How to Exploit Diversity for Scientific Gain

Using Individual Differences to Constrain Cognitive Theory

Edward K. Vogel & Edward Awh

University of Oregon

ABSTRACT—*People often show considerable systematic variability in their ability to perform many different cognitive tasks. In this article, we argue that by combining an individual-differences approach with an experimental-cognitive-neuroscience approach one can often further constrain potential theories of the underlying cognitive mechanisms. In support of this proposal, we outline three basic benefits of using an individual-differences approach: validating neurophysiological measures, demonstrating associations among constructs, and demonstrating dissociations among apparently similar constructs. To illustrate these points, we describe recent work by us and other researchers that utilizes each of these techniques to address specific questions within the domain of visual working memory. It is our hope that some of these techniques for utilizing individual variability may be applied to other domains within cognitive neuroscience.*

KEYWORDS—*individual differences; attention; working memory*

People vary considerably across countless dimensions: physical characteristics and political and religious beliefs, as well as specific skills and aptitudes. This variability can also be observed at a finer level in terms of how individual brains work: Some people have crisp clear memories of long-ago events, while others can't even recall what they did this morning; some can focus attention on an object or task for an extended period of time, while others are easily distracted by anything other than what they are trying to accomplish. This rich diversity in cognitive ability arises through a mix of genetic and environmental contributions and can be thought of as the results of "nature's experiments" (Cronbach, 1957). However, in the context of most standard cognitive neuroscience studies, this variability across individuals is typically treated as a nuisance or as error variance, potentially obscuring differences between levels of their independent variables. Treating individual differences in this way makes sense for cognitive neuroscientists attempting to understand how cognitive constructs such as perception, attention, and memory operate at the general level. Most cognitive neuroscientists are interested in understanding how everyone thinks, not trying to catalog and characterize the entire range of abilities across the population or understand how and why a given individual thinks differently from another. In this article, we argue that these are not mutually exclusive goals, and that by characterizing individual differences in ability within the context of a sound experimental design, one can often learn a great deal more about how a cognitive process operates at a basic level.

In the 50 years since Cronbach's (1957) classic article, there have been many studies that have successfully combined an individual-differences approach with a standard experimental one across several areas of psychology and many domains within the study of cognition (e.g., Kirchoff & Buckner, 2006; Thompson-Schill, Braver, & Jonides, 2005; Wilmer & Nakayama, 2007; Yovel & Kanwisher, 2005). Because the study of individual differences in cognition covers a very broad area, our more manageable goal for this article is to detail how we and others have recently used this combined approach to address specific questions within the subdomain of visual working memory. In particular, we focus on three primary virtues or benefits of the individual-differences approach, and how we have used each of these to gain traction on some specific issues within the visual-working-memory domain. We use specific issues in visual working memory as test cases for describing how the rich data on individual variability can be exploited to help constrain theory, with the hope that some of these tricks can be exported to other domains within cognitive neuroscience.

Address correspondence to Edward Vogel, 1227 Department of Psychology, University of Oregon, Eugene, OR 97403-1227; e-mail: vogel@uoregon.edu.

USING INDIVIDUAL DIFFERENCES TO VALIDATE NEUROPHYSIOLOGICAL MEASURES

In the mid-1990s, Luck and Vogel (1997) were interested in measuring how many visual objects a person could hold in working memory at the same time. To this end, they developed a task in which participants are shown a brief array of simple objects (e.g., colored squares; see Fig. 1a) that they must remember over a blank gap of about 1 second. After this blank gap, the participants are presented either with a test array that is identical to the first array or with an array in which a single item has changed color, and they simply report whether the two arrays are the same or different. Using this change-detection procedure, Luck and Vogel estimated that visual-working-memory capacity is limited to approximately 3 to 4 items. Interestingly, although the average capacity in these initial studies consistently hovered around 3 to 4 items, there were actually large and consistent differences in performance across individuals, ranging from about 1.5 objects up to over 6 objects (see Fig. 1b). At the time, these individual differences were essentially disregarded by the authors as noise, with the real interest being focused on the surprisingly low mean of 3 to 4 objects.

It is now clear that significant theoretical progress can be made by examining the neural activity that mediates performance in a cognitive measure. For example, while it is usually assumed that change-detection scores are determined by the number of items that can be maintained during the memory period, performance might also be influenced by interference during the retrieval stage of the task. However, if task performance is directly associated with neural activity during the retention interval of the task, then behavioral performance can be more confidently attributed to the maintenance of information in working memory. Likewise, this convergence of behavioral and neural data can reinforce the interpretation of the neural data. A well-known difficulty in cognitive neuroscience is determining whether the neural activity of interest is causally related to the cognitive ability of interest or whether it is epiphenomenal. For example, if an increase in brain activation is found when more information is held in memory, this could be caused either by increased activity in brain regions that mediate the process of interest or by a generalized increase in neural activity when the task becomes more effortful. However, if these increases in neural activity are predicted by behavioral performance in the task, then one can more confidently conclude that the neural activity is a valid measure of the underlying cognitive construct.

Along these lines, Vogel and Machizawa (2004) developed an event-related potential (ERP; scalp-recorded electrical brain waves) measure of maintaining information in visual working memory by using a variation of the change-detection task. In this study, when they time-locked the ERPs to the onset of the memory array, they observed a large negative voltage wave beginning after 200 milliseconds that persists throughout the retention interval until presentation of the test array. They referred to this activity as the contralateral delay activity (CDA). An exciting attribute of the CDA was that its amplitude increased as a function of the number of items the participant was remembering on a given trial. However, while the CDA increased for arrays of 1, 2, or 3 objects, it reached a limit somewhere between 3 and 4 objects—showing no further increases for larger array sizes (see Fig. 1c). These results on their own were highly informative, because they demonstrated that this novel neurophysiological data showed a similar characteristic to behavioral performance in this task—namely, that it is limited to representing approximately 3 items at a time. Moreover, the authors took the logic of a behavior–neurophysiology coupling one step further, by testing whether the exact point at which the CDA reached a limit was different for each subject depending upon his or her specific memory capacity. Indeed, they found a strong correlation ($r = .78$) between an individual's memory capacity and the point at which the CDA reached asymptote (see Fig. 1d).

Without the individual-difference analysis, this study would still have made a solid case that the CDA likely reflected visual-working-memory representations. However, the argument would have rested primarily on the finding that the observed neural limit just happened to occur at approximately the same number of items as the typical behavioral limit. By contrast, the individual-difference analysis allowed the investigators to make a much stronger argument that was based upon the entire distribution of capacity scores, not just the mean score. Thus the consideration of individual variability in this study helped to validate the CDA as a neural measure of capacity limitations in visual working memory. Indeed, we see this specific usage of the individual-differences approach as having the most straightforward application to cognitive-neuroscience research. In fact, there are already several neuroimaging studies that have used this approach. For example, Todd & Marois (2005) found that activation in the intraparietal sulcus increased as a function of number of items in a change-detection task. Similarly, this function reached asymptote at around 3 to 4 items and was also predictive of individual differences in visual-working-memory capacity.

USING INDIVIDUAL DIFFERENCES TO HELP DEMONSTRATE ASSOCIATIONS AMONG COGNITIVE CONSTRUCTS

The cognitive constructs of attention and working memory have historically been tightly linked. In fact, some models of working memory have gone as far as proposing that they are essentially the same construct, defining working memory as the active representations in memory that are within the focus of attention (Cowan, 2001). These links have been further strengthened by the results of numerous neuroimaging studies that have shown substantial overlap in the cortical areas active during attention and working-memory tasks (e.g., Awh & Jonides, 2001). However, the demonstration of a coarse anatomical overlap is gen-

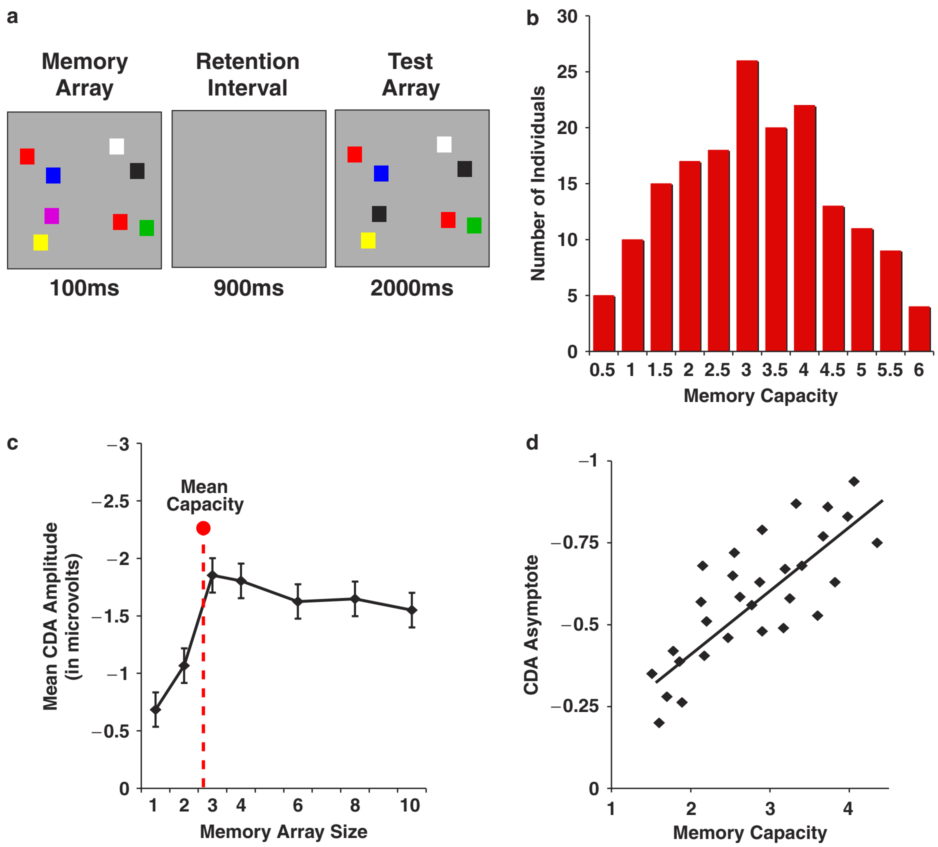

Fig. 1. Measuring visual-working-memory capacity. Panel a shows change-detection stimuli and procedure, in which subjects must attempt to remember colors from the memory array and then, after a brief interval, detect any color changes in the test array. Panel b shows distribution of visual-working-memory capacity estimates from 170 healthy undergraduates, with a mean of 2.88 items and a standard deviation of 1.04 (unpublished data from E.K. Vogel). Panel c shows mean amplitude of contralateral delay activity (CDA; a measure of electrical brain activity) as a function of the number of items in a memory array. Panel d shows correlation of the asymptotic limit of the CDA (derived as the difference in amplitude from 2 items to 4 items) and individual memory capacity. Panel a adapted from Luck & Vogel (1997); panels c and d adapted from Vogel & Machizawa (2004).

erally not sufficient for establishing a functional relationship between two constructs. A common alternative approach to establishing whether two constructs are either functionally isomorphic or at least tightly associated is to measure whether they strongly covary. The examination of individual differences in performance is particularly useful for establishing such relationships because it allows for tests of a given association along the entire range of values of performance. That is, if two constructs are tightly related, then an individual's performance on a task that is primarily limited by one construct should be predictive of his or her performance on a task limited by the related construct.

This general approach has been used successfully many times within the working-memory domain, with individual differences in working-memory capacity being shown to be highly predictive of several relatively distal constructs such as intelligence, reasoning ability, and reading comprehension (Cowan et al., 2005; Daneman & Carpenter, 1980; Miyake, Just, & Carpenter, 1994). Moreover, Kane and Engle have used this approach to measure the more proximal associations between memory capacity (operation span) and an individual's performance on various kinds of attention tasks. For example, their work has demonstrated that individuals with a high memory span tend to perform much better on antisaccade tasks, in which they must look away from the location of an object that just appeared, and Stroop tasks, in which they must report the color of a word while ignoring the conflicting meaning of the word (e.g., the word "black" drawn in red), than low-memory-span individuals do (Kane, Bleckley, Conway, & Engle, 2001; Kane & Engle, 2003). This is particularly impressive because most of these tests appear on the surface to have fairly negligible memory requirements, yet individual differences in performance can be predicted by working-memory ability, so it is clear that some common factor underlies both constructs.

Beyond simply demonstrating an association between the two constructs, Kane and Engle's work has provided new insights into the nature of the limits on working-memory capacity. Specifically, their results suggest that memory capacity may have more to do with how well an individual can selectively attend to information than with how much information he or she can hold at a given time. Following this general logic, Vogel, McCollough, & Machizawa (2005) tested whether an individual's memory capacity predicted how efficiently he or she could control what information was stored in visual working memory. To do this, they measured the amplitude of the CDA component of the ERP while subjects voluntarily attempted to store subsets of items from a memory array. For example, in one experiment, subjects were asked to remember only the red items in a display consisting of a mix of red and blue oriented bars. They found that high-capacity individuals were highly efficient at storing only the relevant items and disregarding the irrelevant items. By contrast, the low-capacity subjects were found to be highly inefficient at excluding the irrelevant items, unnecessarily storing all items in the array into visual working memory. These results are somewhat counterintuitive, because they indicate that the low-capacity subjects often store more information in memory than high-capacity subjects do. However, this extra information is often irrelevant to the current task and hinders access to the relevant information. Thus, the pattern of individual differences in this study reveals more than a general association between attention and memory. These data show that the ability to control what information is stored in memory may be the primary limiting factor in measures of memory ability.

USING INDIVIDUAL DIFFERENCES TO DEMONSTRATE DISSOCIATIONS BETWEEN SIMILAR CONSTRUCTS

While it is common to use individual differences in performance as a means of demonstrating associations between constructs, a powerful but underutilized approach is to use this variability as a way of demonstrating that two similar constructs can be dissociated. This general idea was proposed by Underwood (1975); he argued that if a theory proposes that two variations of a given task are determined by the same underlying construct, then an individual's performance on task A should predict his or her performance on task B. If the two measures of performance do not correlate at all—assuming they are both reliable measures—the theory relating the two tasks should be dropped. We have recently used this general approach as a way of teasing apart the constructs of number and resolution as separate factors that underlie visual-working-memory capacity.

First, a little background information will be helpful. In 2004, Alvarez and Cavanagh were interested in measuring how object complexity influenced how many items could be held in visual working memory (Alvarez & Cavanagh, 2004). To do this, they used a change-detection task in which the to-be-remembered items in a given array were drawn from several categories of visual objects that ranged in complexity or "information load." They found a strong inverse relationship between working-memory performance and complexity. When the objects were simple items such as colored squares, participants could remember on average about 4 of them. However, when the memory items were more complex—such as Chinese characters or 3-D shaded cubes—memory-capacity estimates were in the range of 1 to 1.5 objects. From these results, the researchers concluded that memory capacity is not determined only by the number of objects but also by the total amount of information contained within the objects. Thus, number and resolution were proposed to be intimately intertwined: The higher the resolution, the fewer the number of items that can be held in memory, and vice versa.

A further wrinkle in the complexity issue was provided in a neuroimaging study by Xu and Chun (2006), in which three primary cortical areas were found to show load-sensitive activation in a change-detection task. Of these three, two (the lateral occipital complex and the superior intraparietal sulcus) were highly sensitive to object complexity, reaching asymptotic activation for smaller numbers of complex items than of simple items. By contrast, a third region (the inferior intraparietal sulcus) was completely insensitive to the complexity of the objects, always reaching asymptotic activation at approximately 4 items. These neuroimaging results constrained the existing models of memory capacity by suggesting that both complexity and number of items determine working-memory performance. Further, they provided initial evidence that these constructs might be somewhat distinct factors, because separate cortical areas are sensitive to each.

Following this general line of logic, we (Awh, Barton, & Vogel, 2007) reasoned that poor performance for complex objects might not be the consequence of an inability to maintain multiple objects in memory but, rather, might be due to limitations in discriminating whether the object had changed or not. That is, as the complexity of a given object category increases, the subjective similarity among members of the category also increases. Thus, poorer working-memory performance might be due to more errors in comparing an object in memory with the highly similar object it changed into. To test this hypothesis, we presented observers with memory arrays that were a mix of two categories of complex objects (Chinese characters and 3-D cubes) and manipulated what type of change occurred on a given trial: either a within-category change (i.e., cube to cube or character to character; see Fig. 2a) or a between-category change (i.e., cube to character or character to cube). On the within-category change trials (when sample–test similarity was high), working-memory performance was quite poor, and the results replicated Alvarez & Cavanagh (2004). By contrast, on the between-category trials (when sample–test similarity was low), working memory for these complex items was very good and was equivalent to performance for remembering simple objects (i.e., colored squares). These data suggested that errors in change detection with complex objects were not caused by a failure to represent the items in memory but by the fact that the repre-

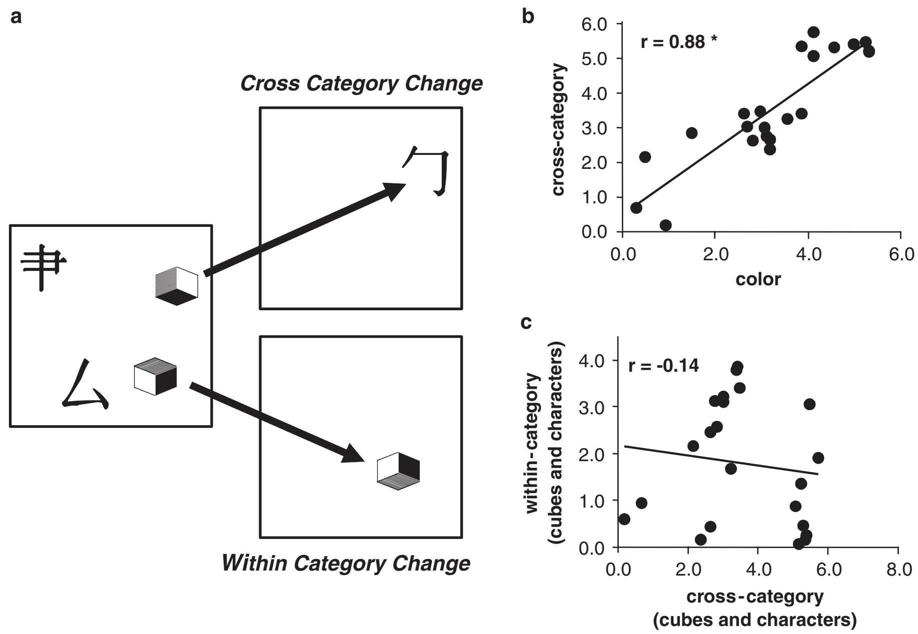

Fig. 2. Change detection and working-memory capacity. Panel a shows cross- and within-category change-detection stimuli (Chinese characters and 3-D cubes) and procedure used in Awh, Barton, and Vogel (2007). Panel b shows correlation between working-memory capacity estimates for colors and for complex items in which the changed item was from a different category. Panel c shows correlation between working-memory capacity estimates for a complex item in which there was a cross-category change or a within-category change. Panel c adapted from Awh, Barton, & Vogel (2007).

sentations did not have sufficient resolution for the detection of very small changes.

Our (Awh et al., 2007) data suggested that performance in the between-category change trials was limited by the total number of items that could be maintained in working memory whereas performance in the within-category condition was instead limited by the resolution of those representations. In this case, the consideration of individual differences became a powerful means by which we could test two critical hypotheses. First, was performance in the between-category change condition limited by the same ability as that which limits the maintenance of simple colored squares? We tested for this expected association and found a strong positive correlation between an individual's memory capacity for colored squares and his or her capacity for complex items that changed across categories ($r = .88$; see Fig. 2b). Second, was performance on the between-category change trials limited by a different construct from that which limits performance on the within-category change trials? We tested for this expected dissociation and found no significant correlation of performance between these two types of trials ($r = .14$; see Fig. 2c), despite the fact that both performance measures were from the same task, blocks of trials, and objects. Moreover, within-category change performance was shown to be a reliable measure because character-to-character performance strongly correlated with cube-to-cube performance ($r = .66$). Thus, while the data showed a clear association between two measures of the "number of items" that could be maintained and two measures of "mnemonic resolution," number and resolution appear to represent distinct aspects of individual memory ability. Of course, these results don't discount the important observation that complexity plays an important role in working-memory capacity, but they do argue against models that propose that number and resolution are the same intertwined construct. If a single shared resource determined both the number and resolution of representations in working memory, then there should have been a strong correlation between these two ability measures. Indeed, these results, when coupled with Xu & Chun's (2006) neuroimaging data, predict that there should be neural measures that are predicted by an individual's mnemonic resolution but that are not related to the number of items that they can hold in memory. Hopefully, future studies exploiting this powerful individual-differences approach will be able to test this prediction.

CONCLUSIONS

The consideration of individual differences in performance can provide a powerful addition to many standard cognitive-neuroscience studies because it can help further constrain existing cognitive theories. We have outlined three general benefits of this approach and discussed how we have utilized these techniques to address specific issues within research on visual working memory. However, these techniques should be applicable to any domain within cognitive neuroscience, so long as the performance measure to be used is a reliable measure of the construct of interest and there is not much restriction of the range of observed values (e.g., ceiling or floor effects). It is our hope that many others will begin to exploit this wealth of systematic diversity across individuals to further specify the cognitive constructs and mechanisms that we all share.

Recommended Reading

Awh, E., Barton, B., & Vogel, E.K. (2007). (See References). A study that used an individual-differences approach to determine whether the "resolution" of items in memory determines the "number of items" that can be held in working memory.

Cowan, N., Elliott, E.M., Saults, J.S., Morey, C.C., Mattox, S., Hismjatullina, A., et al. (2005). (See References). A study that examined the within-subject correlations between many measures of attention, working memory, and intelligence.

Cronbach, L.J. (1957). (See References). A classic article describing the virtues of combining the individual-differences and standard experimental approaches.

Vogel, E.K., & Machizawa, M.G. (2004). (See References). A study showing that an electrical brain wave is highly sensitive to an individual's specific working-memory capacity.

Vogel, E.K., McCollough, A.W., & Machizawa, M.G. (2005). (See References). A study providing evidence that low-capacity individuals are much poorer at controlling what is stored in memory than are high-capacity individuals.

Acknowledgments—The research reported here was supported by grants from the National Science Foundation to EKV and the National Institute of Mental Health to EA.

REFERENCES

Alvarez, G.A., & Cavanagh, P. (2004). The capacity of visual short-term memory is set both by visual information load and by number of objects. *Psychological Science*, *15*, 106–111.

Awh, E., Barton, B., & Vogel, E.K. (2007). Visual working memory represents a fixed number of items regardless of complexity. *Psychological Science*, *18*, 622–628.

Awh, E., & Jonides, J. (2001). Overlapping mechanisms of attention and working memory. *Trends in Cognitive Sciences*, *5*, 119–126.

Cowan, N. (2001). The magical number 4 in short-term memory: A reconsideration of mental storage capacity. *Behavioral and Brain Sciences*, *24*, 87–185.

Cowan, N., Elliott, E.M., Saults, J.S., Morey, C.C., Mattox, S., Hismjatullina, A., et al. (2005). On the capacity of attention: Its estimation and its role in working memory and cognitive aptitudes. *Cognitive Psychology*, *51*, 42–100.

Cronbach, L.J. (1957). The two disciplines of scientific psychology. *American Psychologist*, *12*, 671–684.

Daneman, M., & Carpenter, P.A. (1980). Individual differences in working memory and reading. *Journal of Verbal Learning and Verbal Behavior*, *19*, 450–466.

Kane, M.J., Bleckley, M.K., Conway, A.R.A., & Engle, R.W. (2001). A controlled attention view of working memory capacity. *Journal of Experimental Psychology: General*, *130*, 169–183.

Kane, M.J., & Engle, R.W. (2003). Working memory capacity and the control of attention: The contributions of goal neglect, response competition, and task set to Stroop interference. *Journal of Experimental Psychology: General*, *132*, 47–70.

Kirchoff, B.A., & Buckner, R.L. (2006). Functional-anatomic correlates of individual differences in memory. *Neuron*, *51*, 263–274.

Luck, S.J., & Vogel, E.K. (1997). The capacity of visual working memory for features and conjunctions. *Nature*, *390*, 279–281.

Miyake, A., Just, M.A., & Carpenter, P.A. (1994). Working memory constraints on the resolution of lexical ambiguity: Maintaining multiple interpretations in neutral contexts. *Journal of Memory & Language*, *33*, 175–202.

Thompson-Schill, S.L., Braver, T.S., & Jonides, J. (2005). Individual differences. *Cognitive, Affective and Behavioral Neuroscience*, *5*, 115–116.

Todd, J., & Marois, R. (2005). Posterior parietal cortex activity predicts individual differences in short-term memory capacity. *Cognitive, Affective and Behavioral Neuroscience*, *5*, 144–155.

Underwood, B.J. (1975). Individual differences as a crucible in theory construction. *American Psychologist*, *30*, 128–134.

Vogel, E.K., & Machizawa, M.G. (2004). Neural activity predicts individual differences in visual working memory capacity. *Nature*, *428*, 748–751.

Vogel, E.K., McCollough, A.W., & Machizawa, M.G. (2005). Neural measures reveal individual differences in controlling access to working memory. *Nature*, *438*, 500–503.

Wilmer, J.B., & Nakayama, K. (2007). Two distinct visual motion mechanisms for smooth pursuit: Evidence from individual differences. *Neuron*, *54*, 987–1000.

Xu, Y., & Chun, M. (2006). Dissociable neural mechanisms supporting visual short-term memory for objects. *Nature*, *440*, 91–95.

Yovel, G., & Kanwisher, N. (2005). The neural basis of the behavioral face-inversion effect. *Current Biology*, *15*, 2256–2262.