Supplementary material for "Preference-Based Rank Elicitation using Statistical Models: The Case of Mallows"

A. Auxiliary claim ($\widetilde{r}_{i^*} = 1$) for the proof of Proposition 1

Claim 9. *Assume* $\widetilde{r}_i = r_i = 1$. *Then* $d(\mathbf{r}(i,j), \widetilde{\mathbf{r}}) > d(\mathbf{r}, \widetilde{\mathbf{r}})$.

Proof. First of all,

$$d(\mathbf{r}(i,j), \widetilde{\mathbf{r}}) - d(\mathbf{r}, \widetilde{\mathbf{r}}) = \sum_{(\ell,k):\, r_\ell < r_k} \Big(\mathbb{I}\{(r_\ell(i,j) - r_k(i,j))(\widetilde{r}_\ell - \widetilde{r}_k) < 0\} - \mathbb{I}\{(r_\ell - r_k)(\widetilde{r}_\ell - \widetilde{r}_k) < 0\} \Big)$$

Note that the term in the sum equals to 0 whenever $\{\ell, k\} \cap \{i, j\} = \emptyset$, and whenever $\ell \in \{i, j\}$ and $r_k > r_j$. Thus it is nonzero only if $r_k = r_j$ or if $r_k < r_j$ and $\ell = i$. Therefore,

$$\begin{aligned} d(\mathbf{r}(i,j), \widetilde{\mathbf{r}}) - d(\mathbf{r}, \widetilde{\mathbf{r}}) =& \Big(\mathbb{I}\{(r_j - r_i)(\widetilde{r}_i - \widetilde{r}_j) < 0\} - \mathbb{I}\{(r_i - r_j)(\widetilde{r}_i - \widetilde{r}_j) < 0\} \Big) & (6) \\ &+ \sum_{k:\, 1 < r_k < r_j} \Big(\mathbb{I}\{(r_j - r_k)(\widetilde{r}_i - \widetilde{r}_k) < 0\} - \mathbb{I}\{(r_i - r_k)(\widetilde{r}_i - \widetilde{r}_k) < 0\} \Big) & (7) \\ &+ \sum_{\ell:\, 1 < r_\ell < r_j} \Big(\mathbb{I}\{(r_\ell - r_i)(\widetilde{r}_\ell - \widetilde{r}_j) < 0\} - \mathbb{I}\{(r_\ell - r_j)(\widetilde{r}_\ell - \widetilde{r}_j) < 0\} \Big) & (8) \end{aligned}$$

where (6) corresponds to the case when $\ell = i$ and $k = j$, (7) corresponds to the case when $\ell = i$ and $k \neq j$, and finally, (8) corresponds to the case when $\ell \neq i$ and $k = j$. Note that $r_i = \widetilde{r}_i = 1$ implies that (6) equals to 1 and that (7) equals to $r_j - 2$. Noting also that (8) cannot exceed $r_j - 2$ (because it is the sum of $r_j - 2$ terms of absolute value at most 1), the claim follows. □

B. Implementation based on quick sort for finding the most probable ranking

Algorithm 6 MALLOWSQUICK(δ)

for $j = 1 \to M$ **do** $r_j = 0$
Set $A = \{1, \ldots, M\}$
$\mathbf{r} =$ MQREC($\mathbf{r}, A, \delta, 1$)
return r

Procedure 7 MQREC($\mathbf{r}, A, \delta, c$)

if $\#A = 1$ **then**
    Pick index i from A
    $r_i = c$
else
    Pick a random index $i \in A$ and set $A = A \setminus \{i\}$
    (A_+, A_-) = MALLOWSHALVING(i, A, δ)
    $r_i = \#A_+ + c$
    if $\#A_+ > 0$ **then**
        $\mathbf{r} =$ MQREC($\mathbf{r}, A_+, \delta, c$)
    if $\#A_- > 0$ **then**
        $\mathbf{r} =$ MQREC($\mathbf{r}, A_-, \delta, r_i$)
return r

Procedure 8 MALLOWSHALVING(i, A, δ)

for $j \in A$ **do** ▷ Initialization
  $\widehat{p}_{i,j} = 0, n_{i,j} = 0, c_{i,j} = \text{Inf}$
$A_- = \emptyset, A_+ = \emptyset$
repeat
  for $j \in A$ **do**
    Observe $o = \mathbb{I}\{r_i < r_j\}$
    $\widehat{p}_{i,j} = \widehat{p}_{i,j} + o, n_{i,j} = n_{i,j} + 1$
    $c_{i,j} = \sqrt{\frac{1}{2n_{i,j}} \log \frac{2M^2 n_{i,j}^2}{\delta}}$
    if $1/2 < \widehat{p}_{i,j} - c_{i,j}$ **then** ▷ $\widetilde{r}_j > \widetilde{r}_i$ w.h.p.
      $A = A \setminus \{j\}, A_+ = A_+ \cup \{j\}$
    if $1/2 > \widehat{p}_{i,j} + c_{i,j}$ **then** ▷ $\widetilde{r}_j < \widetilde{r}_i$ w.h.p.
      $A = A \setminus \{j\}, A_- = A_- \cup \{j\}$
until $A = \emptyset$
return (A_+, A_-)

C. Upper bound for KL-divergence in the case of Mallow's ϕ-model

Assume that there are given two Mallow's ϕ-models $\mathbb{P}(.|\phi, \widetilde{\mathbf{r}})$ and $\mathbb{P}(.|\widehat{\phi}, \widetilde{\mathbf{r}})$ with the same central ranking $\widetilde{\mathbf{r}}$. We will concisely write $d_{\mathbf{r}}$ for $d(\widetilde{\mathbf{r}}, \mathbf{r})$. Then we have

$$\begin{aligned}
\mathrm{KL}(\mathbb{P}(.|\phi, \widetilde{\mathbf{r}}), \mathbb{P}(.|\widehat{\phi}, \widetilde{\mathbf{r}})) &= -\sum_{\mathbf{r} \in S_M} \mathbb{P}(\mathbf{r}|\phi, \widetilde{\mathbf{r}}) \log \mathbb{P}(\mathbf{r}|\widehat{\phi}, \widetilde{\mathbf{r}}) + \sum_{\mathbf{r} \in S_M} \mathbb{P}(\mathbf{r}|\phi, \widetilde{\mathbf{r}}) \log \mathbb{P}(\mathbf{r}|\phi, \widetilde{\mathbf{r}}) \\
&= -\sum_{\mathbf{r} \in S_M} \frac{\phi^{d_{\mathbf{r}}}}{Z(\phi)} \log \frac{\widehat{\phi}^{d_{\mathbf{r}}}}{Z(\widehat{\phi})} + \sum_{\mathbf{r} \in S_M} \frac{\phi^{d_{\mathbf{r}}}}{Z(\phi)} \log \frac{\phi^{d_{\mathbf{r}}}}{Z(\phi)} \\
&= -\frac{1}{Z(\phi)} \sum_{\mathbf{r} \in S_M} \phi^{d_{\mathbf{r}}} \left[d_{\mathbf{r}} \log \widehat{\phi} - \log Z(\widehat{\phi}) \right] + \frac{1}{Z(\phi)} \sum_{\mathbf{r} \in S_M} \phi^{d_{\mathbf{r}}} \left[d_{\mathbf{r}} \log \phi - \log Z(\phi) \right] \\
&= -\frac{\log \widehat{\phi}}{Z(\phi)} \sum_{\mathbf{r} \in S_M} d_{\mathbf{r}} \phi^{d_{\mathbf{r}}} + \frac{\log Z(\widehat{\phi})}{Z(\phi)} \sum_{\mathbf{r} \in S_M} \phi^{d_{\mathbf{r}}} + \frac{\log \phi}{Z(\phi)} \sum_{\mathbf{r} \in S_M} d_{\mathbf{r}} \phi^{d_{\mathbf{r}}} - \frac{\log Z(\phi)}{Z(\phi)} \sum_{\mathbf{r} \in S_M} \phi^{d_{\mathbf{r}}} \\
&= \frac{\log \phi - \log \widehat{\phi}}{Z(\phi)} \sum_{\mathbf{r} \in S_M} d_{\mathbf{r}} \phi^{d_{\mathbf{r}}} + \log Z(\widehat{\phi}) - \log Z(\phi) \qquad (9) \\
&\leq \frac{M(M-1)}{2} \log \frac{\phi}{\widehat{\phi}} + \log \frac{Z(\widehat{\phi})}{Z(\phi)} \qquad (10)
\end{aligned}$$

where (9) follows from $Z(\phi) = \sum_{\mathbf{r} \in S_M} \phi^{d_{\mathbf{r}}}$, and (10) is true, because $d_{\mathbf{r}} \leq \frac{M(M-1)}{2}$ for any $\mathbf{r} \in S_M$.

D. Sample complexity analysis for the KL divergence case

For the reading convenience we restate the theorem here.

Theorem 10. *Assume that the ranking distribution is Mallows with parameters ϕ and $\widetilde{\mathbf{r}}$. Then, for any $\epsilon > 0$ and $0 < \delta < 1$,* MALLOWSKLD *returns parameter estimates $\widehat{\mathbf{r}}$ and $\widehat{\phi}$ for which $\mathrm{KL}(\mathbb{P}(\cdot \,|\, \phi, \widetilde{\mathbf{r}}), \mathbb{P}(\cdot \,|\, \widehat{\phi}, \widehat{\mathbf{r}})) < \epsilon$, and the number of pairwise comparisons requested by the algorithm is*

$$\mathcal{O}\left(\frac{M \log_2 M}{\rho^2} \log \frac{M \log_2 M}{\delta \rho} + \frac{1}{D(\epsilon)^2} \log \frac{1}{\delta D(\epsilon)} \right) ,$$

where $\rho = \frac{1-\phi}{1+\phi}$ *and*

$$D(\epsilon) = \frac{\phi}{6(\phi+1)^2}\left(1 - \frac{2}{\exp\left(\frac{\epsilon}{M(M-1)}\right)+1}\right) .$$

Proof. As a first step, the MALLOWSKL algorithm calls the MALLOWSMPR algorithm that returns a ranking $\widehat{\mathbf{r}}$ which is equal to $\widetilde{\mathbf{r}}$ with probability at least $1-\delta/2$. And then, two options i and j are selected for which $\widehat{r}_i = \widehat{r}_i + 1$. As a second step, the MALLOWSKL algorithm is comparing options i to option j in order to obtain an estimate for ϕ. Let us denote the length of the confidence interval of parameter ϕ by $C = \phi_U - \phi_L$ where

$$\widehat{p}_{i,j} - c_{i,j} \le \frac{1}{1+\phi} \le \widehat{p}_{i,j} + c_{i,j} \Rightarrow$$
$$\underbrace{\frac{1}{\widehat{p}_{i,j} + c_{i,j}} - 1}_{=\phi_L} \le \phi \le \underbrace{\frac{1}{\widehat{p}_{i,j} - c_{i,j}} - 1}_{=\phi_U}$$

By setting the confidence interval $c_{i,j}$ to $\sqrt{\frac{1}{2n_{i,j}}\log\frac{8n_{i,j}^2}{\delta}}$, we have

$$\mathbb{P}\left(|p_{i,j} - \widehat{p}_{i,j}| \ge c_{i,j}\right) \le 2\exp(-2c_{i,j}^2 n_{i,j}) = \frac{\delta}{4n_{i,j}^2}$$

where we used the Hoeffding-bound (Hoeffding, 1963). This means that $p_{i,j} \in [\widehat{p}_{i,j} - c_{i,j}, \widehat{p}_{i,j} - c_{i,j}]$ for every time step with probability at least $1-\delta/2$, thus $\phi \in [\phi_L, \phi_U]$ again holds with probability at least $1-\delta/2$ for every time step. Based on (5), we know that, when the MALLOWSKL algorithm terminates, it holds that $\mathrm{KL}(\mathbb{P}(.|\phi, \widetilde{\mathbf{r}}), \mathbb{P}(.|\phi', \widetilde{\mathbf{r}})) < \epsilon$ with probability at least $1-\delta$ for any $\phi' \in [\phi_L, \phi_U]$, because $\phi \in [\phi_L, \phi_U]$ with probability at least $1-\delta/2$ and $\widehat{\mathbf{r}} = \widetilde{\mathbf{r}}$ holds again with probability at least $1-\delta/2$. Therefore the algorithm is correct.

In order to calculate the sample complexity, let us upper bound (5) by using C. Recall that with probability at least $1-\delta/2$, we have

$$\begin{aligned}
\mathrm{KL}(\mathbb{P}(.|\phi, \widetilde{\mathbf{r}}), \mathbb{P}(.|\widehat{\phi}, \widetilde{\mathbf{r}})) &\le \frac{M(M-1)}{2}\log\frac{\phi}{\widehat{\phi}} + \log\frac{Z(\widehat{\phi})}{Z(\phi)} \\
&\le \frac{M(M-1)}{2}\log\frac{\phi_U}{\phi_L} + \log\frac{Z(\phi_U)}{Z(\phi_L)} \\
&\le \frac{M(M-1)}{2}\log\frac{\phi + C}{\phi - C} + \log\frac{Z(\phi+C)}{Z(\phi-C)} \qquad (11)
\end{aligned}$$

where (11) follows from the fact that $Z(.)$ is monotone increasing function. Now let us assume that the first term in (11) is $\le \epsilon/2$, which results in the following upper bound for C:

$$\begin{aligned}
\frac{M(M-1)}{2}\log\frac{\phi+C}{\phi-C} &\le \frac{\epsilon}{2} \\
\frac{\phi+C}{\phi-C} &\le \exp\left(\frac{\epsilon}{M(M-1)}\right) \\
C\left[1 + \exp\left(\frac{\epsilon}{M(M-1)}\right)\right] &\le \phi\left[1 - \exp\left(\frac{\epsilon}{M(M-1)}\right)\right] \\
C &\le \phi\left(1 - \frac{2}{\exp\left(\frac{\epsilon}{M(M-1)}\right)+1}\right) \qquad (12)
\end{aligned}$$

Next, let us upper bound $\log Z(\phi + C)$ from (11). Writing concisely $d_{\mathbf{r}}$ for $d(\widetilde{\mathbf{r}}, \mathbf{r})$, we have

$$\log Z(\phi + C) = \log \sum_{\mathbf{r}\in S_M} (\phi+C)^{d_{\mathbf{r}}} < \log \sum_{\mathbf{r}\in S_M} (\phi + C) = \log n!\,(\phi+C) = \log n! + \log(\phi + C)$$

And, in a similar way, one can lower bound $\log Z(\phi - C)$ as

$$\begin{aligned}\log Z(\phi - C) &= \log \sum_{\mathbf{r} \in S_M} (\phi - C)^{d_{\mathbf{r}}} > \log \sum_{\mathbf{r} \in S_M} (\phi - C)^{M(M-1)/2} = \log n! \, (\phi - C)^{M(M-1)/2} \\ &= \log n! + \frac{M(M-1)}{2} \log (\phi - C)\end{aligned}$$

Now, we can upper bound the second term of (11) as

$$\begin{aligned}\log \frac{Z(\phi + C)}{Z(\phi - C)} &\leq \log n! + \log (\phi + C) - \log n! - \frac{M(M-1)}{2} \log (\phi - C) \\ &\leq \log (\phi + C) - \log (\phi - C)\end{aligned}$$

Now let us assume that the second term in (11) is $\leq \epsilon/2$, which results in the following upper bound for C:

$$C \leq \phi \left(1 - \frac{2}{\exp\left(\frac{\epsilon}{2}\right) + 1} \right) \ . \tag{13}$$

Now we can rewrite C as

$$\begin{aligned} C = \phi_U - \phi_L &= \frac{1}{\widehat{p}_{i,j} - c_{i,j}} - 1 - \frac{1}{\widehat{p}_{i,j} + c_{i,j}} + 1 \\ &\leq \frac{1}{p_{i,j} - 3c_{i,j}} - \frac{1}{p_{i,j} + 3c_{i,j}} \\ &\leq \frac{1}{1/(\phi+1) - 3c_{i,j}} - \frac{1}{1/(\phi+1) + 3c_{i,j}} \qquad (14) \\ &\leq \frac{6c_{i,j}}{1/(\phi+1)^2 - 9c_{i,j}^2} \\ &< 6c_{i,j}(\phi+1)^2 \end{aligned}$$

where (14) follows from Corollary 3.

Summarizing the calculation above, the width of the confidence interval C for $\widehat{\phi}$ needs to satisfy the inequalities given in (12) and (13) for $\mathrm{KL}(\mathbb{P}(.|\phi, \widetilde{\mathbf{r}}), \mathbb{P}(.|\widehat{\phi}, \widetilde{\mathbf{r}})) < \epsilon$. Since (12) is uniformly tighter than (13), therefore the following needs to be satisfied

$$\begin{aligned} C < 6c_{i,j}(\phi+1)^2 &\leq \phi \left(1 - \frac{2}{\exp\left(\frac{\epsilon}{M(M-1)}\right) + 1} \right) \\ \sqrt{\frac{1}{2n_{i,j}} \log \frac{8n_{i,j}^2}{\delta}} = c_{i,j} &\leq \underbrace{\frac{\phi}{6(\phi+1)^2} \left(1 - \frac{2}{\exp\left(\frac{\epsilon}{M(M-1)}\right) + 1} \right)}_{= \text{constant and denoted by } D(\epsilon)} \end{aligned} \tag{15}$$

The MALLOWSKL first calls the MALLOWSMPR algorithm in order to find the central ranking with $\delta/2$, thus the expected number of pairwise comparison of MALLOWSKL is at most

$$\mathcal{O}\left(\frac{M \log_2 M}{\rho^2} \log \frac{M \log_2 M}{\delta \rho} + \frac{1}{D(\epsilon)^2} \log \frac{1}{\delta D(\epsilon)} \right) \tag{16}$$

where the first term is the sample complexity of MALLOWSMPR based on Theorem 5, and the second term is obtained from (15). This completes the proof. □

E. Experiments for identifying the most preferred item

E.0.1. Numerical experiments

In Figure 4, we conducted the same experiments like in Section 7.1, but with various M.

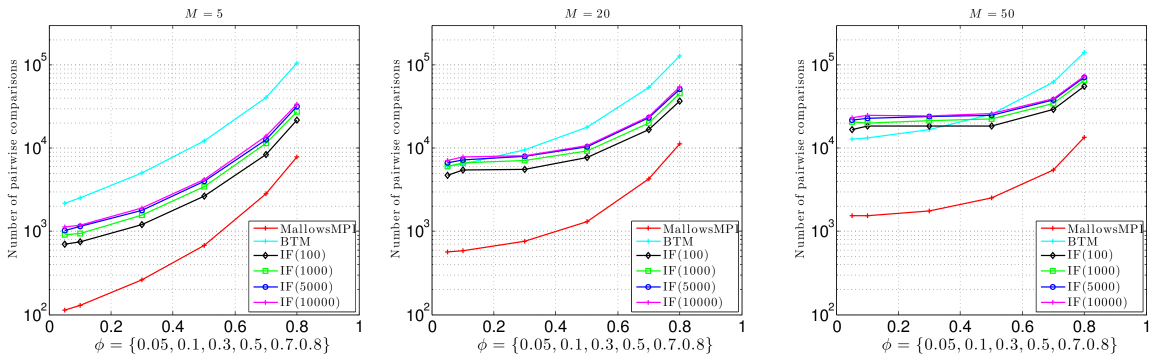

(a) $M = 5$ (b) $M = 20$ (c) $M = 50$

Figure 4. Empirical sample complexity for various parameter setting of ϕ. The number of items M is $\{5, 20, 50\}$. The results are averaged out for 100 repetition. The confidence parameter was set to 0.05, and thus, the accuracy was significantly higher than $1 - \delta$ in every case.

E.0.2. Results on real data

We conducted experiments on real data to asses the efficiency of our method if the model assumption is violated to some extent. We used various datasets taken from the PrefLib ranking data repository[6]. The most important statistics of the datasets are shown in Table 1. Each dataset consists of full and partial ranking[7]. To asses to what extent the data fit to a Mallows ϕ-model, we calculated a statistic based goodness-of-fit statistic as follows. We compute an estimate for all $p_{i,j}$ based on the data which we denote $\widehat{\widehat{p}}_{i,j}$. Then we fit a model by using the method of Cheng et al. (2009). Based on this fitted model, we computed the $p_{i,j}$ values according to Theorem 2. The statistic we calculated then is $\chi^2 = \sum_{i \neq j} (\widehat{\widehat{p}}_{i,j} - p_{i,j})^2 / \widehat{\widehat{p}}_{i,j}$. This statistic reflects to how the pairwise marginal probabilities computed based on the data fits to the ones that are computed based on the fitted model. Clearly, the more close χ^2 is to 1, the more the data fits to a Mallow's ϕ-model (at least in terms of pairwise marginals).

Table 1. The most important statistics of the ranking datasets.

Data set	ID	Items	Rankings	χ^2
1	ED-7-71	7	239	0.17
2	ED-7-35	6	32	0.87
3	ED-7-76	5	110	0.91
4	ED-7-51	6	60	1.52
5	ED-7-64	4	33	1.52
6	ED-7-37	10	446	1.58
7	ED-7-50	10	82	6.65

We run the MallowsMPI, BeatTheMean and IF(10000) on the datasets used 100 times, and we plotted the average of their sample complexity which is shown in Figure 5. The results reveal a few general trends. First, the MallowsMPI achieves lower sample complexity on almost every datasets. Second, the improvement of MallowsMPI in terms of sample complexity is more pronounced on the datasets which meets more with our modeling assumption, namely, it fits

[6]http://www.preflib.org

[7]Each dataset we considered contain complete, transitive, and anti-symmetric relation over a group of objects. We handle the outcome of the comparison of two items with the same rank as tie. Essentially, this means that these outcomes are treated in a neutral way.

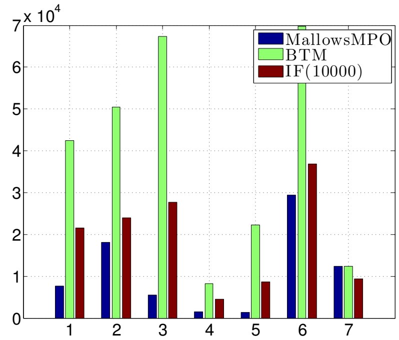

Figure 5. The empirical sample complexity of various methods on real data.

better to Mallow's ϕ-model. Note that, the better the fit, the more close χ^2) value is to 1. Finally, the BEATTHEMEAN achieves the worst sample complexity which might be explained by the fact that its modeling assumption is the most relaxed among the methods we tested in our experiments.

E.1. Merge sort vs. Quick sort

Our MALLOWSMPR algorithm is based on top-down, two-way merge sort, but as we pointed out, other sorting algorithm could be also extended so as it will be amenable to find the most probable ranking in our online learning framework. We described the extended version of quick sort in Appendix B which is called MALLOWSQUICK. The quick sort algorithm is considered one of the most efficient one in practice among the sorting algorithms, therefore we compared the quick sort based algorithm called MALLOWSQUICK to the merge sort based one called MALLOWSMPR. In Figure 6, we plotted the sample complexity achieved by MALLOWSMPR and MALLOWSQUICK. We run the algorithms with various underlying Mallows model with parameters from a wide range ($M \in \{50, 100, 200\}$ and $\phi \in \{0.1, 0.3, 0.5, 0.7, 0.8\}$). The center ranking was selected uniformly random in each run. We found that the MALLOWSQUICK achieves marginally higher sample complexity than MALLOWSMPR. This can be explained by that the confidence interval in MALLOWSQUICK is calculated with δ/M^2, because the worst case performance of quick sort is $\mathcal{O}(M^2)$, whereas the confidence interval for MALLOWSMPR is calculated with only $\mathcal{O}(M \log_2 M)$.

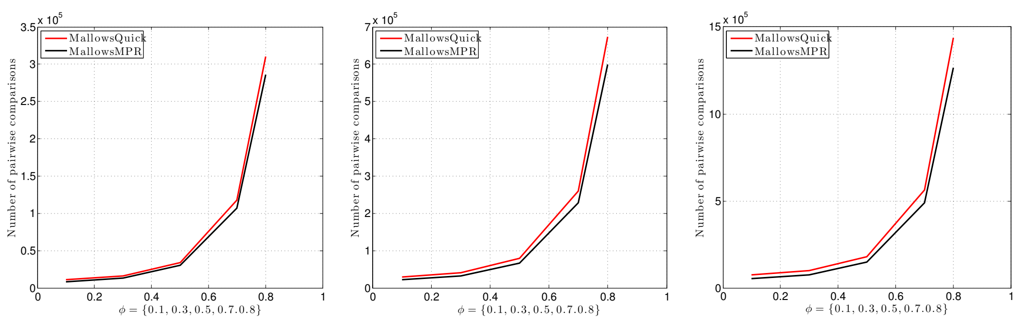

(a) $M = 50$ (b) $M = 100$ (c) $M = 150$

Figure 6. Empirical sample complexity of MALLOWSMPR and MALLOWSQUICK for various parameter setting of ϕ. The number of options M was set to 50, 100 and 150. The results are averaged out for 100 repetition. The confidence parameter was set to 0.05 for every run.