

Querying with Conflicts of Interest

Nischal Aryal
Oregon State University
aryaln@oregonstate.edu

Arash Termehchy
Oregon State University
termehca@oregonstate.edu

Marianne Winslett
University of Illinois
winslett@illinois.edu

ABSTRACT

Conflicts of interest often arise between data sources and their users regarding how the users' information needs should be interpreted by the data source. For example, an online product search might be biased towards presenting certain products higher than in its list of results to improve its revenue, which may not follow the user's desired ranking expressed in their query. The research community has proposed schemes for data systems to implement to ensure unbiased results. However, data systems and services usually have little or no incentive to implement these measures, e.g., these biases often increase their profits. In this paper, we propose a novel formal framework for querying in settings where the data source has incentives to return biased answers intentionally due to the conflict of interest between the user and the data source. We propose efficient algorithms to detect whether it is possible for users to extract relevant information from biased data sources. We propose methods to detect biased information in the results of a query efficiently. We also propose algorithms to reformulate input queries to increase the amount of relevant information in the returned results over biased data sources. Using experiments on real-world datasets, we show that our algorithms are efficient and return relevant information over large data.

PVLDB Reference Format:

Nischal Aryal, Arash Termehchy, and Marianne Winslett. Querying with Conflicts of Interest. PVLDB, 14(1): XXX-XXX, 2020.
doi:XX.XX/XXX.XX

PVLDB Artifact Availability:

The source code, data, and/or other artifacts have been made available at URL_TO_YOUR_ARTIFACTS.

1 INTRODUCTION

Due to financial, social, and political factors, the incentives of data source owners (data sources for short) and their users are often not perfectly aligned. For example, data sources that provide information about products may benefit financially from convincing users to buy certain products. In fact, Amazon, Google, and Apple have been accused of returning results that are biased in this manner [12, 22, 24, 26]. As an example, Apple allegedly reworked its ranking algorithm so that Apple Music was ranked first and Spotify 23rd in the App Store when a user searched for music apps [26]. Similarly, a search for podcast apps returned Apple Podcasts in the first place, followed by 12 irrelevant Apple apps, and then finally

podcast apps from other publishers. A search for "TV" in 2018-9 would return Apple apps at the top, and Netflix ranked in the 100s [26]. As experts commented at the time: "I find it hard to believe that organically there are certain Apple apps that rank better than higher-reviewed, more downloaded competitors," and "there's just a ton to be gained commercially by dominating" search results [26].

This conflict of interest is inherent at some data sources, such as shopping websites, which charge sellers a 'referral fee' of considerable fraction of the price of items sold on its site [2], so that it profits more when users buy higher-priced versions of a given product. Cost-conscious users who re-sort the results by price may have to wade through many undesirable results before reaching the first good match.

Example 1.1. On the day of this writing, a search for 'headset' on a well-known shopping website (**SW** for short) returns headphones priced \$24.99 to \$39.95 in its top 'Results' section, and all items on the remainder of the page are indeed headphones. When re-sorted by price, the top hit on the first page is a charging cord, followed by earbuds, microphones, cables, and one \$1.98 headset. Only on the seventh page of results do headphones start to outnumber other accessories.

Conflicts of interest may also arise from non-financial considerations. To maximize user engagement, some social media sites may promote extreme content, which is not necessarily in anyone else's best interest. More generally, data sources may return information that is biased toward or against controversial topics and causes, political groups, or demographic groups, e.g., not returning or giving a low rank to information related to these [36, 37]. Perhaps the best known recent example is the removal of all topics related to climate change from US government websites [38]. Data sources that belong to foreign adversaries can also promote content that serves the objectives of the data source owner.

In the aforementioned settings, data sources interpret the information in user queries to take actions that do not exactly meet users' preferences and objectives and, in fact, might largely deviate from what users deem as desirable. Current proposals to combat undesirable data source actions usually require the data source to follow and implement certain protocols [32]. For example, researchers have proposed methods to remove bias and ensure fairness for data sources to implement [21, 33-37]. However, data sources often do not have any incentive to implement these proposals because these are usually against their objectives and business models. Some might also argue that these restrictions limit the freedom to disseminate information or conduct business and may not be ethical or have long-term adverse impacts on our society.

In such a situation, the data source and users do share some interests: the data source would like to satisfy the information needs behind users' queries well enough to keep them engaged. Users can leverage this partial common interest by modifying their

This work is licensed under the Creative Commons BY-NC-ND 4.0 International License. Visit <https://creativecommons.org/licenses/by-nc-nd/4.0/> to view a copy of this license. For any use beyond those covered by this license, obtain permission by emailing info@vldb.org. Copyright is held by the owner/author(s). Publication rights licensed to the VLDB Endowment.
Proceedings of the VLDB Endowment, Vol. 14, No. 1 ISSN 2150-8097.
doi:XX.XX/XXX.XX

queries so that the data source interpretations of the submitted queries are those desired by users.

Example 1.2. Consider again the user in Example 1.1 who looks for affordable headphones on SW. If the results are biased towards relatively more expensive products, she can add a price cap to her query. For example, when we reissued the SW headset query with a \$19 maximum price, the *second* result returned cost \$21, but the others on the first page cost less than \$19. Thus, SW does not take price caps literally, but they do influence the results.

But two can play at this game. The data source may be aware of the fact that users know about its undesirable actions, e.g., users have observed undesirable biases in their previous experiences of using the data source. Thus, it may reason that users aim at guiding its interpretations and results towards the users’ desired results by modifying their original queries. Since it is not in the data source’s best interest to accept the modified query as precisely the user’s true information need, the data source may try to recover the original intent behind the modified query. For instance, it might also use its prior belief about users’ intents, e.g., gained from external sources or previous interactions with users, to interpret the true intent behind the modified query. Consequently, users may reason about the data source reasoning process and try to modify their queries more drastically to influence the data source interpretation. The user and data source reason and make decisions based on possible reasoning and strategies of each other iteratively to find the best method to formulate and interpret queries.

Example 1.3. In Example 1.2, given that the user has added a degree of bias in the modified query, the data source may reason that the user has overstated their preference for headphones of a particular price, just to offset the data source’s bias. Hence, it might avoid placing such headphones in the top positions, as we saw when SW still ranked a \$21 headset second even after we limited the price to \$19 in Example 1.1. The user might counteract this by further reducing her submitted maximum price for a headset.

In some of the aforementioned settings, a brute-force approach for the user to get her desired information without allowing the data source to return undesirable results is to submit a query that retrieves most or all of the data and post-process that to extract her desired information. In extreme, in Example 1.3, the user who searches the data source might submit a query that returns all products. This is not usually practical, because of the enormous amount of data managed by these data sources. Furthermore, data sources might not return such a large dataset as it is against their objective, e.g., they might ask users to resubmit a query that returns a relatively small subset of the data. In addition, data sources usually limit the rate of retrieved information, making it difficult for users to download a very large dataset using multiple queries. Moreover, users often lack the computational resources, time, and skills to store and post-process a large dataset.

There are four important problems in this monopolistic setting. First, we would like to know under what conditions the aforementioned recursive reasoning of the user and the data source will converge to a stable state in which the user finds a query that gives her useful information. The user might be able to leverage the fact

that the data sources would like to keep the user engaged, i.e., partial common interest, to convince the data source to return her desired results to some degree. Second, we would like to inform users about the untrustworthy information in the results returned by the data source, e.g., whether the returned top- k results are in fact the top- k results for the query on the entire data. It is challenging as the data source may not show all relevant information to users so they can check the validity of the information. Third, given a user query, we would like to know how much reliable information the user can get from the data source. Finally, given a query, we would like to find another query that can convince the data source to return more relevant and reliable information than the original query.

In this paper, we propose a framework that models strategic information sharing between a user and a data source when there are conflicts of interest between them. The framework models the user and the data source as agents with different objectives and uses the current body of research in game theory on how agents with different goals establish informative communication [8, 9]. Using this framework, we propose efficient algorithms that solve the aforementioned problems efficiently. Our contributions are as follows:

- We propose a formal framework that models the communication of the user and the data source with conflicts of interest using queries that return ranked lists of results (Section 2). Our framework defines the stable states (fixed points) in the recursive reasoning of the users, e.g., client or middleware systems, and data source with conflict of interest for answering a query.
- If the conflict of interest between the user and the data source is too high, the data source may ignore the user query and will return its own preferred ranking regardless of the preference expressed in the input query. We propose efficient algorithms that check whether the communication has stable states in which the input query can influence the decision of the data source about returning the answers (Section 3).
- We propose an efficient algorithm that detects untrustworthy results in the answer to a query (Section 4).
- We prove that the problem of finding the query that returns the maximum amount of information for an original user intent is NP-hard. Hence, we propose efficient algorithms to solve this problem for some general classes of utility functions (Section 5).
- We conduct extensive empirical studies using a wide variety of real-world datasets. Our empirical studies indicate that our algorithms scale to large data (Section 6).

2 FRAMEWORK

2.1 Data Model & Query Language

A schema \mathcal{R} is a set of relation symbols R_i , $1 \leq i \leq n$, where each R_i has a fixed arity a_i . Each relation symbol R_i is associated with a set of attribute symbols $(A_{i,j}, 1 \leq j \leq a_i)$. To simplify our exposition, we define the domain of possible values for each attribute $\mathbf{dom}(A_{i,j})$ in each relation as a subset of or equal to a single set \mathbf{dom} . Each tuple of the relation symbol R_i is then an element of the Cartesian product $\prod_{j=1}^{a_i} \mathbf{dom}(A_{i,j})$. A relation of R_i is a finite set of such tuples. An instance of \mathcal{R} is a set of relations for its relation symbols.

Let \mathcal{Q} denote a query language. An a -ary query $q \in \mathcal{Q}$, $a > 0$, in schema \mathcal{R} is a function that maps each instance I of \mathcal{R} to a totally ordered list where each member of the list belongs to dom^2 . We denote the result of the query q over I by $q(I)$. Let $\text{elm}(l)$ denote the set of tuples in a totally ordered list l . For $e, e' \in \text{elm}(q(I))$, we use $e \prec e'$ ($e \sim e'$) to show that e precedes e' (e and e' share the same position) in $q(I)$. For brevity, by abuse of notation, we show $e \in \text{elm}(q(I))$ as $e \in q(I)$. A pairwise order $e \prec e'$ is in $q(I)$ if the order is present in the result of $q(I)$. For brevity, we show $e \prec e'$ is in $q(I)$ as $e \prec e' \in q(I)$.

Users can express queries in our framework using query languages, such as SQL, that return a set or a ranked list of results, e.g., using ranking constructs such as the *Order By Case* clause [13, 14, 14, 25, 27, 31]. They can also use proposed constructs to express preference relationships in query languages [14, 15, 23, 29].

Table 1: Headphones

id	model	brand	rating	price	best seller
e_1	Tune 510BT	JBL	4.5	25.20	Y
e_2	Q20i	Soundcore	4.6	39.99	N
e_3	A10 Pro	Sisism	4.8	21.99	N
e_4	Sports	Haoyuyan	3.9	299.99	N

2.2 Intents, Queries, and Results

User’s Intents and Queries. We call the user’s true information need, which is expressed in \mathcal{Q} , their *intent*. The data source does *not* know the user’s intent. The user picks their intents from a prior distribution π , which is known to the data source. Given intent τ , the user submits the query $q \in \mathcal{Q}$ to the data source. The query q may *not* be equivalent to τ , e.g., to hide the user’s intent τ to offset the data source bias in its returned results for q .

Example 2.1. Consider a user who queries SW to search for a headphone. A user’s intent τ might be to view headphones ordered by customer rating: `SELECT DISTINCT brand, model, rating, price FROM Products WHERE type = ‘headphones’ ORDER BY rating DESC`. To express this in SW’s real-world interface, the user would first type ‘headphones’ into the search bar, then ask for them to be sorted by rating ($\mathbf{q}_{\text{rating}}$). As of this writing, when we submit $\mathbf{q}_{\text{rating}}$ to SW, the first page of interpretation β results includes 16 models representing a variety of brands and price points, as partially shown in Table 1. The query panel on the results page says that it is showing only products with ratings of 4 and up, though the Haoyuyan brand is rated 3.9.

Data Source Interpretation. The data source interprets the query $q \in \mathcal{Q}$ submitted by the user to obtain its own desired query β . If there is a conflict of interest between the user and the data source, the interpreted query β may *not* be equivalent to q , e.g., if the data source is biased towards ranking certain tuples high.

2.3 Strategies

User strategy. The user strategy $P^r(q|\tau)$ is a (stochastic) mapping from \mathcal{Q} to \mathcal{Q} . Given the intent τ , P^r returns queries that the user submits to the data source.

Example 2.2. Continuing with Example 2.1, to offset SW’s bias, we submitted a different query $\mathbf{q}'_{\text{rating}}$ that asked for non-major

brands (non-JBL) sorted by rating. This resulted in 4.8 and 4.9-rated products from Skullcandy being listed 12th and 14th on the first page, even though all the other products on that page were only rated 4.7. Similarly, the first item on the second page was a 4.8 Skullcandy headset that was not listed on the previous page. However, when we searched for ‘headphones’ with SW’s default search order, no Skullcandy products appeared in the first 10 pages of results. This inspires us to imagine a situation where a data source is biased against products from a brand B and systematically demotes them in its interpretation of the `ORDER BY` clause, e.g., `ORDER BY CASE WHEN brand = ‘B’ THEN rating - .2 ELSE rating END DESC`. In such a situation, if the user knows about this bias and is quite fond of products from B , a better submitted query for her intent is $\mathbf{q}''_{\text{rating}} `ORDER BY CASE WHEN brand = ‘B’ THEN .2 + rating ELSE rating END DESC`. This lists all B products first, in rating order, followed by the other brands’ models in descending order of rating. In the SW interface, the user could approximate this by first using the search bar to query for “headphones”, sorting by rating, and then checking the box to view Skullcandy headphones first. When we did this, the first page of results were all Skullcandy products, though they were not actually presented in rating order.$

Data source strategy. The *data source strategy* $P^s(\beta|q)$ is a (stochastic) mapping from \mathcal{Q} to \mathcal{Q} . The data source uses P^s to interpret the intent behind the submitted query q . For example, if the data source is biased, given q , it computes a biased query β and returns the result of β on its database instance.

2.4 Interaction Setting

Utility Functions. The *utility functions* of the user and the data source measure the desirability of the interpretation of the intent of the user for the user and the data source, respectively. We denote the utility functions of the user and the data source by $U^r(\tau, \beta)$ and $U^s(\tau, \beta)$, respectively, which map the query pairs in $\mathcal{Q} \times \mathcal{Q}$ to real values. If there is some conflict of interest between the user and the data source, U^r and U^s take their maximum values for different values of β for the same intent τ : U^r takes its maximum when β and τ are equivalent, and U^s will be at its maximum for some β' that is not equivalent to τ and returns results that satisfy the biases of the data source. The exact instantiation of U^r and U^s varies depending on the application and the domain of the data.

Example 2.3. Consider the intent τ from Example 2.1 but *no access to the underlying tuples in the Headphone instance*. Assume only two possible data source interpretations: β_1 : `ORDER BY CASE WHEN brand=‘Skullcandy’ THEN rating - .2 ELSE rating END` and β_2 : `ORDER BY rating DESC`. Interpretation β_2 better matches the user’s intent than β_1 . To output a real-valued score, suppose the utility function measures absolute errors in ranking positions (Spearman Footrule distance [30]). Concretely, assume access to the domain of `rating` based on a known schema \mathcal{R} , where the true ranks under τ of the data source’s four headphones are [1, 2, 3, 4] and under β_1 are [4, 3, 2, 1]. Then $U^r(\tau, \beta_1) = -(|1 - 4| + |2 - 3| + |3 - 2| + |4 - 1|) = -8$. Similarly, $U^r(\tau, \beta_2) = -2$, so interpretation β_2 provides higher user utility.

Interaction Setting. An *interaction setting* (interaction for short) is the tuple $(\mathcal{R}, \tau, U^r, U^s)$ where \mathcal{R} is the schema of the database,

τ is the user intent and U^r and U^s are the utility functions of the user and the data source, respectively. U^r and U^s are the same for all user intents on \mathcal{R} .

Common knowledge. We assume that the user knows the schema, but *not* the database instance maintained in the data source. Both data source and user know all utility functions. Users may learn the utility function of the data source by querying the data source and using their past interactions with the data source.

2.5 Equilibria of Interaction

Stable states. As explained in Section 1, the user and the data source use their available information to reason about the strategies of the other party to maximize their utility function. We would like to identify whether this reasoning will converge to any eventual stable state in which the user and data source settle on a pair of fixed strategies where none can increase their utility by deviating from their strategies. Knowing these states will enable us to analyze the strategies selected for the user and the data source and design algorithms to find desired strategies for users in these settings.

Example 2.4. Consider the following scenario. The Headphones instance I has only two tuples e_1 (JBL, rating 4.5, \$25.20) and e_2 (Soundcore, rating 4.6, \$39.99). The user has one of two intents: price-conscious $\tau(I)$, $e_1 \prec e_2$ or rating-conscious $\tau'(I)$: $e_2 \prec e_1$. The data source chooses between two interpretations, $\beta(I)$: $e_2 \prec e_1$ and $\beta'(I)$: $e_1 \prec e_2$. Suppose the user gets $U^r = 2$ if the interpretation matches her intent and $U^r = 0$ otherwise. The source earns an extra commission $c > 0$ when the expensive headphone e_2 is ranked first, and suffers a loss $L > 0$ when the user is dissatisfied with the results. The source utility is commission minus dissatisfaction loss. This yields the payoff matrix:

	$\beta (e_2 \prec e_1)$	$\beta' (e_1 \prec e_2)$
$\tau (e_1 \prec e_2)$	$(0, c - L)$	$(2, 0)$
$\tau' (e_2 \prec e_1)$	$(2, c)$	$(0, -L)$

where each cell shows (U^r, U^s) . If $c < L$, then the source prefers β' under τ (since $0 > c - L$) and prefers β under τ' (since $c > -L$). Hence, there is a stable state. The user strategy is for intent τ to submit a price query (e.g., ORDER BY price ASC) and to submit a rating query (e.g., ORDER BY rating DESC) for τ' . Based on the utility function, the source will follow these queries by choosing β' after the former and β after the latter. However, if instead the commission is very large $c \geq L$, then the source weakly prefers β even under τ (since $c - L \geq 0$), and strictly prefers β under τ' . Thus, the source will return β ; this may prompt a new cycle of reasoning.

Bayesian equilibrium. Since the data source does not know the user's intent, we follow the approach used in the game theory literature and use the concept of *Bayesian equilibrium* to characterize the stable state of the user and data source reasoning [9]. The *belief of the data source* $\Phi(\tau|q)$ is a stochastic mapping from \mathcal{Q} to \mathcal{Q} . The belief of the data source maps every query submitted by the user q to a probability distribution over the set of possible intents. The *expected utility of the user* for the query q is

$$\mathbb{E}[U^r(\tau, \beta) | q] = \sum_{\beta \in \mathcal{Q}} U^r(\tau, \beta) P^s(\beta | q).$$

The *expected utility of the data source* for the query q is

$$\mathbb{E}[U^s(\tau, \beta') | q] = \sum_{\tau \in \mathcal{Q}} U^s(\tau, \beta') \Phi(\tau | q).$$

where $\Phi(\tau | q) = \frac{P^r(q|\tau)\pi(\tau)}{\sum_{y \in \mathcal{Q}} P^r(q|y)\pi(y)}$.

Definition 2.5. An equilibrium for interaction $(\mathcal{R}, \tau, U^r, U^s)$ is a pair of user and data source strategies $(P^r(q|\tau), P^s(\beta|q))$ such that

- if $P^r(q|\tau) > 0$, $q = \arg \max_{q'} \mathbb{E}[U^r(\tau, \beta) | q']$.
- if $P^s(\beta|q) > 0$, $\beta = \arg \max_{\beta'} \mathbb{E}[U^s(\tau, \beta') | q]$.

3 POSSIBILITY OF INFLUENCE

A data source may be so biased towards returning certain order of tuples for a given query such that no user's strategy can modify its decision, i.e., no matter what the input query expressed the user's desired order between tuples, the data source will return tuples based on its own preferred order. For example, if a shopping website is significantly biased toward returning products with higher prices, no matter what the user's submitted query, e.g., returning products based on their quality ratings, the website will return a fixed order of products based on their prices. In this section, we investigate these settings and propose efficient algorithms to detect them.

3.1 Influential Interactions

An equilibrium (P^r, P^s) is **influential** in the interaction $(\mathcal{R}, \tau, U^r, U^s)$ if and only if there are at least two queries $q, q' \in \mathcal{Q}$, $P^r(q|\tau) > 0$ and $P^r(q'|\tau) > 0$, such that $P^s(\beta|q) \neq P^s(\beta|q')$. Otherwise, the equilibrium is non-influential. If an equilibrium is non-influential, the data source selects an interpretation that maximizes its expected utility function based on prior π over the user's intent in the equilibrium. That is, the data source finds it optimal to return answers purely according to its prior belief.

From the definition of equilibrium, it follows that every interaction has a non-influential equilibrium in which the user will send the same query for all their intents. An interaction is **influential** if it has at least one influential equilibrium and is non-influential otherwise. For instance, in Example 2.4 the stable state is an influential equilibrium, whereas the other is an uninfluential equilibrium. Given an interaction setting, before extracting reliable information from data source results or finding the desired strategy for the user, we should check whether the interaction is influential.

Problem 1. Given the interaction setting $(\mathcal{R}, \tau, U^r, U^s)$, determine whether the interaction is influential.

Example 2.4 illustrates a setting where the objectives of the user and the data are extremely opposed, therefore, no influential interaction is possible. The following result formalizes this intuition. To simplify our exposition, we assume that intents, submitted queries, and interpretation on schema \mathcal{R} do not return any duplicate tuple over any instance of \mathcal{R} . Our results extend to the cases where they return duplicate tuples. Two intents on schema \mathcal{R} are **set-equivalent** if they return the same set of answers over all instances of \mathcal{R} .

THEOREM 3.1. *The interaction $(\mathcal{R}, \tau, U^r, U^s)$ is influential if and only if there are set-equivalent intents $\tau \neq \tau'$ and interpretations*

$\beta \neq \beta'$ where the following conditions hold for both user and data source utility functions $U^t \in \{U^r, U^s\}$:

$$U^t(\tau, \beta) \geq U^t(\tau, \beta') \quad \text{and} \quad U^t(\tau', \beta') \geq U^t(\tau', \beta).$$

or

$$U^t(\tau, \beta') \geq U^t(\tau, \beta) \quad \text{and} \quad U^t(\tau', \beta) \geq U^t(\tau', \beta').$$

The proofs of our theoretical results are in Section 8. Theorem 3.1 provides necessary and sufficient conditions on user and data source’s utility functions that enable influential interaction. We can use this result to check whether the interaction is influential by checking the intents and interpretations in \mathcal{Q} . However, it will be too inefficient for expressive intent and query languages.

3.2 Detecting Influential Interactions Efficiently

In this section, we investigate settings for which we can efficiently detect whether an influential equilibrium exists. We first place some mild and natural conditions on utility functions. Given these conditions, we provide efficient algorithms to detect whether influential equilibria are possible.

Bias Function. The utility of the data source may vary between tuples. Consider a shopping website that sells products from various brands, including the website’s owner. The data source may prefer tuples that contain information about its own products due to higher profitability. We formalize this heterogeneity using *bias functions*. Given relation schema \mathcal{R} , the bias function $b(e)$ maps each tuple e into the domain of \mathcal{R} to real values (\mathbb{R}). If the data source is not biased to tuple e , we have $b(e) = 0$. Since the data source is biased towards (or against) some tuples, $b(e) \neq 0$ for some tuples e . The exact instantiation of the bias function varies according to the application. For instance, one may define a bias function as a mapping of one or multiple attributes of tuples, e.g., brand.

Additive Utility. Given an intent τ and a data source interpretation β over an instance I of schema \mathcal{R} , we assume that the utility function for the data source is defined as:

$$U^s(\tau, \beta) = \sum_{e \in \tau(I)} u^s(r(e, \tau(I)), r(e, \beta(I)), b(e)) \quad (1)$$

Here, $r(e, q(I))$ is a function that returns the position of a tuple $e \in q(I)$ and $u^s(\cdot, \cdot, \cdot)$ is a function that maps pairs of positions and a bias value to a real value. We make the same additivity assumption for the user’s utility where the bias function maps every tuple to zero, as the user is not biased to any particular tuple. Due to the biases of the data source against or toward some tuples, there might be tuples in $\beta(I)$ that do not belong to $\tau(I)$. If the data source does not return some tuples in $\tau(I)$, we set $r(e, \beta(I))$ to a fixed value greater than the number of tuples in $\tau(I)$. For the rest of the paper, we assume that all utility functions are in the form of Equation 1. Since users usually do not know the exact instance I , they cannot accurately compute the value of U^s for I . However, as we will show, they can reason about the utilities given the schema of I .

Example 3.2. Consider the interaction from Example 2.3 with concrete utility functions. The user gets maximum utility when the result perfectly matches their intent. We model this by function:

$$u^r(r(e, \tau(I)), r(e, \beta(I))) = -(r(e, \beta(I)) - r(e, \tau(I)))^2$$

This function is maximized at 0 (when interpretation rank equals intent rank), explaining why the user prefers the interpretation β_2 . To explain the data source’s preference for the interpretation β_1 , we define bias function b based on the brand attribute:

$$b(e) = \begin{cases} 2 & \text{if } e.\text{brand} = \text{'Skullcandy'} \\ 0 & \text{otherwise} \end{cases}$$

Therefore, the source’s utility for a Skullcandy tuple e becomes $u^s(\cdot) = -\left(r(e, \beta(I)) - (r(e, \tau(I)) + 2)\right)^2$. Thus the source’s ideal output rank for Skullcandy tuples is shifted two positions lower (numerically larger) than the intent rank, so it prefers β_1 over β_2 .

Partial Common Interest. Although the data source may be biased towards returning some tuples more than other tuples, the results returned by the data source should not be completely independent of the user intent. In particular, the returned results should have connections to the user intent when the biases are equal. More precisely, if the position of the tuple e in intent τ is higher than its position in intent τ' , the data source should place e in the answer to τ in the same or higher position than the position of e in τ' .

Definition 3.3. Given intent τ and interpretations β and β' on an instance I of the schema \mathcal{R} and the data source utility function U^s , U^s is **supermodular** if for all tuples e where $r(\beta(I), e) < r(\beta'(I), e)$, we have $u^s(r(\tau(I), e), r(\beta(I), e), b(e)) - u^s(r(\tau(I), e), r(\beta'(I), e), b(e)))$ is non-increasing in $r(\tau, I, e)$.

In the definition of supermodularity, we compare the positions of a tuple in different intents, as modifying the position of a tuple in an intent creates a new intent. A data source with a supermodular utility function may *not* follow the user’s preferred relative ranking between different tuples in an intent: it may still be more biased to returning some tuples in positions higher than other tuples. Intuitively, the user prefers results that match their preferences the most: the more the user prefers the tuple e , the higher position the user would like to see it in the returned result. Hence, it is natural to assume that U^r is also supermodular, defined similarly to Definition 3.3. For the rest of the paper, we assume utility functions of the data source and user are supermodular.

Example 3.4. Consider the supermodular utility functions in Example 3.2. For each tuple, the data source’s optimal position in interpretation is $r(e, \beta(I)) = r(e, \tau(I)) + 2$, while the user prefers interpretations with $r(e, \beta(I)) = r(e, \tau(I))$. Therefore, there is some commonality of interest for each tuple’s position since both agents’ optimal interpretation positions are increasing in $r(e, \tau(I))$, but also some conflict since the data source’s optimal interpretation position is always +2 lower.

Symmetric Bias. The data source may be equally biased towards or against the set of tuples of an intent. For example, the intent may ask for a ranked list of products produced by a certain provider based on their quality where the data source is biased against products made by that provider. In this case, the data source may try to discourage the user by returning and placing products from other providers, which do not belong to the answers of the intent over the current database instance, in high positions of the returned result. We show that in these cases the interaction is influential.

PROPOSITION 3.5. Given the interaction setting $(\mathcal{R}, \tau, U^r, U^s)$, if $b(e)$ is equal for all tuples e , the interaction has influential equilibria.

THEOREM 3.6. Given interaction setting $(\mathcal{R}, \tau, U^r, U^s)$, let U^s be

$$U^s(\tau, \beta) = - \sum_{e \in \tau(I)} L(r(e, \tau(I)) - (r(e, \beta(I)) + b(e)))$$

over instance I of \mathcal{R} , where $L : \mathbb{R} \rightarrow \mathbb{R}_+$ is a convex function. Assume that U^r has the same structure. If there is an interpretation β' such that for every e , we have $\beta'(I) = \arg \min_{\beta(I)} L(r(e, \tau(I)) - (r(e, \beta(I)) + b(e)))$, the interaction does not have any influential equilibrium.

Intuitively, Theorem 3.6 shows that if the data source's bias is sufficiently large such that there is an interpretation that provides the maximum utility for the data source regardless of the user's desired intent, no influential equilibrium is possible, i.e., users cannot convince the data source to return other results. If L is differentiable, we can check the condition in Theorem 3.6 efficiently.

The subsequent result follows directly from Theorem 3.6 and recognizes some non-influential interactions efficiently for a large group of utility functions, e.g., quadratic utility functions where $L(\cdot)$ is $(r(e, \tau(I)) - (r(e, \beta(I)) + b(e)))^2$. It is particularly useful when the user is interested in the top- k returned results.

COROLLARY 3.7. Given the interaction $(\mathcal{R}, \tau, U^r, U^s)$, let U^s be

$$U^s(\tau, \beta) = - \sum_{e \in \tau(I)} L(r(e, \tau(I)) - (r(e, \beta(I)) + b(e)))$$

for instance I of \mathcal{R} , where $L(v) : \mathbb{R} \rightarrow \mathbb{R}_+$ is convex, non-decreasing in $|v|$, and minimized at $v = 0$. Assume that U^r has the same structure. If for every $b(e)$, $|b(e)| \geq k - \frac{3}{2}$, the interaction does not have an influential equilibrium where k is the number of tuples in $\tau(I)$.

Given the range of $b(\cdot)$, we can check the condition of Corollary 3.7 in constant time.

4 DETECTING TRUSTWORTHY ANSWERS

A data source may be partially biased towards returning a certain order of tuples for a given query. For instance, a shopping website may return some products from its own owners, some of which may be irrelevant to the user's query, while still returning some relevant products to keep the user engaged. In this section, we investigate these settings and present algorithms to detect reliable information from data sources' interpretation of users' queries.

Trustworthy Information. We first define the most basic unit of information in our setting.

Definition 4.1. Given an intent τ and an interpretation β over an instance I of schema \mathcal{R} , a returned tuple $e \in \beta(I) \cap \tau(I)$ is **untrustworthy** if there exists a tuple $e' \in \tau(I) \setminus \{e\}$ such that $r(e', \tau(I)) < r(e, \tau(I))$ and either (i) $e' \notin \beta(I)$, or (ii) $r(e', \beta(I)) > r(e, \beta(I))$. Otherwise, e is **trustworthy**.

Intuitively, a tuple is untrustworthy if the data source's interpretation misranks the position of a tuple in the intent. As discussed in Section 3.2, the data source may misrank tuples due to bias and may also return an incomplete result, e.g., only the top- k tuples that maximize its utility. Specifically, a returned tuple e is untrustworthy if it may have displaced some tuple e' that the intent ranks above it, i.e., $r(e', \tau(I)) < r(e, \tau(I))$. This can happen either because the data

source returns e' but ranks it below e in $\beta(I)$ (a promotion of e), or because the data source omits e' from the returned list altogether (a demotion by omission)

Problem 2. Given the interaction setting $(\mathcal{R}, \tau, U^r, U^s)$ and an interpretation result $\beta(I)$, detect all untrustworthy tuples in $\beta(I)$.

Possible Intent Ranks. As discussed in Section 2.5, the data source does not know the user's intent τ , so it maintains a prior over the user's intent. Specifically, it maintains a prior over the ranks of tuples in $\tau(I)$. Concretely, for each tuple $e \in \tau(I)$, we assume the rank $r(e, \tau(I))$ is a discrete random variable supported on $[z] = \{1, 2, \dots, z\}$ and is uniformly distributed on $[z]$ under the prior.

Indifference Boundary. As discussed in Section 3.2, the supermodularity of the data source's utility function provides some incentive to place tuples that are in higher positions in the intent in relatively higher positions in the interpretation. Specifically, an interpretation that misranks the relative order of tuples in the intent incurs a larger utility loss for the data source when the intent ranks of these tuples are farther apart. However, the bias of the data source may nevertheless lead to interpretations that misrank these tuple that provide higher utility. Especially when two tuples are sufficiently close in the intent, as the data source's utility loss from misranking them is smaller and may be offset by the bias. To characterize when bias provides sufficient incentive for the data source to misrank tuples, we describe how the data source and user compare their utility from alternative interpretations.

Definition 4.2. Given intent τ and interpretations β, β' on instance I of the schema \mathcal{R} , the user is **indifferent** between β and β' if $U^r(\tau, \beta) = U^r(\tau, \beta')$. Given a posterior belief Φ over intents, the data source is **indifferent** between β and β' at Φ if $\mathbb{E}_{\tau \sim \Phi}[U^s(\tau, \beta)] = \mathbb{E}_{\tau \sim \Phi}[U^s(\tau, \beta')]$.

Intuitively, if the user receives the same utility for any two distinct interpretations, the user is indifferent between them. However, when the user is not indifferent, meaning they strictly prefer one interpretation over the other, they will strategically select the query that convinces the data source to return results based on the preferred interpretation. Therefore, the condition where the user is indifferent effectively defines the decision boundary in the space of possible intents, or equivalently, possible positions of tuples in the intent, based on which the user submits different queries. The data source's indifference condition is defined analogously, except that the data source evaluates utilities under its posterior belief Φ because it does not observe the true intent (Section 2.5). Next, we provide the structure of indifference conditions for quadratic utility functions.

PROPOSITION 4.3. Given the interaction setting $(\mathcal{R}, \tau, U^r, U^s)$ suppose the user and data source have the following utility functions for each tuple e :

$$u^r(r(e, \tau(I)), r(e, \beta(I))) = -(r(e, \tau(I)) - r(e, \beta(I)))^2,$$

$$u_e^s(r(e, \tau(I)), r(e, \beta(I))) = -(r(e, \tau(I)) - (r(e, \beta(I)) + b(e)))^2.$$

over instance I of \mathcal{R} . The user is indifferent between interpretations β, β' , iff $\sum_{e \in E} \gamma_e r(e, \tau) = \alpha^r$, where the normal vector is $\gamma_e = 2(r(e, \beta(I)) -$

$r(e, \beta'(I))$). Let Φ be a belief over intents and $\Phi_e := \mathbb{E}_{\tau \sim \Phi} [r(e, \tau(I))]$. The data source is indifferent at belief Φ iff $\sum_{e \in E} \gamma_e \Phi_e = \alpha^s, \alpha^s := \alpha^r + 2 \sum_{e \in E} b(e)(r(e, \beta(I)) - r(e, \beta'(I)))$.

The proof follows by equating the utilities for β and β' , then simplifying the resulting expression. Proposition 4.3 says that the agent's indifference condition has an affine form. Specifically, a hyperplane in the space of possible intent ranks $\{r(e, \tau)\}_{e \in \tau(I)}$ for the user, and a hyperplane in posterior-mean space $\{\Phi_e\}_{e \in \tau(I)}$ for the data source, with the same normal but different intercepts due to the bias. The results extend to any additive, supermodular utility functions where the difference $u(r(e, \tau(I)), r(e, \beta(I))) - u(r(e, \tau(I)), r(e, \beta'(I)))$ is linear in $r(e, \tau(I))$.

Misrankings. Based on Definition 4.1, a returned tuple e is untrustworthy exactly when there exists e' that the intent ranks above e (i.e., $r(e', \tau(I)) < r(e, \tau(I))$) but the interpretation does not place e' above e (either e' is omitted, or it is returned below e). This is a *pairwise violation*: it is witnessed by the pair (e', e) . Thus, even if the data source performs a complex reranking (e.g a sequence of swaps), it suffices to analyze pairwise swaps as canonical local witnesses that isolate a single potentially inverted pair.

LEMMA 4.4. *In the setting of Proposition 4.3, suppose β, β' differ only on $e \neq e'$ and satisfy $r(e, \beta(I)) - r(e, \beta'(I)) = -(r(e', \beta(I)) - r(e', \beta'(I))) \neq 0$. Then the user's indifference condition depends on the difference between the tuples' ranks in the intent:*

$$r(e, \tau(I)) - r(e', \tau(I)) = \delta, \quad \text{where} \quad \delta = \frac{\alpha^s}{\gamma_e} - (b(e) - b(e')) \quad (2)$$

Intuitively, Lemma 4.4 says that when the two interpretations β and β' swap the position of the two tuples, the user's preference between the two interpretations depends only on the *relative positions of the tuples in the intent* and the data source's bias. We refer to δ from Equation 2 as the **indifference threshold**. In the following results, we use this indifference threshold to detect trustworthy information in the interaction.

PROPOSITION 4.5. *In the setting of Lemma 4.4, let $\delta_0 = \frac{\alpha^s}{\gamma_e}$ denote the indifference threshold when $b(e) = b(e')$ and let $\delta = \delta_0 - (b(e) - b(e'))$ denote the bias-shifted indifference threshold. Then the data source interpretation is untrustworthy for $e \in \beta(I)$ if there exists a tuple $e' \in \tau(I)$ such that the tuples rank difference in the intent satisfies $r(e, \tau(I)) - r(e', \tau(I)) \in (\min\{\delta_0, \delta\}, \max\{\delta_0, \delta\}]$.*

Baseline Algorithm. Proposition 4.5 gives a baseline algorithm to detect untrustworthy information in the interaction. For each returned tuple $e \in \beta(I)$ and each tuple in the result of the intent on the schema domain \mathcal{R} such that $e' \notin \beta(I)$ or $r(e', \beta(I)) > r(e, \beta(I))$, the algorithm computes the unbiased and biased indifference thresholds $\delta_0 = \frac{\alpha^s}{\gamma_e}$ and $\delta = \delta_0 - (b(e) - b(e'))$. If the relative positions of these tuples in the intent satisfy the condition in the proposition, then the $e \in \beta(I)$ is trustworthy. This algorithm is inefficient since for each tuple in the interpretation, it has to compute its positions in the intent based on the schema and scales with the size of the schema domain.

THEOREM 4.6. *Given quadratic utility functions for the user and the data source, interpretation result $\beta(I)$ and $\delta \in \{1, \dots, z-1\}$, a*

Algorithm 1 EFFICIENTLY DETECTING TRUSTWORTHY ANSWERS

```

1: Input: interaction setting  $(\mathcal{R}, \tau, U^r, U^s)$ , interpretation result  $\beta(I)$ 
2: Output: set  $C$  of trustworthy tuples in  $\beta(I)$ 
3:  $C \leftarrow \emptyset$ 
4:  $S \leftarrow \emptyset$ 
5: Compute  $[b_{\min}, b_{\max}]$  based on  $b(\cdot)$  and  $\mathcal{R}$ 
6: for  $\delta = 1$  to  $z-1$  do
7:   compute  $g \leftarrow g(z, \delta)$  and  $s \leftarrow s(z, \delta)$  ▷ Theorem 4.6
8:    $b_{\delta}^- \leftarrow \max\{g-1, s\}, b_{\delta}^+ \leftarrow g$ 
9:   if  $b_{\delta}^- < b_{\delta}^+$  then
10:    add  $(\delta, b_{\delta}^-, b_{\delta}^+)$  to  $S$  ▷ store feasible  $(\delta, b_{\delta}^-, b_{\delta}^+)$ 
11: for all  $e \in \beta(I)$  do
12:   trustworthy  $\leftarrow$  true
13:   for all  $(\delta, b_{\delta}^-, b_{\delta}^+) \in S$  do ▷ Interval for possible  $b(e')$ :  $b(e') \in [a, c)$ 
14:     $a \leftarrow b(e) - b_{\delta}^+, c \leftarrow b(e) - b_{\delta}^-$ 
15:    if  $\max\{a, b_{\min}\} < \min\{c, b_{\max}\}$  then
16:     trustworthy  $\leftarrow$  false
17:     break ▷ early exit
18:   if trustworthy then
19:     $C \leftarrow C \cup \{e\}$ 
20: return  $C$ 

```

tuple $e \in \beta(I)$ with bias $b(e)$ is untrustworthy if there exists any tuple e' such that $e' \notin \beta(I)$ or $r(e', \beta(I)) > r(e, \beta(I))$ and its bias $b(e')$ satisfies:

$$b(e') \in [b(e) - g(z, \delta), b(e) - \max\{g(z, \delta) - 1, s(z, \delta)\}]. \quad (3)$$

where $g(z, \delta) = \frac{-\delta^3 + 3\delta^2 z + \delta z^2 + \delta + z^2 - z}{d(z, \delta)}, s(z, \delta) = \frac{(z-\delta)(z-\delta+1)(2\delta+z-1)}{d(z, \delta)}$ and $d(z, \delta) = 3(z^2 - z + \delta(2z + 1) - \delta^2) > 0$

Based on Theorem 4.6 a tuple $e \in \beta(I)$ is untrustworthy whenever there exists a tuple e' (either not in $\beta(I)$ or ranked strictly lower than e in $\beta(I)$) whose bias satisfies the condition in Equation 3 for some δ . Specifically, the theorem provides two requirements for $b(e')$ values. First, $g(z, \delta)$ characterizes which bias gaps make δ a feasible indifference threshold. Second, $s(z, \delta)$ is the minimum bias difference between the two tuples $b(e) - b(e')$ such that the data source ranks e ahead of e' to maximize its utility.

Efficient Algorithm. Algorithm 1 implements the condition in Theorem 4.6 using only the bias range $[b_{\min}, b_{\max}]$ of the data source's bias function $b(\cdot)$ based on the schema domain. It first computes $g(z, \delta)$ and $s(z, \delta)$ for all $\delta \in \{1, \dots, z-1\}$ and keeps only those δ for which the theorem's interval is nonempty, i.e., $b_{\delta}^- = \max\{g(z, \delta) - 1, s(z, \delta)\} < b_{\delta}^+ = g(z, \delta)$. For a returned tuple $e \in \beta(I)$ and fixed δ the algorithm checks Equation 3. Since individual values $b(e')$ are unknown, the algorithm certifies $e \in \beta(I)$ as trustworthy whenever this interval has empty intersection with the feasible bias range $[b_{\min}, b_{\max}]$ for every δ retained above. If the intersection is empty for all retained δ , then no feasible bias value can satisfy the theorem's condition, and e is trustworthy. If the intersection is nonempty for some δ , then e is potentially untrustworthy.

Complexity. Computing $g(z, \delta), s(z, \delta)$ for all δ and constructing S takes $O(z)$ time. The per-tuple check evaluates at most $|S| \leq z-1$ overlap tests, each in $O(1)$ time. Thus the total running time is $O(z + k|S|) \subseteq O(kz)$ for $k = |\text{lm}(\beta(I))|$ tuples.

5 INFLUENTIAL STRATEGIES

The data source may steer results toward sponsored, high-margin, or otherwise biased products. As a result, directly expressing an

intent, e.g., searching for a particular low-priced product, may yield a ranked list dominated by promoted alternatives. A user may, therefore, submit a query that strategically adjusts the ranking of products revealed to the data source to recover relevant results. However, the effectiveness of users' strategies may depend on their interpretation by the data source and the data source's degree of bias. In this section, we investigate these settings and propose algorithms for finding strategies that lead the interaction to an equilibrium and improve the user's utility.

5.1 Detecting Influential Query

We first characterize what information about the user's intent can be credibly communicated to the data source. Building on this, we present an algorithm for finding a query that leads the interaction to an influential equilibrium.

Problem 3. *Given the interaction setting $(\mathcal{R}, \tau, U', U^s)$ find a query that leads to an influential equilibrium.*

In Section 4, we identified untrustworthy rankings in the data source's interpretation. A user aware of such misrankings may counteract this bias by adjusting the ranking of tuples in their queries. For instance, if a shopping website has an incentive to promote a product e' (or demote e), the user's query may rank the product e higher than e' only when e is ranked *sufficiently higher* in the intent. We formalize this notion using relative rank constraints.

Relative rank constraints. Given an intent τ over an instance I of schema \mathcal{R} , tuples $e, e' \in \tau(I)$, and an integer $\delta \geq 1$, a relative rank constraint is the inequality $r(e', \tau(I)) - r(e, \tau(I)) \geq \delta$. Intuitively, the inequality says that e is at least δ positions above e' in the intent.

THEOREM 5.1. *Consider the same setting as Theorem 4.6. For an ordered pair $(e, e') \in \tau(I)$, let $b = b(e) - b(e')$, and let $g(z, \delta)$ be as in Theorem 4.6. Define $\delta^*(e, e')$ to be the (largest) integer solution $\delta \in \{1, \dots, z-1\}$ to $g(z, \delta) - 1 < b \leq g(z, \delta)$, and set $\delta^*(e, e') = \perp$ if no such integer solution exists. If $\delta^*(e, e') \neq \perp$, then the corresponding relative rank constraint is: $r(e', \tau(I)) - r(e, \tau(I)) \geq \delta^*(e, e')$*

Theorem 5.1 provides a closed form solution to find the relative rank constraint between tuples. Informally, the theorem says that the data source's larger bias in favor of e' requires the user to provide stronger evidence that e truly outranks e' in the intent, i.e., a larger δ^* . The following result follows from the theorem.

PROPOSITION 5.2. *If there exists a pair e, e' such that $\delta^*(e, e') \neq \perp$, then the interaction has an influential equilibrium.*

Example 5.3. Suppose an intent τ over Headphones instance I places e_2 far above e_3 , e.g., $r(e_2, \tau(I)) = 1$ and $r(e_3, \tau(I)) = 3$. Consider the ordered pair (e_2, e_3) and assume the data source bias difference is $b = b(e_2) - b(e_3)$ with $b < 0$, i.e., the data source is relatively more biased to promote e_3 . Suppose Theorem 5.1 yields $\delta^*(e_2, e_3) = 2$. The corresponding δ^* ranking constraint is $r(e_3, \tau(I)) - r(e_2, \tau(I)) \geq 2$. Since $r(e_3, \tau(I)) - r(e_2, \tau(I)) = 2$, a user with intent τ can submit this relative ranking information to influence the data source's interpretation. Now consider a different intent τ' on the same instance where e_2 is only one position above e_3 , e.g., $r(e_2, \tau'(I)) = 2$

and $r(e_3, \tau'(I)) = 3$. Then $r(e_3, \tau'(I)) - r(e_2, \tau'(I)) = 1 < 2$, so the above constraint does not hold under τ' and a user with intent τ' would submit the complementary information (i.e., $r(e_3, \tau'(I)) - r(e_2, \tau'(I)) < 2$) instead.

Definition 5.4. Given an intent τ over an instance I of schema \mathcal{R} , a δ -query q^δ is any conjunction of relative rank constraints of the form $r(e', \tau(I)) - r(e, \tau(I)) \geq \delta$, where $\delta \in \mathbb{Z}$. Equivalently, for some set $\mathcal{V} \subseteq \tau(I) \times \tau(I) \times \mathbb{Z}$,

$$q^\delta = \bigwedge_{(e, e', \delta) \in \mathcal{V}} (r(e', \tau(I)) - r(e, \tau(I)) \geq \delta).$$

LEMMA 5.5. *For integer-valued ranks, $\neg(r(e', \tau(I)) - r(e, \tau(I)) \geq \delta)$ iff $r(e, \tau(I)) - r(e', \tau(I)) \geq 1 - \delta$.*

Algorithm. Algorithm 2 constructs the query q^δ that the input intent τ can submit to influence the data source's interpretation. As discussed in Section 3.2, user's usually do not know the exact instance I , therefore the algorithm uses the information based on the domain of schema \mathcal{R} . We therefore run the algorithm over the *rank domain* induced by the schema and the intent's ranking attributes. For instance, if τ has a ranking attribute A (e.g., ORDER BY rating), we take the rank domain to be $\text{dom}(A)$; if τ ranks by attributes A_1, \dots, A_l , we take the rank domain to be the product $\text{dom}(A_1) \times \dots \times \text{dom}(A_l)$. We treat τ as an order over the domain of the cartesian product and work with its induced rank function $r(\cdot, \tau)$ on this domain.

The algorithm iterates over strictly ordered pairs $e \prec_\tau e'$ in this domain and invokes the subroutine (Algorithm 3) to compute $\delta^*(e, e')$ based on Theorem 5.1. While Theorem 5.1 provides a closed form for the quadratic case, the approach generalizes to other additive, supermodular utility functions with additive bias. For other supermodular utility functions, one may replace the solution for $g(\cdot)$. To avoid materializing all values $g(\cdot)$, the Algorithm 3 evaluates $g(\cdot)$ on demand and finds δ^* via binary search over $\delta \in \{1, \dots, z-1\}$. The output query q^δ is a conjunction of relative rank constraints, where \mathcal{V} contains exactly one constraint per pair considered. Importantly, q^δ is constructed to be *true under the input intent τ* . Intuitively, each pair (e, e') with $\delta^*(e, e') \neq \perp$, the algorithm adds either the forward constraint $r(e', \tau) - r(e, \tau) \geq \delta^*$ if it holds, or otherwise adds the logically equivalent complementary constraint $r(e, \tau) - r(e', \tau) \geq 1 - \delta^*$ (Lemma 5.5).

Example 5.6. Suppose given a query over Headphones instance I , Algorithm 2 outputs the query $q^\delta = (r(e_2, \tau(I)) - r(e_1, \tau(I)) \geq 1) \wedge (r(e_4, \tau(I)) - r(e_3, \tau(I)) \geq 1)$. This conjunction enforces e_1 to be ranked above e_2 and e_3 to be ranked above e_4 , but imposes no constraint between the pairs $\{e_1, e_2\}$ and $\{e_3, e_4\}$. Hence q^δ is consistent with multiple ranked lists. For instance, $q_1^\delta(I) : e_1 \prec e_2 \prec e_3 \prec e_4$, $q_2^\delta(I) : e_3 \prec e_4 \prec e_1 \prec e_2$ both satisfy q^δ .

Semantics of q^δ . The output of Algorithm 2 is a query q^δ that is a conjunction of relative rank constraints. Informally, q^δ covers a **set of rankings** consistent with its constraints. Depending on the constraints in q^δ , the set of rankings may be (1) empty, (2) singleton, or (3) contains multiple rankings. To provide intuition, note that ranks lie in the finite set $[z]$. Hence, each inequality expands into a finite disjunction over differences in ranks. In particular, for any ordered pair (e, e') and any $\delta \in \{1, \dots, z-1\}$, $r(e', \tau) - r(e, \tau) \geq$

$\delta \equiv \bigvee_{t=\delta}^{z-1} (r(e', \tau) - r(e, \tau) = t)$. Equivalently, since $r(e', \tau) \in [z]$, the constraint can be written as a disjunction over the possible ranks of e' conditional on the rank of e i.e. $r(e', \tau) - r(e, \tau) \geq \delta \equiv \bigvee_{j=1}^z (r(e, \tau) = j \wedge r(e', \tau) \geq j + \delta)$, where terms with $j + \delta > z$ are unsatisfiable. For example, if $\delta = 2$ and $r(e, \tau) = 4$, then any $r(e', \tau) \in \{6, 7, \dots, z\}$ satisfies the inequality. This explains the three cases. Constraints with $\delta > 1$ impose only minimum separation, so they often leave slack and admit many rankings. If enough constraints use $\delta^* = 1$, they can fix a total order and yield a singleton. If constraints conflict (e.g., force $j + \delta > z$), the set is empty.

Algorithm 2 Detect Influential Queries

```

1: Input: interaction setting  $(\mathcal{R}, \tau, U^r, U^s)$ 
2: Output: query  $q^\delta$ 
3:  $\mathcal{V} \leftarrow \emptyset$  ▷ Set of relative rank constraints
4: for each ordered pair  $(e, e') \in \tau$  do ▷ Iterate in  $\tau$ -order
5:    $\ell \leftarrow \text{CONSTRAINT-QUADRATIC}(e, e', \tau, b(\cdot))$ 
6:   if  $\ell \neq \perp$  then
7:      $\mathcal{V} \leftarrow \mathcal{V} \cup \{\ell\}$ 
8:  $q^\delta \leftarrow \bigwedge_{\ell \in \mathcal{V}} \ell$  ▷ Definition 5.4
9: return  $q^\delta$ 

```

Algorithm 3 CONSTRAINT-QUADRATIC

```

1: Input:  $e, e', \tau, b(\cdot)$ 
2: Output: relative rank constraint (or  $\perp$ )
3:  $\Delta b \leftarrow b(e) - b(e')$  ▷ Theorem 5.1
4:  $\delta^* \leftarrow$  binary search over  $\delta \in \{1, \dots, z-1\}$  using  $g(z, \delta)$ 
5: if  $\delta^* = \perp$  then
6:   return  $\perp$ 
7: if  $r(e', \tau) - r(e, \tau) \geq \delta^*$  then
8:   return  $(r(e', \tau) - r(e, \tau) \geq \delta^*)$ 
9: else
10:  return  $(r(e, \tau) - r(e', \tau) \geq 1 - \delta^*)$  ▷ Lemma 5.5

```

Complexity. Algorithm 2 iterates over ordered pairs of tuples in the results of the intent over the schema domain. In the worst case (when τ is a total order) resulting in a search space of size $\binom{m}{2} = O(m^2)$, where $m = \prod_{i=1}^l |\text{dom}(A_i)|$ is the size of the domain of the ranking attributes A_1, \dots, A_l in intent τ . For each pair, the algorithm locates δ^* satisfying $\Delta b \in (g(z, \delta) - 1, g(z, \delta)]$. If δ^* is found by scanning $\delta \in \{1, \dots, z-1\}$, the per-pair cost is $O(z)$ and the total runtime is $O(m^2 z)$. If $g(z, \delta)$ is nondecreasing in δ , binary search requires $O(\log z)$ evaluations of $g(\cdot)$ per pair, yielding $O(m^2 \log z)$ total time. The output query q^δ contains at most $O(m^2)$ constraints.

Bucketization. The complexity of Algorithm 2 depends on the size of the rank domain induced by the intent's ranking attributes, which can be large when attributes have high-cardinality domains. Following standard practice [3, 16], we therefore bucketize large attribute domains into a small number of groups. For example, continuous attributes such as rating can be discretized into 3–4 buckets. This choice is also motivated by practice. In many data sources, bias is naturally expressed at the level of *groups* of tuples, e.g., owners' brands, sponsored sellers, or coarse price bands, rather than at the level of individual tuples [16]. We provide the bucketization details in Section 6.2 and experimentally study how this choice affects the scalability of the algorithm and the user's utility.

Limitations on query language. Supporting queries like q^δ requires a language that can express constraint-based ranking specifications or a data source that accepts a union of ranking queries. As discussed above, q^δ denotes the set of rankings consistent with its constraints. This set determines whether q^δ can influence the data source's interpretation. If it is empty, q^δ is uninfluential. If it covers the entire space of queries, e.g., all permutations of tuples in the user's intent, then q^δ is again uninfluential. Only when q^δ rules out some rankings but not all can it convey nontrivial information and influence the data source's posterior. Then users can submit q^δ , and it leads the interaction to an influential equilibrium (Proposition 5.2). However, our current query language (Section 2.1) restricts users to single SQL queries whose ranking is determined by an ORDER BY clause. Such queries necessarily produce a single-ranked list and cannot directly encode relative rank constraints. To bridge this gap, we can select one ranking from the non-trivial set of consistent rankings and q^δ and submit its SQL encoding. Due to the prevalence of SQL, in the following sections, we focus on interactions with SQL queries and leave extensions to other query languages as future work.

5.2 Improving User Utility

In the previous section, we presented a query that convinces the data source to return some relevant results despite the data source's bias. However, it may still return only a small subset of the tuples relevant to the user's intent. In practice, the user may refine this query by adjusting the ranking further. In this section, we investigate such refinement and provide an algorithm to find a query that improves the user's expected utility.

Example 5.7. Consider the Headphones instance in Table 1 and let the user's intent τ be to retrieve the highest-rated headphones. Suppose the data source returns top-3 tuples and has per-tuple utility function $u_e^s = (r(e, \tau(I)) - b(e)) r(e, \beta(I))$, with bias $b(e) = 1.9$. Given a query q , let $\Phi_e = \mathbb{E}[r(e, \tau(I)) \mid q]$ denote the posterior mean rank of the tuple e . Given the posterior mean, the data source maximizes $\mathbb{E}[u_e^s \mid q] = (\Phi_e - 1.9) r(e, \beta(I))$. Hence the data source assigns e position $r(e, \beta(I)) = 1$ if $\Phi_e < 1.9$, and $r(e, \beta(I)) = 4$ (not returned) if $\Phi_e > 1.9$ to maximize utility. Now consider the query equivalent to the user's intent q : ORDER BY rating DESC, which induces the ranking $e_3 \prec e_2 \prec e_1 \prec e_4$. Based on theory of order statistics [10], conditioning a uniform prior over the $4!$ possible intent orders given q yields posteriors: $\Phi_{e_3} = 1, \Phi_{e_2} = 2, \Phi_{e_1} = 3, \Phi_{e_4} = 4$. Therefore only e_3 is returned. Now consider another query q' that results in $q'(I) : e_3 \sim e_2 \prec e_1 \prec e_4$. Under q' , the posterior is $\Phi'_{e_3} = \Phi'_{e_2} = \frac{1+2}{2} = 1.5, \Phi'_{e_1} = 3, \Phi'_{e_4} = 4$. Thus both e_3 and e_2 are returned.

Extending the above example, we begin by introducing a class of query refinements that are useful for improving the returned results.

Definition 5.8. Given queries q, q' over an instance I of schema \mathcal{R} , q' is a **super-rank query** of q if:

- (1) $\text{elm}(q(I)) \subseteq \text{elm}(q'(I))$.
- (2) For all $e, e' \in q(I)$, if $r(e, q(I)) = r(e', q(I))$, then $r(e, q'(I)) = r(e', q'(I))$.

- (3) For every $e \in q'(I) \setminus q(I)$ and every $e' \in q(I)$, we have $r(e, q'(I)) > r(e', q'(I))$.

We denote the set of all super-rank queries of q by $\mathcal{Q}_{\text{super}}(q)$.

Intuitively, executing the super-rank query ($q' \in \mathcal{Q}_{\text{super}}$) over some instance may (i) change a sub-list of ranked tuples to a set of tuples with the same position or (ii) preserve some of the orders (iii) append additional (ranked) tuples at the end of the list. Each query is a super-rank of itself.

Example 5.9. Given a query q that returns $e_1 \sim e_3 \prec e_4 \prec e_2$ over Headphones instance I . Results of some super-rank queries of q are:

- (a) set : $e_1 \sim e_3 \sim e_4 \sim e_2$
(b) preserve order : $e_1 \sim e_3 \prec e_4 \prec e_2$
(c) append : $e_1 \sim e_3 \prec e_4 \prec e_2 \prec e_5$

Super-rank strategy. As discussed in Section 5.1, under our query language, the user must submit a single query that induces a single ranked list. Accordingly, given an intent τ , we run Algorithm 2 to obtain the query q^δ (Definition 5.4), and then fix a base query q^{base} by selecting any executable query that is consistent with q^δ .

Definition 5.10. A user strategy $P^r(\cdot | \cdot)$ is a **super-rank strategy** if for every intent τ and query q , we have $P^r(q | \tau) > 0$ only if $q \in \mathcal{Q}_{\text{super}}(q^{\text{base}})$.

Next, we show that the interaction between the user and the data source, when the user leverages a super-rank strategy, always leads to an equilibrium.

THEOREM 5.11. *Given the interaction setting $(\mathcal{R}, \tau, U^r, U^s)$ suppose the user and data source have the following utility functions for each tuple e :*

$$u^r(r(e, \tau(I)), r(e, \beta(I))) = -(r(e, \tau(I)) - r(e, \beta(I)))^2,$$

$$u_e^s(r(e, \tau(I)), r(e, \beta(I))) = -(r(e, \tau(I)) - (r(e, \beta(I)) + b(e)))^2.$$

over instance I of \mathcal{R} . Then, if the user adopts a super-rank strategy, the interaction has an influential equilibrium.

Example 5.12. Consider the two super-ranks of q from Example 5.9. $q_1^{\text{sup}}(I): e_1 \sim e_3 \sim e_4 \sim e_2$ and $q_2^{\text{sup}}(I): e_1 \sim e_3 \sim e_4 \prec e_2$. The query that equally prefers three headphone brands (q_1^{sup}) and the one that has a higher preference for specific headphone brands (q_2^{sup}) are super-rank queries and both lead to equilibria in the interaction.

The above result indicate that to lead the reasoning to some equilibrium, it is sufficient for the user to choose certain types of strategies. *Nonetheless, users may prefer some equilibria over others.* As illustrated in Example 5.12, using some super-rank queries, the user can overstate their preference for some tuples that might otherwise not be returned by the data source. However, a super-rank strategy that maps an intent to a query that places too many tuples in higher positions than their true positions in the intent or adds too many tuples to the intent may reduce the expected posterior of all tuples. For instance, in Example 5.7, again, assume that the data source has only four headphones. If the user submits a query that returns all headphones in the data source, the expected posterior of every tuple will be $\frac{5}{2}$. If the data source returns only tuples where

posterior means are above 2.5, it will not return any tuple. Thus, the user strategy must find the right amount of modification to the intent to offset the bias and maximize extracted information.

Fully Influential Strategy. In a fully influential equilibrium, the data source returns all tuples relevant to the intent behind the submitted user's query, which is desirable for users:

Definition 5.13. An equilibrium $(P^r(\cdot | \cdot), P^s(\cdot | \cdot))$ is **fully influential** if for every query q and intent τ such that $P^r(q | \tau) > 0$, τ is a super-rank of β for all β where $P^s(\beta | q) > 0$. A user strategy P^r is a **fully influential strategy** if it leads to a fully influential equilibrium.

To achieve a fully influential equilibrium, the user must convince the data source that all its desired tuples are highly relevant to their intent. But, as we have explained in the preceding paragraphs, it might not always be possible to achieve that. Thus, one might aim for an equilibrium that *extracts as much useful information as possible from the data source*.

Definition 5.14. An equilibrium $E_{\text{max}} = (P_{\text{max}}^r(\cdot | \cdot), P_{\text{max}}^s(\cdot | \cdot))$ is **maximally influential** if for every $E = (P^r(\cdot | \cdot), P^s(\cdot | \cdot))$ in which P^r is a super-rank strategy, for every query q , and every β_{max} with $P_{\text{max}}^s(\beta_{\text{max}} | q) > 0$, we have that β_{max} is at least as influential as every β with $P^s(\beta | q) > 0$. A user strategy P_{max}^r is a **maximally influential strategy** if it leads to a maximally influential equilibrium.

Next, we present algorithms for finding the user strategy that led to a maximally influential equilibrium (Definition 5.14).

Problem 4. *Given the interaction setting $(\mathcal{R}, \tau, U^r, U^s)$ find a query q^* , that leads to a maximally influential equilibrium.*

Baseline Algorithm. A brute-force approach is to first get the q^{base} derived from τ and then enumerate all possible super-rank queries $\mathcal{Q}_{\text{super}}(q^{\text{base}})$. Then, for each super-rank query (q^{sup}), get the data source's interpretation (β), and compute the user's utility $U^r(q^{\text{sup}}, \beta)$, finally returning one that maximizes user utility. However, the size of the set $\mathcal{Q}_{\text{super}}$ can be super-exponential (Proposition 5.15), particularly because super-rank queries can append an arbitrary number of additional tuples. A simple optimization is to separate the set of super rank queries into two subsets: (i) those that append additional tuples and (ii) those that do not. We can first find the best superrank in the second set and then append additional tuples to it. This will reduce the search space significantly. However, this subset of super-ranks is also exponential, and the problem remains computationally hard. We prove that finding the maximally influential query is, in the general case, NP-hard.

PROPOSITION 5.15. *Let m be the number of tuples in the results of q^{base} over an instance I of schema \mathcal{R} , and let r be the number of tuples that may be appended. The algorithm outputs super-exponential number of equilibria and runs in super-exponential time in r .*

THEOREM 5.16. *The problem of finding the maximally influential strategy is NP-hard in general*

Merge query. Despite the general problem being NP-hard, for additive utility functions, the problem has an optimal substructure that is suited for dynamic programming (DP). A super-rank of q^{base} is obtained by incremental changing a sub-list of ranked tuples to a set of tuples with the same position.

Definition 5.17. Given queries q, q' over an instance I of schema R , let m denote the number of distinct rank positions in $q(I)$. We call q' a **merge query** of q if:

- (1) $\text{elm}(q(I)) = \text{elm}(q'(I))$
- (2) there exist integers $1 \leq i \leq j \leq m$ such that, for every $e \in q(I)$,

$$r(e, q'(I)) = \begin{cases} r(e, q(I)) & \text{if } r(e, q(I)) < i, \\ i & \text{if } i \leq r(e, q(I)) \leq j, \\ r(e, q(I)) - (j - i) & \text{if } r(e, q(I)) > j. \end{cases}$$

Intuitively q' is obtained from q by merging the consecutive rank positions $i, i + 1, \dots, j$ into a single tied position at rank i . We denote the set of merge queries of q by $\mathcal{Q}_{\text{merge}}(q)$.

LEMMA 5.18. Given a query q let m denote the number of rank positions in the results of q over schema \mathcal{R} . For each interval $1 \leq i \leq j \leq m$, let $q_{i:j} \in \mathcal{Q}_{\text{merge}}(q)$ denote the merge query obtained by merging exactly positions i, \dots, j of q . Let u be the expected per-tuple user utility under $q_{i:j}$. For any merge query $q' \in \mathcal{Q}_{\text{merge}}(q)$, let $\Pi(q') = \{[i_1, j_1], \dots, [i_t, j_t]\}$ be the unique partition of $\{1, \dots, m\}$ into consecutive intervals that q' merges into ties. Then the user's expected utility by submitting q' can be written as $\sum_{\ell=1}^t n_{i_\ell:j_\ell} \bar{u}(i_\ell, j_\ell) + C$, where $n_{i:j} = |\{e \in q : i \leq r(e, q) \leq j\}|$ and C is independent of q' .

Intuitively, Lemma 5.18 says that the expected user utility contributed by merging any consecutive position interval $[i, j]$ depends only on that interval, and not on how other positions are merged.

Efficient Algorithm. Algorithm 4 computes a maximally influential query. Given an intent τ , we first run Algorithm 2 to obtain the constraint query q^δ , and then fix an executable base query q^{base} by selecting any query whose induced ranked list is consistent with q^δ (Definition 5.10). As in Algorithm 2, we do not assume access to the realized instance; instead we work over the m rank positions induced by the ranking attributes A_1, \dots, A_l in τ , where $m = \prod_{i=1}^l |\text{dom}(A_i)|$.

A merge query $q' \in \mathcal{Q}_{\text{merge}}(q)$ coarsens this ranking by selecting consecutive position intervals and tying all tuples whose original ranks fall in each selected interval. Equivalently, each q' corresponds to a unique partition of $\{1, \dots, m\}$ into consecutive intervals. By Lemma 5.18, the expected user utility decomposes as a sum of interval scores. Thus, optimizing over $\mathcal{Q}_{\text{merge}}(q)$ reduces to choosing the partition with maximum total score. For each interval $[i, j]$, the algorithm computes $\text{Score}(i, j) = n_{i:j} \cdot u$, $n_{i:j} = |\{e \in q : i \leq r(e, q) \leq j\}|$ where u is the expected per-tuple user utility when exactly positions i, \dots, j are merged (and all other positions remain as in q). To compute u we evaluate the data source's interpretation to the merge query $q_{i:j} \in \mathcal{Q}_{\text{merge}}(q)$ using its utility function $\beta_{i:j} \in \arg \max_{\beta'} \mathbb{E}[U^s(\tau, \beta') \mid q_{i:j}]$. However, some utility functions may have multiple maxima, and therefore the arg max may not be unique. That is, the data source could be indifferent between several optimal interpretations (β) that provide the same

maximal utility. For such utility functions, we assume a consistent *tie-breaking rule*: if multiple interpretations $\{\beta_1, \beta_2, \dots\}$ yield the same maximal utility, the data source chooses the one most favorable to the user. Since the source does not know the user's true intent τ , it selects the interpretation that is closest to the user's query, i.e. β that maximizes $U^r(q, \cdot)$. This is a reasonable assumption since the data source wants to keep the user engaged.

Finally, given $\text{Score}(i, j)$, the algorithm applies the DP recurrence in Proposition 5.19 to compute the optimal partition of $\{1, \dots, m\}$ into consecutive intervals and then backtracks to construct the corresponding merge query $q^* \in \mathcal{Q}_{\text{merge}}(q)$.

PROPOSITION 5.19. Let $\text{OPT}(j)$ denote the maximum total utility attainable by merging a prefix of positions $1, \dots, j$. Then

$$\text{OPT}(0) = 0, \quad \text{OPT}(j) = \max_{1 \leq i \leq j} \{\text{OPT}(i-1) + \text{Score}(i, j)\}.$$

Algorithm 4 Maximally Influential Query

```

1: Input: interaction setting  $(\mathcal{R}, \tau, U^r, U^s)$ 
2: Output:  $q^* \in \mathcal{Q}_{\text{merge}}(q^{\text{base}})$ 
3:  $q^{\text{base}} \leftarrow$  Algorithm 2  $(\mathcal{R}, \tau, U^r, U^s)$ 
4: function  $\text{SCORE}(i, j)$  ▷ Score for merging positions  $i..j$  (Lemma 5.18)
5:    $q_{i:j} \leftarrow$  merge- $[i, j]$  query in  $\mathcal{Q}_{\text{merge}}(q)$ 
6:    $\beta \in \arg \max_{\beta'} \mathbb{E}[U^s(\tau, \beta') \mid q_{i:j}]$ 
7:    $n \leftarrow |\{e \in q : i \leq r(e, q) \leq j\}|$ 
8:    $u \leftarrow$  expected per-tuple user utility under  $(q_{i:j}, \beta)$ 
9:   return  $n \cdot u$ 
10:  $\text{DP}[0] \leftarrow 0$  ▷ DP over prefixes (Proposition 5.19)
11: for  $j = 1$  to  $m$  do
12:    $(\text{DP}[j], \text{parent}[j]) \leftarrow \max_{1 \leq i \leq j} \{\text{DP}[i-1] + \text{SCORE}(i, j)\}$ 
13:  $\Pi^* \leftarrow$  backtrack intervals from  $\text{parent}[\cdot]$ 
14:  $q^* \leftarrow$  merge query in  $\mathcal{Q}_{\text{merge}}(q)$  that merges exactly intervals in  $\Pi^*$ 
15: return  $q^*$ 

```

Complexity. As in Algorithm 2, we work over the rank domain induced by the intent's ranking attributes A_1, \dots, A_l , whose size is $m = \prod_{i=1}^l |\text{dom}(A_i)|$. Algorithm 4 performs dynamic programming over the m induced rank positions. The DP considers all intervals (i, j) with $1 \leq i \leq j \leq m$, yielding $O(m^2)$ transitions, and backtracking takes $O(m)$ time. Tabulating the interval scores $\text{Score}(i, j)$ also ranges over $O(m^2)$ intervals. If each $\text{Score}(i, j)$ can be computed in $O(1)$ time (e.g., using cached posterior/utility summaries per interval), the overall running time is $O(m^2)$; otherwise it is $O(m^2)$ plus the cost of computing all $\text{Score}(i, j)$ values.

6 EMPIRICAL EVALUATION

We empirically evaluate our proposed solutions to demonstrate their usefulness and effectiveness. We describe the experiment setup and then present a quantitative study assessing the efficiency of our algorithm, considering various datasets, queries, attributes, and the size of the possible domain.

6.1 Experiment Setup

Platform. All experiments were performed on a macOS Sequoia 15.6 machine with a Apple M1 Pro chip and 16GB RAM.




Figure 1: Impact of number of attributes on execution time of Detecting Influential (DI) and Maximally Influential (MI)

Table 2: Summary of datasets used in the experiments.

Dataset	Tuples	Attributes	Ranking Attributes ($ A $)
Amazon	14M	11	rating (41), sales in last month (30), price (29961)
PriceRunner	35K	7	seller (306), product model (12849), category (10)
Flights	300K	12	airlines (6), days from trip (49), price (12157)
Census	49K	12	race (5), sex (2), education level (16)
COMPAS	6.8K	12	age (65), prior counts (37), violence score (10)

Datasets. We evaluate our approach on five real-world datasets of varying sizes, domains, and conflict-of-interest scenarios, summarized in Table 2.

Amazon: This dataset contains over 14M product listings from Amazon [11]. The dataset contains attributes such as product category, price, rating, reviews, number of sales in the last month, and isBestSeller. Since Amazon may have an incentive to offload low-rated and less-selling products to clear inventory, we test for conflict of interest based on rating, sales in last month, and price.

PriceRunner: The dataset is collected from PriceRunner¹, a popular product comparison platform [1]. It includes 35,311 product offers from 10 categories (e.g., Mobile Phones) provided by different sellers. We test for conflict of interest based on attributes: seller, product model, and product category. For instance, the system may have an incentive to favor products from certain sellers.

Flights: The dataset was scraped from EaseMyTrip², a flight booking service provider [6]. It contains 300,261 distinct flight bookings from February 11th to March 31st, 2022. We test for conflict of interest based on attributes: airlines, days from trip, and price. For instance, the booking service may be biased towards expensive flights from certain airlines.

We also consider popular datasets used in prior work related to bias [16, 18].

Census: The US Census dataset comprises 49,000 tuples [5]. We test for conflict of interest based on race, sex, and education level. For example, the system may favor individuals with a higher education level (e.g., PhD).

COMPAS: The COMPAS dataset was collected and published by ProPublica³. It contains the demographics, violence scores produced

by the COMPAS software, and criminal offense information for 6,889 individuals. We test for conflict of interest based on age, prior offense counts, and violence decile score.

Queries. For each dataset, we generated a diverse set of SQL queries of the form `SELECT A_1, ..., A_l WHERE Sel(q) ORDER BY A_1, ..., A_l` where `Sel(q)` is a random conjunction of up to three non-key predicates and $l \in \{1, 2, 3\}$ varies the number of ranking attributes. Different subsets of filtering attributes were randomly selected for the WHERE clause. In the query projection, we include the same set of ranking attributes for simplicity, since our focus is on the conflict based on the ranking attributes. For consistency, we assume that the ranking order in the intent is decreasing for numeric attributes (e.g., `ORDER BY Rating DESC`) and use the inherent order in the original dataset for categorical attributes.

Utility Function. We use the quadratic utility functions throughout our experiments. As discussed in Section 3.2, the utility function of the data source might vary between tuples or sets of tuples. To vary the functions $u_e^s(\cdot, \cdot)$ and $u_e^c(\cdot, \cdot)$ in Equation 1 while maintaining consistency across different datasets, we randomly assign bias values in $[0, 1]$ and scale them based on the domain of different datasets. All our reported results are averaged over 5 runs.

Bucketization As discussed in Section 3.2, users usually do not know the exact instance I , therefore our algorithms reason over the domain of schema \mathcal{R} . Therefore, attributes with high-cardinality domains, such as price or product model, impact the scalability of our algorithms. Furthermore, a data source’s bias often applies to groups of items (e.g., expensive vs. cheap) rather than specific values, which motivates grouping domain values into a smaller number of bins. We explore a range of ranking attributes with different domain sizes (Table 2). Following standard practice [16], continuous attributes like age were bucketized into 3-4 groups based on their domain and values. For high-cardinality categorical attributes, we adopted another common approach: keeping the top- k most frequent values and mapping the rest to a single other bin [3]. For instance, in the PriceRunner dataset, we set $k = 1$ and defined bins over the top-10 possible domain values of sellers and models for each category. For datasets with small attribute domains like Census, no bucketization was necessary. We note that the selection of bucketization may affect the credibility of queries and the running time of our algorithms; therefore, we study this effect in Section 6.2.

¹<https://www.pricerunner.com/>

²<https://www.easymytrip.us/>

³<https://www.propublica.org/article/how-we-analyzed-the-compas-recidivism-algorithm>

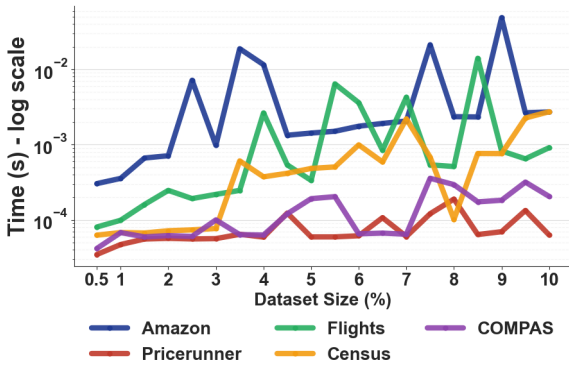


Figure 2: Impact of z on time to detect untrustworthy tuples

6.2 Experiment Results

In this Section, we first present the results of our algorithm in Section 4, whose goal is to detect trustworthy answers from data sources’ results. Then present results on algorithms in Section 5.2 for finding influential queries.

6.2.1 Detecting Trustworthy Answers. We present results on the scalability of Algorithm 1. The complexity of the algorithm depends on the number of tuples relevant to the user’s intent (z) and the number of tuples returned by the data source (k). For instance, at the time of this writing, Amazon displays top-16 results per page; we therefore set $k = 80$ to simulate a user viewing the first five pages of results. To maintain consistency across datasets of different sizes, we vary z as a percentage of each dataset’s total size. As shown in Figure 2, the runtime typically scales linearly with z . However, as discussed in Section 4, we observe some sub-linear performance because the algorithm often terminates early.

6.2.2 Influential Strategies. As explained in Section 5, the brute-force approach to find the maximally influential strategy is NP-hard due to the super-exponential complexity of the set of possible equilibria. Therefore, we only evaluate the efficient maximally influential algorithm (MI; Algorithm 4) and the algorithm for detecting influential query (DI; Algorithm 2).

Number of Attributes: The first set of experiments studies the effect of the size of the possible domain on the running time of the algorithms. To this end, we varied the number of attributes in the query from 1 to 3. The number of attributes (along with their cardinality) determines the search space. Figure 1 shows runtime as the number of ORDER BY attributes increases. Both algorithms slow down with larger datasets because the search space grows as the Cartesian product of attribute domains. The Amazon dataset incurs the highest cost owing to its large domain sizes, yet maximally influential (MI) still completes within the 15-minute cap for 3 attributes. The algorithm for finding the maximally influential strategy has a subroutine call Algorithm 2 and therefore has a higher complexity than detecting one influential query (DI).

Impact of Bucketization: To study the effect of this bucketization, we focus on the attributes with the largest pre-bucketization domains (Table 2): the continuous attribute price for Amazon and Flights, the categorical attribute product model for PriceRunner,

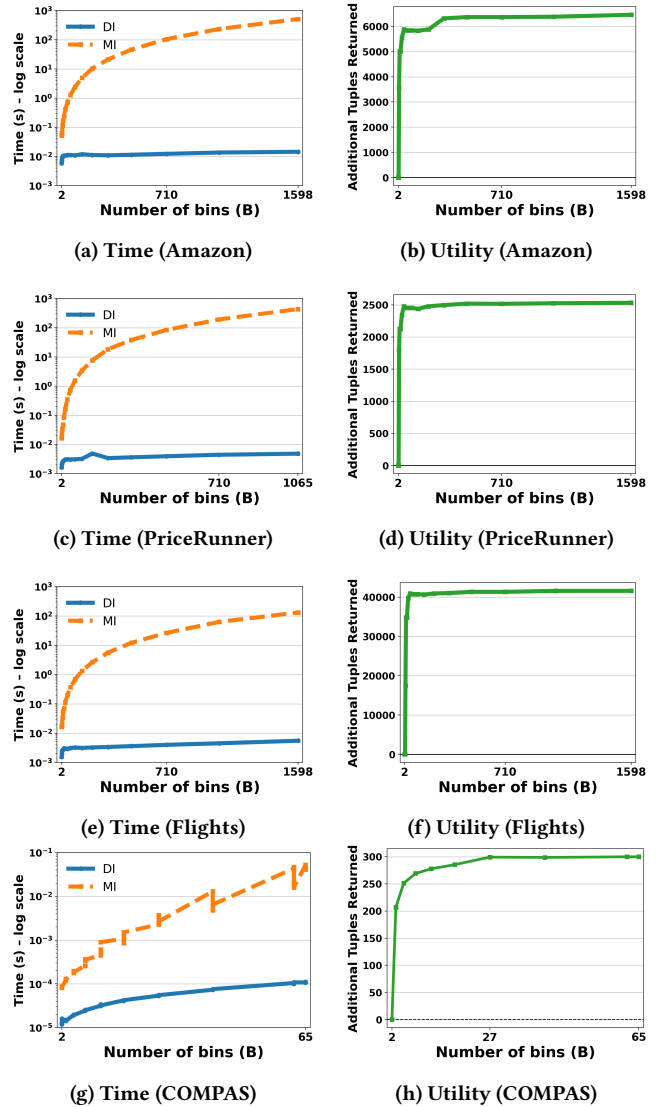


Figure 3: Impact of bucketization on running time and user utility across Amazon, PriceRunner, Flights, and COMPAS.

and age for COMPAS. Census, which is a popular dataset used in prior works related to bias [16, 18], did not require bucketization.

We consider two extremes for the data source’s bias: (i) *group-level* bias over coarse attribute bands (e.g., the most expensive bin) and (ii) *point-level* bias concentrated on specific values (e.g., \$99.99). To interpolate between these extremes while controlling runtime, we use a binary-tree (dyadic) refinement of the ordered domain. Specifically, we start with $B = 2$ equal contiguous bins and iteratively split each bin, yielding $B \in \{2, 4, 8, \dots\}$. For each bin level B , we run Algorithm 2 and Algorithm 4, while evaluating the resulting query on the *full* (unbucketized) domain. We enforce a 10-minute time budget per bin level.

Figure 3 summarizes the trade-offs across four datasets. For each dataset, we report execution time and user utility. The time plots

separate the cost of detecting an influential query from the end-to-end cost of computing the maximally influential query. The utility plots measure the effectiveness of the final query q^* : we report the gain in user utility relative to the original query (baseline) as a function of the additional relevant tuples recovered. As expected, the runtime grows with the number of bins B , since it increases both the number of pairwise checks and the DP search space. At the same time, finer binning identifies more credible ranking information, leading to queries that recover more relevant tuples. Empirically, the achieved utility is a monotone non-decreasing lower bound on the full-domain optimum.

7 RELATED WORK

Responsible Data Management. As explained in Section 1, our work shares the same goal as efforts in responsible data management and science [32]. In particular, researchers have proposed several fair ranking schemes [17, 21, 33–35]. Instead of designing a fair ranking scheme, our goal is to investigate whether informative querying is possible in the presence of a conflict of interest and to enable users to extract reliable information from biased data sources. Also, unlike our setting, these algorithms often require complete information about the underlying dataset. Researchers have investigated the trade-off between the efforts of agents, e.g., school applicants, and their ranks in a ranked-based reward system, e.g., school admission, and proposed methods to design robust ranking systems [20]. However, we focus on the challenges of querying a data source where there are conflicts of interest between the user and the data source. Unlike our settings, the efforts of agents in [20] may be highly costly, e.g., time on extracurricular activities.

Strategic Communication. There is substantial research on strategic communication between agents with non-identical interests in economics and game theory [4, 7–9, 19]. Our problem and approach have important differences with the settings of strategic communication in these fields. First, we propose computational methods to find the influential equilibria and strategies for ranking queries. Second, we address the challenges of querying with conflict of interest over large (structured) data.

8 APPENDIX

This section includes the proofs of our theoretical results.

8.1 Proofs for Section 3.2

PROOF OF THEOREM 3.1. The proof is simple, but we provide a longer proof to illustrate concepts.

Fix two set-equivalent intents $\tau \neq \tau'$ and two distinct interpretations $\beta \neq \beta'$. Consider the induced 2×2 game in which the user can send two queries q, q' and the data source chooses between β, β' after observing the query.

(\Leftarrow) Suppose first that for both $t \in \{r, s\}$,

$$U^t(\tau, \beta) \geq U^t(\tau, \beta') \quad \text{and} \quad U^t(\tau', \beta') \geq U^t(\tau', \beta).$$

Define strategies

$$P^r : \tau \mapsto q, \tau' \mapsto q', \quad P^s : q \mapsto \beta, q' \mapsto \beta'.$$

Since the two intents send different on-path queries (send with positive probability), Bayes' rule implies

$$\Pr(\tau \mid q) = 1, \quad \Pr(\tau' \mid q') = 1.$$

Hence, after observing q , the data source's expected utility from choosing β is $U^s(\tau, \beta)$, while from choosing β' it is $U^s(\tau, \beta')$. By assumption,

$$U^s(\tau, \beta) \geq U^s(\tau, \beta'),$$

so β is a best response after q . Similarly, after observing q' ,

$$U^s(\tau', \beta') \geq U^s(\tau', \beta),$$

so β' is a best response after q' .

Now consider the user's incentives. If intent τ follows P^r and sends q , it obtains utility $U^r(\tau, \beta)$. If it deviates and sends q' , the data source responds with β' , yielding utility $U^r(\tau, \beta')$. Since

$$U^r(\tau, \beta) \geq U^r(\tau, \beta'),$$

intent τ has no profitable deviation (First bullet in Definition 2.5). Likewise, type τ' obtains $U^r(\tau', \beta')$ from sending q' and would obtain $U^r(\tau', \beta)$ by deviating to q , and

$$U^r(\tau', \beta') \geq U^r(\tau', \beta),$$

so intent τ' also has no profitable deviation. Therefore (P^r, P^s) is an equilibrium, and it is influential since the two intents induce different interpretations.

If instead for both $t \in \{r, s\}$,

$$U^t(\tau, \beta') \geq U^t(\tau, \beta) \quad \text{and} \quad U^t(\tau', \beta) \geq U^t(\tau', \beta'),$$

the same argument applies after swapping the labels of β and β' . Hence the interaction is influential.

(\Rightarrow) Suppose the interaction is influential. Then there exists an equilibrium in which two set-equivalent intents $\tau \neq \tau'$ induce two distinct interpretations $\beta \neq \beta'$. Let the corresponding on-path queries be q and q' . Since the equilibrium is influential, the user's on-path behavior separates the two intents. Without loss of generality, suppose τ sends q and τ' sends q' , and the data source responds with β after q and with β' after q' .

Because these queries are sent on path by distinct intents, Bayes' rule implies

$$\Pr(\tau \mid q) = 1, \quad \Pr(\tau' \mid q') = 1.$$

Sequential rationality (second bullet in Definition 2.5) of the data source therefore requires

$$U^s(\tau, \beta) \geq U^s(\tau, \beta') \quad \text{and} \quad U^s(\tau', \beta') \geq U^s(\tau', \beta),$$

since β must be a best response (optimal) after q and β' must be a best response after q' .

Similarly, incentive compatibility for the user implies

$$U^r(\tau, \beta) \geq U^r(\tau, \beta') \quad \text{and} \quad U^r(\tau', \beta') \geq U^r(\tau', \beta),$$

because intent τ could deviate from q to q' and thereby induce β' , while intent τ' could deviate from q' to q and thereby induce β .

Thus the first set of inequalities in the statement holds for both U^r and U^s . If the equilibrium instead matches τ with β' and τ' with β , then the second set of inequalities holds. Therefore, the interaction is influential if and only if one of the two stated conditions is satisfied. \square

PROOF OF PROPOSITION 3.5. It follows from Theorem 1 in [8]. \square

PROOF OF THEOREM 3.6. Fix any query q and any tuple $e \in E$. Let $X_e := r(e, \tau(I)) + b_e$. Since $r(e, \tau(I)) \in [z]$, we have $X_e \in [1+b_e, z+b_e]$ almost surely under any belief. For any interpretation rank $\tilde{r}(e, \beta(I)) \in [z]$, each conditional expected utility term is

$$\mathbb{E}\left[L\left(r(e, \tau(I)) - (\tilde{r}(e, \beta(I)) + b_e)\right) \mid q\right] = \mathbb{E}\left[L\left(\tilde{r}(e, \beta(I)) - X_e\right) \mid q\right].$$

By the hypothesis, for every $r \in [z]$ and every $\tilde{r}(e, \beta(I)) \in [z]$,

$$L\left(r - (r_e^{\beta^*} + b_e)\right) \leq L\left(\tilde{r}(e, \beta(I)) + b_e\right).$$

Substituting $r = r(e, \tau(I))$ gives

$$L(r_e^{\beta^*} - X_e) \leq L(\tilde{r}(e, \beta(I)) - X_e)$$

Taking conditional expectations given q yields

$$\mathbb{E}\left[L(r_e^{\beta^*} - X_e) \mid q\right] \leq \mathbb{E}[L(\tilde{r}(e, \beta(I)) - X_e) \mid q] \quad \text{for all } \tilde{r}(e, \beta(I)) \in [z].$$

Thus, for tuple e the same interpretation rank $r_e^{\beta^*}$ is optimal after every query q . Since U^s is additive across tuples, the data source has an optimal interpretation that assigns each e the same rank $r_e^{\beta^*}$ independent of q . Therefore the data source response cannot depend on the query, and no equilibrium can be influential. \square

PROOF OF COROLLARY 3.7. Fix any query q and $e \in \tau(I)$. Since $r(e, \tau(I)) \in [k]$, let $X_e := r(e, \tau(I)) + b(e) \in [1+b(e), k+b(e)]$. If $b(e) \geq k - \frac{3}{2}$, then $X_e \geq k - \frac{1}{2}$, so k is the closest feasible rank to X_e and, since L is nondecreasing in $|\cdot|$, $r(e, \beta(I)) = k$ is optimal after every q . If $b(e) \leq \frac{3}{2} - k$, then $X_e \leq \frac{3}{2}$, so 1 is closest and $r(e, \beta(I)) = 1$ is optimal after every q . Thus each tuple has a query-independent optimal interpretation rank; by additivity, the entire interpretation is independent of q . Hence no equilibrium can be influential. \square

8.2 Proofs for Section 5.1

PROOF OF PROPOSITION 4.3. The proof is by algebraic simplification.

For (i), write $r_e := r(e, \tau)$, $\beta_e := r(e, \beta)$, $\beta'_e := r(e, \beta')$. Then

$$u^r(r_e, \beta_e) - u^r(r_e, \beta'_e) = -(r_e - \beta_e)^2 + (r_e - \beta'_e)^2 = 2r_e(\beta_e - \beta'_e) - (\beta_e^2 - \beta'^2_e).$$

Summing over e gives

$$U^r(\tau, \beta) - U^r(\tau, \beta') = \sum_{e \in E} \gamma_e r_e - \sum_{e \in E} (\beta_e^2 - \beta'^2_e),$$

so indifference is equivalent to $\sum_{e \in E} \gamma_e r(e, \tau) = \alpha^r$.

For (ii),

$$u^s(r_e, \beta_e) - u^s(r_e, \beta'_e) = \gamma_e r_e - \left((\beta_e + b_e)^2 - (\beta'_e + b_e)^2 \right).$$

Summing and taking expectation under Φ yields

$$\mathbb{E}_{\tau \sim \Phi}[U^s(\tau, \beta) - U^s(\tau, \beta')] = \sum_{e \in E} \gamma_e \Phi_e - \alpha^s,$$

where $\alpha^s = \sum_e ((\beta_e + b_e)^2 - (\beta'_e + b_e)^2) = \alpha^r + 2 \sum_e b_e(\beta_e - \beta'_e)$. Thus the data source is indifferent iff $\sum_{e \in E} \gamma_e \Phi_e = \alpha^s$. \square

PROOF OF LEMMA 4.4. Under the corollary's assumptions, $\gamma_{\bar{e}} = 0$ for $\bar{e} \notin \{e, e'\}$ and $\gamma_{e'} = -\gamma_e \neq 0$, so the user indifference condition in Proposition 4.3 reduces to $r(e, \tau) - r(e', \tau) = \frac{\alpha^r}{\gamma_e}$. Moreover, only e, e' contribute to the intercept, and using $r(e', \beta) - r(e', \beta') = -(\beta_e - \beta'_e)$ gives

$$\alpha^s = \alpha^r + 2(b_e - b_{e'}) (r(e, \beta) - r(e, \beta')) = \alpha^r + \gamma_e (b_e - b_{e'}).$$

Thus $\alpha^r = \alpha^s - \gamma_e (b_e - b_{e'})$, and substituting yields

$$r(e, \tau) - r(e', \tau) = \frac{\alpha^s}{\gamma_e} - (b_e - b_{e'}).$$

This establishes the claimed threshold form (2). \square

PROOF OF PROPOSITION 4.5. For an intent τ , define $d(\tau) = r(e, \tau(I)) - r(e', \tau(I))$. By Lemma 4.4, the user's indifference between β and β' is determined solely by whether $d(\tau)$ is above or below the threshold. The unbiased case uses threshold δ_0 and the biased case uses threshold δ . If $d(\tau) \leq \min\{\delta_0, \delta\}$, then $d(\tau) \leq \delta_0$ and $d(\tau) \leq \delta$, so τ is on the same (weak) side of both boundaries and induces the same interpretation in both bias and unbiased cases. If $d(\tau) > \max\{\delta_0, \delta\}$, then $d(\tau) > \delta_0$ and $d(\tau) > \delta$, so τ is again on the same side of both boundaries and induces the same interpretation in both cases. Therefore, the induced interpretation can change only when the relative position of the tuples lies between the two thresholds. \square

PROOF OF THEOREM 4.6. Let $r = r(e, \tau(I))$ and $r' = r(e', \tau)$, where smaller rank is better and a fixed $\delta \in \{1, \dots, z-1\}$. Each δ partitions the space of all possible ranks of two tuples $(r, r') \in [z]^2$ into two sets \mathcal{H} and \mathcal{H}' . Intuitively, \mathcal{H} corresponds to rank pairs where the difference $r' - r$ is at least δ , and \mathcal{H}' is its complement. We consider the δ -threshold partition that favors e :

$$\mathcal{H}(\delta) = \{(r, r') \in [z]^2 : r' - r \geq \delta\}, \quad \mathcal{H}'(\delta) = [z]^2 \setminus \mathcal{H}(\delta).$$

Thus, on $\mathcal{H}(\delta)$, tuple e outranks e' by at least δ intent ranks.

Conditioning the (uniform) prior on \mathcal{H} and \mathcal{H}' yields posterior mean ranks. Let $\Phi_e(\mathcal{H}) = \mathbb{E}[r \mid (r, r') \in \mathcal{H}]$ and $\Phi_{e'}(\mathcal{H}) = \mathbb{E}[r' \mid (r, r') \in \mathcal{H}]$. A direct count over the discrete triangle $\{r' - r \geq \delta\}$ gives

$$\Phi_e(\mathcal{H}) = \frac{z - \delta + 2}{3}, \quad \Phi_{e'}(\mathcal{H}) = \frac{2z + \delta + 1}{3}.$$

For the complement means $\Phi_e(\mathcal{H}') = \mathbb{E}[r \mid (r, r') \in \mathcal{H}']$ and $\Phi_{e'}(\mathcal{H}') = \mathbb{E}[r' \mid (r, r') \in \mathcal{H}']$, using total sums on $[z]^2$ and simplifying yields

$$\Phi_e(\mathcal{H}') = \frac{\delta^3 - 3(z+1)\delta^2 + (3z^2 + 6z + 2)\delta + 2z(z^2 - 1)}{d(z, \delta)},$$

$$\Phi_{e'}(\mathcal{H}') = \frac{-\delta^3 + 3\delta z^2 + (3z^2 + 1)\delta + z(z^2 - 1)}{d(z, \delta)}.$$

In particular, the difference across the complement region simplifies exactly to:

$$\Phi_e(\mathcal{H}') - \Phi_{e'}(\mathcal{H}') = s(z, \delta). \quad (4)$$

Assume an unconstrained, real-valued data source interpretation $r(\cdot, \beta) \in \mathbb{R}$ (no rounding/clipping) for simplicity, with quadratic utility $u_x^s(r_x^r, r_x^\beta) = -(r_x^r - (r_x^\beta + b_x))^2$. The unique optimal posterior interpretation after observing $X \in \{\mathcal{H}, \mathcal{H}'\}$ is:

$$r(x, \beta_x) = \Phi_x(X) - b_x, \quad x \in \{e, e'\}. \quad (5)$$

Although these scores are real-valued, the interpretation result $\beta(I)$ induces a discrete ordering by sorting tuples by $r(\cdot, \beta_X)$ (and returning the top- k). Hence, on \mathcal{H}' , tuple e is ranked ahead of e' iff $r(e, \beta_{\mathcal{H}'}) < r(e', \beta_{\mathcal{H}'})$.

Under $u'(r_x^\tau, r_x^\beta) = -(r_x^\tau - r_x^\beta)^2$, user indifference between the two posterior interpretations $\beta_{\mathcal{H}}$ and $\beta_{\mathcal{H}'}$ yields an affine boundary in (r, r') , which can be written as a threshold on $(r' - r)$. In a threshold equilibrium over the discrete integer grid $[z]^2$, any real indifference threshold $t \in \mathbb{R}$ induces exactly the same partition $\mathcal{H}(\delta)$ whenever $t \in (\delta - 1, \delta]$. Evaluating the user's indifference condition yields

$$\begin{aligned} t &= \frac{r(e', \beta_{\mathcal{H}}) + r(e', \beta_{\mathcal{H}'}) - (r(e, \beta_{\mathcal{H}}) + r(e, \beta_{\mathcal{H}'}))}{2} \\ &= \delta - g(z, \delta) + (b(e) - b(e')), \end{aligned}$$

where substituting the posterior means and the optimal data source's posterior interpretation (Equation 5) gives $g(z, \delta)$ as stated. Hence, the required separation condition $t \in (\delta - 1, \delta]$ holds if and only if

$$b(e) - b(e') \in (g(z, \delta) - 1, g(z, \delta)]. \quad (6)$$

If $\delta > 1$, the complement region $\mathcal{H}'(\delta)$ contains the nonempty strip $\mathcal{S}(\delta) = \{(r, r') \in [z]^2 : 1 \leq r' - r \leq \delta - 1\} \subseteq \mathcal{H}'(\delta)$, on which $r < r'$ (i.e., e truly outranks e') and also includes all intents with $r > r'$ (where e' truly outranks e). Tuple e is non-credible witnessed by e' if, under the posterior interpretation induced by $\mathcal{H}'(\delta)$, the data source ranks e ahead of e' :

$$r(e, \beta_{\mathcal{H}'}) < r(e', \beta_{\mathcal{H}'}) \iff b(e) - b(e') > \Phi_e(\mathcal{H}') - \Phi_{e'}(\mathcal{H}')$$

Substituting the difference across the complement region from Equation (4) into this inequality yields: $b(e) - b(e') > s(z, \delta)$. Intersecting this with Equation 6 gives the bias interval:

$$b(e) - b(e') \in \left(\max\{g(z, \delta) - 1, s(z, \delta)\}, g(z, \delta) \right]$$

□

8.3 Proofs for Section 5.2

PROOF OF THEOREM 5.11. The proof is by construction, extending the results of Proposition 5.2. Fix an ordered pair (e, e') with $\delta := \delta^*(e, e') \in \{1, \dots, z - 1\}$ and consider the two regions of the space of possible rank differences.

$$\mathcal{H} = \{(r, r') \in [z]^2 : r' - r \geq \delta\}, \quad \mathcal{H}' = [z]^2 \setminus \mathcal{H}.$$

By Proposition 5.2, there exists an influential equilibrium supported on the two queries that separate \mathcal{H} from \mathcal{H}' . Under our query language, the user must submit a single ranked list. Since both \mathcal{H} and \mathcal{H}' are nonempty, we can select two executable queries $q_{\mathcal{H}}$ and $q_{\mathcal{H}'}$ whose induced rankings are consistent with \mathcal{H} and \mathcal{H}' , respectively. We define the data source's on-path responses to $q_{\mathcal{H}}$ and $q_{\mathcal{H}'}$ to match the equilibrium responses guaranteed by Proposition 5.2. This preserves the best responses and incentive constraints (two bullets in Definition 2.5), so it yields an influential equilibrium under executable queries. Finally, allowing the user to submit any super-rank of q^{base} does not eliminate this equilibrium. We construct equilibrium strategies, the user strategy can place all probability on the two on-path queries $q_{\mathcal{H}}$ and $q_{\mathcal{H}'}$, while any other super-rank queries are assigned probability 0 and the data source's interpretation to them is specified off-path. Hence, an influential equilibrium exists under a super-rank strategy. □

PROOF OF PROPOSITION 5.15. Let m be the number of tuples in q^{base} and let $r = |R(I)|$ be the number of tuples that may be appended. The algorithm outputs exactly $2^{m-1} a(r)$ super-ranks, where $a(r) = \sum_{k=1}^r k! S(r, k)$ is the ordered Bell (Fubini) number. Consequently the running time is $\Theta(2^{m-1} a(r) (m+r))$ and is *super-exponential* in r (since $a(r) \geq r!$). We analyze the two phases of the algorithm separately and then combine them.

(*Merge-only count.*) The tuples of q^{base} form a total preorder with $m-1$ adjacent boundaries. Any merge-only coarsening corresponds to choosing which boundaries to delete; there are 2^{m-1} choices, each yielding a unique coarsening.

(*Append count.*) Appending r labeled items as a ranked/tied suffix is in bijection with weak orders (total preorders) on an r -set. The number of such weak orders is $a(r) = \sum_{k=1}^r k! S(r, k)$, where $S(r, k)$ counts partitions of the r items into k unlabeled blocks and the factor $k!$ orders these blocks into ranks. [28]

(*Product.*) The suffix is independent of the chosen merge-only coarsening, so the total number of super-ranks is the product $2^{m-1} a(r)$.

(*Running time.*) The algorithm generates each coarsening once (BFS over adjacent merges), and for each coarsening enumerates all $a(r)$ weak orders of the r appendable items. Materializing one output takes $O(m+r)$ time (copying/coalescing m tuples and writing the r -item suffix). Hence the time is $\Theta(2^{m-1} a(r) (m+r))$.

(*Super-exponential growth in r .*) Since $S(r, r) = 1$, the $k = r$ term gives $a(r) \geq r!$, while $a(r) \leq \sum_{k=1}^r k^r \leq r^{r+1}$. Thus $a(r)$ (and therefore the total) grows faster than any c^r for fixed $c > 1$, i.e., super-exponentially in r . □

PROOF OF THEOREM 5.16. We show NP-hardness by a Karp reduction from SUBSET-SUM.

SUBSET-SUM. Given positive integers w_1, \dots, w_n and a target W , decide whether there exists $S \subseteq \{1, \dots, n\}$ such that $\sum_{j \in S} w_j = W$.

Decision variant (MAXINF). Given a base ranking $q^{\text{min}} = e_1 \prec e_2 \prec \dots \prec e_m$, utilities U^s, U^r (specified below), and a threshold $T \in \mathbb{Z}$, decide whether there exists a *merge-only* super-rank query $q \in \mathcal{Q}_{\text{merge}}^*(q^{\text{min}})$ such that $U^r(q, \beta^*(q)) \geq T$. Here $\mathcal{Q}_{\text{merge}}^*(q^{\text{min}})$ denotes the set of super-rank queries obtained from q^{min} by merging any subset of adjacent boundaries (i.e., partitioning e_1, \dots, e_m into contiguous tie-blocks), and $\beta^*(q) = \arg \max_{\beta} U^s(q, \beta)$.

Since $\mathcal{Q}_{\text{merge}}^*(q^{\text{min}}) \subseteq \mathcal{Q}_{\text{super}}(q^{\text{min}})$, NP-hardness for MAXINF over $\mathcal{Q}_{\text{merge}}^*$ implies NP-hardness for the full super-rank problem.

Reduction. Given an instance (w_1, \dots, w_n, W) , set $m := n + 1$ and construct

$$q^{\text{min}} = e_1 \prec e_2 \prec \dots \prec e_m.$$

Any $q \in \mathcal{Q}_{\text{merge}}^*(q^{\text{min}})$ is uniquely encoded by a binary vector $b = (b_1, \dots, b_n) \in \{0, 1\}^n$, where $b_j = 0$ means the boundary between ranks j and $j+1$ is merged (tied), and $b_j = 1$ means it is kept as a cut.

Fix an arbitrary interpretation β^\dagger (e.g., a canonical interpretation). Define the data source utility to make β^\dagger the unique best response:

$$U^s(q, \beta) := \begin{cases} 1 & \text{if } \beta = \beta^\dagger, \\ 0 & \text{otherwise.} \end{cases}$$

Hence $\beta^*(q) = \beta^\dagger$ for all q .

Define the user utility (computable in polynomial time from b) by

$$U^r(q, \beta^\dagger) := - \left| \sum_{j: b_j=0} w_j - W \right|.$$

Finally set the threshold $T := 0$.

Correctness. For any merge-only super-rank query q encoded by b ,

$$U^r(q, \beta^*(q)) \geq 0 \iff - \left| \sum_{j: b_j=0} w_j - W \right| \geq 0 \iff \sum_{j: b_j=0} w_j = W.$$

Thus, letting $S = \{j : b_j = 0\}$, there exists q with $U^r(q, \beta^*(q)) \geq 0$ if and only if the SUBSET-SUM instance has a solution.

Complexity. The construction of q^{\min} and the explicit utilities takes $\text{poly}(n)$ time. Moreover, given b , evaluating U^r requires computing $\sum_{j: b_j=0} w_j$, which is $\text{poly}(n)$. Therefore SUBSET-SUM reduces to MAXINF in polynomial time, proving NP-hardness. \square

PROOF OF LEMMA 5.18. Fix $q' \in \mathcal{Q}_{\text{merge}}(q)$ and write $\Pi(q') = \{[i_1, j_1], \dots, [i_\ell, j_\ell]\}$ for its merged-interval partition. By additivity of U^r , the expected utility decomposes as a sum over the sets of tuples whose ranks lie in each interval $[i_\ell, j_\ell]$. Under the i.i.d. prior, tuples within the same merged interval are conditionally exchangeable, so each contributes the same expected per-tuple utility $\bar{u}_{q'}(i_\ell, j_\ell)$, yielding a contribution $n_{i_\ell: j_\ell} \bar{u}_{q'}(i_\ell, j_\ell)$ from that interval. Finally, merging intervals disjoint from $[i, j]$ only changes ties *within* those disjoint intervals and imposes no additional ordinal constraints involving tuples whose original ranks lie in $[i, j]$. Hence the posterior (and the data source's tuplewise posterior interpretation) for tuples in $[i, j]$ is the same under q' as under the single-interval merge query $q_{i:j}$. Therefore $\bar{u}_{q'}(i, j) = u_{ij}$ for every $[i, j] \in \Pi(q')$, and summing over the intervals gives the stated form. \square

PROOF OF PROPOSITION 5.19. Any optimal solution on the prefix $[1..j]$ ends with some last run $[i..j]$. By Lemma 5.18, its contribution is $\text{Score}(i, j)$, and the best possible contribution before it is $\text{OPT}(i-1)$. Maximizing over the choice of i yields the recurrence. Conversely, picking the maximum i for each j constructs an optimal partition. \square

REFERENCES

- [1] Leonidas Akritidis. 2020. Product Classification and Clustering. UCI Machine Learning Repository. DOI: <https://doi.org/10.24432/C5M91Z>.
- [2] Amazon. 2025. Amazon selling fees. <https://sell.amazon.com/pricing> Accessed: 2025-01-21.
- [3] Abolfazl Asudeh, H. V. Jagadish, Gerome Miklau, and Julia Stoyanovich. 2018. On Obtaining Stable Rankings. *CoRR* abs/1804.10990 (2018). arXiv:1804.10990 <http://arxiv.org/abs/1804.10990>
- [4] Matthew Backus, Thomas Blake, and Steven Tadelis. 2019. On the Empirical Content of Cheap-Talk Signaling: An Application to Bargaining. *Journal of Political Economy* 127, 4 (None 2019), 1599–1628. <https://doi.org/10.1086/701699>
- [5] Barry Becker and Ronny Kohavi. 1996. Adult. UCI Machine Learning Repository. DOI: <https://doi.org/10.24432/C5XW20>.
- [6] Shubham Bhatwal. 2021. Flight Price Prediction. Kaggle. <https://www.kaggle.com/datasets/shubhambathwal/flight-price-prediction>.
- [7] Andreas Blume, Ernest K. Lai, and Wooyoung Lim. 2020. Chapter 13 Strategic information transmission: a survey of experiments and theoretical foundations. In *Handbook of Experimental Game Theory*. Edward Elgar Publishing, Cheltenham, UK. <https://doi.org/10.4337/9781785363337.00022>
- [8] Archishman Chakraborty and Rick Harbaugh. 2007. Comparative cheap talk. *Journal of Economic Theory* 132, 1 (2007), 70–94. <https://doi.org/10.1016/j.jet.2005.07.003>
- [9] Vincent P. Crawford and Joel Sobel. 1982. Strategic Information Transmission. *Econometrica* 50, 6 (1982), 1431–1451. <http://www.jstor.org/stable/1913390>
- [10] H. A. David and H. N. Nagaraja. 2003. *Order Statistics* (3rd ed.). John Wiley and Sons, New York. <https://doi.org/10.1002/0471722162>
- [11] Aaron Frias. 2025. Amazon Products Dataset. Kaggle. <https://www.kaggle.com/datasets/aaronfrias/amazon-products-dataset/data>.
- [12] Sidney Fussell. 2019. Is Amazon's search engine biased? It's hard to tell. *Atlantic* (18 September 2019).
- [13] Ulrich Güntzer, Wolf-Tilo Balke, and Werner Kießling. 2000. Optimizing Multi-Feature Queries for Image Databases. In *Proceedings of the 26th International Conference on Very Large Data Bases (VLDB '00)*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 419–428.
- [14] Ihab F. Ilyas, George Beskales, and Mohamed A. Soliman. 2008. A survey of top-k query processing techniques in relational database systems. *ACM Comput. Surv.* 40, 4, Article 11 (Oct. 2008), 58 pages. <https://doi.org/10.1145/1391729.1391730>
- [15] Werner Kießling and Gerhard Köstler. 2002. Chapter 91 - Preference SQL – Design, Implementation, Experiences. In *VLDB '02: Proceedings of the 28th International Conference on Very Large Databases*, Philip A. Bernstein, Yannis E. Ioannidis, Raghu Ramakrishnan, and Dimitris Papadias (Eds.). Morgan Kaufmann, San Francisco, 990–1001. <https://doi.org/10.1016/B978-155860869-6/50098-6>
- [16] Jinyang Li, Yuval Moskovitch, and H. V. Jagadish. 2023. Detection of Groups with Biased Representation in Ranking. In *2023 IEEE 39th International Conference on Data Engineering (ICDE)*. IEEE Computer Society, Los Alamitos, CA, USA, 2167–2179. <https://doi.org/10.1109/ICDE55515.2023.00168>
- [17] Jinyang Li, Yuval Moskovitch, and H. V. Jagadish. 2023. Detection of Groups with Biased Representation in Ranking. arXiv:2301.00719 [cs.LG] <https://arxiv.org/abs/2301.00719>
- [18] Jinyang Li, Yuval Moskovitch, Julia Stoyanovich, and H. V. Jagadish. 2023. Query Refinement for Diversity Constraint Satisfaction. *Proc. VLDB Endow.* 17, 2 (Oct. 2023), 106–118. <https://doi.org/10.14778/3626292.3626295>
- [19] Lingxiao Li and Jia Xie. 2025. Comparative Cheap Talk: Evidence from the Housing Market. *International Real Estate Review* 28, 2 (None 2025), 117–148. <https://doi.org/None>
- [20] Lydia T. Liu, Nikhil Garg, and Christian Borgs. 2022. Strategic ranking. *Proceedings of Machine Learning Research* 151 (2022), 2489–2518. Publisher Copyright: Copyright © 2022 by the author(s); 25th International Conference on Artificial Intelligence and Statistics, AISTATS 2022 ; Conference date: 28-03-2022 Through 30-03-2022.
- [21] Michael Mathioudakis, Carlos Castillo, Giorgio Barnabò, and Sergio Celis. 2020. Affirmative Action Policies for Top-k Candidates Selection: With an Application to the Design of Policies for University Admissions. In *Proceedings of the 35th Annual ACM Symposium on Applied Computing*. Association for Computing Machinery, New York, NY, USA, 1640–1649. <https://doi.org/10.1145/3341105.3373878>
- [22] Dana Mattioli. 2019. Amazon Changed Search Algorithm in Ways That Boost Its Own Products. *Wall Street Journal* (16 September 2019). <https://www.wsj.com/articles/amazon-changed-search-algorithm-in-ways-that-boost-its-own-products-11568645345>
- [23] Niccolò Meneghetti, Denis Mindolin, Paolo Ciaccia, and Jan Chomicki. 2015. Output-sensitive Evaluation of Prioritized Skyline Queries. In *Proceedings of the 2015 ACM SIGMOD International Conference on Management of Data (Melbourne, Victoria, Australia) (SIGMOD '15)*. Association for Computing Machinery, New York, NY, USA, 1955–1967. <https://doi.org/10.1145/2723372.2723736>
- [24] Tripp Mickle. 2019. Apple Dominates App Store Search Results, Thwarting Competitors. *Wall Street Journal* (23 July 2019). <https://www.wsj.com/articles/apple-dominates-app-store-search-results-thwarting-competitors-11563897221>
- [25] MySQL. [n.d.]. MySQL. <https://dev.mysql.com/doc/refman/8.4/en/case.html>. Accessed: 2025-02-10.
- [26] Jack Nicas and Keith Collins. 2019. How Apple's apps topped rivals in the App Store it controls. *The New York Times* (9 September 2019). <https://www.nytimes.com/interactive/2019/09/09/technology/apple-app-store-competition.html>
- [27] Oracle. [n.d.]. Oracle. https://docs.oracle.com/cd/B13789_01/appdev.101/b10807/13_elements004.htm. Accessed: 2025-02-10.
- [28] B.C. Rennie and A.J. Dobson. 1969. On stirling numbers of the second kind. *Journal of Combinatorial Theory* 7, 2 (1969), 116–121. [https://doi.org/10.1016/S0021-9800\(69\)80045-1](https://doi.org/10.1016/S0021-9800(69)80045-1)
- [29] D. Kossmann S. Borzsony and K. Stocker. 2001. The Skyline Operator. In *Proceedings of the 17th International Conference on Data Engineering (ICDE '01)*. IEEE Computer Society, USA, 421.
- [30] C. Spearman. 1961. The Proof and Measurement of Association Between Two Things. In *Studies in individual differences: The search for intelligence*. Appleton-Century-Crofts, East Norwalk, CT, US, 45–58. <https://doi.org/10.1037/11491-005>
- [31] SQLServer. [n.d.]. SQLServer. <https://learn.microsoft.com/en-us/sql/t-sql/language-elements/case-transact-sql?view=sql-server-ver16>. Accessed: 2025-02-10.

- [32] Julia Stoyanovich, Serge Abiteboul, Bill Howe, H. V. Jagadish, and Sebastian Schelter. 2022. Responsible data management. *Commun. ACM* 65, 6 (May 2022), 64–74. <https://doi.org/10.1145/3488717>
- [33] Behzad Tabibian, Vicenç Gomez, Abir De, Bernhard Schölkopf, and Manuel Gomez-Rodriguez. 2020. On the design of consequential ranking algorithms. Jonas Peters and David Sontag (Eds.), Vol. 124. PMLR, 171–180. <http://proceedings.mlr.press>
- [34] Meike Zehlke, Francesco Bonchi, Carlos Castillo, Sara Hajian, Mohamed Megahed, and Ricardo Baeza-Yates. 2017. FA*IR: A Fair Top-k Ranking Algorithm. In *Proceedings of the 2017 ACM on Conference on Information and Knowledge Management (Singapore, Singapore) (CIKM '17)*. Association for Computing Machinery, New York, NY, USA, 1569–1578. <https://doi.org/10.1145/3132847.3132938>
- [35] Meike Zehlke and Carlos Castillo. 2020. Reducing Disparate Exposure in Ranking: A Learning To Rank Approach. In *Proceedings of The Web Conference 2020 (Taipei, Taiwan) (WWW '20)*. Association for Computing Machinery, New York, NY, USA, 2849–2855. <https://doi.org/10.1145/3366424.3380048>
- [36] Meike Zehlke, Ke Yang, and Julia Stoyanovich. 2022. Fairness in Ranking, Part I: Score-Based Ranking. *ACM Comput. Surv.* 55, 6, Article 118 (dec 2022), 36 pages. <https://doi.org/10.1145/3533379>
- [37] Meike Zehlke, Ke Yang, and Julia Stoyanovich. 2022. Fairness in Ranking, Part II: Learning-to-Rank and Recommender Systems. *ACM Comput. Surv.* 55, 6, Article 117 (dec 2022), 41 pages. <https://doi.org/10.1145/3533380>
- [38] Karen Zraik. 2024. Farmers Sue Over Deletion of Climate Data From Government Websites. *The New York Times* (25 February 2024). <https://www.nytimes.com/2025/02/24/climate/agriculture-farmer-website-data-lawsuit.html>