Philosophical Psychology
Vol. 22, No. 1, February 2009, 61–75

Routledge
Taylor & Francis Group

Closing the gap between ideal and real behavior: Scientific vs. engineering approaches to normativity

Sergei Gepshtein

Early normative studies of human behavior revealed a gap between the norms of practical rationality (what humans ought to do) and the actual human behavior (what they do). It has been suggested that, to close the gap between the descriptive and the normative, one has to revise norms of practical rationality according to the Quinean, engineering view of normativity. On this view, the norms must be designed such that they effectively account for behavior. I review recent studies of human perception which pursued normative modeling and which found good agreement between the normative prescriptions and the actual behavior. I make the case that the goals and methods of this work have been incompatible with those of the engineering approach. I argue that norms of perception and action are observer-independent properties of biological agents; the norms are discovered using methods of natural sciences rather than the norms are designed to fit the observed behavior.

Keywords: Bayesian; Decision; Economics; Engineering; Haptics; Illusion; Normativity; Psychophysics; Rationality; Vision

1. Introduction

In 1964, George A. Miller, a proponent of mathematical thinking in psychology, had the following to say of the relationship between normative and descriptive theories of human behavior:

Sergei Gepshtein is a research scientist at the Theoretical Neuroscience Group of RIKEN Brain Science Institute, Japan and a staff scientist at the Vision Center Laboratory of the Salk Institute for Biological Studies, CA, USA. Correspondence to: Sergei Gepshtein, Vision Center Laboratory, The Salk Institute for Biological Studies, 10010 N. Torrey Pines Road, La Jolla, CA 92037, USA. Email: sergei@salk.edu

ISSN 0951-5089 (print)/ISSN 1465-394X (online)/09/0100061-15
DOI: 10.1080/09515080802703695

> For many years psychologists regarded the work of their economic colleagues as a good example of bad psychology. Economists seemed too much concerned with the rational and ethical problems of what people ought to do. Psychologists, fancying themselves natural scientists, not logicians or moral philosophers, insisted on searching instead for the laws that govern what people actually do. The vast gap between what they actually do and what they ought to do is overwhelmingly clear to everyone, so it seemed that the normative approach of the economist had to be rejected a priori in psychology. . . . Beginning around 1950, however, the psychological aspects of economic theory began to grow increasingly interesting to many psychologists. (p. 32)

The tide noticed by Miller brought the normative approach into the mainstream of psychology and neuroscience. Today, normative theories are routinely used as testbeds for assumptions about biological mechanisms of sensorimotor behavior, learning, and decision-making. Among the modern normative approaches, the center stage is occupied by models based on the Statistical Decision Theory (Geisler & Kersten, 2002; Green & Swets, 1966; Knill & Richards, 1996; Maloney, 2002; Rao, Olshausen, & Lewicki, 2002) of which we will see several examples below.

In a recent issue of *Philosophical Psychology*, Hardy-Vallée and Thagard (2008) ponder Miller's gap between what agents ought to do (the norms) and what they actually do. The authors take a metanormative perspective. They ask a normative question about building normative models: How ought one to choose norms for evaluation of behavior? Hardy-Vallée and Thagard explore a prototypical case of decision-making—the ultimatum game—and conclude that the traditional view (the ''standard picture'') of rationality is incomplete because it lacks an adequate account of effectiveness of norms. Effectiveness of norms is the degree to which norms account for observed behavior. Hardy-Vallée and Thagard are concerned that the effectiveness is generally assumed rather than evaluated, which on their view leads to a distorted picture of rationality.

As an alternative to the standard picture of rationality, Hardy-Vallée and Thagard defend an engineering approach, owing to the naturalistic epistemology of philosopher Willard V. O. Quine who viewed normative theories as a branch of engineering (Quine, 1986). From this perspective, norms must be selected (or perhaps designed) to insure their effectiveness: ''engineers do not build bridges or launch rockets, but spell out what one should do in order to attain this goal'' (Hardy-Vallée & Thagard, 2008, p. 180). From their explicit premise that engineering is normative and science is descriptive, the authors maintain that an adequate account of rationality will emerge from the engineering approach to normativity.

In this note I make the case that the concern about effectiveness of norms is being answered by the modern normative approaches to human perception and action, but the emerging answer is inconsistent with the engineering view of normativity. I argue that norms of perception and action studied in the modern work are observer-independent properties of biological agents: the norms are discovered, as were the laws of classical physics; rather than the norms are fitted to observed behavior, as statistical models are fitted to experimental data.

2. Illusion of Irrationality

Before I turn to a broader topic of genesis of norms, I will scrutinize the belief that errors of perception or judgment by biological agents necessarily manifest suboptimal (irrational) behavior. This belief transpires in Hardy-Vallée and Thagard's (2008) discussion of how actual behavior can deviate from normative prescriptions. Central to the authors' argument is the distinction of external vs. internal rationality:

> An 'internal' (or subjective) assessment of rationality is an evaluation of the coherence of intentions, actions and plans. . . . An 'external' (or objective) assessment of rationality is an evaluation of the effectiveness of a rule or procedure. . . . An action can be rational from the first perspective but not from the second one, and vice versa. Hence subjects' poor performance in probabilistic reasoning can be internally rational (subjects may have good reason to choose a certain prospect) without being externally rational (their behavior is still suboptimal). The Gambler's fallacy is and always will be a fallacy: it is possible, however, that fallacious reasoners follow rational rules, maximizing an unorthodox utility function. This distinction thus specifies two nonexclusive ways in which someone can behave irrationally. One can be externally irrational if the outcome of an action is suboptimal; in this case the attribution of irrationality requires data about the agent and the outcome. One can also be internally irrational if, regardless of the outcomes of the action, the agent's desires and action performed are incoherent. (Hardy-Vallée & Thagard, 2008, pp. 181–182)

Hardy-Vallée and Thagard evaluate human performance in perceptual illusions as "externally irrational:"

> Studying optical illusions, for instance, [Weiss, Simoncelli, and Adelson (2002)] conclude for that they are 'the best solution of a rational system designed to operate in the presence of uncertainty' (p. 598). Geisler and Kersten (2002) explain that Weiss et al.'s assumptions about the probability and likelihood distributions of velocities . . . 'incorporated [implicitly] into the visual system arise through a combination of evolution and perceptual learning' (p. 509). . . . This is where the distinction between internal and external assessment of rationality is important: research on perception does not prove that optical illusions are not illusions. They are still externally irrational but appear as internally rational, that is, produced by a rational Bayesian mechanism. It is still wrong—externally irrational—to infer that two lines are of different length in the Müller-Lyer illusion even for someone well-versed in the psychology and psychophysics of perception because, when we measure these two lines, they are of different length. (Hardy-Vallée & Thagard, 2008, pp. 182–183)

The authors reject the claim that perceptual illusions can be rational inferences. Hardy-Vallée and Thagard believe that non-veridical (incorrect) perception must be irrational on some account; they use the notion of external irrationality to support their case.

Since the disagreement about rationality in illusions appears to hinge on the distinction between veridicality and rationality, I will next trace this distinction in two studies of perceptual illusions. Both studies represent a probabilistic normative approach to modeling biological systems. Such probabilistic models take into account the ambiguities inherent in biological measurements (e.g., Knill & Richards, 1996) and also the imperfections of "noisy" biological computations (e.g., Blakemore, 1990).[1] These models are formal statements that incorporate rules of logic and mathematics, as required by the standard picture of rationality cited by Hardy-Vallée and Thagard. At the same time, these models represent a naturalistic approach to normativity: they take into account basic properties of biological systems, which is key to meaningful comparison of normative prescriptions with biological phenomena. Although predictions of such models are "noisy" (they vary stochastically from trial to trial) and may exhibit a systematic bias (as is the case in perceptual illusions), they nevertheless implement ideal rational agents: ideal because the models capture only the essential properties of agents; rational because they follow rules with the explicit goal to optimize behavior.

2.1. Motion Illusions

My first illustration is the study of Weiss et al. (2002) who argued that a family of motion illusions is a result of optimal computations by the human visual system. The ideal-observer model of Weiss et al. is an implementation of the Bayesian Decision Theory. The model predicts how an optimal visual system ought to interpret the generally ambiguous optical stimulation. Because of the ambiguity, the visual system often has to guess (infer) properties of stimulation from indirect cues. In Bayesian models, the inference is guided by the prior history of stimulation—the "prior distribution"—which represents agent's knowledge about the environment.[2] The optimal guess is the one that is most likely to yield a correct interpretation of the stimulus in light of this knowledge. Thus, the prior distribution used by Weiss et al. instantiates their assumptions that (1) low velocities prevail in the natural viewing conditions, and (2) the visual system takes into account the prevalence of low velocities. For example, when the illumination is low and the immediate stimulation provides little information for estimating velocities of moving objects, the visual system ought to rely on its prior knowledge that in most cases object velocities are low.

According to the theory of Weiss et al. (2002), the probabilistic inference should lead to correct perception under most conditions of stimulation. But because of the assumption that velocities are generally low, the visual system may sometimes underestimate velocity. Illusions arise, for instance, when a fast-moving object is presented under reduced illumination (as when you drive in fog). Weiss et al. and others (Hürlimann, Kiper, & Carandini, 2002; Stocker & Simoncelli, 2006) confirmed predictions of this model in psychophysical experiments with human observers.

The work by Weiss et al. (2002) supports the view that biological vision uses the prior history of stimulation to resolve perceptual ambiguities. This strategy is rational even though sometimes it yields incorrect percepts. Both the veridical (correct) and the non-veridical (illusory) perception predicted by the ideal observer manifest a rational biological system.

2.2. *Ventriloquism*

My second illustration is ventriloquism, a perceptual illusion known since antiquity. Ventriloquists produce voice without moving their lips, such that the audience attributes the source of voice to an object external to the ventriloquist: usually a puppet or dummy moved to simulate speech production.

Alais and Burr (2004) used the normative framework of optimal cue combination (Landy, Maloney, Johnston, & Young, 1995; Maloney & Landy, 1989; Yuille & Bülthoff, 1996) to investigate how humans estimate the location of sound from concurrent visual and auditory "cues." The model used by Alais and Burr (which I will describe in some detail in the next section) prescribed how the uncertain visual and auditory information ought to be combined to maximize the precision of estimated object location. By this prescription, contributions of single-modality (visual or auditory) estimates into the combined inter-sensory estimate depend on reliabilities (precisions) of the single-modality estimates.[3] Thus, whenever visual information is more reliable than auditory information, the combined visual-auditory estimate ought to depend more on vision than audition. The model also predicts the precision of combined estimates. In many cases, the expected precision of combined estimates is higher than the precision of visual and auditory estimates measured separately.

Alais and Burr (2004) presented human observers with two concurrent stimuli: visual (a blurred spot of light on a screen) and auditory (stereophonic sounds from two speakers on the sides of the screen). The degree of visual blur was varied: the more blur, the less precise the visual estimates of spot location, such that visual precision was sometimes better and sometimes worse than auditory precision. When visual precision was better, observers attributed the source of sound to the location of the spot of light (as in ventriloquism), in agreement with predictions of the model. The measured precision of combined visual-auditory localization was also in good agreement with the predictions.

The results of Alais and Burr (2004) are evidence that attribution of auditory signals to a source of visual signal is a rational behavior by a neural system that maximizes precision of sensory estimates. The illusion of ventriloquism showcases how neural computations that yield correct perception in a vast majority of cases can make mistakes in some cases, just as in the study of motion perception by Weiss et al. (2002).

2.3. *Fallible Optimality*

Perceptual illusions are spectacular examples of how optimization can lead to occasional errors by biological systems that otherwise are overwhelmingly reliable. The illusions are not unique in this respect. For example, sensory adaptation is thought to improve (presumably optimize) agent's performance in face of the variable environment (Barlow, 1990; Clifford & Wenderoth, 1999). Yet adaptation sometimes impairs sensory performance (e.g., Krekelberg, van Wezel, & Albright, 2006). Similarly, selective attention is thought to improve (presumably optimize) agent's performance in face of demanding tasks. But attention, too, can sometimes impair performance (e.g., Yeshurun & Carrasco, 1998, 2000).

Evidently, biological optimization occasionally backfires. Why? A part of the answer is the uncertainty that plagues both sensory and motor systems, due to the ambiguity of sensory stimulation (Knill & Richards, 1996) and due to the noise associated with planning and execution of even simple motor acts (Maloney, Trommershäuser, & Landy, 2007). This issue is taken on by probabilistic normative models of the sort we saw above. Another part of the answer—which still awaits a normative formulation—is the basic fact that neural computational resources are limited so agents cannot optimize performance for many tasks at once. Instead, performance is optimized for tasks that are most common or most pressing. This line of thought adds an economic dimension to the normative approach; ''economic'' in the sense of optimal allocation of scarce resources. The need for explicit modeling of allocation and reallocation of neural resources is generally recognized in studies of neural systems (e.g., Laughlin & Sejnowski, 2003; Lennie, 2003; Sakitt & Barlow, 1982; Sperling & Dosher, 1986; Stocker & Simoncelli, 2006), and new theoretical efforts are under way to include the economic considerations in the scope of normative theories of perception (e.g., Gepshtein, Tyukin, & Kubovy, 2007).

An external observer impressed by the sporadic errors of perception may be tempted to describe the perceiving agent as irrational, or ''externally irrational'' as Hardy-Vallée and Thagard (2008) do. But from the normative perspective presented herein, this assessment is incorrect. Here, behavior is classified as rational when it agrees with predictions of a normative model; behavior is classified as irrational (to a variable and measurable degree) when it disagrees with predictions, whether or not behavior is successful.

Notice that, from this normative perspective, rationality is attributed to the biological agent rather than to the frame of reference used for evaluation. In other words, rationality of biological agents or systems is their observer-independent property. Biological behavior may *appear* irrational to an external observer who uses an inappropriate normative framework, as illusions may appear irrational to an observer who fails to distinguish veridicality of perception from its rationality. Such failures render the observer incorrect rather than they render the observed behavior irrational.

Notice also that any behavior can be described as ''externally irrational'' because one can always approach it using an incomplete or inappropriate normative framework, in which the predicted performance will be superior to the performance of biological agents. The fact that every conceivable behavior can be gratuitously classified as ''externally irrational'' sheds further doubt on usefulness of the distinction between internal and external rationality.

3. Genesis of Normative Models

I have illustrated how recent work on perceptual illusions helped to shrink the gap between the predictions of normative models of perception and the actual human perception. It seems that this work has resolved the concern voiced by Hardy-Vallée and Thagard, that the earlier normative models (the "standard picture" of rationality) had ignored the effectiveness of norms. Indeed, results of experiments are now routinely compared with prescriptions of normative models, testing effectiveness of normative models. But is this work consistent with the engineering (Quinean) approach to normativity? More generally, do modern normative models of biological perception and action evolve according to an engineering prescription? To answer these questions I will now take a closer look at how candidate norms (i.e., hypotheses about biological norms) are nominated and tested. I will then compare the structure of reasoning used in normative studies with the structure of standard scientific reasoning.

3.1. Norms of Sensory Fusion

The aforementioned framework of optimal cue integration rests on the assumption that sensory systems seek to maximize precision of sensory estimates. The assumption is formalized using the principle of maximum likelihood, which is why these models are sometimes called Maximum Likelihood Estimation (MLE) ideal observers.[4] MLE observers take into account all the sensory information relevant to the task and thus achieve the highest possible precision (the lowest variability) of the combined estimates.

The normative framework of cue combination was first used to study how human vision combined information from different sources—"depth cues"—about the spatial layout of visual scenes (Landy et al., 1995; Young, Landy, & Maloney, 1993). Examples of such cues are texture gradient, motion parallax, and binocular parallax. This framework was later expanded to study interactions of vision with other sensory modalities. A result of this effort has been an increasingly general framework for research of sensory interactions. This work has had a constructive influence on the descriptive studies in that it compels researchers to revisit previous results and also seek new results, using more rigorous experimental methods than in the earlier work.

An example of how the normative project has a constructive influence on the descriptive project is the recent work on visual-haptic interaction.[5] Early descriptive studies showed that visual information often dominated perception when observers could concurrently use visual and haptic signals: observers' reports of object properties depended on visual rather than haptic information. This phenomenon was dubbed "visual capture" (Hay, Pick, & Ikeda, 1965; Rock & Victor, 1964). But some other studies showed that touch could dominate perception when information from touch was more appropriate for the task than information from vision.[6] The optimal cue-combination framework presented an opportunity to recast these results as special cases of a simple quantitative model—the MLE observer—in which the contribution of every sensory modality into the combined percept is weighted by modality's precision: a more precise modality contributes more than a less precise modality. Thus, the MLE observer predicted that sight should dominate perception when visual precision is higher than haptic precision, touch should dominate when haptic precision is higher than visual precision, and the contributions of modalities in the intermediate conditions are described by a predicted function. Predictions of the model were based on the estimates of precision of individual modalities, measured in separate experiments, making the predictive models completely constrained, i.e., leaving *no free parameters* for fitting the predictions to the data.

Psychophysical experiments confirmed predictions of the ideal observer for visual-haptic interaction: Ernst and Banks (2002) reduced precision of visual information using visual noise and found that the contribution of noisy visual information decreased by the amount predicted by the optimal model. Gepshtein and Banks (2003) asked whether visual and haptic information was combined differently under different viewing conditions, because the relative precision of vision and touch naturally depended on the viewing angle. The results showed that here too information from vision and touch was weighted as prescribed by the normative model. (The study of ventriloquism described in the previous section is an application of this approach to the combination of visual and auditory signals.)

It turned out, however, that the optimality of visual-haptic combination was restricted. Thus, Gepshtein, Burge, Ernst, and Banks (2005) compared how visual and haptic signals combined when the signals coincided in space and when they were spatially separate. When the signals coincided, perception followed the prescriptions of the ideal observer, as in Gepshtein and Banks (2003). In particular, the precision of combined visual-haptic percepts was better than precision by sight alone or touch alone. But when the signals were spatially separated, the optimal behavior broke down. The precision of combined percept was only as good as by sight alone or touch alone. In other words, it turned out that the human nervous system applied the optimal method of sensory fusion selectively, only for signals that coincided in space.

The findings that (a) visual-haptic combination was consistent with predictions of the ideal observer for coincident signals, and (b) the combination was inconsistent with predictions for non-coincident signals, could be interpreted as evidence of suboptimality of visual-haptic combination.

Yet the findings can be interpreted differently, using a broader normative framework, supporting the view that both coincident and non-coincident signals are combined optimally. To see that, note that signals that originate from the same object normally have a common spatial source. It is therefore plausible that a frugal nervous system will allocate its computational resources for sensory combination only where the combination is needed in the everyday behavior, rather than for signals that arrive from arbitrary spatial locations. This possibility is supported by recent evidence that visual and haptic signals were combined optimally when they were physically separated in space, but were presented to observers such that they appeared as parts of the same object (Helbig & Ernst, 2007; Takahashi, Diedrichsen, & Watt, 2008).

It is therefore plausible that the results of Gepshtein et al. (2005) about combination of non-coincident sensory signals manifest a frugal nervous system which is specialized for perception of objects and which allocates its computational resources optimally. In other words, the improved precision in coincident signals and the lack of improvement in non-coincident signals are both manifestations of optimal behavior.

In spirit of the above normative approach, this possibility can be tested further as follows:

1. A mathematical expression must be derived describing how spatial distances between signals relate to the likelihood that signals belong to the same object. (I will call this expression an "objecthood function.") The objecthood function can be found using the rules of projective geometry and the statistical measurements of optical stimulation.
2. Then, the objecthood function must be incorporated in an MLE model so one can derive new predictions and test them. An agreement between the predictions and results of experiments would support the broader normative model.
3. Importantly, the broader normative model would hold only if the objecthood function is *general*, i.e., if it is consistent with experimental results in contexts other than visual-haptic combination, such as in studies of spatial perceptual organization. If alternative objecthood functions exist, then quantitative comparison of the alternative predictions is in order.

3.2. Normative Reasoning

The above review of visual-haptic combination is an illustration of how candidate normative models of biological perception are advanced and tested. To see the structure of reasoning used in this work, recall a form of inference called *abductive reasoning*. Sometimes it is called "abductory" or "retroductive" reasoning. Abductive reasoning was first described by Charles S. Pierce who defined it as follows:

> The surprising fact, C, is observed;
> But if A were true, C would be a matter of course,
> Hence, there is reason to suspect that A is true. (Buchler, 1955, p. 151)

Peirce held that the standard method of scientific discovery constitutes a combination of abductive and inductive reasoning. Suppose that term A in the above syllogism leads to predictions other than C (along with C itself), such that the "reason to suspect that A is true" can be tested in separate experiments. This way new hypotheses are formulated and tested. The latter step—of experimental hypothesis testing—is an instance of standard inductive reasoning.

The logic of normative work on sensory fusion is readily explained using the Peircean argument as follows. The normative theorist is aware of the empirical results from studies of cue combination (term C in Peirce's syllogism). The theorist realizes that if the human nervous system implemented the maximum-likelihood computation (term A) then many empirical results from previous studies (term C) would follow necessarily. This step is that of abductive reasoning. The next step is inductive. The hypothesis that the nervous system implements maximum-likelihood computations leads to new predictions (e.g., that precision of combined estimates can be higher than precision of component estimates) which are then tested in experiments.

Similarly, the hypothesis that sensory systems are specialized for perception of objects is derived from the facts that are "surprising," i.e., cannot be explained by theories that lack a notion of objecthood. This line of thought can be tested as proposed in the previous section, through a combination of abduction (the hypothesis of optimal recourse allocation) and induction (the general testing of the objecthood function), in accord with the Peircean view of standard scientific reasoning.

3.3. Scientific and Engineering Normativity

Two aspects of norm development are important for the task of contrasting the scientific and engineering views of normativity: generality of normative models and their principledness:

- *Generality*: a candidate normative model is accepted when its predictions are confirmed in experiments. The experiments test predictions both narrowly, in the domain for which the model is first proposed, and broadly, for other consequences of the model. Importantly, if the broad test fails, the model is rejected or revised.

  For example, the framework of optimal cue integration has been testes in studies of many sensory cues: first within and then between sensory modalities. Several predictions were tested: concerning the magnitudes of sensory estimates and also the distributions of errors of estimates. The testing was rigorous: all model parameters were measured in experiments, leaving no free parameters for fitting model predictions to the data.

- *Principledness*: normative models incorporate assumptions about the goals and constraints of studied systems as formal statements. The formalism is reduced to basic principles, such as the principle of maximum likelihood in work on cue integration, or probability theorems in work on perceptual inference.[7]

In contrast, neither generality nor principledness is essential for the engineering approach to normativity. Hardy-Vallée and Thagard (2008) make that clear:

> Norms are justified by their effectiveness. If a procedure succeeds best in attaining a particular goal in a certain context, it is therefore, a normatively correct procedure in that context. (p. 177)

The place of normativity in engineering has been clarified by Houkes (2002) who argued that a "minimal condition for calling activities 'engineering'" is their focus on design of useful artifacts ("engineers do not merely . . . develop theories or perform calculations, but they design and modify artificial structures, such as bridges, airports, or banana flavorings" [Houkes, 2002, p. 260]). The focus on design of useful artifacts determines the nature of normativity in engineering:

> Technical artifacts can be described in two ways: from one perspective, they are physical structures; from the other, they are functional objects that stand in a relation to other objects and human practices. The normativity inherent in engineering may only be understood from this dual description or nature of artifacts: the rules made by the engineer specify proper use of the artifact in common circumstances, and 'proper use' can be at least partially explained by

> means of the physical structure. One should not drive over a bridge with a truckload of concrete piers, given the maximal stress on the structure; one should not put aluminum containers in a microwave oven, given the high reflectance rate of the metal. (Houkes, 2002, p. 260)

In other words, the engineer's goal is to effectively manage physical properties of technological artifacts. This goal requires mastery of many tools used by the scientist, but the scientist and the engineer use these tools differently.

Because of the emphasis on effectiveness, the engineering approach to normativity is similar to fitting statistical models to experimental data. A model is judged successful when it fits the data closely. Scientists, too, seek to develop theories and models that fit ("explain") the data, but here the fit of model and data is achieved by means other than in engineering. Close fit alone is not a mark of success. In fact, the scientific community tends to reject normative theories derived according to the engineering prescription in which only effectiveness matters. For example, in Bayesian modeling of perception it is expected that the prior distribution of estimates is justified on grounds other than a model's good fit to the data. If such justification is not available, the model is denounced as unprincipled; it is viewed as a statistically sophisticated descriptive model of the data rather than a candidate for inclusion in a broader normative theory. It is therefore misleading to think of scientific theories as artifacts designed by the scientist to fit experimental data, similar to how bridges are designed by the engineer to control traffic and fruit flavorings are designed to imitate taste of fruit.

3.4. Closing Miller's Gap

Why did the gap emerge between the norms of behavior and the actual behavior in the "standard picture" of human rationality? I hope the above illustrations made it clear that the reason is not an inadequacy of the scientific approach to normativity. Miller's gap has shrunk because our understanding of the constraints of biological systems has improved, and we have improved the methods used for evaluating behavior. Thus, probabilistic models mentioned above have been taking into account the uncertainties implicit to everyday perception, such that laboratory tests of model predictions engaged the same biological processes that are normally engaged in the natural behavior outside of the laboratory. This approach implements the influential notion of "ecological validity" of Brunswick (1955) who anticipated the current interest in how neural systems adapt to statistical properties of the environment. On this view, early normative models of human rationality failed to explain behavior because they disregarded the basic fact that behavior is optimized to its natural context.

The importance of proper context for evaluating rationality in biological systems is underscored by recent studies of human action planning under uncertainty. These studies have shown that humans can achieve nearly optimal motor performance in face of uncertainties implicit to planning and execution of simple movements (e.g., Trommershäuser, Maloney, & Landy, 2003). The modeled uncertainties have increasingly approximated those in the natural behavior (Gepshtein, Seydell, & Trommershäuser, 2007). In contrast, early studies of rationality presented subjects with arbitrary tasks, offering no justification for the probabilities attached to particular outcomes of behavior. For example, Maloney et al. (2007) compared human performance under uncertainties implicit to everyday motor behavior (such as hand reaching) with performance in cognitive "paper-and-pencil" tasks of the traditional decision-making literature. The authors concluded that results of the latter are not representative of the former, such that

> Human capacity for decision making bears the same relation to the economic tasks of classical decision making as human language competence bears to solving the Sunday crossword puzzle. (Maloney et al., 2007, p. 312)

In other words, closing Miller's gap requires that not only the norms undergo naturalization, but so also does the context in which the norms are evaluated.

4. Conclusions

Both scientific and engineering approaches to normativity strive to develop models that reliably predict behavior (i.e., are efficient) and both rest on general principles of mathematics and physics (i.e., both can be principled). But the priorities of two approaches are different. The priority of engineering is efficiency. The priorities of science are generality and principledness. Because of the focus on efficiency, a root metaphor for engineering accounts of behavior is statistical model. This approach encourages the engineer to use different models to fit different data, yielding multiple narrow models as the scope of explananda grows. In contrast, the scientist seeks to ground models of behavior in a small number of basic principles, such that a few principles explain as wide range of data as possible (*lex parsimoniae*). A root metaphor for scientific account of behavior is natural law. Because of the burden of *lex parsimoniae*, scientific models are bound to develop slower but remain simpler than the engineering models. Thanks to this simplicity, the scientific approach is likely to give us more comprehensive (and more comprehensible) theories of biological behavior than the engineering approach.

Acknowledgements

I wish to thank John Jacobson, Michael Kubovy, and Cees von Leeuwen for valuable comments, and the National Institute of Natural Sciences of Japan and the Swartz Foundation for generous financial support.

Notes

[1] One reason visual measurement are inherently ambiguous is the optical projection from the three-dimensional scenes to the effectively two-dimensional retinal surface, such that the

same retinal image may correspond to different stimuli. An example of noisy biological computation is the absorption of light by retinal receptors, which is a stochastic process.

[2] This knowledge can be implicit, e.g., implemented in automatic neural computations, possibly inaccessible to awareness. This knowledge is usually described by saying that nervous systems "take into account" or "represent" regularities of the environment.

[3] The notion of reliability has a technical meaning; it is defined as the inverse of the variance of the distribution of estimates. Precision of estimation is defined as the standard deviation of this distribution.

[4] The notion of "ideal observer" can be used narrowly or broadly. In the narrow sense, ideal observer models disregard the constraints of biological computation, so the researcher may compare performance of a biological system with a mathematical ideal. In the broad sense, ideal observer models incorporate *some* biological constraints, such as the decision noise in the Statistical Decision Theory (Geisler, 1989; Green & Swets, 1966). Here I use the latter approach.

[5] "Haptic" information is sensory information gathered through active touch, using tactile and proprioceptive signals.

[6] For example, touch is more reliable than vision for estimating roughness of surfaces (Lederman & Abbott, 1981).

[7] This is not to say that the cue-integration models and the Bayesian inference models are incompatible. On the contrary, the two kinds of models belong to the same decision-theoretic framework and are readily combined (e.g., Hillis, Watt, Landy, & Banks, 2004).

References

Alais, D., & Burr, D. (2004). The ventriloquist effect results from near-optimal bimodal integration. *Current Biology*, *14*, 257–262.

Barlow, H.B. (1990). A theory about the functional role and synaptic mechanism of visual after-effects. In C. Blakemore (Ed.), *Vision: Coding, and efficiency* (pp. 363–375). Cambridge: Cambridge University Press.

Blakemore, C. (1990). *Vision: Coding, and efficiency*. Cambridge: Cambridge University Press.

Brunswik, E. (1955). Representative design and probabilistic theory in a functional psychology. *Psychological Review*, *62*, 193–221.

Buchler, J. (Ed). (1955). *Philosophical writings of Peirce*. New York: Dover Publications.

Clifford, C.W., & Wenderoth, P. (1999). Adaptation to temporal modulation can enhance differential speed sensitivity. *Vision Research*, *39*, 4324–4332.

Ernst, M.O., & Banks, M.S. (2002). Humans integrate visual and haptic information in a statistically optimal fashion. *Nature*, *415*, 429–433.

Geisler, W.S. (1989). Sequential ideal-observer analysis of visual discrimination. *Psychological Review*, *96*, 267–314.

Geisler, W.S., & Kersten, D. (2002). Illusions, perception and Bayes. *Nature Neuroscience*, *5*, 508–510.

Gepshtein, S., & Banks, M.S. (2003). Viewing geometry determines how vision and haptics combine in size perception. *Current Biology*, *13*, 483–488.

Gepshtein, S., Burge, J., Ernst, M.O., & Banks, M.S. (2005). The combination of vision and touch depends on spatial proximity. *Journal of Vision*, *5*(11), 1013–1023.

Gepshtein, S., Seydell, A., & Trommershäuser, J. (2007). Optimality of human movement under natural variations of visual-motor uncertainty. *Journal of Vision*, *7*(5), 1–18.

Gepshtein, S., Tyukin, I., & Kubovy, M. (2007). The economics of motion perception and invariants of visual sensitivity. *Journal of Vision*, *7*(8), 1–18.

Green, D.M., & Swets, J.A. (1966). *Signal detection theory and psychophysics*. New York: Wiley.

Hardy-Vallée, B., & Thagard, P. (2008). How to play the ultimatum game: An engineering approach to metanormativity. *Philosophical Psychology*, *21*(2), 173–192.

Hay, J.C., Pick Jr, H.L., & Ikeda, K. (1965). Visual capture produced by prism spectacles. *Psychonomic Science*, *2*, 215–216.

Helbig, H.B., & Ernst, M.O. (2007). Optimal integration of shape information from vision and touch. *Experimental Brain Research*, *179*(4), 595–606.

Hillis, J.M., Watt, S.J., Landy, M.S., & Banks, M.S. (2004). Slant from texture and disparity cues: optimal cue combination. *Journal of Vision*, *4*(12), 967–992.

Houkes, W.N. (2002). Normativity in Quine's naturalism: The technology of truth-seeking?. *Journal for General Philosophy of Science*, *33*, 251–267.

Hürlimann, F., Kiper, D., & Carandini, M. (2002). Testing the Bayesian model of motion perception. *Vision Research*, *42*, 2253–2257.

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

Krekelberg, B., van Wezel, R.J., & Albright, T. D. (2006). Adaptation in macaque MT reduces perceived speed and improves speed discrimination. *Journal of Neurophysiology*, *95*, 255–270.

Landy, M.S., Maloney, L.T., Johnston, E.B., & Young, M. (1995). Measurement and modeling of depth cue combination: In defense of weak fusion. *Vision Research*, *35*, 389–412.

Laughlin, S.B., & Sejnowski, T.J. (2003). Communication in neuronal networks. *Science*, *301*, 1870–1874.

Lederman, S.J., & Abbott, S.G. (1981). Texture perception: Studies of intersensory organization using a discrepancy paradigm, and visual versus tactual psychophysics. *Journal of Experimental Psychology: Human Perception and Performance*, *7*, 902–915.

Lennie, P. (2003). The cost of cortical computation. *Current Biology*, *13*, 493–497.

Maloney, L.T. (2002). Statistical decision theory and biological vision. In D. Heyer & R. Mausfeld (Eds.), *Perception and the physical world* (pp. 145–189). New York: Wiley.

Maloney, L.T., & Landy, M.S. (1989). A statistical framework for robust fusion of depth information. In W.A. Pearlman (Ed.), *Visual communications and image processing IV* (Vol. 1199, pp. 1154–1163). Proceedings of Society of Photo-optical Instrumentation Engineers (SPIE).

Maloney, L.T., Trommershäuser, J., & Landy, M.S. (2007). Questions without words: A comparison between decision making under risk and movement planning under risk. In W. Gray (Ed.), *Integrated models of cognitive systems* (pp. 297–313). New York: Oxford University Press.

Miller, G.A. (1964). *Mathematics and psychology*. New York: Wiley.

Quine, W.V.O. (1986). Reply to Morton White. In L.E. Hahn & P.A. Schilpp (Eds.), *The philosophy of W. V. Quine* (pp. 663–665). La Salle: Open Court.

Rao, R.P.N., Olshausen, B.A., & Lewicki, M.S. (2002). *Probabilistic models of the brain: Perception and neural function*. Cambridge, MA: MIT Press.

Rock, I., & Victor, J. (1964). Vision and touch: An experimentally created conflict between the two senses. *Science*, *143*, 594–596.

Sakitt, B., & Barlow, H.B. (1982). A model for the economical encoding of the visual image in cerebral cortex. *Biological Cybernetics*, *43*, 97–108.

Sperling, G., & Dosher, B.A. (1986). Strategy and optimization in human information processing. In K. Boff, L. Kaufman, & J. Thomas (Eds.), *Handbook of perception and performance,* 1, (pp. 2–65). New York: Wiley.

Stocker, A. A., & Simoncelli, E. P. (2006). Noise characteristics and prior expectations in human visual speed perception. *Nature Neuroscience*, *9*, 578–585.

Takahashi, C., Diedrichsen, J., & Watt, S. (2008). The brain integrates visual and haptic information from different spatial locations when using a tool. *Journal of Vision*, *8*(6), 1060.

Trommershäuser, J., Maloney, L.T., & Landy, M.S. (2003). Statistical decision theory and the selection of rapid, goal-directed movements. *Journal of the Optical Society of America A, Optics, image science, and vision*, *20*, 1419–1433.

Weiss, Y., Simoncelli, E.P., & Adelson, E.H. (2002). Motion illusions as optimal percepts. *Nature Neuroscience*, *5*(6), 598–604.

Yeshurun, Y., & Carrasco, M. (1998). Attention improves or impairs visual performance by enhancing spatial resolution. *Nature*, *396*, 72–75.

Yeshurun, Y., & Carrasco, M. (2000). The locus of attentional effects in texture segmentation. *Nature Neuroscience*, *3*, 622–627.

Young, M.J., Landy, M.S., & Maloney, L.T. (1993). A perturbation analysis of depth perception from combinations of texture and motion cues. *Vision Research*, *33*, 2685–2696.

Yuille, A.L., & Bütlhoff, H.H. (1996). Bayesian decision theory and psychophysics. In D.C. Knill & W. Richards (Eds.), *Perception as Bayesian inference* (pp. 123–161). Cambridge: Cambridge University Press.