Diversity, Ability, and Expertise in Epistemic Communities

Author Names Removed for Blind Review

Online Supplement

In "Diversity, Ability, and Expertise in Epistemic Communities," we explored how three variations on the Hong-Page "diversity trumps ability" model of group inquiry show that the story about the interplay between diversity and ability is more nuanced than many related discussions let on. In that paper, we show how Hong and Page's result importantly depends on the landscape being completely random, how the number of heuristics affects the outcome when the landscapes are not random, and how the method of communication that groups use affects their performance. Here we explore two more variations on the model: (1) how groups composes of some of the best-performing agents and some random agents perform, and (2) how the size of the groups affect the original result. This document is a meant as a supplement to our original paper, so readers are encouraged to consult that paper first.

S1. Mixed Groups

In the original Hong-Page model, only groups of the best performing agents and groups of random agents are considered. Interesting results occur when we compare groups with some proportion of each. The difference between tournament and relay group dynamics, discussed more thoroughly in our original paper, also plays out on these mixed groups. Returning to our initial pool of 12 heuristics, we show results for relay dynamics in Figure 9.

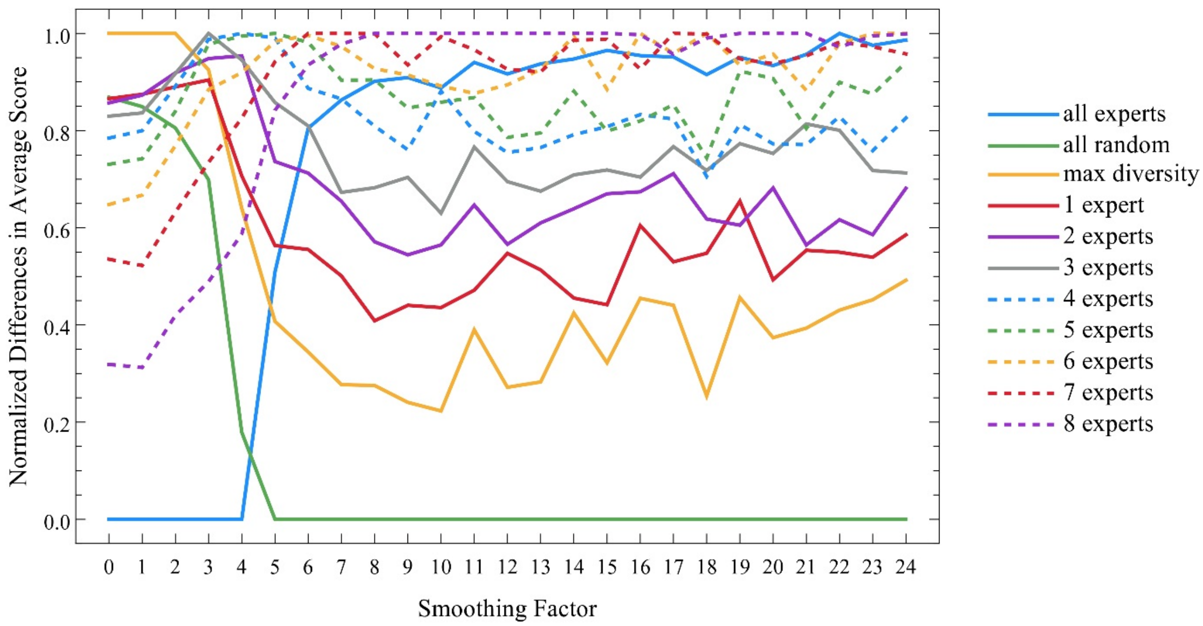

Figure 9. Normalized averages for pure and mixed groups of size 9 using a relay group dynamics

In Figure 9 that group that scores the highest average at any smoothness factor is normalized to 1, that which scores the lowest is normalized to 0. The dramatic cross-over between best performers and randoms thus plays out in the cross-over between smoothness 5 and 6. For 5 and below it is best performers that do the worst; for 6 and above it is the random group that performs worst.

We track not merely groups of best performers and random heuristics but others as well. Our analysis of the Hong-Page result, consistent with indications in their work, is that what groups of the 'best-performing' have going against them is redundancy. They are too much alike, thereby losing the exploratory spread of individuals selected randomly. If that analysis is right, a 'group' consisting of a single best agent should have even higher redundancy and so should do worse. That is indeed the result for the range shown: a single best agent does so far worse than any of those groups shown in Figure 9 that we left it off rather than distort the readability of the chart.[1]

Contrary to Thompson (2014), we have suggested that it is not the randomness of random groups that is an epistemic virtue, but the extent to which their heuristics jointly cover the available space. If that analysis is right, a 'maximum diversity' group should do better than a random group. Hong and Page (2004, 16386) give a measure of diversity that doesn't entail that maximally diverse groups cover the heuristic space as much as they can. If instead, so instead, we construct groups with maximum diversity by choosing a random heuristic for the first member but then constraining successive choices for later members of the group so as to duplicate heuristic numbers as little as possible. If the number 11 is already 'taken' in the assignment of heuristics to three previous group members, for example, it is removed as a candidate for further assignments. Figure 9 shows the performance of such a group significantly outperforming a group of random agents at every point.

[1] Even this can be complicated by additional factors. In relay dynamics (though not simultaneous), with a heuristic pool of 24 and smoothness factor between 20 and 30 the single best performers outperforms the random group.

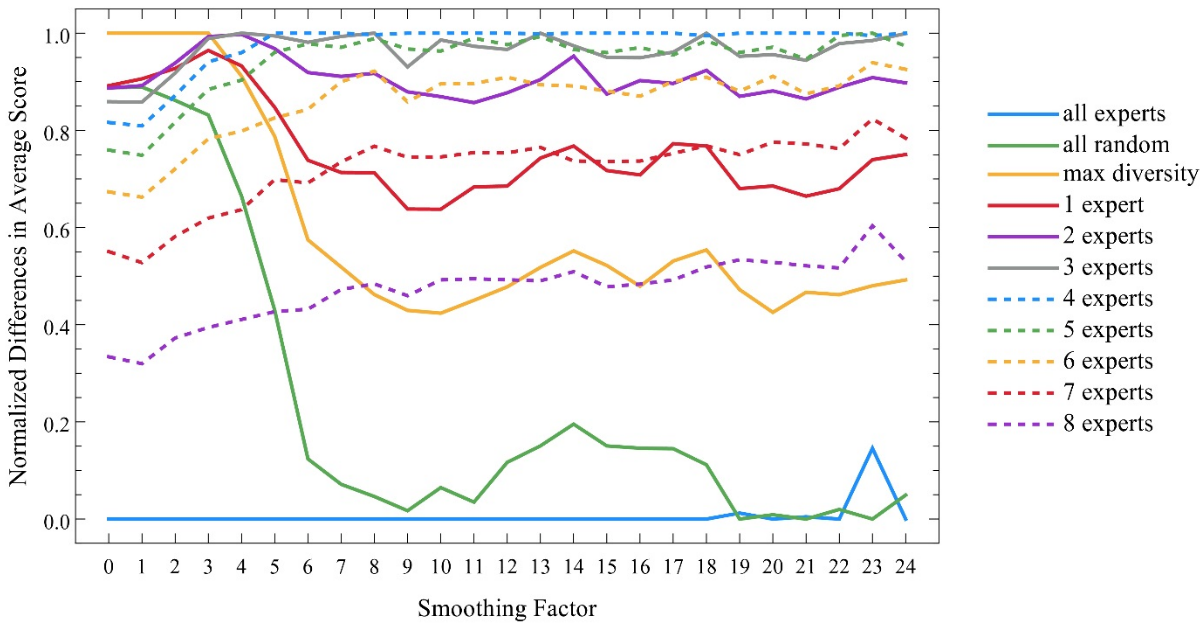

Figure 10 Normalized averages for pure and mixed groups of size 9 using a tournament group dynamics

The other groups shown are composed of 1 best performer with 8 random, 2 best performers with 7 random, and so forth. For smoothness factors above 8 the pattern is clear. Even 1 best performer among randoms does better than our max diversity group. Groups with 2 best performers do better still, and the 'added value' of best performers in place of randoms increases to roughly 7 or 8, though groups with 1 or 2 random members still do better than groups composed of best performers alone. The difference between tournament and relay dynamics is clear when we construct a similar group for tournaments results, shown in Figure 10.

It remains true in tournament play that a single best performer does so far worse than others that he is left out in order to avoid distorting the chart. It also remains true that a 'maximum diversity' group outperforms a random group at all points.

In tournament dynamics, however, groups of 9 best performer do far worse than in relay dynamics, occupying the normalized bottom of the chart at almost all points. Random groups do slightly better than best performer at most points, though they tend to join them at the bottom with higher smoothing factors.

The performance of mixed groups in tournaments is particularly interesting and importantly different than in the case of relay dynamics. In tournament dynamics, mixed groups with at least one best performer and at least one random agent do better than either just best performers or just randoms at all points. If one traces groups with 1 best performer, 2 best performers, 3 best performers, or 4, these score progressively higher values in the graph. From that point, however, increasing the percentages of best performers proves a disadvantage: groups with 5, 6, 7, or 8 best performers do worse than those with 3 or 4. We have found very similar results showing advantage to roughly half-mixed groups with heuristic pools of 24 in place of 12.

S2. Group Size

Are there other parameters that advantage either diversity or expertise? In stating their original result, Hong and Page require that groups be 'good sized'; the groups of 9 we have used throughout are very close to the groups of 10 they use in simulation. In limited confirmation of that requirement, it appears that larger groups—at least up to a limit—favor groups of random heuristics over groups of the best performing agents. Results hold for both relay and tournament dynamics.

We replicated the smoothing factor and max heuristic graphs of the previous sections for groups of 3 and 6 rather than 9. Previous results for 9 are included for comparison. Up to groups of 9, it is clear that increased group size in relay dynamics advantages the relative performance of groups of random heuristics over groups of the best performing agents. Figures 14 through 16 show similar comparisons for simultaneous tournament dynamics.

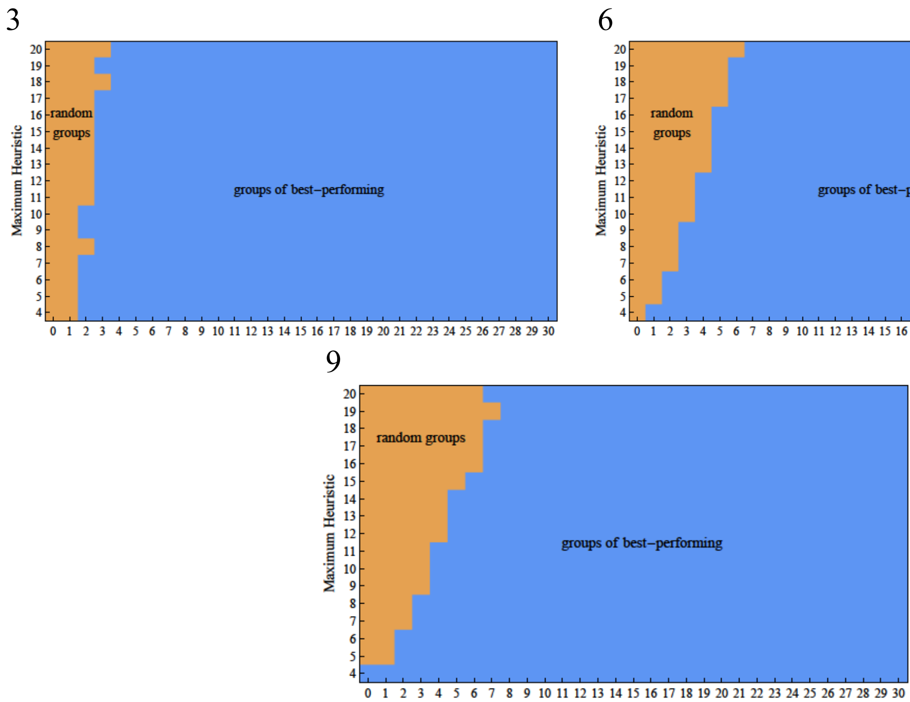

Figure 11 Relay: areas in which groups of random heuristics do best (brown) and areas in which groups of the best-performing do best (blue) across a parameter sweep of landscape smoothness and max heuristic for groups of 3, 6, and 9

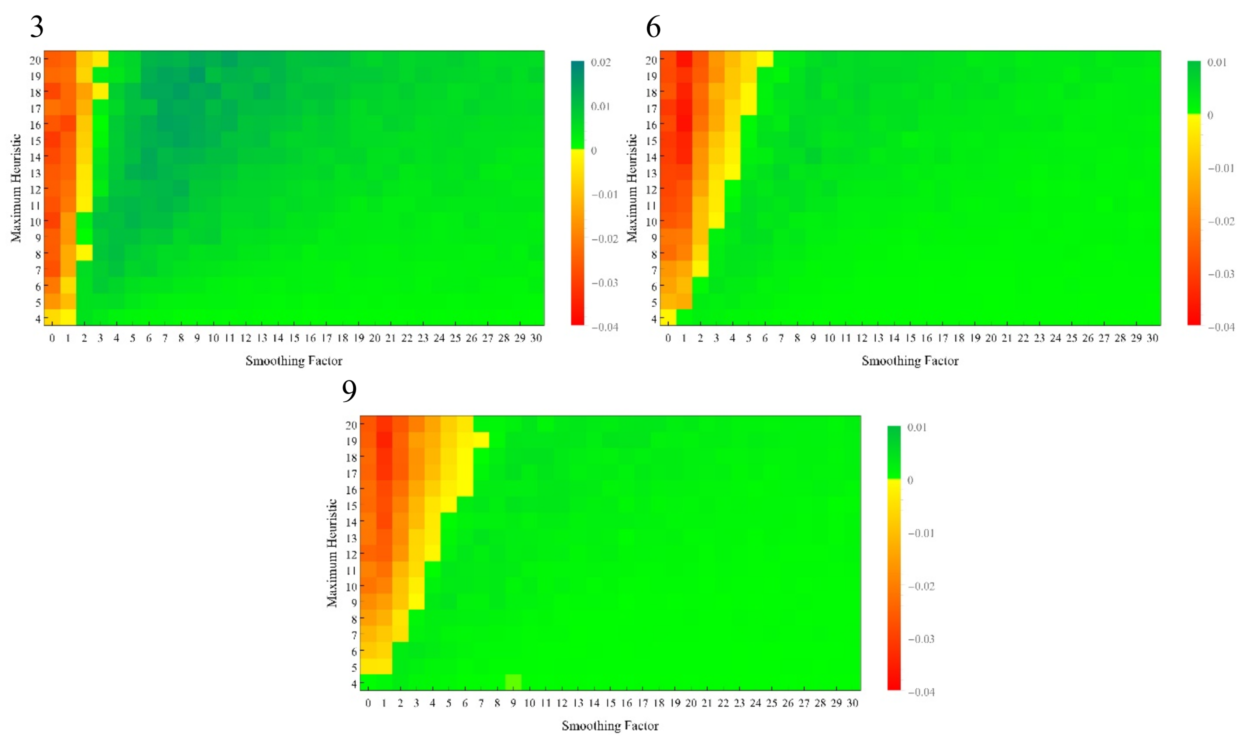

Figure 12 Relay: Differences in averages for groups of random heuristics and groups of the best-performing for groups of 3, 6, and 9.

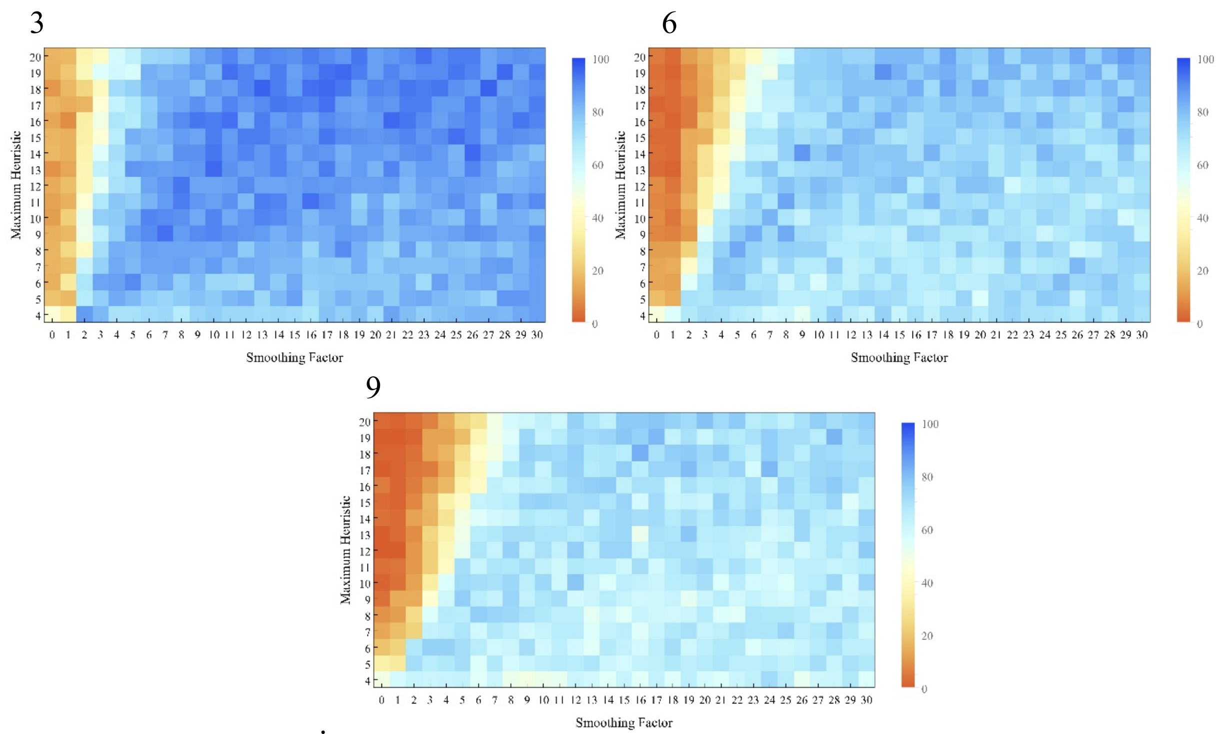

Figure 13 Relay: Percentages of runs in which groups of the best-performing do better than groups of random heuristics.

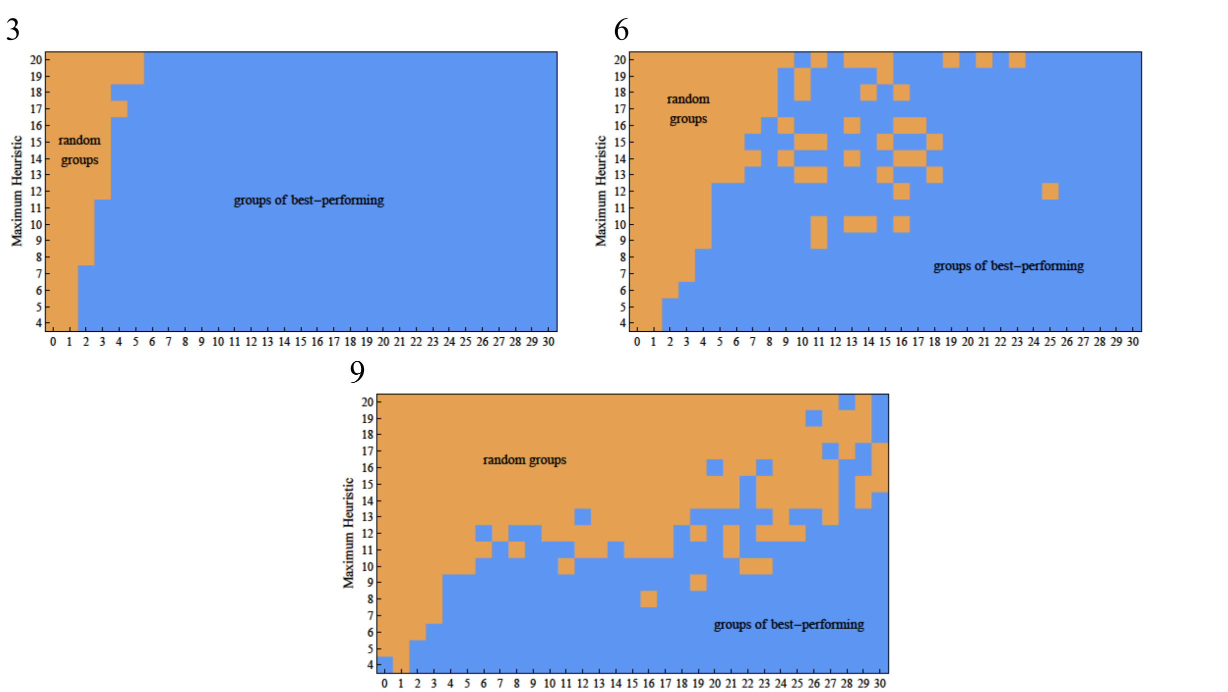

Figure 14 Tournament: areas in which groups of random heuristics do best (brown) and areas in which groups of the best-performing do best (blue) across a parameter sweep of landscape smoothness and max heuristic for groups of 3, 6, and 9

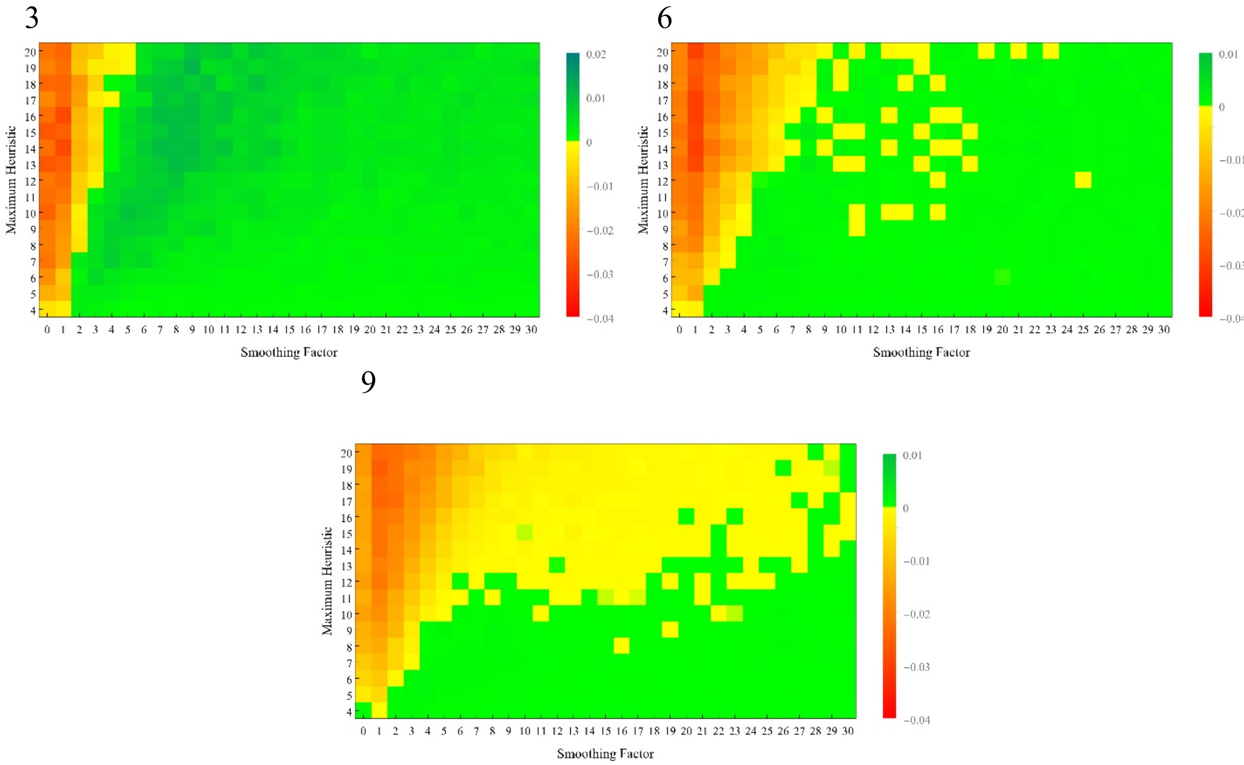

Figure 15 Tournament: Differences in averages for groups of random heuristics and groups of the best-performing for groups of 3, 6, and 9.

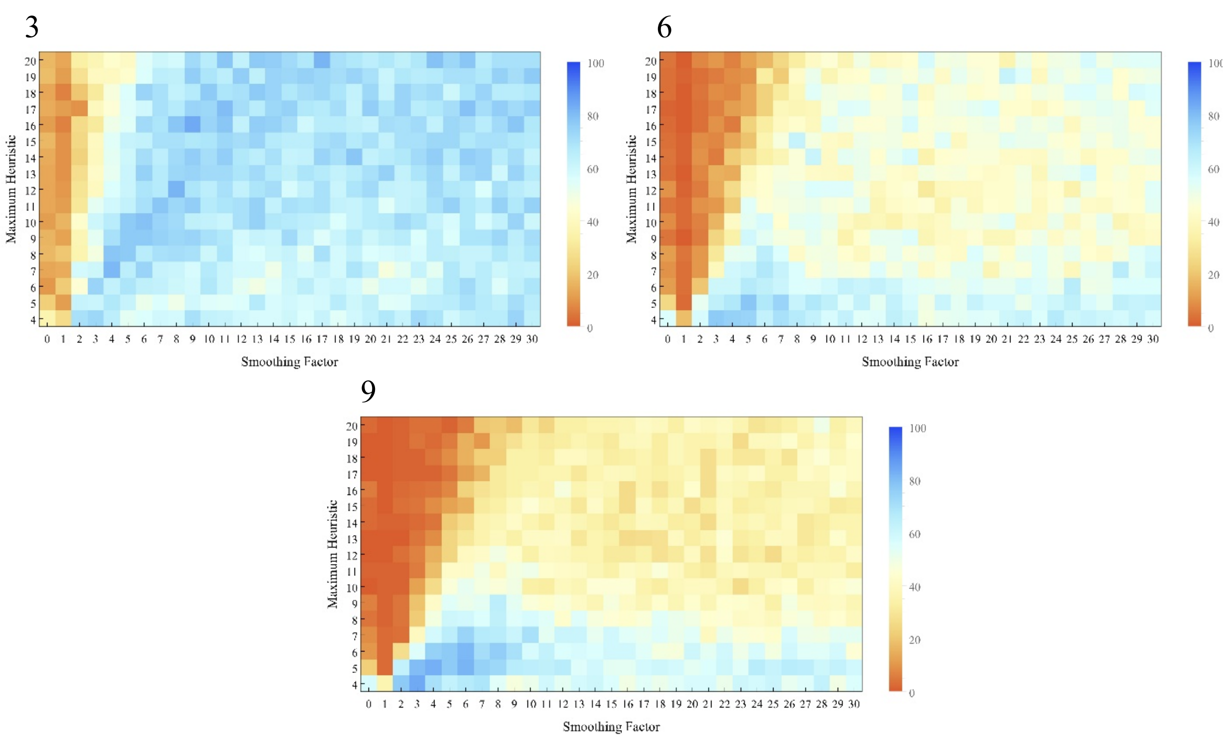

Figure 16 Tournament: Percentages of runs in which groups of the best-performing do better than groups of random heuristics.

The smaller the group, these figures indicate, the greater the advantage of the best performers. The larger the group, all things considered, the greater the advantage for random heuristics. Here again heuristic coverage offers an explanation. Larger groups of best performers will have a small reduction in redundancy, since our best performers are genuinely distinct, but larger groups of random heuristics can still be expected to outstrip the best-performing in terms of coverage on the available heuristic set. At least one reason why larger groups favor random heuristics over best performing heuristics is because the ratio of coverage for randoms over best performers increases with group size.

References

Hong, Lu & Scott E. Page (2004). Groups of diverse problem solvers can outperform groups of high-ability problem solvers. *Proceedings of the National Academy of Sciences. 101*, 16385-16389.

Thompson, Abigail (2014). Does diversity trump ability?: An example of the misuse of mathematics in the social sciences. *Notices of the American Mathematical Society*, 61, 1024-1030.