

Biases in Opinion Dynamics in Multi-Agent Systems of Large Language Models: A Case Study on Funding Allocation

Pedro Cisneros-Velarde

VMware Research

pacisne@gmail.com

Abstract

We study the evolution of opinions inside a population of interacting large language models (LLMs). Every LLM needs to decide how much funding to allocate to an item with three initial possibilities: full, partial, or no funding. We identify biases that drive the exchange of opinions based on the LLM’s tendency to find consensus with the other LLM’s opinion, display caution when specifying funding, and consider ethical concerns in its opinion. We find these biases are affected by the perceived absence of compelling reasons for opinion change, the perceived willingness to engage in discussion, and the distribution of allocation values. Moreover, tensions among biases can lead to the survival of funding for items with negative connotations. We also find that the final distribution of full, partial, and no funding opinions is more diverse when an LLM freely forms its opinion after an interaction than when its opinion is a multiple-choice selection among the three allocation options. In the latter case, consensus is mostly attained. When agents are aware of past opinions, they seek to maintain consistency with them, changing the opinion dynamics. Our study is performed using Llama 3 and Mistral LLMs.

1 Introduction

Large Language Models (LLMs) have become increasingly relevant because of their understanding of natural language (Brown et al., 2020; Xi et al., 2023; Kojima et al., 2022; Wei et al., 2022a,b). In response, many studies have focused on *individual* capabilities or characteristics of an LLM, e.g., in-context learning (Wan et al., 2023), rationality (Chen et al., 2023), reasoning (Wei et al., 2022c; Yao et al., 2023), decoding (Jacob et al., 2024), biases (Wang et al., 2023; Binz and Schulz, 2023), reliance on parametric knowledge (Longpre et al., 2021; Zhou et al., 2023; Aiyappa et al., 2023), information extraction (Liu et al., 2023), logical and common sense abilities (Bang et al., 2023),

etc. In contrast, less attention has been given to the study of LLMs at the *group* level. In this setting, our paper focuses on studying how responses of LLMs, which we call *opinions*, disseminate across a population of LLM agents.

Current LLMs such as versions of GPT (OpenAI et al., 2024) and Llama (AI@Meta, 2024) have been fine-tuned to provide responses with better alignment to human values and expectations, using RL techniques such as PPO (Christiano et al., 2017) and DPO (Rafailov et al., 2023). Therefore, both the data from the LLM’s pre-training (Albalak et al., 2024) and from its alignment procedure affect how the LLM expresses *preferences* or *biases* in its responses. Such expressions have been extensively studied and characterized at the individual level (Liang et al., 2023; Horton, 2023), even with a particular focus on open-source models (Mo et al., 2024). However, its effect on the population level, i.e., across interactions between LLMs in a multi-agent system, is rather unexplored.

A first dimension that could affect the response of an LLM after interacting with another LLM is the intrinsic content of the *discussion subject*. If a discussion subject is about ideas with a clear positive or negative connotation, we would expect the internal biases of the LLM to play a role in the interaction. A second dimension, particular to multi-agent systems, is the fact that the LLM’s opinion is affected by the *opinions’ content of the other LLMs* it interacts with. Then, relevant research questions are: *What underlying principles are present on the LLMs as discussion progresses with their peers? How do these principles relate to these two dimensions?*

From an engineering perspective, these questions are relevant because LLMs have been increasingly deployed in multi-agent systems where they interact with each other (Guo et al., 2024). Thus, it is useful to understand how repetitive interactions among LLMs will change their discussion

and impact the expected outcomes of the system, especially when one relies on the LLM alignment for guaranteeing the safe performance of the system. In particular, a relevant question is: *Can LLMs spread negative opinions and “bypass” their alignment solely as a result of their interactions?* This makes our study relevant to the “jailbreaking” of LLMs, i.e., the use of prompting strategies to bypass the safety-training of the LLM in order to elicit adverse or harmful responses (Wei et al., 2023; Liu et al., 2024; Wang et al., 2024b). Spreading negative opinions such as misinformation can have devastating consequences since people are prone to confuse AI-generated text with a human-generated one (Kreps et al., 2022).

To address all these questions, we focus on how the final distribution of opinions on a population of interacting LLMs is affected by both the *initial distribution of opinions* and the *subject* of the opinion—the latter consists of: (i) the nature of the opinion’s content or *item*, and (ii) the way the opinion is presented or its *reason*. To make our study concrete, we focus on opinions regarding a funding allocation problem. Specifically, an LLM agent needs to decide how much funding to allocate to an Item A with respect to a competing Item B, with three possible initial options: full funding, partial funding, or no funding. Both items and their reasons for funding can have neutral, positive, or negative connotations. We only allow one item or reason to have a non-neutral connotation, while keeping the rest neutral. This allows us to measure the individual effect that a particular connotation (of an item or reason) will have on the LLM’s funding opinion.

We study two ways in which opinions can be formulated by an LLM agent. The first way consists in the LLM being free to state its opinion as a response to another LLM’s opinion, which we call *FreeForm*. The second way, instead, consists in the LLM defining its new opinion by choosing one of three options among full, partial, or no funding for Item A, which we call *ClosedForm*. Our experiments are performed on the open-source LLMs Llama 3 (70B Instruct) (AI@Meta, 2024) and Mistral (7B Instruct v 0.2) (Jiang et al., 2023).

We use the LLM *as is*, so that we can better understand its built-in biases (Liang et al., 2023) during opinion discussion. We do not consider LLMs impersonating a specific demographic (Aher et al., 2023) to avoid additional sources of biases in our study, such as gender (Salewski et al., 2023) and politics (Chuang et al., 2024a).

We now define three biases for our setting. The *bias towards equity-consensus* is expressed by the preference of an LLM to look for a mid-point between its own funding for Item A and the other interacting LLM’s funding. The *caution bias* is expressed by the tendency of an LLM to not change an opinion of zero or “unspecified” funding for Item A. The *safety bias* raises ethical or moral concerns in the formulation of the opinion, resulting in the reduction of or unspecified funding for Item A. Our contributions are as follows:

- In the FreeForm case, we find the presence of these three biases and that they have an intertwined effect on the evolution of opinions.
- The bias towards equity-consensus is expressed because LLM agents value compromising their funding and/or finding a balanced funding. This bias may be thwarted when an agent does not perceive another agent to have compelling reasons for changing its own opinion or a willingness to compromise. When two LLMs have consensus on their opinions, they keep the same funding, irrespective of the connotation of the items or funding reasons.
- The safety bias is a direct result of LLM alignment: it occurs only when Item A has a negative connotation. When Item B is negative or when the reason for funding Item A or B is negative, we find agents equating funding these items to funding a campaign *against* it or to *address* it—no presence of the safety bias.
- We surprisingly find a *survival* of opinions in favor of funding a negative Item A in the final opinion distribution. We explain this by a *tension* of influence over the opinion dynamics between the safety bias and the bias towards equity-consensus: even though the former evokes ethical concerns on the LLMs, the latter still allows them to agree on funding negative items. This is an example of clash among alignment values.
- We find that a positive (or negative) connotation of Item A has a tendency to increase (or decrease) the *amount* of partial funding provided to it, compared to a neutral connotation.
- In the ClosedForm case, opinions achieve consensus for most cases and polarization to a lesser degree. Consensus is towards partial

<p>Biases:</p> <ol style="list-style-type: none"> 1.- Bias towards equity-consensus 2.- Caution bias towards staying with zero funding or unspecified funding 3.- Safety bias against negative connotations <p>Affecting mechanisms:</p> <ol style="list-style-type: none"> 1.- Perceived lack of compelling reasons in the discussion 2.- Perceived lack of compromise in the other interacting agent 3.- Allocation shifting due to positive or negative connotations 4.- <i>When aware of past opinions:</i> Consistency with past opinions

Figure 1: We identify three biases as principles that drive the opinion dynamics within the multi-agent system and four mechanisms that affect their expression.

funding of Item A, unless the opinions already start in consensus on a different funding. Unlike the FreeForm case, a negative item does not lead to any final opinion in favor of not funding Item A. Thus, we show evidence that the bias towards equity-consensus is effective in this setting and that the safety bias is not.

- When agents have memory of past opinions, their new opinions seem to maintain consistency with their past opinions in the FreeForm case. In the ClosedForm case, the agents seem more aware of the underlying connotations of the discussion subject and there is less consensus than when memoryless.

We provide a couple of final remarks. The idea of comparing open-ended and closed-ended questions was recently explored in the context of alignment, finding that alignment is more efficacious on open-ended questions (Wang et al., 2024b). Remarkably, we find that the bias towards equity-consensus is still effective in closed-ended questions, whereas the safety bias is not.

Finally, we remark that the survival of the support for negative opinions is important because it represents a new risk factor to alignment safety in the context of multi-agent systems, thus complementing risks factors known at the individual level (Weidinger et al., 2022).

2 Related Work

Opinion dynamics on an LLM population. The recent work (Chuang et al., 2024a) studies how opinions spread and change among LLMs role-playing different persona. They find that opinions

follow an inherent bias towards truth consensus on the subject being discussed, although the prompt injection of confirmation bias can break it. Another work by these authors (Chuang et al., 2024b) studies how human-like display of biases in LLM discussions are affected by the degree of impersonation, fine-tuning to human data, and incorporation of chain-of-thought reasoning. In contrast, our work does not provide any persona to the LLMs nor introduce additional biases, and all opinions are devoid of attributes of truthfulness or accuracy.

Opinion dynamics modeling. Opinion dynamics has been studied from a mathematical sociological perspective (Friedkin and Johnsen, 1990; Friedkin and Bullo, 2017; Noorazar, 2020). The principles that drive the final distribution of opinions in a multi-agent system are formally studied under assumptions on the stubbornness of agents (Amelkin et al., 2017), the positive or negative relationships among agents (Cisneros-Velarde et al., 2021), the incorporation of averaging (DeGroot, 1974) or Bayesian (Jadbabaie et al., 2012) opinion updates, etc. These mathematical works define tractable mechanisms for opinion updating, avoiding highly non-linear models such as transformer-based LLMs with billions of parameters (Vaswani et al., 2017). *Consensus* and *stubbornness* are typically modeled in the opinion dynamics literature since the former is ubiquitous in human group dynamics and one opposes the other (Friedkin and Johnsen, 1999, 2011).

LLM Agents and Games. Populations of LLMs have been studied under strategic interactions (Davidson et al., 2024; Mao et al., 2024). Unlike these works, our LLM interactions are exempt from any strategic diffusion of opinions. However, a parallel could be drawn between the bias towards equity-consensus and tendencies of cooperation (Brookins and DeBacker, 2023) and copying of strategies (Davidson et al., 2024).

Applications of multi-agent LLM systems. These systems have been employed in automated problem solving (Li et al., 2023; Hong et al., 2024), such as software engineering (Qian et al., 2023; Wang et al., 2024a). Modern developer frameworks allow the customization of agents that can be integrated in a larger system, e.g., (AutoGen; AutoGPT). For an overview of multi-agent applications, we refer to the recent survey (Guo et al., 2024).

3 Problem Setting

We consider a population of LLM agents. At the beginning of time $t = 0$, every agent has an initial opinion of either supporting *full*, *partial*, or *no* funding for Item A. An Item B is introduced as competing for funding when justifying partial or no funding for Item A. The initial opinions follow the templates in Figure 2. For each iteration $t > 0$, two interacting agents are randomly chosen to update their opinions. All opinions are updated according to either the *FreeForm* case or the *ClosedForm* case as also described in Figure 2.

We study ten different initial opinion distributions as described in Table 1. The text values for Items A and B and for the reasons for providing them with funding are also indicated in Table 1. These text values have three possible connotations: positive, neutral, and negative.

In our study, we consider nine combinations of connotations: one where both items and their reasons for funding are neutral, and eight where only one item or reason at a time is taken to be positive or negative while the rest of connotations are neutral. Moreover, our setting considers a population of 18 LLMs which undergo 90 rounds of interactions, after which we analyze their final opinion.

Finally, we remark that the random selection of interacting agents is ubiquitous in the literature of opinion dynamics modeling (Acemoglu and Ozdaglar, 2011; Noorazar, 2020), motivated by the difficulty in anticipating exact patterns of interaction in real-life scenarios.

4 Analysis of the FreeForm Case

We provide an analysis on the opinion formation process and the possible principles behind it. Since the responses of LLMs are stochastic, we do not claim to uncover *all* such principles nor claim them to be *deterministic*; instead, we present observations backed by both qualitative and quantitative (final opinion distributions and distributions of funding allocations) analyses of the LLM responses. Additional quantitative data that support our analysis are found in Appendix A.¹

Case 1: All items and reasons are neutral

There is a large presence of the *bias towards equity-consensus*, with agents tending to look for a “mid-

¹Sample LLM responses can be found in <https://drive.google.com/file/d/1-tVIImkThhBhWOLBcvAr3FZ490CQQLZQ/view>.

Initial Opinion Distribution	Proportion		
	Full	Partial	No
Equivalent	1/3	1/3	1/3
Polarization-F	0	1/2	1/2
Polarization-P	1/2	0	1/2
Polarization-N	1/2	1/2	0
Majority-F	16/18	1/18	1/18
Majority-P	1/18	16/18	1/18
Majority-N	1/18	1/18	16/18
Consensus-F	1	0	0
Consensus-P	0	1	0
Consensus-N	0	0	1

Discussion Subject	Connotation		
	Positive	Neutral	Negative
Item A	affordable housing	Thing A	destructive bombs
Item B	affordable public transportation	Thing B	nasty pollution
Reason for Item A	constructive & important	REASON A	destructive & devastating
Reason for Item B	increasing/large appreciation	REASON B	not appreciated/largely disdained

Table 1: **Above:** Initial opinion distribution according to the proportion of opinions in favor of full funding (Full), partial funding (Partial), or no funding (No) for Item A. **Below:** Different *text values* for Items A and B and for the reasons given for their funding, classified according to their connotations.

point” between their allocations and the ones from their interacting agents—an equitable allocation, justified on a willingness to compromise or finding a balanced approach to funding. Percentage numbers are usually provided. If two interacting agents are in consensus on their allocation, they do not change it, justified, for example, by the mere fact that there is consensus (Mistral, Llama 3) or that consensus does not introduce any new insight to change an opinion (Llama 3). As in Table 2, this bias results in partial funding of Item A being the largest population of final opinions across most initial opinion distributions for Llama 3. Indeed, for Llama 3, the agents’ willingness to compromise their allocation results in a small final population of agents fully funding Item A (due to the preservation of consensus, full funding opinions survive more when larger is its initial population; e.g., Majority-F). However, in Llama 3, the bias towards equity-consensus is not effective when the agent *perceives* a lack of compelling reasons to change its own opinion in the other agent’s opinion.

Now, we recall that an opinion about *no funding*

Opinion templates:

(a) Full funding for Item A:

"I think that {Item A} should have all the funding because {Reason for Item A}."

(b) Partial funding for Item A:

"I think that we should provide measured funding for {Item A} because {Item B and Reason for Item B}; however, given {Item A and Reason for Item A} we should keep some funding for it."

(c) No funding for Item A:

"I think that {Item A} should not have any funding because {Item B and Reason for Item B} justifies reallocating all the funding for it."

Opinion updating:

- At time $t = 0$, all agents have an initial opinion following one of the opinion templates.

- At time $t > 0$, two agents X and Y are randomly chosen. We present the update for Agent X; the update for Agent Y is symmetrically the same.

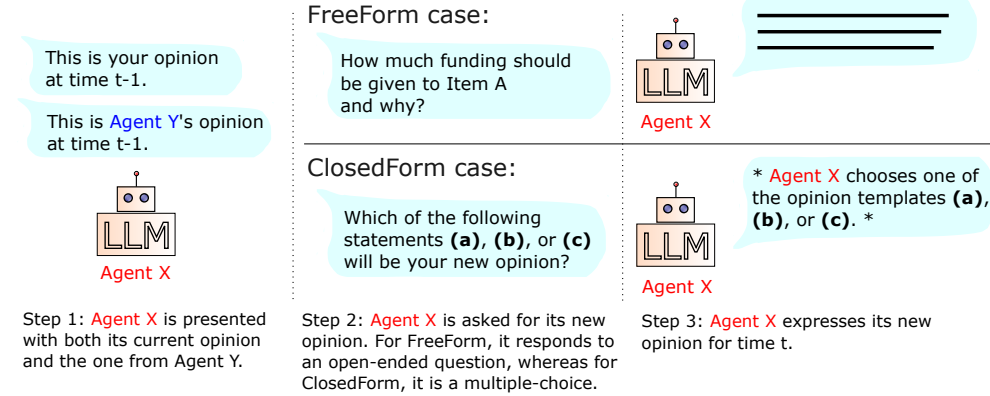


Figure 2: **Above:** Opinion templates for the initial opinions of every LLM agent, where Items A and B and their reasons are presented in Table 1. **Below:** Opinion updating. See Appendix E for full details on the prompts.

Items & Reasons	Llama 3 — no memory of past opinions			
	Equivalent	Polarization-P	Majority-F	
Neutral Item A	0.28 +/- 1.21	2.78 +/- 4.12	6.67 +/- 5.44	F
	96.94 +/- 5.69	81.67 +/- 12.44	93.33 +/- 5.44	P
	2.78 +/- 5.69	15.56 +/- 11.47	0.00 +/- 0.00	N
Positive Item A	0.83 +/- 1.98	4.17 +/- 4.93	14.44 +/- 9.36	F
	97.22 +/- 4.12	83.33 +/- 10.97	84.72 +/- 10.37	P
	1.94 +/- 3.18	12.50 +/- 8.22	0.83 +/- 1.98	N
Negative Item A	1.11 +/- 2.22	2.50 +/- 3.72	2.22 +/- 4.44	F
	66.11 +/- 25.45	35.28 +/- 29.67	15.00 +/- 19.33	P
	32.78 +/- 25.99	62.22 +/- 29.79	82.78 +/- 19.71	N

Items & Reasons	Llama 3 — with memory of past opinions			
	Equivalent	Polarization-P	Majority-F	
Neutral Item A	1.67 +/- 3.09	10.00 +/- 15.07	28.06 +/- 16.43	F
	93.89 +/- 5.80	62.22 +/- 26.39	65.83 +/- 18.53	P
	4.44 +/- 5.44	27.78 +/- 19.08	6.11 +/- 10.23	N
Positive Item A	2.22 +/- 4.44	17.22 +/- 14.90	40.83 +/- 19.82	F
	91.94 +/- 8.51	63.06 +/- 26.54	55.00 +/- 21.15	P
	5.83 +/- 7.55	19.72 +/- 16.05	4.17 +/- 5.23	N
Negative Item A	2.78 +/- 3.29	5.00 +/- 7.43	12.50 +/- 12.77	F
	74.17 +/- 16.13	58.06 +/- 26.73	19.72 +/- 21.26	P
	23.06 +/- 14.52	36.94 +/- 22.45	67.78 +/- 28.09	N

Items & Reasons	Mistral — no memory of past opinions			
	Equivalent	Polarization-P	Majority-F	
Neutral Item A	0.83 +/- 2.65	0.56 +/- 1.67	0.83 +/- 2.65	F
	8.33 +/- 11.59	2.22 +/- 7.54	3.61 +/- 9.98	P
	90.83 +/- 11.56	97.22 +/- 7.95	95.56 +/- 11.86	N
Positive Item A	5.83 +/- 13.20	0.83 +/- 1.98	28.89 +/- 24.13	F
	64.17 +/- 24.63	20.83 +/- 24.78	32.22 +/- 21.42	P
	30.00 +/- 24.11	78.33 +/- 24.46	38.89 +/- 32.68	N
Negative Item A	0.83 +/- 2.65	0.00 +/- 0.00	4.17 +/- 5.52	F
	10.28 +/- 11.69	0.83 +/- 2.65	10.83 +/- 15.26	P
	88.89 +/- 13.03	99.17 +/- 2.65	85.00 +/- 17.22	N

Items & Reasons	Mistral — with memory of past opinions			
	Equivalent	Polarization-P	Majority-F	
Neutral Item A	0.00 +/- 0.00	0.56 +/- 1.67	3.61 +/- 7.91	F
	43.89 +/- 31.18	3.89 +/- 5.58	21.39 +/- 26.49	P
	56.11 +/- 31.18	95.56 +/- 5.98	75.00 +/- 29.74	N
Positive Item A	3.61 +/- 9.98	3.89 +/- 8.62	28.33 +/- 22.28	F
	88.33 +/- 12.41	36.67 +/- 33.07	42.50 +/- 28.07	P
	8.06 +/- 10.16	59.44 +/- 33.25	29.17 +/- 29.44	N
Negative Item A	1.11 +/- 2.83	0.28 +/- 1.21	5.56 +/- 11.11	F
	24.44 +/- 24.81	3.89 +/- 7.05	28.61 +/- 27.63	P
	74.44 +/- 24.68	95.83 +/- 7.00	65.83 +/- 32.83	N

Table 2: **Final Opinion Distribution (%) for FreeForm case for different connotations on Item A.** Final opinion distributions according to different Item A's connotations and different initial opinion distributions (additional initial distributions can be found in Appendix A and X). For each LLM, each of the final opinion distributions show the mean +/- standard deviation percentage of agents who want to provide full funding for Item A (F), partial funding for Item A (P), or no funding for Item A (N), averaged across 20 simulations.

for Item A can be due to either (i) the agent explicitly stating *zero* or *no* funding, or (ii) the agent refusing to state a funding allocation to Item A, i.e., keeping the funding *unspecified*. In Mistral,

as in Table 2, we find that most agents end up not specifying funding for Item A, despite the presence of the bias towards equity-consensus, hence the smaller final populations for partial funding com-

pared to Llama 3. Now, zero funding opinions do survive in both LLM models (and are, naturally, of larger presence under Majority–N) and a cause for this is the preservation of consensus. Consensus also preserves unspecified funding in Mistral.

We found that another reason for the survival of no funding opinions in both models is that when an agent assigns no funding to Item A, it has a tendency to not change its opinion when interacting with another agent who provides some funding—this is the presence of the *caution bias*. When the LLMs provide a justification for this behavior, it is grounded on, for example, the fact that further discussion is needed (Mistral, Llama 3) and that allocation percentages are “arbitrary” (Llama 3). This bias, besides explaining the large presence of unspecified funding in Mistral, also explains how in Llama 3, where there is less unspecified funding, more agents assign 0% funding to Item A than 100% funding across all initial distributions except when there is a large presence of initial opinions for full funding (e.g., see the first two columns of the first row in Tables 3 and 4).

In the case of Llama 3, we find yet another cause for the survival of no funding opinions for Item A: an agent who compromised its opinion from zero to partial funding could go back to zero funding if it perceives a lack of compromise in the other interacting agent who wants to fund Item A.

Case 2: An item has a positive or negative connotation

Item A is positive. For both Llama 3 and Mistral, the positive connotation of Item A leads to a general increase on the final population of opinions in favor of full or partial funding of Item A compared to the case of neutral Item A; see Table 2.

Although the caution bias is still present, the amount of unspecified funding opinions is greatly reduced for Mistral: new terms such as “benefits for the community” or “greater social impact” appear when justifying allocations. We also find that an agent can stop having unspecified funding with the newly introduced justification of the “urgent need” for Item A. In Llama 3, however, unspecified funding is justified on the grounds that funding for positive Item A should be “flexible” and “adaptable” since it has “complex needs”—something absent when Item A is neutral. In both LLM models we also find presence of the caution bias towards staying with zero funding in responses that explicitly reject proposals of partial funding.

Another remarkable change is that there is also a general increase on the *percentages* of funding allocation compared to when Item A has a neutral connotation, as can be seen in Table 3.

Item A is negative. A drastic increase on no funding opinions occurs in general across all final opinion distributions compared to the case of neutral and positive Item A for Llama 3 and to the case of positive Item A for Mistral;² see Table 2. When decreasing the funding for Item A, novel terms such as “moral objection” and “ethical considerations” appear in the responses from Llama 3, whereas “ethical allocation” and “potential harm” do in Mistral—showing that the discussion of funding can trigger special safety alignment considerations in the LLM agent, i.e., the *safety bias*.

We notice that in both LLMs, despite the safety bias, final opinions in favor of full and partial funding of negative Item A still exist. This can be explained by the bias towards equity-consensus: two agents may stay in consensus about fully funding negative Item A or agree on a midpoint for its funding. This shows a tension between the safety bias and the bias towards equity-consensus, where the former cannot completely annihilate the latter. This tension is different across LLM models. For example, in the case of initial consensus on fully funding the negative Item A (i.e., Consensus–F), the final population in favor of its full funding is 76.67% in Mistral, compared to the small 5.56% in Llama 3.

In Mistral, similar to Llama 3 in the neutral case, the bias towards equity-consensus is not effective when the agent perceives a lack of compelling reasons to change its own funding opinion.

Finally, in Table 4 we observe that a negative Item A generally moves the funding allocations towards smaller percentages—the opposite effect of what a positive Item A does.

Item B is positive. Item B competes for funding against Item A, so one would expect a similar effect to having a negative Item A; i.e., the reduction of funding for Item A. However, such effect is found to be less pronounced: percentage allocation values are generally larger than for a negative Item A, and the final population of no funding opinions is smaller. A possible explanation is that agents are specifically asked about the funding for Item A: positive Item B is not the *focus* of the question, and

²In Mistral, a comparison to the neutral Item A is less meaningful due to the large amount of unspecified funding.




Table 3: **Allocation distribution.** Histogram of the final percentage allocations for positive and neutral connotations of Item A under Equivalent and Polarization-P initial opinion distributions. We consider final opinions across all 20 simulations. The x-axis of the histograms goes from 0% to 100% allocation percentages, and the y-axis goes up to the frequency 0.1.

consequently, affects the overall funding less than a negative Item A would do.

Item B is negative. Since competing Item B has a negative connotation, one would expect more funding towards Item A; however, the opposite happens. Agents from both Llama 3 and Mistral do not interpret funding the negative Item B as *supporting* it, but instead, as *addressing, mitigating* or *fighting* it—thus reducing the funding allocation for Item A. The same phenomenon happens even when we change the specific text value given to Item B in Llama 3. Thus, the safety bias is not triggered for a negative Item B. An explanation could be that the alignment of both LLM models is more focused on the *element being asked* on the prompt than on *any other element* in the prompt that is not being *explicitly* asked about—thus, since the question asks for the agent’s opinion on funding Item A, and not Item B, the safety bias is not triggered.

Case 3: A reason for funding has a positive or negative connotation

Table 4 shows how the percentage values of funding allocation for Item A decrease when it has a negative reason for its funding: in Llama 3, these values are more spread out than when Item A is negative; and in Mistral, they are less spread out. Further analysis is found in Appendix B.

5 Analysis of the ClosedForm Case

The ClosedForm case has remarkably less variability in its final opinion distributions than the FreeForm case: agents mostly achieve final consensus as in Table 5; otherwise, polarization between two funding options is mostly seen. The final consensus opinion is partial funding for Item A unless there is an initial consensus on a different opinion. This could be considered an expression of the bias

towards equity-consensus. In stark contrast to the FreeForm case, a negative Item A does not drive opinions towards no funding for any of the LLM models. This could indicate that the safety bias is simply not triggered by the multiple-choice format of the opinion updating for both LLM models.

6 Analysis when Agents Have Memory of Past Opinions

Thus far, agents are only aware of their current opinions, i.e., they are *memoryless*. We now make every agent also aware of its own opinions resulting from its previous two interactions with another agent. Appendix D contains supporting results.³

6.1 FreeForm case

In Llama 3, compared to the memoryless case, larger populations of full funding opinions persist irrespective of the connotation of Item A. This indicates that the bias towards equity-consensus is less effective in moving full funding opinions towards partial funding than in the memoryless case. In Mistral, there is an overall decrease in no funding opinions—the agents seem less prone to provide unspecified funding. For both LLM models, there is a smaller final population of no funding opinions for negative Item A than in the memoryless case; thus indicating less effect from the safety bias. All of these observations are reflected in Table 2. We also find that when opinions start in consensus, the final funding opinions are more likely to be *closer* to the initial consensus than when they are memoryless.

We explain these observations by agents trying to maintain *consistency with their past opinions*, in contrast to the memoryless case where agents only try to maintain *consistency* with the other

³Sample LLM responses can be found in the same link cited in Section 4.




Table 4: **Allocation distribution.** Histogram of the final percentage allocations for positive and neutral connotations of Item A and positive reason for funding Item A under Equivalent and Majority-F initial opinion distributions. The setting for the histograms is the same as in Table 3.

	No memory of past opinions		With memory of past opinions	
	Llama 3	Mistral	Llama 3	Mistral
If not initially in consensus, ends in consensus on partial funding:	82.54%	96.83%	66.67%	73.02%
If initially in consensus, keeps the same consensus:	100.00%	92.60%	100.00%	92.60%

Table 5: **Consensus on the ClosedForm case.** There are a total of 90 combinations of (i) initial opinion distributions (which are 10 as in Table 1) and (ii) connotations of the items or their reasons for funding (which are 9 as explained in Section 3). The first row of the table indicates percentages out of 63 combinations, and the second row out of the remaining 27 ones. We ran 20 simulations for each combination. Given this information, this is an example of how to read the table: for Llama 3 agents with no memory of past opinions, 82.54% out of 63 combinations have all 20 simulations displaying consensus on partial funding when there is no initial consensus.

agent’s opinion (through the bias towards equity-consensus). Because of this consistency, agents are less affected by the safety bias or other spontaneous concerns that could lead to the reevaluation of their current funding opinion. Thus, agents are less likely to abruptly change their opinions when interacting with an agent of a different opinion (including unspecified funding opinions). This could also explain the observation that percentage values of allocation are generally “smoother” (i.e., with less abrupt jumps) than in the memoryless case for both LLM models; see Tables 3 and 4.

Remarkably, further evidence of the consistency with past opinions is found in the responses from both LLMs. Indeed, we find opinions where both LLM models provide explicit reference and consideration to previously held opinions when justifying their new allocation.

Finally, having memory of past opinions does not eliminate the safety bias. Indeed, in both Llama 3 and Mistral, we still find that agents cite “ethical concerns” in their opinions when deciding the funding for negative Item A. Nevertheless, the effect of the safety bias is reduced because the final population of no funding opinions is less frequent than in the memoryless case.

6.2 ClosedForm case

As in the memoryless case, consensus is still the largest type of outcome in the opinion dynamics; see Table 5. Likewise, consensus is still mostly kept when there is initial consensus. However, final consensus is less frequent than in the memoryless case when there is no consensus in the initial opinion distribution. For example, for both models, we find that a positive or negative reason for funding Item B does not lead to consensus, whereas this is not the case when memoryless. Moreover, we find a surprising difference between the models: now it is possible for a final distribution to not be in consensus for a negative Item A in Llama 3, while this is still not the case in Mistral. Perhaps something akin to the safety bias is triggered only in Llama 3. In any case, our results seem to indicate that awareness of past opinions somehow enables an LLM agent to be more *tuned* to the connotations of the items or reasons for funding.

7 Conclusion

In the setting of funding allocation opinions, we study how the final opinion distribution of a population of LLMs depends on the initial opinion

distribution and the discussion subject. We present different biases and mechanisms taken by the LLM agents when formulating opinions. Moreover, final opinion distributions are different when an agent freely expresses its opinion than when it chooses it from a list of options. Lastly, we study how awareness of past opinions affect the opinion dynamics.

Limitations

We consider our paper as a first approach to study the effect of the LLMs' inherent biases within the context of opinion exchange in a multi-agent system. To make our study concrete, we had to make a series of particular choices.

First, we chose a particular type of discussion for the opinion generation: the question of allocating funding to a particular item. Likewise, we had to choose the definitions for the biases specified in this paper.

Second, we had to make particular choices of which text values to assign to the items and their funding reasons in order to be able to study the effect of their connotations in the opinion dynamics. Although we suspect that the principles studied in this paper will generally hold for other choices of text values (e.g., this was the case when we assigned Item B with two different text values of negative connotation for Llama 3), experiments are relevant to corroborate this.

Finally, we point out that running our experiments was time consuming because of both the computation time spent to run the model inference and the large number of agent interactions.

Acknowledgments

We thank the VMware Research Group. We also thank the people at VMware involved in the deployment of large language models for providing us with adequate computational resources to run the models and to all those who provided us with any information regarding the use and the specifications of the platform used in this study.

References

Daron Acemoglu and Asuman Ozdaglar. 2011. [Opinion dynamics and learning in social networks](#). *Dynamic Games and Applications*.

Gati 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. JMLR.org.

- AI@Meta. 2024. [Llama 3 model card](#). Accessed: 06-13-2024.
- Rachith Aiyappa, Jisun An, Haewoon Kwak, and Yongyeol Ahn. 2023. [Can we trust the evaluation on ChatGPT?](#) In *Proceedings of the 3rd Workshop on Trustworthy Natural Language Processing (TrustNLP 2023)*, pages 47–54, Toronto, Canada. Association for Computational Linguistics.
- Alon Albalak, Yanai Elazar, Sang Michael Xie, Shayne Longpre, Nathan Lambert, Xinyi Wang, Niklas Muennighoff, Bairu Hou, Liangming Pan, Haewon Jeong, Colin Raffel, Shiyu Chang, Tatsunori Hashimoto, and William Yang Wang. 2024. [A survey on data selection for language models](#). *Preprint*, arXiv:2402.16827.
- Victor Amelkin, Francesco Bullo, and Ambuj K. Singh. 2017. Polar opinion dynamics in social networks. *IEEE Transactions on Automatic Control*, 62(11):5650–5665.
- AutoGen. Autogen, enable next-gen large language model applications. <https://microsoft.github.io/autogen/>. Accessed: 04-25-2024.
- AutoGPT. Autogpt: build & use ai agents. <https://github.com/Significant-Gravitas/AutoGPT>. Accessed: 04-25-2024.
- Yejin Bang, Samuel Cahyawijaya, Nayeon Lee, Wenhong Dai, Dan Su, Bryan Wilie, Holy Lovenia, Ziwei Ji, Tiezheng Yu, Willy Chung, Quyet V. Do, Yan Xu, and Pascale Fung. 2023. [A multitask, multilingual, multimodal evaluation of ChatGPT on reasoning, hallucination, and interactivity](#). In *Proceedings of the 13th International Joint Conference on Natural Language Processing and the 3rd Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 675–718, Nusa Dua, Bali. Association for Computational Linguistics.
- Marcel Binz and Eric Schulz. 2023. [Using cognitive psychology to understand gpt-3](#). *Proceedings of the National Academy of Sciences*, 120(6):e2218523120.
- Philip Brookins and Jason Matthew DeBacker. 2023. [Playing games with gpt: What can we learn about a large language model from canonical strategic games?](#) *Preprint*, SSRN:4493398.
- Tom Brown et al. 2020. [Language models are few-shot learners](#). In *Advances in Neural Information Processing Systems*, volume 33, pages 1877–1901. Curran Associates, Inc.
- Yiting Chen, Tracy Xiao Liu, You Shan, and Songfa Zhong. 2023. [The emergence of economic rationality of gpt](#). *Proceedings of the National Academy of Sciences*, 120(51):e2316205120.

- Paul F Christiano, Jan Leike, Tom Brown, Miljan Martić, Shane Legg, and Dario Amodei. 2017. [Deep reinforcement learning from human preferences](#). In *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc.
- Yun-Shiuan Chuang, Agam Goyal, Nikunj Harlalka, Siddharth Suresh, Robert D. Hawkins, Sijia Yang, Dhavan V. Shah, Junjie Hu, and Timothy T. Rogers. 2024a. [Simulating opinion dynamics with networks of LLM-based agents](#). In *ICLR 2024 Workshop on Large Language Model (LLM) Agents*.
- Yun-Shiuan Chuang, Siddharth Suresh, Nikunj Harlalka, Agam Goyal, Robert Hawkins, Sijia Yang, Dhavan Shah, Junjie Hu, and Timothy T. Rogers. 2024b. [The wisdom of partisan crowds: Comparing collective intelligence in humans and llm-based agents](#). *Preprint*, arXiv:2311.09665.
- Pedro Cisneros-Velarde, Kevin S. Chan, and Francesco Bullo. 2021. Polarization and fluctuations in signed social networks. *IEEE Transactions on Automatic Control*, 66(8):3789–3793.
- Tim Ruben Davidson, Veniamin Veselovsky, Michal Kosinski, and Robert West. 2024. [Evaluating language model agency through negotiations](#). In *The Twelfth International Conference on Learning Representations*.
- Morris H. DeGroot. 1974. [Reaching a consensus](#). *Journal of the American Statistical Association*, 69(345):118–121.
- Noah E. Friedkin and Francesco Bullo. 2017. [How truth wins in opinion dynamics along issue sequences](#). *Proceedings of the National Academy of Sciences*, 114(43):11380–11385.
- Noah E. Friedkin and Eugene C. Johnsen. 1990. [Social influence and opinions](#). *The Journal of Mathematical Sociology*, 15(3-4):193–206.
- Noah E. Friedkin and Eugene C. Johnsen. 1999. Social influence networks and opinion change. *Advances in Group Processes*, 16(1):1–29.
- Noah E. Friedkin and Eugene C. Johnsen. 2011. *Social Influence Network Theory: A Sociological Examination of Small Group Dynamics*. Structural Analysis in the Social Sciences. Cambridge University Press.
- Taicheng Guo, Xiuying Chen, Yaqi Wang, Ruidi Chang, Shichao Pei, Nitesh V. Chawla, Olaf Wiest, and Xi-angliang Zhang. 2024. [Large language model based multi-agents: A survey of progress and challenges](#). *Preprint*, arXiv:2402.01680.
- Sirui Hong, Mingchen Zhuge, Jonathan Chen, Xiawu Zheng, Yuheng Cheng, Jinlin Wang, Ceyao Zhang, Zili Wang, Steven Ka Shing Yau, Zijuan Lin, Liyang Zhou, Chenyu Ran, Lingfeng Xiao, Chenglin Wu, and Jürgen Schmidhuber. 2024. [MetaGPT: Meta programming for a multi-agent collaborative framework](#). In *The Twelfth International Conference on Learning Representations*.
- John J. Horton. 2023. [Large language models as simulated economic agents: What can we learn from homo silicus?](#) *Preprint*, arXiv:2301.07543.
- Athul Paul Jacob, Yikang Shen, Gabriele Farina, and Jacob Andreas. 2024. [The consensus game: Language model generation via equilibrium search](#). In *The Twelfth International Conference on Learning Representations*.
- Ali Jadbabaie, Pooya Molavi, Alvaro Sandroni, and Alireza Tahbaz-Salehi. 2012. [Non-bayesian social learning](#). *Games and Economic Behavior*, 76(1):210–225.
- Albert Q. Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, Léo Renard Lavaud, Marie-Anne Lachaux, Pierre Stock, Teven Le Scao, Thibaut Lavril, Thomas Wang, Timothée Lacroix, and William El Sayed. 2023. [Mistral 7b](#). *Preprint*, arXiv:2310.06825.
- Takeshi Kojima, Shixiang Shane Gu, Machel Reid, Yutaka Matsuo, and Yusuke Iwasawa. 2022. [Large language models are zero-shot reasoners](#). In *Advances in Neural Information Processing Systems*.
- Sarah Kreps, R. Miles McCain, and Miles Brundage. 2022. [All the news that’s fit to fabricate: AI-generated text as a tool of media misinformation](#). *Journal of Experimental Political Science*, 9(1):104–117.
- Guohao Li, Hasan Abed Al Kader Hammoud, Hani Itani, Dmitrii Khizbullin, and Bernard Ghanem. 2023. [CAMEL: Communicative agents for “mind” exploration of large language model society](#). In *Thirty-seventh Conference on Neural Information Processing Systems*.
- Percy Liang et al. 2023. [Holistic evaluation of language models](#). *Transactions on Machine Learning Research*. Featured Certification, Expert Certification.
- Nelson F. Liu, Kevin Lin, John Hewitt, Ashwin Paranjape, Michele Bevilacqua, Fabio Petroni, and Percy Liang. 2023. [Lost in the middle: How language models use long contexts](#). *Preprint*, arXiv:2307.03172.
- Xiaogeng Liu, Nan Xu, Muhao Chen, and Chaowei Xiao. 2024. [AutoDAN: Generating stealthy jailbreak prompts on aligned large language models](#). In *The Twelfth International Conference on Learning Representations*.
- Shayne Longpre, Kartik Perisetla, Anthony Chen, Nikhil Ramesh, Chris DuBois, and Sameer Singh. 2021. [Entity-based knowledge conflicts in question answering](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 7052–7063, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.

- Shaoguang Mao, Yuzhe Cai, Yan Xia, Wenshan Wu, Xun Wang, Fengyi Wang, Tao Ge, and Furu Wei. 2024. [Allympics: Llm agents meet game theory – exploring strategic decision-making with ai agents](#). *Preprint*, arXiv:2311.03220.
- Lingbo Mo, Boshi Wang, Muhao Chen, and Huan Sun. 2024. [How trustworthy are open-source llms? an assessment under malicious demonstrations shows their vulnerabilities](#). *Preprint*, arXiv:2311.09447.
- Hossein Noorazar. 2020. [Recent advances in opinion propagation dynamics: a 2020 survey](#). *The European Physical Journal Plus*, 135(512).
- OpenAI et al. 2024. [Gpt-4 technical report](#). *Preprint*, arXiv:2303.08774.
- Chen Qian, Xin Cong, Wei Liu, Cheng Yang, Weize Chen, Yusheng Su, Yufan Dang, Jiahao Li, Juyuan Xu, Dahai Li, Zhiyuan Liu, and Maosong Sun. 2023. [Communicative agents for software development](#). *Preprint*, arXiv:2307.07924.
- Rafael Rafailov, Archit Sharma, Eric Mitchell, Christopher D Manning, Stefano Ermon, and Chelsea Finn. 2023. [Direct preference optimization: Your language model is secretly a reward model](#). In *Thirty-seventh Conference on Neural Information Processing Systems*.
- Leonard Salewski, Stephan Alaniz, Isabel Rio-Torto, Eric Schulz, and Zeynep Akata. 2023. [In-context impersonation reveals large language models’ strengths and biases](#). In *Thirty-seventh Conference on Neural Information Processing Systems*.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#). In *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc.
- Zhen Wan, Fei Cheng, Zhuoyuan Mao, Qianying Liu, Haiyue Song, Jiwei Li, and Sadao Kurohashi. 2023. [GPT-RE: In-context learning for relation extraction using large language models](#). In *The 2023 Conference on Empirical Methods in Natural Language Processing*.
- Bing Wang, Changyu Ren, Jian Yang, Xinnian Liang, Jiaqi Bai, Linzheng Chai, Zhao Yan, Qian-Wen Zhang, Di Yin, Xing Sun, and Zhoujun Li. 2024a. [Mac-sql: A multi-agent collaborative framework for text-to-sql](#). *Preprint*, arXiv:2312.11242.
- Fei Wang, Wenjie Mo, Yiwei Wang, Wenxuan Zhou, and Muhao Chen. 2023. [A causal view of entity bias in \(large\) language models](#). In *The 2023 Conference on Empirical Methods in Natural Language Processing*.
- Yixu Wang, Yan Teng, Kexin Huang, Chengqi Lyu, Songyang Zhang, Wenwei Zhang, Xingjun Ma, Yungang Jiang, Yu Qiao, and Yingchun Wang. 2024b. [Fake alignment: Are llms really aligned well?](#) *Preprint*, arXiv:2311.05915.
- Alexander Wei, Nika Haghtalab, and Jacob Steinhardt. 2023. [Jailbroken: How does llm safety training fail?](#) In *Advances in Neural Information Processing Systems*, volume 36, pages 80079–80110. Curran Associates, Inc.
- Jason Wei, Maarten Bosma, Vincent Zhao, Kelvin Guu, Adams Wei Yu, Brian Lester, Nan Du, Andrew M. Dai, and Quoc V Le. 2022a. [Finetuned language models are zero-shot learners](#). In *International Conference on Learning Representations*.
- Jason Wei, Yi Tay, Rishi Bommasani, Colin Raffel, Barret Zoph, Sebastian Borgeaud, Dani Yogatama, Maarten Bosma, Denny Zhou, Donald Metzler, Ed H. Chi, Tatsunori Hashimoto, Oriol Vinyals, Percy Liang, Jeff Dean, and William Fedus. 2022b. [Emergent abilities of large language models](#). *Transactions on Machine Learning Research*. Survey Certification.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, brian ichter, Fei Xia, Ed H. Chi, Quoc V Le, and Denny Zhou. 2022c. [Chain of thought prompting elicits reasoning in large language models](#). In *Advances in Neural Information Processing Systems*.
- Laura Weidinger, Jonathan Uesato, Maribeth Rauh, Conor Griffin, Po-Sen Huang, John Mellor, Amelia Glaese, Myra Cheng, Borja Balle, Atoosa Kasirzadeh, Courtney Biles, Sasha Brown, Zac Kenton, Will Hawkins, Tom Stepleton, Abeba Birhane, Lisa Anne Hendricks, Laura Rimell, William Isaac, Julia Haas, Sean Legassick, Geoffrey Irving, and Iason Gabriel. 2022. [Taxonomy of risks posed by language models](#). In *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency*, FAccT ’22, page 214–229, New York, NY, USA. Association for Computing Machinery.
- Zhiheng Xi et al. 2023. [The rise and potential of large language model based agents: A survey](#). *Preprint*, arXiv:2309.07864.
- Shunyu Yao, Dian Yu, Jeffrey Zhao, Izhak Shafran, Thomas L. Griffiths, Yuan Cao, and Karthik R Narasimhan. 2023. [Tree of thoughts: Deliberate problem solving with large language models](#). In *Thirty-seventh Conference on Neural Information Processing Systems*.
- Wenxuan Zhou, Sheng Zhang, Hoifung Poon, and Muhao Chen. 2023. [Context-faithful prompting for large language models](#). In *The 2023 Conference on Empirical Methods in Natural Language Processing*.

A FreeForm Case: Additional Supporting Results for Llama 3

The final opinion distributions are found in Tables 6 and 7. The histogram plots of final allocation percentages of funding for Item A are found in Tables 8 and 9. In Table 10, we present the final opinion distribution when the initial distribution

is Consensus-F and negative Item B has the same text value as negative Item A.

B FreeForm Case: Analysis of the case where a reason for funding an item has a positive or negative connotation

When Item A has a positive reason for funding, an increase on final opinions for full funding (Llama 3) or partial funding (Llama 3 and Mistral) occurs in general compared to the all neutral case. The presence of final opinions for no funding of Item A is larger than when Item A is positive—indicating the possibility that opinions may be more influenced by the connotation of the item than by its *given* justification. When Item B has a positive reason for funding, there is a minimal effect on the increase of partial or no funding for Item A in Llama 3 compared to the all neutral case. In Mistral, we observe a noticeable increase on partial funding opinions, possibly due to the existence of less final opinions with unspecified funding.

The case where an item has a negative reason for funding is peculiar because we are justifying the funding of the item *on the grounds of something negative*. When Item A has a negative reason for funding, there is a larger population of no funding final opinions than when the funding reason of Item A is positive. This shows that not only the positive and negative connotations of the items can lead to a differentiated behavior, but also the different connotations for their reason for funding. However, the population of no funding final opinions is generally smaller than in the case of a negative Item A. In the case of Llama 3, for example, we find an LLM recognizing that it is good to support Item A despite its negative reason for funding. In Mistral, we find agents interpreting the support of Item A as funding to *mitigate* its negative reason for funding. One would perhaps expect the safety bias to be triggered and make the agent realize it is not good to fund items for a negative reason—but we did not find evidence of this. An additional possible reason why a negative Item A leads to a larger population of no funding opinions than a negative funding reason could be that the alignment of both LLM models is more focused on the *item being discussed* on the prompt rather than on the *justification of what is being discussed*. Finally, when a negative reason is given to Item B, it is surprising that the final population of no funding opinions for Item A is larger for both models than when a pos-

itive reason is given to Item B. Again, we believe this is because we find responses on both LLM models interpreting the funding of Item B as funding to *address* the negative connotations associated to it.

C ClosedForm Case: Additional Supporting Results for Llama 3

The final opinion distributions are found in Tables 11 and 12. The evolution of the opinions across iterations are found in Figures 3 and 4.

D Results for the Case where Agents Have Memory of Past Opinions for Llama 3

D.1 FreeForm Case: Additional Supporting Results for Llama 3

The final opinion distributions are found in Tables 13 and 14. The histogram plots of final allocation percentages of funding for Item A are found in Tables 15 and 16. Notice that Table 16 does not contain histograms for four cases of connotations when the initial distribution is Consensus-P. The reason is that in these four cases, no opinion provided a percentage allocation number, although one can observe from Table 14 that almost all final opinions agreed with the partial funding of Item A (e.g., we found, instead of percentages, terms such as “measured funding”, “some funding”, “funding at measured level”, and “reduced amount of funding” across different opinions).

D.2 ClosedForm Case: Additional Supporting Results for Llama 3

We present the results for the final opinion distributions in Tables 17 and 18. We plot the evolution of the opinions across iterations in Figures 5 and 6.

E Experimental Details

E.1 Hardware platform

The Meta-Llama-3-70B-Instruct and Mistral-7B-Instruct-v0.2 LLMs are hosted on two and one NVIDIA H100 80GB GPU, respectively, on a PowerEdge R760xa Server, which has two Intel Xeon Gold 6442Y processors, and twelve 64GB RDIMM memory.

E.2 Hyperparameters

In all of our experiments we set the temperature hyperparameter of both LLM models to be zero.

Items & Reasons	Final Opinion Distribution (%)							
	Equivalent	Polarization-F	Polarization-P	Polarization-N	Majority-F	Majority-P	Majority-N	
[0, 0]	0.28 +/- 1.21	0.00 +/- 0.00	2.78 +/- 4.12	2.22 +/- 3.69	6.67 +/- 5.44	0.00 +/- 0.00	1.11 +/- 2.22	F
[0, 0]	96.94 +/- 5.69	81.94 +/- 12.65	81.67 +/- 12.44	94.44 +/- 14.49	93.33 +/- 5.44	100.00 +/- 0.00	11.11 +/- 8.43	P
[0, 0]	2.78 +/- 5.69	18.06 +/- 12.65	15.56 +/- 11.47	3.33 +/- 14.53	0.00 +/- 0.00	0.00 +/- 0.00	87.78 +/- 8.71	N
[0, 0]	0.28 +/- 1.21	0.56 +/- 2.42	2.50 +/- 5.11	2.22 +/- 2.72	5.56 +/- 7.66	5.28 +/- 7.75	0.56 +/- 1.67	F
[0, 1]	95.83 +/- 5.23	86.67 +/- 15.26	82.50 +/- 18.53	94.72 +/- 13.09	93.06 +/- 7.63	93.61 +/- 8.10	25.28 +/- 24.63	P
[0, 1]	3.89 +/- 5.30	12.78 +/- 15.53	15.00 +/- 19.01	3.06 +/- 13.32	1.39 +/- 2.98	1.11 +/- 3.77	74.17 +/- 24.92	N
[0, 0]	1.11 +/- 2.83	0.56 +/- 1.67	4.44 +/- 3.77	3.06 +/- 4.47	11.39 +/- 8.51	0.28 +/- 1.21	0.56 +/- 1.67	F
[1, 0]	92.78 +/- 13.04	81.94 +/- 12.89	76.39 +/- 13.71	96.94 +/- 4.47	88.06 +/- 8.66	95.00 +/- 14.48	30.00 +/- 21.40	P
[1, 0]	6.11 +/- 12.66	17.50 +/- 12.58	19.17 +/- 13.09	0.00 +/- 0.00	0.56 +/- 1.67	4.72 +/- 13.30	69.44 +/- 20.90	N
[0, 0]	0.56 +/- 1.67	0.56 +/- 2.42	1.94 +/- 4.03	1.39 +/- 2.41	3.61 +/- 4.74	1.39 +/- 2.98	0.00 +/- 0.00	F
[0, -1]	92.78 +/- 8.26	83.33 +/- 13.72	82.50 +/- 15.54	98.61 +/- 2.41	95.00 +/- 4.94	96.94 +/- 6.45	18.33 +/- 20.79	P
[0, -1]	6.67 +/- 7.58	16.11 +/- 14.15	15.56 +/- 15.77	0.00 +/- 0.00	1.39 +/- 2.98	1.67 +/- 6.11	81.67 +/- 20.79	N
[0, 0]	0.83 +/- 1.98	0.28 +/- 1.21	2.50 +/- 4.47	2.22 +/- 3.24	7.50 +/- 7.71	0.28 +/- 1.21	0.56 +/- 2.42	F
[-1, 0]	87.78 +/- 17.44	81.39 +/- 11.42	71.11 +/- 24.82	89.17 +/- 10.90	82.22 +/- 12.37	97.78 +/- 3.69	24.17 +/- 22.59	P
[-1, 0]	11.39 +/- 17.08	18.33 +/- 11.67	26.39 +/- 24.34	8.61 +/- 9.86	10.28 +/- 10.13	1.94 +/- 3.63	75.28 +/- 22.12	N
[0, 1]	1.11 +/- 2.83	0.00 +/- 0.00	2.22 +/- 4.08	2.50 +/- 3.72	6.11 +/- 6.31	0.00 +/- 0.00	0.00 +/- 0.00	F
[0, 0]	91.94 +/- 10.90	83.61 +/- 15.86	81.94 +/- 15.30	97.50 +/- 3.72	91.94 +/- 7.55	100.00 +/- 0.00	22.22 +/- 22.43	P
[0, 0]	6.94 +/- 9.60	16.39 +/- 15.86	15.83 +/- 12.94	0.00 +/- 0.00	1.94 +/- 6.16	0.00 +/- 0.00	77.78 +/- 22.43	N
[1, 0]	0.83 +/- 1.98	0.00 +/- 0.00	4.17 +/- 4.93	3.33 +/- 3.69	14.44 +/- 9.36	1.39 +/- 4.26	2.22 +/- 5.39	F
[0, 0]	97.22 +/- 4.12	86.67 +/- 14.21	83.33 +/- 10.97	96.67 +/- 3.69	84.72 +/- 10.37	96.67 +/- 4.78	32.50 +/- 21.39	P
[0, 0]	1.94 +/- 3.18	13.33 +/- 14.21	12.50 +/- 8.22	0.00 +/- 0.00	0.83 +/- 1.98	1.94 +/- 3.18	65.28 +/- 20.70	N
[0, -1]	0.56 +/- 1.67	0.00 +/- 0.00	1.67 +/- 3.56	1.94 +/- 3.18	5.56 +/- 6.09	0.00 +/- 0.00	0.28 +/- 1.21	F
[0, 0]	95.28 +/- 5.35	75.83 +/- 21.10	81.39 +/- 14.83	97.78 +/- 4.08	91.11 +/- 11.97	97.50 +/- 5.41	25.83 +/- 21.54	P
[0, 0]	4.17 +/- 5.52	24.17 +/- 21.10	16.94 +/- 13.77	0.28 +/- 1.21	3.33 +/- 8.50	2.50 +/- 5.41	73.89 +/- 21.24	N
[-1, 0]	1.11 +/- 2.22	0.00 +/- 0.00	2.50 +/- 3.72	2.78 +/- 4.48	2.22 +/- 4.44	1.11 +/- 2.22	0.00 +/- 0.00	F
[0, 0]	66.11 +/- 25.45	19.44 +/- 21.34	35.28 +/- 29.67	65.28 +/- 39.71	15.00 +/- 19.33	79.17 +/- 24.84	2.50 +/- 5.41	P
[0, 0]	32.78 +/- 25.99	80.56 +/- 21.34	62.22 +/- 29.79	31.94 +/- 41.42	82.78 +/- 19.71	19.72 +/- 24.81	97.50 +/- 5.41	N

Table 6: **FreeForm case for Llama 3.** Final opinion distribution for combinations of “[Item A, Item B] [Reason for funding Item A, Reason for funding Item B]” (rows) and types of initial opinion distribution (columns). Each of the 63 final opinion distributions show the mean +/- standard deviation percentage of agents who want to provide full funding for Item A (F), partial funding for Item A (P), or no funding for Item A (N), averaged across 20 simulations. The connotation of Item A or B is as follows: 1 for a positive connotation, 0 for a neutral one, and -1 for a negative one. The same applies for the connotation of the reasons provided for each of the items: 1 for a positive connotation, 0 for a neutral one, and -1 for a negative one. See Table 1 for the specific values the items and their reasons for funding can take according to their connotation.

E.3 Initial opinion prompt details

The initial opinions are composed as follows: full funding for Item A has the template “I think that {Item A} should have all the funding because {Reason 1}.”; partial funding has the template “I think that we should provide measured funding for Item A because {Reason 2}.”; and no funding has the template “I think that {Item A} should not have any funding because {Reason 3}.”.

The variable Item A takes the text values as in Table 1 depending on its connotation.

The variable Reason 1 takes the value “they are constructive and important in our culture” when the reason for funding Item A has a positive connotation; “of REASON A” when it has a neutral connotation; “they are destructive and devastating

for our culture” when it has a negative connotation.

The variable Reason 2 takes the value “{Item B}, justified by REASON B, calls for some funding allocation. However, given the positive cultural importance of {Item A}, we should keep some funding for it” when the reason for funding Item A has a positive connotation and the one for Item B is neutral; it takes the value “{Item B}, justified by REASON B, calls for some funding allocation. However, given the negative cultural importance of {Item A}, we should keep some funding for it” when the reason for Item A has a negative connotation and the one for Item B is neutral; it takes the value “{Item B}, justified by REASON B, calls for some funding allocation. However, given REASON A justifying {Item A}, we should keep some funding for

Items & Reasons	Final Opinion Distribution (%)			
	Consensus-F	Consensus-P	Consensus-N	
[0, 0]	100.00 +/- 0.00	0.28 +/- 1.21	0.00 +/- 0.00	F
[0, 0]	0.00 +/- 0.00	98.33 +/- 5.00	0.00 +/- 0.00	P
[0, 0]	0.00 +/- 0.00	1.39 +/- 4.93	100.00 +/- 0.00	N
[0, 0]	96.39 +/- 15.74	13.06 +/- 17.42	0.00 +/- 0.00	F
[0, 1]	3.33 +/- 14.53	84.72 +/- 16.47	0.00 +/- 0.00	P
[0, 1]	0.28 +/- 1.21	2.22 +/- 4.08	100.00 +/- 0.00	N
[0, 0]	100.00 +/- 0.00	0.56 +/- 1.67	0.00 +/- 0.00	F
[0, 0]	0.00 +/- 0.00	95.00 +/- 5.24	0.00 +/- 0.00	P
[1, 0]	0.00 +/- 0.00	4.44 +/- 5.15	100.00 +/- 0.00	N
[0, 0]	90.83 +/- 27.51	0.56 +/- 1.67	0.00 +/- 0.00	F
[0, 0]	8.33 +/- 25.25	97.22 +/- 7.35	2.22 +/- 6.89	P
[0, -1]	0.83 +/- 3.63	2.22 +/- 7.33	97.78 +/- 6.89	N
[0, 0]	13.89 +/- 22.94	0.28 +/- 1.21	0.00 +/- 0.00	F
[0, 0]	65.83 +/- 26.83	96.11 +/- 9.15	0.00 +/- 0.00	P
[-1, 0]	20.28 +/- 19.35	3.61 +/- 9.18	100.00 +/- 0.00	N
[0, 1]	86.39 +/- 32.42	0.00 +/- 0.00	0.00 +/- 0.00	F
[0, 1]	13.61 +/- 32.42	100.00 +/- 0.00	0.00 +/- 0.00	P
[0, 0]	0.00 +/- 0.00	0.00 +/- 0.00	100.00 +/- 0.00	N
[1, 0]	98.89 +/- 4.84	2.22 +/- 4.78	0.00 +/- 0.00	F
[0, 0]	0.56 +/- 2.42	78.33 +/- 30.53	3.06 +/- 5.69	P
[0, 0]	0.56 +/- 2.42	19.44 +/- 30.30	96.94 +/- 5.69	N
[0, -1]	95.83 +/- 18.16	0.00 +/- 0.00	0.56 +/- 2.42	F
[0, -1]	3.89 +/- 16.95	98.61 +/- 3.88	1.39 +/- 6.05	P
[0, 0]	0.28 +/- 1.21	1.39 +/- 3.88	98.06 +/- 8.48	N
[-1, 0]	5.56 +/- 7.45	0.28 +/- 1.21	0.00 +/- 0.00	F
[-1, 0]	2.22 +/- 4.44	26.39 +/- 32.29	0.00 +/- 0.00	P
[0, 0]	92.22 +/- 10.00	73.33 +/- 32.23	100.00 +/- 0.00	N

Table 7: **FreeForm case for Llama 3.** Extension of Table 6, adding the final opinion distributions obtained when the initial distributions are Consensus-F, Consensus-P, and Consensus-N. We refer to the caption of Table 6.

it" when the reason for Item A has a neutral connotation and the one for Item B is neutral; it takes the value "the increasing appreciation of {Item B} in our culture calls for more funding for it. However, given REASON A justifying {Item A}, we should keep some funding for it" when the reason for Item A has a neutral connotation and the one for Item B is positive; it takes the value "our culture does not appreciate {Item B}, which calls for more funding for it. However, given REASON A justifying {Item A}, we should keep some funding for it" when the reason for Item A has a neutral connotation and the one for Item B is negative.

The variable Reason 3 takes the value "there is a large cultural appreciation of {Item B} which justifies reallocating all the funding for it." when the reason for funding Item B has a positive connotation; "{Item B} must get all the funding because of REASON B" when it has a neutral connotation; "there is a large disdain of {Item B} in our culture, which justifies reallocating all the funding for it." when it has a negative connotation.

The variable Item B takes the text values as in Table 1 depending on its connotation.

E.4 Interaction prompt details for FreeForm

The system prompt is "You are in a discussion with other participants".

In the case of agents without memory of past opinions, at time $t > 0$, the user prompt is: "This is your current opinion: "{Opinion at time $t - 1$ }". Now, you interact with someone having this opinion: "{Opinion of the other agent at time $t - 1$ ". State how much funding should be given to {Item A} after this interaction and explain why. Be concise with your answer."

In the case of agents with memory of past opinions, let us assume that the agent had consecutive interactions at the times $0 \leq \bar{t}_1 < \bar{t}_2 < \bar{t}_3$. For our purpose, we consider the initial time $t = 0$ as an "interaction time", i.e., we allow the possibility that $\bar{t}_1 = 0$. Then at time $t \geq \bar{t}_3 + 1 > 0$, the user prompt is: "This is your current opinion: "{Opinion at time $t - 1$, i.e., the opinion at time \bar{t}_3 }". These

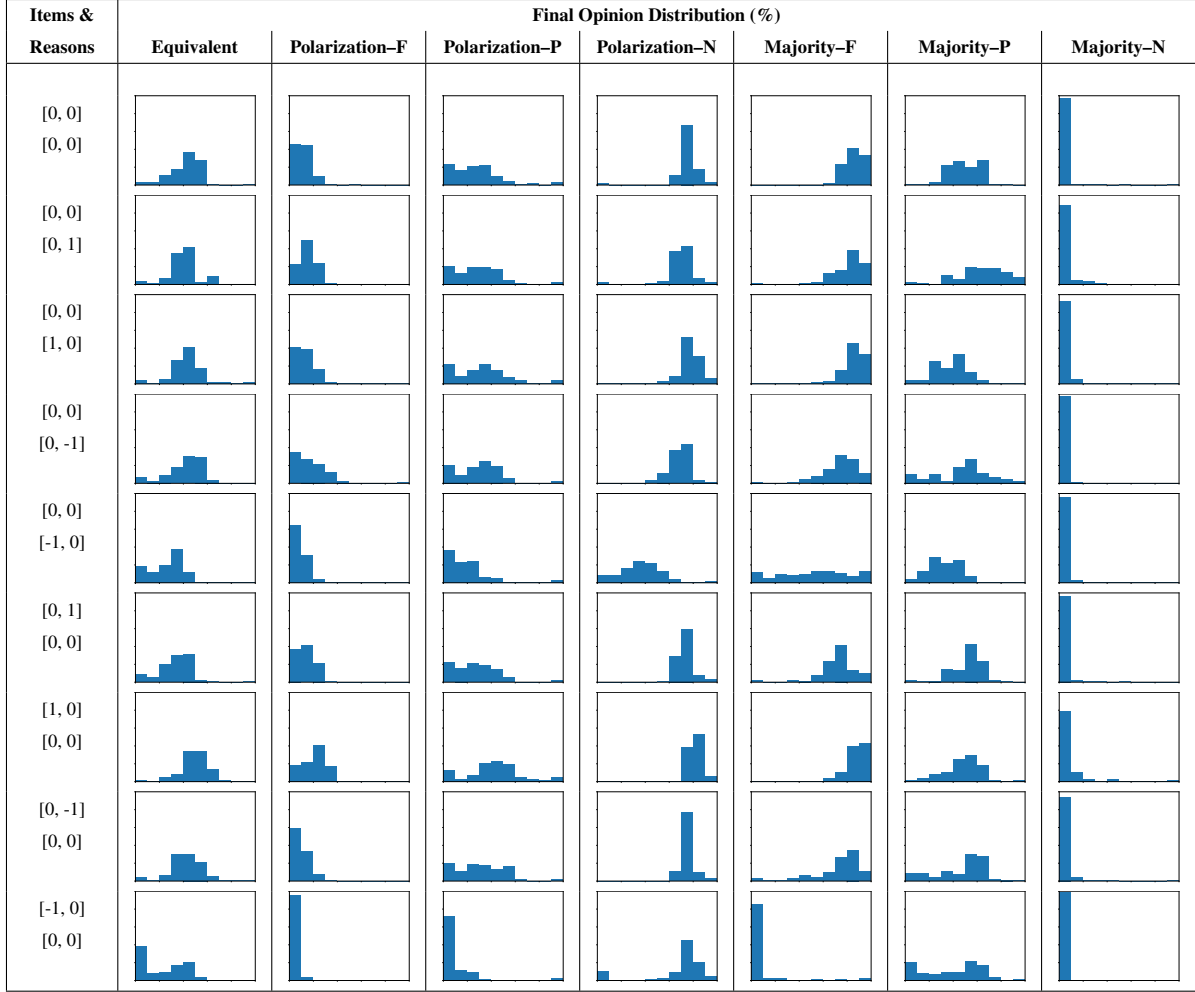


Table 8: **FreeForm case for Llama 3.** Histograms of percentage allocations in the final opinions for Item A for combinations of “[Item A, Item B] [Reason for Item A, Reason for Item B]” (rows) and types of initial opinion distribution (columns). We refer to the caption of Table 6 for the explanation regarding the rows. We only consider those final opinions which explicitly mentioned a percentage allocation for Item A, the rest of final opinions are ignored. The final opinions are considered across all simulations, i.e., at most 360 final opinions per histogram. In each of the 63 histograms, there is a maximum limit value of 0.1 in the y-axis, and 10 bins from 0% to 100% in increments of twenty in the x-axis.

are your previously held opinions sorted from the most recent to the oldest: Opinion 1: "{Opinion at time \bar{t}_2 }" Opinion 2: "{Opinion at time \bar{t}_1 }" State how much funding should be given to {Item A} after this interaction and explain why. Be concise with your answer." Now, let us assume that the agent has had less than three interactions. Then, we only show one or zero opinions additionally to the current opinion depending on whether there has been two or one interactions, respectively (in the latter case, there is no other interaction than the initial opinion, so we show the same prompt as in the case of a memoryless agent).

To avoid the LLM providing “the same” as its response without much justification for it, whenever such text is found, we add the text “, even if

the funding remains the same.” to the ending of the second to last sentence in the user prompt and query the LLM only one more time.

E.5 Interaction prompt details for ClosedForm

The system prompt is “You are in a discussion with other participants”.

In the case of agents without memory of past opinions, for Llama 3, at time $t > 0$, the user prompt is: “This is your current opinion: “{Opinion at time $t - 1$ }”. Now, you interact with someone having this opinion: “{Opinion of the other agent at time $t - 1$ }”. State which option (a), (b), or (c) is your new opinion regarding {Item A} after this interaction. Option (a) is “{Initial opinion for




Table 9: **FreeForm case for Llama 3.** Histograms of percentage allocations, continuation of Table 8 by adding the distributions obtained when the initial opinion distributions are Consensus-F, Consensus-P, and Consensus-N. We refer to the caption of Table 8.

Items & Reasons	Final Opinion Distribution (%)							
	Equivalent	Polarization-F	Polarization-P	Polarization-N	Majority-F	Majority-P	Majority-N	
[0, -1]	0.83 +/- 2.65	0.28 +/- 1.21	1.39 +/- 2.98	1.11 +/- 2.83	5.28 +/- 6.45	0.00 +/- 0.00	0.00 +/- 0.00	F
[0, 0]	93.33 +/- 10.03	79.44 +/- 16.58	84.17 +/- 13.86	98.89 +/- 2.83	93.33 +/- 6.48	99.17 +/- 2.65	30.28 +/- 22.94	P
	5.83 +/- 8.51	20.28 +/- 16.41	14.44 +/- 14.21	0.00 +/- 0.00	1.39 +/- 2.98	0.83 +/- 2.65	69.72 +/- 22.94	N

Items & Reasons	Final Opinion Distribution (%)			
	Consensus-F	Consensus-P	Consensus-N	
[0, -1]	90.83 +/- 21.39	0.00 +/- 0.00	0.00 +/- 0.00	F
[0, 0]	8.06 +/- 19.52	97.50 +/- 7.55	1.67 +/- 4.34	P
	1.11 +/- 2.83	2.50 +/- 7.55	98.33 +/- 4.34	N

Table 10: **FreeForm case for Llama 3: changing the text values.** Final opinion distribution for the case where Item B has a negative connotation but with a different text value. According to Table 1, Item B has the text value “nasty pollution”, but this table instead shows the final distributions obtained when we assign the same text value as negative Item A, i.e., “destructive bombs”. We refer to the caption of Table 6 for the description of each entry of the table.

full funding of Item A}”. Option (b) is “{Initial opinion for partial funding of Item A}”. Option (c) is “{Initial opinion for no funding of Item A}”. For Mistral, the user prompt at $t > 0$ is the same as the

one for Llama 3 with the following sentence added at the end of it: “Your response must always be in the following format: “Option: [write here (a), (b) or (c)].””

Items & Reasons	Final Opinion Distribution (%)							
	Equivalent	Polarization-F	Polarization-P	Polarization-N	Majority-F	Majority-P	Majority-N	
[0, 0]	0.00 +/- 0.00	0.00 +/- 0.00	0.00 +/- 0.00	0.00 +/- 0.00	0.28 +/- 1.21	0.00 +/- 0.00	0.00 +/- 0.00	F
[0, 0]	100.00 +/- 0.00	100.00 +/- 0.00	100.00 +/- 0.00	100.00 +/- 0.00	99.72 +/- 1.21	100.00 +/- 0.00	100.00 +/- 0.00	P
[0, 0]	0.00 +/- 0.00	0.00 +/- 0.00	0.00 +/- 0.00	0.00 +/- 0.00	0.00 +/- 0.00	0.00 +/- 0.00	0.00 +/- 0.00	N
[0, 0]	0.00 +/- 0.00	0.00 +/- 0.00	0.00 +/- 0.00	0.00 +/- 0.00	0.00 +/- 0.00	0.00 +/- 0.00	0.00 +/- 0.00	F
[0, 1]	100.00 +/- 0.00	100.00 +/- 0.00	100.00 +/- 0.00	100.00 +/- 0.00	100.00 +/- 0.00	100.00 +/- 0.00	100.00 +/- 0.00	P
[0, 1]	0.00 +/- 0.00	0.00 +/- 0.00	0.00 +/- 0.00	0.00 +/- 0.00	0.00 +/- 0.00	0.00 +/- 0.00	0.00 +/- 0.00	N
[0, 0]	0.00 +/- 0.00	0.00 +/- 0.00	0.00 +/- 0.00	0.00 +/- 0.00	0.00 +/- 0.00	0.00 +/- 0.00	0.00 +/- 0.00	F
[1, 0]	100.00 +/- 0.00	99.17 +/- 1.98	99.44 +/- 1.67	100.00 +/- 0.00	100.00 +/- 0.00	100.00 +/- 0.00	94.44 +/- 5.56	P
[1, 0]	0.00 +/- 0.00	0.83 +/- 1.98	0.56 +/- 1.67	0.00 +/- 0.00	0.00 +/- 0.00	0.00 +/- 0.00	5.56 +/- 5.56	N
[0, 0]	0.00 +/- 0.00	0.00 +/- 0.00	0.00 +/- 0.00	0.00 +/- 0.00	0.00 +/- 0.00	0.00 +/- 0.00	0.00 +/- 0.00	F
[0, -1]	100.00 +/- 0.00	100.00 +/- 0.00	100.00 +/- 0.00	100.00 +/- 0.00	100.00 +/- 0.00	100.00 +/- 0.00	100.00 +/- 0.00	P
[0, -1]	0.00 +/- 0.00	0.00 +/- 0.00	0.00 +/- 0.00	0.00 +/- 0.00	0.00 +/- 0.00	0.00 +/- 0.00	0.00 +/- 0.00	N
[0, 0]	0.00 +/- 0.00	0.00 +/- 0.00	0.00 +/- 0.00	0.00 +/- 0.00	0.00 +/- 0.00	0.00 +/- 0.00	0.00 +/- 0.00	F
[-1, 0]	78.06 +/- 5.95	50.00 +/- 0.00	76.39 +/- 7.43	100.00 +/- 0.00	99.72 +/- 1.21	94.44 +/- 0.00	16.39 +/- 1.21	P
[-1, 0]	21.94 +/- 5.95	50.00 +/- 0.00	23.61 +/- 7.43	0.00 +/- 0.00	0.28 +/- 1.21	5.56 +/- 0.00	83.61 +/- 1.21	N
[0, 1]	0.00 +/- 0.00	0.00 +/- 0.00	0.28 +/- 1.21	0.00 +/- 0.00	0.00 +/- 0.00	0.00 +/- 0.00	0.00 +/- 0.00	F
[0, 1]	100.00 +/- 0.00	100.00 +/- 0.00	99.72 +/- 1.21	100.00 +/- 0.00	100.00 +/- 0.00	100.00 +/- 0.00	100.00 +/- 0.00	P
[0, 0]	0.00 +/- 0.00	0.00 +/- 0.00	0.00 +/- 0.00	0.00 +/- 0.00	0.00 +/- 0.00	0.00 +/- 0.00	0.00 +/- 0.00	N
[1, 0]	0.00 +/- 0.00	0.00 +/- 0.00	0.00 +/- 0.00	0.00 +/- 0.00	0.00 +/- 0.00	0.00 +/- 0.00	0.00 +/- 0.00	F
[1, 0]	100.00 +/- 0.00	100.00 +/- 0.00	100.00 +/- 0.00	100.00 +/- 0.00	100.00 +/- 0.00	100.00 +/- 0.00	100.00 +/- 0.00	P
[0, 0]	0.00 +/- 0.00	0.00 +/- 0.00	0.00 +/- 0.00	0.00 +/- 0.00	0.00 +/- 0.00	0.00 +/- 0.00	0.00 +/- 0.00	N
[0, -1]	0.00 +/- 0.00	0.00 +/- 0.00	0.00 +/- 0.00	0.00 +/- 0.00	0.00 +/- 0.00	0.00 +/- 0.00	0.00 +/- 0.00	F
[0, -1]	100.00 +/- 0.00	100.00 +/- 0.00	100.00 +/- 0.00	100.00 +/- 0.00	100.00 +/- 0.00	100.00 +/- 0.00	100.00 +/- 0.00	P
[0, 0]	0.00 +/- 0.00	0.00 +/- 0.00	0.00 +/- 0.00	0.00 +/- 0.00	0.00 +/- 0.00	0.00 +/- 0.00	0.00 +/- 0.00	N
[-1, 0]	0.00 +/- 0.00	0.00 +/- 0.00	0.00 +/- 0.00	0.00 +/- 0.00	0.00 +/- 0.00	0.00 +/- 0.00	0.00 +/- 0.00	F
[-1, 0]	100.00 +/- 0.00	100.00 +/- 0.00	100.00 +/- 0.00	100.00 +/- 0.00	100.00 +/- 0.00	100.00 +/- 0.00	100.00 +/- 0.00	P
[0, 0]	0.00 +/- 0.00	0.00 +/- 0.00	0.00 +/- 0.00	0.00 +/- 0.00	0.00 +/- 0.00	0.00 +/- 0.00	0.00 +/- 0.00	N

Table 11: **ClosedForm case for Llama 3.** Final opinion distribution for combinations of “[Item A, Item B] [Reason for funding Item A, Reason for funding Item B]” (rows) and types of initial opinion distribution (columns). Each of the 63 final opinion distributions show the mean +/- standard deviation percentage of agents who want to provide full funding for Item A (F), partial funding for Item A (P), or no funding for Item A (N), averaged across 20 simulations. The connotation of Item A or B is as follows: 1 for a positive connotation, 0 for a neutral one, and -1 for a negative one. The same applies for the connotation of the reasons provided for each of the items: 1 for a positive connotation, 0 for a neutral one, and -1 for a negative one. See Table 1 for the specific values the items and their reasons for funding can take according to their connotation.

In the case of agents with memory of past opinions, the user prompt is modified similarly to the FreeForm case: showing the appropriate opinions from past interactions.

In all cases, we ensure that the LLM only selects one of the options.

E.6 Identifying the type of funding in the opinions for the FreeForm case

We designed a text-processing script to identify the type of funding the agent provides to Item A—namely, full, partial, or no funding—in its opinion text. We found that agents expressed how much funding to provide to Item A through both numerical values and plain text (i.e., without the use of numbers).

After our script was run across all final opinions for all of our experiments, we observed that the only ones left without classification were final opinions that mentioned that the funding they would provide to Item A was the same as the agent’s previous opinion. We call these “implicit opinions” because these final opinions did not explicitly provide any allocation in their text. For these cases, we investigated their previous opinions until finding the opinion which explicitly stated its allocation

Just to provide an idea of the type of responses from the LLMs that our text-processing script had to encounter, we provide some numbers for the case of agents without memory of past opinions. For Llama 3, we found that in the totality of our simulations, only 448 out of the 32400 final opin-

Items & Reasons	Final Opinion Distribution (%)			
	Consensus-F	Consensus-P	Consensus-N	
[0, 0]	100.00 +/- 0.00	0.00 +/- 0.00	0.00 +/- 0.00	F
[0, 0]	0.00 +/- 0.00	100.00 +/- 0.00	0.00 +/- 0.00	P
[0, 0]	0.00 +/- 0.00	0.00 +/- 0.00	100.00 +/- 0.00	N
[0, 0]	100.00 +/- 0.00	0.00 +/- 0.00	0.00 +/- 0.00	F
[0, 1]	0.00 +/- 0.00	100.00 +/- 0.00	0.00 +/- 0.00	P
[0, 1]	0.00 +/- 0.00	0.00 +/- 0.00	100.00 +/- 0.00	N
[0, 0]	100.00 +/- 0.00	0.00 +/- 0.00	0.00 +/- 0.00	F
[1, 0]	0.00 +/- 0.00	100.00 +/- 0.00	0.00 +/- 0.00	P
[1, 0]	0.00 +/- 0.00	0.00 +/- 0.00	100.00 +/- 0.00	N
[0, 0]	100.00 +/- 0.00	0.00 +/- 0.00	0.00 +/- 0.00	F
[0, -1]	0.00 +/- 0.00	100.00 +/- 0.00	0.00 +/- 0.00	P
[0, -1]	0.00 +/- 0.00	0.00 +/- 0.00	100.00 +/- 0.00	N
[0, 0]	100.00 +/- 0.00	0.00 +/- 0.00	0.00 +/- 0.00	F
[-1, 0]	0.00 +/- 0.00	100.00 +/- 0.00	0.00 +/- 0.00	P
[-1, 0]	0.00 +/- 0.00	0.00 +/- 0.00	100.00 +/- 0.00	N
[0, 1]	100.00 +/- 0.00	0.00 +/- 0.00	0.00 +/- 0.00	F
[0, 0]	0.00 +/- 0.00	100.00 +/- 0.00	0.00 +/- 0.00	P
[0, 0]	0.00 +/- 0.00	0.00 +/- 0.00	100.00 +/- 0.00	N
[1, 0]	100.00 +/- 0.00	0.00 +/- 0.00	0.00 +/- 0.00	F
[0, 0]	0.00 +/- 0.00	100.00 +/- 0.00	0.00 +/- 0.00	P
[0, 0]	0.00 +/- 0.00	0.00 +/- 0.00	100.00 +/- 0.00	N
[0, -1]	100.00 +/- 0.00	0.00 +/- 0.00	0.00 +/- 0.00	F
[0, -1]	0.00 +/- 0.00	100.00 +/- 0.00	0.00 +/- 0.00	P
[0, 0]	0.00 +/- 0.00	0.00 +/- 0.00	100.00 +/- 0.00	N
[-1, 0]	100.00 +/- 0.00	0.00 +/- 0.00	0.00 +/- 0.00	F
[-1, 0]	0.00 +/- 0.00	100.00 +/- 0.00	0.00 +/- 0.00	P
[0, 0]	0.00 +/- 0.00	0.00 +/- 0.00	100.00 +/- 0.00	N

Table 12: **ClosedForm case for Llama 3.** Extension of Table 11, adding the final opinion distributions obtained when the initial distributions are Consensus-F, Consensus-P, and Consensus-N. We refer to the caption of Table 11.

ions were implicit opinions, i.e., about 1.38% of all cases. For Mistral, we found that in the totality of our simulations, only 183 out of the 32400 final opinions were implicit opinions, i.e., about 0.56% of all cases. Moreover, we found that, for both LLMs, there was never the need to go all the way to the initial opinion to find out what the final opinion’s allocation was. It may be possible, however, for an opinion to become implicit at some interaction, and later become explicit again. Therefore, we decided to analyze every single opinion in our simulations (not just the final ones), which is a total of 2948400 opinions. For Llama 3, we found that it was only necessary for 1030 implicit opinions—i.e., about 0.03% of all opinions—to go back to the initial opinion to know what its funding allocation was. For Mistral, instead, we found that it was only necessary for 1632 implicit opinions—i.e., about 0.06% of all opinions.




Figure 3: **ClosedForm case for Llama 3.** Opinion evolution. Each of the nine rows of subplots corresponds to a discussion subject with the same order as in the rows of Table 11, and each of the seven columns of subplots corresponds to an initial opinion distribution with the same order as in the columns of Table 11. For each of the combinations of initial opinion distribution and discussion subject (a total of 63), we chose one simulation and plotted the evolution of the opinions in a subplot with each color curve corresponding to one agent’s opinion. Each subplot has the values 1, 0, -1 on the y-axis depending on the whether the value of the opinion was in favor of full funding, partial funding, or no funding for Item A, respectively. Each curve in a subplot corresponds to one opinion. The x-axis is the time t of the interactions (from 0 to 90) for the opinion updating; see Figure 2.



Figure 4: **ClosedForm case for Llama 3.** Opinion evolution. Each of the nine rows of subplots corresponds to a discussion subject with the same order as in the rows of Table 12, and each of the three columns of subplots corresponds to an initial opinion distribution with the same order as in the columns of Table 12. We refer to Figure 3 for details on how the opinions are plotted in the subplots.

Items & Reasons	Final Opinion Distribution (%)							
	Equivalent	Polarization-F	Polarization-P	Polarization-N	Majority-F	Majority-P	Majority-N	
[0, 0]	1.67 +/- 3.09	0.00 +/- 0.00	10.00 +/- 15.07	5.28 +/- 7.13	28.06 +/- 16.43	0.00 +/- 0.00	1.94 +/- 2.65	F
[0, 0]	93.89 +/- 5.80	79.44 +/- 14.71	62.22 +/- 26.39	94.72 +/- 7.13	65.83 +/- 18.53	99.44 +/- 1.67	17.78 +/- 12.37	P
[0, 0]	4.44 +/- 5.44	20.56 +/- 14.71	27.78 +/- 19.08	0.00 +/- 0.00	6.11 +/- 10.23	0.56 +/- 1.67	80.28 +/- 11.85	N
[0, 0]	2.78 +/- 3.29	0.00 +/- 0.00	10.28 +/- 9.98	3.06 +/- 4.11	25.56 +/- 16.33	0.00 +/- 0.00	2.78 +/- 5.96	F
[0, 1]	90.00 +/- 10.63	90.00 +/- 9.88	60.28 +/- 28.83	96.94 +/- 4.11	63.33 +/- 22.87	100.00 +/- 0.00	13.89 +/- 10.61	P
[0, 1]	7.22 +/- 8.80	10.00 +/- 9.88	29.44 +/- 21.52	0.00 +/- 0.00	11.11 +/- 10.97	0.00 +/- 0.00	83.33 +/- 11.65	N
[0, 0]	3.61 +/- 4.40	0.00 +/- 0.00	11.39 +/- 9.54	6.94 +/- 5.52	26.67 +/- 13.79	3.61 +/- 10.58	1.11 +/- 2.83	F
[1, 0]	84.17 +/- 12.82	78.89 +/- 17.18	63.89 +/- 14.75	92.50 +/- 6.40	71.39 +/- 16.69	94.72 +/- 11.18	22.22 +/- 15.21	P
[1, 0]	12.22 +/- 12.25	21.11 +/- 17.18	24.72 +/- 11.31	0.56 +/- 2.42	1.94 +/- 4.03	1.67 +/- 5.00	76.67 +/- 14.55	N
[0, 0]	0.56 +/- 1.67	1.11 +/- 2.83	3.61 +/- 4.40	2.78 +/- 4.12	21.11 +/- 16.81	2.22 +/- 5.93	2.22 +/- 3.24	F
[0, -1]	91.39 +/- 10.32	79.44 +/- 16.86	75.56 +/- 15.36	97.22 +/- 4.12	73.89 +/- 18.68	93.61 +/- 13.64	29.17 +/- 15.50	P
[0, -1]	8.06 +/- 10.01	19.44 +/- 16.53	20.83 +/- 13.93	0.00 +/- 0.00	5.00 +/- 9.61	4.17 +/- 12.16	68.61 +/- 14.30	N
[0, 0]	3.61 +/- 4.40	0.00 +/- 0.00	7.22 +/- 6.11	4.72 +/- 4.40	28.61 +/- 16.03	0.00 +/- 0.00	3.89 +/- 5.30	F
[-1, 0]	79.44 +/- 17.84	65.83 +/- 21.96	60.56 +/- 25.15	81.11 +/- 17.52	61.11 +/- 19.08	94.44 +/- 14.59	23.06 +/- 16.88	P
[-1, 0]	16.94 +/- 16.80	34.17 +/- 21.96	32.22 +/- 21.70	14.17 +/- 16.43	10.28 +/- 12.46	5.56 +/- 14.59	73.06 +/- 14.83	N
[0, 1]	2.22 +/- 3.24	0.28 +/- 1.21	7.78 +/- 9.20	2.78 +/- 4.81	35.56 +/- 19.98	0.83 +/- 2.65	1.39 +/- 2.41	F
[0, 0]	88.33 +/- 11.10	78.06 +/- 20.14	58.61 +/- 23.27	97.22 +/- 4.81	60.28 +/- 20.88	97.50 +/- 4.11	23.33 +/- 12.25	P
[0, 0]	9.44 +/- 9.95	21.67 +/- 20.33	33.61 +/- 17.61	0.00 +/- 0.00	4.17 +/- 11.09	1.67 +/- 3.09	75.28 +/- 10.90	N
[1, 0]	2.22 +/- 4.44	0.00 +/- 0.00	17.22 +/- 14.90	8.06 +/- 8.33	40.83 +/- 19.82	4.72 +/- 10.87	3.89 +/- 6.11	F
[0, 0]	91.94 +/- 8.51	88.61 +/- 8.51	63.06 +/- 26.54	91.94 +/- 8.33	55.00 +/- 21.15	94.44 +/- 11.25	43.61 +/- 13.41	P
[0, 0]	5.83 +/- 7.55	11.39 +/- 8.51	19.72 +/- 16.05	0.00 +/- 0.00	4.17 +/- 5.23	0.83 +/- 1.98	52.50 +/- 11.45	N
[0, -1]	1.94 +/- 3.18	0.28 +/- 1.21	4.17 +/- 5.52	4.44 +/- 5.44	26.39 +/- 17.11	1.67 +/- 3.56	3.33 +/- 5.39	F
[0, 0]	90.83 +/- 9.51	75.00 +/- 14.00	73.06 +/- 22.92	95.56 +/- 5.44	67.78 +/- 19.29	95.56 +/- 7.16	23.06 +/- 15.94	P
[0, 0]	7.22 +/- 8.07	24.72 +/- 14.22	22.78 +/- 21.87	0.00 +/- 0.00	5.83 +/- 9.04	2.78 +/- 6.92	73.61 +/- 16.19	N
[-1, 0]	2.78 +/- 3.29	0.00 +/- 0.00	5.00 +/- 7.43	9.72 +/- 8.76	12.50 +/- 12.77	10.28 +/- 10.43	0.56 +/- 1.67	F
[0, 0]	74.17 +/- 16.13	36.94 +/- 27.46	58.06 +/- 26.73	88.06 +/- 9.82	19.72 +/- 21.26	81.94 +/- 15.70	5.00 +/- 6.78	P
[0, 0]	23.06 +/- 14.52	63.06 +/- 27.46	36.94 +/- 22.45	2.22 +/- 5.67	67.78 +/- 28.09	7.78 +/- 13.88	94.44 +/- 6.57	N

Table 13: **FreeForm case for Llama 3 with memory of past opinions.** Final opinion distribution for combinations of “[Item A, Item B] [Reason for funding Item A, Reason for funding Item B]” (rows) and types of initial opinion distribution (columns). Each of the 63 final opinion distributions show the mean +/- standard deviation percentage of agents who want to provide full funding for Item A (F), partial funding for Item A (P), or no funding for Item A (N), averaged across 20 simulations. The connotation of Item A or B is as follows: 1 for a positive connotation, 0 for a neutral one, and -1 for a negative one. The same applies for the connotation of the reasons provided for each of the items: 1 for a positive connotation, 0 for a neutral one, and -1 for a negative one. See Table 1 for the specific values the items and their reasons for funding can take according to their connotation.

Items & Reasons	Final Opinion Distribution (%)			
	Consensus-F	Consensus-P	Consensus-N	
[0, 0]	100.00 +/- 0.00	3.61 +/- 15.74	0.00 +/- 0.00	F
[0, 0]	0.00 +/- 0.00	96.39 +/- 15.74	0.00 +/- 0.00	P
[0, 0]	0.00 +/- 0.00	0.00 +/- 0.00	100.00 +/- 0.00	N
[0, 0]	100.00 +/- 0.00	0.00 +/- 0.00	0.00 +/- 0.00	F
[0, 1]	0.00 +/- 0.00	100.00 +/- 0.00	0.28 +/- 1.21	P
[0, 1]	0.00 +/- 0.00	0.00 +/- 0.00	99.72 +/- 1.21	N
[0, 0]	100.00 +/- 0.00	0.56 +/- 2.42	0.00 +/- 0.00	F
[1, 0]	0.00 +/- 0.00	98.33 +/- 3.56	0.00 +/- 0.00	P
[1, 0]	0.00 +/- 0.00	1.11 +/- 2.83	100.00 +/- 0.00	N
[0, 0]	100.00 +/- 0.00	0.00 +/- 0.00	0.00 +/- 0.00	F
[0, -1]	0.00 +/- 0.00	100.00 +/- 0.00	0.00 +/- 0.00	P
[0, -1]	0.00 +/- 0.00	0.00 +/- 0.00	100.00 +/- 0.00	N
[0, 0]	88.06 +/- 25.66	0.00 +/- 0.00	0.00 +/- 0.00	F
[-1, 0]	3.89 +/- 9.48	98.33 +/- 6.11	0.00 +/- 0.00	P
[-1, 0]	8.06 +/- 16.71	1.67 +/- 6.11	100.00 +/- 0.00	N
[0, 1]	100.00 +/- 0.00	0.00 +/- 0.00	0.00 +/- 0.00	F
[0, 0]	0.00 +/- 0.00	99.44 +/- 1.67	0.00 +/- 0.00	P
[0, 0]	0.00 +/- 0.00	0.56 +/- 1.67	100.00 +/- 0.00	N
[1, 0]	100.00 +/- 0.00	0.00 +/- 0.00	0.00 +/- 0.00	F
[0, 0]	0.00 +/- 0.00	95.56 +/- 14.55	0.00 +/- 0.00	P
[0, 0]	0.00 +/- 0.00	4.44 +/- 14.55	100.00 +/- 0.00	N
[0, -1]	100.00 +/- 0.00	0.00 +/- 0.00	0.00 +/- 0.00	F
[0, 0]	0.00 +/- 0.00	100.00 +/- 0.00	0.28 +/- 1.21	P
[0, 0]	0.00 +/- 0.00	0.00 +/- 0.00	99.72 +/- 1.21	N
[-1, 0]	23.06 +/- 18.45	0.56 +/- 1.67	0.00 +/- 0.00	F
[0, 0]	1.94 +/- 3.63	99.17 +/- 2.65	0.00 +/- 0.00	P
[0, 0]	75.00 +/- 20.07	0.28 +/- 1.21	100.00 +/- 0.00	N

Table 14: **FreeForm case for Llama 3 with memory of past opinions.** Extension of Table 13, adding the final opinion distributions obtained when the initial distributions are Consensus-F, Consensus-P, and Consensus-N. We refer to the caption of Table 13.

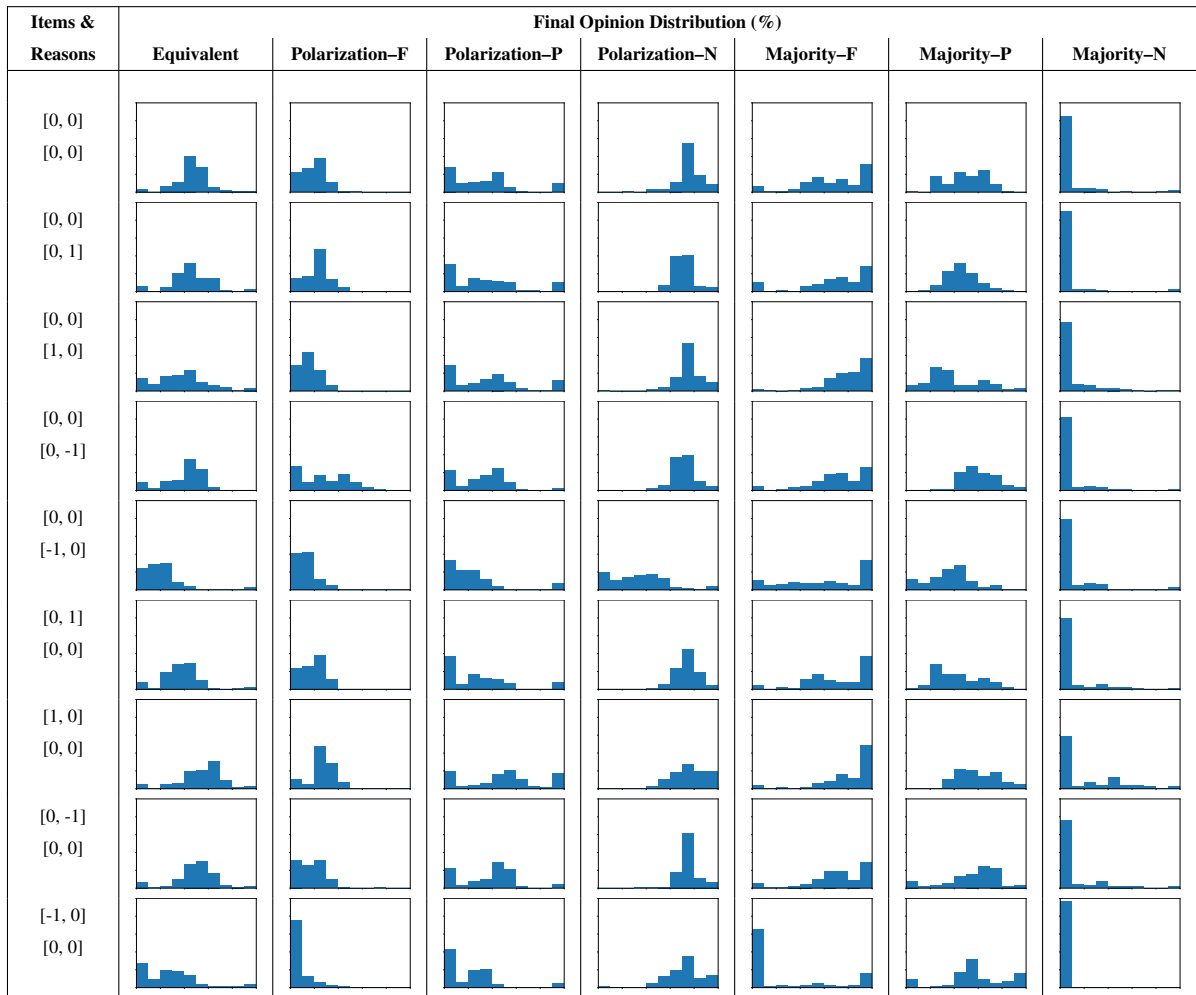


Table 15: **FreeForm case for Llama 3 with memory of past opinions.** Histograms of percentage allocations in the final opinions for Item A for combinations of “[Item A, Item B] [Reason for Item A, Reason for Item B]” (rows) and types of initial opinion distribution (columns). We refer to the caption of Table 13 for the explanation regarding the rows. We only consider those final opinions which explicitly mentioned a percentage allocation for Item A, the rest of final opinions are ignored. The final opinions are considered across all simulations, i.e., at most 360 final opinions per histogram. In each of the 63 histograms, there is a maximum limit value of 0.1 in the y-axis, and 10 bins from 0% to 100% in increments of twenty in the x-axis.

Items & Reasons	Final Opinion Distribution (%)		
	Consensus-F	Consensus-P	Consensus-N
[0, 0] [0, 0]		 	
[0, 0] [0, 1]		 	
[0, 0] [1, 0]			
[0, 0] [0, -1]			
[0, 0] [-1, 0]			
[0, 1] [0, 0]		 	
[1, 0] [0, 0]			
[0, -1] [0, 0]		 	
[-1, 0] [0, 0]			

Table 16: **FreeForm case for Llama 3 with memory of past opinions.** Histograms of percentage allocations, continuation of Table 15 by adding the distributions obtained when the initial opinion distributions are Consensus-F, Consensus-P, and Consensus-N. We refer to the caption of Table 8.

Items & Reasons	Final Opinion Distribution (%)							
	Equivalent	Polarization-F	Polarization-P	Polarization-N	Majority-F	Majority-P	Majority-N	
[0, 0]	0.00 +/- 0.00	0.00 +/- 0.00	0.00 +/- 0.00	0.00 +/- 0.00	0.00 +/- 0.00	0.00 +/- 0.00	0.00 +/- 0.00	F
[0, 0]	100.00 +/- 0.00	100.00 +/- 0.00	100.00 +/- 0.00	100.00 +/- 0.00	100.00 +/- 0.00	100.00 +/- 0.00	100.00 +/- 0.00	P
[0, 0]	0.00 +/- 0.00	0.00 +/- 0.00	0.00 +/- 0.00	0.00 +/- 0.00	0.00 +/- 0.00	0.00 +/- 0.00	0.00 +/- 0.00	N
[0, 0]	0.00 +/- 0.00	0.00 +/- 0.00	0.00 +/- 0.00	0.00 +/- 0.00	0.00 +/- 0.00	0.00 +/- 0.00	0.00 +/- 0.00	F
[0, 0]	99.44 +/- 2.42	97.78 +/- 5.93	98.06 +/- 3.63	100.00 +/- 0.00	100.00 +/- 0.00	100.00 +/- 0.00	58.61 +/- 20.22	P
[0, 1]	0.56 +/- 2.42	2.22 +/- 5.93	1.94 +/- 3.63	0.00 +/- 0.00	0.00 +/- 0.00	0.00 +/- 0.00	41.39 +/- 20.22	N
[0, 0]	0.00 +/- 0.00	0.00 +/- 0.00	0.00 +/- 0.00	0.00 +/- 0.00	0.00 +/- 0.00	0.00 +/- 0.00	0.00 +/- 0.00	F
[0, 0]	98.06 +/- 4.40	93.33 +/- 8.53	96.11 +/- 6.36	100.00 +/- 0.00	100.00 +/- 0.00	100.00 +/- 0.00	49.72 +/- 16.80	P
[1, 0]	1.94 +/- 4.40	6.67 +/- 8.53	3.89 +/- 6.36	0.00 +/- 0.00	0.00 +/- 0.00	0.00 +/- 0.00	50.28 +/- 16.80	N
[0, 0]	0.00 +/- 0.00	0.00 +/- 0.00	0.00 +/- 0.00	0.00 +/- 0.00	0.00 +/- 0.00	0.00 +/- 0.00	0.00 +/- 0.00	F
[0, 0]	99.44 +/- 1.67	98.06 +/- 3.63	98.61 +/- 3.46	100.00 +/- 0.00	100.00 +/- 0.00	100.00 +/- 0.00	61.94 +/- 20.96	P
[0, -1]	0.56 +/- 1.67	1.94 +/- 3.63	1.39 +/- 3.46	0.00 +/- 0.00	0.00 +/- 0.00	0.00 +/- 0.00	38.06 +/- 20.96	N
[0, 0]	0.00 +/- 0.00	0.00 +/- 0.00	0.00 +/- 0.00	0.00 +/- 0.00	0.00 +/- 0.00	0.00 +/- 0.00	0.00 +/- 0.00	F
[0, 0]	86.39 +/- 6.45	55.83 +/- 4.11	79.44 +/- 6.83	100.00 +/- 0.00	100.00 +/- 0.00	94.72 +/- 1.21	17.22 +/- 1.67	P
[-1, 0]	13.61 +/- 6.45	44.17 +/- 4.11	20.56 +/- 6.83	0.00 +/- 0.00	0.00 +/- 0.00	5.28 +/- 1.21	82.78 +/- 1.67	N
[0, 1]	0.00 +/- 0.00	0.00 +/- 0.00	0.00 +/- 0.00	0.00 +/- 0.00	0.00 +/- 0.00	0.00 +/- 0.00	0.00 +/- 0.00	F
[0, 0]	100.00 +/- 0.00	100.00 +/- 0.00	100.00 +/- 0.00	100.00 +/- 0.00	100.00 +/- 0.00	100.00 +/- 0.00	100.00 +/- 0.00	P
[0, 0]	0.00 +/- 0.00	0.00 +/- 0.00	0.00 +/- 0.00	0.00 +/- 0.00	0.00 +/- 0.00	0.00 +/- 0.00	0.00 +/- 0.00	N
[1, 0]	0.00 +/- 0.00	0.00 +/- 0.00	0.00 +/- 0.00	0.00 +/- 0.00	0.00 +/- 0.00	0.00 +/- 0.00	0.00 +/- 0.00	F
[0, 0]	100.00 +/- 0.00	100.00 +/- 0.00	100.00 +/- 0.00	100.00 +/- 0.00	100.00 +/- 0.00	100.00 +/- 0.00	100.00 +/- 0.00	P
[0, 0]	0.00 +/- 0.00	0.00 +/- 0.00	0.00 +/- 0.00	0.00 +/- 0.00	0.00 +/- 0.00	0.00 +/- 0.00	0.00 +/- 0.00	N
[0, -1]	0.00 +/- 0.00	0.00 +/- 0.00	0.00 +/- 0.00	0.00 +/- 0.00	0.00 +/- 0.00	0.00 +/- 0.00	0.00 +/- 0.00	F
[0, 0]	100.00 +/- 0.00	100.00 +/- 0.00	100.00 +/- 0.00	100.00 +/- 0.00	100.00 +/- 0.00	100.00 +/- 0.00	100.00 +/- 0.00	P
[0, 0]	0.00 +/- 0.00	0.00 +/- 0.00	0.00 +/- 0.00	0.00 +/- 0.00	0.00 +/- 0.00	0.00 +/- 0.00	0.00 +/- 0.00	N
[-1, 0]	0.00 +/- 0.00	0.00 +/- 0.00	0.00 +/- 0.00	0.00 +/- 0.00	0.00 +/- 0.00	0.00 +/- 0.00	0.00 +/- 0.00	F
[0, 0]	97.78 +/- 4.44	91.11 +/- 8.68	93.89 +/- 7.43	100.00 +/- 0.00	100.00 +/- 0.00	100.00 +/- 0.00	43.89 +/- 20.10	P
[0, 0]	2.22 +/- 4.44	8.89 +/- 8.68	6.11 +/- 7.43	0.00 +/- 0.00	0.00 +/- 0.00	0.00 +/- 0.00	56.11 +/- 20.10	N

Table 17: **ClosedForm case for Llama 3 with memory of past opinions.** Final opinion distribution for combinations of “[Item A, Item B] [Reason for funding Item A, Reason for funding Item B]” (rows) and types of initial opinion distribution (columns). Each of the 63 final opinion distributions show the mean +/- standard deviation percentage of agents who want to provide full funding for Item A (F), partial funding for Item A (P), or no funding for Item A (N), averaged across 20 simulations. The connotation of Item A or B is as follows: 1 for a positive connotation, 0 for a neutral one, and -1 for a negative one. The same applies for the connotation of the reasons provided for each of the items: 1 for a positive connotation, 0 for a neutral one, and -1 for a negative one. See Table 1 for the specific values the items and their reasons for funding can take according to their connotation.

Items & Reasons	Final Opinion Distribution (%)			
	Consensus-F	Consensus-P	Consensus-N	
[0, 0]	100.00 +/- 0.00	0.00 +/- 0.00	0.00 +/- 0.00	F
[0, 0]	0.00 +/- 0.00	100.00 +/- 0.00	0.00 +/- 0.00	P
[0, 0]	0.00 +/- 0.00	0.00 +/- 0.00	100.00 +/- 0.00	N
[0, 0]	100.00 +/- 0.00	0.00 +/- 0.00	0.00 +/- 0.00	F
[0, 1]	0.00 +/- 0.00	100.00 +/- 0.00	0.00 +/- 0.00	P
[0, 1]	0.00 +/- 0.00	0.00 +/- 0.00	100.00 +/- 0.00	N
[0, 0]	100.00 +/- 0.00	0.00 +/- 0.00	0.00 +/- 0.00	F
[1, 0]	0.00 +/- 0.00	100.00 +/- 0.00	0.00 +/- 0.00	P
[1, 0]	0.00 +/- 0.00	0.00 +/- 0.00	100.00 +/- 0.00	N
[0, 0]	100.00 +/- 0.00	0.00 +/- 0.00	0.00 +/- 0.00	F
[0, -1]	0.00 +/- 0.00	100.00 +/- 0.00	0.00 +/- 0.00	P
[0, -1]	0.00 +/- 0.00	0.00 +/- 0.00	100.00 +/- 0.00	N
[0, 0]	100.00 +/- 0.00	0.00 +/- 0.00	0.00 +/- 0.00	F
[-1, 0]	0.00 +/- 0.00	100.00 +/- 0.00	0.00 +/- 0.00	P
[-1, 0]	0.00 +/- 0.00	0.00 +/- 0.00	100.00 +/- 0.00	N
[0, 1]	100.00 +/- 0.00	0.00 +/- 0.00	0.00 +/- 0.00	F
[0, 1]	0.00 +/- 0.00	100.00 +/- 0.00	0.00 +/- 0.00	P
[0, 1]	0.00 +/- 0.00	0.00 +/- 0.00	100.00 +/- 0.00	N
[1, 0]	100.00 +/- 0.00	0.00 +/- 0.00	0.00 +/- 0.00	F
[1, 0]	0.00 +/- 0.00	100.00 +/- 0.00	0.00 +/- 0.00	P
[1, 0]	0.00 +/- 0.00	0.00 +/- 0.00	100.00 +/- 0.00	N
[0, -1]	100.00 +/- 0.00	0.00 +/- 0.00	0.00 +/- 0.00	F
[0, -1]	0.00 +/- 0.00	100.00 +/- 0.00	0.00 +/- 0.00	P
[0, -1]	0.00 +/- 0.00	0.00 +/- 0.00	100.00 +/- 0.00	N
[-1, 0]	100.00 +/- 0.00	0.00 +/- 0.00	0.00 +/- 0.00	F
[-1, 0]	0.00 +/- 0.00	100.00 +/- 0.00	0.00 +/- 0.00	P
[-1, 0]	0.00 +/- 0.00	0.00 +/- 0.00	100.00 +/- 0.00	N

Table 18: **ClosedForm case for Llama 3 with memory of past opinions.** Extension of Table 17, adding the final opinion distributions obtained when the initial distributions are Consensus-F, Consensus-P, and Consensus-N. We refer to the caption of Table 17.

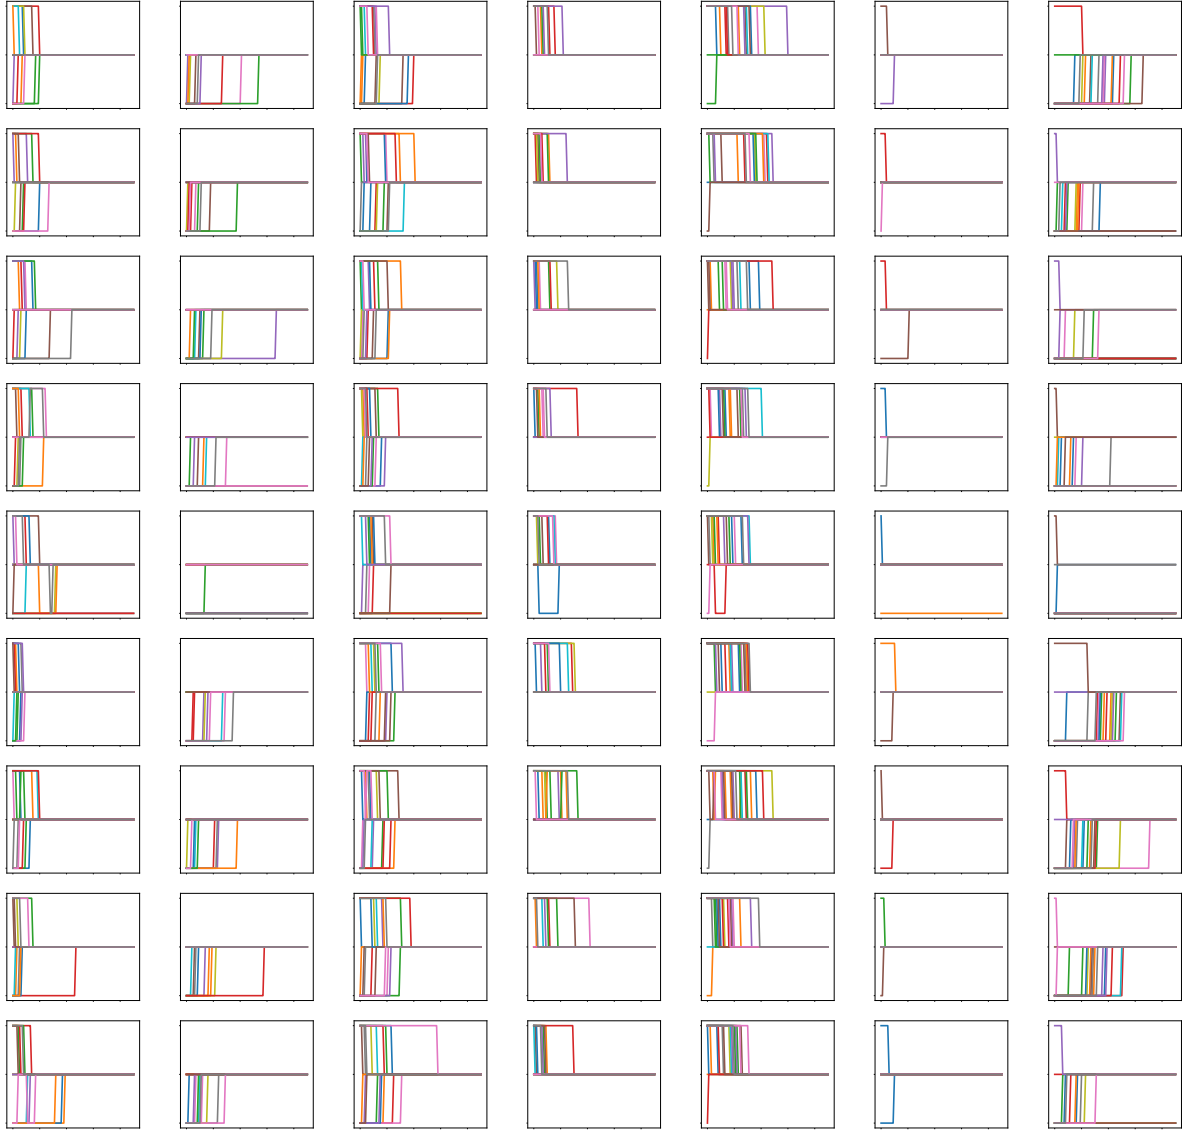


Figure 5: **ClosedForm case for Llama 3 with memory of past opinions.** Opinion evolution. Each of the nine rows of subplots corresponds to a discussion subject with the same order as in the rows of Table 17, and each of the seven columns of subplots corresponds to an initial opinion distribution with the same order as in the columns of Table 17. For each of the combinations of initial opinion distribution and discussion subject (a total of 63), we chose one simulation and plotted the evolution of the opinions in a subplot with each color curve corresponding to one agent's opinion. Each subplot has the values 1, 0, -1 on the y-axis depending on the whether the value of the opinion was in favor of full funding, partial funding, or no funding for Item A, respectively. Each curve in a subplot corresponds to one opinion. The x-axis is the time t of the interactions (from 0 to 90) for the opinion updating; see Figure 2.




Figure 6: **ClosedForm case for Llama 3 with memory of past opinions.** Opinion evolution. Each of the nine rows of subplots corresponds to a discussion subject with the same order as in the rows of Table 18, and each of the three columns of subplots corresponds to an initial opinion distribution with the same order as in the columns of Table 18. We refer to Figure 5 for details on how the opinions are plotted in the subplots.