

Figure 1: Pd for the cross-company experiment. Naive Bayes used to create prediction models. Within each plot the baseline results on the unprivatized data (“orig”) are shown on the left-hand-side and the results from the PriestPrivacy approach of Taneja et al. [?] are shown on the right-hand-side (“s10”, “s20”, “s40”), along with k-anonymity where k is 2 and 4 (k2 and k4). In the middle are the results of applying MORPH to data that CLIFF has reduced to X% of the original size. Of particular interest is the “m10” results; i.e. use CLIFF to reduce the data to 10% of the original and then then apply MORPH. Note that in all cases, “m10” results in higher pds and lower pfs (compared to both both the original data and in data processed by PriestPrivacy).

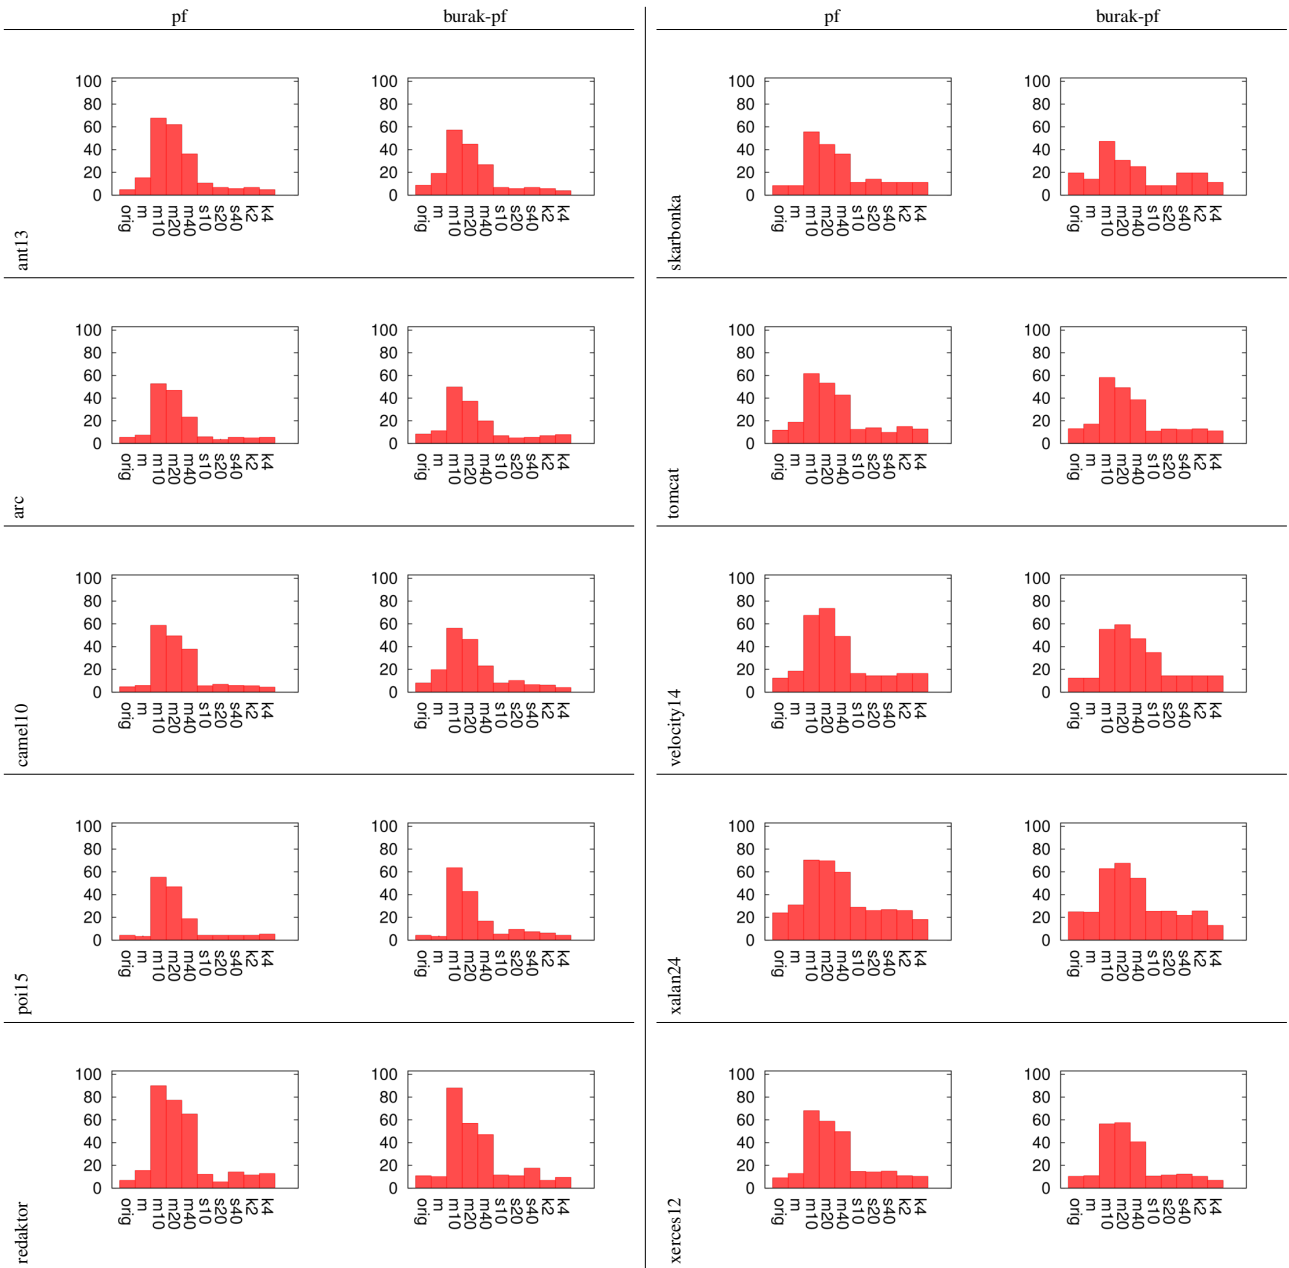


Figure 2: pf for the cross-company experiment. Naive Bayes used to create prediction models. Within each plot the baseline results on the unprivatized data (“orig”) are shown on the left-hand-side and the results from the PriestPrivacy approach of Taneja et al. [?] are shown on the right-hand-side (“s10”, “s20”, “s40”), along with k-anonymity where k is 2 and 4 (k2 and k4). In the middle are the results of applying MORPH to data that CLIFF has reduced to X% of the original size. Of particular interest is the “m10” results; i.e. use CLIFF to reduce the data to 10% of the original and then then apply MORPH. Note that in all cases, “m10” results in higher pds and lower pfs (compared to both both the original data and in data processed by PriestPrivacy).



Figure 3: f for the cross-company experiment. Naive Bayes used to create prediction models. Within each plot the baseline results on the unprivatized data (“orig”) are shown on the left-hand-side and the results from the PriestPrivacy approach of Taneja et al. [?] are shown on the right-hand-side (“s10”, “s20”, “s40”), along with k-anonymity where k is 2 and 4 (k2 and k4). In the middle are the results of applying MORPH to data that CLIFF has reduced to X% of the original size. Of particular interest is the “m10” results; i.e. use CLIFF to reduce the data to 10% of the original and then then apply MORPH. Note that in all cases, “m10” results in higher pds and lower pfs (compared to both both the original data and in data processed by PriestPrivacy).



Figure 4: g for the cross-company experiment. Naive Bayes used to create prediction models. Within each plot the baseline results on the unprivatized data (“orig”) are shown on the left-hand-side and the results from the PriestPrivacy approach of Taneja et al. [?] are shown on the right-hand-side (“s10”, “s20”, “s40”), along with k-anonymity where k is 2 and 4 (k2 and k4). In the middle are the results of applying MORPH to data that CLIFF has reduced to X% of the original size. Of particular interest is the “m10” results; i.e. use CLIFF to reduce the data to 10% of the original and then then apply MORPH. Note that in all cases, “m10” results in higher pds and lower pfs (compared to both both the original data and in data processed by PriestPrivacy).