Modeling human clarification question production through expected regret

Polina Tsvilodub[1*], Michael Franke[1*], Robert D. Hawkins[2]

[1] University of Tübingen; [2]Stanford University

polina.tsvilodub@uni-tuebingen.de

Wh-questions are ubiquitous in everyday conversations, and there are many ways in which speakers respond to them. For example, suppose a customer asks a bartender: (1) 'What drinks do you have?' The bartender could list all 30 beverages on the menu (i.e., provide a costly *exhaustive* answer), or opt to mention just a few relevant options (i.e., provide a *mention-some* answer). Decision-theoretic accounts of question answering in pragmatics posit that cooperative answerers mention information that maximizes the questioner's expected utility with respect to their *latent goal* they might guess based on common knowledge (e.g., getting a cocktail) [1, 2]. But when faced with uncertainty about the questioner's goal (e.g., do they prefer cocktails or soft drinks), the bartender might also answer with a (clarification) question: (2) 'Are you looking for a cocktail or a soft drink?'. Recent probabilistic models of question-answering have suggested that answerers trade off two factors when choosing their answer: (1) uncertainty about the questioner's goal, and (2) the cost of different responses [3, 4]. We propose that answerers navigate this trade-off through *expected regret* [5]: they consider how much they might regret giving any available answer (i.e., exhaustive r_{exh} or mention-some r_{ms}) if wrong about the questioner's goal (here: $G = \{g_1, g_2\}$). When expected regret is high — either because uncertainty is high or because an exhaustive answer is costly — answerers should prefer to ask clarification questions.
We formalize the answerer's beliefs about the questioner goals as $P(G)$. The answerer considers a default answer $r_j \in R = \{r_{exh}, r_{ms_1}, r_{ms_2}\}$ with the highest expected utility: $r^* = \arg\max_{r_j} \mathop{\mathbb{E}}_{g_i \in G}[U(g_i, r_j)]$. We assume the utilities of answers listed in Table 1, where exhaustive answers have lower utility due to higher production cost (since options for both goals have to be listed). We then calculate the *expected regret* of r^*: *ExpRegret*$(r^*) = \sum_{g_i \in G} P(g_i) \cdot$ *Regret*(g_i, r^*), where *Regret*$(g_i, r^*) = \arg\max_{r_j} U(g_i, r_j) - U(g_i, r^*)$. We model the probability of asking a clarification question r_{cq} proportionally to scaled expected regret: $P(r_{cq}) = (1 + \exp(-\textit{ExpRegret}(r^*)))^{-1}$. The probability of answering with $r_j \in \{r_{exh}, r_{ms}\}$ is then $P(r_j) = (1 - P(r_{cq})) \cdot$ *softmax*$(EU(r_j))$. We test whether our model accounts for the two contextual factors that could influence the rate of r_{cq} by varying (1) the uncertainty about the questioner's goals through the parameter $P(g_1) = 1 - \epsilon$, and (2) the cost of the exhaustive answer as $U(r_{exh}) = 1 - \delta$ (see Table 1). The predicted probability of choosing r_{cq} is shown in Figure 1 (left), and fit to human data in Figure 1 (right), indicating that the model predicts a higher clarification question rate with higher uncertainty ($\epsilon \to 0.5$) and lower utility of r_{exh} ($\delta \to 1$).
We tested these predictions in a web-based 2×2 factorial experiment. Participants (N=41) saw four trials as in example trial (I). Specific options within two basic-level categories were listed (two or four options per basic-level category; *small* or *large* option space). In context, questioners were equally likely to prefer either basic-level category (*high* uncertainty), or almost certainly preferred one (*low* uncertainty). Participants selected a response among a clarification question, an exhaustive response and two mention-some responses. Figure 1 (right) confirms that participants preferred asking clarification questions under high uncertainty about the questioner goal when the exhaustive answer was costly (marginally corroborated by a Bayesian logistic regression, $\beta = -1.62[-3.72, 0.00]$). Future work should extend to modeling the interaction between uncertainty and cost of exhaustive answers, and explore other information-theoretic measures for modeling speakers' decisions.

Goal	$P(G)$	$U(r_{ms_1})$	$U(r_{ms_2})$	$U(r_{exh})$
g_1	$1-\epsilon$	1	0	$1-\delta$
g_2	ϵ	0	1	$1-\delta$

Table 1: General structure of the answerer's beliefs about the questioner's goals G and payoffs U for the answerer associated with different default responses. We assume two possible goals.

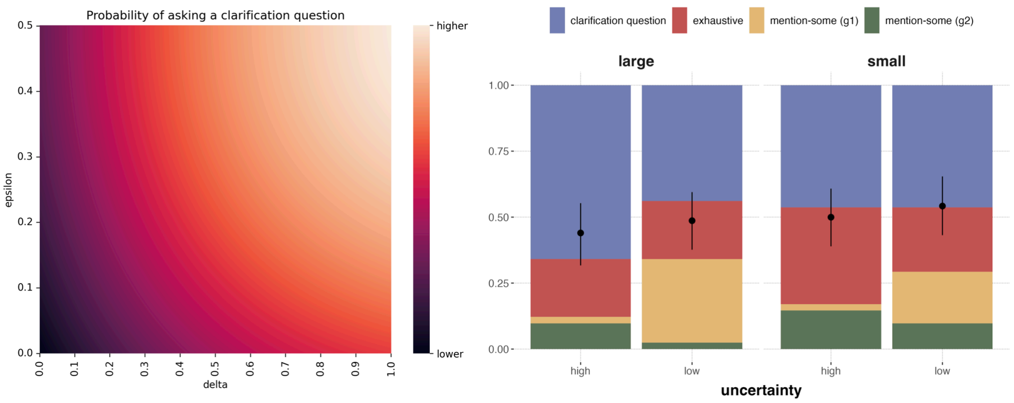

Figure 1: **Left**: Probability of asking a clarification question, depending on the uncertainty about the questioner's goals $G = \{g_1, g_2\}$, with $P(g_1) = 1-\epsilon$ and $P(g_2) = \epsilon$, and the utility of the exhaustive answer $U(r_{exh}) = 1-\delta$. **Right**: Proportions of different responses (color) selected by participants (y-axis), given different uncertainty conditions (x-axis) and different total number of available options (eight in total in the left facet, four in total in the right facet). Black dots indicate posterior means and 95-% CrIs of the rate of clarification questions when fitting ϵ and δ to human data, along with a temperature and cost parameter for $P(r_{cq})$. Means of Bayesian posterior predictive p-values for all conditions are well above 0, suggesting qualitative fit. An alternative model where $P(r_{cq}) = (1 + \exp(-\sum_{r_j} P(r_j) \cdot \mathit{ExpRegret}(r_j)))^{-1}$ was not credibly different from the reported model.

(I) Example trial (high uncertainty, small option space condition): You are a bartender in a bar. The bar offers cocktails and soft drinks. Today the bar has the following menu (presented in a picture in the experiment): Cocktails: Mojito, Tequila Sunrise. Soft drinks: Ginger beer, orange juice. From experience you know that customers are equally likely to order cocktails and soft drinks. A customer walks in and asks: "What drinks do you have?" You reply: ( ) Do you prefer cocktails or soft drinks? ( ) If you are looking for soft drinks, we have ginger beer and orange juice. ( ) We have a Mojito, ginger beer, a Tequila Sunrise, and orange juice. ( ) If you are looking for cocktails, we have Mojito and Tequila Sunrise.

[1] van Rooy, R. (2003). Questioning to resolve decision problems. *Linguistics and Philosophy*, *26*(6), 727–763.

[2] Benz, A. (2006). *Utility and relevance of answers*. Springer.

[3] Hawkins, R. D., & Goodman, N. D. (2017). *Why do you ask? The informational dynamics of questions and answers* [PsyArXiv].

[4] Hawkins, R. D., Tsvilodub, P., Bergey, C. A., Goodman, N. D., & Franke, M. (2025). Relevant answers to polar questions. *Philosophical Transactions B*, *380*(1932), 20230505.

[5] Loomes, G., & Sugden, R. (1982). Regret theory: An alternative theory of rational choice under uncertainty. *The economic journal*, *92*(368), 805–824.