To cite this chapter:

Rutten, I., Voorspoels, W., Koster, E. H. W., & Vanpaemel, W. (2017). A model-based approach for assessing attentional biases in people with depressive symptoms. In G. Gunzelmann, A. Howes, T. Tenbrink, & E. J. Davelaar (Eds.), *Proceedings of the 39th annual conference of the cognitive science society* (pp. 3022-3027). Cognitive Science Society.

A Model-based Approach for Assessing Attentional Biases in People with Depressive Symptoms

Isa Rutten (isa.rutten@kuleuven.be)
Quantitative Psychology and Individual Differences Research Group, Tiensestraat 102,
University of Leuven, 3000 Leuven, Belgium (“ same affiliation for 2 other authors)

Wouter Voorspoels (wouter.voorspoels@kuleuven.be)
Laboratory of Experimental Psychology, (“)

Ernst H.W. Koster (ernst.koster@ugent.be)
Department of Experimental Clinical and Health Psychology, Henri Dunantlaan 2,
Ghent University, B-9000 Ghent, Belgium

Wolf Vanpaemel (wolf.vanpaemel@kuleuven.be)
Quantitative Psychology and Individual Differences Research Group, (“)

Abstract

Biased attention is assumed to play an important role in the etiology and maintenance of depression and depressive symptoms. In this paper, we used data from a categorization task and an associated model to assess the attentional bias of people with varying levels of depressive symptoms. Attentional bias was operationalized as the parameter estimate in a prototype model of categorization. For estimation, we used a Bayesian hierarchical mixture approach. We expected to find a positive correlation between depressive symptoms and an AB for negative material and a negative correlation between depressive symptoms and a bias toward positive material. Despite good model fit, Bayesian regression analyses revealed weak or moderate evidence in favor of the null model assuming no association between attentional preferences and depressive symptoms, both for negative and positive material.

Keywords: psychology; cognitive science ; attention; concepts and categories; Bayesian modeling; mood disorder

Introduction

Biased attention takes an important position in cognitive theories explaining the etiology and maintenance of psychological disorders. In particular, depression has been theorized to be linked to biased attention for negative information (Beck, 1976). Multiple methods have been developed to assess attentional biases (AB). The most popular approaches include reaction time assessments, in which response latencies for negative, positive, and neutral material are compared, and translated into AB indices. Eye tracking techniques, comparing fixation durations on negative, positive, and neutral stimuli, are also a popular approach. Caution is recommended however, in both approaches. When relying on reaction time assessments, one has to consider the ambiguity and lack of reliability often associated with response latencies (Rodebaugh et al., 2016). Eye tracking techniques appear to yield good reliability estimates, but only if certain conditions are met (Rodebaugh et al., 2016). Moreover, as these techniques focus exclusively on overt attentional processes, they are less informative with regard to attentional resources that are allocated covertly, without saccadic eye movements. Perhaps not surprisingly, results obtained with these existing approaches are not consistent (Peckham, McHugh, & Otto, 2010).

The goal of this paper is to explore novel methodologies for assessing attentional biases associated with depressive symptoms. In particular, inspired by Viken, Treat, Nosofsky, McFall, and Palmeri (2002), we tested the applicability of a categorization approach to assess AB’s in the context of depression (see also Kruschke & Vanpaemel, 2015). Participants were presented pictures of human faces varying in emotional expression (emotional stimulus dimension) and hair color (neutral stimulus dimension). They were asked to classify these stimuli in two different categories, each represented by a prototype stimulus of that category. The two prototypes reflected extreme levels of the stimulus dimensions and were each other’s opposites. For example, prototype A had a light hair color, and a very sad facial expression, versus prototype B with a dark hair color and a slightly sad facial expression. In this way, participants could either choose to focus on the hair color (or neutral) dimension or on the facial expression (or emotional) dimension to classify the pictures. These data, taking the form of classification counts of each category, per stimulus, were then used to estimate participant-specific parameters in a prototype model (e.g., Nosofsky, 1987). One of these parameters corresponds to the attentional weight (AW) for one stimulus dimension, reflecting the relative attentional preference for that stimulus dimensions compared to the other stimulus dimension.

We expected to find a positive correlation between depressive symptoms and an AB for negative material and a negative correlation between depressive symptoms and a bias toward positive material (Peckham et al., 2010). In addition

to depressive symptoms, we assessed anxiety symptoms, as depression and anxiety are known for their comorbidity, and our aim was to isolate the relation of AB to depressive symptoms. Brooding or depressive rumination, and a negative mood, were also investigated in this study, since both variables have already been related to an attentional bias for negative material (Bradley, Mogg, & Lee, 1997; Koster, De Lissnyder, Derakshan, & De Raedt, 2011). Finally, the personality trait, neuroticism was assessed, in order to explore whether this important risk factor for depression, could be related to an AB for negative information.

Method

Every participant received two versions of a classification task, each version was administered in two (within-participant) conditions. In all versions and conditions, participants were asked to classify facial stimuli according to two prototype stimuli. The stimuli were made up of only two different stimulus dimensions: brightness of hair color and intensity of emotional expression. In the happy condition**,** the emotional expression ranged from a slightly to a very happy facial expression, whereas in the sad condition, it ranged from slightly to very sad.

We report all data exclusions, all included questionnaires or measures, and all study conditions.

Participants

A total of 309 first-year psychology students participated in this study in exchange for course credits (262 women, mean age = 18.53, SD = 1.90, with a range from 17 to 39). The sample size was determined by the number of participants showing up during the two weeks of data collection, available for all first-year students of the Psychology department of the University of Leuven (Belgium).

Materials

Self-report Measures The Center for Epidemiologic Studies - Depression Scale (CES-D; Radloff, 1977) was used to assess depressive symptoms (score: 0 - 60). The Hospital Anxiety and Depression Scale (HADS; Spinhoven et al., 1997), was included to assess comorbid depression (score: 0 - 21) and anxiety symptoms (score: 0 - 21). We also included the Rumination Response Scale (RRS; Nolen-Hoeksema, & Morrow, 1991) to assess brooding (score: 5 - 20), and the Ten Item Personality Inventory (TIPI; Gosling, Rentfrow, Swann, 2003) to assess neuroticism (score: 1 - 7). Finally, the current mood of participants was measured with a 5-point Likert scale ("How do you feel at this moment? 1 (very unhappy) - 2 (slightly unhappy) - 3 (neutral) - 4 (slightly happy) - 5 (very happy"). Dutch versions of these questionnaires were administered.

Stimuli The stimuli were adopted from the Karolinska Directed Emotional Faces (KDEF; Lundqvist, Flykt, & Öhman, 1998). Applying Fotomorph 13.9 (Softland SRL) and GIMP (2008), the pictures were adjusted in such way they only systematically differed from each other on two dimensions: The intensity of facial expression and the brightness of the hair color. On the basis of the modified stimuli, a negative and positive stimulus set were created. In the sad condition, the intensity of the facial expression ranged between very sad and slightly sad, whereas in the happy condition, the intensity ranged from very happy to slightly happy. By creating five different levels of emotional expression (mild intensity – strong intensity) and five different levels of brightness of hair color (light – dark) and combining all possible levels, we obtained a stimulus set consisting of 25 different pictures for each condition. Prototypes were extreme on both dimensions (for example, a stimulus having a very sad emotional expression of level 5, and very dark hair color of level 5). Figure 1 presents example stimuli.

Task In both conditions, participants were asked to classify all pictures of the condition-specific stimulus sets into category A or B. Category A was represented by one of the four prototypes, and category B by the inverse prototype, within the same emotional condition. For example, in the sad condition, if prototype A had level 5 of sadness (very sad) and level 5 of hair color (very dark), then prototype B had level 1 of sadness (slightly sad), and level 1 of hair color (very light). The prototypes stayed on participants' computer screen during the entire task. In each trial, participants had to classify one picture into one of the two categories. No feedback was provided, so participants could freely choose how to classify the stimuli. In making the classification decisions, participants thus could either base their classifications on the pictures' hair color or facial expression, reflecting their attentional focus on these dimensions.



Figure 1. The two upper pictures represent two possible prototypes. The stimuli at the bottom are random examples of the negative stimulus set, to be classified in category A or B. Adapted from "The Karolinska Directed Emotional Faces" – KDEF", by D. Lundqvist et al., 1998, CD ROM from Department of Clinical Neuroscience, Psychology section, Karolinska Institutet, ISBN 91-630-7164-9. Copyright 2015 by D. Lundqvist.

For each condition, there were two versions of the task, differing in the prototype pairs used. Version 1 consisted of prototype A1, with level 5 of sadness, and level 5 of hair color, and prototype B1, with level 1 of sadness, and level 1 of hair color. Version 2 consisted of prototype A2 with level 1 of sadness, and level 5 of hair color, and prototype B2 with level 5 of sadness and level 1 of hair color.

Procedure

Each participant ran through both conditions, and in each condition, they performed two versions of the same task. The order of the conditions, and task versions within the conditions, was counterbalanced between participants.

In each task, participants categorized two blocks in which all 23 non-prototype stimuli were presented in a random order. Thus, within each task version, each stimulus was classified twice. After completing the categorization tasks, participants were asked to indicate their current mood state, and to fill out the CES-D, HADS, RRS, and TIPI.

Model

The categorization data were analyzed using Nosofsky's (1987) weighted prototype classification model:

$$P(A \mid i) = \frac{\eta_{iA}}{\eta_{iA}+\eta_{iB}}$$

The model assumes that the probability of classifying stimulus *i* in category *A*, $P(A \mid i)$, is driven by the perceived similarity between stimulus *i* and prototype *A*, η_{iA}, divided by the overall perceived similarity between stimulus i and prototype A and prototype B. The perceived similarity between stimulus *i* and prototype *A* is assumed to be an exponential decay function determined by a sensitivity parameter *c*, and the weighted distance between the stimulus and the prototype, d_{iA}:

$$\eta_{iA} = e^{-cd_{iA}}$$

The sensitivity parameter *c* reflects how clearly the stimuli could be discriminated from each other. The weighted distance between the stimuli and prototype *A*, d_{iA}, was computed as follows:

$$d_{iA} = w_a \left| x_{ia} - x_{Aa} \right| + (1 - w_a) \left| x_{ih} - x_{Ah} \right|$$

where x_{ia} represents the coordinate of stimulus *i* on dimension *affect*, and x_{ih} the coordinate of stimulus *i* on dimension *hair*, in psychological space. x_{Aa} is the coordinate of prototype *A* on dimension *affect*, and x_{Ah} the coordinate of prototype *A* on dimension *hair*. w_a is the AW for the affect dimension, and $1 - w_a$ is the AW for the hair dimension.

The coordinates of the stimuli and prototypes were obtained in a separate study in which a different group of 32 participants rated the two dimensions of all stimuli on a 10-point Likert scale, ranging from 0 (light hair color/mild facial expression) to 9 (dark hair color/ intense facial expression). For both the sad and happy conditions, the split-half reliabilities of respectively the affect and hair color dimensions were .96 and .99. The coordinates of the stimuli were calculated by taking the mean rating on each dimension for every stimulus.

The prototype model was extended hierarchically and with a mixture component in a Bayesian framework (see e.g., Bartlema, Lee, Wetzels, & Vanpaemel, 2014). The hierarchical extension was used to accommodate continuous, qualitative individual differences, and to shrink extreme values, whereas the mixture component was included to accommodate discrete, quantitative individual differences. In particular, it allowed differentiating between three groups of participants: The first group consists of people whose behavior was captured better by a guessing model, which assumed that the response probability for each stimulus was .5. Identifying these participants avoids contamination of our parameter estimates by participants for whom the prototype model was not sufficiently appropriate, that is, participants who appeared to be guessing (see, e.g., Voorspoels, Rutten, Bartlema, Tuerlinckx & Vanpaemel, in press, and Zeigenfuse & Lee, 2010 for a similar approach).

Among the participants assigned to the prototype group, we allowed two subgroups: the 'affect group' and the 'neutral/hair group' with the 'affect group' having a higher group-level AW for the affect dimension, as compared to the 'hair group'. In particular, we restricted the mean AW for the affect dimension in the 'affect group' to be higher than the mean AW for affect in the 'hair group'.

Results

Model Analyses

The model was implemented in JAGS (Plummer, 2011). We ran 3 chains with 36 000 iterations each, after discarding 4000 iterations for burn in. We performed separate analyses for the sad and happy conditions**.** The data from the two task versions (different prototypes) within each condition were jointly modelled in order to obtain a single AW in each condition. The model analyses identified two clearly distinguishable groups (a group focusing on affect and a group focusing on hair color) in both conditions.

In the sad condition, the 'affect group' contained 141 participants, whose AW was larger for the affect dimension (the group-level posterior had a mean *w_sadness* = .88), compared to the 'neutral/hair group', containing 163 participants (mean *w_sadness* = .21).

In the happy condition, the 'affect group' contained (again) 141 participants, whose AW was larger for the affect dimension (the group-level mean attention to happiness was: *w_happiness* = .87), compared to the 'neutral/hair group' (mean *w_happiness* = .19), consisting of 167 participants.

Also, a number of participants were not distinguishable from guessers. In the sad condition, we identified five guessers, with one of them being also, the only, guesser in the happy condition.

To evaluate the fit of the model, we inspected the group-level posterior predictive for all categorization tasks. To give an idea of how well the model performs, Figure 2 presents the posterior predictive for one of the two task versions in the sad condition**.** Each panel shows a schematic representation of the stimulus space (five by five stimuli, represented by the squares), with the corner stimuli being prototypes. The panels depict the latent groups. In each square, the background color provides an indication of the model's posterior prediction for the corresponding stimulus and group, as a gradient between the top left prototype (orange) and lower right prototype (blue). The stronger the color matches the prototype, the more firmly the model predicts classification in the corresponding category. In each square, the circle color is an indication of the average observed classification count of the corresponding stimulus, across all participants in the group. Thus, matching colors between background and circle provide insight in the match between the model's posterior predictions and the observed data.

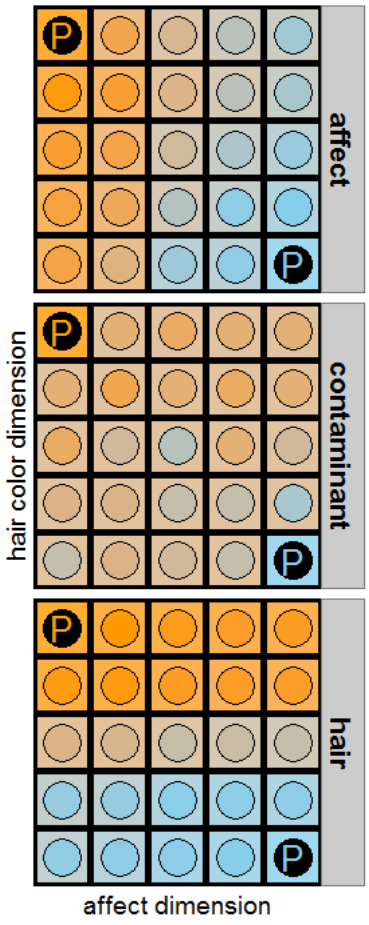


Figure 2. Posterior predictive check, see text for details.

Inspection of Figure 2 reveals that the model successfully captures the patterns within each latent group. Also, the predicted and observed stimulus groupings are sensible considering the AWs applied in each group, with the affect group categorizing the stimuli according to the affect dimension, and the hair group according to the hair color dimension.

Regression Analyses

After excluding five participants who were assigned to the guessing group, multiple regression analyses were performed to investigate the relationship between attentional preferences, as operationalized by the individual-level estimates of w_a, and depressive symptoms, anxiety symptoms, brooding, neuroticism, and current mood. All variables showed sufficient variability (see Table 1 for the descriptive statistics). Variance inflation factors (VIF; Hair, Anderson, Tatham, Black, 1995) indicated only a low degree of multicollinearity in our model (largest VIF was 3.22, well below the threshold of 10).

Table 1: Descriptive statistics.

	SD	mean	range
CES-D	10.06	17.83	0 - 52
HADS_D	3.63	4.41	0 - 18
HADS_A	3.97	7.07	0 - 20
TIPI_E	1.41	4.36	1 - 7
RRS	3.52	10.47	5 - 20
Mood	0.60	3.11	1 - 5
w_sadness	0.37	0.51	0.04 - 0.99
w_happiness	0.37	0.49	0.05 - 0.98

Note. w_sadness is w_a in the sad condition, whereas w_happiness is w_a in the happy condition.

The regression analyses were performed using the BayesFactor package (Morey & Rouder, 2013). A BF compares the evidence for the null model with the evidence for the alternative model. Given the null model in the numerator, a BF > 1 indicates evidence in favor of the null model, whereas a BF < 1 indicates evidence in favor of the alternative model.

To test the effect of each predictor (e.g., depressive symptoms), we compared a restricted model containing all predictors, except for the predictor of interest (null model), against a full model containing all predictors (alternative model). As a sensitivity analysis, both medium and ultrawide scale factors were used to calculate the Bayes factors. As can be seen in Table 2, the Bayes factors showed evidence in favor of the null model for all predictors in both conditions (with the exception of current mood in the sad condition when a medium scale factor was used). The strength of the evidence depends on the exact choice of scale factor. When the scale factor is medium, the evidence in favor of the null

model is weak, and thus these results are best interpreted as being inconclusive. With increasing scale factor, the evidence in favor of the null grows stronger, but it is never very strong. Overall, these analyses suggest that no strong evidence for meaningful associations could be found between the symptom and traits scores and the AWs. The near zero partial correlation coefficients (using the ppcor package; Seongho, 2015), confirm this picture, as can be found in Table 2.

Table 2: Overview of the BFs (> 1 indicates support for the null model), and the corresponding partial correlation coefficients.

	sad			happy		
	BFm	BFu	PC	BFm	BFu	PC
CES-D	1.89	3.23	-.00	2.59	4.60	.02
HADS_D	2.50	4.38	.03	2.84	5.07	.03
HADS_A	2.75	4.87	.03	3.05	5.49	.02
TIPI_E	2.82	4.99	-.02	3.04	5.48	-.02
RRS	2.32	4.05	-.03	2.53	4.48	.04
mood	0.87	1.39	-.09	1.17	1.93	-.08

Note. BFm = Bayes factor with medium r scale, BFu = Bayes factor with ultrawide r scale. PC = partial correlation coefficient.

Discussion

Applying a categorization approach to the assessment of attentional biases in people with varying levels of depressive symptoms revealed weak to moderate evidence for the absence of an association between severity of depressive symptoms and an attentional bias to sadness or happiness. Similar BFs were observed for the other predictors of interest: anxiety symptoms, brooding, neuroticism, and current mood. In the light of the small Bayes factors, especially when using a medium scale factor, we cannot make strong statements about rejecting the alternative model, or accepting the null model.

These findings are in line with previous inconsistencies in results obtained with the frequentist approaches investigating attentional biases in the context of depression (Peckham et al., 2010). In some studies p values above .05, whereas in other studies, p values below .05 were found, without a clear explanation as to when to expect significant results and when not.

An important limitation of the current study that could explain the results, is the recruited sample. A high functioning student sample was recruited. Though we could observe a reasonable amount of variability in depression scores, the sample may have been too healthy to detect attentional biases related to depressive symptomatology. A next step is to apply this approach to data obtained in a sample of more severely depressed participants.

We believe the modelling approach demonstrated here, has a number of advantages that might prove useful in helping to solve the elusiveness of attentional biases in the context of depression. First, attentional preferences were extracted from a model, that excluded people whose behavior could better be predicted by a guessing model, instead of the prototype model. This means that data resulting from random behavior were filtered out. Second, by considering a specific parameter in a model to conceptualize attentional biases, other processes that might influence participants' behavior in the task, such as people's discrimination abilities (*c* parameter), were factored out. Third, assessing attentional processes by considering their impact on categorization behavior could be quite insightful, given that classification decisions reflect the way in which people organize and structure their world. Attentional bias indices obtained by analyzing categorization behavior can give us an idea about how strongly attentional preferences influence the way in which people perceive and organize their environment.

Acknowledgments

The research leading to the results reported in this paper was supported in part by the Research Fund of KU Leuven (GOA/15/003) and by the Interuniversity Attraction Poles programme financed by the Belgian government (IAP/P7/06).

References

Bartlema, A., Lee, M., Wetzels, R., & Vanpaemel, W. (2014). A Bayesian hierarchical mixture approach to individual differences: Case studies in selective attention and representation in category learning. *Journal of Mathematical Psychology, 59*, 132-150. http://dx.doi.org/10.1016/j.jmp.2013.12.002

Beck, A. T. (1976). *Cognitive therapy and the emotional disorders*. New York: International University Press.

Bradley, B. P., Mogg, K., & Lee, S. C. (1997). Attentional biases for negative information in induced and naturally occurring dysphoria. *Behaviour Research and Therapy*, *35*(10), 911-927. http://dx.doi.org/10.1016/S0005-7967(97)00053-3

Fotomorph 13.9. Softland SRL. http://fotomorph.soft112.com/

GIMP, G. (2008). *Image manipulation program.* User Manual, Edge-Detect Filters, Sobel, The GIMP Documentation Team.

Gosling, S. D., Rentfrow, P. J., & Swann, W. B. Jr., (2003). A very brief measure of the Big-Five personality domains. *Journal of Research in Personality, 37*, 504-528. doi:10.1016/S0092-6566(03)00046-1

Hair, J. F. Jr., Anderson, R. E., Tatham, R. L. & Black, W. C. (1995). *Multivariate Data Analysis* (3rd ed). New York: Macmillan.

Koster, E. H., De Lissnyder, E., Derakshan, N., & De Raedt, R. (2011). Understanding depressive rumination from a cognitive science perspective: The impaired disengagement hypothesis. *Clinical Psychology Review*,

31(1), 138-145. http://dx.doi.org/10.1016/j.cpr.2010.08.005

Kruschke, J. K., & Vanpaemel, W. (2015). Bayesian estimation in hierarchical models. In J. Busemeyer, J. Townsend, Z. J. Wang, & A. Eidels (Eds.), *The Oxford Handbook of Computational and Mathematical Psychology* (pp. 279-299). Oxford: Oxford University Press.

Lundqvist, D., Flykt, A., & Öhman, A. (1998). The Karolinska Directed Emotional Faces - KDEF, CD ROM from Department of Clinical Neuroscience, Psychology section, Karolinska Institutet, ISBN 91-630-7164-9.

Morey, R. D. & Rouder, J. N. (2013). Bayes Factor: Computation of Bayes factors for common designs. R package version 0.9.4. Retrieved from http://CRAN.R-project.org/package=BayesFactor

Nolen-Hoeksema, S., & Morrow, J. (1991). A prospective study of depression and posttraumatic stress symptoms after a natural disaster: The 1989 Loma Prieta Earthquake. *Journal of Personality and Social Psychology, 61*(1), 115-121. http://dx.doi.org/10.1037/0022-3514.61.1.115

Nosofsky, R. M. (1987). Attention and learning processes in the identification and categorization of integral stimuli. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 13*(1), 87–108. http://dx.doi.org/10.1037/0278-7393.13.1.87

Peckham, A. D., McHugh, R. K., & Otto, M. W. (2010). A meta-analysis of the magnitude of biased attention in depression. *Depression and Anxiety, 27*(12), 1135 - 1142. doi: 10.1002/da.20755

Plummer, M. (2011). JAGS Version 3.1. 0 user manual. International Agency for Research on Cancer, 2.

Radloff, L. S. (1977). Center for Epidemiological Studies Depression Scale. *Applied Psychological Measurement 1*, 385-401.

Rodebaugh, T. L., Scullin, R. B., Langer, J. K., Dixon, D. J., Huppert, J. D., Bernstein, A., ... & Lenze, E. J. (2016). Unreliability as a threat to understanding psychopathology: The cautionary tale of attentional bias. *Journal of Abnormal Psychology, 125*(6), 840-851. http://dx.doi.org/10.1037/abn0000184

Seongho K. (2015). ppcor: Partial and semi-partial (part) correlation. R package version 1.1. https://CRAN.R-project.org/package=ppcor

Spinhoven, P. H., Ormel, J., Sloekers, P. P. A., Kempen, G. I. J. M., Speckens, A. E. M., & Van Hemert, A. M. (1997). A validation study of the Hospital Anxiety and Depression Scale (HADS) in different groups of Dutch subjects. *Psychological Medicine, 27*(2), 363-370.

Viken, R. J., Treat, T. A., Nosofsky, R. M., McFall, R. M., & Palmeri, T. J. (2002). Modeling individual differences in perceptual and attentional processes related to bulimic symptoms. *Journal of Abnormal Psychology, 111*(4), 598. http://dx.doi.org/10.1037/0021-843X.111.4.598

Voorspoels, W., Rutten, I., Bartlema, A., Tuerlinckx, F., & Vanpaemel, W. (in press). Sensitivity to the prototype in children with high-functioning autism spectrum disorder: An example of Bayesian cognitive psychometrics. *Psychonomic Bulletin & Review*.

Zeigenfuse, M.D., & Lee, M.D. (2010). A general latent-assignment approach for modeling psychological contaminants. *Journal of Mathematical Psychology*, 54, 352-362. doi: 10.1016/j.jmp.2010.04.001