

Can you *experience* ‘top-down’ effects on perception?: The case of race categories and perceived lightness

Chaz Firestone · Brian J. Scholl

Published online: 18 December 2014
© Psychonomic Society, Inc. 2014

Abstract A recent surge of research has revived the notion that higher-level cognitive states such as beliefs, desires, and categorical knowledge can directly change what we see. The force of such claims, however, has been undercut by an absence of visually apparent demonstrations of the form so often appealed to in vision science: such effects may be revealed by statistical analyses of observers’ responses, but you cannot literally *experience* the alleged top-down effects yourself. A singular exception is an influential report that racial categorization alters the perceived lightness of faces, a claim that was bolstered by a striking visual demonstration that Black faces appear darker than White faces, even when matched for mean luminance. Here, we show that this visually compelling difference is explicable in terms of purely low-level factors. Observers who viewed heavily blurred versions of the original Black and White faces still judged the Black face to be darker and the White face to be lighter even when these observers could not perceive the races of the faces, and even when they explicitly judged the faces to be of the same race. We conclude that the best subjectively appreciable evidence for top-down influences on perception does not reflect a genuinely *top-down* effect after all: instead, such effects arise from more familiar (if subtle) bottom-up factors within visual processing.

Keywords Lightness perception · Race perception · Modularity · Top-down effects · Cognitive penetrability

C. Firestone (✉) · B. J. Scholl (✉)
Department of Psychology, Yale University, Box 208205, New Haven, CT 06520-8205, USA
e-mail: chaz.firestone@yale.edu
e-mail: brian.scholl@yale.edu
URL: <http://www.yale.edu/perception/>

Introduction

In contrast to the traditional “modular” view of perception, according to which the unconscious inferences that determine what we see are driven largely or only by the patterns of light striking the eyes (Fodor, 1983; Pylyshyn, 1999), a surge of recent research has suggested that higher-level states such as intentions, desires, social attitudes, and categorical knowledge can directly affect what we see. For example, it has been reported that desirable objects such as money or chocolate appear closer than do neutral or undesirable objects (Balcetis & Dunning, 2010); that wearing a heavy backpack makes hills look steeper (Bhalla & Proffitt, 1999); and that reflecting on an unethical deed makes the world look darker (Banerjee, Chatterjee, & Sinha, 2012). Such empirical reports, among hundreds of others, have revived claims (previously popular during the “New Look” movement from the middle of the last century) that perception is a ‘constructive’ process that consults the rest of the mind for input in delivering a visual percept of the environment (for recent reviews, see Collins & Olson, 2014; Dunning & Balcetis, 2013; Lupyan, 2012; Proffitt, 2006; Vetter & Newen, 2014; for philosophical commentary, see Raftopoulos & Zeimbekis, 2014).

Top-down demonstrations?

These studies advertise their results as effects on perception *per se*; the relevant manipulations are said to literally alter our visual experiences. However, an awkward fact about these claims is that nearly all such examples fail to be subjectively noticeable—a fact that one can appreciate for oneself in everyday life, given the general nature of such claims. For example, if you put a \$1 bill next to a \$100 bill, the latter does not appear closer; if you stand at the base of a hill and put on a heavy backpack, the hill does not look noticeably steeper; and if you think now of a past transgression, the room does not

seem noticeably dimmer. Of course, such observations do not by themselves rule out genuinely perceptual interpretations of such phenomena, given how much visual processing occurs unconsciously. But the absence of subjectively appreciable evidence is, at the very least, unusual within the broader context of vision science—where data and experiments about what we see are routinely accompanied by “demonstrations” in which interested observers can experience the relevant phenomena for themselves in often-dramatic fashion. Indeed, many foundational discoveries about the visual system have been (and continue to be) made before collecting even a single data point. For example, no experiments are needed to convince us of the reality of illusory contours, change-blindness, or apparent motion; they are, as it were, “apparent”.

The lack of visual demonstrations is also an *inconvenient* fact about claims of top-down influences on perception, because it leaves them open to other challenges. When such claims rely largely or only on subjects’ reports, one may wonder whether subtle non-perceptual influences could explain the experimental results. For example, follow-up experiments have shown that subjects who are given a heavy backpack and are then asked to estimate a hill’s slant can correctly guess the experimenters’ intentions, and eliminating such task demands (by crafting a deceptive cover story for the backpack) eliminates the backpack’s effect on slant estimates (Durgin et al., 2009, Durgin, Klein, Spiegel, Strawser, & Williams, 2012). By contrast, worries about such alternate explanations dissipate when a given perceptual phenomenon can be experienced directly by anyone who so desires.

Levin and Banaji (2006): *Black faces look darker, controlling for mean luminance*

A singular exception to the lack of subjectively appreciable examples of (alleged) top-down effects on perception is the remarkable report of an apparent influence of race categories on lightness perception. Levin and Banaji (2006) created highly controlled composite images of Black and White faces¹ that were matched for mean luminance, but they noticed (and then experimentally verified) that the faces nevertheless appeared to have different lightnesses; in particular, the Black face looked darker than the White face, even when the faces were equally luminant on average (see Fig. 1a). (Following standard usage in vision science, “luminance” refers to the objective intensity of light per unit area radiating from an image; “lightness” refers to the visually perceived shade of a surface—and so may vary independently of luminance, as in all lightness illusions; cf. Adelson, 2000.) As is not unusual with such visually compelling demonstrations, Levin and Banaji note that the inspiration for their studies came quite

¹ Following Levin and Banaji (2006), we use “Black” and “White” to refer to the corresponding race categories.

by accident, when one author noticed the effect while viewing the face stimuli as used (for other purposes) in the other author’s previous research.

Our primary reason for focusing on this report in the present paper is that it is “singular”—by which we mean that it is, to our knowledge, the only purported top-down effect on perception that readers can experience for themselves simply by looking at the stimuli. (See [Discussion](#) for commentary on what “counts” as top-down in this context.) Perhaps for this reason, Levin and Banaji’s report is often cited as the strongest case for a top-down effect on perception (e.g., Collins & Olson, 2014). Another reason for focusing on this report is its explicit regard for salient alternate explanations for such results, including task demands or mere differences in where subjects focus while viewing the images. In each case, Levin and Banaji explicitly discuss how their results may withstand such challenges—which in other treatments are too often mentioned only in passing, if at all. Indeed, without compelling demonstrations, the issue of task demands might have loomed large: subjects in Levin and Banaji’s (2006) studies were told in advance that the experiments were about “how people perceive the shading of faces of different races” (p 504)—instructions that could have led subjects to adjust their responses even without a difference in actual percepts (for discussion, see Rosenthal & Rubin, 1978). However, the compelling visual demonstration assuages such worries, since the difference in lightness between the Black and White faces is so visually compelling.

The current studies

Does the difference in lightness between the Black and White faces reflect a genuinely *top-down* influence on perception? Even though they are matched for mean luminance, there are also several low-level differences between the stimuli. For example, the Black face appears to have a darker jaw (while the White face has darker eyes), and the Black face also appears to be either glossier or more strongly illuminated, as evidenced by lighter patches on its cheekbones and brow (while the White face appears matte and perhaps not as directly lit).

Could such low-level stimulus properties explain the difference in perceived lightness, beyond the races of the faces as such? Of course, this question cannot be answered a priori. On the one hand, subtle factors of this sort do often influence perceived lightness — as when a patch appears darker than surrounding patches because it appears to be under illumination rather than in shadow (Adelson, 2000). On the other hand, such effects arise only from some special properties, and of course the Black and White faces had to differ along some low-level dimensions by definition (since otherwise they would be the same image!). So the key empirical question, not previously explored, is whether the low-level differences

in these particular images give rise to different perceived lightnesses, independently of race.

Rather than try to predict and manipulate the precise low-level differences that might play a role, we tested the possibility of *all possible* low-level explanations at once, using a different strategy. We generated heavily blurred versions of the faces (Fig. 1b) such that most observers did not detect any racial difference between the images, and then asked whether those observers still judged the blurry “Black” stimulus to be darker than the blurry “White” stimulus (despite still being equated for mean luminance). If the two faces appear to have different lightnesses even when they do not appear to have different races (and even for subjects who explicitly guess the face images to have the same race), then there must be some low-level differences between the stimuli, independent of race, that influence perceived lightness in practice.

Experiment 1a: lightness judgments of blurred faces

Method

Participants

One hundred observers were recruited online through Amazon Mechanical-Turk and were monetarily reimbursed. A power analysis based on Levin and Banaji’s (2006) Experiment 1 indicated that only 26 observers would be required to achieve 80 % power to detect the lightness-distortion effect. We conservatively increased this sample size, however, given that (1) our stimuli were degraded, (2) our task involved only a single lightness judgment (instead of a repeated-measures design), (3) we expected to exclude a non-trivial number of subjects (see below), and (4) the cost of additional observers in such an online sample is marginal.

Stimuli

We constructed blurred images (Fig. 1b; 178×223 pixels) by submitting the original Black and White face stimuli (Fig. 1a; from Experiment 1 of Levin & Banaji, 2006) to the “Glass” effect in Adobe Illustrator, which (as we independently

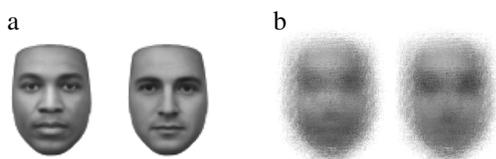


Fig. 1 **a** The face stimuli (matched in mean luminance) from Experiment 1 of Levin and Banaji (2006). **b** The blurred versions of the same stimuli used in the present study, which preserved the match in mean luminance, while obscuring the race information for most observers

verified with the resulting bitmap images) preserved the match in mean luminance and contrast.²

Procedure

Observers viewed the two blurry images side-by-side (on a white background, in a random left-to-right order, with their nearest features horizontally separated by 100 pixels) and made a two-alternative forced-choice (2AFC) judgment comparing their lightnesses. Observers were randomly assigned to answer the question “Which of these images is lighter?” or “Which of these images is darker?”. To prevent suspicion that lightness was the property of interest (since such task demands can contaminate alleged top-down effects; Durgin et al., 2009; Firestone & Scholl, 2014), observers were also asked (on the same page, in a randomized order across observers) which image was blurrier, larger, and more aesthetically pleasing.

After these judgments, the images disappeared, and observers answered two free-response questions to determine how they had interpreted the images before being probed about specific properties. First, observers described the images they had seen in two sentences. Then, after being explicitly told that the images were blurry faces, they provided a two-sentence description of the faces “to help someone pick them out” “from within a set of many other faces”. The images then reappeared, and observers gave their “best guess” about whether the races of the faces looked the “same” or “different”. This question was again embedded within additional (randomly ordered) same/different questions about other aspects of the images: age, gender, shape, facial expression, and facial hair. Finally, observers reported whether they ever suspected that racial categorization was our primary interest, and also whether they had previously participated in an experiment involving blurry or scrambled faces.

Results

Answers to the free-response questions suggested that observers extracted little or no accurate race information from the blurred images. When first asked to describe them, only 3 % of observers mentioned race—and all three explicitly described the faces either as having the same race, or being images of the very same face. (For example, one observer wrote: “They were of a white man in his late 30s. They were both very blurry”.) When asked to describe them again so that someone else could pick them out of a lineup, only 12 % of observers mentioned race in any form—and again, the vast majority of these (10/12 observers) appeared to believe the faces had the same race or belonged to the same person.

² For a step-by-step walkthrough that describes how to make these blurred images, see <http://www.yale.edu/perception/blurring/>.

(For example, one observer wrote: “Young, female caucasian with strong jaw line. Thin cheeks and taught [sic] brow.”) Moreover, the faces were blurry enough that several observers rejected the premise of this question outright. For example, one observer wrote: “Good luck. I couldn’t make out enough detail to tell anything more than that they were faces.” Finally, in response to the same/different question about race (asked while the faces were again in full view), 83 % of observers reported that the faces looked to have the same race. (In fact, numerically, the races were judged to be the most similar of the six properties we asked about. For example, only 46 % of observers answered that the facial hair on the faces was the same—indicating that they were not simply selecting “same” for every question.) Thus, most observers indeed failed to categorize the blurry faces as Black and White (if they categorized their races at all).

Lightness judgments

Of primary interest was whether observers judged the blurry “Black” image to be darker than the blurry “White” image, even when observers could not perceive the races of the faces (according to maximally conservative criteria). Considering only those observers who (1) judged that the races of the faces looked the “same”, (2) indicated that they were never suspicious that the study’s purpose involved race and/or perception of lightness, (3) confirmed that they had never been in a study like this before, and (4) passed several attention checks, 72 % of observers reported that the blurry Black image was darker or that the blurry White image was lighter (47/65 observers; $\chi^2(1, n = 65) = 12.94, P < 0.001$).³

Experiment 1b: direct replication

Given the ease of online data collection and the importance of conducting direct replications, we ran the exact same experiment again, on an independent group of 100 online observers. Considering only those observers who met the exclusion criteria established earlier, 65 % indicated that the blurry Black image was darker or that the blurry White image was lighter (41/63 observers; $\chi^2(1, n = 63) = 5.73, P = 0.017$).

³ We report the results with these exclusion criteria in order to be maximally conservative, but the effect did not depend on the exclusions. (Simply considering all observers who completed the survey, 69 % indicated that the blurry Black image was darker or that the blurry White image was lighter; 69/100, $\chi^2(1, n = 100) = 14.44, P < 0.001$). This is also true of the results from Experiment 2.

Experiment 2a: explicit racial categorization

The foregoing experiments showed that observers judged the stimuli to be differentially light and dark even when they explicitly judged the images to have the “same” race. Still, it may be possible that some observers perceived the faces’ races but nevertheless failed to report this perception. In particular, it is possible that they still perceived some hint of racial difference, and that this ‘hint’ made one of the faces look darker, but that on balance this hint was not enough to push them toward a “different” response on our race measure. For this reason, Experiment 2a replicated the basic design and logic of Experiment 1a, with the exception that observers were now forced to explicitly categorize the race of each face, by selecting from a list of options.

Method

This experiment was identical to Experiment 1a, except as noted here. Five hundred new observers participated. (We very conservatively increased the sample size, expecting to exclude a much greater proportion of observers than in the previous experiments). Instead of giving “same”/“different” judgments for age, gender, shape, facial expression, facial hair, and race, observers saw a drop-down menu underneath each blurry image (which were presented one at a time in a randomized order) with four options for each property, and they selected one of the options for each image. (For example, the options for “expression” were “happy”, “sad”, “fearful”, and “surprised”). For the crucial race judgment, the options were “Caucasian”, “African-American”, “East Asian”, or “Hispanic”. (Obviously this is not a complete list of possible races, but we wanted to limit options to increase statistical power while providing more options than just the faces’ true races). There were no free-response questions.

Results

As before, observers appeared not to have extracted accurate race information from the images. Even when forced to select a race for each image, only 18 % of observers categorized the blurry White face as “Caucasian” and the blurry Black face as “African-American” (and even this may be an overestimate since, after all, we contend that one image really does look darker than the other).⁴ For comparison, a greater proportion of observers categorized at least one of the faces as female (49 %) than accurately categorized the faces’ races. (And the observers were not simply answering randomly. For example,

⁴ The breakdown of race judgments was as follows: for the White image, 72.6 % Caucasian, 8.3 % African-American, 6.5 % East Asian, 12.6 % Hispanic; for the Black image, 45.2 % Caucasian, 28.9 % African-American, 5.7 % East Asian, 20.2 % Hispanic.

95 % of observers judged the faces to be under 40 years of age, whereas only 5 % judged them to be over 40 years of age — when chance would be 50 % for each). We even received an unsolicited personal note from one observer recommending an improvement for the study: “You should allow a ‘none’ option for the races. I didn’t really perceive any, but I felt like I was forced to choose among the options.” (And even this observer judged the Black image to be darker).

Lightness judgments

Of primary interest was whether observers judged the blurry “Black” image to be darker than the blurry “White” image, even when observers could not perceive the races of the faces as measured by their explicit race judgments in a maximally conservative way. Considering only those observers who (1) gave the same explicit race judgment for each face, (2) indicated that they were never suspicious that the study’s purpose involved race and/or perception of lightness, (3) confirmed that they had never been in a study like this before, and (4) passed an attention check, 61 % of observers reported that the blurry Black image was darker or that the blurry White image was lighter (135/223 observers; $\chi^2(1, n = 223) = 9.91, P = 0.002$).

Experiment 2b: direct replication

As in the first set of experiments, we again ran a direct replication of Experiment 2a. The effect replicated, and indeed was even stronger: considering those observers who met the exclusion criteria from Experiment 2a, 66 % indicated that the blurry Black image was darker or that the blurry White image was lighter (122/186 observers; $\chi^2(1, n = 186) = 18.09, P < 0.001$).

Discussion

There may be no more foundational question in the study of visual perception than that posed by Koffka (1935): “Why do things look as they do?” Though vision science often relies on indirect measures and careful experimentation, subjectively appreciable demonstrations rightly remain a ‘gold standard’: in the science of how we see, there is no more compelling evidence than a demonstration that anyone can personally experience. It is for this reason that we emphasize the lack of such demonstrations throughout the very broad literature on top-down effects: though such effects purport to alter perception, they rarely (if ever) are observable first-hand.

What counts as “top-down”?

At first glance, there may seem to be many cases where cognition exerts a subjectively appreciable influence on perception. For example, observers can often voluntarily control which interpretation of a multistable figure (such as a Necker cube) they see. However, such phenomena are typically mediated by changes in eye-movements and/or covert attention. With Necker cubes, for example, voluntary switches from one interpretation to another are occasioned largely or only by changes in the location of gaze (Kawabata et al., 1978) or the focus of spatial attention (Peterson & Gibson, 1991; Toppino, 2003). Such phenomena thus do not ‘count’ for present purposes as “top-down” effects, since they involve (merely) changing the input to perception rather than changing how visual processing works. They thus seem no more relevant to questions of modularity, encapsulation, and cognitive penetrability than are more pedestrian examples of “changing perception” simply by voluntarily closing your eyes.

Does race influence lightness?

By contrast, the striking demonstration at the heart of Levin and Banaji’s (2006) report suggests something quite different—that the visual system has access to the mind’s rich and nuanced knowledge about race categories, and that this knowledge is brought to bear in computing the lightness of faces. However, we have shown here that the images continue to appear differentially light even when race information is unavailable—an effect that must be due to low-level stimulus differences. And so these same low-level differences might also explain the striking demo from Levin and Banaji (2006). And it is especially striking that the demonstration with blurred faces is compelling—as in Fig. 1b—given that we would expect blurring to make the effect weaker for independent reasons (since blurring also necessarily attenuates low-level cues to lightness). (Indeed, blurring *must* weaken and eventually eliminate the effect in the limit, since with maximal blurring the two stimuli would become identical gray patches.)

Of course, our results do not (and could not) definitely identify the factor responsible for the initial demonstration: it could be that the difference in perceived lightness with blurred faces is due to low-level stimulus features, while the same effect with unblurred faces is due to race categories. But clearly additional evidence would be required to make such a case, and we conclude here that the initial demonstration of Levin and Banaji (2006) provides no evidence for a subjectively appreciable top-down effect on perception. (Similarly, perhaps our results were driven by *unconscious* race categorization. This cannot be ruled out—though it would be an odd sort of implicit race judgment that could influence explicit lightness judgments but not explicit race judgments. Our point

is simply that the lightness illusion at the heart of these results can be readily explained without any appeal to race, and such alternate explanations would have to be ruled out for such phenomena to provide evidence for top-down effects.)

Though our focus here is only on whether there are subjectively appreciable top-down effects, it is important to note that Levin and Banaji (2006) were explicitly sensitive to many of these concerns—and that, in general, their paper was a paragon of methodological and theoretical sophistication. In fact, to address possible low-level explanations, they report other experiments using uniformly gray line-drawings of faces. However, it is telling that this method, in eliminating low-level image differences, also completely eliminated the subjectively appreciable perceptual differences; the Black and White line-drawing faces just do not look differentially light. And though Levin and Banaji (2006) still reported a (much smaller) distortion effect using the line drawings, the lack of an associated “demo” leaves these results open to alternative explanations.⁵ Moreover, Levin and Banaji (2006) noted that there was “no hint of a correlation” (p 504) between attitudes toward the races (e.g., feeling “cool” toward Black or White individuals) and the lightness distortion effect, even though racial attitudes do correlate with degree of focus on a face’s racial qualities (e.g., Fazio & Dunton, 1997; see also MacLin & Malpass, 2001, for a demonstration of how manipulations of racial categorization need not involve subjectively appreciable lightness distortions.) In any case, while further work may continue to explore whether race influences perceived lightness, our conclusion here is simply that any such effect (like nearly all top-down effects) does not involve a subjectively appreciable demonstration that can only be explained in such terms.

Conclusion

The possibility of top-down effects on perception remains an open and exciting prospect that could revolutionize our understanding of how perception works. However, such effects are unlikely to achieve the status of other perceptual phenomena without the subjectively appreciable “demonstrations” so

⁵ For example, without the associated “demonstration”, task demands become an especially acute concern given the explicit instructions telling subjects that the experiments were about “how people perceive the shading of faces of different races” (p 504). An additional reason to suspect a response-bias account of these results is that Levin and Banaji (2006, Experiment 1) found lightness-distortion effects even where they “shouldn’t” have—not only when observers judged the lightness of the faces by matching a grayscale patch to a reference face, but also when matching *copies of the faces themselves* to a reference face. This pattern of results implies an “El Greco fallacy” (Firestone, 2013; Firestone & Scholl, 2014), because the means of reproduction (i.e., the adjustable face) should have been distorted in the same way as the stimulus reproduced (i.e., the reference face), and the effects should have cancelled out. That an effect was still observed suggests a non-perceptual explanation.

common in vision science. Here we have shown that the strongest counterexample we know of may not truly count as a case of *experiencing* a top-down effect, leaving no subjectively appreciable evidence that higher-level factors affect what we see.

Acknowledgments For helpful conversation and/or comments on earlier drafts, we thank Dan Levin, Gary Lupyan and also the members of the Yale Perception and Cognition laboratory. We also thank Dan Levin for sharing his stimuli with us, and Gary Lupyan for suggesting the use of the *Glass* distortion as a means of blurring the images.

References

- Adelson, E. H. (2000). Lightness perception and lightness illusions. In M. Gazzaniga (Ed.), *The new cognitive neurosciences* (2nd ed., pp. 339–351). Cambridge: MIT Press.
- Balcetis, E., & Dunning, D. (2010). Wishful seeing: More desired objects are seen as closer. *Psychological Science*, *21*, 147–152.
- Banerjee, P., Chatterjee, P., & Sinha, J. (2012). Is it light or dark? Recalling moral behavior changes perception of brightness. *Psychological Science*, *23*, 407–409.
- Bhalla, M., & Proffitt, D. R. (1999). Visual-motor recalibration in geographical slant perception. *Journal of Experimental Psychology: Human Perception and Performance*, *25*, 1076–1096.
- Collins, J., & Olson, I. (2014). Knowledge is power: How conceptual knowledge transforms visual cognition. *Psychonomic Bulletin & Review*, *21*, 843–860.
- Dunning, D., & Balcetis, E. (2013). Wishful seeing: How preferences shape visual perception. *Current Directions in Psychological Science*, *22*, 33–37.
- Durgin, F. H., Baird, J. A., Greenburg, M., Russell, R., Shaughnessy, K., & Waymouth, S. (2009). Who is being deceived? The experimental demands of wearing a backpack. *Psychonomic Bulletin & Review*, *16*, 964–969.
- Durgin, F. H., Klein, B., Spiegel, A., Strawser, C. J., & Williams, M. (2012). The social psychology of perception experiments: Hills, backpacks, glucose and the problem of generalizability. *Journal of Experimental Psychology: Human Perception and Performance*, *38*, 1582–1595.
- Fazio, R. H., & Dunton, B. C. (1997). Categorization by race: The impact of automatic and controlled components of racial prejudice. *Journal of Experimental Social Psychology*, *33*, 451–470.
- Firestone, C. (2013). On the origin and status of the “El Greco fallacy”. *Perception*, *42*, 672–674.
- Firestone, C., & Scholl, B. J. (2014). “Top-down” effects where none should be found: The El Greco fallacy in perception research. *Psychological Science*, *25*, 38–46.
- Fodor, J. A. (1983). *The modularity of mind: An essay on faculty psychology*. Cambridge: MIT Press.
- Kawabata, N., Yamagami, K., & Noaki, M. (1978). Visual fixation points and depth perception. *Vision Research*, *18*, 853–854.
- Koffka, K. (1935). *Principles of Gestalt psychology*. New York: Harcourt, Brace.
- Levin, D. T., & Banaji, M. R. (2006). Distortions in the perceived lightness of faces: The role of race categories. *Journal of Experimental Psychology: General*, *135*, 501–512.

- Lupyan, G. (2012). Linguistically modulated perception and cognition: The label-feedback hypothesis. *Frontiers in Cognition*, 3, 54.
- MacLin, O. H., & Malpass, R. S. (2001). Racial categorization of faces: The ambiguous race face effect. *Psychology, Public Policy, and Law*, 7, 98–118.
- Peterson, M. A., & Gibson, B. S. (1991). Directing spatial attention within an object: Altering the functional equivalence of shape descriptions. *Journal of Experimental Psychology: Human Perception and Performance*, 17, 170–182.
- Proffitt, D. R. (2006). Embodied perception and the economy of action. *Perspectives on Psychological Science*, 1, 110–122.
- Pylyshyn, Z. W. (1999). Is vision continuous with cognition? The case for cognitive impenetrability of visual perception. *Behavioral and Brain Sciences*, 22, 341–365.
- Raftopoulos, A., & Zeimbekis, J. (Eds.). (2014). *Cognitive penetrability*. Cambridge: MIT Press (in press).
- Rosenthal, R., & Rubin, D. B. (1978). Interpersonal expectancy effects: The first 345 studies. *Behavioral and Brain Sciences*, 1, 377–386.
- Toppino, T. C. (2003). Reversible-figure perception: Mechanisms of intentional control. *Perception & Psychophysics*, 65, 1285–1295.
- Vetter, P., & Newen, A. (2014). Varieties of cognitive penetration in visual perception. *Consciousness & Cognition*, 27, 62–75.