

Shaking up the mind's ground floor: the cognitive penetration of visual attention*

Wayne Wu

[Forthcoming *The Journal of Philosophy*]

Figures for which copyright permission not granted have been redacted

1. Introduction

Philosophers have returned to the cognitive penetration of vision in recent years in part due to its epistemic implications.¹ If belief contents can penetrate visual experience, perhaps because one sees what one believes—that is, belief contents provide visual content—then vision cannot be a neutral tribunal for belief. It is not clear, however, that we have good evidence that cognition penetrates vision and thus that the epistemic consequences are actual.² Claims about cognitive

*This material was first presented in the workshop, *Thinking, Attending and Sensing*, in March 2014 at the University of Utah, organized by Dustin Stokes. I'm grateful to Dustin and the participants for discussion, especially Marisa Carrasco and Christopher Mole. The material was later presented at the University of Geneva, in December 2014. Thanks to the audience for feedback. Susanna Siegel later provided helpful thoughts on the paper. Zina Ward helped improve readability of the empirical sections. Finally, thanks to a referee for this journal for helpful comments that motivated many improvements.

¹ Susanna Siegel, "Cognitive Penetrability and Perceptual Justification," *Noûs*, XLVI, 2 (2012): 201–22.

² On some attempts to demonstrate cognitive penetration see Fiona Macpherson, "Cognitive Penetration of Colour Experience: Rethinking the Issue in Light of an Indirect Mechanism," *Philosophy and Phenomenological Research*, LXXXIV, 1 (2012): 24–62 and Dustin Stokes, "Perceiving and Desiring: A New Look at the Cognitive Penetrability of Experience," *Philosophical Studies* CLVIII, 3 (2012): 479–92 among others.

penetration in humans are claims about interactions between systems subject to empirical verification. The experimental data philosophers have drawn on are to my mind insufficient to establish cognitive penetration in a way that meets empirical standards. In what follows, I draw on a concrete neural-computational model to argue for cognitive penetration where the epistemic consequences are actual: one's beliefs depend on what one notices or attends to, but attention is biased by cognition. Such cognitive influence raises pervasive challenges for any epistemic agent trying to discover the truth.

Visual attention has long been ruled out as a possible example of the cognitive penetration of vision, but this dismissal relies on a faulty conception of attention as a gate to the visual system that determines its inputs.³ The dismissal is appropriate for *overt* visual attention: movement of the eyes to fixate on targets of interest. Here cognition controls visual experience by controlling what gets into the mind via the eyes, but this is an uninteresting case of cognitive influence. Rather, it is *covert* visual attention, attention independent of eye movement, that is the primary form of visual attention and which drives overt attention. Covert visual attention is cognitively penetrated and indeed, likely subject to disparate forms of informational bias that have substantial epistemic consequences.

Section 2 presents cognitive penetration as an empirical thesis about computation and information, and highlights the central role of computational models in establishing the thesis. A sufficient condition for penetration is provided

³ For criticism of this assumption, see Chris Mole, "Attention and Cognitive Penetration," in John Zembekis and Athanassios Raftopoulos, eds., *The Cognitive Penetrability of Perception* (Oxford: Oxford University Press, 2015), pp. 218-38.

that identifies a standard for experimental cognitive science in testing for cognitive penetration. Section 3 characterizes attention as a psychological phenomenon tied to action while section 4 explores the neural basis of visual attention and presents an explanatorily powerful computational theory of that basis, *divisive normalization*. Section 5 argues that intention cognitively penetrates visual attention: the visual system uses information about the target of intention to inform selective processing of visual stimuli in order to guide task performance. Finally, section 6 discusses how attention might be disparately penetrated and highlights the epistemic consequences due to the biasing of attention. Tying epistemic bias to attention makes available a new set of concepts to understand epistemic agency.

2. Cognitive Penetration

Jerry Fodor characterized informational encapsulation as follows:

Imagine a computational system with a proprietary...database. Imagine that this device operates to map its characteristic inputs onto its characteristic outputs...and that, in the course of doing so, its informational resources are restricted to what its proprietary database contains. That is, the system is “encapsulated” with respect to information that is not in its database.⁴

Similarly, Zenon Pylyshyn wrote:

⁴ Jerry A. Fodor, *The Mind Doesn't Work That Way: The Scope and Limits of Computational Psychology* (Cambridge: MIT Press, 2001), p. 63.

if a system is cognitively penetrable then the function it computes is sensitive, in a semantically coherent way, to the organism's goals and beliefs, that is, it can be altered in a way that bears some logical relation to what the person knows.⁵

I present these ideas to illustrate the computational core of informational penetration. "Informationally encapsulated" and "informationally penetrated" should be understood as two place predicates: *X is informationally encapsulated from/informationally penetrated by Y*. Our primary interest is where X is vision (indeed visual attention), Y cognition. A focal point will be the computations performed by the visual system, so assessing the cognitive penetration of vision requires a computational theory of relevant visual processes. As cognitive penetration is an empirical thesis about influences on computation, to satisfactorily establish it requires specification of a *computational model* that explains what penetration comes to. Without a plausible model, any claim of cognitive penetration is underspecified.

For Fodor and Pylyshyn, "information" means representational content as tied to semantic value, but I use the term in its statistical sense (e.g. Shannon mutual information): information reduces uncertainty about a random variable. Relatedly, information can be tied to correlation as in Dretske's notion of indication or Grice's

⁵ Zenon W. Pylyshyn, "Is Vision Continuous with Cognition?: The Case for Cognitive Impenetrability of Visual Perception," *Behavioral and Brain Sciences* XXII, 3 (1999): 343.

notion of natural meaning, though these ideas are tied to reducing uncertainty as well.⁶ Does this amount to changing the subject? I will argue in the final section that the computational notion of cognitive penetration I discuss has epistemic significance. We need not exactly toe the Fodor/Pylyshyn line to engage with the substantial issues concerning epistemology and cognitive architecture that exercised them.

Nevertheless, in speaking of information, two plausible assumptions allow us to track relevant semantic relations. The first is a weak supervenience claim: that in the case of veridical experience of a target (this can be object, property or location), the neural basis of that experience carries information about the target. Thus, any change in experience implies a change in informational (neural) content. This correlation is guaranteed by the causal relation between the target and the relevant parts of the visual system that give rise to the experience, and it is weak in that it allows that when one is under an illusion or hallucination of some target, the experience is still about it even if, by definition, there can be no information about what is not there.

This correlation is plausible given a second assumption, namely that one focuses on information processing in relevant parts of the visual *ventral stream*. The ventral stream plays a necessary role in visual experience in humans. In brief, damage to ventral stream areas can give rise to visual agnosia, including the inability to see objects. One area in humans seems crucial to normal experience of

⁶ Fred Dretske, *Explaining Behavior: Reasons in a World of Causes* (Cambridge, MA: MIT Press, 1991). Paul Grice, "Meaning," in *Studies in the Way of Words* (Cambridge, MA: Harvard University Press, 1987), pp. 213–23.

objects, namely the lateral occipital complex (LOC), and the monkey (macaque) homolog of LOC is thought to be located in the monkey inferotemporal cortex (IT).⁷ Macaque monkey vision is widely studied, and provides a model for human vision. Other regions of the ventral stream play important roles in supporting visual experience such as area V4 which plays roles in figure/ground segmentation, color constancy, and feature representation.⁸ Similarly, the middle temporal area MT, which lies in the dorsal stream, is important for motion processing and damage to it can lead to akinetopsia (the inability to see motion).⁹ As we shall see, V4, MT, and IT exhibit the attentional modifications that will be critical to our discussion: receptive field remapping (shrinking). Given these two assumptions, computation over information about a visual target in relevant areas of the ventral stream will track visual experience of it. In the case we shall consider, computation of information about a target will track the experience of visual *attention* to it. Specifically, one begins with seeing two objects and comes to attentionally select one in order to deal effectively with it.

Visual computations map certain visual inputs onto certain visual outputs. David Marr argued that vision maps a two-dimensional retinal image input onto a representation of the three-dimensional world. This functional characterization

⁷ On LOC and patient DF, see A. David Milner and Melvyn A. Goodale, *The Visual Brain in Action*, 2nd ed. (Oxford: Oxford University Press, 2006). On IT as the monkey homolog of LOC, see Doris Y. Tsao et al., "Faces and Objects in Macaque Cerebral Cortex," *Nature Neuroscience* VI, 9 (2003): 989–95.

⁸ While V4 was for a long time taken to be a color area, its function is very complex. See Anna W. Roe et al., "Toward a Unified Theory of Visual Area V4," *Neuron* LXXIV, 1 (2012): 12–29.

⁹ Mark Nawrot, "Disorders of Motion and Depth," *Neurologic Clinics* XXI, 3 (2003): 609–29.

provides Marr's *computational theory* of vision.¹⁰ But "computation" is often used in a mechanistic sense, referring to how the mapping is achieved by underlying algorithms including the representations or contents over which the algorithmic procedures operate as well as the neural circuits that realize the computation. One sort of computational model gives a mathematical description (set of equations) that describes the phenomenon along with a plausible implementation of the computation in a neural circuit. I later consider such a model for attention.

Here then is a sufficient condition for the cognitive penetration (SCP) of vision by cognitive systems where visual computations operate over visual input I to generate output O among possible outputs O_n in light of cognitive content/information R :

(SCP) If cognitive systems contain information R such that vision computes over R where this computation explains why the visual system yields O rather than some other output O_n , given visual input I , then cognition cognitively penetrates vision.

That is, information from cognition is used by vision to produce a specific output. If the condition is met, then cognition informationally penetrates visual computations by providing information that vision uses. Since claims of cognitive penetration are empirical theses, demonstrating that the antecedent is satisfied requires the following:

¹⁰ David Marr, *Vision* (San Francisco: W. H. Freeman and Company, 1982).

1. Specification of the nature of the visual computation,
2. Specification of the way in which visual computations depends on cognitive contents.
3. Specification of relevant experimental evidence in favor of the model as fleshed out by (1) and (2).

The philosophical clarification of the concept of cognitive penetration sets a task for cognitive science in (1)-(3). I shall focus on (1) and (2) which requires a computational model; (3) is a matter of ongoing empirical work. In the next section, I introduce a functional conception of attention that has its root in experimental practice in cognitive science. This identifies the relevant aspect of vision that will exhibit cognitive penetration: visual attention. The computational model for visual attention will be described in section 4.

3. Attention and selection for task

William James defined attention as follows:

Everyone knows what attention is. It is the taking possession by the mind, in clear and vivid form, of one out of what seem several simultaneously possible objects or trains of thought. Focalization, concentration, of consciousness are of

its essence. It implies withdrawal from some things in order to deal effectively with others.¹¹

Yet many are skeptical about defining what attention is, to the point that one theorist, in a paper entitled, “There is no such thing as attention”, has argued that all attempts to do so have failed.¹² These skeptics have overlooked a constant feature of attention that is sewn into the fabric of experimental practice. Most cognitive scientists speak of attention as *selection for processing*, but this does not even provide a sufficient condition since dumb manufacturing machines select items for further processing. Appending “information” to talk of processing does not help since the different cones in the retina can be understood to select information for further processing (due to their selective absorption spectra), but one would not speak of such selection as attention. What then is distinctive of attentional selection?

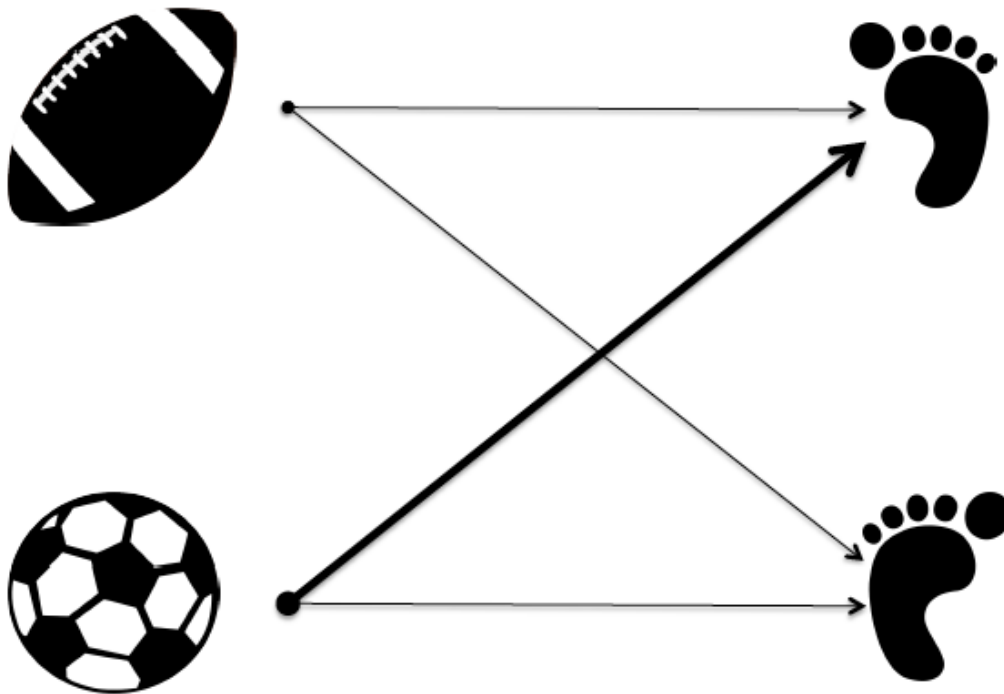
On my view, attention is *selection for action*.¹³ “Action” is taken in the broadest sense to include mental as well as bodily action. Let me illustrate a subject’s selection of a target for action in terms of the *Many-Many Problem*: the idea that in many action contexts, an agent confronts more targets than can be acted on at once and for each, there are different responses. In Figure 1, the agent sees an American football and a soccer ball and can kick each one either with the left or right foot.

¹¹ William James, *The Principles of Psychology, Volume 1* (Boston, MA: Henry Holt and Co., 1890).

¹² Britt Anderson, “There Is No Such Thing as Attention,” *Frontiers in Theoretical and Philosophical Psychology II* (2011): 246. There is much I agree with in Anderson’s critique.

¹³ Wayne Wu, *Attention* (Abingdon: Routledge, 2014).

Insert Figure 1 Here



The actions available at a time are given in this *behavior space* where the arrows define four possible actions. This space is a *psychological space* in that the inputs are visually experienced, so the arrows connect visual experience to motor output. Let us say that the agent now decides (intends) to and kicks the soccer ball with the right foot. The darker arrow in the figure represents the actual action and, by implication, the computational processes linking visual experience to motor output. We can theorize that the causal role of intention in generating action is to influence the selection of the appropriate behavioral path, precisely the path that the agent

intends to carry out: one kicks the ball with the right foot because one intends to do just that. Intentions facilitate the action by representing the input-output mapping identified by the arrow representing the action.¹⁴ This solves the Many-Many Problem by biasing a solution that amounts to what the agent does, in the face of many behavioral options.

Attention is found in the “for action” of the visual experience of the soccer ball. While the agent sees both balls, it is only the experience of the soccer ball that guides the action. So, action-relevant processing selects that information and ignores the basketball to generate a specific action. To pump intuitions, identify two objects before you and then fix your eyes on a point between them. For practice, shift attention between them without moving your eyes (many find this difficult). Now, shift attention to the point of fixation and, while holding the eyes fixed, reach for one of the objects. Attention must shift back to that object to guide reach even when the eye does not move. Similarly, while maintaining fixation, move the eyes to one of the objects. The science of eye movement suggests that covert attention—attention independent of eye movement—programs overt attention, the eye movement.¹⁵ This example of overt attention (eye movement in a task) presents the core of an experimental example to be discussed below (Yarbus). The current point is that solving the Many-Many Problem by acting requires selection of some relevant target in order to guide action. As James noted, one must attend to the item and select it from among others in order to deal effectively with it.

¹⁴ For more on these issues, see Wu, *op. cit.* chapter 3.

¹⁵ See Katherine Armstrong, “Covert Spatial Attention and Saccade Planning,” in Christopher Mole, Declan Smithies, and Wayne Wu eds., *Attention: Philosophical and Psychological Essays* (Oxford: Oxford University Press, 2011), 78–96.

In what follows, I rely only on a consequence of the selection for action view: when a subject selects a target in order to guide action in an experimental task, the subject attends to that target. This is an *empirical sufficient condition* for attention that links selection for task performance to attention. Its use in our discussion can be independently motivated, distinct from the selection for action view of attention, as it is a condition assumed in the main experimental paradigms used to study attention: visual search, dichotic listening, multiple object tracking, and spatial cueing among others. Here is the experimental logic. To study attention, an experimenter must control how the subject attends. The experimenter designs a task whereby attention to a target is controlled by making the target relevant to performing the task. That is, to adequately execute the task, subjects must select information from the target to inform their response. For example, in dichotic listening, one is presented with two verbal streams to each ear and asked to shadow (i.e. verbally repeat) just one stream, say the stream presented to the left ear. One must select just those words to inform repeating them. The assumption is that when one does so, one is attending to that stream.

Consider next a *delayed match to sample* task used by Robert Desimone and co-workers to study attention effects on visual processing in the macaque monkey visual system. As we will look carefully at visual neural response during performance of this task, it will be important to understand it. The task requires that a subject remember an initially presented “sample” that must later be matched (or not) to one of two test stimuli presented after a delay period. If there is a match (i.e. one of the test pair matches the earlier sample), the subject identifies the match by

moving the eye to it. Electrical activity from visual neurons is recorded while the subject (awake behaving monkeys) performs the task. Specifically, Desimone et al. recorded electrical activity from cells in area V4 and IT when the two test stimuli were placed in the cells' receptive fields. The receptive field can be understood as that region of external space such that placing an appropriate stimulus within that space causes the cell to generate action potentials, i.e. spikes of electrical discharges. The dotted circles in Figure 2 identify the location of the receptive field. To manipulate attention, the experimenters drew on the empirical sufficient condition for attention: "The task used to focus the animal's attention on a particular location was a modified version of a "match-to-sample" task."¹⁶ Thus, before moving their eyes, the animals must attend to the test stimuli within the neuron's receptive field so as to determine whether a match is present.

Insert Figure 2 Here

Two stimuli are presented in succession while the animal maintains fixation (FP, the fixation point): first a sample (the cue; second panel) and then, after a delay of 1.5 seconds, a test array of two objects, one of which can be a match of the cue. If there is a match (top two rows), the animal reports this by making an eye movement (saccade) to the match from the fixation point. The animal thus comes to fixate on the match. Note that there are two "Target (Match)-Present" conditions, one (top row) involving a target that is the *preferred (good)* stimulus for the neuron whose activity is recorded (good stimulus because it drives the neuron to generate the

¹⁶ Leonardo Chelazzi et al., "Responses of Neurons in Macaque Area V4 during Memory-Guided Visual Search," *Cerebral Cortex*, XI, 8 (August 2001): 761–72.

largest number of spikes per unit time, i.e. high firing rate) and the other (middle row) involving a target that is not a *preferred* stimulus (i.e. is *bad*) which drives a lower firing rate. In “Target-Absent” conditions (bottom row), the animal reports no match by maintaining fixation (not moving the eye).


Two salient deployments of attention occur: (1) attention to select the cue to store in working memory during the delay period and (2) attention to program the eye movement because the animal must attend to the relevant target to prepare the saccade. Covert attention is found in the animal’s covertly selecting the match so as to move the eye to it. These are just cases of selection of a target to perform a task which neuroscientists and psychologists take to be a form of attention (the empirical sufficient condition for attention). Note that recordings to be discussed below concern the second deployment of attention to prepare the saccade.

Because selection for task accords with the animal’s intentions to perform the task, the previous case exemplifies what I call *top-down, goal-directed* attention, specifically intention-influenced attention:

S top-down attends to X *in an intention-oriented manner* if and only if S attends to X as a result of an intention to Φ in respect of X.

“ Φ ” can stand for *attend to X* in which case one intentionally attends to X, but the intention can be to perform an action directed to X, say to saccade to it. The influence of intention on attention can be subtle yet pervasive, as seen in Alfred Yarbus’s study of goal-directed eye movements.

Insert Figure 3¹⁷



Yarbus presented human subjects with a painting and asked them to perform various tasks including to remember what the individuals were wearing; to remember what objects were in the room; and to estimate how long the father was away (panels b-d respectively). In measuring eye movements, Yarbus tracked overt attention, and by implication, the shifts in covert attention that program those movements. Interestingly, the eye movements clearly tracked the subjects' intentions to carry out the tasks. Each pattern seems to make sense, given the objects that must be examined in order to perform each task. Given that eye

¹⁷ Figure from M. F. Land, "Eye Movements and the Control of Actions in Everyday Life," *Progress in Retinal and Eye Research* 25 (2006): 296–324.

movements are programmed by covert attention, the data show that covert attention is sensitive to the subject's intention even if the subject does not intend to attend in the patterns observed. Attention shifts in a way consistent with the subject's broader intention to respond to Yarbus' commands. So, when subjects perform tasks, they do so with specific intentions and this seems to set the deployment of attention to targets in a task-relevant way. We can treat this as an empirical hypothesis to explain the eye tracking data, a hypothesis we will flesh out below by deploying a computational model of attentional modulation in the visual system.


Returning to the experiment of Desimone et al., I make two assumptions about intentions, granting that monkeys can have intentions. This attribution of intention has always been part of the neuroscientist's description of monkey behavior (if necessary, one can shift to action-directed representations as the monkey functional homolog of intentions). The assumptions are: (a) that the monkey's behavior is explained by an intention to act on a target as instructed, including withholding eye movement; (b) that the neural mechanisms observed in the monkey's visual system are also found in humans (monkeys vision is a model for human vision). Thus, the monkeys studied by Desimone et al. exemplify intention-directed attention, a selection of a specific item to guide task performance (moving the eye to report a match). But what is the visual system doing during such action? In the next section, I describe the recorded neural activity and a computational mechanism that explains it.

4. Divisive Normalization

James emphasized that attention withdraws from some things to deal effectively with others. This is evocative of the demands of the Many-Many Problem: multiple objects yield multiple behavioral possibilities, but action often requires selection of a subset of those possibilities. The empirical sufficient condition for attention points out that such selection for task yields attention. Moreover, selection for task performance is reflected in neural activity given the subject's response to the Many-Many Problem: the brain shows selectivity that serves action and realizes the subject's attending to task-relevant targets.

Consider a common signature of attention in the activity of single neurons: *gain modulation*. This modulation is essentially the amplification of a signal carrying information. Figure 4 shows a simulation of the response of a single visual neuron to a set of motion stimuli, each moving in a different direction.

Insert Figure 4 Here [use PDF file]



The lower dotted horizontal line identifies the *baseline* or *spontaneous* activity of the neuron, namely its firing rate when no stimulus is presented in its receptive field. The lowest bell curve is plotted on the basis of presenting the neuron with a each member of the set of moving stimuli, with motion in a particular direction, and measuring the response to that stimulus. For the lower curve, using the empirical sufficient condition for attention, the experimenter directs the animal's attention outside of the receptive field (e.g. to a different task relevant object). Thus, the motion stimulus in the receptive field is not attended. The graph shows that the good or preferred stimulus for this neuron is where the direction of motion is at 0 degrees, in this case, vertical upward motion. The response drops off to near baseline when the motion is downward (180°). If the experimenter now makes the motion stimulus in the receptive field task relevant so that the animal must attend

to it, measured response to each stimulus increases (the modulation of neural gain). This yields the top curve which shows that the signal is amplified during attention.¹⁸

Gain modulation, the difference between the two bell curves, is correlated with attention to the stimulus in question. It is as if the activity at each point in the “unattended” bottom curve is multiplied by a constant factor when attention is engaged, generating the corresponding point on the top “attended” curve (hence, this is called *multiplicative gain*). It is important to emphasize, however, that the mere presence of gain modulation does not entail attention. It is not a sufficient condition for it, for one could increase neural response by appropriate increases in the intensity (luminance) of each stimulus, generating the two curves. In that case, the gain modulation would not be due to attention but to stimulus luminance. So, the difference between the two curves in Figure 4 is not uniquely associated with attention. It is the empirical sufficient condition for attention that secures the relevance of attention when neural signal amplification occurs. That is, the experimenters concluded that the observed gain modulation is an attention effect because, in addition to the modulation, they deployed that empirical sufficient condition to direct the animal’s attention to within the neuron’s receptive field.

I shall focus on the visual computation postulated to underlie gain modulation and which explains another form of attentional modulation, namely *receptive field remapping*. Receptive field remapping concerns cases where two stimuli are present in the neuron’s receptive field. Thus far, we have considered only

¹⁸ For actual data demonstrating this, see Carrie J. McAdams and John H. R. Maunsell, “Effects of Attention on the Reliability of Individual Neurons in Monkey Visual Cortex,” *Neuron*, XXIII, 4 (1999): 765–73.

single stimuli, but it would be natural to predict that a neuron's response to two stimuli will be the sum of the neuron's response to each individual stimulus. The actual response, however, is often the weighted average of the two individual responses. Thus, for a number of neurons, it is division not addition that best describes neural activity in response to multiple stimuli. The response of the neuron is divided and not summed.

Where attention is directed to one of the two objects, the neural response over time comes to be driven by the attended stimuli, as seen from neural data recorded during the task described in Figure 2:

Figure 5¹⁹

Here, we are looking at the time course of neural response from a V4 neuron. Time 0 indicates the presentation of the two-object test array after the delay period where the animal must move the eye to the match (recall: right boxes, Figure 2). The top and bottom open circle curves show the plot of neural response over time for the good and poor stimuli to which the neuron responds, each placed by itself in the receptive field (the good stimulus gives the strongest (higher) response). The middle two curves (dark and dashed lines) show the neural response when both

¹⁹ From Chelazzi et al. op. cit.

objects are present in the receptive field where the animal attends to only one of the two stimuli. Initially, one sees what is essentially the weighted average of the response (the initial response to two stimuli is in between the peaks of response to each individual stimulus at 100 ms). This means that in V4, both objects are initially registered by the neuron. As time goes on, the middle curves split, with one more closely tracking the bottom curve (dashed line; here, attention to the poor stimulus) and the other more closely tracking, indeed overlapping, the top curve (solid line; attention to the good stimulus). Since neural response comes to be as if driven by only one of the two objects in its receptive field (“ignoring one to deal effectively with the other”), one can speak of the receptive field as shrinking (remapped) around the attended stimulus.²⁰ Note that the remapping of the receptive field occurs before the action (saccade) to the attended stimulus (time of saccade indicated by the black bar at about 250 ms), leaving open the possibility that it might serve programming the target location for the saccade.²¹

We can think of the resolution of neural activity in Figure 5 in terms of *biased competition*. There is a resource that stimuli compete for, namely a neuron’s capacity to generate spikes (generating spikes requires energy, a limited resource). Each stimulus aims to garner the neuron’s spikes, i.e. to be signaled to downstream processes. The neural signature of competition is the weighted average response seen with multiple stimuli. Bias is exemplified in the shift of neural activity that yields receptive field shrinking/remapping. Attention is said to “emerge” from such

²⁰ Jeffrey Moran and Robert Desimone, “Selective Attention Gates Visual Processing in the Extrastriate Cortex,” *Science*, CCXXIX, 4715 (August 23, 1985): 782–84.

²¹ See Milner and Goodale *op. cit.* for an account of how attentional modulations in the ventral stream might serve action computations in the dorsal stream.

biased competition.²² The effects of biased competition are found throughout the visual system such as in inferotemporal (IT) cortex within which neural activity comes to correlate with (“represent”) objects.²³ Figure 5 gives a snapshot of shifts in neural activity that occur throughout the visual system when a target is attended to. The idea then is that the resulting subject-level visual state, one’s attending to an experienced object, is realized by selective neural modulation within the ventral stream where specific stimuli seem to be neurally selected. Psychological selectivity in attention has a basis in neural selective modulation in the brain, all tied to the task being performed.

We can think of neural activation in terms of neural computation. Neural computation is given in terms of mathematical/statistical models of the neuron’s behavior. In remapping, the activity of a neuron is a function of its interaction with other neurons, the links between them constituting a circuit that implements division or divisive normalization. Divisive normalization has been called a *canonical* computation in that it is deployed in a variety of brain regions and for a variety of phenomena.²⁴ It is not restricted to attention, but in attention relevant circuits, it models many of the observed attentional modulations in neural activity.²⁵

²² Robert Desimone and John Duncan, “Neural Mechanisms of Selective Visual Attention,” *Annual Review of Neuroscience* XVIII (1995): 193–222.

²³ Leonardo Chelazzi et al., “Responses of Neurons in Inferior Temporal Cortex during Memory-Guided Visual Search,” *Journal of Neurophysiology* LXXX, 6 (December 1998): 2918–40.

²⁴ Matteo Carandini and David J. Heeger, “Normalization as a Canonical Neural Computation,” *Nature Reviews Neuroscience*, XIII, no. 1 (2012): 51–62.

²⁵ Geoffrey M Boynton, “A Framework for Describing the Effects of Attention on Visual Responses,” *Vision Research* XLIX, 10 (2009): 1129–43; Joonyeol Lee and John H. R. Maunsell, “A Normalization Model of Attentional Modulation of Single Unit

That is to say, the equations accurately describe neural activity under attention.

While models of divisive normalization differ in mathematical details, all emphasize division. I shall begin with the mathematical model presented by Lee and

Maunsell.²⁶ The following equation describes the neuron's response for two stimuli:

$$R_{1,2} = \left[\frac{N_1 \cdot (I_1)^u + N_2 \cdot (I_2)^u}{N_1 + N_2} \right]^{1/u}$$

Equation 1


The equation models a neuron's response $R_{1,2}$ to two stimuli, O_1 and O_2 , in the neuron's receptive field. This neuron, M , responds to O_1 and O_2 in a way sensitive to two factors: N , the normalization term, and I , the input term. These terms are due to populations of neurons that constitute M 's *normalization pool*. Metaphorically, think about the normalization pool as the neuron's peer group to which it listens for advice about the stimuli. At the neural level, normalization pools provide information to the neuron about the input in question, and can effectively suppress M 's response. M 's response is thus determined by the input in its receptive field and by the activity of other neurons in its normalization pool. A schematic circuit is given in Figure 6 for a neuron in visual area MT that is sensitive to moving stimuli.

This circuit realizes the computational mechanism described in Equation 1:

Responses," *PLoS ONE*, IV, 2 (2009): e4651; John H. Reynolds and David J. Heeger, "The Normalization Model of Attention," *Neuron*, LXI, 2 (January 29, 2009): 168–85.
²⁶ Lee and Maunsell, *op. cit.* There are important differences although the basic idea of division remains. I shall largely ignore these. Where the models agree is that there is a route for attention to affect the computation. The location of the entry point is one difference.

Figure 6²⁷

Insert here



In Lee and Maunsell’s model, the normalization term N takes on values between 0 and 1. Attention influences the magnitude of N . As they note: “The effect of attention is introduced by letting attention modulate the normalization associated with the attended stimulus...In this way, attention acts only through the normalization mechanism.”²⁸ This influence of attention on N is captured by the following equation:

Equation 2 $N_{\text{attended}} = (1-s)(1 - e^{-\beta\alpha C}) + s$

²⁷ From Lee and Maunsell, *op. cit.*

²⁸ Lee and Maunsell, *op. cit.* p. 4651

Just note that β is the contribution of attention to N where β can take values from 1 to infinity: 1 where no attention is directed at the stimulus and greater than 1 otherwise. As attention increases β , the value of N gets closer to 1. Looking at equation 1, a prediction is that where attention to O_1 (say) is greater with O_2 being ignored, the response $R_{1,2}$ will be driven more by I_1 than by I_2 as observed in Figure 5. Where attention is directed to one of two stimuli within a neuron's receptive field, the response of the neuron is determined by the attended stimulus. If we think of attention as a form of selection, then where only one object is attended to, divisive normalization takes two inputs and returns only one as output. Attention then serves as a gate within the visual system, filtering one object for further processing, namely that object on which the subject intends to act.²⁹ Given that we are speaking of areas that are implicated in visual experience, including many areas in the ventral stream, the idea is that divisive normalization in these areas implements the subject's visually attending to a visually experienced object. Divisive normalization currently provides the most unified computational account of the neural basis of visual attention. This discharges item (1) in our sufficient condition for cognitive

²⁹ Lee and Maunsell, *op. cit.*, write: "We assume that a neuron with a receptive field containing two stimuli receives a direct, tuned input with a strength that depends on how well the stimulus matches the preferred stimulus for the cell" (4651). On their account, where the response is maximum, the neuron conveys the information that the stimulus in the receptive field is maximally similar to its preferred stimulus; where the response is minimal, the neuron conveys the information that the stimulus is maximally dissimilar to its preferred stimulus (the metric of similarity here is left open); where the response is in between these values, the neuron carries information about "mixed" similarity, a case that involves more uncertainty (hence less information) about what is in the environment. There are alternative accounts of how to understand the neuron's function

penetration: providing a model of visual (attentional) computation. How then does intention enter into the story?

5. Intention cognitively penetrates vision

Our empirical hypothesis, based on the Yarbus results and the Many-Many Problem, is that intention sets attention. What one attends to depends on the task at hand which in turn depends on one's intention to perform that task. If divisive normalization is the neural computational basis of subject-level attention to a task-relevant object, then the neural realization of top-down, intention-directed attention should also be sensitive to intention. It is this point that suggests cognitive penetration: for the monkeys, their intention to act on a specific object shifts the visual computation of divisive normalization from a weighted average reflecting two stimuli to a response driven by only one stimulus, namely the task relevant stimulus. In other words, think of the visual system as keeping an eye out for information about the animal's current goals so as to shift its processing to better serve those goals.

But what is the content of the intention? In the experiment of Figure 2, the monkey is trained to perform the match-to-sample task and during the experiment, intends to perform the task. One option: when presented with a stimulus, the monkey intends to act on *it*. As Figure 2 demonstrates, the monkey must retain the sample in working memory due to the delay period when no stimulus is presented. Once the target array is presented, the monkey must determine which of the two

objects matches the sample and move the eye to it. The monkey's intention then might develop in this way: initially intending to match two presented test stimuli (general intention), it comes to intend to move the eye to *that* target (namely the one in the test array that matches the sample; recall, on some trials the monkey must not move the eye). As with humans, such intentions can rapidly develop, sensitive to task context.

It is natural to think of the intention at issue as demonstrative. Thus, the intention might be to move the eye to *that* target, a perceptual demonstrative to the match. For example, locate a medium sized object before you. Got it? Now, grab it. I submit that execution of that task requires a demonstrative intention to grab the relevant object attended. Such perception-based demonstratives require attention to fix reference,³⁰ but now a puzzle arises: if intention shapes attention to the match target, as seen in the activity of the V4 neuron in Figure 5, how can intention also be demonstrative since that requires prior attention to the match? That is, we need attention to a target to intend to act on that target, but that intention seems to be needed for (top-down) attention to the target. Is this not a problematic circle?

The solution is to recognize two stages of attention: attention for fixing demonstrative thought and attention for movement. If these stages are distinctly realized, then the problematic circularity is eliminated. Recent work by Mante et

³⁰ John Campbell, *Reference and Consciousness* (Oxford: Oxford University Press, 2002); Imogen Dickie, "Visual Attention Fixes Demonstrative Reference by Eliminating Referential Luck," in Christopher Mole, Declan Smithies, and Wayne Wu eds., *Attention: Philosophical and Psychological Essays*, (New York: Oxford University Press, 2011), 292–322; Declan Smithies, "What Is the Role of Consciousness in Demonstrative Thought?," *Journal of Philosophy* XVIII, 1 (2011): 5–34.

al.³¹ provides evidence that there are attentional mechanisms that are *post-*perceptual and not reliant on resolving competition within the visual system. Mante et al. observed selectivity in prefrontal cortex, a higher-order nonsensory area linked to cognition and planning, where this selectivity did not depend on prior selectivity in sensory cortex, such as the modulations depicted in Figure 5, yet still reflected selectivity needed to serve the task. They write: “Our results show that the modulation of sensory responses is not necessary to select among sensory inputs... Multiple selection mechanisms may exist within the brain” (83).

Given that there is evidence for different mechanisms of selectivity, some of which are realized in prefrontal regions that do not require sensory modulation, and some which are realized by sensory modulation as in Figure 5, the puzzle can be resolved. The attentional selectivity needed to ground a demonstrative intention might rely on the selectivity mechanisms observed in prefrontal areas tied to planning where these mechanisms do not rely on modulations of sensory processing but are post-perceptual. Once the demonstrative intention to an object is formed, this intention can then feedback, via the divisive normalization mechanism, to influence sensory processing, as indicated in Figure 5, which then serves eye movement. Thus, the two instances of attention in the task, attention to fix intention and attention as a result of intention, can be realized in different ways. Two distinct deployments of attention are at issue, and no problematic circularity arises.

In a different formulation of divisive normalization, Heeger and Reynolds posit an *attention field* as the source of a signal that shifts divisive normalization as

³¹ Valerio Mante et al., “Context-Dependent Computation by Recurrent Dynamics in Prefrontal Cortex,” *Nature D*, 7474 (November 7, 2013): 78–84.

the β term does on the Lee and Maunsell model. In both models, a place is made for the influence of attention on visual computation. There is good reason, however, to reconceptualize the attention field (also the β term) as an *intention* field. That is, in both models, it is intention, not attention, that directs shifts in neural computation. After all, the proposed models are models *of attention*, meant to elucidate the machinery that realizes the subject's attention state or at least modulations of relevance to attention. Attention is the explanandum to be explained as effect of shifts in neural computation, and thus not to be deployed as the cause of the neural activity. Indeed, this idea is a central claim of the biased-competition model of attention, and divisive normalization is an implementation of biased competition (the bias is captured in the attention field or β ; the competition is the division by the respective normalization pools). Thus, Desimone and Duncan noted that:

The approach we take differs from standard view of attention, in which attention functions as a mental spotlight enhancing the processing (and perhaps binding together the features) of the illuminated item. Instead, the model we develop is that attention is an emergent property of many neural mechanisms working to resolve competition for visual processing and control of behavior.³²

³² Desimone and Duncan, *op. cit.* Similar point can be made about the premotor theory of attention where attention is the effect of selective motor preparation of eye movements to a spatial location and not the cause of such spatial selectivity. In my conversation with several prominent endorsers of the divisive normalization model, they have expressed willingness to eliminate talk of attention as a cause at the relevant stages in the model.

If attention is meant to be the output of the divisive normalization computation we have been considering, then we should not see it as one of the inputs shifting the underlying computation (this would be a problematic explanatory circularity). Rather, given that we are discussing *top-down* intention-driven forms of attention, a natural source of the bias is just the intention that shapes the selectivity required for performing the task. In other words, intention provides the relevant signal to shift visual processing in a task-relevant way. It provides the input for the β term in the Maunsell and Lee model or the attention field in the Reynolds and Heeger model.

Granting this, in what sense does the input of the *intention field* carry information (in the statistical sense) about intention? The input resolves uncertainty about the target of intention, namely which object within the visual field is task-relevant, the intended target of action. From the visual neuron's perspective, it is receiving information about what the target of action is as represented by the subject's intention. On receiving that information, the neuron's response shifts to serve that intention. The content of intention is provided to the divisive normalization mechanism precisely to induce the task-relevant shift in processing, the selection of the intended target. Neural filtering (receptive field remapping) can serve the intended task by shifting visual neuronal response so that it carries information from just the task-relevant input. If visual processing is thus sensitive to intention, the response of the neuron should be appropriately modulated so that it filters the object of intention via remapping of the receptive field. This sort of remapping might occur throughout the visual hierarchy as some evidence currently

suggests.³³ Thus, to return to equation 1, if I_1 corresponds to the object that is the target of intention, then the magnitude of N_1 as against N_2 should be appropriately increased so that I_1 comes to dominate in contributing to the neural response $R_{1,2}$. This shift will occur because of the contribution of β to N_1 where β carries information about the subject's current intentions. Such an interpretation of the circuit and computation makes biological sense: for visual processing to be task-appropriate, it needs information regarding the intended task and its targets. This discharges (2) in our sufficient condition for cognitive penetration as I have argued how intention (cognition) might directly affect visual neural computation for attention. We have, in passing, discussed (3) some empirical evidence for the divisive normalization model though there is more work to be done.³⁴

Recall then that sufficient condition for cognitive penetration:

(SCP) If cognitive systems contain information R such that vision computes over R where this computation explains why the visual system yields O rather than some other output O_n , given visual input I, then cognition cognitively penetrates vision.

³³ For a review of the evidence, see Sabine Kastner and Leslie. G. Ungerleider, "The Neural Basis of Biased Competition in Human Visual Cortex," *Neuropsychologia*, XXXIX, 12 (2001): 1263–76.

³⁴ For recent experimental support, see Katrin Herrmann et al., "When Size Matters: Attention Affects Performance by Contrast or Response Gain," *Nature Neuroscience*, XIII, 12 (December 2010): 1554–59, doi:10.1038/nn.2669.

Intention is an action-directed cognitive state that often represents an object as to be acted on. As we saw in the Yarbus case, a subject's intention affects what objects the subject attends to. It determines what objects are task relevant. Further, in the case of intentions that represent specific targets as in the eye movement actions we have discussed, information about the intended target is transmitted to the visual system in its computation of divisive normalization such that the visual output depends on what target is intended. If this computation is both central to attentional effects in the visual system and if divisive normalization is a canonical neural computation, then in many task-relevant stages in visual processing, computations will be sensitive to intention. In our case, divisive normalization establishes the state of the subject's visual attention, namely the subject's visual selection of one of several seen objects to perform a task. It does so by shifting the neural responses from responding to all the objects within the receptive field to just the object on which the subject intends to act. Across the visual system, the net effect will be that of all the objects within the subject's field of view, a certain experienced object will be attended to. In this way, intention affects visual computations as per SCP and thus, intention cognitively penetrates vision. This case strikes me as currently the best case for cognitive penetration, for it relies on a well-described computational model of the response of neurons under conditions of attention that accounts for the known data. We have then both behavioral data suggesting penetration (e.g. Yarbus and the monkey data discussed), a specific computational model that explicates what penetration might amount to, and increasing empirical support for that model. But why might cognitive penetration of attention matter?

6. Epistemic Consequences: an agenda for a psychological epistemology

Cognitive penetration is philosophically significant in part because it is epistemically significant. One significant consequence arises if cognitive penetration amounts to seeing what one believes, roughly, where the content of belief penetrates visual experience so as to provide contents for experience. This sort of penetration would be epistemically problematic, raising worries about vision as a neutral tribunal for belief. This is not the sort of cognitive penetration exemplified by top-down attention, yet top-down attention that is biased by intention is epistemically consequential in ways that I shall now describe. It is uncontroversial that bias, in a colloquial sense, is epistemically consequential affecting what beliefs we form and what evidence we consider. The epistemic significance of the cognitive penetration of visual attention is precisely that attention is a target of such bias. Put another way, the colloquial biases we all have can influence the neural biases that set attention. The top-down biasing in divisive normalization is an important source of the epistemic effects of bias in the more general sense. Once we recognize how attention is neurally biased, the apparatus used to understand it can illuminate epistemic bias. I will focus on theoretical reasoning but the same points apply to practical reasoning.³⁵

³⁵ There is quite a bit of discussion on these issues in philosophy. For one recent publication that focuses on similar issues raised here and for many references therein, see Susanna Siegel, "Can Selection Effects on Experience Influence Its Rational Role?," in Tamar Gendler ed., *Oxford Studies in Epistemology Volume 4*, (Oxford: Oxford University Press, 2013). Siegel discusses work by Keith Payne

Let us now consider two contrasts: top-down versus bottom-up attention and automatic versus controlled attention. We can expand the definition of top-down attention as follows:

S top-down perceptually attends to X if and only if S's perceptual attention to X is the result of a non-perceptual mental state.

Bottom-up perceptual attention can be defined as the absence of top-down influences as just specified: attention as occurring independently of a non-perceptual mental state such as when attention is driven by a suddenly appearing stimulus (stimulus driven attention or capture of attention). The definition relies on a hierarchy that takes perception as at the mind's ground floor with everything non-perceptual above perception: cognition, motor systems, emotions, moods, desires, intentions etc. Top-down attention then captures a broad set of top-down influences (biases) on visual (perceptual) computations relevant to attention. An empirical possibility to be explored is that the intention field or β term is a more general receptor of biasing signals from all non-perceptual states. Here, intuitions will hopefully suffice: other states like emotions are well known to direct attention, say one's fear of a snake in one's path that pulls attention to it so as to help one avoid it,

where implicit stereotypes affect perceptual judgments (B Keith Payne, "Prejudice and Perception: The Role of Automatic and Controlled Processes in Misperceiving a Weapon," *Journal of Personality and Social Psychology* LXXXI, no. 2 (August 2001): 181–92.). One explanation of the effect Siegel dubs *anti-selection of experience for uptake*, namely that relevant experience is biased against (as James would say, one withdraws from it) leading to an inappropriate response. Siegel's talk of anti-selection is suggestive of attention.

or one's desire that drives attention to the desired item, say when one is hungry and focuses on plates of food at a buffet. Given how these states shift attention, it is plausible to hypothesize that the emotional and conative states shift visual computation in a similar way as effected by intention, namely as within the divisive normalization model presented earlier. There could be multiple and disparate influences on attention. Whether this is so, of course, is an empirical matter. What gives flesh to the proposal is the divisive normalization model.

Now consider control versus automaticity in attention:

S's attention to X is controlled in respect of its feature F iff *S's attending to X with feature F is a result of an intention to do so in the F way.*

Intuitively, your attention involves control when its having feature F (in the limiting case, its being directed towards some target) is a result of your intending to attend in that way. So, control in attention is attending as you intend: if I intend to attend to the snake, then my attention's having the snake as its target is a controlled feature. Automaticity then can be defined

S's attention to X is automatic in respect of its feature F iff *S's attending to X with feature F is not a result of an intention to do so in the F way.*

That is, attention is automatic relative to F iff it is not controlled relative to F. So, if my attention is pulled to the snake despite my wanting to look away, then attention

is automatic. Automaticity is tied to the absence of intentions in respect of F.³⁶ The point is that when James speaks of the withdrawal from one object concomitant with attention to another, that withdrawal is often automatic. One doesn't need to intend to ignore one object when one intends to attend to another. Sometimes, withdrawal just (necessarily/automatically) happens.

Let's not quibble over definitional details. The issue is the forest, not the trees. Consider an epistemic agent who aims to figure out what to believe but is confronted with disparate sources of evidence. Epistemic agents face a Many-Many Problem: given many sources of evidence, what should they believe? Consider a simple situation where perceptible reasons p and q (e.g. considering a set of experimental data) seem to be relevant to the belief that r , the belief that *not* r and the belief that s among others. We can then map an *epistemic behavior space* that identifies links between putative reasons (evidence) provided by perception and potential perceptual beliefs (again, this is a simplified case). The links identify paths that when taken yield beliefs based upon perceived evidence. Some of these paths will be reasonable in that the evidence supports the belief; some will not in that support is absent though of course, the agent could mistakenly take there to be support. As in the Many-Many Problem for bodily movement, taking one path is a form of selection, here for epistemic action: in selecting some perceptual evidence over others, as when examining a slew of laboratory data, the agent is deploying attention to select that evidence to inform the task of fixing belief.

³⁶ For more discussion, see Wayne Wu, "Mental Action and the Threat of Automaticity," in Tillman Vierkant, Julian Kiverstein, and Andy Clark eds., *Decomposing the Will* (Oxford: Oxford University Press, 2013).

The biases tied to attention amount to biases tied to selectivity in epistemic agency. I conclude by considering some notable cases: Case 1: consider an agent who recognizes that evidence p supports r and evidence q supports s where r and s are two competing scientific hypotheses. An epistemic vice in this case is to willfully ignore q and focus on r . In this way, the subject intentionally attends to p and withdraws from q , perhaps because r is the subject's pet theory. Here, attention is top-down and its concomitant withdrawal is intentional. Clearly, such an agent fails a basic epistemic virtue. Yet consider Case 2: as James noted, by attending to something, this necessitates withdrawal from other things. Another agent might attend to p which supports their favored theory r and in that automatic way fail to notice q which supports the conflicting hypothesis. In attending appropriately to evidence for hypothesis r , the agent fails to notice relevant evidence for the competitor. This is, of course, an epistemic failure but not an intentional one. Rather, unlike Case 1, the agent in Case 2 lacks a certain capacity, a susceptibility to notice relevant evidence. So, we are likely to castigate the agent in Case 1 more than the agent in Case 2. But note that Case 2 can be fleshed out in different ways: Case 2A, where the subject just has a limited capacity for what can be attended to, so that the ignoring of additional evidence is automatic due to cognitive capacity limitations; Case 2B, where the subject has biases that are not rooted in intentions to ignore evidence but rather due to other states such as desires, values, or habits. These top-down, non-intention influences can also drive attention so that the filtering out of additional evidence occurs, even if these influences are in the defined sense automatic. The nature, then, of the top-down influence on attention matters. The

subject in 2B has something to correct, biases that have negative epistemic consequences despite sufficient attentional capacity; the subject in 2A simply has a limited capacity to process information, something that might simply be a product of a more limited cognitive endowment. Where a specific epistemic agent stands with respect to bias and capacities for attention will be in part an empirical matter; there is no reason to expect that each epistemic agent is similar to others with respect to appropriate and inappropriate biases as well as attentional and agential capacities.

Focusing on appropriate evidence is ideal, but as we have seen, forming a belief on the basis of such focus can be an action with different sources, leading to different assessments of one's epistemic standing. Individuating such cases is the payoff of highlighting attention's role in agency and the different ways it can be influenced and deployed. Consider, as one more example the capture of attention. Inattentional blindness studies demonstrate the different powers certain classes of stimuli have in disrupting top-down attention. As Arien Mack and Irving Rock demonstrated, a dot placed outside of a subject's zone of attention in demanding tasks is not very effective in drawing the subject's attention away from her current focus. In contrast, a subject's name is very effective in capturing attention.³⁷ So consider Case 3: a subject is appropriately reflecting on evidence relevant to hypothesis r and is focusing on evidence p . We further hope that the agent can be sensitive to additional evidence and come to notice it, say the presence in her collected data of evidence q , evidence that might support a contrary hypothesis. This

³⁷ Arien Mack and Irvin Rock, *Inattentional Blindness* (Cambridge, MA: MIT Press, 1998).

requires, in certain cases, the power of that evidence to capture the subject's attention, to circuit break current theoretical activity as it were.³⁸ Given the results of inattention blindness, we expect speakers to have different capacities to notice things outside of attention's current focus. Will the subject notice relevant evidence outside of the evidence attended? Here, the capacity of such evidence to automatically and bottom-up, as it were, drive attention will depend on a host of factors, biases of negative and positive sources—prejudices, incorrect valuation, suspect desires, and epistemic habits to name a few. A subject so negatively disposed to the alternative hypothesis, perhaps one endorsed by a hated rival, might simply fail to notice relevant evidence simply because the capacity for bottom-up attention (attentional capture) is limiting.

What this all points to, I think, is that being sensitive to evidence is in many ways a type of skill. An agent confronted with a plethora of evidence, much of it pulling in different directions, must find a way to sort it out. An ideal case is a calm reflector who has time to weigh the evidence and come to some reasonable conclusion in light of rational reflection. In such cases, beliefs are appropriately based on the evidence, a link I claim requires attention to the evidence due to the Many-Many Problem. But once we see attention's role and also the different sorts of biases and states that can influence its deployment, we can see that real-life agency, epistemic or otherwise, represents a fine balance between control and automatic features of that process and between top-down and bottom-up elements. Like a

³⁸ Maurizio Corbetta and Gordon L. Shulman, "Control of Goal-Directed and Stimulus-Driven Attention in the Brain," *Nature Reviews: Neuroscience* III (2002): 20115.

skilled gymnast or concert pianist, skilled epistemic agents must find a way to balance controlled and automatic features in a way that is also sensitive to the right top-down sources, the agent's intentions, desires, emotions and so forth. This can take learning and training. These top-down influences must also leave room for noticing new information, bottom-up as it were, when the context requires.

I began with the idea of shaking up the mind's ground floor, the perceptual tribunal of empirical beliefs, through intention. That shaking up is, I think, empirically supported by the informational penetration of the sort discussed earlier, and has potentially disparate sources and implementations. The bias that we worry about in epistemology and action interacts with the computational biases that realize attention. How we cope with such seismic shifts in perception is just the challenge of agency, one we can meet with training, skill and appropriate attention.