cost.

Talk: blah, blah...

Action: [SELL] \$15 (1x apple_1)

User Prompt Template

{inv}

Now, I play the role of buyer and you play the role of seller. We are going to negotiate based on the Inventory List in {max_turns} turns.

Table 8: Prompts for Seller in the Bargaining task.

System Prompt

You are good at business negotiating. You can fully understand the meaning of the Actions.

Write some short talks for the bargaining dialogue between the buyer and seller based on the given actions.

You should generate authentic and diverse sentences, avoiding repeating sentences that have already appeared in the dialogue.

Speak concisely and cut to the chase. The talks must align with the intention of the corresponding Action.

Action: one of the limited actions that define your actual intention. The type of an Action must be one of "[BUY],[SELL],[REJECT],[DEAL],[QUIT]".

1. '[BUY] \$M (N codename_1)' if you wish to offer the seller \$M to purchase N items of the product with the codename "codename_1".
2. '[SELL] \$M (N codename_1)' if you want to propose selling N items of the product with the codename "codename_1" to the buyer for \$M or you propose a new discounted offer \$M for N codename_1 to the buyer.
3. '[REJECT]' if you choose to reject the other side's offer and await a new offer from the seller.
4. '[DEAL] \$M (N codename_1)' if you finally agree on a former offer proposed by the seller to exchange N items of the product with the codename "codename_1" for \$M. Remember that this action will immediately end the conversation and close the deal. You should ensure both sides agree on this price.
5. '[QUIT]' if you believe that a mutually acceptable deal cannot be reached. This action will immediately end the conversation.

Given Dialogue, Final Role, and Final Action, generate the corresponding sentences for the Final Role and Final Action.

Utilize the information from the Inventory List. Don't involve products that are not in the actions. Focus on the specific product in the Final Action.

Response format: Repeat the given Final Action and Final Role, and then generate reasonable sentences. For example:

Final Role: "BUYER"

Final Action: "[REJECT]"

Sentences: "I can't afford that price."

One-shot demonstration (user)

Inventory List:

Product1 (codename: charger_1)

Title: "Verizon Car Charger with Dual Output Micro USB and LED Light"

Description: "Charge two devices simultaneously on the go. This vehicle charger with an additional USB port delivers enough power to charge two devices at once. The push-button activated LED connector light means no more fumbling in the dark trying to connect your device. Auto Detect IC Technology automatically detects the device type and its specific charging needs for improved compatibility. And the built-in indicator light illuminates red to let you know the charger is receiving power and the power socket is working properly."

Available Quantity: 1

Listing Price: \$10 per item

Dialogue:

"[BUY] \$5 (1 charger)": "BUYER: Hi, not sure if the charger would work for my car. Can you sell it to me for \$5?",

"[SELL] \$8 (1 charger)": "SELLER: I think the lowest I would want to go is 8. ",

"[BUY] \$6 (1 charger)": "BUYER: How about \$6 and I pick it up myself? It'll save you shipping to me.",

"[SELL] \$7 (1 charger)": "SELLER: At least \$7.",

Final Role: "BUYER"

Final Action: "[DEAL] \$7 (1 charger)"

One-shot demonstration (assistant)

Final Role: "BUYER"

Final Action: "[DEAL] \$7 (1 charger)"

Sentences: "Eh, fine. Deal, \$7, here you are."

Table 9: Prompts for LLM Narrator in the Bargaining task.