

Figure 7 — What the split choice is worth (measured, not assumed); axis runs 0-1 so the no-skill baseline stays visible
the Dummy baseline (0.500 / 0.577) and the 1-NN baseline are non-fitted references, not part of the trainable-model range

