The Strategy Aggregation Effect in Judgment

Henrik Olsson

Santa Fe Institute

Author Note

Correspondence concerning this article should be addressed to Henrik Olsson, Santa Fe Institute, 1399 Hyde Park Rd, Santa Fe, NM 87501, olsson@santafe.edu

Abstract

Predicting a criterion that is probabilistically related to pieces of information, or cues, is a paradigmatic judgment task that has been investigated both, in research trying to identify the individual judgment and decision making strategies people use, and in the wisdom-of-crowds literature where the focus is on how aggregation can improve accuracy. I combine these two lines of research to investigate how the performance of individual and aggregated linear strategies are affected by different environmental and group aggregation factors and how performance differences between individual and aggregated linear strategies can be understood in a unified framework. I show that constrained linear strategies are more accurate for individual judgments, but when these judgments are averaged, an unconstrained linear strategy is more accurate. This *strategy aggregation effect* can be understood by analyzing a decomposition of the mean squared error into bias, variance, and covariance. Because of their lower bias but higher variance, unconstrained linear strategies perform worse for individual judgments, but better for averaged judgments where aggregation minimizes variance. In simulations with artificial and real environments, I further show that this aggregation effect does not occur if there are correlations between individual judgments. Here, constrained linear strategies always outperform an unconstrained linear strategy, because the larger covariance component of the unconstrained linear strategy outweighs its lower bias. I end with real-world implications of the results for cognitive strategies and decision environments in group and organizational settings.

Keywords: individual judgment, group judgment, constrained linear models, wisdom of crowds, mean squared error decomposition

Disclosures and Acknowledgments

The author has no conflict of interest and is grateful to Mirta Galesic and Klaus Fiedler for helpful comments on a previous version of the manuscript.

The Strategy Aggregation Effect in Judgment

Predicting a criterion that is probabilistically related to pieces of information, or cues, is a paradigmatic judgment task that takes many forms: An analyst examining financial indicators to predict corporate bankruptcy, a manager using behavior in interviews to assess job candidates, or a physician looking at symptoms that indicate the severity of a disease. What is the best strategy to use when making such judgments? This question has been extensively studied in the literature that investigates what strategies individuals are using to make judgments and decisions.

Depending on the structure of the environment and the uncertainties involved in knowing this structure, the answers have ranged from integrating and weighting of all available information (Anderson, 1980; Hammond, 1996) to the use of more constrained linear strategies, such as equal weighting of information, or cues (Dana & Dawes, 2004; Dawes, 1979; Dawes & Corrigan, 1974) and single cue heuristics (Gigerenzer, Hertwig, & Pachur, 2011; Gigerenzer, Todd, & the ABC Research Group, 1999; Hogarth & Karelaia, 2007). This line of research has mostly been concerned with the identification and performance of individual strategies. Another line of research has investigated the same or similar judgment tasks from the perspective of how aggregation of judgments can improve the accuracy of judgments. Starting with the classic study by Galton (1905), where he averaged the predictions of weight judgments of a dressed ox at a county fair, most studies in this wisdom-of-crowd's line of research have mostly sidestepped the issue of how characteristics of individual judgment strategies might affect the accuracy of averaged predictions (for overviews, see e.g., Page, 2007; Sunstein, 2006; Surowiecki, 2004).

In this paper I combine these two lines of research to investigate if linear judgment strategies that work well for individual judgments also work well when these judgments are aggregated. Specifically I ask two questions, 1) How is the performance of individual and

aggregated linear strategies affected by different environmental and group aggregation factors?, and 2) Can performance differences between individual and aggregated linear strategies be understood in a unified framework?

I document a *strategy aggregation effect*, where a constrained linear strategy (i.e., equal weighting of cues or a single cue strategy) that has a high predictive accuracy for individual judgments, works less well for aggregated judgments. When judgments are aggregated, a less constrained linear strategy (i.e., a multiple linear regression like strategy) can achieve a higher predictive accuracy than more constrained linear strategies. This strategy aggregation effect is modulated by several environmental and group aggregation factors. Most notably, the strategy aggregation effect is strongest when the judgment task is difficult, the experience with the environment is low, and the correlation between judgments is low. When the correlation between judgments is high, the strategy aggregation effect is reduced.

I further show that this aggregation effect can be understood by analyzing a decomposition of the mean squared error (MSE) into bias, variance, and covariance (Ueda & Nakano, 1996). In a setting where we want to assess the predictive accuracy of an individual strategy, bias is the (squared) expected difference between the judgments and the criterion and the variance is simply the expected variance of those judgments. When judgments are averaged, MSE can be further decomposed with a covariance component, which is the expected pair-wise covariance between individual judgments (for a formal description, see the Appendix). The bias and variance in this decomposition is the average of the individual biases and variances. MSE decompositions have been used in machine learning (e.g., Geman, Bienenstock, & Doursat, 1992) and judgment and decision making research (e.g., Davis-Stober, Budescu, Broomell, & Dana, 2015), but have yet to be used to connect the characteristics of individual judgment

strategies with the performance of aggregated judgments.

In the following, I reinterpret the results from research on individual linear judgment strategies and wisdom-of-crowds research in the bias-variance-covariance framework to arrive at predictions when a less, or more, constrained linear strategy will achieve higher, or lower, predictive accuracy. From these two lines of research I identify five environmental, and two group aggregation factors that have been shown to influence the accuracy in individual judgment tasks and in aggregated judgments. I relate these factors to the different components in the MSE decomposition. Specifically, for individual judgments a constrained linear strategy will have lower variance than an unconstrained linear strategy. If the bias component of the constrained linear strategy is low enough, the constrained linear strategy will outperform an unconstrained linear strategy. This is especially likely when the sample size is small relative to the number of cues and when the task is difficult, making optimal linear weights hard to estimate. For aggregated judgments, the higher variance of the unconstrained linear strategy matters less, because the variance is reduced by aggregation. The lower bias of the unconstrained linear strategy will make the unconstrained linear strategy perform better than a constrained linear strategy. I describe the results from simulations showing when the strategy aggregation effect occurs in artificial and real environments, how different environmental and group aggregation factors affect the strategy aggregation effect, and how analyzing the MSE decomposition can be used to understand why the strategy aggregation effect occurs. I conclude by discussing real-world implications of the results in the context of how organizations should structure their decision environments.

Unconstrained and Constrained Linear Judgment Strategies in Individual Judgment

The most thoroughly investigated class of judgment and decision making strategies are linear integration strategies (Payne, Bettman, & Johnson, 1993; Hogarth & Karelaia, 2007;

Juslin, Olsson, & Olsson, 2003). The task that confronts participants in this line of research is to predict the criterion value based on a number of cues. If the criterion is binary, the task is equivalent to a categorization task. If the criterion is continuous and the task is to choose between different options, it is usually referred to as a decision-making task. If the criterion is continuous and the task is to predict the criterion value, it is usually referred to as a judgment or estimation task. Linear integration strategies have been investigated in all of these judgment and decision making tasks. In their simplest form linear integration strategies take cue values as input, weigh them and add them together to form a prediction, either in the form of a choice or a continuous judgment. These strategies range from multiple linear regression to strategies with fixed weights (e.g., fixed equal weight strategies). They can also emulate the predictions from heuristics such as take-the-best (Gigerenzer & Goldstein, 1996).

Roughly, linear integration strategies can be divided into unconstrained and constrained strategies. Unconstrained linear strategies (or proper linear models where the weights are estimated with ordinary least squares), use all available information in the task. If the task conditions for finding the optimal weights are satisfied, they will recover the population weights when the sample size goes to infinity. Constrained linear strategies, on the other hand, might use only a subset of the available information and typically do not try to find optimal linear weights. Examples include a linear strategy that gives equal weight to all pieces of information (Equal Weight or EW strategy) or uses only one cue out of many possible (Single Cue or SC strategy).

When and Why Do Constrained Linear Strategies Outperform Unconstrained Regression in Individual Judgment?

Previous research has shown that constrained linear strategies, such as EW or SC, can outperform ordinary least squares regression (LR) in terms of predictive accuracy (e.g., Dana & Dawes, 2004; Dawes, 1979; Dawes & Corrigan, 1974). The simple explanation for this finding is

that constraining the weights in a linear strategy reduces the variance of the predictions. For example, Davis-Stober et al. (2010) showed that a constrained linear estimator developed in an ordinary least squares framework always incurs less variance than an unconstrained linear estimator. Unconstrained estimation of the parameters in a linear strategy can make the strategy overfit, that is, an apparently good fit is more indicative of noise and idiosyncrasies in the data than true patterns. If the overfit strategy is tested on new data, it will perform poorly. In general, a constrained linear estimator will improve over LR if the constrained estimator can offset its (usually higher) bias with lower variance.

Many characteristics of the task environment affect the performance of individual constrained and unconstrained linear strategies. Five of the most important ones are 1) the number of cues (e.g., Davis-Stober et al., 2010; Einhorn & Hogarth, 1975), 2) the inter-correlation among cues (Davis-Stober et al., 2010), 3) the (linear) predictability of the environment (R^2), 4) the sample size (i.e., the number of instances used to estimate the parameters), and 5) the dispersion of environmental cue weights (Martignon & Hoffrage, 2002). These characteristics are listed in Table 1 together with how each of them was implemented in the simulations I report later. In general, individual constrained linear strategies have an advantage when the number of cues is large, the inter-correlation among cues are positive, predictability is low, and sample size is small (see Davis-Stober et al., 2010 for proofs and simulations). If the dispersion of environmental cue weights is large (e.g., there is one cue that is a much better predictor than all the other cues taken together), SC would perform better, or equal to, LR when the cue-weight dispersion is small (given that SC can identify a cue with high predictive validity).

Aggregating Judgments: Wisdom of crowds

In the 1920s and 1930s, psychologists were excited about the finding that averaging

individual judgments led to more accurate estimates than those of the average individual judge. (e.g., Overstreet, 1925). After years of debating why these groups outperformed the average individual, most psychologists dismissed these effects as merely illustrating “an obvious statistical principle of reducing the error variance” (Lorge, Fox, Davitz, & Brenner, 1958, p. 345) that cannot tell us much about psychological processes. Soon, averaged judgments of so called “statisticized” groups, whose members did not interact or have face-to-face contact, were used as mere baselines against which interacting groups and best members of the groups were compared (for a review see Gigone & Hastie, 1997).

The publication of Surowiecki’s popular book *The Wisdom of Crowds* (2004), however, increased the interest in the performance of statisticized groups. The power of the wisdom of crowds has been demonstrated in medical diagnosis, prediction markets, game shows, political polls, and forecasting of political and economic events in the U.S. and elsewhere (Clemen 1989, Gigone & Hastie 1997; Wallsten et al. 1997; Graefe, 2014; 2016; Gutsche et al., 2014; Murr, 2011; Lewis-Beck & Skalaban, 1989; Miller et al., 2012; Mellers et al., 2014, Tetlock & Gardner, 2015).

Recently, there has been considerable advancement in understanding when and why aggregation produces more, or less, accurate judgments (Davis-Stober, Budescu, Dana, & Broomell, 2014; Davis-Stober et al., 2015). These advancements rely mainly on the decomposition of mean squared error into components of bias, variance, and covariance. One limitation with the current understanding of the performance of aggregated judgments is that the criticism raised in 1920’s still applies, in the sense that there is only a loose connection between the statistical explanation and the strategies that people use to make their judgments. For example, there are theoretical results that give insights in how to compose an optimally wise

crowd, but these results rest on fixed assumptions on different types of individuals in terms of, for example, bias and variance of individual judgments (Davis-Stober et al., 2015; Lamberson & Page, 2012). However, it is still not well known how different judgment strategies give rise to differences in bias and variance and how these differences interact with properties of the judgment task, such as the predictability of the environment, and properties of the group aggregation, such as the size of the group.

When and Why Do Constrained Linear Strategies Outperform Unconstrained Linear Strategies in Aggregated Judgments?

The ability of constrained linear strategies to control variance is the main mechanism that makes them outperform unconstrained linear strategies in individual judgments. In aggregated judgments, however, the size of the group controls the variance, with the mean variance diminishing with group size, (i.e., $\frac{1}{m} \overline{var}$, where $\overline{var}$ is the mean of the variance of the individual strategies). If the group size is large, the bias and covariances between the individual judgments will influence the accuracy the most. The impact of the covariance increases with the size of the group (i.e., $(1 - \frac{1}{m}) \overline{cov}$ where $\overline{cov}$ is the mean of the pair-wise covariances between individual strategy judgments)¹. Note that the covariances can be zero, which means that disagreement between judges would increase the performance of aggregated even more. Thus, the two main group aggregation factors that affects the performance of aggregated judgments are 1) group size, and 2) the covariance between the judgments (see Table 1).

When group size is large, a constrained linear strategy will outperform the unconstrained linear strategy if the sum of the bias and covariance components is smaller for the constrained

---

¹ This is the case for the MSE decomposition I use in this paper. There are, however, other decompositions that assumes that the criterion is a random variable and this assumption adds an ability component to the decomposition that consists of the covariance between the judgments and the criterion. I discuss different decompositions in the Discussion and the Appendix.

linear strategy than for the unconstrained linear strategy. However, this inequality does not specify when different components can be expected to be large or small for the different strategies.

In some circumstances, the covariance of the unconstrained linear strategy will be higher than for a constrained linear strategy: The constrained linear strategy's lower variance will limit the covariances between individual judgments. Thus, one would expect that the factors that influence the covariance among members (e.g., increasing the overlap in experience with the environment) would have a greater impact on the performance of the unconstrained linear strategy than a constrained linear strategy.

Investigating Individual and Aggregated Strategy Performance in Different Task Environments

To investigate how environmental and group aggregation factors interact with different individual linear strategies to affect aggregated judgments, I conduct a series of simulations in artificial and real environments. I use simulations because there are no analytical results that relate variance and covariance to the MSE for the constrained linear estimator (Davis-Stober et al., 2010 derived results for the variance of the weights in their constrained linear estimator framework, but these results do not translate to the variance of predictions, see the Appendix).

I will limit my investigations to the paradigmatic and well-studied judgment task of predicting criterion values from a set of cues, most similar to a task where a physician tries to determine disease severity based on a range of symptoms (e.g., Hammond, 1955; Karelaia & Hogarth, 2008). This class of tasks encompasses a wide range of real world judgment tasks. Classical examples in the literature include predicting Wechsler IQ scores from Rorschach cues (Hammond, 1955), predicting teachers' judgments of end of year reading achievement from socioeconomic status and various indicators of language comprehension at the beginning of the

school year (Cooksey, Freebody, & Davidson, 1986), and predicting advertising sales based on information related to the number of advertising pages in a certain time period (Ashton, 1982), to name just a few. In the same way as the majority of studies in this tradition, I will assume that both environments and individual judgment processes can be described with linear functions (Einhorn, Kleinmuntz, & Kleinmuntz, 1979; see Karelaia & Hogarth, 2008, for a discussion of the linear assumption in judgment). This is obviously a simplification, but a reasonable first approximation especially since there is evidence that people tend to start and stick with a linear judgment strategy even when it would be better to switch to another strategy (Olsson, Enkvist, & Juslin, 2006). In addition, the behavior of unconstrained and constrained individual linear strategies in linear environments are well understood, making it easier to choose factors in the simulations that are known to influence the predictive ability of the strategies (Davis-Stober et al., 2010).

Specifically, the simulations most closely resemble a real world setting in which an external “aggregator” receives judgments from individuals that make independent judgments from available information that might be shared between analysts. For example, a supervisor of a group of analysts must find ways to make the best possible aggregation of the analysts’ judgments in order to use the judgments for predictions. Individual analysts might have access to instances that have different number of cues (number of cues in the simulation) with varying correlations between them (cue correlation in the simulations). An individual analyst might also encounter judgment environments with varying levels of difficulty (R^2 in the simulations) and varying degrees of instances available for making a judgment (training set size in the simulations). The importance of different cues in the environment might also differ. In some environments, cues might be equally predictive of the criterion, and in others it might be that one

cue is better than all other cues (variance of beta weights in the simulations). From the perspective of the analyst's supervisor, the main concerns are how to aggregate the analysts judgments (arithmetic mean in the simulations), determine how many analysts are needed for making adequate judgments (group size in the simulations), and how to minimize the correlations between judgments (training set overlap in the simulation).

Strategies

In this section I describe the three judgment strategies used in the simulations for the artificial and real environments.

Linear Regression (LR). In the unconstrained LR it is assumed that the judgment $\hat{y}$ is the sum of the weighted cue values $c_1 \dots c_j$,

$$\hat{y} = \sum_{j=1}^J w_j \cdot c_j, \quad (3)$$

where the weights w are free parameters estimated with ordinary least squares. This strategy is equivalent to multiple linear regression.

Equal Weights (EW). In the constrained EW it is assumed that the judgment $\hat{y}$ is the sum of the equally weighted cue values $c_1 \dots c_j$,

$$\hat{y} = w \cdot \sum_{j=1}^J c_j, \quad (4)$$

where the single weight w is a free parameter estimated with the constrained linear estimator framework developed by Davis-Stober et al. (2010).

Single Cue (SC). In the constrained SC it is assumed that the judgment $\hat{y}$ is the sum of a single weighted cue value c ,

$$\hat{y} = w \cdot c_j, \quad (5)$$

where the single weight w is a free parameter estimated with the constrained linear estimator framework (Davis-Stober et al., 2010). The cue is selected to be the most predictive according to the ordinary least squares weights. This strategy is equivalent to univariate linear regression.

It is important to consider what should be exogenously given in the constrained linear strategies and what should be estimated from data. For example, getting the cue directions right can have a big influence on the predictions of constrained linear strategies (Katsikopoulos, Schooler, & Hertwig, 2010). Should the assumption be that people are using the correct cue directions or should the assumption be that people estimate the cue directions from the sample at hand and incorporate them in the weights? One solution assumes some sort of prior (possibly noisy) knowledge of the cue directions. This approach could work in a task where participants are asked about cue directions before they start the task. But if the goal is to compare the performance of different judgment strategies, a principled way of assigning cue directions in constrained linear strategies is needed. When performance comparisons between decision making strategies are made, the usual way of doing it is to estimate the directions of the cues in the training set and apply this directional knowledge in the test set (e.g., Brighton, 2006). I will follow the same strategy in the simulations I present below, by using the sign of the OLS weights as an indicator of cue directions for all strategies.

Simulations with Artificial Environments

I first conduct simulations in artificially created environments. I investigate how the five environmental and the two group aggregation factors discussed in the previous sections affect the performance of the constrained and unconstrained linear strategies (see Table 1). I manipulate the degree of covariance by having different overlapping experiences with the environment (i.e., overlapping training sets that affects the covariance between the individual judgments).

I fit groups of individual unconstrained and constrained linear strategies to part of the data. I then use the fitted parameter values to predict the criterion values on unseen instances. The average group judgment is the arithmetic mean of the judgments produced by all strategies of a specific type with no information sharing assumed between individuals apart from different

levels of overlapping training sets. All group members have the same kind of strategy, but the estimated weights will differ due to differences in the training sets. Following the linear judgment literature, I use R^2 as measure of the predictability of the environment for both the artificial and real environments (see e.g., Hogarth & Karelaia, 2007; Karelaia & Hogarth, 2008).

Environments. I generated artificial environments that varied in the five different environmental characteristics described earlier (see Table 1 for the values used in the simulations).

All cue values were generated from a multivariate normal distribution with a normally distributed error term with a zero mean and a standard deviation of one. The environments had a single-factor structure in the sense that all entries in the $p \times p$ inter-cue correlation matrix (i.e., what is usually called the design matrix X) were equal to r . For each of the four combinations of p and r , I used the procedure described in Waller and Jones (2011) to randomly generate 50 beta weight vectors for each of the ten different R^2 values). In the simulations where I estimated the bias-variance-covariance, I selected a subsets of beta weights to illustrate how the MSE components interact to produce the strategy aggregation effect. I chose three different beta vectors for each of the four combinations of p and r that produced an R^2 of .35 based on the beta vectors' minimum, maximum, and median absolute variance (see Table S1 in the Supplementary materials for the values of these beta values and the mean and standard deviations of all beta values used in the simulation). That is, from the large set of beta values that I used for the main simulation I choose the one with the minimum, maximum, and median dispersion as measured by the variance of the beta values for each instance. The reason for choosing different variances is that the variance of the beta weights should be related to the performance of the different strategies that I investigate. For example, if the variance is high, the LR strategy should work

better than the EW strategy. The results for the minimum absolute variance and the maximum absolute variance are presented in the Supplementary Materials (Figures S9-S12).

I chose the training set sizes to span sample sizes where the unconstrained linear strategy has difficulties in estimating the true environmental beta weights (i.e., sample sizes close to the number of cues) to sample sizes where the unconstrained linear strategy has a better chance of estimating them (i.e., sample sizes much larger than the number of cues). In the main simulation, I compared individual performance with a group size of 300.

Simulation Procedure. To answer the first research questions (i.e., 1) How is the performance of individual and aggregated linear strategies affected by different environmental and group aggregation factors?), I estimated the MSE of each strategy by generating training set items for the linear environments described by the beta weights and fitting a group of instantiations of each of the strategies to these training set items. The parameters were then fixed and used to predict 2,000 newly generated test set items. This was repeated 1000 times with newly generated cue values that with added normally distributed noise produced the criterion values for the training and test sets.

In the no overlap condition, all individual instantiations of a strategy were trained on independently generated training sets. In the overlap conditions that affect the covariance between the individual strategies, different percentages of the training set items were shared among the individual instantiations of a strategy. The individual strategy MSE was calculated as the mean MSE between the individual predictions for the individual instantiations of a strategy and the criterion values in the test set. The group MSE was calculated as the MSE between the mean prediction of the individual instantiations of a strategy and the criterion values in the test set.

To answer the second research question (i.e., 2) Can performance differences between individual and aggregated linear strategies be understood in a unified framework?, I estimated the bias-variance-covariance with the subset of the beta weights described above (i.e., those with minimum, maximum, and median variance) from the main MSE simulation, as estimating these for all the MSE simulation conditions is computationally infeasible. The bias was calculated on individual test set items as the deviation between the mean prediction across the 1,000 iterations and the criterion value for that item. The variance was calculated as the variance of the 1,000 predictions and the covariance as the average covariance across individuals for that item. The covariance is the MSE minus the bias and variance components. All components were then averaged to get one value for each component for the whole test set.

Note that the MSE for LR can be calculated directly for environments where the ordinary least squares assumptions hold and both the cue values and the criterion are jointly normal. The bias will be zero and if the sample size is large relative to the number of cues, the variance can be calculated by multiplying the known error term in the linear model equation with the number of cues divided by the sample size. However, no such easy calculations are available for constrained linear strategies. In addition, by introducing correlations between cues, the variance needs to be estimated by $tr[(X^T X)^{-1}]$ (see the Appendix for further details).

Results

1. How Is the Performance of Individual and Aggregated Linear Strategies Affected by Different Environmental and Group Aggregation Factors?

The results for the first research question are organized in three sections, each describing one the three factors that had the most influence on aggregated strategy performance: environmental predictability, covariance between individual judgments, sample size, and group size (see Table 1).

Environmental predictability. Figure 1 shows simulation results for no overlap between training sets and Figure 2 for 50% training set overlap (the results for the other overlap conditions are presented in Figure S1-S4 in the Supplemental Materials). Each of the four panels shows the results for different combination of the number of cues (3 or 9) and correlation between cues (0 or .5). The results are presented as the proportion of wins, that is, proportion of replications in which the constrained strategies (EW - circles and SC - diamonds) had a lower MSE than the unconstrained strategy (LR), for 10 different levels of R^2 . The proportion of replications in which the unconstrained and constrained strategies perform equally well is, less than 2% of replications across all simulations. Proportion of wins higher than .5 means that the constrained linear strategy outperforms LR in the majority of cases. Proportion of wins lower than .5 means that LR outperforms the constrained linear strategy. Figures 1 and 2 show the results for a training set size of 2 times the number of cues ($2p$).

Across all simulations (Figures 1,2, S1-S12), the strategy aggregation occurs when R^2 is below .55. That is, constrained linear strategies outperform LR in the individual condition, but LR outperforms them in the group condition. The strategy aggregation effect, however, is moderated by the covariance between individual judgments.

Covariance between individual judgments. When training sets do not overlap, Figure 1 shows that the strategy aggregation effect occurs for most values of R^2 (up to .55 to .75, depending on the number of cues and cue correlations): Constrained linear strategies outperform LR in the individual condition, but LR outperforms them in the group condition. When R^2 becomes sufficiently high (between .55 and .75 and up), the strategy aggregation effect does not occur: LR outperforms the constrained linear strategies in both individual and group conditions. Similar pattern occurs for different combinations of the number of cues and cue correlations.

When training sets overlap, Figure 2 shows that the strategy aggregation effect occurs only for medium-sized values of R^2 . To illustrate, compare the upper left panels in Figure 2 and Figure 1. When $R^2=.55$, the proportion of wins for SC in Figure 2 is .98 in the individual condition and .32 in the group condition. In contrast, in Figure 1 these values are .98 and 0, showing a much stronger aggregation effect. When R^2 is either low or high, the strategy aggregation effect does not occur: the constrained linear strategies win in both the individual and group conditions when R^2 is low, and LR wins when R^2 is high. Again, these patterns are approximately similar for different combinations of the number of cues and cue correlations.

Figures S1-S4 in the Supplemental Materials show that there is a gradual increase in the strength of the strategy aggregation effect with decreasing covariance between the individual judgments. When the overlap is .167 (Figure S1), the strategy aggregation effect still occurs for lower values of R^2 (up to .55 to .75, depending on the number of cues and cue correlations). When the overlap is .333 (Figure S2) the strategy aggregation effect occurs when R^2 is around .25-.55 and when the overlap is .667 (Figure S3) the pattern is similar to when the overlap is .5 and occurs when R^2 is around .45-.65. When the overlap is .833 (Figure S4), there are only two instances where the proportion of wins for the constrained linear strategies is above .5 and the proportion of wins in the group condition is below .5 ($p=3, r=0.5, R^2=0.65$ and $p=9, r=0, R^2=.55$).

Sample size. Next, I investigate the effect of sample size, implemented as training set size in the simulations, on the prevalence of the aggregation effect. Figures 1 and 2 show results of simulations in which the training set equals 2 times the number of cues ($2p$). Figures S5 to S12 show results for a number of other training set sizes. When the training set size is smaller ($1.3p$) the location of aggregation effects in both no overlap (Figure S5) and overlap (Figure S6)

conditions moves toward higher values of R^2 . Similar patterns as in Figure 1 and Figure 2 occur when the training set size is $1.7p$ (Figures S7 and S8). As the training set size increases to $5p$, the location of the strategy aggregation effect moves towards lower values of R^2 for no overlap condition (Figure S9), and almost disappears in the overlap condition (Figure S10). For the largest training set size ($10p$), in the no overlap condition the strategy aggregation effect occurs only for a R^2 of .05 (Figure S11), and in the overlap condition it does not occur at all (Figure S12).

Group size. A question of practical importance is how large the group needs to be in order to observe an aggregation effect. Figure 3 shows proportion of wins for the constrained linear strategies for different group sizes and two R^2 s (.15 and .45), without overlap between training sets. The group size needed for the strategy aggregation effect is well below 300, the group size used in Figures 1 and 2. For $R^2=.45$, the strategy aggregation effect occurs already for group sizes between 3 and 10. For $R^2=.15$, the results depend on the number of cues: when the number of cues is 3, the strategy aggregation effect occurs for group sizes between 20 and 50, and when the number of cues is 9, the strategy aggregation effect occurs for group sizes between 10 and 20.

Summary. The results show a strategy aggregation effect: a constrained linear strategy is more accurate for individual judgments, but unconstrained linear strategies are more accurate for aggregated judgments. Across all simulations, the strategy aggregation occurs when the environmental predictability is low and when the covariance between individual strategies is low. In contrast, when the covariance between individual judgments is higher the strategy aggregation effect only occurs for a limited range of environmental predictabilities. The sample size also has a large impact on the prevalence of the strategy aggregation effect, with smaller

sample sizes producing a stronger strategy aggregation effect. The group size that is needed for the strategy aggregation effect to occur is small, but is moderated by environmental predictability. When environmental predictability is high, the strategy aggregation effect is observed for group sizes between three and ten. When environmental predictability is low, the strategy aggregation effect is observed for group sizes between ten and fifty.

2. Can Performance Differences between Individual and Aggregated Linear Strategies be Understood in a Unified Framework?

Analysis of bias, variance, and covariance. To gain further insights why the strategy aggregation effect occurs in the no overlap condition and why there is much less evidence for it in the overlap conditions, the results of the simulations that estimates the bias-variance-covariance components of the MSE are presented in Figures 4 and 5. The height of the bars represents MSE, and their segments represent bias, variance, and covariance components. The standard error of the estimates for the different components are less than .02 (derived from the differences between cross-validation folds). The panels show this decomposition for different combinations of number of cues and cue correlations, for the beta vector that had the median absolute variance (see Table 1) producing an R^2 of .35.

Covariance between individual judgments. When training sets do not overlap, Figure 4 shows that EW and SC outperform LR in the individual condition because LR incurs too much variance compared to the other two strategies. The bias component of LR, however, is smaller than for the other two strategies. Therefore, in the group condition (where the variance component becomes vanishingly small; not shown as it is the same as Figure 4 but without the variance segment), bias is the only component that contributes to MSE when there is no overlap between training sets and, as a result, LR outperforms the constrained linear strategies.

When training sets overlap, Figure 5 shows the bias advantage that LR has is offset by

LR incurring a higher degree of covariance than the other two strategies. For example, the upper right panel shows that LR has a lower bias than both EW and SC, but the LR covariance is more than twice as large as the covariance for EW and SV.

Dispersion of environmental cue weights. I also investigated the bias, variance, and covariance components with different dispersions of environmental beta weights (Supplementary Materials, Figures S13 to S16). A higher variance of a set of beta weights implies that there is one cue that is better than the other cues in predicting the criterion and this should be a good fit for SV. When a set of beta weights have a higher variance (Figures S15 and S16), SC has a lower MSE than EW in five out of eight combinations of number of cues and inter-cue correlations. When the set of beta weights have a lower variance, EW has a lower MSE than SC in seven out of eight conditions when the beta weights have a lower variance (Figures S13 and S14). The results for LR and EW do not change much across different beta variances.²

Summary: The results show that a bias-variance-covariance analysis contributes to the understanding of the performance of unconstrained and constrained linear strategies. When the covariance between individual strategies is low, the strategy aggregation effect occurs because the variance component becomes vanishingly small and bias is the only component that contributes to MSE. As a result, the unconstrained strategy LR outperforms the constrained strategies EW and SC in the group condition. When the covariance between individual strategies is high, however, the strategy aggregation effect is reduced because the bias advantage that LR has is offset by LR incurring a higher degree of covariance than EW and SV.

---

² I also investigated how the performance of SC relates to the dispersion of beta weights by analyzing how many times SV picks the correct cue, that is, the cue with the highest environmental beta value. Overall, SV picks the correct cue in 75% of the cases. This value is higher for higher values of R^2 (over 98% when $R^2=.95$) and lower for lower values of R^2 (under 10% when $R^2=0.05$). The majority of individual SV strategies use the same cue when the dispersion is high (70%, $R^2=0.35$), but when the dispersion is low the majority uses different cues (65%, $R^2=0.35$).

Simulations with Real Environments

Although these artificial environment simulations were reasonably comprehensive, they are still limited to quite specific circumstances (e.g., normally and identically distributed variables and additive linear integration of those variables to produce the criterion). I next investigate whether the strategy aggregation effect can be observed in real environments and focus on the second research question on how a bias-variance-covariance analysis contributes to the understanding of performance in these environments.

Environments. To investigate if the strategy aggregation effect can occur with real data, I chose five environments with different levels of R^2 . I use R^2 derived from the predictions of a multiple regression model as an estimate of the difficulty of the environments, but this does not mean that there are no nonlinearities or factors that interact in these environments.

In contrast to many other strategy performance studies, I chose to use both subjective data (i.e., survey and psychological test data) and environments with objectively given cues and criteria. I only chose data sets that had enough objects to reliably estimate bias, variance, and covariance in cross-validation. (see Waldorp, 2008, for an assessment of the sample sizes required to estimate bias and variance from real world data). I also ran a subset of the simulations with five and ten-fold cross validation with almost identical results). In contrast to the artificial environments, cross-validation is needed as there is a limited number of objects to sample from. I estimated R^2 with multiple linear regression on each dataset. The major characteristics of the environments are summarized in Table 2 with the environments sorted from lowest R^2 to highest R^2 .

The first environment (Gcb, $R^2=.05$, http://personality-testing.info/_rawdata/) is based on data from 2,495 respondents answering 15 items from the generic conspiracy belief scale (GCB),

a 10-item version of the Big Five inventory, and self-rated word knowledge. I used the mean GCB score as the criterion variable and the 11 cues were the 10 big five items and self-rated word knowledge.

The second environment (Atus, $R^2=.12$) is based on data from 9,017 female respondents in the 2006 and 2007 American Time Use Study (ATUS) and analyzed in Zick, Stevens, and Bryant (2011). The criterion is the Body Mass Index (BMI) of survey respondents and the 18 cues are primary eating time (time spent solely eating), secondary eating time, secondary drinking time, food preparation time, physically active time, sleep time, television/video time, age, Black, Hispanic, other (i.e., not Black or Hispanic), married/cohabitating, education, employed, poor health, occupation with high physical activity, number of kids < age 6, number of kids age 6-17.

The third environment (Bf, $R^2=.44$, http://personality-testing.info/_rawdata/), is based on data from 9,881 respondents answering a 10-item version of the Big Five inventory. I used one of the two extraversion questions as criterion and the nine cues were the remaining items. This environment should be particularly well suited for SV, as the second extraversion item is highly correlated with the first, while the rest of the items have low correlations with the extraversion items.

The fourth environment (Abalone, $R^2=.54$, <https://archive.ics.uci.edu/ml/datasets/Abalone>) is a classic machine learning data set with 4,177 instances. The task is to predict the age of abalone from physical measures. I used number of shell rings (which gives the age) as the criterion and the seven cues were sex, length, diameter, height, whole weight, shucked weight, viscera weight, and shell weight.

The fifth environment (Calhousing, $R^2=.64$,

http://www.dcc.fc.up.pt/~ltorgo/Regression/cal_housing.html) is a dataset with 20,640 objects where the criterion is median house value. The eight cues are median income in city, housing median age in a city, total rooms, total bedrooms, population, households, latitude, and longitude.

Simulation Procedure. MSE, bias, variance and covariance were simultaneously estimated with three-fold cross-validation with 300 individual instantiations of each strategy fitted to 10,000 training sets. The parameters were then fixed and used to predict the criterion values in the holdout fold. There was one training set size (two times the number of cues) and two training set overlap conditions. I compared individual performance with a group size of 300. In the no overlap condition, all 300 individual strategies had independently generated training sets. In the overlap condition, 50% of the training set items were shared among the 300 individual strategies. I calculated the individual strategy MSE as the mean MSE between the individual predictions for the 300 individual strategies and the criterion values in the test set. The group MSE was calculated as the MSE between the mean prediction of the 300 individual strategies and the criterion values in the test set. All MSE components were calculated on the object level, then averaged over all objects in the test set and finally averaged over all three folds.

Results

Here I present results relevant to the question if the strategy aggregation effect can be observed in real world environments and how a bias-variance-covariance analysis contributes to the understanding of performance in these environments.

Figure 6 shows the estimates of bias, variance, and covariance for the five environments. The height of the bars represents the MSE. The environments are ordered from low R^2 (top row)

to high R^2 (bottom row) in the same way as in Table 2. Note that the bias component is the same across group sizes and the no overlap and overlap conditions, so only variance and covariance play a role in elucidating differences in performances between the strategies.

When training sets do not overlap, Figure 6 shows a strategy aggregation effect in three out of five environments. A comparison between the subplots for the individual condition (first column in Figure 6) with the group condition where there is no overlap (second column), shows a switch from LR being the worst performing strategy in three out five environments to becoming the best performing strategy in all five environments. In the individual condition in the Bf environment (third row), LR performs better than EW, but worse than SV. In the group condition, LR performs better than both EW and SV. In the calhousing environment (fifth row), which has the highest R^2 , LR outperforms EW and SC in both the individual condition and the group condition. The bias-variance-covariance components in Figure 6 illustrates why the strategy aggregation effect occurs. In the individual condition LR is plagued with high variance but has a lower bias compared to constrained strategies. The aggregation in the group condition effectively reduces the variance to zero and bias is the only remaining component that contributes to MSE.

When training sets do not overlap, Figure 6 shows a strategy aggregation effect only for one out of five environments. A comparison between the subplots for the individual condition (first column in Figure 6) with the group condition where there is a training set overlap (third column), shows that the strategy aggregation effect only occurs for the abalone environment (fourth row). For Gcb (first row) and Atus (second row), LR is now the worst performing strategy both in the individual and group conditions. The reason LR performs worse when there is overlap is that LR is more susceptible to increased covariance. LR is a more flexible strategy,

in the sense that it is more sensitive to the particulars of the training set, and if the training sets overlap the predictions of different instantiations of LR will become more similar to each other. In contrast, predictions from EW and SC are less sensitive to the particulars of the training set.

Summary. The simulations with real world environments replicate the strategy aggregation effect found in artificial environments and show that a bias-variance-covariance analysis contributes to the understanding of the performance of unconstrained and constrained linear strategies in real world environments. When the covariance between individual strategies is low, the results show the strategy aggregation effect in three out of five environments. In contrast, when the covariance between individual strategies is higher the strategy aggregation effect only occurs in one environment. The reason there is a strategy aggregation effect when there is a lower covariance between the judgments is the same as for the artificial environments: the unconstrained strategy LR is plagued with high variance but has a lower bias compared to the unconstrained strategies EW and SV. The aggregation in the group condition reduces variance to close to zero and bias is the only component left in the MSE. The reason the strategy aggregation effect does not occur when there is a higher covariance between individual strategies is the same as for the artificial environments: LR is more susceptible to increased covariance than EW and SV and the covariance overshadows the lower bias of LR.

Discussion

In the wisdom-of-crowds literature, it is well known that the diversity of predictions is important to achieve high accuracy of aggregated judgments. As is well understood in machine learning research on ensembles of algorithms for prediction, unconstrained or more flexible strategies can achieve high accuracy. In research on individual judgment strategies, however, constrained or less flexible strategies achieve high levels of accuracy when the judgment task is

difficult and the amount of information to make the judgments is limited. My results show that these two insights are in contradiction only on the surface: Both recommendations achieve the same goal of minimizing variance, but in different ways. Increasing group size minimizes the variance component of unconstrained linear strategies, while the constraints inherent in strategies such as EW, limits the space of possible predictions that the strategy can produce and thereby reduces variance.

By connecting the bias-variance profile of individual strategies with the resulting bias-variance-covariance profile of the group for unconstrained and constrained linear judgment strategies, I showed in simulations on artificial and real environments that the main reason for the strategy aggregation effect is that the unconstrained linear strategy has lower bias but higher variance than constrained linear strategies. This, in turn, leads to the unconstrained linear strategy performing worse in the individual judgment setting but better in the group setting where the aggregation of the predictions minimizes the variance component. I also documented the limiting conditions for the strategy aggregation effect to occur. The task environment must be difficult (as measured by R^2) and the training set sizes must be small relative to the number of cues. The most drastic reduction of the strategy aggregation effect occurs when there are high levels of covariance between the individual judgments. In this case, the unconstrained linear strategy incurs more covariance than the constrained linear strategies and the covariance component outweighs the lower bias of the unconstrained linear strategy relative to the unconstrained linear strategies. Additional simulations with artificial environments show that the critical group size for observing the strategy aggregation effect is approximately 15.

The results from the simulations in the artificial environments were consistent with the results for the artificial environment. The strategy aggregation effect occurs mostly for

environments with lower values of R^2 . For higher values of R^2 , the unconstrained linear strategy outperforms the constrained linear strategy in both individual and group judgments, even when the training set size is small (i.e., two times the number of cues). The simulations in the real environments showed that the strategy aggregation effect can also occur outside of artificial environments tailored to the assumptions of linear regression.

Implications for Research on Individual and Group Judgment

The results in this paper confirm the argument made by Gigerenzer and Brighton (2009) that for individual judgments people might be better off using heuristics with higher bias and lower variance than more flexible low bias and high variance strategies. The results also replicate previous research showing that equal weights and single variable strategies outperform unconstrained linear regression when the sample size is small and the predictability of the environment is low (Davis-Stober et al., 2010). The normative implication of these results is that individuals should use constrained linear strategies if they face judgment situations where information is scarce and the environmental predictability is low. If information is abundant and environmental predictability is high, they are better off using unconstrained strategies.

From an aggregation perspective, however, it is better if individuals use unconstrained linear strategies. For example, the results suggest that in an organizational context where judgments are expected to be aggregated it is better that the organization encourages individuals to use unconstrained linear strategies. Of course, what is better from an organization's perspective might not be better from an individual's perspective. For example, as long as the individual payoff is higher than the shared payoff, it is better for an individual to continue to use a constrained linear strategy.

In my simulations, I find that shared training sets, which induce correlations between

individual judgments, cancel the aggregation effect. In other words, in the presence of correlations the constrained linear strategies always outperformed the unconstrained one. In real world settings we would most certainly observe correlations between individual judgments, and one could then ask whether strategy aggregation effect matters in those settings.

Real World Implications

To understand the relevance of the strategy aggregation effect in real world settings, consider the main conclusion from research on the wisdom of crowds: diversity is necessary for a crowd to perform well (e.g., Page, 2007). Diversity in the context of the bias-variance-covariance decomposition is associated with low covariance between individuals. If an organization or any group wants to harness the wisdom of crowds they need to structure the organizational environment so that the dependencies between individuals' judgments are as low as possible. Organizations can increase diversity in a number of ways. For example, by limiting the information flow between individuals the amount of shared information can be lowered and, consequently, the correlation between individuals' judgments could be kept at a lower level. By increasing diversity, however, the organization creates the condition where the strategy aggregation effect is observed. My results suggest that these efforts for improving the wisdom of crowds by reducing correlations will also reinforce the need for using unconstrained linear strategies.

Limitations and Future Research

I will discuss three main limitations of the analyses presented in this this paper, which also present the opportunities for further research: the use of MSE as a measure of predictive accuracy, the use of linear judgment strategies in continuous tasks, and the focus on normative questions about strategy performance.

Other measures of predictive accuracy. In this paper I used MSE as a measure of

predictive accuracy. I focused on MSE for two reasons. First, I wanted to keep the correspondence between the objective function being minimized in the fitting of the parameters and the evaluation of the predictive accuracy. Second, the MAD and MAPE cannot be decomposed into bias, variance, and covariance and prevents an analysis in terms of those components.

Going beyond linear judgment strategies, linear environments, and continuous judgments. There is of course a large number of strategies that have been proposed to capture how people make judgments and there are infinite number of ways that cues can be related to the criterion in real world environments. In this paper, I chose to focus on the paradigmatic weighted additive strategy in the form of multiple linear regression and linear environments in the artificial simulations. The use of linear strategies and linear environments is a natural starting point as most of the research focused on these. There is also the question of how one should choose the functional form that relates cues to criterion, as there is an infinite amount of functions that could be used. The results from the simulations with real world environments, however, were consistent with the artificial environment simulations.

A next step is to investigate the performance of individual and group predictions from multiple linear regression and other strategies in a variety of environments. One prominent alternative to linear judgment strategies are exemplar models (Medin & Schaffer, 1978; Nosofsky, 1986)), which also have been shown to capture people's judgments, especially when one goes beyond simple linear environments. Implementing exemplar models, however, is not trivial. For example, a full implementation of the generalized context model for continuous judgments would not work in the present simulations. It would fit the training set perfectly and not be able to make predictions on the test set. For it to work, the parameters need to be

constrained and the question then becomes how to constrain the parameters. In addition, constraining the parameters restricts the space of possible predictions and therefore minimizes variance. The computational cost is also high for exemplar models with many free parameters.

Another obvious extension of the work presented here is to analyze decision and categorization tasks. By implementing same or similar models in all tasks we can investigate if the task as such would invite more or less of the aggregation effect. One challenge is how one should estimate MSE components as they are defined in different ways in different decompositions. A particular challenge is the covariance component, as this is undefined in MSE decompositions of 0-1 loss functions.

Going beyond normative questions about strategy performance. In this paper, I have used computer simulations to investigate questions about the performance of unconstrained and constrained strategies in individual and group judgment. There is, however, a large number of descriptive questions that relate to the use of different strategies in individual and group judgment. One question that has been actively investigated in the literature on adaptive cognition is the so called strategy selection problem (for example Rieskamp & Otto, 2006): How does an individual know when to switch from one strategy to another when environmental conditions change? Many approaches to this problem assume that people somehow calculate the costs and benefits based on experiences with one or several decision environments (Beach & Mitchell, 1978; Payne et al., 1993), possibly through reinforcement learning mechanisms (Rieskamp & Otto, 2006). In order for these approaches to work in the group setting, other mechanisms are probably needed. For example, pure reinforcement learning mechanisms that operates on the feedback of individual judgments might need to be augmented with some sort of meta-cognitive insight: “I know that my judgments are being aggregated, therefore I should abandon the strategy

that worked well for me individually and switch to another”. It would be worthwhile to explore whether people are actually adapting their strategy in response to the social context and if so, how.

References

- Ashton, A. H. (1982). An empirical study of budget-related predictions of corporate executives. *Journal of Accounting Research*, 20, 440–449.
- Beach, L. R., & Mitchell, T. R. (1978). A contingency model for the selection of decision strategies. *Academy of Management Review*, 3, 439–449.
- Brighton, H. (2006). Robust inference with simple cognitive models. In C. Lebiere & R. Wray (Eds.), *AAAI spring symposium: Cognitive science principles meet AI-hard problems* (pp. 17–22). Menlo Park, CA: American Association for Artificial Intelligence.
- Dana, J., & Dawes, R. M. (2004). The superiority of simple alternatives to regression for social science predictions. *Journal of Educational and Behavioral Statistics*, 29, 317–331.
- Davis-Stober, C. P., Budescu, D. V., Dana, J., & Broomell, S. B. (2014). When is a crowd wise? *Decision*, 1, 79–101.
- Davis-Stober, C. P., Budescu, D. V., Broomell, S. B., & Dana, J. (2015). The composition of optimally wise crowds. *Decision Analysis*, 12, 130–143.
- Davis-Stober, C. P., Dana, J., & Budescu, D. V. (2010). A constrained linear estimator for multiple regression. *Psychometrika*, 75, 521–541.
- Dawes, R. M. (1979). The robust beauty of improper linear models in decision making. *American Psychologist*, 34, 571–582.
- Dawes, R. M., & Corrigan, B. (1974). Linear models in decision making. *Psychological Bulletin*, 81, 95–106.
- Einhorn, H. J., & Hogarth, R. M. (1975). Unit weighting schemes for decision making. *Organizational Behavior and Human Performance*, 13, 171–192.
- Geman, S., Bienenstock, E., & Doursat, R. (1992). Neural networks and the bias/variance

- dilemma. *Neural Computation*, 4, 1–58.
- Gigerenzer, G., & Brighton, H. (2009). Homo heuristics: Why biased minds make better inferences. *Topics in Cognitive Science*, 1, 107–143.
- Gigerenzer, G., Hertwig, R., & Pachur, T. (Eds.). (2011). *Heuristics: The foundations of adaptive behavior*. New York: Oxford University Press.
- Gigerenzer, G., Todd, P. M., & the ABC Research Group. (1999). *Simple heuristics that make us smart*. New York: Oxford University Press.
- Hogarth, R. M., & Karelaia, N. (2007). Heuristic and linear models of judgment: Matching rules and environments. *Psychological Review*, 114, 733–758.
- Katsikopoulos, K. V., Schooler, L. J., & Hertwig, R. (2010). The robust beauty of ordinary information. *Psychological Review*, 117, 1259–1266.
- Lamberson, P. J., & Page, S. E. (2012). Optimal forecasting groups. *Management Science*, 58, 805–810.
- Martignon, L., & Hoffrage, U. (2002). Fast, frugal, and fit: simple heuristics for paired comparison. *Theory and Decision*, 29–71.
- Page, S. E. (2007). *The difference: How the power of diversity creates better groups, firms, schools, and societies*. Princeton, NJ: Princeton University Press.
- Payne, J. W., Bettman, J. R., & Johnson, E. J. (1993). *The adaptive decision maker*. New York: LRbridge University Press.
- Rieskamp, J., & Otto, P. E. (2006). SSL: A theory of how people learn to select strategies. *Journal of Experimental Psychology: General*, 135, 207–236.
- Sunstein, C. R. (2006). *Infotopia: How many minds produce knowledge*. New York: Oxford University Press.

Surowiecki, J. (2004). *The wisdom of crowds*. New York: Doubleday.

Ueda, N., & Nakano, R. (1996). Generalization error of ensemble estimators. *IEEE International Conference on Neural Networks*, 90–95.

Waller, N., & Jones, J. (2011). Investigating the performance of alternate regression weights by studying all possible criteria in regression models with a fixed set of predictors. *Psychometrika*, 76, 410–439.

Wolpert, D. H. (1997). On bias plus variance. *Neural Computation*, 9, 1211–1243.

Zick, C. D., Stevens, R. B., & Bryant, W. K. (2011). Time use choices and healthy body weight: A multivariate analysis of data from the American Time use Survey. *International Journal of Behavioral Nutrition and Physical Activity*, 8, 84.

Table 1

Factors That Affect Individual and Aggregated Performance of Strategies and How They Were Implemented in the Simulations

Factors affecting strategy performance	Simulation implementation
Individual performance	
Number of cues	$p=[3, 9]$
Inter-cue correlations	$r=[0, .5]$
Environmental predictability	$R^2=[.05, .15, .25., .35, .45, .55, .65, .75, .85, .95]$
Sample size	Training set size = $[1.3, 1.7, 2, 5, 10] \times p$
Dispersion of environmental cue weights	Different variances of beta weights (Table S1)
Aggregated performance	
Group size	$[3, 10, 20, 50, 100, 300]$
Covariance between individual judgments	Training set overlap $[0\%, 16.7\%, 33.3\%, 50\%, 66.7\%, 83.3\%]$

Table 2

Environment Characteristics

Environment	R^2	Number of objects	Number of cues	Criterion
Gcb	.05	2495	11	Mean generic conspiracy belief scale score
Atus	.12	9017	18	Body Mass Index
Bf	.44	9881	9	Extraversion score
Abalone	.54	4177	7	Number of rings
Calhousing	.64	20640	8	Median house value

Figure 1. Proportion of wins in the no overlap condition for EW (circles) and SV (diamonds) for 10 different levels of R^2 in artificial environments. Unfilled circles and diamonds represent the individual model condition (Group size=1) and filled circles and diamonds represent the group condition (Group size=300). The four panels show the results for different number of cues ($p=4$ and $p=9$) and cue correlations ($r=0$ and $r=.5$). Proportion of wins higher than .5 means that the constrained strategy outperforms LR in the majority of replications. Proportion of wins lower than .5 means that LR outperforms the constrained strategy in the majority of replications.

Figure 2. Proportion of wins in the overlap condition for EW (circles) and SV (diamonds) for 10 different levels of R^2 in artificial environments. Unfilled circles and diamonds represent the individual condition (Group size=1) and filled circles and diamonds represent the group condition (Group size=300). The four panels show the results for different number of cues ($p=4$ and $p=9$) and cue correlations ($r=0$ and $r=.5$). Proportion of wins higher than .5 means that the constrained strategy outperforms LR in the majority of replications. Proportion of wins lower than .5 means that LR outperforms the constrained strategy in the majority of replications.

Figure 3. Proportion of wins in the no overlap condition for EW (circles) and SV (diamonds) for six different group sizes and two levels of R^2 in artificial environments. Unfilled circles and diamonds represent the $R^2=.15$ condition and filled circles and diamonds represent the $R^2=.45$ condition. The four panels show the results for different number of cues ($p=4$ and $p=9$) and cue correlations ($r=0$ and $r=.5$).

Figure 4. Bias (light gray) and variance (dark gray) for the individual condition when there is no training set overlap. The height of the bars is equal to MSE.

Figure 5. Bias (light gray), variance (dark gray), and covariance (white) for the group condition when there is a 50% training set overlap. The height of the bars is equal to MSE.

Figure 6. Bias (light gray), variance (dark gray), and covariance (white) for the 5 real-world environments. The rows are environments in the same order as in Table 2. The columns (left to right) present results for the individual condition, the group condition when there is no training set overlap, and the condition when there is a 50% training set overlap. The height of the bars represents MSE. The environments are ordered from low R^2 (top row) to high R^2 (bottom row).

Appendix

In this Appendix, I give the equations for the decomposition of the mean squared error (MSE) that I use in this paper, describe alternative MSE decompositions, and describe the specific bias variance decomposition for unconstrained linear models.

The Bias Variance Decomposition

I use the most widely applied MSE decomposition for comparing the predictive performance of different models. The mathematical derivation of the bias-variance decomposition can be found in machine-learning and statistical textbooks, but it is perhaps most thoroughly set out and discussed in the seminal article by Geman, Bienenstock, and Doursat (1992) and later extended to aggregated predictions by Ueda and Nakano (1996).³ This decomposition assumes that the test set is fixed, and the training set is random. In practice, this means that the different components of the decomposition are estimated by repeated sampling from a large training set. The sample of instances is used to estimate the parameters of the models, which are then fixed and used to predict items in the test set. The estimates of the different components are at the item level and are averaged to produce overall estimates for a specific test set. The decomposition can be expressed as:

$$MSE = \overline{bias}^2 + \frac{1}{M} \overline{var} + \left(1 - \frac{1}{M}\right) \overline{cov} + \sigma_\epsilon^2. \quad (1)$$

The bias and variance components are averages over the M group members, indicated by bars above these terms. The decomposition also includes a covariance term, capturing the dependency between group members expressed as their average pair-wise covariance. Finally, the error term σ_ϵ^2 is usually referred to as “irreducible” error and cannot be estimated separately

---

³ There are also decompositions for zero-one loss functions in classification tasks (Domingos, 2000; James, 2003; Kohavi & Wolpert, 1996; Kong & Dietterich, 1995).

on real world data. Equation 1 also makes it clear that when the size of the group increases the impact of the variance component decreases (the $1/M$ term) and the impact of the covariance component increases (the $1-1/M$ term).

Alternative MSE Decompositions

In general, for a single estimator MSE can be decomposed in the following way (Wolpert, 1997):⁴

$$MSE = bias^2 + var - 2cov_{XY} + var(Y). \quad (2)$$

The covariance term here is the covariance between judgments and the criterion, when the covariance increases the total error becomes lower. Specified in this way, we are faced with a bias-variance-covariance tradeoff.

This decomposition is also directly applicable to any set of data, assuming that the criterion is a random variable. Bias is estimated as the difference between the mean prediction across items and the mean of the criterion variable. Variance is estimated as the variance of the predictions across items and covariance is estimated as the covariance between the predictions and the criterion. The last term does not refer to irreducible error as in Equation 1, it simply refers to the variance of the criterion variable Y . This decomposition contrasts with the decomposition in Equation 1 where the components are estimated within items. Also note that directly applying Equation 2 to a fixed data set does not give an estimate of the prediction error. For this, additional procedures are needed, such as cross-validation.

Equation 2 can be extended to apply to group judgments by adding a covariance term that captures the pair-wise covariance among the group members,

---

⁴ This is a simplified presentation of this decomposition. For a full discussion of how different conditioning events changes the decomposition, see Wolpert (1997).

$$MSE = \overline{bias}^2 + \frac{1}{M} \overline{var} + \left(1 - \frac{1}{M}\right) \overline{cov} - 2\overline{cov}_{XY} + var(Y). \quad (3)$$

This decomposition was recently used by Davis-Stober et al. (2015), generalizing the results in Lamberson and Page (2012) regarding the composition of optimally wise crowds.

The main reason I did not use the decomposition in Equation 3 is that it obscures the role of bias at the item level, at least for the artificial simulations and the strategies investigated in this paper. If one would apply Equation 3 to predictions from linear regression in the artificial environments I use in my simulations, the bias component would become vanishingly small as training set size or group size becomes sufficiently large.

The Bias Variance Decomposition for Unconstrained Linear Models

In my simulations I use environments in which the relations between cues and criterion are linear and all cues, criterion, and the error term is normally distributed. For this setup there exists analytical solutions to the bias and variance components. In this paper I have relied on simulations because deriving solutions to the bias, variance, and covariance terms in Eq. 1 for a constrained linear estimator that does not rely on plugging in empirical estimates is a difficult task (e.g., small sample estimates of the covariance matrix of the cues are needed, see Roberts, 1999). There are solutions for the components of the constrained linear estimator that I use, but only for the bias and variance of the beta weights (Davis-Stober et al., 2010). These results do not carry over to the MSE of predictions.

For the constrained estimator case, even if analytic solutions can be derived, simulations are still needed to estimate the MSE components on real world data. For the unconstrained linear model, however, the solutions are quite simple and can be found in advanced textbooks on multiple linear regression (see also Rosset & Tibshirani, 2019 for derivations under different fixed and random assumptions). When the true function is a linear function of x , $f(x) = x^T \beta$,

$\epsilon \sim N(0, \sigma^2)$, and ϵ is independent of x , the ordinary least squares (OLS) estimator, $\hat{\beta}$, is $(X^T X)^{-1} X^T Y$. The covariance matrix of the parameter vector $\hat{\beta}$ is $(X^T X)^{-1}$. Further, the R^2 of the true function is $1 - \sigma_\epsilon^2$. A linear model estimated with OLS when (x, y) is jointly normal, has zero bias and the MSE of the predictions can be expressed in terms of the covariance matrix of $\hat{\beta}$,

$$MSE = \text{tr}[(X^T X)^{-1}] + \sigma_\epsilon^2.$$

This simply means that the expected MSE of the predictions is the sum of the variances in the covariance matrix of the estimated parameters $\hat{\beta}$. If the number of data points in the training set, n , is much larger than the number of variables, p , (i.e., when $n - p \rightarrow \infty$), the variance can be expressed as $\frac{p}{n} \sigma_\epsilon^2$. If $n - p$ is small, a small sample correction is needed, resulting in $\frac{p}{n - p - 1} \sigma_\epsilon^2$.

Supplemental Materials

Table S1
Means and Variances For All Environmental Beta Weights and Examples of Beta Weight Vectors with Different Dispersions

<i>Means and variances of beta vectors</i>										
	R^2									
	0.05	0.15	0.25	0.35	0.45	0.55	0.65	0.75	0.85	0.95
Mean beta for p=3	-	-	-	-	-	-	-	-	-	-
	0.017	0.012	0.010	0.023	0.011	0.022	0.012	0.056	0.051	-0.019
Mean beta for p=9	-	-	-	-	-	-	-	-	-	-
	0.002	0.003	0.000	0.004	0.006	0.005	0.005	0.001	0.003	-0.002
Mean variance of beta for p=3	-	-	-	-	-	-	-	-	-	-
	0.020	0.060	0.097	0.143	0.178	0.228	0.263	0.285	0.335	0.371
Mean variance beta of for p=9	-	-	-	-	-	-	-	-	-	-
	0.007	0.023	0.036	0.052	0.068	0.082	0.099	0.108	0.130	0.144
<i>Examples of beta vectors with different dispersions</i>										
	Cues									
	C1	C2	C3	C4	C5	C6	C7	C8	C9	
Min variance beta vector for p=3, $R^2=.35$	-	-	-	-	-	-	-	-	-	-
	0.367	0.287	0.364							
Median variance beta vector for p=3, $R^2=.35$	-	-	-	-	-	-	-	-	-	-
	0.155	0.201	0.394							
Max variance beta vector for p=3, $R^2=.35$	-	-	-	-	-	-	-	-	-	-
	0.621	0.023	0.038							
Min variance beta vector for p=9, $R^2=.35$	-	-	-	-	-	-	-	-	-	-
	0.363	0.217	0.205	0.197	0.160	0.158	0.144	0.130	0.044	
Median variance beta vector for p=9, $R^2=.35$	-	-	-	-	-	-	-	-	-	-
	0.080	0.081	0.082	0.140	0.155	0.155	0.223	0.324	0.368	
Max variance beta vector for p=9, $R^2=.35$	-	-	-	-	-	-	-	-	-	-
	0.285	0.230	0.106	0.083	0.000	0.005	0.019	0.179	0.692	

Note: The top panel shows the overall means and variances of the 50 beta weight vectors used in the simulations for each R^2 . The means and variances are first calculated for each cue across beta vectors and then averaged. The bottom panel shows examples of beta vectors with different dispersions (variances). The beta vectors with the minimum, median, and maximum variances are determined by first calculating the absolute variance across each of the 50 beta weight vectors for an R^2 of .35 and then selecting the vectors with the minimum, median, and maximum values, respectively.

Figure S1. Proportion of wins when there is a .167 training set overlap for EW (circles) and SV (diamonds) for 10 different levels of R^2 and a training set size of 1.3 times the number of cues. Unfilled circles and diamonds represent the individual model condition (Group size=1) and filled circles and diamonds represent the group condition (Group size=300). The four panels show the results for different number of cues ($p=4$ and $p=9$) and cue correlations ($r=0$ and $r=.5$). Proportion of wins higher than .5 means that the constrained strategy outperforms CAM in the majority of replications. Proportion of wins lower than .5 means that CAM outperforms the constrained strategy in the majority of replications.

Figure S2. Proportion of wins when there is a .333 training set overlap for EW (circles) and SV (diamonds) for 10 different levels of R^2 and a training set size of 1.3 times the number of cues. Unfilled circles and diamonds represent the individual model condition (Group size=1) and filled circles and diamonds represent the group condition (Group size=300). The four panels show the results for different number of cues ($p=4$ and $p=9$) and cue correlations ($r=0$ and $r=.5$). Proportion of wins higher than .5 means that the constrained strategy outperforms CAM in the majority of replications. Proportion of wins lower than .5 means that CAM outperforms the constrained strategy in the majority of replications.

Figure S3. Proportion of wins when there is a .667 training set overlap for EW (circles) and SV (diamonds) for 10 different levels of R^2 and a training set size of 1.3 times the number of cues. Unfilled circles and diamonds represent the individual model condition (Group size=1) and filled circles and diamonds represent the group condition (Group size=300). The four panels show the results for different number of cues ($p=4$ and $p=9$) and cue correlations ($r=0$ and $r=.5$). Proportion of wins higher than .5 means that the constrained strategy outperforms CAM in the majority of replications. Proportion of wins lower than .5 means that CAM outperforms the constrained strategy in the majority of replications

Figure S4. Proportion of wins when there is a .833 training set overlap for EW (circles) and SV (diamonds) for 10 different levels of R^2 and a training set size of 1.3 times the number of cues. Unfilled circles and diamonds represent the individual model condition (Group size=1) and filled circles and diamonds represent the group condition (Group size=300). The four panels show the results for different number of cues ($p=4$ and $p=9$) and cue correlations ($r=0$ and $r=.5$). Proportion of wins higher than .5 means that the constrained strategy outperforms CAM in the majority of replications. Proportion of wins lower than .5 means that CAM outperforms the constrained strategy in the majority of replications

Figure S5. Proportion of wins in the no overlap condition for EW (circles) and SV (diamonds) for 10 different levels of R^2 and a training set size of 1.3 times the number of cues. Unfilled circles and diamonds represent the individual model condition (Group size=1) and filled circles and diamonds represent the group condition (Group size=300). The four panels show the results for different number of cues ($p=4$ and $p=9$) and cue correlations ($r=0$ and $r=.5$). Proportion of wins higher than .5 means that the constrained strategy outperforms CAM in the majority of replications. Proportion of wins lower than .5 means that CAM outperforms the constrained strategy in the majority of replications.

Figure S6. Proportion of wins in the overlap condition for EW (circles) and SV (diamonds) for 10 different levels of R^2 and a training set size of 1.3 times the number of cues. Unfilled circles and diamonds represent the individual model condition (Group size=1) and filled circles and diamonds represent the group condition (Group size=300). The four panels show the results for different number of cues ($p=4$ and $p=9$) and cue correlations ($r=0$ and $r=.5$). Proportion of wins higher than .5 means that the constrained strategy outperforms CAM in the majority of replications. Proportion of wins lower than .5 means that CAM outperforms the constrained strategy in the majority of replications.

Figure S7. Proportion of wins in the no overlap condition for EW (circles) and SV (diamonds) for 10 different levels of R^2 and a training set size of 1.7 times the number of cues. Unfilled circles and diamonds represent the individual model condition (Group size=1) and filled circles and diamonds represent the group condition (Group size=300). The four panels show the results for different number of cues ($p=4$ and $p=9$) and cue correlations ($r=0$ and $r=.5$). Proportion of wins higher than .5 means that the constrained strategy outperforms CAM in the majority of replications. Proportion of wins lower than .5 means that CAM outperforms the constrained strategy in the majority of replications.

Figure S8. Proportion of wins in the overlap condition for EW (circles) and SV (diamonds) for 10 different levels of R^2 and a training set size of 1.7 times the number of cues. Unfilled circles and diamonds represent the individual model condition (Group size=1) and filled circles and diamonds represent the group condition (Group size=300). The four panels show the results for different number of cues ($p=4$ and $p=9$) and cue correlations ($r=0$ and $r=.5$). Proportion of wins higher than .5 means that the constrained strategy outperforms CAM in the majority of replications. Proportion of wins lower than .5 means that CAM outperforms the constrained strategy in the majority of replications.

Figure S9. Proportion of wins in the no overlap condition for EW (circles) and SV (diamonds) for 10 different levels of R^2 and a training set size of 5 times the number of cues. Unfilled circles and diamonds represent the individual model condition (Group size=1) and filled circles and diamonds represent the group condition (Group size=300). The four panels show the results for different number of cues ($p=4$ and $p=9$) and cue correlations ($r=0$ and $r=.5$). Proportion of wins higher than .5 means that the constrained strategy outperforms CAM in the majority of replications. Proportion of wins lower than .5 means that CAM outperforms the constrained strategy in the majority of replications.

Figure S10. Proportion of wins in the overlap condition for EW (circles) and SV (diamonds) for 10 different levels of R^2 and a training set size of 5 times the number of cues. Unfilled circles and diamonds represent the individual model condition (Group size=1) and filled circles and diamonds represent the group condition (Group size=300). The four panels show the results for different number of cues ($p=4$ and $p=9$) and cue correlations ($r=0$ and $r=.5$). Proportion of wins higher than .5 means that the constrained strategy outperforms CAM in the majority of replications. Proportion of wins lower than .5 means that CAM outperforms the constrained strategy in the majority of replications.

Figure S11. Proportion of wins in the no overlap condition for EW (circles) and SV (diamonds) for 10 different levels of R^2 and a training set size of 10 times the number of cues. Unfilled circles and diamonds represent the individual model condition (Group size=1) and filled circles and diamonds represent the group condition (Group size=300). The four panels show the results for different number of cues ($p=4$ and $p=9$) and cue correlations ($r=0$ and $r=.5$). Proportion of wins higher than .5 means that the constrained strategy outperforms CAM in the majority of replications. Proportion of wins lower than .5 means that CAM outperforms the constrained strategy in the majority of replications.

Figure S12. Proportion of wins in the overlap condition for EW (circles) and SV (diamonds) for 10 different levels of R^2 and a training set size of 10 times the number of cues. Unfilled circles and diamonds represent the individual model condition (Group size=1) and filled circles and diamonds represent the group condition (Group size=300). The four panels show the results for different number of cues ($p=4$ and $p=9$) and cue correlations ($r=0$ and $r=.5$). Proportion of wins higher than .5 means that the constrained strategy outperforms CAM in the majority of replications. Proportion of wins lower than .5 means that CAM outperforms the constrained strategy in the majority of replications.

Figure S13. Bias (light gray) and variance (dark gray) for the individual model condition, no training set overlap, and Beta weights with minimum absolute variance in each condition. The height of the bars is equal to MSE.

Figure S14. Bias (light gray), variance (dark gray), and covariance (white) for the group condition, 50% training set overlap, and Beta weights with minimum absolute variance in each condition. The height of the bars is equal to MSE.

Figure S15. Bias (light gray) and variance (dark gray) for the individual model condition, no training set overlap, and Beta weights with maximum absolute variance in each condition. The height of the bars is equal to MSE.

Figure S16. Bias (light gray), variance (dark gray), and covariance (white) for the group condition, 50% training set overlap, and Beta weights with maximum absolute variance in each condition. The height of the bars is equal to MSE.