A Illustration of Quantities Introduced in Section 2

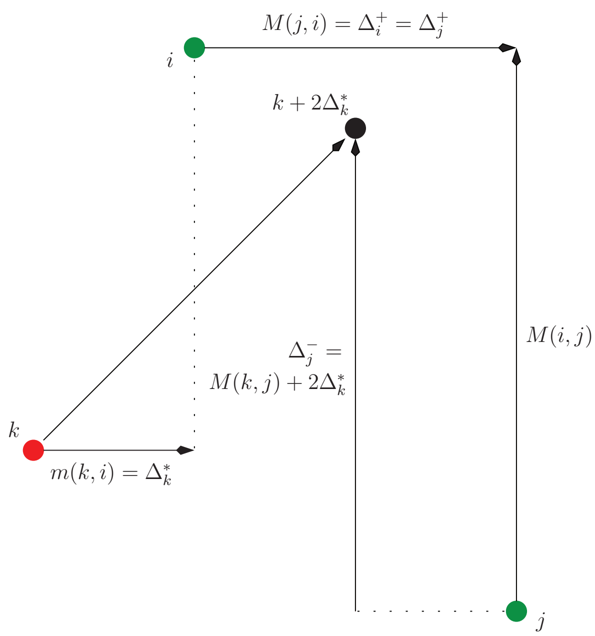

Figure 6: Example for the illustration of the quantities introduced in Section 2.

Figure 6 gives an example with two objectives, where operating points i and j are on the Pareto front, while point k is not. Moving k to the right by $m(k,i)$, it would not be strongly dominated by i anymore. This is also the distance of k to the Pareto front, that is, $\Delta^* = m(k,i)$.

The quantity $M(j,i)$ gives the amount by which point i has to be moved to dominate j. Since in our example $M(j,i) < M(i,j)$, this is also the (first component of the) required accuracy[7] for i and j, that is, $\Delta_i^+ = \Delta_j^+ = M(j,i)$.

Finally, moving k by $2\Delta_k^*$, the modified point would be on the Pareto front. Then the (second component of the) required accuracy for j would be $\Delta_j^- = M(k,j) + 2\Delta_k^*$, since if j is moved by this amount it would appear to dominate the modified operating point k.

B Proof of the Lower Bound

Proof sketch for Theorem 17. We construct small modifications of the given operating points, such that the output of the algorithm has to change to remain correct. We argue that these small changes can only be detected reliably if the operating points are sampled sufficiently often, as given by the theorem.

We use distributions $\mathcal{D}_i$ with fully correlated objective values: $\mathcal{D}_i\{\boldsymbol{y}_i + \frac{1}{4}(1,\dots,1)\} = p_i$ and $\mathcal{D}_i\{\boldsymbol{y}_i - \frac{1}{4}(1,\dots,1)\} = 1 - p_i$, where p_i is chosen to give the desired mean. For the given operating points $p_i = \frac{1}{2}$. By choosing $p_i \in [\frac{1}{4}, \frac{3}{4}]$ the operating point can be increased or decreased by up to $\frac{1}{8}$.

We now consider the modifications of the operating points that require a change of the output of the algorithm:
Let $i, j \in P^*$ with $\Delta_i^+ = M(i,j)$. Then without modification both i and j have to be in the output P of the algorithm. With the modification $\boldsymbol{y}_i' = \boldsymbol{y}_i - 3\tilde{\Delta}_i^{\epsilon_0}$, we have $m(i,j) \geq 2\epsilon_0$ such that i must not appear in the output P.
If $i, j \in P^*$ with $\Delta_i^+ = M(j,i)$, then with the modification $\boldsymbol{y}_i' = \boldsymbol{y}_i + 3\tilde{\Delta}_i^{\epsilon_0}$, we have $m(j,i) \geq 2\epsilon_0$ such that j must not appear in P.
Let $i \notin P^*$. Since the unmodified suboptimal point i may or may not appear in a correct output P, we consider two cases, depending on the probability that $i \in P$. First we assume that for the unmodified $\boldsymbol{y}_i$, with probability $\geq \frac{1}{2}$ the algorithm outputs a P with $i \notin P$. Then we consider the modification $\boldsymbol{y}_i' = \boldsymbol{y}_i + 2\tilde{\Delta}_i^{\epsilon_0}$. This makes the modified operating point Pareto optimal such that it has to appear in P. Secondly we assume that with probability $\geq \frac{1}{2}$ the unmodified $\boldsymbol{y}_i$ is in P (because i is close to the Pareto front). Then the modified point $\boldsymbol{y}_i' = \boldsymbol{y}_i - 2\tilde{\Delta}_i^{\epsilon_0}$ has distance $\geq 2\epsilon_0$ from the Pareto front and must not appear in the output P.

We see that the algorithm has to detect any of these changes to provide a correct output. From a standard argument it follows, that with probability δ the algorithm cannot distinguish between a modified and an unmodified operating point i, if the operating point is not sampled $\Omega\left(\frac{1}{(\tilde{\Delta}_i^{\epsilon_0})^2}\log(\frac{1}{\delta})\right)$ times, see e.g. Even-Dar et al. (2002). Summing over all operating points gives the lower bound of the theorem. □

[7] Recall that the required estimation accuracy for a point j on the Pareto front is defined as $\Delta_j = \min\{\Delta_j^+, \Delta_j^-\}$.