

What You Can Do with Coordinated Samples

Edith Cohen^{*†}

Haim Kaplan[†]

Abstract

Sample coordination, where similar instances have similar samples, was proposed by statisticians four decades ago as a way to maximize overlap in repeated surveys. Coordinated sampling had been since used for summarizing massive data sets.

The usefulness of a sampling scheme hinges on the scope and accuracy within which queries posed over the original data can be answered from the sample. We aim here to gain a fundamental understanding of the limits and potential of coordination. Our main result is a precise characterization, in terms of simple properties of the estimated function, of queries for which estimators with desirable properties exist. We consider unbiasedness, nonnegativity, finite variance, and bounded estimates.

Since generally a single estimator can not be optimal (minimize variance simultaneously) for all data, we propose *variance competitiveness*, which means that the expectation of the square on any data is not too far from the minimum one possible for the data. Surprisingly perhaps, we show how to construct, for any function for which an unbiased nonnegative estimator exists, a variance competitive estimator.¹

1 Introduction

Many data sources, including IP or Web traffic logs from different time periods or locations, measurement data, snapshots of data depositories that evolve over time, and document/feature and market-basket data, can be viewed as a collection of *instances*, where each instance is an assignment of numeric values from some set V to a set of items (the set of items is the same for different instances but the value of each item changes).

When the data is too massive to manipulate or even store in full, it is useful to obtain and work with a random sample of each instance. Two common sampling schemes are Poisson sampling (each item is sampled independently with probability that depends only on its value) and bottom- k (order) sampling. The samples are efficient to compute, also when instances are presented as streams or are distributed across multiple servers. It is convenient to specify these sampling schemes through a *rank function*, $r : [0, 1] \times V \rightarrow \mathbb{R}$, which maps seed-value pairs to a number $r(u, v)$ that is non-increasing with u and non-decreasing with v . For each item h we draw a *seed* $u(h) \sim U[0, 1]$ uniformly at random and compute the rank value $r(u(h), v(h))$, where $v(h)$ is the value of h . With Poisson sampling, item h is sampled $\iff r(u(h), v(h)) \geq T(h)$, where $T(h)$ are fixed thresholds, whereas a bottom- k sample includes the k items with highest ranks.² Poisson PPS samples (Probability Proportional to Size [26], where each item is included with probability proportion to its value) are obtained using the rank function $r(u, v) = v/u$ and a fixed $T(h)$ across items. Priority (sequential Poisson) samples [33, 20, 38] are bottom- k samples utilizing the PPS ranks

^{*}Microsoft Research, SVC edith@cohenwang.com

[†]The Blavatnik School of Computer Science, Tel Aviv University, Tel Aviv, Israel. haimk@cs.tau.ac.il

¹This is the full version of a RANDOM 2013 paper

²The term bottom- k is due to historic usage of the inverse rank function and lowest k ranks [35, 36, 11, 12, 13]

items:	1	2	3	4	5	6	7	8
Instance1:	1	0	4	1	0	2	3	1
Instance2:	3	2	1	0	2	3	1	0
PPS sampling probabilities for T=4 (sample of expected size 3):								
Instance1:	0.25	0.00	1.00	0.25	0.00	0.50	0.75	0.25
Instance2:	0.75	0.50	0.25	0.00	0.50	0.75	0.25	0.00

Figure 1: Two instances with 8 items and respective PPS sampling probabilities for threshold value 4, so item with value v is sampled with probability $\min\{1, v/4\}$. To obtain two coordinated PPS samples of the instances, we associate an independent $u(i) \sim U[0, 1]$ with each item $i \in [8]$. We then sample $i \in [8]$ in instance $h \in [2]$ if and only if $u(i) \leq v_h(i)/4$, where $v_h(i)$ is the value of i in instance h . When coordinating the samples this way, we make them as similar as possible. In the example, item 1 will always (for any drawing of seeds) be sampled in instance 2 if it is sampled in instance 1 and vice versa for item 7.

$r(u, v) = v/u$ and successive weighted sampling without replacement [35, 21, 11] corresponds to bottom- k samples with the rank function $r(u, v) = -v/\ln(u)$.

Samples of different instances are *coordinated* when the set of random seeds $u(h)$ is shared across instances. Scalable sharing of seeds when instances are dispersed in time or location is facilitated through random hash functions $u(h) \leftarrow H(h)$, where the only requirement for our purposes is uniformity and pairwise independence. Figure 1 contains an example data set of two instances and the PPS sampling probabilities of each item in each instance, and illustrates how to coordinate the samples. Note that the sample of one instance does not depend on values assumed in other instances, which is important for scalable deployment.

Why coordinate samples? Sample coordination was proposed in 1972 by Brewer, Early, and Joice [3], as a method to maximize overlap and therefore minimize overhead in repeated surveys [37, 34, 36]: The values of items change, and therefore there is a new set of PPS sampling probabilities. With coordination, the sample of the new instance is as similar as possible to the previous sample, and therefore the number of items that need to be surveyed again is minimized. Coordination was subsequently used to facilitate efficient processing of large data sets. Coordinated samples of instances are used as synopses which facilitate efficient estimation of multi-instance functions such as distinct counts (cardinality of set unions), sum of maxima, and similarity [5, 4, 7, 18, 31, 22, 23, 6, 19, 12, 2, 25, 13, 17]. Estimates obtained over coordinated samples are much more accurate than possible with independent samples. Used this way, coordinated sampling can be casted as a form of Locality Sensitive Hashing (LSH) [29, 24, 28]. Lastly, coordinated samples can sometimes be obtained much more efficiently than independent samples. One example is computing samples of the d -neighborhoods of all nodes in a graph [7, 10, 31, 11, 12, 8]. Similarity queries between neighborhoods are useful in the analysis of massive graph datasets such as social networks or Web graphs.

Our aim here is to study the potential and limitations of estimating multi-instance functions from coordinated samples of instances. The same set of samples can be used to estimate multiple queries, including queries that are specified only after sampling is performed. The sample may also be primarily used to estimate subset statistics over one instance at a time, such as sums, averages, and moments. While we focus here on queries that span multiple instances, we keep this in mind, and therefore do not aim for a sampling scheme optimized for a particular query (although some of our results can be applied this way), but rather, to optimize the estimator given the sampling scheme and query.

Sum aggregates: Most queries in the above examples can be casted as *sum aggregates* over selected items h of an *item function* $f(v)$ applied to the item weight tuple in different instances $v(h) = (v_1(h), v_2(h), \dots)$. In particular, distinct count (set union) is a sum aggregate of $OR(v)$, max-sum aggregates $\max(v) =$

$\max_i v_i$, min-sum aggregates $\min(\mathbf{v}) = \min_i v_i$, and L_p^p (p th power of L_p -difference) is the sum aggregate of the exponentiated range function $\text{RG}_p(\mathbf{v}) = |\max(\mathbf{v}) - \min(\mathbf{v})|^p$. In our example of Figure 1, L_2^2 of items [4] is $(1 - 3)^2 + (2 - 0)^2 + (4 - 1)^2 + (1 - 0)^2 = 18$, and is computed by summing the item function $\text{RG}_2(v_1, v_2) = (v_1 - v_2)^2$ over these items. The L_1 of items {1, 3} is $|1 - 3| + |4 - 1| = 5$, using the item function $\text{RG}(v_1, v_2) = |v_1 - v_2|$, and the max-sum of items {6, 7, 8} is $\max\{2, 3\} + \max\{3, 1\} + \max\{1, 0\} = 7$, which uses the item function $\max\{v_1, v_2\}$. Other queries in our examples which are not in sum aggregate form can be approximated well by sum aggregates: The Jaccard similarity is a ratio of min-sum and max-sum, and thus the ratio of min-sum estimate with small additive error and a max-sum estimate with small relative error [17] is a good Jaccard similarity estimator. The L_p difference is the p th root of L_p^p , and can be estimated well by taking the p th root of a good estimator for L_p^p , which is a sum aggregate of RG_p . When the domain of query results is nonnegative, as is the case in the examples, the common practice, which we follow, is to restrict the estimates, which often are plugged in instead of the exact result, to be nonnegative as well.

Sum estimators – one item at a time: To estimate a sum aggregate, we can use a linear estimator which is the sum of single-item estimators, estimating the item function $f(\mathbf{v}(h))$ for each selected item h . We refer to such an estimator as a *sum estimator*. When the single-item estimators are unbiased, from linearity of expectation, so is the sum estimate. When the single-item estimators are unbiased and sampling of different items is pairwise independent (respectively, negatively correlated, as with bottom- k sampling), the variance of the sum is (resp., at most) the sum of variances of the single-item estimators. Therefore, the relative error of the sum estimator decreases with the number of selected items we aggregate over. We emphasize that unbiasedness of the single-item estimators (together with pairwise independence or negative correlations between items) is critical for good estimates of the sum aggregate since a variance component that is due to bias “adds up” with aggregation whereas otherwise the relative error “cancels out” with aggregation.³ The Horvitz-Thompson (HT) estimator [27] is a classic sum estimator which is unbiased and nonnegative. To estimate $f(\mathbf{v})$, the HT estimator outputs 0 when the value is not sampled and the inverse-probability estimate $f(\mathbf{v})/p$ when the value is sampled, where p is the sampling probability. The HT estimator is applicable to some multi-instance functions [17].

From here on, we restrict our attention to estimating item functions $f(\mathbf{v}) \geq 0$ where each entry of \mathbf{v} is Poisson sampled and focus on unbiased and nonnegative estimators for $f(\mathbf{v})$. We provide the model in detail in Section 2. See Appendix B for further discussion on using sum estimators with bottom- k samples.

The challenge we address: Throughout the 40 year period in which coordination was used, estimators were developed in an ad-hoc manner, lacking a fundamental understanding of the potential and limits of the approach. Prior work was mostly based on adaptations of the HT estimator for multiple instances. The HT estimator is applicable provided that for any \mathbf{v} where $f(\mathbf{v}) > 0$, there is a positive probability for an outcome that both reveals $f(\mathbf{v})$ and allows us to determine a probability p for such an outcome. These conditions are satisfied by some basic functions including $\max(\mathbf{v})$ and $\min(\mathbf{v})$. There are functions, however, for which the HT estimator is not applicable, but nonetheless, for which nonnegative and unbiased estimators exist. Moreover, the HT estimator may not be optimal even when it is applicable.

As a particular example, the only L_p difference for which a “satisfactory” estimator was known was

³More concretely, the sample of any particular item is likely to have all or most entries missing, in which case, the variance of any nonnegative unbiased single-item estimator is $\Omega((1/p)f(\mathbf{v})^2)$, where p is the probability that the outcome reveals “enough” on $f(\mathbf{v})$. In such a case, a fixed 0 estimate (which is biased) will have variance $f(\mathbf{v})^2$, which is lower than the variance of any unbiased estimator when p is small, but also clearly useless and will result in large error for the sum aggregate. To understand this more precisely, let p be the probability that the outcome provides a lower bound of at least $f(\mathbf{v})/2$. To be nonnegative, the contribution to the expectation from the other $(1 - p)$ portion of outcomes can not exceed $f(\mathbf{v})/2$, so the remaining contribution of the p portion of outcomes must be at least $f(\mathbf{v})/2$, giving the lower bound on the variance.

the L_1 difference [17]. Stated in terms of estimators for item functions, prior to our work, there was no unbiased and nonnegative estimator known for $\text{RG}_p(\mathbf{v}) = |\max(\mathbf{v}) - \min(\mathbf{v})|^p$ for any $p \neq 1$. For $p = 1$, the known estimator used the relation $\text{RG}(\mathbf{v}) = \max(\mathbf{v}) - \min(\mathbf{v})$, separately estimating the maximum and the minimum and showing that when samples are coordinated then the estimate for the maximum is always at least as large as the one for the minimum and therefore the difference of the estimates is never negative. But even for this $\text{RG}(\mathbf{v})$ estimator, there was no understanding whether it is “optimal” and more so, what optimality even means in this context. Moreover, the ad hoc construction of this estimator does not extend even to slight variations, like $\max\{v_1 - v_2, 0\}$, which sum-aggregates to the natural “one sided” L_1 difference.

Contributions highlights

Characterization: We provide a complete characterization, in terms of simple properties of the function f and the sampling scheme parameters, of when estimators with the following combinations of properties exist for f :

- unbiasedness and nonnegativity.
- unbiasedness, nonnegativity, and *finite variances*, which means that for all \mathbf{v} , the variance given data \mathbf{v} is finite.
- unbiasedness, nonnegativity, and *bounded estimates*, which means that for each \mathbf{v} , there is an upper bound on all estimates that can be obtained when the data is \mathbf{v} . Bounded estimates implies finite variances, but not vice versa.

The J estimator Our characterization utilizes a construction of an estimator, which we call *the J estimator*, which we show has the following properties: The J estimator is unbiased and nonnegative if and only if an unbiased nonnegative estimator exists for f . The J estimator has a finite variance for data \mathbf{v} or is bounded for data \mathbf{v} if and only if an (nonnegative unbiased) estimator with the respective property for \mathbf{v} exists.

Variance competitiveness: Ideally, we would like to have a single estimator, which subject to desired properties of the estimator, minimizes the variance for all data. Such estimators are known in the statistics literature as UMVUE (uniform minimum variance unbiased) estimators [32]. Generally however, an (unbiased, nonnegative, linear) estimator with minimum variance on all data vectors [30] may not exist. Simple examples show that this is also the case for our coordinated sampling model even when restricting our attention to particular natural functions.

We are therefore introducing and aiming for a notion of *variance competitiveness*, which means that for *any* data vector, the variance of our estimator is not “too far” from the minimum variance possible for that vector by a nonnegative unbiased estimator. More precisely, an estimator is c -competitive if for all data \mathbf{v} , the expectation of its square is within a factor of c from the minimum possible for \mathbf{v} by an estimator that is unbiased and nonnegative on all data. Variance competitiveness bridges the gap between the ideal UMVUE estimators (which generally do not exist) and the practice of estimator selections with no “worst-case” guarantees.

v -optimality: To study competitiveness, we need to compare the variance on each data vector to the minimum possible, and to do so, we need to be able to express the “best possible” estimates. We say that an estimator is v -optimal if amongst all estimators that are unbiased and nonnegative on all data, it minimizes variance for the data \mathbf{v} . We express the v -optimal estimates, which are the values a v -optimal estimator assumes on outcomes that are consistent with data \mathbf{v} , and the respective v -optimal variance. We show that

the v -optimal estimates are uniquely defined (almost everywhere). The v -optimal estimates, however, are inconsistent for different data since as we mentioned earlier, it is not generally possible to obtain a single unbiased nonnegative estimator that minimizes variance for all data vectors. They do allow us, however, to analyse the competitiveness of estimators, and in particular, that of the J estimator.

Competitiveness of the J estimator: We show that the J estimator is competitive. In particular, this shows the powerful and perhaps surprising result that whenever for any particular data vector v there exists an estimator with finite variance that is nonnegative and unbiased on all data, then there is a *single* estimator, that for all data, has expectation of the square that is $O(1)$ of the minimum possible.

2 Coordinated Shared-Seed Sampling for a Single Item

Using the terminology in the introduction, we are now looking at a single item, its values $v = (v_1, \dots, v_r)$ in different instances, the sampling scheme projected on this one item, and estimating the item function f .

We denote that data domain by $\mathbf{V} \subseteq \mathbb{R}_{\geq 0}^r$. The *sampling scheme* is specified by non-decreasing continuous functions $\tau = (\tau_1, \dots, \tau_r)$ on $[0, 1]$ such that the infimum of the range of τ_i is at most $\inf_{v \in \mathbf{V}} v_i$.

The sampling of v is performed by drawing a random value $u \sim U[0, 1]$, which we refer to as the *seed*. We then include the i th entry in the outcome S if and only if v_i is at least $\tau_i(u)$:

$$i \in S \iff v_i \geq \tau_i(u).$$

We use the notation $S(u, v)$ for the outcome of the sampling scheme applied to data v and seed u .

Sampling is PPS if $\tau_i(u)$ are linear functions: there is a fixed vector τ^* such that $\tau_i(u) \equiv u\tau_i^*$, in which case entry i is included with probability $\min\{1, v_i/\tau_i^*\}$. Our use of the term PPS in this single-item context is compatible with PPS sampling of each instance i using threshold τ_i^* . In this case, we look at the “projected” sampling scheme on the i th entry of our item’s tuple.

Observe that our model assumes *weighted* sampling, where the probability that an entry is sampled depends (and is non-decreasing) with its value. Transiting briefly back to sampling of instances, weighted sampling results in more accurate estimation of quantities (such as averages of sums) where larger values contribute more. It is also important for boolean domains ($\mathbf{V} = \{0, 1\}^r$) when most items have 0 values and in this case, enables us to sample only “active” items.

We assume that the seed u and the functions τ are available to the estimator, and in particular, treat the seed as provided with the outcome. When an entry is sampled, we know its value and also can compute the probability that it is sampled. When an entry is not sampled, we know that its value is at most $\tau_i(u)$ and we can compute this upper bound from the seed u and the function τ_i . Putting this information together, for each outcome $S(u, v)$, we can define the set $V^*(S)$ of all data vectors consistent with the outcome. This set captures all the information we can glean from the sample on the data.

$$V^*(S) \equiv V^*(u, v) = \{z \mid \forall i \in [r], (i \in S \wedge z_i = v_i) \vee (i \notin S \wedge z_i < \tau_i(u))\}.$$

Structure of the set of outcomes. From the outcome, which is the set of sampled entries and the seed ρ , we can determine $V^*(u, v)$ also for all $u \geq \rho$. We also have that for all $u \geq \rho$ and $z \in V^*(\rho, v)$, $V^*(u, z) = V^*(u, v)$. Fixing v , the sets $V^*(u, v)$ are non-decreasing with u and the set S of sampled entries is non-increasing, meaning that $V^*(u, v) \subset V^*(\rho, v)$ and $S(u, v) \supset S(\rho, v)$ when $u < \rho$.

The containment order of the sets $V^*(S)$ is a tree-like partial order on outcomes. For two outcomes, the sets $V^*(S)$ are either disjoint, and unrelated in the containment order, or one is fully contained in another,

Figure 2: Illustration of the containment order on all possible outcomes $V^*(S)$. Example has data domain $\mathbf{V} = \{0, 1, 2\} \times \{0, 1, 2\}$ and seed mappings $\tau_1 = \tau_2 \equiv \tau$. The root of the tree corresponds to outcomes with $u \in (\tau^{-1}(2), 1]$. In this case, the outcome reveals no information on the data and $V^*(S)$ contains all vectors in \mathbf{V} . When $u \in (\tau^{-1}(1), \tau^{-1}(2)]$ the outcome identifies entries in the data that are equal to “2”. When $u \in (0, \tau^{-1}(1)]$, the outcome reveals the data vector.

and succeeds it in the containment order. The outcome $S(u, \mathbf{v})$ precedes $S(\rho, \mathbf{v})$ in the containment order if and only if $u > \rho$. When V is finite, the containment order is a tree order, as shown in Figure 2.

For \mathbf{z} and \mathbf{v} , the set of all u such that $\mathbf{z} \in S(u, \mathbf{v})$, if not empty, is a suffix of $(0, 1]$ that is open to the left.

Lemma 2.1

$$\begin{aligned} \forall \rho \in (0, 1] \forall \mathbf{v} \\ \mathbf{z} \in V^*(\rho, \mathbf{v}) \implies \exists \epsilon > 0, \forall x \in (\rho - \epsilon, 1], \mathbf{z} \in V^*(x, \mathbf{v}) \end{aligned}$$

Proof Correctness for all $x \in [\rho, 1]$ follows from the structure of the set of outcomes: Since $V^*(x, \mathbf{v}) \supset V^*(\rho, \mathbf{v})$ for all $x \geq \rho$ then $\mathbf{z} \in V^*(\rho, \mathbf{v}) \implies \mathbf{z} \in V^*(x, \mathbf{v})$.

Consider now the set S of entries that satisfy $v_i \geq \tau_i(\rho)$. Since $\mathbf{z} \in V^*(\rho, \mathbf{v})$, we have $\forall i \in S, z_i = v_i$ and $\forall i \notin S, \max\{z_i, v_i\} < \tau_i(\rho)$. Since τ_i is continuous and monotone, for all $i \notin S$, there must be an $\epsilon_i > 0$ such that $\tau_i(\rho - \epsilon_i) > \max\{z_i, v_i\}$. We now take $\epsilon = \min_{i \notin S} \epsilon_i$ to conclude the proof. ■

■

3 Estimators and properties

Let $f : \mathbf{V}$ be a function mapping \mathbf{V} to the nonnegative reals. An *estimator* \hat{f} of f is a numeric function applied to the outcome. We use the notation $\hat{f}(u, \mathbf{v}) \equiv \hat{f}(S(u, \mathbf{v}))$. On continuous domains, an estimator must be (Lebesgue) integrable. An estimator is *fully specified* for \mathbf{v} if specified on a set of outcomes that have probability 1 given data \mathbf{v} . Two estimators \hat{f}_1 and \hat{f}_2 are *equivalent* if for all data \mathbf{v} , $\hat{f}_1(u, \mathbf{v}) = \hat{f}_2(u, \mathbf{v})$ with probability 1.

An estimator \hat{f} is *nonnegative* if $\forall S, \hat{f}(S) \geq 0$ and is *unbiased* if $\forall \mathbf{v}, \mathbf{E}[\hat{f}|\mathbf{v}] = f(\mathbf{v})$. An estimator has *finite variance* on \mathbf{v} if $\int_0^1 \hat{f}(u, \mathbf{v})^2 du < \infty$ (the expectation of the square is finite) and is *bounded* on \mathbf{v} if $\sup_{u \in (0, 1]} \hat{f}(u, \mathbf{v}) < \infty$. If a nonnegative estimator is bounded on \mathbf{v} , it also has finite variance for \mathbf{v} . We say that an estimator is bounded or has finite variances if the respective property holds for all $\mathbf{v} \in \mathbf{V}$.

***v*-optimality.** We say that an unbiased and nonnegative estimator is *v-optimal*, that is, optimal with respect to a data vector \mathbf{v} , if it has minimum variance for \mathbf{v} . We refer to the estimates that a *v*-optimal estimator assumes on outcomes consistent on data \mathbf{v} as the *v-optimal estimates* and to the minimum variance attainable for \mathbf{v} as the *v-optimal variance*.

Variance competitiveness. An estimator \hat{f} is *c-competitive* if

$$\forall \mathbf{v}, \int_0^1 \left(\hat{f}(u, \mathbf{v}) \right)^2 du \leq c \inf_{\hat{f}'} \int_0^1 \left(\hat{f}'(u, \mathbf{v}) \right)^2 du,$$

where the infimum is over all unbiased nonnegative estimators \hat{f}' of f . For any unbiased estimator, the expectation of the square is closely related to the variance:

$$\text{VAR}[\hat{f}|\mathbf{v}] = \int_0^1 (\hat{f}(u, \mathbf{v}) - f(\mathbf{v}))^2 du = \int_0^1 \hat{f}(u, \mathbf{v})^2 du - f(\mathbf{v})^2 \quad (1)$$

When minimizing the expectation of the square, we also minimize the variance. Moreover, c -competitiveness means that

$$\forall \mathbf{v}, \text{VAR}[\hat{f}|\mathbf{v}] \leq c \inf_{\hat{f}'} \text{VAR}[\hat{f}'|\mathbf{v}] + (c - 1)f(\mathbf{v})^2 \quad (2)$$

for all data vectors \mathbf{v} for which a nonnegative unbiased estimator with finite variance on \mathbf{v} exists, the variance of the estimator is at most c times the \mathbf{v} -optimal variance plus an additive term of $(c - 1)$ times $f(\mathbf{v})^2$.

An important remark is due here. In the typical scenario, discussed in the introduction, the sample is likely to provide little or no information on $f(\mathbf{v})$, the variance is $\Omega(f(\mathbf{v})^2)$, and hence competitiveness as we defined it in terms of the expectation of the square translates to competitiveness of the variance. Otherwise, when for some data in the domain the sample is likely to reveal the value, it is not possible to obtain a universal competitiveness result in terms of variance. (One such example is RG on PPS samples, looking at data tuples where the maximum has sampling probability 1.) Interestingly, for RG_2 it is possible to get a bounded ratio in terms of variance. More details are in the companion experimental paper [16].

4 The lower bound function and its lower hull

For a function f , we define the respective *lower bound function* \underline{f} and the lower hull function H_f . We then characterize, in terms of properties of \underline{f} and H_f when nonnegative unbiased estimators exist for f and when such estimators exist that also have finite variances or are bounded.

The lower bound function $\underline{f}(S)$: For $Z \subset \mathbf{V}$, we define $\underline{f}(Z) = \inf\{f(v) \mid v \in Z\}$ to be the tightest lower bound on the values of f on Z . We use the notation $\underline{f}(S) \equiv \underline{f}(V^*(S))$, $\underline{f}(\rho, \mathbf{v}) \equiv \underline{f}(V^*(\rho, \mathbf{v}))$. When \mathbf{v} is fixed, we use $\underline{f}^{(\mathbf{v})}(u) \equiv \underline{f}(u, \mathbf{v})$.

From Lemma 2.1, we obtain that $\forall \mathbf{v}$, $\underline{f}^{(\mathbf{v})}(u)$ is *left-continuous*, that is:

Corollary 4.1 $\forall \mathbf{v} \forall \rho \in (0, 1]$, $\lim_{\eta \rightarrow \rho^-} \underline{f}^{(\mathbf{v})}(\eta) = \underline{f}^{(\mathbf{v})}(\rho)$.

Lemma 4.2 A nonnegative unbiased estimator \hat{f} must satisfy

$$\forall \mathbf{v}, \forall \rho, \int_{\rho}^1 \hat{f}(u, \mathbf{v}) du \leq \underline{f}^{(\mathbf{v})}(\rho) \quad (3)$$

Proof Unbiased and nonnegative \hat{f} must satisfy

$$\forall \mathbf{v}, \forall \rho \in (0, 1]. \int_{\rho}^1 \hat{f}(u, \mathbf{v}) du \leq \int_0^1 \hat{f}(u, \mathbf{v}) du = f(\mathbf{v}) . \quad (4)$$

From definition of \underline{f} , for all $\epsilon > 0$ and ρ , there is a vector $\mathbf{z}^{(\epsilon)} \in S(\rho, \mathbf{v})$ such that $f(\mathbf{z}^{(\epsilon)}) \leq \underline{f}(\rho, \mathbf{v}) + \epsilon$. Recall that for all $u \geq \rho$, $S(u, \mathbf{v}) = S(u, \mathbf{z}^{(\epsilon)})$, hence, using, (4),

$$\int_{\rho}^1 \hat{f}(u, \mathbf{v}) du = \int_{\rho}^1 \hat{f}(u, \mathbf{z}^{(\epsilon)}) du \leq f(\mathbf{z}^{(\epsilon)}) \leq \underline{f}(\rho, \mathbf{v}) + \epsilon .$$

Taking the limit as $\epsilon \rightarrow 0$ we obtain $\int_{\rho}^1 \hat{f}(u, \mathbf{v}) du \leq \underline{f}(\rho, \mathbf{v})$. ■

■

The lower hull of the lower bound function and v -optimality: We denote the function corresponding to the lower boundary of the convex hull (lower hull) of $\underline{f}^{(v)}$ by $H_f^{(v)}$. Our interest in the lower hull is due to the following relation (The proof is postponed to Section A):

Theorem 4.1 *An estimator \hat{f} is v -optimal if and only if for $u \in [0, 1]$ almost everywhere*

$$\hat{f}(u, \mathbf{v}) = -\frac{dH_f^{(v)}(u)}{du}.$$

Moreover, when an unbiased and nonnegative estimator exists for f , there also exists, for any data \mathbf{v} , a nonnegative and unbiased v -optimal estimator.

We use the notation $\hat{f}^{(v)}(u) = -\frac{dH_f^{(v)}(u)}{du}$ for the v -optimal estimates on outcomes consistent with \mathbf{v} . Since the lower bound function is monotone non-increasing, so is $H_f^{(v)}$, and therefore $H_f^{(v)}$ is differentiable almost everywhere and $\hat{f}^{(v)}$ is defined almost everywhere. Figure 3 illustrates an example lower bound function and the corresponding lower hull.

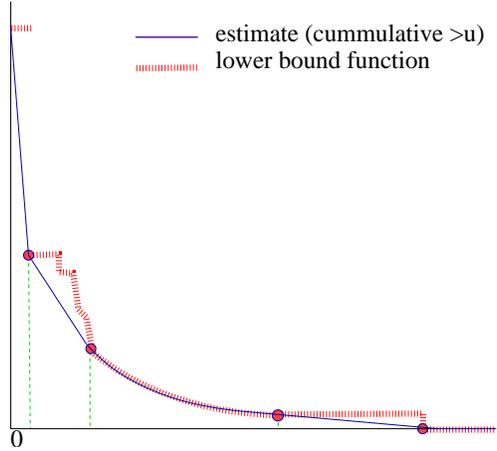


Figure 3: Lower bound function $\underline{f}^{(v)}(u)$ for $u \in (0, 1]$ and corresponding lower hull $H_f^{(v)}(u)$, which is also the integral of the nonnegative estimator with minimum variance on \mathbf{v} : $\int_u^1 \hat{f}^{(v)}(x) dx$. The figure visualizes the lower bound function which is always left-continuous and monotone non increasing. The lower hull is continuous and also monotone non-increasing.

5 Characterization

Theorem 5.1 *$f \geq 0$ has an estimator that is*

- *unbiased and nonnegative* \iff

$$\forall \mathbf{v} \in \mathbf{V}, \lim_{u \rightarrow 0^+} \underline{f}^{(\mathbf{v})}(u) = f(\mathbf{v}). \quad (5)$$

- *unbiased, nonnegative, and finite variances* \iff

$$\forall \mathbf{v} \in \mathbf{V}, \int_0^1 \left(\frac{dH_f^{(\mathbf{v})}(u)}{du} \right)^2 du < \infty. \quad (6)$$

- *unbiased, nonnegative, and bounded* \iff

$$\forall \mathbf{v} \in \mathbf{V}, \lim_{u \rightarrow 0^+} \frac{f(\mathbf{v}) - \underline{f}^{(\mathbf{v})}(u)}{u} < \infty. \quad (7)$$

We establish sufficiency in Theorem 5.1 by constructing an estimator $\hat{f}^{(J)}$ (the J estimator) that is unbiased and nonnegative when (5) holds, bounded when (7) holds, and has finite variances if (6) holds. The proof of the theorem is provided in Section 7, following the presentation of the J estimator in the next section.

6 The J estimator

Fixing \mathbf{v} , we define an estimator $\hat{f}^{(J)}(u, \mathbf{v})$ incrementally, starting with $u = 1$ and such that the value at u depends on values at $u' > u$. We first define $\hat{f}^{(J)}(u, \mathbf{v})$ for all $u \in (\frac{1}{2}, 1]$ by $\hat{f}^{(J)}(u, \mathbf{v}) = 2f(1, \mathbf{v})$. At each step we consider intervals of the form $(2^{-j-1}, 2^{-j}]$, setting the estimate to the same value for all outcomes $S(u, \mathbf{v})$ for $u \in (2^{-j-1}, 2^{-j}]$. Assuming the estimator is defined for $u \geq 2^{-j}$, we extend the definition to the interval $u \in (2^{-j-1}, 2^{-j}]$ as follows.

$$\begin{aligned} \hat{f}^{(J)}(u, \mathbf{v}) &= 0, \text{ if } \underline{f}(2^{-j}, \mathbf{v}) = \int_{2^{-j}}^1 \hat{f}^{(J)}(u, \mathbf{v}) du \\ \hat{f}^{(J)}(u, \mathbf{v}) &= 2^{j+1} \left(\underline{f}(2^{-j}, \mathbf{v}) - \int_{2^{-j}}^1 \hat{f}^{(J)}(u, \mathbf{v}) du \right), \text{ otherwise} \end{aligned}$$

Lemma 6.1 *The J estimator is well defined, is unbiased and nonnegative when (5) holds, and satisfies*

$$\forall \rho \forall \mathbf{v}, \int_{\rho}^1 \hat{f}^{(J)}(u, \mathbf{v}) du \leq \underline{f}(\rho, \mathbf{v}) \quad (8)$$

$$\forall \rho \forall \mathbf{v}, \int_{\rho}^1 \hat{f}^{(J)}(u, \mathbf{v}) du \geq \underline{f}(4\rho, \mathbf{v}). \quad (9)$$

Proof We first argue that the constructions, which are presented relative to a particular choices of the data \mathbf{v} , produce a consistent estimator. For that, we have to show that for every outcome $S(\rho, \mathbf{v})$, the assigned value is the same for all vectors $\mathbf{z} \in V^*(\rho, \mathbf{v})$. Since $\underline{f}(\rho, \mathbf{v}) = \underline{f}(\rho, \mathbf{z})$ for all $\mathbf{z} \in V^*(\rho, \mathbf{v})$, in particular this holds for $\rho = 2^{-j}$, so the setting of the estimator for $u \in (2^{-j-1}, 2^{-j}]$ is the same for all $S(u, \mathbf{z})$ where $\mathbf{z} \in V^*(2^{-j}, \mathbf{v})$ (also when $\mathbf{z} \notin V^*(2^{-j-1}, \mathbf{v})$). Therefore, the resulting estimator is consistently defined.

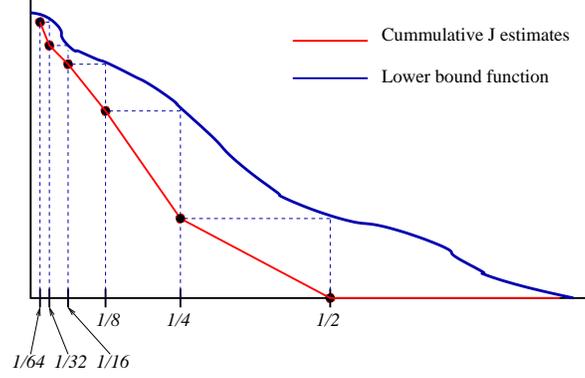


Figure 4: Lower bound function $\underline{f}^{(v)}(u)$ for $u \in (0, 1]$ and cumulative J estimates on outcomes consistent with $v \int_u^1 \hat{f}^{(J)}(x, v) dx$. The J estimate $\hat{f}^{(J)}(u, v)$ is the negated slope.

We show that the construction maintains the following invariant for $j \geq 1$:

$$\underline{f}(2^{-j+1}, v) = \int_{2^{-j}}^1 \hat{f}^{(J)}(u, v) du. \quad (10)$$

From the first step of the construction,

$$\int_{1/2}^1 \hat{f}^{(J)}(u, v) du = \int_{1/2}^1 2\underline{f}(1, v) du = \underline{f}(1, v).$$

So (10) holds for $j = 1$. Now we assume by induction that (10) holds for j and establish that it holds for $j + 1$. If $\underline{f}(2^{-j}, v) = \underline{f}(2^{-j+1}, v)$, then by the definition of the J estimator $\int_{2^{-j-1}}^{2^{-j}} \hat{f}^{(J)}(u, v) du = 0$ and we get

$$\int_{2^{-j-1}}^1 \hat{f}^{(J)}(u, v) du = \int_{2^{-j}}^1 \hat{f}^{(J)}(u, v) du = \underline{f}(2^{-j+1}, v) = \underline{f}(2^{-j}, v).$$

Otherwise, by definition, $\int_{2^{-j-1}}^{2^{-j}} \hat{f}^{(J)}(u, v) du = \underline{f}(2^{-j}, v) - \underline{f}(2^{-j+1}, v)$ and hence $\int_{2^{-j-1}}^1 \hat{f}^{(J)}(u, v) du = \underline{f}(2^{-j}, v)$ and (10) holds for $j + 1$.

From monotonicity, $\underline{f}(2^{-j+1}, v) \leq \underline{f}(2^{-j}, v)$ and when substituting (10) in the definition of the estimator we obtain that the estimates are always nonnegative.

To establish (8) we use (10), the relation $2^{\lfloor \log_2 \rho \rfloor} \leq \rho < 2^{1+\lfloor \log_2 \rho \rfloor}$, and monotonicity of $\underline{f}(u, v)$, to obtain

$$\int_{\rho}^1 \hat{f}^{(J)}(u, v) du \leq \int_{2^{\lfloor \log_2 \rho \rfloor}}^1 \hat{f}^{(J)}(u, v) du = \underline{f}(2^{1+\lfloor \log_2 \rho \rfloor}, v) \leq \underline{f}(\rho, v).$$

Similarly, we establish (9) using (10), and the relation $2^{-1+\lceil \log_2 \rho \rceil} \leq \rho \leq 2^{\lceil \log_2 \rho \rceil}$:

$$\int_{\rho}^1 \hat{f}^{(J)}(u, v) du \geq \int_{2^{\lceil \log_2 \rho \rceil}}^1 \hat{f}^{(J)}(u, v) du = \underline{f}(2^{1+\lceil \log_2 \rho \rceil}, v) \geq \underline{f}(4\rho, v)$$

Lastly, unbiasedness follows from (5) and combining (8) and (9):

$$\underline{f}(\rho, v) \geq \int_{\rho}^1 \hat{f}^{(J)}(u, v) du \geq \underline{f}(4\rho, v).$$

when we take the limit as $\rho \rightarrow 0$. ■

■

Computing the J estimate from an outcome S : From the outcome we know the seed value ρ and the lower bound function $\underline{f}^{(v)}(u)$ for all $u \geq \rho$ (recall that the lower bound on this range is the same for all data $\mathbf{v} \in V^*(S)$, so we do not need to know the data \mathbf{v}). We compute $i \leftarrow \lfloor -\log_2 \rho \rfloor$ and use the invariant (10) in the definition of J, obtaining the J estimate $2^{i+1}(\underline{f}(2^{-i}, \mathbf{v}) - \underline{f}(2^{-i+1}, \mathbf{v}))$.

Example: We demonstrate the application of the J estimator through a simple example. The data domain in our example includes pairs (v_1, v_2) of nonnegative reals. We are interested in $f(v_1, v_2) = (\max\{v_1 - v_2, 0\})^2$, which sum aggregates to (the square of) the “one sided” Euclidean distance. The data is PPS sampled with threshold $\tau = 1$ for both entries, therefore, the sampling probability of entry i is $\min\{1, v_i\}$. Sampling is coordinated, which means that for a seed $\rho \in U[0, 1]$, entry i is sampled if and only if $v_i \geq \rho$. The outcome S includes the values of the sampled entries and the seed value ρ . If no entry is sampled, or only the second entry is sampled, the lower bound function for $x \geq \rho$ is 0 and the J estimate is 0. If only the first entry is sampled, the lower bound function, for $x \geq \rho$, and accordingly, the J estimate are

$$\begin{aligned} \underline{f}^{(v)}(x) &= \max\{0, v_1 - x\}^2 \\ \hat{f}^{(J)}(S) &= 2^{\lfloor -\log_2 \rho \rfloor + 1} \left(\max\{0, v_1 - 2^{-\lfloor -\log_2 \rho \rfloor}\}^2 - \max\{0, v_1 - 2^{1-\lfloor -\log_2 \rho \rfloor}\}^2 \right). \end{aligned}$$

If both entries are sampled, the lower bound function for $x \geq \rho$ and the J estimate are

$$\begin{aligned} \underline{f}^{(v)}(x) &= \max\{0, v_1 - \max\{v_2, x\}\}^2 \\ \hat{f}^{(J)}(S) &= 2^{\lfloor -\log_2 \rho \rfloor + 1} \left(\max\{0, v_1 - \max\{v_2, 2^{-\lfloor -\log_2 \rho \rfloor}\}\}^2 - \max\{0, v_1 - \max\{v_2, 2^{1-\lfloor -\log_2 \rho \rfloor}\}\}^2 \right). \end{aligned}$$

6.1 Competitiveness of the J estimator

Theorem 6.1 *The estimator $\hat{f}^{(J)}$ is $O(1)$ -competitive.*

Proof We will show that

$$\forall \mathbf{v}, \int_0^1 \left(\hat{f}^{(J)}(u, \mathbf{v}) \right)^2 du \leq 84 \int_0^1 \left(\hat{f}^{(v)}(u) \right)^2 du.$$

Let $\rho = 2^{-j}$ for some integer $j \geq 0$. Recall the construction of $\hat{f}^{(J)}$ on an interval $(\rho/2, \rho]$. The value is fixed in the interval and is either 0, if $\int_\rho^1 \hat{f}^{(J)}(u, \mathbf{v}) du = \underline{f}^{(v)}(\rho)$ or is $2 \frac{\underline{f}^{(v)}(\rho) - \int_\rho^1 \hat{f}^{(J)}(u, \mathbf{v}) du}{\rho}$. Using (9) in Lemma 6.1, we obtain that for $u \in (\rho/2, \rho]$,

$$\hat{f}^{(J)}(u, \mathbf{v}) \leq 2 \frac{\underline{f}^{(v)}(\rho) - \int_\rho^1 \hat{f}^{(J)}(u, \mathbf{v}) du}{\rho} \leq 2 \frac{\underline{f}^{(v)}(\rho) - \underline{f}^{(v)}(4\rho)}{\rho}$$

Thus,

$$\int_{\rho/2}^\rho \left(\hat{f}^{(J)}(u, \mathbf{v}) \right)^2 du \leq \frac{\rho}{2} \frac{4}{\rho^2} \left(\underline{f}^{(v)}(\rho) - \underline{f}^{(v)}(4\rho) \right)^2 = \frac{2}{\rho} \left(\underline{f}^{(v)}(\rho) - \underline{f}^{(v)}(4\rho) \right)^2 \quad (11)$$

We now bound the expectation of $\hat{f}^{(v)}$ on $u \in (\rho/2, 4\rho)$ from below.

$$\begin{aligned} \int_{\rho/2}^{4\rho} \hat{f}^{(v)}(u) du &= \int_{\rho}^{4\rho} \hat{f}^{(v)}(u) du + \int_{\rho/2}^{\rho} \hat{f}^{(v)}(u) du \\ &\geq \int_{\rho}^{4\rho} \hat{f}^{(v)}(u) du + \frac{\rho}{2} \hat{f}^{(v)}(\rho) \end{aligned} \quad (12)$$

$$\geq \int_{\rho}^{4\rho} \hat{f}^{(v)}(u) du + \frac{\rho}{2} \frac{\underline{f}^{(v)}(\rho) - \int_{\rho}^1 \hat{f}^{(v)}(u) du}{\rho} \quad (13)$$

$$\begin{aligned} &= \int_{\rho}^{4\rho} \hat{f}^{(v)}(u) du + \frac{1}{2} \left(\underline{f}^{(v)}(\rho) - \int_{\rho}^{4\rho} \hat{f}^{(v)}(u) du - \int_{4\rho}^1 \hat{f}^{(v)}(u) du \right) \\ &= \frac{1}{2} \int_{\rho}^{4\rho} \hat{f}^{(v)}(u) du + \frac{1}{2} \left(\underline{f}^{(v)}(\rho) - \int_{4\rho}^1 \hat{f}^{(v)}(u) du \right) \geq \frac{1}{2} \left(\underline{f}^{(v)}(\rho) - \underline{f}^{(v)}(4\rho) \right) \end{aligned} \quad (14)$$

Inequality (12) follows from monotonicity of the function $\hat{f}^{(v)}$. Inequality (13) from the definition of $\hat{f}^{(v)}$ as the negated derivative of the lower hull of $\underline{f}^{(v)}$ $\hat{f}^{(v)}(\rho) \geq \frac{\underline{f}^{(v)}(\rho) - \int_{\rho}^1 \hat{f}^{(v)}(u) du}{\rho}$ More precisely, we can use the explicit definition of the v -optimal estimates (24) provided in Section A,

$$\hat{f}^{(v)}(\rho) = \inf_{0 \leq \eta < \rho} \frac{\underline{f}^{(v)}(\eta) - \int_{\rho}^1 \hat{f}(u, v) du}{\rho - \eta} \geq \inf_{0 \leq \eta < \rho} \frac{\underline{f}^{(v)}(\rho) - \int_{\rho}^1 \hat{f}(u, v) du}{\rho - \eta} = \frac{\underline{f}^{(v)}(\rho) - \int_{\rho}^1 \hat{f}(u, v) du}{\rho}.$$

Lastly, inequality (14) uses $\int_{4\rho}^1 \hat{f}^{(v)}(u) du \leq \underline{f}^{(v)}(4\rho)$, which follows from nonnegativity of $\hat{f}^{(v)}$ and Lemma 4.2.

Dividing both side by 3.5ρ we obtain a lower bound on the *average* value of $\hat{f}^{(v)}(u)$ in the interval $[\rho/2, 4\rho]$.

$$\frac{1}{3.5\rho} \int_{\rho/2}^{4\rho} \hat{f}^{(v)}(u) du \geq \frac{1}{7\rho} \left(\underline{f}^{(v)}(\rho, v) - \underline{f}^{(v)}(4\rho) \right) \quad (15)$$

We next show that the value $\hat{f}^{(J)}(u, v)$ on $u \in (\rho/2, \rho]$ is at most some constant times the expected value of the square of $\hat{f}^{(v)}$ on $u \in (\rho/2, 4\rho)$.

$$\int_{\rho/2}^{4\rho} \hat{f}^{(v)}(u)^2 du \geq \int_{\rho/2}^{4\rho} \left(\frac{1}{3.5\rho} \int_{\rho/2}^{4\rho} \hat{f}^{(v)}(u) du \right)^2 du \geq \quad (16)$$

$$\int_{\rho/2}^{4\rho} \left(\frac{1}{7\rho} \left(\underline{f}^{(v)}(\rho) - \underline{f}^{(v)}(4\rho) \right) \right)^2 du \geq \quad (17)$$

$$3.5\rho \left(\frac{1}{7\rho} \right)^2 \left(\underline{f}^{(v)}(\rho) - \underline{f}^{(v)}(4\rho) \right)^2 \geq \quad (18)$$

$$\frac{1}{28} \int_{\rho/2}^{\rho} \hat{f}^{(J)}(u, v)^2 du$$

Inequality (16) uses the fact that for any random variable X , $(\mathbb{E}[X])^2 \leq \mathbb{E}[X^2]$ applied to $\hat{f}^{(v)}$ for $u \in$

$(\rho/2, 4\rho]$. Inequality (17) follows from (15). Lastly, Inequality (18) follows from (11). We obtain

$$\begin{aligned} \int_0^1 \hat{f}^{(J)}(u)^2 du &= \sum_{i=0}^{\infty} \int_{2^{-i-1}}^{2^{-i}} \hat{f}^{(J)}(u)^2 du \\ &\leq 28 \sum_{i=0}^{\infty} \int_{2^{-i-1}}^{\min\{1, 2^{-i+2}\}} \hat{f}^{(v)}(u)^2 du \leq 28 \cdot 3 \int_0^1 \hat{f}^{(v)}(u)^2 du \end{aligned}$$

■

7 Proof of Theorem 5.1

Proof • “ \Rightarrow ” (5): From Lemma 4.2, an unbiased and nonnegative estimator \hat{f} must satisfy (3). Fixing \mathbf{v} in (3) and taking the limit as $\rho \rightarrow 0$ we obtain that $\mathbf{E}[\hat{f}|\mathbf{v}] = \int_0^1 \hat{f}(u, \mathbf{v}) du \leq \lim_{u \rightarrow 0} \underline{f}^{(v)}(u, \mathbf{v})$. Combining with unbiasedness: $\mathbf{E}[\hat{f}|\mathbf{v}] = f(\mathbf{v})$ we obtain (5).

• “ \Leftarrow ” (5): Follows immediately from Lemma 6.1.

• “ \Rightarrow ” (7): We bound from below the contribution to the expectation of unbiased and nonnegative \hat{f} of outcomes $S(u, \mathbf{v})$ for $u \leq \rho$: $\int_0^\rho \hat{f}(u, \mathbf{v}) = \int_0^1 \hat{f}(u, \mathbf{v}) - \int_\rho^1 \hat{f}(u, \mathbf{v}) \geq f(\mathbf{v}) - \underline{f}^{(v)}(u)$. The last inequality follows from unbiasedness and nonnegativity (3). Hence, the average value $\hat{f}(u, \mathbf{v})$ when $u < \rho$ must be at least $\frac{f(\mathbf{v}) - \underline{f}^{(v)}(\rho)}{\rho}$, and thus, considering all possible values of $\rho > 0$, we obtain that \hat{f} can be bounded only if it satisfies (7).

• “ \Leftarrow ” (7): Note that (7) \implies (5), and therefore the conditions of Lemma 6.1 are satisfied and the J estimator is well defined, nonnegative, and unbiased. It remains to show that given (7), or the equivalent statement

$$\forall \mathbf{v} \exists c < \infty \forall u, f(\mathbf{v}) - \underline{f}^{(v)}(u) \leq cu, \quad (19)$$

the J estimator is bounded. Fix \mathbf{v} and let c be as in (19).

$$\hat{f}^{(J)}(\rho, \mathbf{v}) \leq 2 \frac{f^{(v)}(\rho/2) - \int_{2\rho}^1 \hat{f}^{(J)}(u, \mathbf{v}) du}{\rho} \quad (20)$$

$$\leq 2 \frac{f(\mathbf{v}) - \underline{f}^{(v)}(8\rho)}{\rho} \quad (21)$$

$$= 16 \frac{f(\mathbf{v}) - \underline{f}^{(v)}(8\rho)}{8\rho} \leq 16c \quad (22)$$

Inequality (20) is from the definition of the J estimator. Inequality (21) uses definition of the lower bound function and (9). Lastly, (22) follows from our assumption (19).

• “ \iff ” (6): From Theorem 4.1, for all \mathbf{v} , (6), which is square-integrability of $\hat{f}^{(v)}(u)$, is *necessary* for existence of a nonnegative unbiased estimator with finite variance for \mathbf{v} . Sufficiency follows from the proof of Theorem 6.1, which shows that for all \mathbf{v} , the expectation of the square of the J estimator is at most a constant times the minimum possible, and (2), which states that the variance is bounded if and only if the

expectation of the square is bounded. ■

Conclusion

We developed a precise understanding of the queries we can estimate accurately over coordinated shared-seed samples. We defined variance competitiveness, and showed that $O(1)$ competitiveness is always attainable when estimators with finite variances exist. Our work uses a fresh, CS-inspired, and unified approach to the study of estimators that is particularly suitable for data analysis from samples and sets the ground for continued work.

In a follow up work [15], aiming for good performance in practice, we seek estimators that are *variance optimal (admissible [32])*, that is, can not be strictly improved. We show that there is a range of variance optimal and competitive estimators and that this choice can be leveraged to obtain natural estimators with desired properties or customized estimators to patterns in our data: The L^* estimator, which is defined for any function for which a nonnegative unbiased estimator exists, is the unique monotone⁴ variance-optimal estimator and is guaranteed to have a ratio of at most 4. Customization is facilitated by deriving *order-optimal estimators* which optimize variances with respect to (any) specified priorities over the data domain. The L^* estimator prioritizes data on which the estimated function is smaller whereas the U^* estimator prioritizes large values. We demonstrate the potential for applications in [16], for L_p difference estimation and in [9], for sketch-based similarity estimation in massive graphs.

A natural question is to bound the best possible competitive ratio. That is, the supremum over instances (data domain, shared-seed sampling scheme, and function) for which unbiased nonnegative estimators with finite variances exist of the best ratio attainable for the instance. The L^* estimator in [15] gives an upper bound of 4 and we were able to provide a tighter upper bound of 3.375 and a lower bound of 1.44. An interesting question is to close this gap.

In the future, we hope to build on our understanding of coordinated sampling to extend our treatment of independent sampling in [14]. We envision automated tools that provide estimates according to specifications of the sampling scheme, query, competitive ratio, and prioritization of patterns in the data.

References

- [1] N. Alon, N. Duffield, M. Thorup, and C. Lund. Estimating arbitrary subset sums with few probes. In *Proceedings of the 24th ACM Symposium on Principles of Database Systems*, pages 317–325, 2005.
- [2] K. S. Beyer, P. J. Haas, B. Reinwald, Y. Sismanis, and R. Gemulla. On synopses for distinct-value estimation under multiset operations. In *SIGMOD*, pages 199–210. ACM, 2007.
- [3] K. R. W. Brewer, L. J. Early, and S. F. Joyce. Selecting several samples from a single population. *Australian Journal of Statistics*, 14(3):231–239, 1972.
- [4] A. Z. Broder. On the resemblance and containment of documents. In *Proceedings of the Compression and Complexity of Sequences*, pages 21–29. IEEE, 1997.
- [5] A. Z. Broder. Identifying and filtering near-duplicate documents. In *Proceedings of the 11th Annual Symposium on Combinatorial Pattern Matching*, volume 1848 of *LNCS*, pages 1–10. Springer, 2000.

⁴Estimate value is non-increasing with information

- [6] J. W. Byers, J. Considine, M. Mitzenmacher, and S. Rost. Informed content delivery across adaptive overlay networks. *IEEE/ACM Trans. Netw.*, 12(5):767–780, October 2004.
- [7] E. Cohen. Size-estimation framework with applications to transitive closure and reachability. *J. Comput. System Sci.*, 55:441–453, 1997.
- [8] E. Cohen. All-distances sketches, revisited: Scalable estimation of the distance distribution and centralities in massive graphs. Technical Report cs.DS/1306.3284, arXiv, 2013.
- [9] E. Cohen, D. Delling, F. Fuchs, A. Goldberg, M. Goldszmidt, and R. Werneck. Scalable similarity estimation in social networks: closeness, node labels, and random edge lengths. In *COSN*, 2013.
- [10] E. Cohen and H. Kaplan. Spatially-decaying aggregation over a network: model and algorithms. *J. Comput. System Sci.*, 73:265–288, 2007.
- [11] E. Cohen and H. Kaplan. Summarizing data using bottom-k sketches. In *Proceedings of the ACM PODC’07 Conference*, 2007.
- [12] E. Cohen and H. Kaplan. Tighter estimation using bottom-k sketches. In *Proceedings of the 34th VLDB Conference*, 2008.
- [13] E. Cohen and H. Kaplan. Leveraging discarded samples for tighter estimation of multiple-set aggregates. In *Proceedings of the ACM SIGMETRICS’09 Conference*, 2009.
- [14] E. Cohen and H. Kaplan. Get the most out of your sample: Optimal unbiased estimators using partial information. In *Proc. of the 2011 ACM Symp. on Principles of Database Systems (PODS 2011)*. ACM, 2011. full version: <http://arxiv.org/abs/1203.4903>.
- [15] E. Cohen and H. Kaplan. A case for customizing estimators: Coordinated samples. Technical Report cs.ST/1212.0243, arXiv, 2012.
- [16] E. Cohen and H. Kaplan. How to estimate change from samples. Technical Report cs.DS/1203.4903, arXiv, 2012.
- [17] E. Cohen, H. Kaplan, and S. Sen. Coordinated weighted sampling for estimating aggregates over multiple weight assignments. *Proceedings of the VLDB Endowment*, 2(1–2), 2009. full version: <http://arxiv.org/abs/0906.4560>.
- [18] E. Cohen, Y.-M. Wang, and G. Suri. When piecewise determinism is almost true. In *Proc. Pacific Rim International Symposium on Fault-Tolerant Systems*, pages 66–71, December 1995.
- [19] A. Das, M. Datar, A. Garg, and S. Rajaram. Google news personalization: scalable online collaborative filtering. In *WWW*, 2007.
- [20] N. Duffield, M. Thorup, and C. Lund. Priority sampling for estimating arbitrary subset sums. *J. Assoc. Comput. Mach.*, 54(6), 2007.
- [21] P. S. Efraimidis and P. G. Spirakis. Weighted random sampling with a reservoir. *Inf. Process. Lett.*, 97(5):181–185, 2006.
- [22] P. Gibbons and S. Tirthapura. Estimating simple functions on the union of data streams. In *Proceedings of the 13th Annual ACM Symposium on Parallel Algorithms and Architectures*. ACM, 2001.
- [23] P. B. Gibbons. Distinct sampling for highly-accurate answers to distinct values queries and event reports. In *International Conference on Very Large Databases (VLDB)*, pages 541–550, 2001.
- [24] A. Gionis, P. Indyk, and R. Motwani. Similarity search in high dimensions via hashing. In *Proceedings of the 25th International Conference on Very Large Data Bases, VLDB ’99*, 1999.
- [25] M. Hadjieleftheriou, X. Yu, N. Koudas, and D. Srivastava. Hashed samples: Selectivity estimators for set similarity selection queries. In *Proceedings of the 34th VLDB Conference*, 2008.

- [26] J. Hájek. *Sampling from a finite population*. Marcel Dekker, New York, 1981.
- [27] D. G. Horvitz and D. J. Thompson. A generalization of sampling without replacement from a finite universe. *Journal of the American Statistical Association*, 47(260):663–685, 1952.
- [28] P. Indyk. Stable distributions, pseudorandom generators, embeddings and data stream computation. In *Proc. 41st IEEE Annual Symposium on Foundations of Computer Science*, pages 189–197. IEEE, 2001.
- [29] P. Indyk and R. Motwani. Approximate nearest neighbors: Towards removing the curse of dimensionality. In *Proc. 30th Annual ACM Symposium on Theory of Computing*, pages 604–613. ACM, 1998.
- [30] J. Lanke. On UMV-estimators in survey sampling. *Metrika*, 20(1):196–202, 1973.
- [31] D. Mosk-Aoyama and D. Shah. Computing separable functions via gossip. In *Proceedings of the ACM PODC'06 Conference*, 2006.
- [32] P. Mukhopandhyay. *Theory and Methods of Survey Sampling*. PHI learning, New Delhi, 2 edition, 2008.
- [33] E. Ohlsson. Sequential poisson sampling. *J. Official Statistics*, 14(2):149–162, 1998.
- [34] E. Ohlsson. Coordination of pps samples over time. In *The 2nd International Conference on Establishment Surveys*, pages 255–264. American Statistical Association, 2000.
- [35] B. Rosén. Asymptotic theory for successive sampling with varying probabilities without replacement, I. *The Annals of Mathematical Statistics*, 43(2):373–397, 1972.
- [36] B. Rosén. Asymptotic theory for order sampling. *J. Statistical Planning and Inference*, 62(2):135–158, 1997.
- [37] P. J. Saavedra. Fixed sample size pps approximations with a permanent random number. In *Proc. of the Section on Survey Research Methods, Alexandria VA*, pages 697–700. American Statistical Association, 1995.
- [38] M. Szegedy. The DLT priority sampling is essentially optimal. In *Proc. 38th Annual ACM Symposium on Theory of Computing*. ACM, 2006.

A Pointwise variance optimality

We work with *partial specification* of (nonnegative unbiased) estimators. A partial specification \hat{f} is defined over a subset \mathcal{S} of outcomes that is closed under the containment order $V^*(\mathcal{S})$:

$$\begin{aligned} \forall \mathbf{v} \exists \rho_{\mathbf{v}} \in [0, 1], S(u, \mathbf{v}) \in \mathcal{S} \text{ a.e. for } u > \rho_{\mathbf{v}} \wedge \\ S(u, \mathbf{v}) \notin \mathcal{S} \text{ a.e. for } u \leq \rho_{\mathbf{v}} . \end{aligned}$$

We also require that $\hat{f} : \mathcal{S} \geq 0$, that is, estimates are nonnegative when specified, and that

$$\forall \mathbf{v}, \rho_{\mathbf{v}} > 0 \implies \int_{\rho_{\mathbf{v}}}^1 \hat{f}(u, \mathbf{v}) du \leq f(\mathbf{v}) \quad (23a)$$

$$\forall \mathbf{v}, \rho_{\mathbf{v}} = 0 \implies \int_{\rho_{\mathbf{v}}}^1 \hat{f}(u, \mathbf{v}) du = f(\mathbf{v}) \quad (23b)$$

If $\rho_{\mathbf{v}} = 0$, we say that the estimator is *fully specified* for \mathbf{v} .

We show that a partial specification \hat{f} can always be extended to an unbiased nonnegative estimator:

Lemma A.1 *If f satisfies (5) (has a nonnegative unbiased estimator) and \hat{f} is partially specified then it can be extended to an unbiased nonnegative estimator.*

Proof The only if direction follows from Lemma 4.2. The “if” direction follows from a slight modification of the J estimator construction, in which we “start” to specify the estimator at $\rho_{\mathbf{v}}$ instead of 1. ■

Theorem A.1 *If \hat{f} is a partially specified estimator, then an extension of \hat{f} to the outcomes $S(\rho, \mathbf{v})$ for $\rho \in (0, \rho_{\mathbf{v}}]$ is a partially specified estimator and minimizes $\text{VAR}[\hat{f}|\mathbf{v}]$ (over all unbiased and nonnegative estimators which are extensions of the partial specification \hat{f}).*

$$\begin{aligned} \iff \text{for } u \in (0, \rho_{\mathbf{v}}] \text{ a.e.} \\ \hat{f}(u, \mathbf{v}) = \inf_{0 \leq \eta < u} \frac{f^{(\mathbf{v})}(\eta) - \int_{\eta}^1 \hat{f}(x, \mathbf{v}) dx}{u - \eta} . \end{aligned} \quad (24)$$

Moreover, a solution \hat{f} of (24) always exists and $H(u) = \int_u^1 \hat{f}(x, \mathbf{v}) dx$ for $u \in (0, \rho_{\mathbf{v}}]$ is unique.

Proof From Lemma A.1 an extended \hat{f} to outcomes $S(\eta, \mathbf{v})$ for $\eta \in (0, \rho_{\mathbf{v}}]$ is a partially specified estimator if and only if

$$\forall u \leq \rho_{\mathbf{v}} \int_u^1 \hat{f}(x, \mathbf{v}) dx \leq \underline{f}^{(\mathbf{v})}(u) \quad (25a)$$

$$\int_0^1 \hat{f}(x, \mathbf{v}) dx = f(\mathbf{v}) \quad (25b)$$

From the structure of the constraints (25), if we take a solution $\hat{f}(u, \mathbf{v})$ and shift “weight” to lower u values (decreasing \hat{f} on higher u values and increasing on lower ones so that the total contribution is the same) then the result is also a solution. If the shift results in decrease of higher values and increase of lower values, then the variance decreases.

Therefore, a minimum variance solution of (25) must be monotone nonincreasing with u almost everywhere, meaning, that monotonicity holds if we exclude from $(0, \rho_v]$ a set of points with zero measure.

$$\begin{aligned} \exists \text{zero measure } U^0 \subset (0, \rho_v] \forall u_1, u_2 \in (0, \rho_v] \setminus U_0, \\ u_1 < u_2 \implies \hat{f}(u_1, \mathbf{v}) \geq \hat{f}(u_2, \mathbf{v}) \end{aligned} \quad (25c)$$

Looking now only at extensions satisfying (25a), (25b), and (25c), we can decrease variance if we could "shift weight" from the higher estimates on lower u values to the lower estimates on higher u values. Such shifts, however, may violate (25a). Therefore, an estimator has minimum variance only if all such shifts result in violation, that is, after excluding a zero measure set, between any two points with strictly different estimate values there must be a point η such that $\int_{\eta}^1 \hat{f}(u, \mathbf{v}) du$ "touches" the lower bound at η . Formally,

$$\begin{aligned} \exists \text{ zero measure } U^0 \subset (0, \rho_v] \forall u_1, u_2 \in (0, \rho_v] \setminus U_0 \\ u_1 < u_2 \wedge \hat{f}(u_1, \mathbf{v}) > \hat{f}(u_2, \mathbf{v}) \implies \\ \exists \eta \in (u_1, u_2], \int_{\eta}^1 \hat{f}(u, \mathbf{v}) du = \lim_{u \rightarrow \eta^+} \underline{f}^{(\mathbf{v})}(u). \end{aligned} \quad (25d)$$

We established that (25) ((25a)- (25d)) are *necessary* for an extension to have minimum variance. We also argued that any extension satisfying (25a) and (25b) can be transformed to one satisfying (25c) and (25d) without increasing its variance. Therefore, if we can establish that (25) have a unique solution (up to equivalence), the solution must have minimum variance.

Geometrically, an extension \hat{f} satisfies either (25) or (24) if and only if the function $H(u) = \int_u^1 \hat{f}(x, \mathbf{v}) dx$ is the lower boundary of the convex hull of the lower bound function $\underline{f}^{(\mathbf{v})}(x)$ for $x \in (0, \rho_v]$ and the point $(\rho_v, \int_{\rho_v}^1 \hat{f}(u, \mathbf{v}) du)$. This implies existence and uniqueness of the solution and also that it has the desired properties.

We provide an algebraic proof. We start by exploring the structure which an extension of $\hat{f}(u, \mathbf{v})$ to $u \in (0, \rho_v]$, which satisfies (24), must have. We subsequently use the structure to establish existence and nonnegativity, uniqueness (*a.e.*), that it satisfies (25), and lastly, to show that any function that violates (24) must also violate (25). The combination of these properties implies the claim of the theorem.

Structure: We say that ρ is an *LB (lower bound) point* of $\hat{f}(u, \mathbf{v})$ if the infimum of the solution of (24) at ρ is achieved as $\eta \rightarrow \rho^-$, that is,

$$\hat{f}(\rho, \mathbf{v}) = \lim_{\eta \rightarrow \rho^-} \frac{\underline{f}^{(\mathbf{v})}(\eta) - \int_{\rho}^1 \hat{f}(u, \mathbf{v}) du}{\rho - \eta}.$$

Otherwise, if the infimum is approached when $\eta < \rho - \epsilon$ for some $\epsilon > 0$, we say that ρ is a *gap point*. Note that it is possible for ρ to be both an LB and a gap point if the infimum is approached at multiple places. If

$$\int_{\rho}^1 \hat{f}(u, \mathbf{v}) du < \lim_{\eta \rightarrow \rho^-} \underline{f}^{(\mathbf{v})}(\eta) = \underline{f}^{(\mathbf{v})}(\rho),$$

then ρ must be a gap point, and if equality holds, ρ can be either an LB or a gap point (or both). The equality follows using Corollary 4.1 (left continuity of $\underline{f}^{(\mathbf{v})}$).

We use this classification of points to partition $(0, \rho_v]$ into *LB* and *gap* subintervals of the form $(y, \rho]$, that is, open to the left and closed to the right.

LB subintervals are maximal subintervals containing exclusively LB points (which can double as gap points) that have the form $(y, \rho]$. From left continuity and monotonicity of $\underline{f}^{(\mathbf{v})}$, if ρ is an LB point and not a gap point then there is some $\epsilon > 0$ such that $(x - \epsilon, x]$ are also LB points. Thus all LB points that are not gap points must be part of LB subintervals and these subintervals are open to the left and closed to the right (the point y is a gap point which may double as an LB point). Alternatively, we can identify LB subintervals as maximal such that

$$\forall x \in (y, \rho], \int_x^1 \hat{f}(u, \mathbf{v}) du = \underline{f}^{(\mathbf{v})}(x).$$

Gap subintervals are maximal that satisfy

$$\forall x \in (y, \rho), \int_x^1 \hat{f}(u, \mathbf{v}) du < \lim_{u \rightarrow x^+} \underline{f}^{(\mathbf{v})}(u). \quad (26)$$

Note that a consecutive interval of gap points may consist of multiple back-to-back gap subintervals.

We can verify that the boundary points $b \in \{y, \rho\}$ of both LB and gap subintervals satisfy

$$\int_b^1 \hat{f}(u, \mathbf{v}) du = \lim_{x \rightarrow b^+} \underline{f}^{(\mathbf{v})}(x) \quad (27)$$

Visually, LB intervals are segments where $\int_u^1 \hat{f}(x, \mathbf{v}) dx$ “identifies” with the lower bound function whereas gap intervals are linear segments where it “skips” between two points where it touches (in the sense of limit from the right) the (left-continuous) lower bound function. Figure 3 illustrates a partition of an example lower bound function and its lower hull into gap and LB subintervals. From left to right, there are two gap subintervals, an LB subinterval, a gap subinterval, and finally a trivial LB subinterval (where the LB and estimates are 0).

Existence, nonnegativity, and (25a) :

For some $0 < \rho \leq \rho_v$, let $\hat{f}(u, \mathbf{v})$ be an extension which satisfies (24) and (25a) for all $u \in (\rho, \rho_v]$. We show that there exists $y < \rho$ such that the solution can be extended to $(y, \rho]$, that is, (24) and (25a) are satisfied also for $u \in (y, \rho]$.

Consider the solution of (24) at $u = \rho$. If the infimum is attained at $\eta < \rho$, let y be the infimum over points η that attain the infimum of (24) at ρ . We can extend the solution of (24) to the interval $(y, \rho]$, which is a gap interval (or the prefix of one). The solution is fixed throughout the interval:

$$\forall x \in (y, \rho], \hat{f}(x, \mathbf{v}) = \frac{\lim_{\eta \rightarrow y^+} \underline{f}^{(\mathbf{v})}(\eta) - \lim_{\eta \rightarrow \rho^+} \underline{f}^{(\mathbf{v})}(\eta)}{\rho - y}. \quad (28)$$

Assume now that the infimum of the solution of (24) at ρ is attained as $\eta \rightarrow \rho^-$ and not at any $\eta < \rho$. Let y be the supremum of points $x < \rho$

$$\inf_{0 \leq \eta < x} \frac{\underline{f}^{(\mathbf{v})}(\eta) - \lim_{u \rightarrow x^+} \underline{f}^{(\mathbf{v})}(u)}{x - \eta} < \lim_{\eta \rightarrow x^-} \frac{\underline{f}^{(\mathbf{v})}(\eta) - \lim_{u \rightarrow x^+} \underline{f}^{(\mathbf{v})}(u)}{x - \eta} \quad (29)$$

(We allow the rhs to be $+\infty$). We have $y < \rho$ and can extend the solution to $(y, \rho]$. The extension satisfies $\int_x^1 \hat{f}(u, \mathbf{v}) du = \underline{f}^{(\mathbf{v})}(x)$ and $(y, \rho]$ is an LB subinterval or a prefix of one. The actual estimator values on $(y, \rho]$ are the negated left derivative of $\underline{f}^{(\mathbf{v})}(u)$:

$$\forall x \in (y, \rho], \hat{f}(x, \mathbf{v}) = \lim_{\eta \rightarrow x^-} \frac{\underline{f}^{(\mathbf{v})}(\eta) - \underline{f}^{(\mathbf{v})}(x)}{x - \eta}$$

In both cases (where the extension corresponds to an LB or gap subinterval), the solution on $(y, \rho]$ exists, is nonnegative, and satisfies (25a) for $u \in (y, \rho]$.

Thus, starting from $u = \rho_v$, we can iterate the process and compute a solution on a suffix of $(0, \rho_v]$ that is partitioned into gap and LB intervals. The sum of sizes of intervals, however, may converge to a value $< \rho_v$ and thus the process may not converge to covering $(0, \rho_v]$. To establish existence, assume to the contrary that there is no solution that covers $(0, \rho_v]$. Let x be the infimum such that there are solutions on $(x, \rho_v]$ but there are no solutions for $(x - \epsilon, \rho_v]$ for all $\epsilon > 0$. Let $\hat{f}(u, \mathbf{v})$ be a solution on $u \in (x, \rho_v]$ and consider its partition to LB and gap intervals. Each interval that intersects $[x, x + \epsilon)$ must contain a left boundary point y_ϵ of some subinterval (it is possible that $y_\epsilon = x$). All boundary points satisfy $\int_{y_\epsilon}^1 \hat{f}(x, \mathbf{v}) \leq \underline{f}^{(\mathbf{v})}(y_\epsilon)$. We have

$$\begin{aligned} \int_x^1 \hat{f}(u, \mathbf{v}) du &= \lim_{\epsilon \rightarrow 0} \int_{y_\epsilon}^1 \hat{f}(x, \mathbf{v}) \leq \lim_{\epsilon \rightarrow 0} \underline{f}^{(\mathbf{v})}(y_\epsilon) \\ &= \lim_{\eta \rightarrow x^+} \underline{f}^{(\mathbf{v})}(\eta) \leq \underline{f}^{(\mathbf{v})}(x) \end{aligned}$$

The last inequality follows from monotonicity of $\underline{f}^{(\mathbf{v})}(u)$. Thus, we can apply our extension process starting from $\rho = x$ and extend the solution to $(y, x]$ for some $y < x$, contradicting our assumption and establishing existence.

Uniqueness: Assume there are two different solutions and let $x \in (0, \rho_{\mathbf{v}})$ be the supremum of values on which they differ. Consider one of the solutions. x can not be an interior point of a (gap or LB) subinterval because the estimator is determined on the subinterval from values on the right boundary, which is the same for both solutions. If x is a right boundary point, $\int_x^{\rho_{\mathbf{v}}} \hat{f}(u, \mathbf{v}) du$ is same for both functions and thus the left boundary and the solution on the interval are uniquely determined and must be the same for both functions, contradicting our assumption.

Satisfies (25b) (unbiasedness): From (5) (a sufficient necessary condition to existence of unbiased nonnegative estimator) and (27)

$$\int_0^1 \hat{f}(u, \mathbf{v}) du = \lim_{u \rightarrow 0^+} \underline{f}^{(\mathbf{v})}(u) = f(\mathbf{v}).$$

Satisfies (25c) and (25d): Within an LB subinterval the estimator must be nonincreasing, since otherwise we obtain a contradiction to the definition of an LB point. The fixed values on two consecutive gap intervals $(y, x]$ and $(x, \rho]$ must also be nonincreasing because otherwise the infimum of the solution at ρ could not be obtained at x since a strictly lower value is obtained at y .

The only solution of (25): Let $\hat{f}^{(\mathbf{v})}(u) \equiv \hat{f}(u, \mathbf{v})$ where \hat{f} is a solution of (24). Let \hat{f}' be another extension which satisfies (25a) and (25b) (and thus constitutes a partially specified estimator that is fully specified on \mathbf{v}) but is not equivalent to $\hat{f}^{(\mathbf{v})}$, that is, for some ρ

$$\int_{\rho}^{\rho_{\mathbf{v}}} \hat{f}^{(\mathbf{v})}(u) du \neq \int_{\rho}^{\rho_{\mathbf{v}}} \hat{f}'(u, \mathbf{v}) du .$$

We show that \hat{f}' must violate (25c) or (25d).

We first show that if

$$\int_{\rho}^{\rho_{\mathbf{v}}} \hat{f}^{(\mathbf{v})}(u) du < \int_{\rho}^{\rho_{\mathbf{v}}} \hat{f}'(u, \mathbf{v}) du \tag{30}$$

then monotonicity (25c) must be violated. Let $\bar{\rho}$ be the supremum of points satisfying (30). Then $\bar{\rho}$ must be a gap point and some neighborhood to its left must be gap points. Let $y < \bar{\rho}$ be the left boundary point of this gap interval. Recall that $\hat{f}^{(\mathbf{v})}$ is constant on a gap interval.

Let $\rho \in (y, \bar{\rho})$ be such that (30) holds. There must be measurable set in $[\rho, \bar{\rho}]$ on which $\hat{f}' > \hat{f}^{(\mathbf{v})}$. Since y is a left boundary of a gap interval, from (27),

$$\int_y^1 \hat{f}^{(\mathbf{v})}(u) du = \lim_{u \rightarrow y^+} \underline{f}^{(\mathbf{v})}(u)$$

and therefore, since \hat{f}' satisfies (25a), $\int_y^{\rho_{\mathbf{v}}} \hat{f}' du \leq \int_y^{\rho_{\mathbf{v}}} \hat{f}^{(\mathbf{v})}(u) du$. Thus, there must exist a measurable subset of $[y, \rho]$ on which $\hat{f}' < \hat{f}^{(\mathbf{v})}$. Since $\hat{f}^{(\mathbf{v})}$ is constant in $(y, \bar{\rho}]$, \hat{f}' violates (25c).

Lastly, we show that if

$$\int_{\rho}^{\rho_{\mathbf{v}}} \hat{f}^{(\mathbf{v})}(u) du > \int_{\rho}^{\rho_{\mathbf{v}}} \hat{f}'(u, \mathbf{v}) du \tag{31}$$

then even if (25c) holds then (25d) is violated.

There must exist a measurable subset $U \subset [\rho, \rho_{\mathbf{v}}]$ on which $\hat{f}^{(\mathbf{v})}(u) > \hat{f}'(u, \mathbf{v})$.

Let η be the supremum of points $x \in [0, \rho]$ satisfying $\int_x^{\rho_{\mathbf{v}}} \hat{f}'(u, \mathbf{v}) du = \lim_{u \rightarrow x^+} \underline{f}^{(\mathbf{v})}(u)$. Since this is satisfied by $x = 0$ and the supremum can not be satisfied by ρ , such $\eta < \rho$ is well defined.

Since $\hat{f}^{(\mathbf{v})}$ satisfies (25a), $\int_{\eta}^{\rho_{\mathbf{v}}} \hat{f}^{(\mathbf{v})}(u) du \leq \int_{\eta}^{\rho_{\mathbf{v}}} \hat{f}'(u, \mathbf{v}) du$. Therefore, there must be another measurable set $W \subset [\eta, \rho]$ on which $\hat{f}^{(\mathbf{v})}(u) < \hat{f}'(u, \mathbf{v})$. Because $\hat{f}^{(\mathbf{v})}$ is monotone, the values of \hat{f}' on W are strictly larger than on U . Hence, \hat{f}' violates (25d). ■

Theorem 4.1 follows by applying Theorem A.1 to an empty specification.

B Using sum estimators with bottom- k sampling

We provide details on the application of sum estimators with bottom- k sampling of instances. A technical requirement for applying sum estimators is that the sampling scheme of each item is explicit. This naturally happens with Poisson sampling where an entry (item h in instance i) is sampled $\iff r(u(h), v(h)) \geq T_i(h)$. With bottom- k sampling, however, item inclusions are dependent and therefore the separation a bit more technical, and relies on a technique named rank conditioning [12, 1].

Rank conditioning was applied to estimate subset sums over bottom- k samples of a single instance, and allows us to treat each entry as Poisson sampled. In rough details, for each entry, we obtain a “ $T_i(h)$ substitute” which allows the inclusion of h in the sample of instance i on seeds (and thus ranks) of all other items being fixed. With this conditioning, the inclusion of h in the sample of i is Poisson with threshold T_i , that is, follows the rule $r(u(h), v(h)) \geq T_i$, where T_i is equal to the k th largest rank value of items in instance i with h excluded which is the same as the $(k + 1)$ st largest rank value.