2

Public Objects and Private Qualia

The Scope and Limits of Psychophysics

Donald D. Hoffman

The Standard Framework

Psychophysics is typically defined as the scientific study of relationships between physical stimuli and the perceptual experiences they evoke. The physical stimuli might be tones of differing amplitudes, the evoked perceptual experiences might be the subjective loudnesses of the tones, and the relationship might be expressed as a simple equation, such as a power law, relating amplitude and loudness (e.g., Stevens, 1957). Or the physical stimulus might be the surface of an apple, the evoked perceptual experience might be the visually perceived 3-dimensional (3D) shape of that surface, and the relationship might be expressed as a complex mathematical formula that displays the relative influences of shading cues, texture gradients, and stereoscopic disparities on the perceived 3D shape (e.g., Knill & Richards, 1996).

In standard psychophysical theory, physical stimuli are conceived as being objective, in the sense that they exist whether or not they are perceived. A tone, as a physical acoustic stimulus, is conceived as existing and having a specific amplitude whether or not it is perceived. An apple, as a physical stimulus, is conceived as existing and having a specific shape whether or not it is perceived.

By contrast, in standard psychophysical theory, perceptual experiences are conceived as being subjective, in the sense that they do not exist if they are not perceived, and can vary among observers. An acoustic stimulus has amplitude but no loudness unless it is perceived, and the perceived loudness of a particular acoustic stimulus can differ from one observer to another. An apple has no visual shape without an observer, and its perceived shape can differ from one observer to another.

Because, in the standard view, physical stimuli are objective and perceptual experiences subjective, all lawful relationships that psychophysics discovers between them constitute an empirical bridge between the objective and subjective, and therefore are a critical source of data for the development of scientific theories of the

Handbook of Experimental Phenomenology: Visual Perception of Shape, Space and Appearance, First Edition. Edited by Liliana Albertazzi.

relationship between mind and body, between the physical world and conscious experiences. It was in hope of constructing such a theory that Gustav Theodor Fechner launched the field of psychophysics with his 1860 book *Elemente der Psychophysik*. He understood that laws relating amplitude and loudness, objective 3D shape and subjective visual experiences of 3D shape, are invaluable constraints on theoretical attempts to solve the mind–body problem.

In the standard view, physical stimuli *evoke* the corresponding perceptual experiences: There is a *causal* relation between physical stimuli and perceptual experiences, mediated by the neurophysiology of the sensory systems that transform physical stimuli into perceptual experiences. Thus the lawful relationships that psychophysics discovers are in fact descriptions of causal relationships. The amplitude of an acoustic signal is causally responsible, via the intervening neurophysiology of the auditory system, for the resulting perceived loudness. The shape of the apple is causally responsible, via the intervening neurophysiology of the visual system, for the resulting perceived shape.

The source of these causal relationships is assumed to be physics together with evolution by natural selection (and other evolutionary forces; see, e.g., Orzack & Sober, 2001). There are, in principle, many possible relationships between physical stimuli and perceptual experiences. However, these relationships differ in the fitness that they confer on the organism. Over time, natural selection has shaped our perceptual systems to instantiate relationships that are more adaptive, and has eliminated many other relationships in the process.

There is a standard view about the adaptive relationships between experience and reality that result from natural selection. According to this view, perceptual experiences that more closely approximate true physical properties are ipso facto more fit. The idea is that creatures whose perceptions are more veridical, that is, closer approximations to the objective truth, outcompete rivals whose perceptions are less veridical. Thus, over time, natural selection has shaped perceptual experiences to be good estimates of objective physical properties.

For instance, the textbook *Vision Science* explains that "Evolutionarily speaking, visual perception is useful only if it is reasonably accurate . . . Indeed, vision is useful precisely because it is so accurate. By and large, *what you see is what you get*. When this is true, we have what is called *veridical perception* . . . This is almost always the case with vision . . ." (Palmer, 1999a, p. 6). Knill, Kersten, and Yuille (1996, p. 6) explain that "Visual perception . . . involves the evolution of an organism's visual system to match the structure of the world and the coding scheme provided by nature." Purves, Lotto, Williams, Nundy, and Yang (2001, p. 296) offer a diachronic version of this idea in which ". . . percepts correspond to, and are generated by, the historical significance of proximal stimuli." Lehar (2003, p. 375) says, "The perceptual modeling approach reveals the primary function of perception as that of generating a fully spatial virtual-reality replica of the external world in an internal representation." Noe and Regan (2002) say that "Perceivers are right to take themselves to have access to environmental detail and to learn that the environment is detailed" (p. 576) and that ". . . the environmental detail is present, lodged, as it is, right there before individuals and that they therefore have access to that detail by the mere movement of their eyes or bodies" (p. 578).

This standard view can be captured in the mathematics of Bayesian estimation (see, e.g., Maloney, 2002, for an accessible review). If $W = \{w_1, \ldots, w_n\}$ are the relevant objective states of the world, and $S = \{s_1, \ldots, s_m\}$ are the relevant stimulations at a sensory transducer, then natural selection shapes the perceptual system so that when it receives stimulation s_i its resulting perceptual experience is, usually, of the best estimate w_j of the world state. This estimate can be obtained by computing the conditional probability $P(W|S)$, and then choosing the world state w_i that, for instance, maximizes this probability (the so-called "maximum a posteriori," or MAP, estimate). The conditional probability $P(W|S)$ can be computed from Bayes' theorem (Bayes, 1763):

$$P(W|S) = P(S|W)\, P(W) \,/\, P(S), \tag{2.1}$$

where $P(W)$ is a probability measure called the *prior* that describes the unconditional probability of the various world states, and $P(S|W)$ is a Markovian kernel (Revuz, 1984, p. 12) called the *likelihood function* that describes probabilistically how world states lead to stimulations at the sensory transducer. $P(W|S)$ is called the *posterior* probability. An *ideal Bayesian observer* always perceives the world state w_i that is the best Bayesian estimate for any given transducer stimulation s_i. Natural selection, according to the standard view, has shaped our perceptual experiences to be good approximations, in most cases, to ideal Bayesian observers.

The best Bayesian estimate depends not only on the posterior probability $P(W|S)$, but also on the utilities associated with the various possible perceptions and misperceptions of the world. In the standard view, these utilities are determined by accuracy, with more accurate perceptions having greater utility. Geisler and Diehl (2003), for instance, remark that, "In general, (perceptual) estimates that are nearer the truth have greater utility than those that are wide of the mark." This assumption motivates using the MAP estimate as the perceived world state.

In summary, the standard psychophysical framework is as follows. An objective physical world causes transducer stimulations that, via neural processes, lead to subjective perceptual experiences. Natural selection has shaped these processes so that our perceptual experiences are good Bayesian estimates of the true objective world states that cause them. As a result, our perceptions are usually close to veridical, and "what you see is what you get." Perception is useful because, in the normal case, it accurately reports truth.

A Problem with the Standard Framework

One problem with the standard framework is its assumption that more accurate perceptions have greater utility. This is in general not the case, as the simple example in Figure 2.1 illustrates (see Marion, Stephens, Mark, & Hoffman, 2011, for a similar example).

In this example, there are three states of the world (w_1, w_2, w_3) and two possible sensory stimulations (s_1, s_2). To fix ideas, we can imagine that the three world states are, say, three kinds of food that an organism might eat. The first two columns give

	s_1	s_2	Prior	Utility
w_1	1/3	2/3	1/5	10
w_2	2/3	1/3	2/5	2
w_3	1/3	2/3	2/5	2

Figure 2.1 A simple example showing that more accurate perceptions need not have greater utility.

the likelihood function, $P(S|W)$; for instance, $P(s_1|w_1) = 1/3$. The third column gives the prior probabilities of the world states. The fourth column shows the utilities to the organism of the world states. We can think of each utility as the fitness benefit an organism gets if it eats the food.

A simple computation using Bayes' theorem shows that for s_1 the MAP estimate is w_2, and that for s_2 the MAP estimate is w_3. (All computations mentioned in this section are given explicitly in the Appendix.) Thus an ideal Bayes observer would perceive the food w_2 if presented with stimulus s_1 and would perceive the food w_3 if presented with stimulus s_2. Suppose this observer is now offered a choice between two foods to eat, one that gives it stimulus s_1 and one that gives it stimulus s_2. Then this observer will perceive that it has been offered a choice between the foods w_2 and w_3. It has been shaped, we assume, by natural selection to choose, when possible, the food with greater utility. However, the utility of w_2 is 2, as is the utility of w_3. So if an organism is offered a choice between any two of the three kinds of food, that is, between any two of the three world states, it must choose at random, since it sees all options as having the same utility. This is unfortunate, since one food, namely w_1, has far greater utility than the others. An observer that was more likely to eat w_1, when it was available, would have a fitness advantage over the ideal Bayes observer.

It is easy to construct such an observer. Consider a different organism, presented with the situation described in Figure 2.1, but which is not constrained to estimate the truth. Suppose that instead it simply perceives two colors, say Red and Green. It sees Red when given sensory stimulation s_1 and Green when given s_2. Suppose further, that this observer has been shaped by natural selection to prefer the color that has greater expected utility. A simple computation (see the Appendix) shows that the expected utility of Green is 4 and the expected utility of Red is 3.143. So Green has greater expected utility than Red, and this observer will choose Green over Red. As a result, it will be better than the ideal Bayesian observer at getting as much fitness as possible from its choice of foods. We will call this new observer an *interface* observer (Hoffman, 1998, 2009; Koenderink, 2011; Mark, Marion, & Hoffman, 2010), since it does not see the truth but instead sees symbols, in this case colors, that guide adaptive behavior. The perceptions of an interface observer are tuned to discriminate utility, not reality.

In this example, the ideal Bayesian observer always chooses at random when offered a choice between two foods, whereas the interface observer makes a choice, whenever possible, that maximizes expected utility. So the interface observer is more fit than the ideal Bayesian observer. However, the situation can be far worse for the ideal Bayesian observer. If, for instance, in Figure 2.1 we set the utility of w_2 to be 3 (instead of 2) and leave the utility of w_1 and w_3 alone, then the ideal Bayesian observer will no longer choose at random when offered a choice between two foods. Instead, if one of the foods leads to sensory stimulation s_1 and the other leads to s_2, then the ideal Bayesian observer will always choose the food that leads to s_1, because it sees this as w_2 and now w_2 has greater utility (namely, 3) than w_3 (namely, 2). However, the interface observer will always make the opposite choice. It will choose the food that leads to s_2, because the expected utility of s_2 is 4.25 whereas the expected utility of s_1 is 3.71 (see the Appendix for the computation). So the interface observer still sees Green when given s_2 and Red when given s_1, and always chooses to maximize expected utility. The ideal Bayesian observer always chooses, in this case, to *minimize* expected utility. The interface observer maximizes fitness, the ideal Bayesian observer minimizes fitness. In a head-to-head competition, interface observers would drive ideal Bayesian observers to quick extinction (see Mark et al., 2010, for evolutionary games in which interface observers outcompete critical-realist observers; see Nowak, 2006, for an accessible introduction to evolutionary games).

So, the standard framework for psychophysics assumes that more accurate perceptions have greater utility. The example just discussed shows that this assumption is, in general, false. Moreover, we see that more accurate perceptions can lead to *minimizing* expected utility.

This is no small point. The standard framework for psychophysics depends critically on the assumption that more accurate perceptions have greater utility, and that therefore our perceptual systems estimate the truth. If this is false, then much of the rest of the framework is false. For instance, psychophysical functions are, in the standard view, assumed to give lawful relations between perceptions and the truth. But what is assumed, in the standard view, to be the truth, namely our perceptions of the physical world, has just been called into question. What, then, is the objective truth, and how can we know it, if our perceptions have been shaped by natural selection to discriminate expected utility, not to report the truth? We clearly need a new framework for psychophysics.

The Interface Theory of Perception

Perceptions that are tuned to discriminate expected utility will, in general, be fitter than those that are tuned to discriminate reality. The interface theory of perception is a conceptual framework and mathematical formalism to understand perceptions that are tuned to utility.

The conceptual difficulty we face can be put as follows: If our perceptions aren't true, then how in the world can they be helpful?

Fortunately, technology has given us a metaphor than can help us over this conceptual hurdle: the windows desktop of a PC. The desktop is an interface between

El

the user and the computer that is designed to guide adaptive behaviors toward the computer (e.g., keystrokes, mouse clicks, joystick moves) while allowing the user to be ignorant of the complexity of the hardware and software inside the computer. Icons on the desktop have colors, shapes, and positions, all of which are useful guides to behavior, but none of which are true. Suppose, for instance, that there is a red, square icon in the upper left corner of the screen. This does not mean that the corresponding file in the computer is red, square, and in the upper left corner of the computer. To think otherwise would be to completely misunderstand the purpose of the interface. It is there to hide the complexity of the computer, and to give simple symbols that guide useful behavior. The icons on the desktop are tuned to utility, to be guides for useful behavior; they are not intended to be accurate descriptions of reality. Indeed, one reason we pay good money for desktops is that they free us from the burden of accurately seeing the complexity of the computer. Before desktops were available, computers were the province of just a few talented engineers. Now they are accessible to everyone. So here we have an intuitive case where seeing the truth can be an impediment, and where seeing an interface that hides the truth and guides useful behaviors can be far more helpful.

As a conceptual framework for understanding psychophysics, the interface theory of perception says that our perceptions (and those of all species) are comparable to an interface, one that has been shaped by natural selection. Our perceptual interface has been shaped to guide adaptive behaviors and to hide the complexity of the truth. Space and time are the desktop of the visual interface of *Homo sapiens*, and 3D objects with their colors, shapes, textures, positions, and movements are among the icons of that desktop. The desktop and its icons do not resemble the objective world, whatever its nature might be. Instead they hide that world and serve as simplifying guides to adaptive behavior. The reason our perceptual systems do not, in general, report the truth is because, as our example in the previous section shows, true perceptions are not, in general, more fit, and are therefore eliminated by natural selection. As Pinker (1997, p. 561) puts it, "We are organisms, not angels, and our minds are organs, not pipelines to the truth. Our minds evolved by natural selection to solve problems that were life-and-death matters to our ancestors, not to commune with correctness . . ."

This framework has consequences that are uncomfortable to standard intuitions and assumptions. We assume that space and time are objective properties of the external world. The interface theory of perception says, instead, that space and time are a species-specific adaptation that, in all likelihood, bears no resemblance to any aspect of objective reality. We assume that each object in space-time, say an apple, is public and objective. The interface theory says that the apple I perceive is an icon in my space-time perceptual desktop, and that the apple you perceive is an icon in your space-time perceptual desktop. Your apple and my apple are distinct, just as my headache is distinct from your headache. Something in the objective world triggers both of us to perceive an apple, but whatever that thing might be in the objective world, it almost certainly is nonspatial, nontemporal, and in no way resembles an apple.

One might object that this view has an obvious problem: If that leopard crouching behind the bush is just an icon of your perceptual desktop, why don't you run over and hit it with your fist? After you're dead, and your theory with you, we'll know that leopards are more than mere icons.

The reason not to hit the leopard is the same reason one does not carelessly drag an icon to the trash. Although icons do not resemble the truth, actions with icons do have consequences. If I drag that icon to the trash, I could lose months of work. I'm careful with the icon not because I take it literally, but because I take it seriously. Similarly, I'm careful about the leopard not because I take it literally, but because I take it seriously. Natural selection has shaped my perceptions over millions of years, endowing me with symbols that must be taken seriously. The leopard icon automatically triggers emotions and withdrawal reactions that are highly adaptive. I ignore them at my peril. Indeed, I am the offspring of those who did not ignore them. I must take my perceptual icons seriously, but this does not logically entail that I must take them literally. Since natural selection rewards us based on utility, not reality, chances are that none of our perceptual icons can be taken literally.

One might instead raise a deflationary objection as follows: There really is nothing new being said here. Physicists have told us since at least the advent of quantum mechanics that our perception of the world as having solid objects with shapes and colors is an illusion, and that really those objects are mostly empty space, with tiny particles zipping around at high speeds (see, e.g., Barfield, 1957). So of course our perceptions are not literally true, simply because our perceptions lack the resolution to see the truth. Nevertheless, our perceptions are a great approximation to the truth at the scale of resolution at which we do see. So, when we see the surface shape of an apple, of course that surface is not literally a true representation of reality. But it is a good estimate of the *envelope* of the motions of the atomic and subatomic particles inside the apple that are the truth. We recognize that our perceptions cannot see the whole truth, because these particles are too small and fast. But we still get to keep the ideal Bayesian observer framework, in which our perceptual systems successfully estimate certain statistical properties of the truth, such as the envelope of the particle trajectories that we perceive as the surface of the apple.

The problem with this deflationary objection is that the particles themselves are in space and time, that is, within the desktop of the perceptual interface of *Homo sapiens*, and therefore cannot be the objective reality, only projections of that reality into our species-specific interface. The mistake here parallels the mistake made by someone who says, "I understand that the icons on my computer desktop are not the reality, just convenient symbols. But if I pull out my trusty magnifying glass, and look really close at those icons, I can see individual pixels, and those are the objective reality." The pixels are still on the desktop, they are still part of the user interface, and in no way resemble the reality of (in this example) the computer hardware and software. In the same way, the subatomic particles are still in space and time, and thus are still part of our perceptual interface.

This idea is difficult, because there were no selection pressures for us to understand it. Those of our ancestors who took their perceptions seriously and literally had the same fitness advantage as those who took their perceptions seriously but not literally. The default assumption appears to be to take our perceptions seriously and literally, and logical argument is usually required to unmask the error of this assumption (clever experiments help also, such as Koenderink, Van Doorn, & Todd, 2009;

El

Koenderink, Van Doorn, de Ridder, & Oomes, 2010). The error is to assume that our perceptions are real insights into objective reality, when in fact they are simply a species-specific guide to adaptive behavior. They are shaped by the necessity of successfully raising offspring, not by some necessity to see truth.

Another objection to the interface theory says, "If our perceptions are not true, then why do we all see the same things? Isn't the easiest explanation simply that we are all seeing the same truth about the world?" Indeed it would be an easy explanation. But it is not the only possible explanation. Perceptual agreement does not logically entail perceptual truth. According to the interface theory, different people placed in the same circumstances see substantially the same things because their interfaces are substantially similar (see a similar argument by Da Pos & Albertazzi, 2010, for the case of color perception). There are, to be sure, small differences between individuals, in part because our species is still evolving, and variations are required for natural selection to operate. In contrast, there are probably large differences between many pairs of species, for example, between humans and mantis shrimps, because the interfaces of different species have been shaped to guide adaptive behavior in different niches. The mantis shrimp, for instance, has perhaps 10 different types of color receptors, compared to just 3 for humans, and can see light polarization (Cronin & Marshall, 2002; Cronin, Marshall, & Land, 1994; Marshall & Oberwinkler, 1999), suggesting a much richer color interface for mantis shrimps than for humans. It might be that mantis shrimps experience colors that no human can even imagine (try to imagine a color you have never seen before). Our perceptual interfaces are a source for, and constraint on, our imaginations.

One implication of the interface theory is that the objects we see, being icons of our interface, have no causal powers. One billiard ball hits a second. We conclude that the first billiard ball causes the subsequent motion of the second (see, e.g., Michotte, 1963, for an engaging investigation of the perception of causality). But this conclusion is a mistake. Our perceptual interface hides the true causal structure of objective reality, just as a windows interface hides the causal structure of a computer. If I drag an icon to the trash can, it appears that the movement of the icon to the trash can causes the deletion of the file. But this is simply a useful fiction. There is no feedback from the icon to the file, only feedforward from the computer to the windows interface.

The interface theory thus leads to epiphysicalism, the doctrine that space-time and physical objects within space-time are constructions of perception, and that these constructions have no causal powers. Epiphysicalism contrasts with epiphenomenalism, the doctrine that the physical brain causes our conscious experiences, and that these conscious experiences have no causal powers.

One radical consequence of epiphysicalism is the conclusion that the brain, being a physical object, has no causal powers. None of our perceptions or actions are caused by brain activity. This does not entail that it is pointless to study the brain. To the contrary, such brain imaging and neurophysiological studies are an essential source of data that cannot be obtained elsewhere. Epiphysicalism simply entails that, when interpreting this data, we must take care not to conflate our species-specific perceptual icons with objective truth, and we must not assume that our perceptions in any way resemble the truth.

Formal Models of Perception

We have informally described the difference between the ideal Bayesian observer versus the interface observer, and given a concrete example in which the ideal Bayesian observer always makes choices which minimize expected fitness, while the interface observer, in the same situation, always makes choices that maximize expected fitness. In this section we mathematically characterize the difference between the two models.

We again represent the objective world by some set W, and the possible sensory stimulations by some set S. In consequence of the sensory stimulations, an observer has various perceptual experiences, which we denote by the set E. These sets might have structures, such as a metric, partial order, topology, σ-algebra, or measure. Different theories of perception can be characterized by the restrictions they place on the perceptual map, m: $W \to E$.

Naïve-realist theories of perception say that our perceptual experiences are a direct awareness of the objective world. In this case $W = E$ and the map m is an isomorphism of the structures on W.

Strong critical-realist theories of perception say that our perceptual experiences are a direct awareness of a subset of the objective world. In this case $E \subset W$, and the map m restricted to the domain E is an isomorphism of the structures on W.

Weak critical-realist theories of perception say that our perceptual experiences are not, in general, a direct awareness of a subset of the objective world, but that the structure of our perceptual experiences reflects the structure of the objective world. In this case $E \not\subset W$, and m is a homomorphism of the structures on W.

The interface theory of perception says that our perceptual experiences are not, in general, a direct awareness of a subset of the objective world, and that the structure of our perceptual experiences does not, in general, reflect the structure of the objective world. However, the interface theory maintains that the probabilities of events in perceptual experience are systematically related to probabilities of events in the objective world. In this case, $E \not\subset$ W and m is not a homomorphism of the structures on W. The only restriction on m is that it is a measurable function. Specifically, if W is the σ-algebra on W, and E is the σ-algebra on E, then for every event $A \in E$, we have that $m^{-1}(A) \in W$. This simply means, in other words, that m is a random variable.

The ideal Bayesian observer is a strong critical realist, whose perceptual experiences mirror a subset of objective reality. The map m is a composition of two maps, π and β. The first map, π, projects the objective world W onto sensory stimulations S, that is, π: $W \to S$. The second map, β, uses Bayes' theorem to compute the MAP estimate of the sensory stimulations, that is, β: $S \to E \subset W$, such that for each $s_i \in S$, $\beta(s_i)$ is the MAP estimate of s_i. The perceptual map $m = \beta \circ \pi$, where o denotes composition of functions.

Utility can be modeled as a real-valued function, u, that depends on the objective world W and the possible organisms O, that is, u: $W \times O \to \boldsymbol{R}$. This captures the fact that the same world state can have different utilities for different organisms. Dark chocolate, for instance, is healthy for most humans but deadly for most cats. The utility function u is the source of the pressures that guide evolution by natural selection.

The problem with the ideal Bayesian observer, from an evolutionary perspective, is that its perceptual experiences are not constrained by the actual utilities described by the utility function u. Instead its perceptual experiences are constrained by an artificial utility function that places higher utility on perceptions closer to the truth. Only after its perceptual experiences are determined by this artificial utility function can the ideal Bayesian observer use the real utility function u to make choices. This extra step, using an artificial utility function, is not only unnecessary and a waste of computational resources, it can, as our example above shows, lead the ideal Bayesian observer to choose in a way that always minimizes fitness. Extra work for less fitness is not a solution that will long be tolerated by natural selection.

The interface observer is not constrained to have this extra step of estimating reality using an artificial utility function. In consequence, its perceptual experiences are free to be shaped by natural selection to directly discriminate expected utilities, based on the correct utility function u. Its perceptual map m is a composition of two maps, π and α. The first map, π, projects the objective world W onto sensory stimulations S, that is, $\pi: W \rightarrow S$. The second map, $\alpha: S \rightarrow E \not\subset W$, leads to perceptual experiences E that are tuned to utility, in the sense that for all $e_i \in E$, the subsets $m^{-1}(e_i)$ of the objective world W are constructed such that the interface observer can maximize its expected utility for the choices it faces and the actions it must take.

Consider, for example, the situation depicted in Figure 2.2. In this example, the objective world has a single resource that varies in quantity from 1 to 100, as indicated on the horizontal axis. The utility function is the roughly Gaussian-shaped curve. The interface observer has just three perceptual experiences: Red, Green, and Yellow. As indicated in the figure, the subset m^{-1}(Red) is the set of resource quantities {0, . . ., 38, 82, . . ., 100}. Looking at the utility function, we see that the perceptual experience Red is tuned to report the low utility quantities of the resource. The subset m^{-1}(Green) is the set {39, . . ., 49, 69, . . ., 81}. The perceptual experience Green is tuned to report the intermediate utility quantities of the resource. The perceptual experience Yellow reports the high utility quantities of the resource. This interface observer is not trying to estimate the true state of the world, that is, the exact quantity of the resource. Instead, its perceptions are tuned to discriminate utilities so that it can make choices that maximize its expected utility.

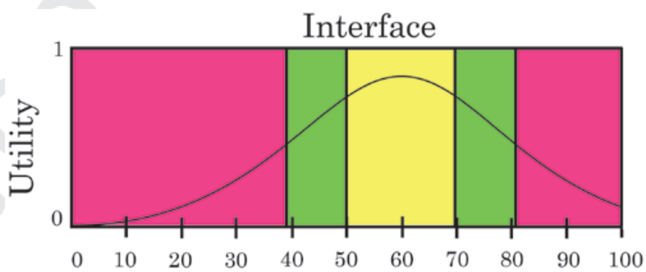

Figure 2.2 An interface observer with three perceptual experiences (Red, Green, Yellow) that are tuned to discriminate utility, rather than reality.

One general principle that follows from our concrete examples is as follows. Suppose that an organism is given k options, whose sensory stimulations are $s_1, \ldots, s_k$, and must choose one of them. The ideal Bayesian observer chooses $max(u(\text{MAP}(s_1)), \ldots, u(\text{MAP}(s_k)))$. The interface observer, shaped by natural selection, chooses $max(U(m^{-1}(e_1)), \ldots, U(m^{-1}(e_k)))$, where U denotes expected utility and $e_i \in E$. In consequence, the utility of the choice of the ideal Bayesian observer is always less than or equal to the utility of the choice made by the interface observer, with equality only happening in the special case where

$$max(u(\text{MAP}(s_1)), \ldots, u(\text{MAP}(s_k))) = max(U(m^{-1}(e_1)), \ldots, U(m^{-1}(e_k))). \qquad (2.2)$$

In the second section of this chapter we presented an example where $max(u(\text{MAP}(s_1)), \ldots, u(\text{MAP}(s_k))) = min(U(m^{-1}(e_1)), \ldots, U(m^{-1}(e_k)))$, which was strictly less than $max(U(m^{-1}(e_1)), \ldots, U(m^{-1}(e_k)))$. In general, then, the interface observer shaped by natural selection is fitter than the ideal Bayesian observer.

One can ask how an optimal interface observer changes as the complexity of its environment changes, as the utility function changes, and as the number of its perceptual categories changes. One can also ask how well this optimal interface observer competes with critical-realist observers as the cost of perceptual information and perceptual processing varies. In most cases, the situation is too complex to derive closed-form solutions. But Monte Carlo simulations (Mark et al., 2010) show that as the interface observer adds more perceptual categories, the categories tend to focus on discriminating among the high utility situations. For moderate increases in the complexity of the environment, critical realists gain fitness, but as the complexity continues to rise the critical realists lose fitness due to the increasing costs of collecting and processing perceptual information.

No Psychophysical Laws

Psychophysical laws are claimed to be lawful relationships between objective physical stimuli and the perceptual experiences they evoke. However, natural selection forges lawful relationships between *utility* and perceptual experiences. In the process, it cares nothing about forging lawful relationships between objective reality and perceptual experiences. In consequence, there are almost surely no psychophysical laws as they are traditionally understood, beyond the minimum requirement that the perceptual mapping m is measurable, that is, that m is a random variable.

How then are we to interpret, for example, Stevens' (1957) classic report of the various "perceptual continua on which psychological magnitude is a power function of the stimulus" (p. 166) and his general observation that "equal stimulus ratios produce equal subjective ratios" (p. 153)? And how are we to interpret the various Bayesian perceptual models, for example, in Knill and Richards (1996), showing lawful relationships between perceived 3D shape and actual 3D shape?

The human perceptual systems construct a complex species-specific interface containing perceptual experiences of many types and at many levels (see, e.g., Kanizsa, 1979, on the constructive nature of perception). There are lawful relationships between the perceptual experiences at different levels. Some "psychophysical" laws are in fact laws relating perceptual experiences to perceptual experiences.

Moreover, technology has made available to us many ways to extend our natural senses to build more detailed and precise perceptual descriptions within the space-time framework of our species-specific interface. The amplitude of an acoustic stimulus and the perceived loudness of that stimulus are both within our interface. The acoustic wave itself travels in space and time, and therefore resides within the space-time interface of *H. sapiens*. The amplitude of an acoustic stimulus is not an observer-independent feature of objective reality. The laws relating amplitude and loudness are therefore not laws relating objective reality and perception, but rather laws relating different levels of our interface. The acoustic level requires augmentation of our natural senses by technology, but the technology itself resides within our space-time interface and gives reports couched within the language of our space-time interface. To compare, a magnifying glass might let us see the pixels of our windows interface in a way not possible with the naked eye, but those pixels are still in the interface. Similarly, microscopes, telescopes, microphones, and numerous other technologies let us explore our interface in ways not possible with the unaided senses, but these technologies do not magically pierce the veil of our interface and reveal the objective reality behind. They remain trapped within the space-time desktop of our perceptual interface (see Mausfeld, 2002, for an engaging discussion of the "physicalist trap" in theories of perception).

Thus, no purported psychophysical law reported to date is a genuine relationship between objective reality and subjective perceptual experience. All purported psychophysical laws are really descriptions of lawful relationships between different aspects or levels of our subjective interface. There are no psychophysical laws, only intra-interface laws, which we might call *psychophenomical* laws. When an ideal Bayesian observer succeeds in quantitatively accounting for an aspect of human perception, this does not confirm that indeed our perceptions are ideal Bayesian estimates of the objective truth. It simply confirms that our perceptual interface is mathematically consistent, and that conditional probabilities can be consistently computed within it. Bayes' theorem is, after all, just the correct formula for computing conditional probabilities, and does not entail the conclusion that our perceptions estimate properties of objective reality.

This has an important implication for the prior probabilities $P(W)$ used in ideal Bayesian observer theory. In the standard story, these priors describe the probabilities of properties or events in the objective world. In order to use Bayes' theorem, as in Equation 2.1 above, to compute the needed posterior probabilities $P(W|S)$, one must have the prior probabilities $P(W)$. How can we know $P(W)$? The standard answer is that we can go out and measure $P(W)$ in the objective world.

That might work if we had a way to see the objective world. But we can only see within the confines of our species-specific interface, and that interface has been shaped by natural selection to discriminate utility, not reality. So the standard attempts to measure $P(W)$ will fail. When Bayesian vision scientists measure what they think are the true priors about objects and their properties, in fact what they are doing is measuring statistics of items within the interface of *H. sapiens*. Such measurements are useful for understanding our perceptual interface (see, e.g., Geisler, Perry, Super,

El

& Gallogy, 2001), but they are not reports of the objective probabilities P(W). Similarly, evidence that human subjects can optimally use utility information to adjust their visually guided pointing (Hudson, Maloney, & Landy, 2008) is not surprising, given the interface theory of perception. Nor is it surprising that perceptual learning is enhanced at moments of reward or behavioral success (Lin, Pype, Murray, & Boynton, 2010; Seitz, Kim, & Watanabe, 2009; Swallow & Jiang, 2010). But these results do not entail that human subjects have access to statistics of the objective world.

No Reductive Functionalism

Another assumption of the standard theory of psychophysics is reductive functionalism: All mental states, including perceptual states, are identical to functional states of the brain. The idea is that brain activity is responsible for all our conscious experiences, but that this is not due to something special about the physics or biology of neurons. Instead, it is due to something about the functional interactions of neurons, something about the program that they instantiate. If we understood this program we could, in principle, also instantiate it on a computer, and the computer would have conscious experiences as well.

It would be a mistake to ask a reductive functionalist how functional states of the brain cause conscious experiences. Their claim is not that functional states *cause* experiences, but that they are *identical* to the experiences. To compare, one would not ask how an unmarried man causes a bachelor; an unmarried man *is* a bachelor.

One reason for adopting reductive functionalism in the standard theory of psychophysics is that it gives a tool for relating measurable behaviors to conscious experiences. For instance, John Locke in his 1690 *Essay Concerning Human Understanding* raised a question now known as the spectrum inversion problem. He asked if it is possible that "the idea that a violet produced in one man's mind by his eyes were the same that a marigold produced in another man's, and vice versa." Could two observers be functionally identical, naming colors the same way, and giving the same responses in every possible test of color vision, and yet have radically different color experiences?

This question continues to be discussed more than three hundred years later (see, e.g., Palmer, 1999b). For a reductive functionalist, the answer is clear. If two observers are functionally identical, then all of their perceptual states are identical. So spectrum inversion of the kind discussed by Locke is not logically possible. Thus we can safely infer some features of perceptual experiences from the measurable functional properties of observers.

Unfortunately, reductive functionalism is provably false (Hoffman, 2006a, 2006b). The proof is straightforward. Let the set of Jack's color experiences be denoted by X and Jill's by Y. Assume that $X = Y$, that is, that Jack and Jill have the same range of color experiences. Suppose that Jill's color experiences are scrambled relative to Jack's, by some function b: $X \to Y$ that is a bijection, that is, a one-to-one and onto mapping. For instance, b might map Jack's reds to Jill's greens, and so on. Let f: $X \to W$, where W is some set, be any functional relations among Jack's color experiences. Then the spectrum inversion question becomes: Is it possible to find functional relations $g: Y \rightarrow W$ among Jill's color experiences that make Jack and Jill functionally indistinguishable? The answer is yes: The solution is $g = f \circ h^{-1}$. This solution works because the map $g \circ h$, which takes $X \rightarrow Y \rightarrow W$, is simply $g \circ h = f \circ h^{-1} \circ h = f$. So, if you scramble the color experiences of Jack by h to get the color experiences of Jill, and then use g as the functional relations on Jill's color experiences, you get exactly the same functional relations f that apply to Jack's original color experiences. So Jack and Jill are functionally identical, even though their color experiences are scrambled. Note that this solution holds for *every* bijective scrambling, h, between the color experiences of Jack and Jill. Moreover, since the sets X and Y could denote any kinds of perceptual experiences, not just color experiences, the proof applies to all perceptual experiences.

Thus perceptual experiences are not identical to any functional relationships among them, and reductive functionalism is false. One can adopt a weaker version of functionalism, namely nonreductive functionalism, which claims that all mental states, including perceptual experiences, are caused by, or somehow arise from, functional relations (see, e.g., Chalmers, 1996). Nonreductive functionalism is immune from the disproof of reductive functionalism just given. However, there are no explicit theories on offer for how functional relations can give rise to perceptual experiences, nor are there any remotely plausible ideas on offer. So nonreductive functionalism remains a philosophical option, not a scientific theory.

The upshot is that psychophysics has, at present, no secure theoretical foundation from which to infer perceptual experiences from measurable behaviors. To be sure, such inferences are routinely made, and must be made, in the course of psychophysical research. But the reductive functionalist framework that used to underwrite such inferences is now defunct, and no new framework has been erected in its place (but see Hoffman, 2008, for one possible direction).

One might ask whether functionalism is incompatible with the standard view, discussed in the first section of this chapter, that physical stimuli *evoke* the corresponding perceptual experiences, that is, that there is a causal relation between physical stimuli and perceptual experiences. The two doctrines are compatible. One can claim that physical stimuli evoke certain patterns of neural activity, and that the functional properties of this activity are identical to, or give rise to, the resulting perceptual experience. Thus, even reductive functionalism is compatible with the view that physical stimuli can cause perceptual experiences: Such stimuli can cause the functional properties of brain states that are identical with such experiences. However, reductive functionalism is false, and nonreductive functionalism is, at present, not even wrong.

Illusions and Hallucinations

The standard story in psychophysics is that perception is useful because it is, in the normal case, a good approximation to the truth. This standard story leads naturally to a standard account of illusions and hallucinations. According to this account, an illusion is an *untrue* perception that is triggered by a specific stimulus, and is experienced by most normal observers when exposed to that stimulus. A hallucination is an *untrue* perception that is seen by few people, and in the absence of an appropriate stimulus. The Necker cube is an illusion, because it is an untrue perception of a 3D cube that is experienced by most normal observers when they are exposed to a 2D line drawing of a cube. Seeing spiders crawling all over one's body while having delirium tremens is a hallucination because it is seen by few people, and is seen in the absence of spiders.

However, as we have seen, the standard psychophysical story is wrong. Perception is not useful because it is a good approximation to the truth. It is useful because it has been shaped by natural selection to discriminate utility, and to hide the causal and structural complexity of reality. In the normal case, no perceptions are true, and none are approximations to truth. Instead, our perceptions constitute a species-specific interface that guides adaptive behavior while hiding the truth.

The question naturally arises: How shall we revise the standard theory of illusions and hallucinations? The standard theory critically depends on the claim that illusions and hallucinations are exceptional cases in which perception turns out to be untrue. We now know that this is not the exceptional case, it is the normal case. Everyday perceptions, ones that we would not want to call illusions or hallucinations, are, in the normal case, untrue.

The key idea for a new definition of illusion and hallucination is that everyday perceptions, ones that we would not want to call illusions or hallucinations, are, in the normal case, adaptive guides to behavior. So we can revise the standard definition of illusion and hallucination by replacing "untrue" with "unadaptive guide to behavior."

Thus an illusion is a perception that is triggered by a specific stimulus, and is experienced by most normal observers when exposed to that stimulus, but is not an adaptive guide to behavior. The Necker cube is an illusion because the 3D perception of a cube is not an adaptive guide to behavior with a flat sheet of paper. The point is not that the flat sheet of paper is the truth, whereas the cube is false. Rather, the flat sheet of paper happens to be, in this instance, a perception that guides adaptive behaviors, whereas the cube is a perception that does not.

A hallucination is a perception that is seen by few people, is seen in the absence of an appropriate stimulus, and is not an adaptive guide to behavior. A perception of spiders crawling over one's body during delirium tremens is a hallucination because it is a perception seen by few people, is seen in the absence of an appropriate stimulus, and leads to behaviors, such as fear and brushing of skin, that in this instance are not adaptive.

Psychophenomics

Perception is a biological adaptation. A mutation that enhances the ability of perception to guide adaptive behaviors is more likely to be retained than one that does not. Other things being equal, a mutation that allows perception to be faster and cheaper is more likely to be retained than one that does not. As a result, perception is shaped to be fast and cheap at discriminating utility. Perceptual experiences are not guides to truth, but guides to behaviors that, when successful, are just a bit more

El

adaptive than those of one's competitors. Space-time, objects, shapes, colors, sounds, smells, tastes, and motions are not a window on truth, but a window on utilities relevant to a specific species trying to navigate its niche. They constitute a species-specific interface that guides adaptive behaviors and hides the causal and structural complexity of the objective world.

The laws discovered so far by psychophysics are not laws relating perceptual experiences to objective reality. Instead they are laws relating one part of the interface to another, one class of perceptual phenomena to another. They are psychophenomical laws.

There may be no psychophysical laws—if such laws are relations between perception and objective reality—save one: The mapping from objective reality to perception is a random variable. However, there may be many interesting laws relating perception and utility (e.g., Mark et al., 2010). But the study of these laws is in its infancy, because the standard view within psychophysics assumes that perception is about estimating truth. Natural selection is harsh with perceptions that waste time on truth, rather than focusing on utility. In effect it says, "You shall know the truth, and the truth shall drive you extinct."

Acknowledgments

I thank Majid Borumand, Pete Foley, Justin Mark, Brian Marion, and Kyle Stephens for helpful comments and discussions. This work was supported in part by a grant from Procter & Gamble.

Appendix

In this Appendix we perform the Bayesian and expected utility computations discussed in the second section of the paper, based on the data given in Figure 2.1.

To compute the MAP estimates, we first need to compute $P(s_1)$ and $P(s_2)$. These are as follows:

$$\begin{aligned} P(s_1) &= P(s_1|w_1)\,P(w_1) + P(s_1|w_2)\,P(w_2) + P(s_1|w_3)\,P(w_3) \\ &= (1/3)(1/5) + (2/3)(2/5) + (1/3)(2/5) \\ &= 7/15 \end{aligned}$$

$$\begin{aligned} P(s_2) &= P(s_2|w_1)\,P(w_1) + P(s_2|w_2)\,P(w_2) + P(s_2|w_3)\,P(w_3) \\ &= (2/3)(1/5) + (1/3)(2/5) + (2/3)(2/5) \\ &= 8/15 \end{aligned}$$

Then we can compute the posterior probabilities of the world states, given s_1.

$$\begin{aligned} P(w_1|s_1) &= P(s_1|w_1)\,P(w_1)/P(s_1) = (1/3)(1/5)/(7/15) = 1/7 \\ P(w_2|s_1) &= P(s_1|w_2)\,P(w_2)/P(s_1) = (2/3)(2/5)/(7/15) = 4/7 \\ P(w_3|s_1) &= P(s_1|w_3)\,P(w_3)/P(s_1) = (1/3)(2/5)/(7/15) = 2/7 \end{aligned}$$

El

Since $P(w_2|s_1)$ is larger than the others, the MAP estimate for stimulus s_1 is w_2. Next we can compute the posterior probabilities of the world states, given s_2.

$$P(w_1|s_2) = P(s_2|w_1)\ P(w_1)/P(s_2) = (2/3)(1/5)/(8/15) = \frac{1}{4}$$

$$P(w_2|s_2) = P(s_2|w_2)\ P(w_2)/P(s_2) = (1/3)(2/5)/(8/15) = \frac{1}{4}$$

$$P(w_3|s_2) = P(s_2|w_3)\ P(w_3)/P(s_2) = (2/3)(2/5)/(8/15) = \frac{1}{2}$$

Since $P(w_3|s_2)$ is larger than the others, the MAP estimate for stimulus s_2 is w_3. Next we can compute the expected utilities of the sensory stimuli s_1 and s_2.

$$\begin{aligned} U(s_1) &= P(w_1|s_1)U(w_1) + P(w_2|s_1)U(w_2) + P(w_3|s_1)U(w_3) \\ &= (1/7)(10) + (4/7)(2) + (2/7)(2) \\ &= 22/7 \end{aligned}$$

$$\begin{aligned} U(s_2) &= P(w_1|s_2)U(w_1) + P(w_2|s_2)U(w_2) + P(w_3|s_2)U(w_3) \\ &= (1/4)(10) + (1/4)(2) + (1/2)(2) \\ &= 4 \end{aligned}$$

We see that s_2 has greater expected utility, and will thus be seen as Green by the interface observer. If we change the utility of w_2 to 3, then the new expected utilities are as follows.

$$\begin{aligned} U'(s_1) &= P(w_1|s_1)U(w_1) + P(w_2|s_1)U(w_2) + P(w_3|s_1)U(w_3) \\ &= (1/7)(10) + (4/7)(3) + (2/7)(2) \\ &= 26/7 \end{aligned}$$

$$\begin{aligned} U'(s_2) &= P(w_1|s_2)U(w_1) + P(w_2|s_2)U(w_2) + P(w_3|s_2)U(w_3) \\ &= (1/4)(10) + (1/4)(3) + (1/2)(2) \\ &= 17/4 \end{aligned}$$

We see that s_2 still has greater expected utility, and will thus still be seen as Green, and will still be chosen, by the interface observer. However, the ideal Bayesian observer will now choose the food that causes stimulus s_1 since the MAP estimate for s_1 is w_2 and w_2 and has greater utility (namely, 3) than the utility of w_3 (namely, 2), which is the MAP estimate of s_2.

References

Barfield, O. (1957). *Saving the appearances: A study in idolatry.* London: Faber and Faber.

Bayes, T. (1763). An essay towards solving a problem in the doctrine of chances. *Philosophical Transactions of the Royal Society of London*, *53*, 370–418.

Chalmers, D. (1996). *The conscious mind.* Oxford: Oxford University Press.

Cronin, T. W., Marshall, N. J., & Land, M. F. (1994). The unique visual system of the mantis shrimp. *American Scientist*, *82*(4), 356–365.

Cronin, T. W., & Marshall, N. J. (2002). The unique visual world of mantis shrimps. In F. Prete (Ed.), *Complex worlds from simpler nervous systems* (pp. 239–268). Cambridge, MA: MIT Press.

Da Pos, O., & Albertazzi, L. (2010) It is in the nature of colors. *Seeing and Perceiving*, *23*, 39–73.

Fechner, G. (1860) *Elemente der Psychophysik*. Leipzig: Breitkopf und Härtel.

Hoffman, D. D. (1998). *Visual intelligence: How we create what we see*. New York: W. W. Norton.

Hoffman, D. D. (2006a). The scrambling theorem: A simple proof of the logical possibility of spectrum inversion. *Consciousness and Cognition*, *15*, 31–45.

Hoffman, D. D. (2006b). The scrambling theorem unscrambled: A response to commentaries. *Consciousness and Cognition*, *15*, 51–53.

Hoffman, D. D. (2008). Conscious realism and the mind–body problem. *Mind & Matter*, *6*, 87–121.

Hoffman, D. D. (2009). The interface theory of perception: Natural selection drives true perception to swift extinction. In S. Dickinson, M. Tarr, A. Leonardis, & B. Schiele (Eds.), *Object categorization: Computer and human vision perspectives* (pp. 48–165). Cambridge: Cambridge University Press.

Hudson, T. E., Maloney, L. T., & Landy, M. S. (2008). Optimal compensation for temporal uncertainty in movement planning. *PLoS Computational Biology*, *4*(7), 1–9.

Geisler, W. S., & Diehl, R. L. (2003). A Bayesian approach to the evolution of perceptual and cognitive systems. *Cognitive Science*, *27*, 379–402.

Geisler, W. S., Perry, J. S., Super, B. J., & Gallogy, D.P. (2001). Edge co-occurrence in natural images predicts contour grouping performance. *Vision Research*, *41*, 711–724.

Kanizsa, G. (1979). *Organization in vision*. New York: Praeger.

Knill, D. C., Kersten, D., & Yuille, A. (1996). Introduction: A Bayesian formulation of visual perception. In D. Knill & W. Richards (Eds.), *Perception as Bayesian inference* (pp. 1–21). Cambridge: Cambridge University Press.

Knill, D., & Richards, W. (Eds.). (1996). *Perception as Bayesian inference*. Cambridge: Cambridge University Press.

Koenderink, J. J. (2011). Vision and information. In L. Albertazzi, G. J. Van Tonder, & D. Vishnawath (Eds.), *Perception beyond inference. The information content of visual processes* (pp. 27–57). Cambridge, MA: MIT Press/Bradford Books.

Koenderink, J. J., van Doorn A. J., de Ridder, H., & Oomes, S. (2010). Visual rays are parallel. *Perception*, *39*, 1163–1171.

Koenderink, J. J., van Doorn, A. J., & Todd, J. T. (2009). Wide distribution of external local sign in the normal population. *Psychological Research*, *73*, 14–22.

Lehar, S. (2003). Gestalt isomorphism and the primacy of subjective conscious experience: A Gestalt Bubble model. *Behavioral and Brain Sciences*, *26*, 375–444.

Lin, J. Y., Pype, A. D., Murray, S. O., & Boynton, G. M. (2010). Enhanced memory for scenes presented at behaviorally relevant points in time. *PLoS Biology*, *8*(3), 1–6.

Locke, J. (1690/1979). *An essay concerning human understanding*. Cambridge: Cambridge University Press.

Maloney, L. T. (2002). Statistical decision theory and biological vision. In D. Heyer & R. Mausfeld (Eds.), *Perception and the physical world: Psychological and philosophical issues in perception* (pp. 145–189). Hoboken, NJ: John Wiley & Sons, Inc.

Marion, B., Stephens, K., Mark, J., & Hoffman, D. D. (2011). Seeing truth can minimize fitness. Manuscript submitted for publication.

Mark, J., Marion, B., & Hoffman, D. D. (2010). Natural selection and veridical perception. *Journal of Theoretical Biology*, *266*, 504–515.
Marshall, J., & Oberwinkler, J. (1999). The colourful world of the mantis shrimp. *Nature*, *401*, 873–874.
Mausfeld, R. (2002). The physicalist trap in perception theory. In D. Heyer & R. Mausfeld (Eds.), *Perception and the physical world: Psychological and philosophical issues in perception* (pp. 75–112). Hoboken, NJ: John Wiley & Sons, Inc.
Michotte, A. (1963). *The perception of causality*. London: Methuen.
Noe, A., & Regan, J. K. (2002). On the brain-basis of visual consciousness: A sensorimotor account. In A. Noe & E. Thompson (Eds.), *Vision and mind: Selected readings in the philosophy of perception* (pp. 567–598). Cambridge, MA: MIT Press.
Nowak, M. A. (2006). *Evolutionary dynamics: Exploring the equations of life*. Cambridge, MA: Belknap/Harvard University Press.
Orzack, S. H., & Sober, E. (2001). *Adaptationism and optimality*. Cambridge: Cambridge University Press.
Palmer, S. (1999a). *Vision science: Photons to phenomenology*. Cambridge, MA: MIT Press.
Palmer, S. (1999b). Color, consciousness, and the isomorphism constraint. *Behavioral and Brain Sciences*, *22*, 923–989.
Pinker, S. (1997). *How the mind works*. New York: W. W. Norton.
Purves, D., Lotto, R. B., Williams, S. M., Nundy, S., & Yang, Z. (2001). Why we see things the way we do: Evidence for a wholly empirical strategy of vision. *Philosophical Transactions of the Royal Society of London B*, *356*, 285–297.
Revuz, D. (1984). *Markov chains*. Amsterdam: North-Holland.
Seitz, A. R., Kim, D., & Watanabe, T. (2009). Rewards evoke learning of unconsciously processed visual stimuli in adult humans. *Neuron*, *61*, 700–707.
Stevens, S. S. (1957). On the psychophysical law. *Psychological Review*, *64*, 153–181.
Swallow, K. M., & Jiang, Y. V. (2010). The attentional boost effect: Transient increases in attention to one task enhance performance in a second task. *Cognition*, *115*, 118–132.

El

El