Accelerated Preference Elicitation with LLM-Based Proxies

David Huang[1] Edwin Lock[2] Francisco Marmolejo-Cossío[1,3,4] David Parkes[1]

[1]Harvard University [2]University of Oxford [3]Boston College [4]IOG

In Many Real-World Markets:

People typically have a fuzzy understanding of their preferences.

Eliciting preferences can incur high cost (cognitive or otherwise).

Examples

- Customers in a supermarket.
- A family deciding what to order for dinner on DoorDash.

The **preference elicitation** literature focuses on eliciting just enough information to find market equilibria. Relies on mathematically precise **valuation** or **demand queries**, and requires polynomially many queries.

Can LLM-powered proxies elicit preferences more efficiently and naturally?

Proxy Auction Model

- An **auctioneer** sells **multiple goods** G.
- Multiple **people** each wish to acquire a **bundle** (subset) of these goods.
- People **communicate** with the auctioneer only through a **proxy**.

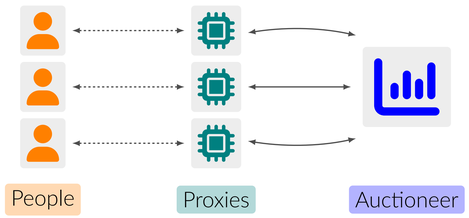

Preferences (XOR bids)

A XOR bid $\theta = (B, v)$ consists of a set of **atomic bundles** B and values $v(b)$ for these bundles $b \in B$.

The bid's induced valuation is $v^\theta(b) \coloneqq \max_{b' \in B | b' \subseteq b} v(b')$, for all bundles $b \subseteq G$.

Fact: XOR bids can represent any valuation satisfying free disposal.

Pricing and Demand

The auctioneer sets **personal price** $\varphi^i(b)$ for each bundle $b \subseteq G$ and each person i.

XOR bid θ **demands** bundle b if it maximizes her quasi-linear utility $v^\theta(b) - \varphi(b)$.

Definition: Prices $\varphi^1, \ldots, \varphi^m$ and **allocation bundles** $b^1, \ldots, m$ are **competitive equilibrium (CE)** for XOR bids $\theta^1, \ldots, \theta^m$, if each bid demands θ^i at prices φ^i.

Fact: CE exist for any XOR bids, and maximize **social welfare** $\sum_i v^{\theta^i}(b^i)$.

The Competitive Equilibrium Combinatorial Auction

The CECA (Lahaie and Parkes, 2004) incrementally elicits information from bidders, via proxies, until it reaches CE.

Step 1. Retrieve a XOR bid θ^i from each proxy $\omega(p_i)$.
Step 2. Compute a CE $(\varphi^{\theta^i}, b^i)$ for these bids using integer linear programming.
Step 3. Ask every proxy whether their person is satisfied with the allocation b^i at prices φ^{θ^i}. If so, **return** the CE.
Step 4. Tell the proxy of every unsatisfied person to update their XOR bid, and go to Step 1.

LLM Proxy Designs

- Each proxy maintains a candidate XOR bid representing her current belief of person's preferences.
- Proxies gather information from their person through **value**, **demand** and **natural language queries**.
- Our LLM proxies maintain **chat history** to guide inference.

Drop-In Proxies: Uses an LLM to decide whether the person is satisfied, or whether to make a further valuation or demand query. Variants also **infer** values of un-queried bundles in regular intervals.

Plus-Questions Proxy: Initially, asks the person questions in a natural language conversation. Then, functions similar to Drop-In Proxies, while additionally using the conversation to guide the inference.

Hybrid Proxy: Combine the Plus-Questions Proxy with the traditional proxy design for "best of both worlds".

Experiments

We explore our LLM proxies experimentally using an **LLM simulation pipeline** to study **auction efficiency** and **long-run performance**. Three different scenarios:

Electronics	Preserves	Transportation
A collection of like-new electronics donated to a local library, ranging from AirPods to iPads and accessories.	Gourmet food items are available, including exotic jams and chutneys with unique flavors, perfect for cooking or gourmet meals.	A variety of electric scooters and bikes are up for auction, catering to urban commuters and fitness enthusiasts.
Apple AirPods (2nd Gen)	Organic Strawberry Jam	Electric Scooter S2 Pro (2024)
Apple AirPods Max	Wild Blueberry Preserves	E-Scooter Elettrica (2023)
Apple iPad (9th Gen)	Apricot & Lavender Conserve	Voltron SP03 E-Scooter (2024)
Apple iPad Air (M2)	Sugar-Free Raspberry Spread	Troik Verve+ 2 (2023)
Apple Pencil (2nd Gen)	Spiced Plum Chutney	Titan Escape 3 (2023)
Apple Pencil Pro	Mango & Passionfruit Jam	Schwin Suburban (2021)

LLM Simulation Pipeline

The LLM pipeline simulates people for our experiments.

- The LLMs simulating people answers natural language questions as well as value and demand queries.
- Simulated preferences are developed by constructing and refining a **seed description** in **natural language**.
- The pipeline then systematically constructs a XOR bid from the seed. This bid grounds the answers to natural language, value and demand queries.

Results

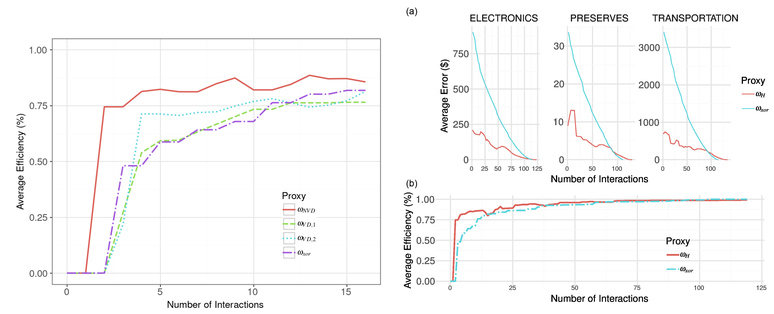

Result 1 (Auction Efficiency): Efficient outcomes can be reached with less communication between the proxy and the person using LLM-based proxies.

Result 2 (Long-Run Performance): The hybrid proxy consistently achieves a better approximation of a person's valuation than the traditional proxy does, in 50 iterations. Moreover, the auction with this proxy converges to maximum social welfare.

Result 3 (Robustness of LLM simulations): The LLM simulation pipeline is robust. The valuations generated by models from different providers are broadly consistent, given the same seeds.

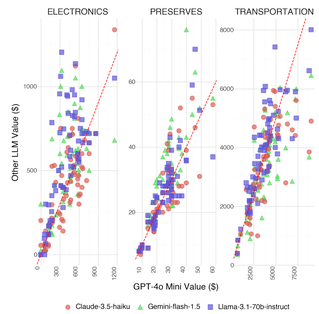

Outlook

Our LLM-powered proxies **reduce communication complexity** through natural language interactions and inference, leading to more **human-centric** and **efficient** auction mechanisms.

Future directions

- exploring practicality and efficacy of LLM proxies in lab experiments with human participants;
- Further work on scalability of LLM proxies to larger auction environments.

References

[LP04] Sebastien M Lahaie and David C Parkes, *Applying learning algorithms to preference elicitation*, Proceedings of the 5th ACM Conference on Electronic Commerce, 2004, pp. 180–188.