

Cumulative differences between paired samples

Isabel Kloumann, Hannah Korevaar, Chris McConnell, Mark Tygert, and Jessica Zhao

Meta Platforms, Inc.

Abstract

The simplest, most common paired samples consist of observations from two populations, with each observed response from one population corresponding to an observed response from the other population at the same value of an ordinal covariate. The pair of observed responses (one from each population) at the same value of the covariate is known as a “matched pair” (with the matching based on the value of the covariate). A graph of cumulative differences between the two populations reveals differences in responses as a function of the covariate. Indeed, the slope of the secant line connecting two points on the graph becomes the average difference over the wide interval of values of the covariate between the two points; that is, slope of the graph is the average difference in responses. (“Average” refers to the weighted average if the samples are weighted.) Moreover, a simple statistic known as the Kuiper metric summarizes into a single scalar the overall differences over all values of the covariate. The Kuiper metric is the absolute value of the total difference in responses between the two populations, totaled over the interval of values of the covariate for which the absolute value of the total is greatest. The total should be normalized such that it becomes the (weighted) average over all values of the covariate when the interval over which the total is taken is the entire range of the covariate (that is, the sum for the total gets divided by the total number of observations, if the samples are unweighted, or divided by the total weight, if the samples are weighted). This cumulative approach is fully nonparametric and uniquely defined (with only one right way to construct the graphs and scalar summary statistics), unlike traditional methods such as reliability diagrams or parametric or semi-parametric regressions, which typically obscure significant differences on account of their parameter settings.

1 Introduction

Analyzing the differences between two populations often involves data of paired observations: each response and corresponding value of an ordinal covariate, both observed for one of the populations, are accompanied by a response from the other population for the same value of the covariate. Such a pair of responses at the same value for the covariate is known as a “matched pair.” This paper presents graphical methods and scalar summary statistics for detecting differences. The graphs reveal differences as a function of the covariate. The summary statistics condense the graphs into scalar metrics measuring the overall differences across all values of the covariate. Appendix A reviews methodology from Tygert (2021a) for extending the analysis from a single scalar covariate to multiple covariates (the earlier work concerned unpaired samples, but turns out to apply just as well to the case of paired samples treated in the present paper). Matched pairs are common both in observational studies, as reviewed by Rubin (1973), Caliendo & Kopeinig (2008), Austin (2011), and others, and in designed experiments such as randomized controlled trials, as reviewed by Section 3.3 of Cox & Reid (2000), Bai *et al.* (2022), and others.

A classic method for analyzing paired samples is to calculate the weighted average difference in responses between the populations; see, for example, Chapter 4, “Analysis of matched pairs using sample means,” of Brown & Hollander (2007). (The weighted average is just the usual arithmetic average if the samples are unweighted or uniformly weighted.) The methods of the present paper are an extension. Specifically, the scalar metric proposed below that generalizes the classical statistic of Kuiper (1962) is always greater than or equal to the absolute value of the weighted average difference in responses, so can detect differences to which the weighted average alone is blind. The proposed Kuiper metric is simply the absolute value of the total weighted difference in responses, totaled over the interval of values of the covariate for which the total has the greatest magnitude. Of course, if the interval of values of the covariate for which the total has the

greatest magnitude is the entire range of values of the covariate, then the Kuiper metric is the same as the absolute value of the weighted average difference. (The totals taken here should be normalized by the total number of observations, in the case that the observations are unweighted, or by the total weight, in the case that the samples are weighted.) Using the interval for which the total has the greatest magnitude avoids cancellation between positive and negative differences that can occur in the weighted average over the entire range of values of the covariate. Such cancellation can arise when differences at high values of the covariate trade off against differences at low values of the covariate, for instance.

The methods of the present paper require the covariate to be ordinal, that is, the values of the covariate must have a natural ordering. The ordering need not be perfect, but the methods require the ordering to be better than purely random. The weighted average difference in responses ignores the ordering, as do several other classical methods, such as the sign test or the signed-rank test of Wilcoxon (1945); details are available, for example, from Hollander *et al.* (2014) or Chapter 12, “Analysis of matched pairs using ranks,” of Brown & Hollander (2007).

Classic methods that can take advantage of the ordering of the values for the ordinal covariate include regression with parametric models, as reviewed, for example, by Draper & Smith (1998) or Chapter 7, “Logistic regression for matched case-control studies,” of Hosmer *et al.* (2013). In contrast, the methods of the present paper are fully nonparametric. Parametric models tend to be more statistically powerful than nonparametric methods when the models are known to be valid, true, and correct. However, the nonparametric methods of the present paper are always valid and have no parameters to tune, making the methods especially easy to use and trust, with no possibility for fudging results (whether intentionally or inadvertently). The methods of the present paper are thus complementary to parametric regression.

Section 3 below extensively compares the nonparametric methods of Section 2 to what are possibly the most popular among all classical methods that take advantage of the ordering: the traditional semi-parametric methods known as “reliability diagrams” or “calibration plots” reviewed in Subsection 2.5. Reliability diagrams histogram the responses of both populations separately, partitioning the values of the covariate into “bins” (also known as “buckets”). Being semi-parametric, reliability diagrams are subject to less manipulation and outright fudging than parametric regression. Yet the comparisons of Section 3 illustrate many problems with reliability diagrams relative to the proposed cumulative graphs. The traditional approach depends on the choice of bins, and that choice often turns out to influence the inferred results dramatically. The cumulative approach of Section 2 below avoids the problems. Section 3 demonstrates the advantages over the simple weighted average discussed above, too.

The advantages of the cumulative approach have been pointed out earlier, including by Ibarra *et al.* (2022) and Tygert (2021a,b,c, 2022). The present paper merely realizes the advantages of the cumulative statistics for paired samples, whereas these earlier works focused on unpaired data. Much earlier, the work of Delgado (1993) introduced an important, archetypal special case of part of what the present paper proposes; namely, Delgado (1993) investigated the Kolmogorov-Smirnov scalar summary statistic for paired samples.

The remainder of the paper has the following structure: Section 2 develops the cumulative methods and reviews the traditional reliability diagrams in Subsection 2.5. Section 3 illustrates the performance of the methods on both synthetic and measured data.¹ Section 4 draws some conclusions. Appendix A reviews the reduction via space-filling curves of the analysis with multiple covariates to the simpler case of a single scalar covariate.

2 Methods

This section presents the methodology of the present paper. Subsection 2.1 sets notational conventions used throughout. Subsection 2.2 constructs cumulative graphical methods for analyzing the differences in responses between two populations, controlled for a scalar score. Subsection 2.3 summarizes the graphs of Subsection 2.2 into scalar statistics. Subsection 2.4 discusses the statistical significance of the graphs of Subsection 2.2 and of the metrics of Subsection 2.3. Subsection 2.5 reviews the classical graphical methods known as “reliability diagrams.” The following section, Section 3, provides several examples of the methods of the present section applied to various data sets.

¹Permissively licensed open-source software implementing the methods in Python — software that also reproduces all figures and statistics reported below — is available at <https://github.com/facebookresearch/metapaired>

Appendix A reviews previously developed methods that enable the direct application of the methods of the present section to the analysis of differences in responses while controlling for multiple covariates (rather than for only a single scalar covariate). The reduction of the problem of conditioning on multiple covariates to the problem of conditioning on a single scalar score relies on space-filling curves. The reduction decouples from the methodology for analyzing the case of a single scalar covariate; Appendix A focuses on the reduction, while the present section focuses on the scalar case, specifically for data sets with paired samples (whereas earlier works studied only unpaired data).

2.1 Notation

We consider real numbers S_1, S_2, \dots, S_m , called “scores,” that are the values of a covariate. We assume without loss of generality that $S_1 < S_2 < \dots < S_m$. For each score, say score S_j , we consider real numbers $R_j^{(1)}, R_j^{(1)}, \dots, R_j^{(n_j)}$ and $Q_j^{(1)}, Q_j^{(1)}, \dots, Q_j^{(n_j)}$, called “responses,” that are the values of outcomes for two populations. Thus, the total sample size is

$$n = \sum_{j=1}^m n_j, \quad (1)$$

where again n_j is the number of pairs of responses $(Q_j^{(1)}, R_j^{(1)}), (Q_j^{(2)}, R_j^{(2)}), \dots, (Q_j^{(n_j)}, R_j^{(n_j)})$ associated with score S_j . For each score, say score S_j , we also consider positive real numbers $W_j^{(1)}, W_j^{(2)}, \dots, W_j^{(n_j)}$, called “weights,” that weight the contribution of the corresponding pairs $(Q_j^{(1)}, R_j^{(1)}), (Q_j^{(2)}, R_j^{(2)}), \dots, (Q_j^{(n_j)}, R_j^{(n_j)})$ associated with score S_j . If the samples are unweighted, then we set $W_j^{(k)}$ to be a constant independent of the values of the indices j and k .

We assume that the scores and weights are not random and that the responses are random and all stochastically independent.

For $j = 1, 2, \dots, m$, the weighted average responses at score S_j are

$$\tilde{R}_j = \frac{\sum_{k=1}^{n_j} R_j^{(k)} W_j^{(k)}}{\sum_{k=1}^{n_j} W_j^{(k)}} \quad (2)$$

and

$$\tilde{Q}_j = \frac{\sum_{k=1}^{n_j} Q_j^{(k)} W_j^{(k)}}{\sum_{k=1}^{n_j} W_j^{(k)}}, \quad (3)$$

and the corresponding total weight is

$$\tilde{W}_j = \sum_{k=1}^{n_j} W_j^{(k)}. \quad (4)$$

We form the cumulative differences

$$C_k = \frac{\sum_{j=1}^k (\tilde{Q}_j - \tilde{R}_j) \tilde{W}_j}{\sum_{j=1}^m \tilde{W}_j} \quad (5)$$

for $k = 1, 2, \dots, m$.

2.2 Graphs

We define abscissae to be the accumulated weights

$$A_k = \frac{\sum_{j=1}^k \tilde{W}_j}{\sum_{j=1}^m \tilde{W}_j} \quad (6)$$

for $k = 1, 2, \dots, m$.

In a plot of C_k versus A_k for $k = 1, 2, \dots, m$, the expected value of the slope from $k = j - 1$ to $k = j$ is

$$\mathbb{E} \left[\frac{C_j - C_{j-1}}{A_j - A_{j-1}} \right] = \mathbb{E} [\tilde{Q}_j - \tilde{R}_j], \quad (7)$$

which is simply the expected difference in the weighted average responses corresponding to score S_j . Thus, *the slope of a secant line connecting two points on the graph of C_k versus A_k becomes the weighted average difference in responses between the two populations over the long range of scores between the two points.*

2.3 Scalar metrics

We define the Kuiper metric to be the range of the graph,

$$D = \max_{0 \leq j \leq m} C_j - \min_{0 \leq j \leq m} C_j, \quad (8)$$

where $C_0 = 0$ and C_j is defined in (5) for $j = 1, 2, \dots, m$. A mathematically equivalent definition of the Kuiper metric is

$$D = \max_{1 \leq i \leq k \leq m} \left| \frac{\sum_{j=i}^k (\tilde{Q}_j - \tilde{R}_j) \tilde{W}_j}{\sum_{j=1}^m \tilde{W}_j} \right|, \quad (9)$$

that is, *the Kuiper distance D is the absolute value of the total weighted difference between the two subpopulations' responses, with the total taken over the interval of indices (or scores) for which the magnitude of the total is greatest.*

Another popular statistic is the Kolmogorov-Smirnov metric,

$$E = \max_{1 \leq j \leq m} |C_j|, \quad (10)$$

which unfortunately lacks the intuitive interpretation of the Kuiper metric that (9) expresses.

2.4 Statistical significance

Under the null hypothesis that the expected value of \tilde{Q}_j is equal to the expected value of \tilde{R}_j for all $j = 1, 2, \dots, m$, the sequence C_1, C_2, \dots, C_m is a driftless random walk with variance for the increment $C_j - C_{j-1}$ being the expected value of $(\tilde{Q}_j - \tilde{R}_j)^2 (\tilde{W}_j)^2 / (\sum_{k=1}^m \tilde{W}_k)^2$. The expected variance of C_m is then

$$\sum_{j=1}^m \mathbb{E} \left[\frac{(\tilde{Q}_j - \tilde{R}_j)^2 (\tilde{W}_j)^2}{(\sum_{k=1}^m \tilde{W}_k)^2} \right] = \mathbb{E} \left[\sum_{j=1}^m \frac{(\tilde{Q}_j - \tilde{R}_j)^2 (\tilde{W}_j)^2}{(\sum_{k=1}^m \tilde{W}_k)^2} \right]. \quad (11)$$

The right-hand side of (11) shows that an unbiased estimator for the variance of C_m based on the observations is

$$\sigma^2 = \sum_{j=1}^m \frac{(\tilde{Q}_j - \tilde{R}_j)^2 (\tilde{W}_j)^2}{(\sum_{k=1}^m \tilde{W}_k)^2}; \quad (12)$$

displaying at the origin of the plots of Subsection 2.2 a triangle whose tip-to-tip height is 4σ gives an indication of the variations that can be expected at roughly the 95% confidence level when assuming the null hypothesis that there is no difference in the expected responses between the populations being compared.

2.5 Review of reliability diagrams

The classical graphical method is known as the “reliability diagram,” as discussed, for example, by Bröcker & Smith (2007). A reliability diagram considers ℓ bins with boundaries $B_0 < B_1 < B_2 < \dots < B_\ell$, where $B_0 = -\infty$ and $B_\ell = \infty$ (whereas $B_1, B_2, \dots, B_{\ell-1}$ are finite real numbers). We form the weighted averages for each bin:

$$\bar{S}_i = \frac{\sum_{j: B_{i-1} < S_j \leq B_i} \sum_{k=1}^{n_j} S_j W_j^{(k)}}{\sum_{j: B_{i-1} < S_j \leq B_i} \sum_{k=1}^{n_j} W_j^{(k)}}, \quad (13)$$

$$\bar{R}_i = \frac{\sum_{j: B_{i-1} < S_j \leq B_i} \sum_{k=1}^{n_j} R_j^{(k)} W_j^{(k)}}{\sum_{j: B_{i-1} < S_j \leq B_i} \sum_{k=1}^{n_j} W_j^{(k)}}, \quad (14)$$

and

$$\bar{Q}_i = \frac{\sum_{j: B_{i-1} < S_j \leq B_i} \sum_{k=1}^{n_j} Q_j^{(k)} W_j^{(k)}}{\sum_{j: B_{i-1} < S_j \leq B_i} \sum_{k=1}^{n_j} W_j^{(k)}} \quad (15)$$

for $i = 1, 2, \dots, \ell$. The reliability diagram is then the graph of \bar{Q}_i versus \bar{S}_i plotted together with the graph of \bar{R}_i versus \bar{S}_i , both for $i = 1, 2, \dots, \ell$.

The traditional choice for the boundaries of the bins is to require $B_2 - B_1 \approx B_3 - B_2 \approx \dots \approx B_{\ell-1} - B_{\ell-2}$. Another sensible choice is to ensure that the sum of the squares of the weights associated with the observations in any of the bins, divided by the square of the sum of the weights in that bin, be similar for all bins. The latter choice ensures that noise due to statistical variations in the responses gets averaged away about the same amount in every bin when forming the averages in (14) and (15). Calculating the boundaries of the bins for this latter choice is possible via the algorithm described in Remark 5 of Tygert (2021b).

3 Results and discussion

This section illustrates the methods of the previous section via their application to several data sets.² Subsection 3.1 applies the methods to synthetic data for which the ground truth is known by construction. Subsection 3.2 applies the methodology to a classic data set from a direct-mail fundraising campaign, combining the methods of the previous section with those reviewed in Appendix A.

3.1 Synthetic examples

This subsection presents Figures 1–3. These figures display the results of processing data for which the ground truth is known by construction:

In Figures 1–3, the top rows of figures display the graphs of cumulative differences from Subsection 2.2. In these top rows, the graphs on the left use the random observations, while the graphs on the right use the exact expected values of the corresponding observations (thus displaying the ground truth with no noise from random observations). The second rows of the figures display reliability diagrams each with 10 bins. The third rows of the figures display reliability diagrams each with 100 bins. The reliability diagrams on the left use the classical binning, with bins roughly equispaced along the scores. The reliability diagrams on the right choose bins such that the ratio of the sum of the squares of the weights in the bin to the square of the sum of the weights in the bin is roughly the same for every bin (this ensures that the variance in each bin is similar). The fourth, bottom rows of the figures are reliability diagrams constructed from the exact expected values of the corresponding observations (thus displaying the ground truth for the reliability diagrams).

Figures 1–3 vary the number m of distinct values among all the scores (S_1, S_2, \dots, S_m) , with $m = 1,000$ in Figure 1, $m = n = 4,000$ in Figure 2, and $m = 400$ in Figure 3. Figures 1–3 all set the sample size n to be 4,000, while varying the number m of unique values among the scores for the n pairs of observations.

In every figure, the reliability diagrams obscure variations that are readily apparent in the cumulative graphs; the captions to the figures discuss in greater detail. The reliability diagrams with fewer bins are much less noisy yet smooth over important variations. The cumulative graphs manage to capture everything all at once, in contrast. In Figures 1 and 2, the scalar statistics of Kuiper and of Kolmogorov and Smirnov successfully reflect the steep slopes in the cumulative graphs; in Figure 3, oscillations dampen the effectiveness of the scalar metrics. The captions to the figures discuss the results further.

3.2 KDD Cup 1998

This subsection presents Figures 4–6. Appendix A reviews the Hilbert curve used in the present subsection. Figures 4–6 consider data from an experiment that solicited donations during a fundraising campaign:

²Permissively licensed open-source Python scripts that reproduce all figures and statistics reported here are available at <https://github.com/facebookresearch/metapaired>

The 1998 Knowledge Discovery and Data-mining (KDD) Cup was a competition based on a data set regarding donations to a national veterans organization in the U.S.³ The data set reports the donations made in different years, including 1995 and 1996; we total all the donations that each individual made in each year and retain the 63,826 individuals who donated at least once. We exclude all those with missing ages (“AGE” in the data), those with missing mean household incomes for the Census block in which the donor lives (“IC3”), and those missing the fractions of householders in the corresponding Census block who are married (“MARR1”), so that we can use these three as covariates. In the data set, ages are integers, average household incomes are integer multiples of a hundred dollars, and fractions married are integer percentages. All covariates get normalized to the unit interval $[0, 1]$ prior to ordering their values via the Hilbert curve reviewed in Appendix A. We perturbed the incomes by about one part in a hundred million in order to ensure that they would be distinct; the previous subsection omits such random perturbation in order to illustrate use of the method of Section 2.4 of Tygert (2022), which is an alternative to random perturbations. The random perturbations guarantee that the scores are all unique, so the total number of scores is equal to the number of individuals, that is, $m = n = 63,826$.

The responses are the total amounts in dollars that every individual donated in each year (1995 and 1996), with those who never donated in either year removed. To account for inflation and focus on variation with respect to the covariates, we normalized every individual’s total donation in each year by dividing by the mean total donation, with the mean averaging over all individuals in that year. In each year, the mean total donation, with the mean averaging over all individuals in that year, is of course equal to 1 following this normalization. The mean difference between 1995 and 1996 in the individuals’ normalized donations, with the mean taken over all individuals, is therefore 0 (as $1 - 1 = 0$).

Figures 4, 5, and 6 condition on the covariates (1) age and income, (2) income and age, and (3) age, fraction married, and income, respectively. (Controlling for covariates via the Hilbert curve displays some dependence on the order in which we condition on the covariates, resulting in Figures 4 and 5 being different, despite conditioning on the same pair of covariates, age and income.) The top rows of the figures plot the cumulative differences in responses between 1995 and 1996 (with each year normalized such that the mean response, averaged over all individuals, is 1). The top rows of Figures 4 and 5 also display points in the two-dimensional plane of the covariates, where the points correspond to the individuals’ values for the covariates. The grayscale shading of each point indicates the score assigned via parameterization with the Hilbert curve. Interactive display of the covariates’ values corresponding to given scores is available in the software accompanying the present paper.⁴ The middle rows of the figures give reliability diagrams each with 10 bins. The bottom rows of the figures give reliability diagrams each with 100 bins. The leftmost reliability diagrams employ the standard binning, partitioning the scores roughly equally, with bins equispaced along the scores. The rightmost reliability diagrams bin such that the ratio of the sum of the squares of the weights in the bin to the square of the sum of the weights in the bin is roughly the same for every bin, so that the variance in each bin is similar.

All three figures reveal roughly the same differences between responses in 1995 and in 1996. In principle Figure 6 could have been very different from Figures 4 and 5, since Figure 6 conditions on an extra covariate, but apparently the extra conditioning made only so much difference. The reliability diagrams on the right look more helpful than those on the left; those on the left use the traditional binning, with the widths of the intervals of scores that the bins span being roughly the same for all bins. The cumulative diagrams appear to be the most helpful, or at least the clearest.

³The data from the 1998 KDD Cup is available at <https://kdd.org/kdd-cup/view/kdd-cup-1998/Data>

⁴Open-source software with interactive plots is available at <https://github.com/facebookresearch/metapaired>

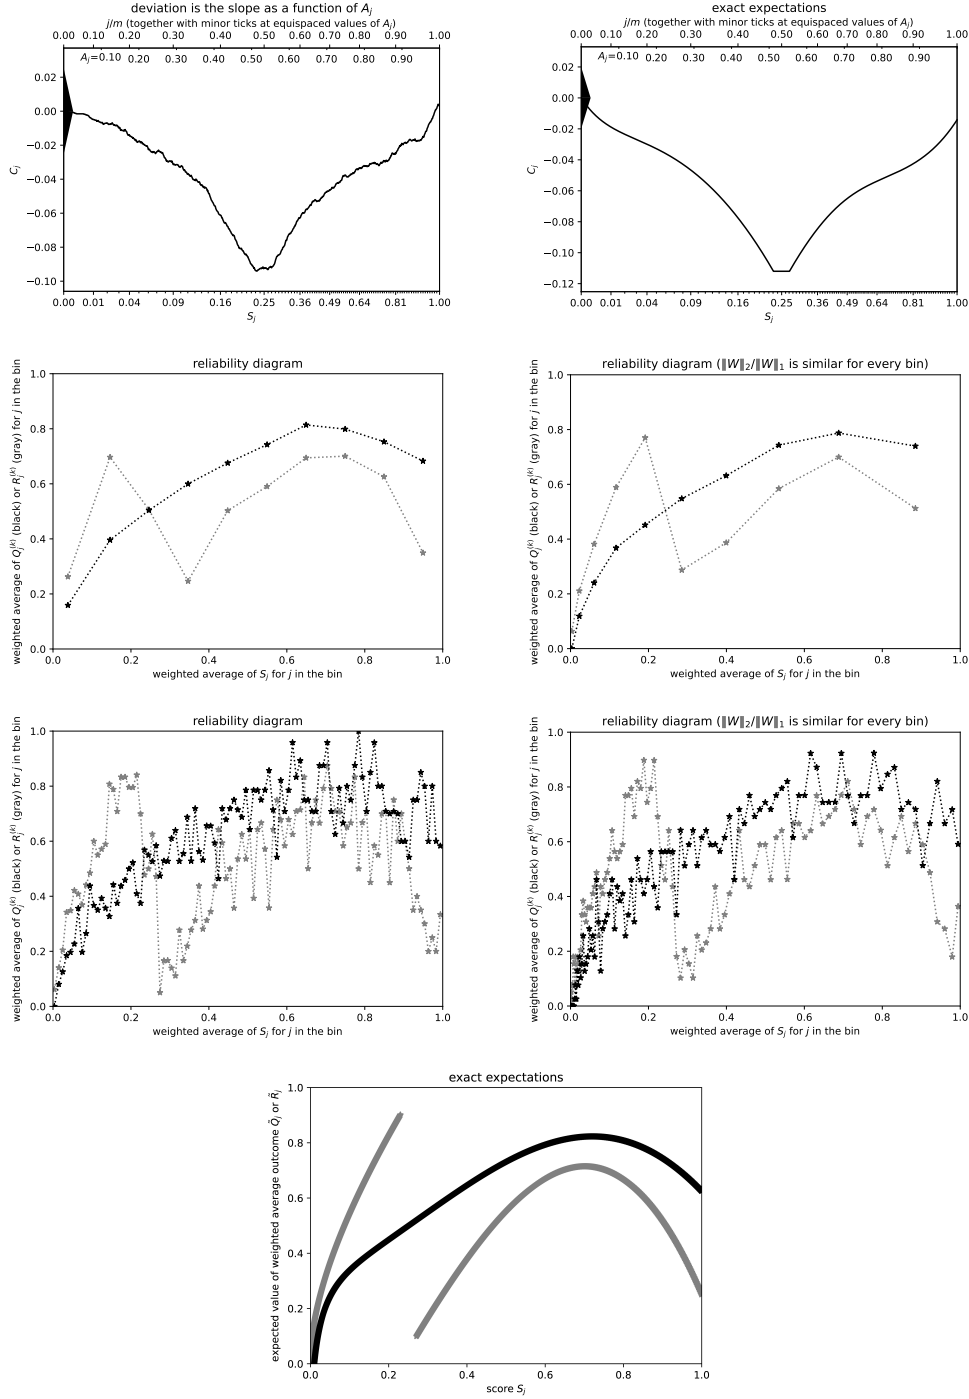


Figure 1: $m = 1,000$, $n = 4,000$; Kuiper's statistic is $0.09816/\sigma = 7.784$, Kolmogorov's and Smirnov's is $0.09403/\sigma = 7.456$. The cumulative graphs clearly reveal a narrow range of scores (right around the median score) that is flat, that is, where there is little difference in the responses between the populations. In contrast, discerning the lack of difference in responses between the populations in this narrow range is very challenging using only the reliability diagrams — the cumulative graphs are far clearer. The weighted average difference in responses is the vertical coordinate C_m at the greatest (rightmost) score S_m in the cumulative plots, which is clearly much smaller (closer to 0) than the full vertical range of the graph; the vertical range ($\max_{0 \leq j \leq m} C_j - \min_{0 \leq j \leq m} C_j$) is the value of the Kuiper metric. Overall, the empirical cumulative graph closely matches the exact expected ground-truth cumulative graph.

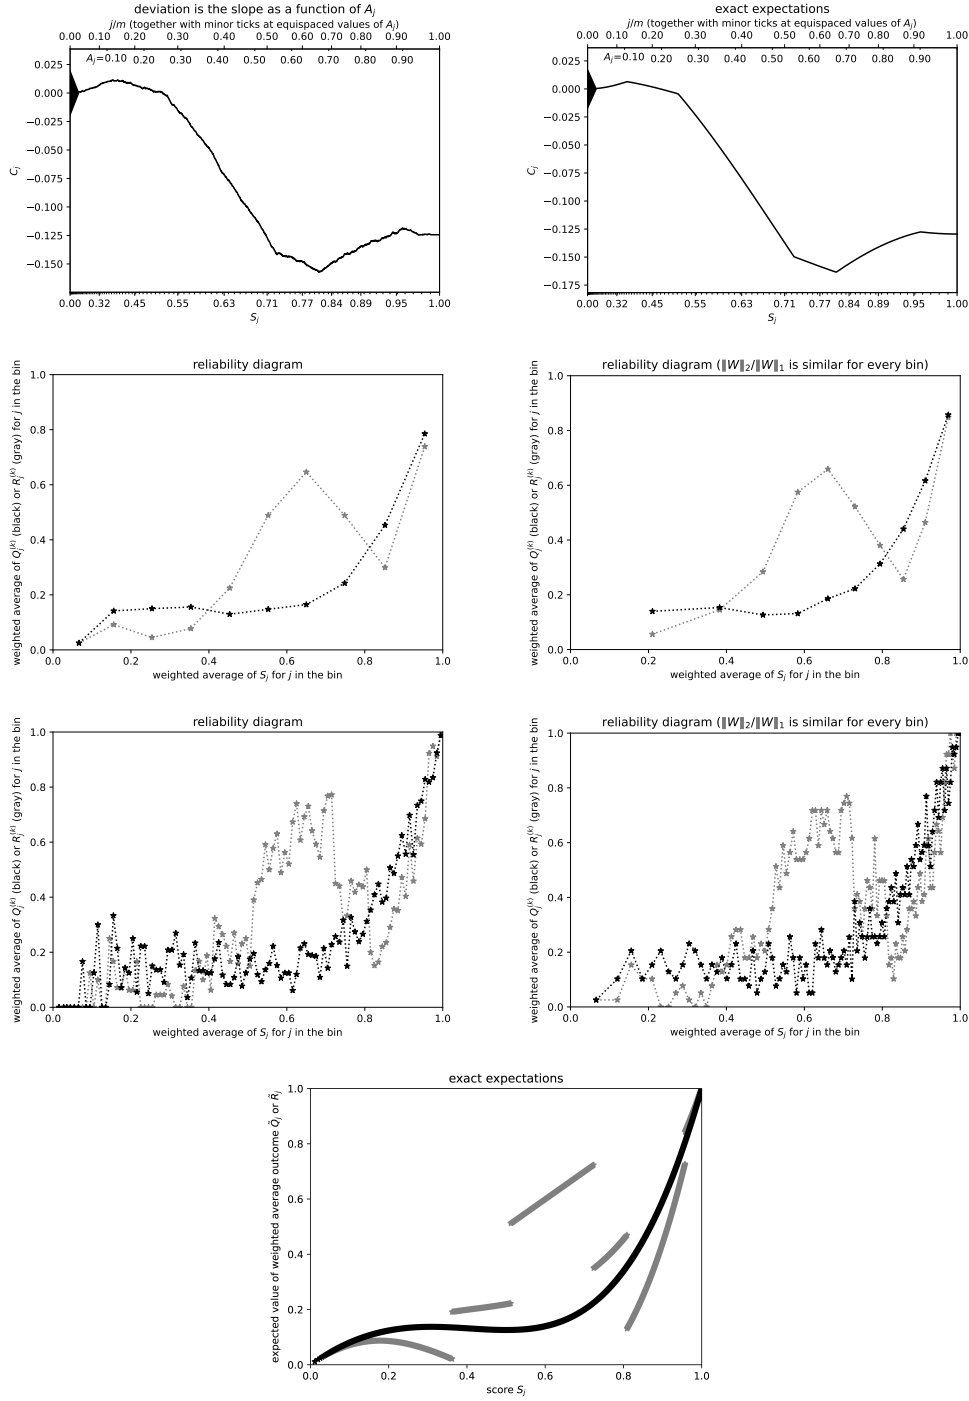


Figure 2: $m = 4,000$, $n = 4,000$; Kuiper's statistic is $0.1686/\sigma = 16.20$, Kolmogorov's and Smirnov's is $0.1571/\sigma = 15.09$. The cumulative graphs have fairly clear kinks at the scores where the reliability diagrams ideally would jump discontinuously in order to match the lowermost plot. However, detecting discontinuous jumps in the reliability diagrams based on the actual random observations is very hard. When the bins are narrow enough to resolve the jumps, the noise on the weighted average response in each bin is large, with the averages jumping all over due to the noise (in addition to the jumps in the underlying expected distribution depicted in the lowermost plot). For the most part, the cumulative graph based on random observations closely matches the exact expected ground-truth cumulative graph.

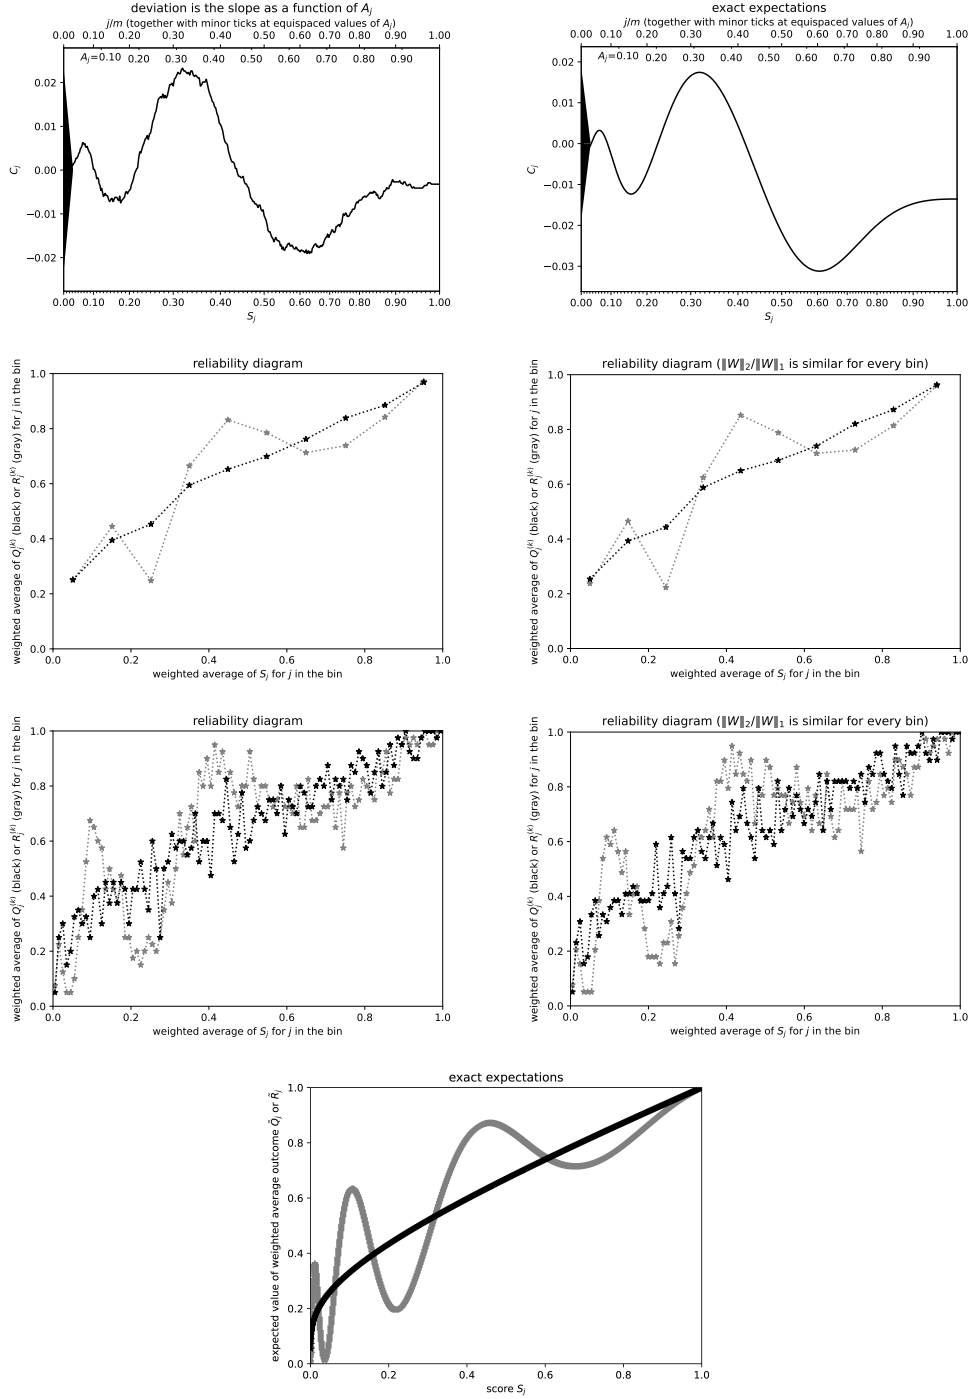


Figure 3: $m = 400$, $n = 4,000$; Kuiper's statistic is $0.04235/\sigma = 3.676$, Kolmogorov's and Smirnov's is $0.02328/\sigma = 2.021$. Recall that the weighted average difference in responses is the vertical coordinate C_m at the greatest (rightmost) score S_m in the cumulative plots. This value is significantly smaller, that is, nearer to 0, than the vertical range of the graph; the vertical range ($\max_{0 \leq j \leq m} C_j - \min_{0 \leq j \leq m} C_j$) is the value of the Kuiper statistic. Whereas any one of the empirical reliability diagrams obscures at least some aspect of the exact ground-truth diagram displayed in the lowermost plot, the empirical cumulative graph closely resembles the exact expected ground-truth cumulative graph.

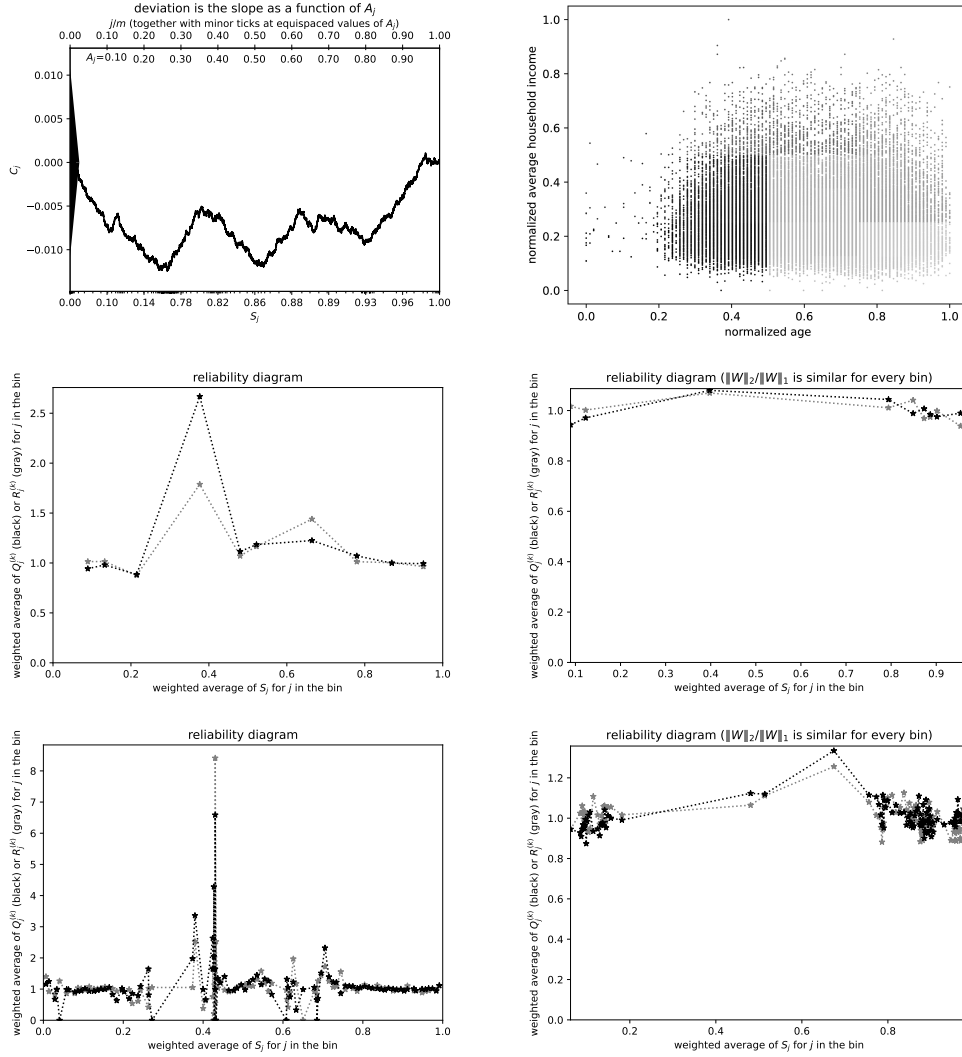


Figure 4: $m = n = 63,826$; Kuiper's statistic is $0.01354/\sigma = 2.514$, Kolmogorov's and Smirnov's is $0.01245/\sigma = 2.312$. Notice how these scalar statistics miss the very steep slopes in the cumulative graph, as the graph oscillates rapidly. The following figure, Figure 5, fixes this failure by reversing the order of the covariates in the parameterization via the Hilbert curve. The present figure and the following figure display similarly high slopes, but re-ordering the scores as in the following figure prevents the oscillation that the scalar metrics cannot capture.

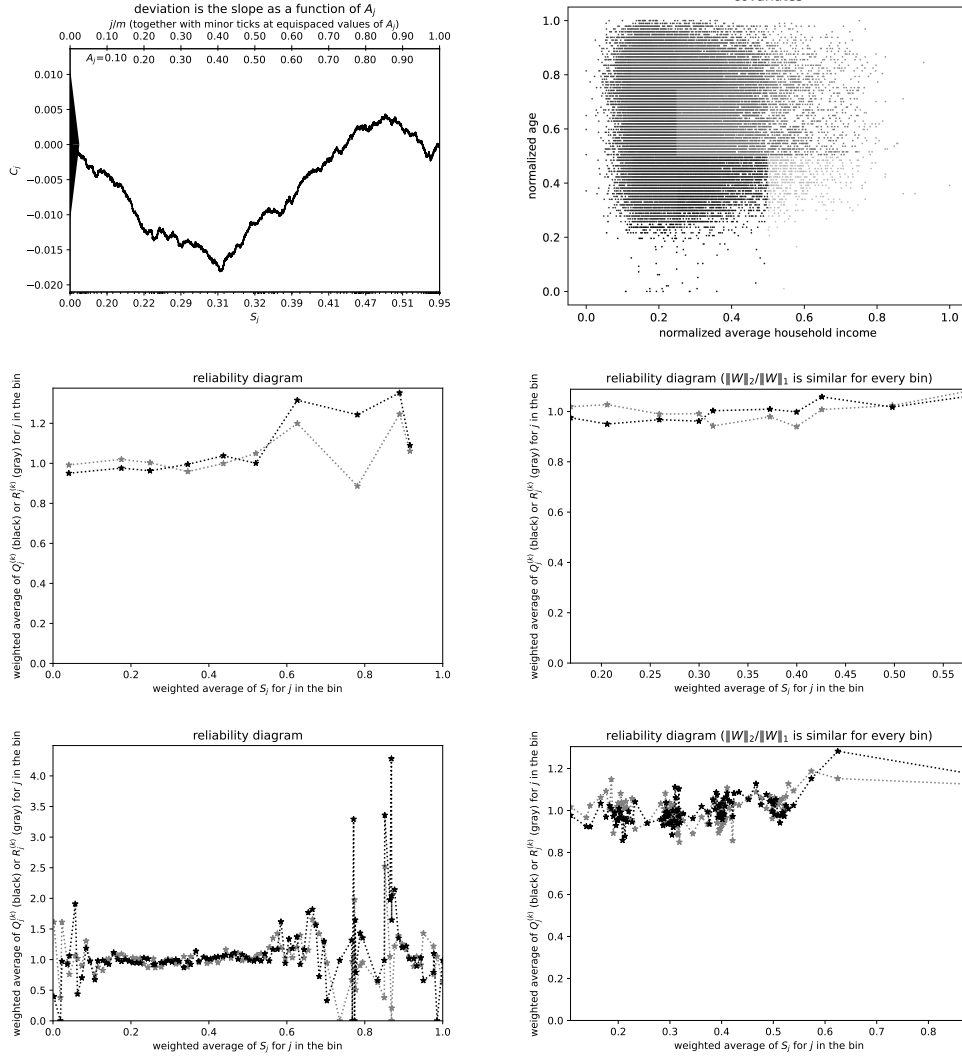


Figure 5: $m = n = 63,826$; Kuiper's statistic is $0.02246/\sigma = 4.170$, Kolmogorov's and Smirnov's is $0.01813/\sigma = 3.366$. This figure is the same as the previous figure, Figure 4, but with the order of the covariates reversed in the parameterization by the Hilbert curve. Both the present figure and the former figure reveal similarly steep slopes; the present figure reduces oscillations in the cumulative graph in comparison to the previous figure. The scalar statistics reflect the steeper slopes better in the present figure, due to the reduction in cancellation from oscillation.

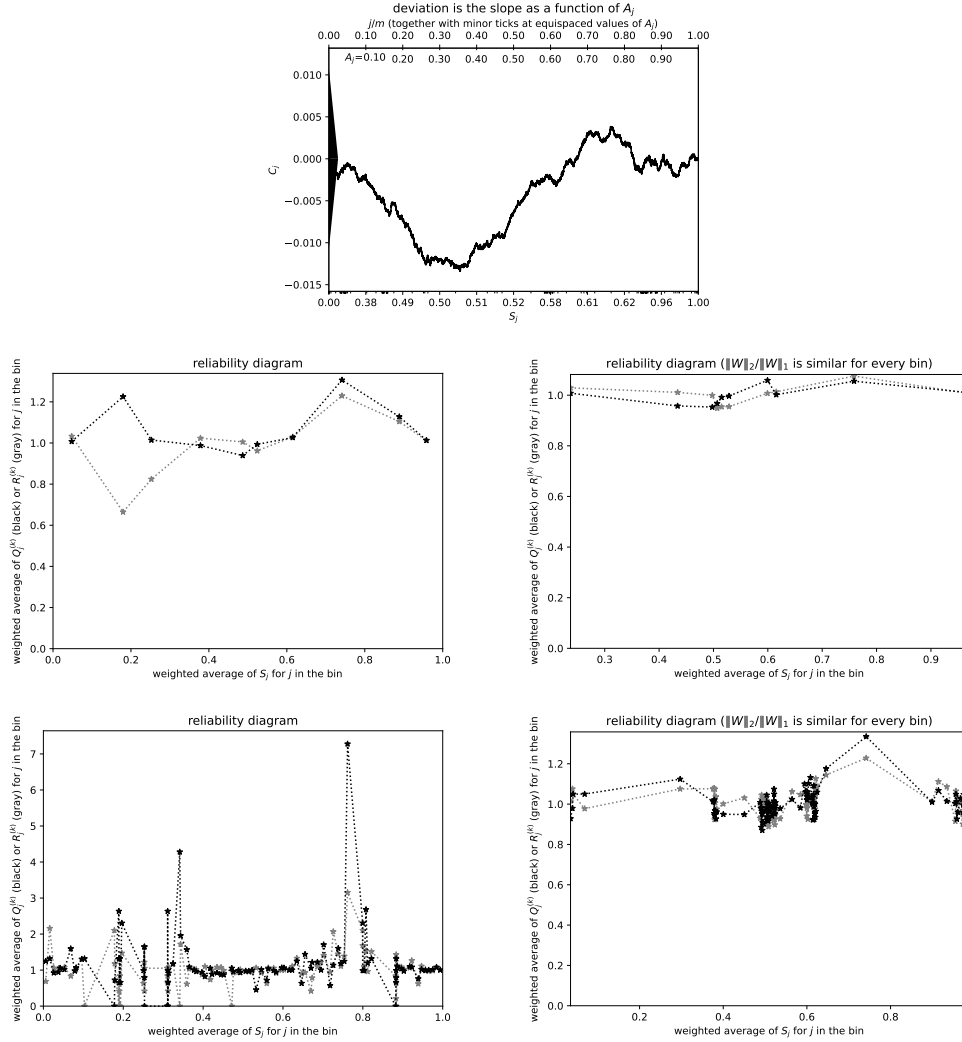


Figure 6: $m = n = 63,826$; Kuiper's statistic is $0.01723/\sigma = 3.199$, Kolmogorov's and Smirnov's is $0.01338/\sigma = 2.485$. This figure is the same as the previous two figures, Figures 4 and 5, but now using all three covariates in the parameterization via the Hilbert curve. The covariates for the present figure are the normalized age of the donor, the normalized average fraction married in the Census block where the donor lives, and the normalized average household income in the Census block where the donor lives. The previous two figures omitted conditioning on the fraction married.

4 Conclusion

Earlier works demonstrated many advantages of the cumulative statistics over conventional methods for analyzing differences between unpaired samples. As shown above, these advantages hold for matched pairs, as well. For instance, the Kuiper metric avoids the cancellation of positive and negative differences that often afflicts the simplistic weighted average. (The Kuiper metric is the absolute value of the total weighted difference in responses, totaled over the interval of values of the covariate for which the absolute value is greatest. The simplistic weighted average is the same as the total taken over the entire range of values of the covariate; the totals here refer to sums normalized such that they become the weighted average when totaled over the entire range. If the samples are unweighted, then the weighted average is just the usual arithmetic average.) The Kuiper statistic can detect differences that the weighted average averages away.

The cumulative methods have no parameters to set that the analyst could bungle or fudge; the cumulative graphs always reveal all significant features in the differences between the pair of populations being analyzed, unlike the canonical reliability diagrams. In contrast, the choices of bins for reliability diagrams often obscure important features, as illustrated above with both measured data sets and synthetic examples for which the ground truth is known by construction. Furthermore, the methods detailed above for a single scalar covariate extend to the case of multiple covariates (as illustrated already in Section 3 above); the following appendix presents this multivariate extension.

Acknowledgements

We would like to thank Kamalika Chaudhuri and Khalid El-Arini for their support.

A Review of space-filling curves

This appendix reviews the construction and applications of Hilbert curves. Hilbert curves reduce the problem of analyzing the differences in responses while controlling for multiple covariates to the problem of analyzing the differences while controlling for only a single scalar covariate. The rest of the present paper treats the latter problem in detail. This appendix reviews the reduction introduced by Tygert (2021a).

Space-filling curves were introduced by Peano (1890). A simplification by Hilbert (1891) soon followed and forms the basis for the method used in the present paper. Related applications are detailed by Moon *et al.* (2001).

A space-filling curve in p dimensions is a continuous function h that maps the unit interval $(0, 1)$ onto the unit hypercube $(0, 1)^p$ (“onto” means that the mapping is surjective, covering every point in the hypercube). Accompanying h is a function g mapping from the unit hypercube $(0, 1)^p$ to the unit interval $(0, 1)$ that is a right inverse for h , that is, $h(g(x)) = x$ for any point x in the unit hypercube $(0, 1)^p$. As reviewed, for example, by Tygert (2021a), g cannot be continuous, unlike h .

The continuity of the function h mapping from the unit interval $(0, 1)$ onto the unit hypercube $(0, 1)^p$ guarantees that, for any points t and u from the unit interval $(0, 1)$, if t and u are near to each other, then $h(t)$ and $h(u)$ are also close. This ensures that, for any real-valued function f on the unit hypercube $(0, 1)^p$, local averages of $f \circ h$ will also be local averages of f ; here, “local” refers to proximity in the usual Euclidean metric and $f \circ h$ denotes the composition of f and h , that is, $(f \circ h)(t) = f(h(t))$ for any point t in the unit interval $(0, 1)$. We can place a total order on the unit hypercube $(0, 1)^p$ by defining that $x < y$ means $g(x) < g(y)$ for any two points x and y in the unit hypercube $(0, 1)^p$. A local average of f under this ordering is also a local average of $f \circ h$ and thus is also a local average of f under the usual Euclidean metric, as observed already in the second sentence of this same paragraph.

A Hilbert curve h is the particular type of space-filling curve that orders points via a depth-first traversal of the canonical 2^p -ary (dyadic) tree on the unit hypercube $(0, 1)^p$. Of course, the canonical dyadic tree in one dimension ($p = 1$) is the binary tree, the canonical dyadic tree in two dimensions ($p = 2$) is the quad-tree, and the canonical dyadic tree in three dimensions ($p = 3$) is the oct-tree. Figure 7 approximates the image of a Hilbert curve via 255 line segments. The full Hilbert curve is the limit of approximations with increasingly many line segments, with every line segment in an approximation being equally long as every other line segment in the same approximation. Since the mapping from one dimension to p dimensions for




Figure 7: An approximation with 255 line segments to the Hilbert curve in $p = 2$ dimensions; the numbers displayed near the points specify their order in the Hilbert curve.

each approximation is continuous and the approximations converge uniformly, the limit — the Hilbert curve — must exist and be continuous. Our implementation uses an approximation with $2^{64} - 1$ line segments, suitable for computation in 64-bit arithmetic with unsigned 8-byte integers.

To form the cumulative graphs and scalar summary statistics of Section 2, we need scalar scores. To obtain the score associated with a given value of a vector of covariates, we apply g to the vector. This reduces the problem of analyzing multiple covariates to the problem of analyzing a single scalar covariate; Section 2 details a solution to the latter problem, thus yielding a solution to the former problem (via the scalar scores obtained from the Hilbert curve), too. As discussed in the fourth paragraph of this appendix, averaging responses locally at the scalar scores turns out to average the same responses locally at the associated values of the multiple covariates, where the “associated values” are those obtained by applying h to the scalar scores. This ensures that the cumulative graphs and statistics take meaningful averages.

(Only the ordering of the scores affects the shapes of the cumulative graphs and the values of the scalar summary statistics. For the numerical values of the scores used to label the lower horizontal axes, we normalize the results to range from 0 to 1. Such normalization has no effect on the ordering of the scores. Specifically, we apply an affine transformation — linear plus a constant — first subtracting the minimum of the results of g applied to all observed vectors of covariates and then dividing by the original maximum minus the original minimum.)

References

- Austin, Peter C. 2011. An introduction to propensity score methods for reducing the effects of confounding in observational studies. *Multivar. Behav. Res.*, **43**(3), 399–424.
- Bai, Yuehao, Romano, Joseph P., & Shaikh, Azeem M. 2022. Inference in experiments with matched pairs. *J. Amer. Statist. Assoc.*, **117**(540), 1726–1737.
- Bröcker, Jochen, & Smith, Leonard A. 2007. Increasing the reliability of reliability diagrams. *Weather Forecast.*, **22**(3), 651–661.
- Brown, Jr., Byron William, & Hollander, Myles. 2007. *Statistics: A Biomedical Introduction*. Wiley Classics Library. Wiley.
- Caliendo, Marco, & Kopeinig, Sabine. 2008. Some practical guidance for the implementation of propensity score matching. *J. Econ. Surv.*, **22**(1), 32–72.
- Cox, David R., & Reid, Nancy. 2000. *The Theory of the Design of Experiments*. Monographs on Statistics and Applied Probability, vol. 86. Chapman & Hall / CRC Press.
- Delgado, Miguel A. 1993. Testing the equality of nonparametric regression curves. *Statist. Probab. Lett.*, **17**(3), 199–204.
- Draper, Norman R., & Smith, Harry. 1998. *Applied Regression Analysis*. 3rd edn. Wiley.
- Hilbert, David. 1891. Ueber die stetige abbildung einer line auf ein flächenstück. *Math. Ann.*, **38**(3), 459–460.
- Hollander, Myles, Wolfe, Douglas A., & Chicken, Eric. 2014. *Nonparametric Statistical Methods*. 3rd edn. Wiley.
- Hosmer, Jr., David W., Lemeshow, Stanley, & Sturdivant, Rodney X. 2013. *Applied Logistic Regression*. 3rd edn. Wiley.
- Ibarra, Imanol Arrieta, Gujral, Paman, Tannen, Jonathan, Tygert, Mark, & Xu, Cherie. 2022. Metrics of calibration for probabilistic predictions. *J. Mach. Learn. Res.*, **23**, 1–54.
- Kuiper, Nicolaas H. 1962. Tests concerning random points on a circle. *Proc. Koninklijke Nederlandse Akademie van Wetenschappen Series A*, **63**, 38–47.
- Moon, Bongki, Jagadish, H. V., Faloutsos, Christos, & Saltz, Joel H. 2001. Analysis of the clustering properties of the Hilbert space-filling curve. *IEEE Trans. Knowl. Data Eng.*, **13**(1), 124–141.
- Peano, Giuseppe. 1890. Sur une courbe, qui remplit toute une aire plane. *Math. Ann.*, **36**(1), 157–160.
- Rubin, Donald B. 1973. Matching to remove bias in observational studies. *Biometrics*, **29**, 159–183.
- Tygert, Mark. 2021a. *Controlling for multiple covariates*. Tech. rept. 2112.00672. arXiv. Available at <https://arxiv.org/abs/2112.00672>.
- Tygert, Mark. 2021b. Cumulative deviation of a subpopulation from the full population. *J. Big Data*, **8**(117), 1–60. Available at <https://arxiv.org/abs/2008.01779>.
- Tygert, Mark. 2021c. A graphical method of cumulative differences between two subpopulations. *J. Big Data*, **8**(158), 1–29. Available at <https://arxiv.org/abs/2108.02666>.
- Tygert, Mark. 2022. *Calibration of P-values for calibration and for deviation of a subpopulation from the full population*. Tech. rept. 2202.00100. arXiv. Available at <https://arxiv.org/abs/2202.00100>.
- Wilcoxon, Frank. 1945. Individual comparisons by ranking methods. *Biometrics Bulletin*, **1**(6), 80–83.