Lawful relation between perceptual bias and discriminability

Xue-Xin Wei[a,b,c] and Alan A. Stocker[d,e,1]

[a]Grossman Center for the Statistics of Mind, Columbia University, New York, NY 10027; [b]Center for Theoretical Neuroscience, Columbia University, New York, NY 10027; [c]Department of Statistics, Columbia University, New York, NY 10027; [d]Department of Psychology, University of Pennsylvania, Philadelphia, PA 19104; and [e]Department of Electrical and Systems Engineering, University of Pennsylvania, Philadelphia, PA 19104

Edited by Wilson S. Geisler, University of Texas at Austin, Austin, TX, and approved August 4, 2017 (received for review November 21, 2016)

Perception of a stimulus can be characterized by two fundamental psychophysical measures: how well the stimulus can be discriminated from similar ones (discrimination threshold) and how strongly the perceived stimulus value deviates on average from the true stimulus value (perceptual bias). We demonstrate that perceptual bias and discriminability, as functions of the stimulus value, follow a surprisingly simple mathematical relation. The relation, which is derived from a theory combining optimal encoding and decoding, is well supported by a wide range of reported psychophysical data including perceptual changes induced by contextual modulation. The large empirical support indicates that the proposed relation may represent a psychophysical law in human perception. Our results imply that the computational processes of sensory encoding and perceptual decoding are matched and optimized based on identical assumptions about the statistical structure of the sensory environment.

perceptual behavior | efficient coding | Bayesian observer | Weber–Fechner | stimulus statistics

Perception is a subjective experience that is shaped by the expectations and beliefs of an observer (1). Psychophysical measures provide an objective yet indirect characterization of this experience by describing the dependency between the physical properties of a stimulus and the corresponding perceptually guided behavior (2).

Two fundamental measures characterize an observer's perception of a stimulus. Discrimination threshold indicates the sensitivity of an observer to small changes in a stimulus variable (Fig. 1*A*). The threshold depends on the quality with which the stimulus variable is represented in the brain (2) (i.e., encoded; Fig. 1*B*); a more accurate representation results in a lower discrimination threshold. In contrast, perceptual bias is a measure that reflects the degree to which an observer's perception deviates on average from the true stimulus value. Perceptual bias is typically assumed to result from prior beliefs and reward expectations with which the observer interprets the sensory evidence (1), and thus is determined by factors that seem not directly related to the sensory representation of the stimulus. As a result, it has long been believed that there is no reason to expect any lawful relation between perceptual bias and discrimination threshold (3).

However, here we derive a direct mathematical relation between discrimination threshold $D(\theta)$ and perceptual bias $b(\theta)$ based on a recently proposed observer theory of perception (4, 7). The key idea is that both the encoding and the decoding process of the observer are optimally adapted to the statistical structure of the perceptual task (Fig. 2*A*). Specifically, we assume encoding to be efficient (8, 9) such that it maximizes the information in the sensory representation about the stimulus given a limit on the overall available coding resources (10). The assumption implies a sensory representation whose coding resources are allocated according to the stimulus distribution $p(\theta)$. This results in the encoding constraint

$$p(\theta) \propto \sqrt{J(\theta)}, \qquad \textbf{[1]}$$

where the Fisher information $J(\theta)$ represents the coding accuracy of the sensory representation (4, 11–13). Fisher information provides a lower bound on the discrimination threshold irrespective of whether the estimator is biased or not, which can be formulated as $D(\theta) \geq c/\sqrt{J(\theta)}$, where c is a constant (5, 6). Rather than assuming a tight bound, we make the weaker assumption that the bound is equally "loose" over the range of the stimulus value θ. Using the encoding constraint (Eq. **1**) above, we can express discrimination threshold in terms of the stimulus distribution as

$$D(\theta) \propto 1/p(\theta). \qquad \textbf{[2]}$$

Similarly, we have shown that the perceptual bias of the Bayesian observer model (Fig. 2*A*) can also be expressed in terms of the stimulus distribution (4, 7). Assuming that uncertainty in the perceptual process is dominated by internal (neural) noise and that the noise is relatively small, we can analytically derive the observer's bias as

$$b(\theta) \propto (1/p(\theta)^2)'. \qquad \textbf{[3]}$$

We can show that the expression holds independently of the details of the assumed loss function for a large family of symmetric loss functions (*Supporting Information*). Magnitude and sign of its proportionality coefficient, however, depend on several factors, including the noise magnitude and the loss function. Finally, by combining Eqs. **2** and **3,** we obtain a direct functional relation between perceptual bias and discrimination threshold in the form of

$$b(\theta) \propto (D(\theta)^2)'; \qquad \textbf{[4]}$$

Significance

We present a law of human perception. The law expresses a mathematical relation between our ability to perceptually discriminate a stimulus from similar ones and our bias in the perceived stimulus value. We derived the relation based on theoretical assumptions about how the brain represents sensory information and how it interprets this information to create a percept. Our main assumption is that both encoding and decoding are optimized for the specific statistical structure of the sensory environment. We found large experimental support for the law in the literature, which includes biases and changes in discriminability induced by contextual modulation (e.g., adaptation). Our results imply that human perception generally relies on statistically optimized processes.

Author contributions: X.-X.W. and A.A.S. designed research, performed research, analyzed data, and wrote the paper.

The authors declare no conflict of interest.

This article is a PNAS Direct Submission.

[1]To whom correspondence should be addressed. Email: astocker@psych.upenn.edu.

This article contains supporting information online at www.pnas.org/lookup/suppl/doi:10.1073/pnas.1619153114/-/DCSupplemental.

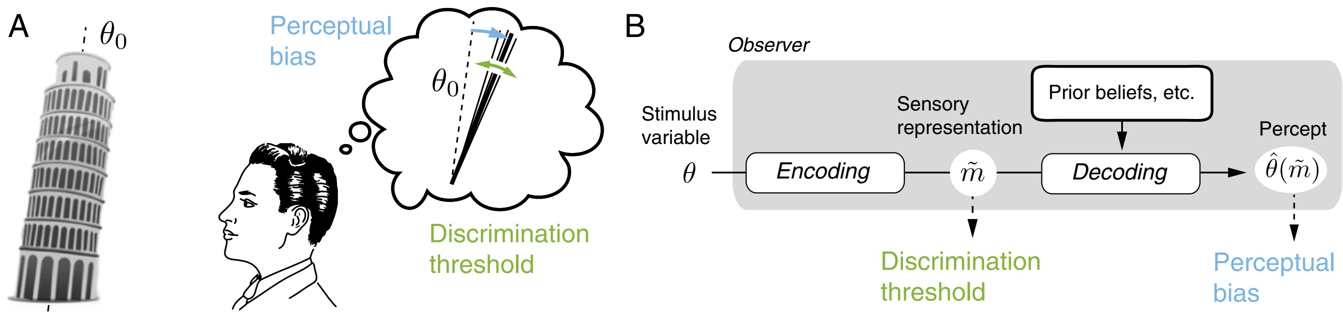

Fig. 1. Psychophysical characterization and modeling of perception. (*A*) Perception of a stimulus variable (e.g., the angular tilt θ_0 of the leaning tower of Pisa) can be characterized by an observer's discrimination threshold and perceptual bias. Discrimination threshold specifies how well an observer can discriminate small deviations around the particular stimulus orientation θ_0 given that there is noise in perceptual processing (green arrow). Perceptual bias specifies how much, on average, the perceived orientations over repeated presentations (thin lines) deviate from the true stimulus orientation (blue arrow). (*B*) Modeling perception as an encoding–decoding processing cascade. Discriminability is limited by the characteristics of the encoding process, i.e., the quality of the internal, sensory representation $\tilde{m}$. Perceptual bias, however, also depends on the decoding process that typically involves cognitive factors such as prior beliefs and reward expectations. Both discrimination threshold and perceptual bias can be characterized with appropriate psychophysical methods (indicated by dashed arrows).

i.e., perceptual bias (as a function of the stimulus variable) is proportional to the slope of the square of the discrimination threshold (Fig. 2*B*).

We tested the surprisingly simple relation against a wide range of existing psychophysical data. Figs. 3 and 4 show data for those perceptual variables for which both discrimination threshold and perceptual bias have been reported over a sufficiently large stimulus range. We grouped the examples according to their characteristic bias–threshold patterns. The first group consists of a set of circular variables (Fig. 3 *A–C*). It includes local visual orientation, probably the most well-studied perceptual variable. Orientation perception exhibits the so-called oblique effect (38), which describes the observation that the discrimination threshold peaks at the oblique orientations yet is lowest for cardinal orientations (14). Based on the oblique effect, Eq. **4** predicts that perceptual bias is zero at, and only at, both cardinal and oblique orientations. Measured bias functions confirm this prediction (15). Other circular variables that exhibit similar patterns are heading direction using either visual or vestibular information (16), 2D motion direction measured with a two alternative forced-choice (2AFC) procedure (17, 18) or by smooth pursuit eye movements (19), pointing direction (20), and motion direction in depth (21, 22). The relation also holds for the more high-level perceptual variable of perceived heading direction of approaching biological motion (human pedestrian) (23) as shown in Fig. 3*C*.

The second group contains noncircular magnitude variables for which discrimination threshold (approximately) follows Weber's law (24) and linearly increases with magnitude (Fig. 3*D*). We predict that these variables should exhibit a perceptual bias that is also linear in stimulus magnitude. Indeed, we found this to be true for spatial frequency [threshold (14, 34, 39), bias (25)] as well as temporal frequency [threshold (40), bias (27)]

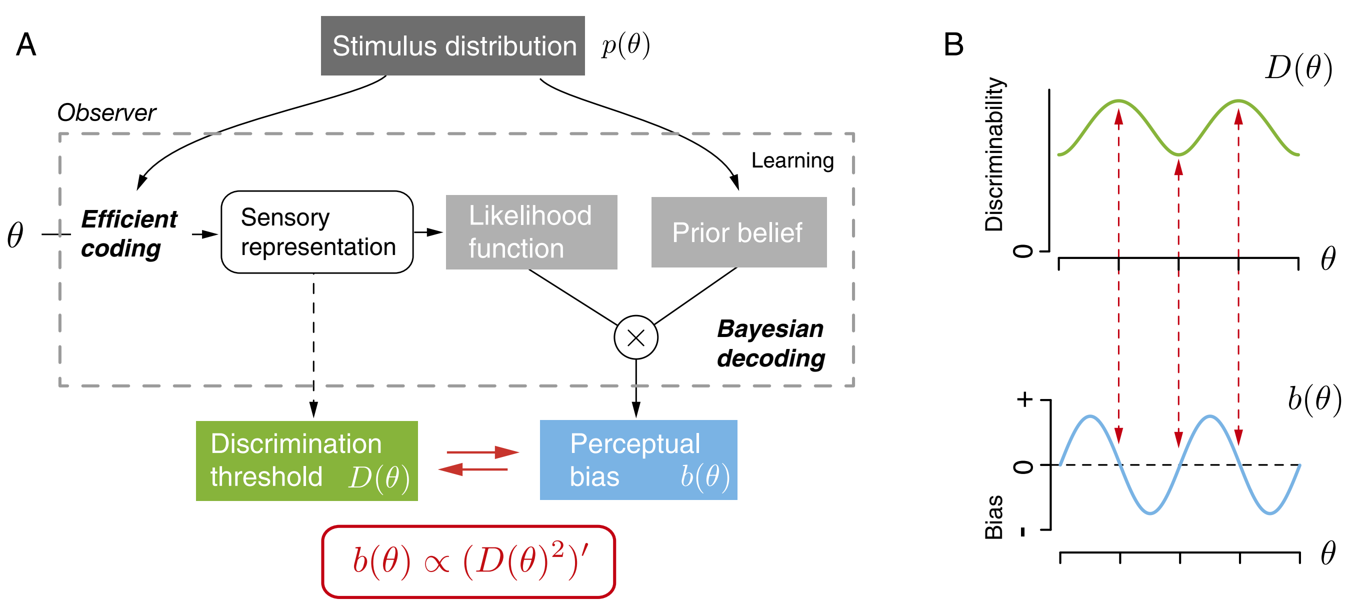

Fig. 2. Observer model that links perceptual bias and discrimination threshold. (*A*) Our theory of perception proposes that encoding and decoding are both optimized for a given stimulus distribution (4). Based on this theory, the encoding accuracy characterized by Fisher information $J(\theta)$ and the bias $b(\theta)$ of the Bayesian decoder are both dependent on the stimulus distribution $p(\theta)$. With Fisher information providing a lower bound on discriminability $D(\theta)$ (5, 6), we can mathematically formulate the relation between perceptual bias and discrimination threshold as $b(\theta) \propto (D(\theta)^2)'$. (*B*) Arbitrarily chosen example highlighting the characteristics of the relation: Bias is zero at the extrema of the discrimination threshold (red arrows) and largest for stimulus values where the threshold changes most rapidly. Thus, the magnitude of perceptual bias typically does not covary with the magnitude of the discrimination threshold.

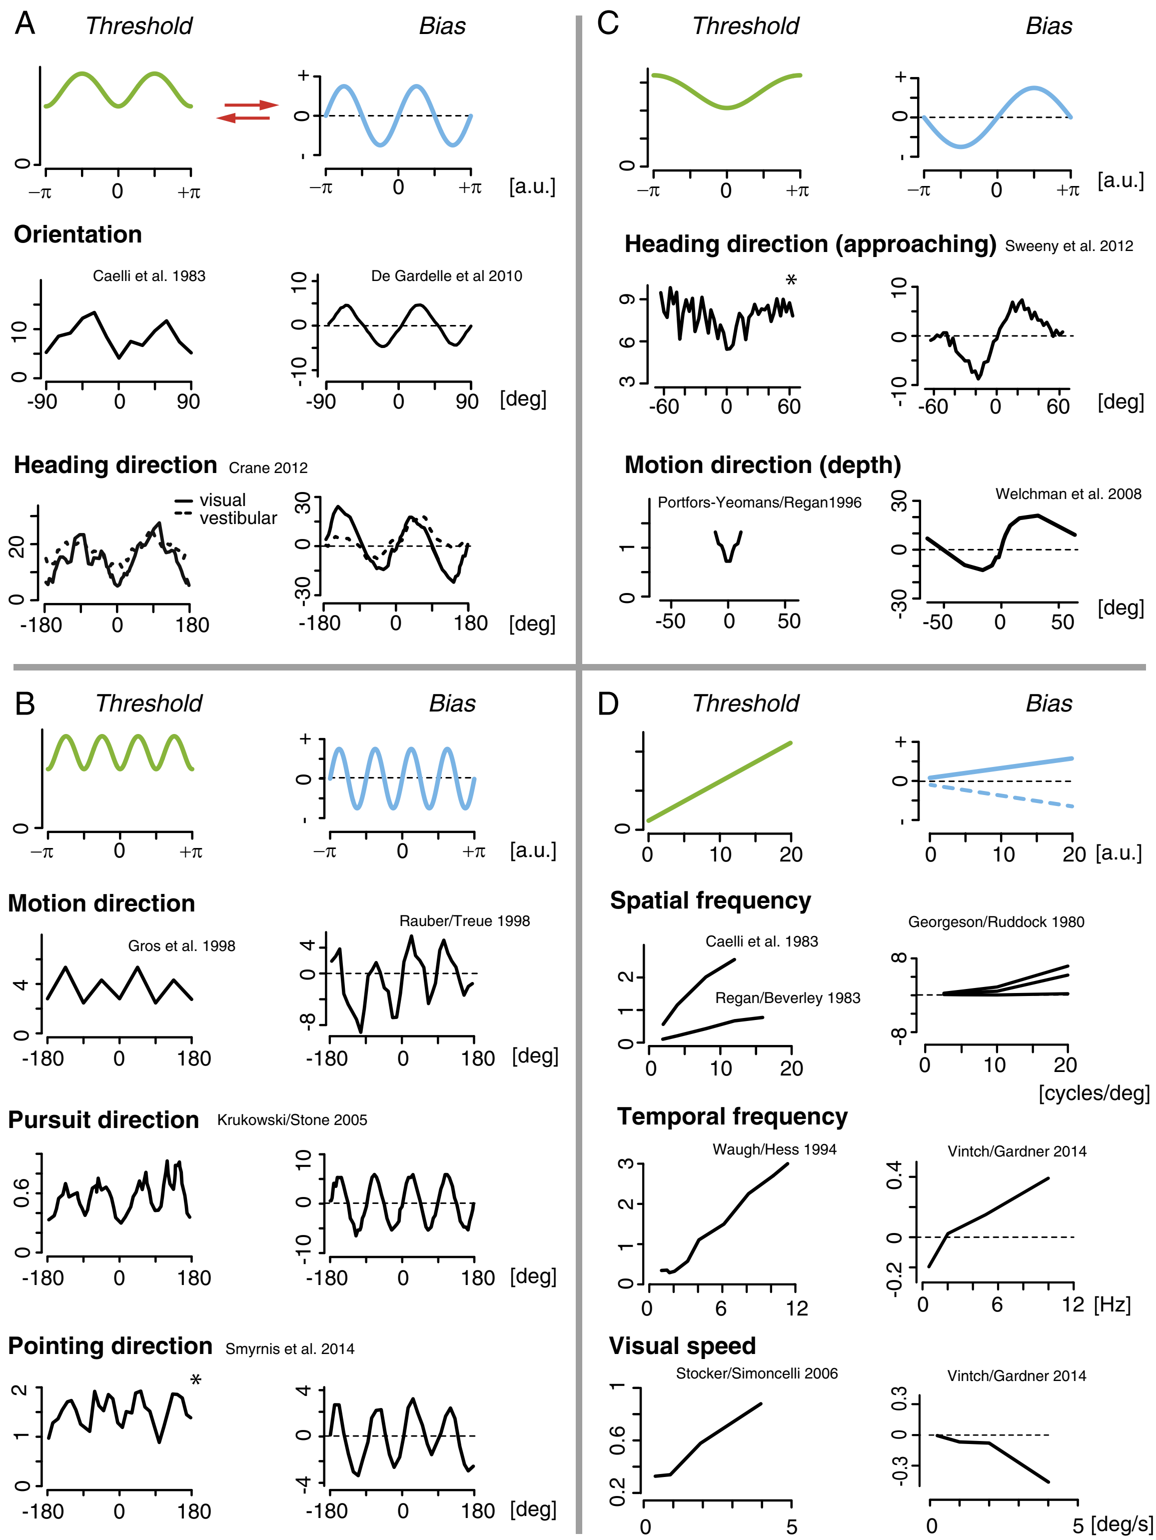

Fig. 3. Predicted and measured bias–threshold patterns in perception. Data are organized into different groups. Green/blue curves represent the predicted discrimination threshold/bias. (*A*) Discrimination threshold (14) and bias (15) data for perceived orientation and heading direction (16) (solid lines for visual stimulation, and dashed lines for vestibular stimulation). (*B*) We found similar patterns for perceived motion direction [measured both with a 2AFC procedure (17, 18) or with smooth pursuit behavior (19)] and pointing direction [where subjects had to estimate the direction of a visually presented arrow (20)]. (*C*) The predicted relation also holds for perceived motion direction in depth (21, 22), as well as for the perception of higher-level stimuli such as the approaching heading direction of a person (23), although the quality of the available data is limited. (*D*) The bias–threshold relation is different for magnitude (i.e., noncircular) variables: Discrimination threshold is typically proportional to the stimulus value [Weber's law (24)], and thus we predict that perceptual bias is also linear in stimulus value. Reported patterns for visually perceived spatial frequency (14, 25, 34), temporal frequency (26, 27), and grating speed (27, 28) match this prediction. All data curves are replotted from their corresponding publications, except bias data for orientation (15) and threshold data for visual speed (28), which both were derived by analyzing the original data. Stars (*) mark data that represent perceptual variability rather than discrimination thresholds (see also Figs. S2 and S3 for more details).

in vision. Visual speed is another example for which discrimination threshold approximately follows Weber's law (28) and bias is also approximately linear with stimulus speed (27), although, in contrast to the other examples, with a negative proportionality coefficient. A possible explanation for the sign difference may be that speed perception is governed by a loss function that differs from the loss functions for the other variables. In perceiving the speed of a moving object, an observer's emphasis might be on estimating the speed "just right" to maximize the chance to, e.g., successfully intercept the object. A loss function that approximates the "all-or-nothing" characteristics of the L_0 norm would serve that goal, and would also predict the negative sign of the proportionality coefficient (see *Supporting Information*). Finally, perceived weight, one of the classical examples for illustrating Weber's law (24), also seems to be consistent with the proposed bias–threshold pattern (41).

The last group contains bias–threshold patterns that are not intrinsic to individual specific variables but are induced by contextual modulation (Fig. 4). Spatial context as in the tilt illusion can induce characteristic repulsive biases in the perceived stimulus orientation away from the orientation of the spatial surround (30). The corresponding change in discrimination threshold (29) well matches the predicted pattern based on our theoretically derived relation. A similar bias–threshold pattern has been reported when probing motion direction instead of orientation (31). Furthermore, similar patterns have been observed for temporal context, i.e., as the result of adaptation. Adaptation-induced biases and changes in discrimination threshold for perceived visual orientation (32, 33) and spatial frequency (34, 35) are qualitatively in agreement with the prediction. At a slightly longer time scale, perceptual learning is also known to reduce discrimination thresholds. We therefore predict that perceptual learning also induces repulsive biases away from the learned stimulus value. This prediction is indeed confirmed by data for learning orientation (36) and motion direction (42), albeit the existing data are sparse. Finally, attention has been known as a mechanism that can decrease discrimination threshold (43). We predict that this decrease should coincide with a repulsive bias in the perceived stimulus variable. Although limited in extent, data from a Vernier gap size estimation experiment are in agreement with this prediction (37) and are further supported by recent results (44). In sum, the derived relation can readily explain a wide array of empirical observations across different perceptual variables, sensory modalities, and contextual modulations.

Based on the strong empirical support, we argue that we have identified a universal law of human perception. It provides a unified and parsimonious characterization of the relation between discrimination threshold and perceptual bias, which are the two main psychophysical measures characterizing the perception of a stimulus variable. Only a very few quantitative laws are known in perceptual science, including Weber–Fechner's law (2, 24) and Stevens' law (45). These laws express simple empirical regularities which provide a compact yet generally valid description of the data. The law we have proposed here shares the same virtue. However, unlike these previous laws, our law is not the result of empirical observations but rather was derived based on theoretical considerations of optimal encoding and decoding (4). Thus, as such, it does not merely describe perceptual behavior but rather reflects our understanding of why perception exhibits such characteristics in the first place. Note that, without the theory, we would not have discovered the general empirical regularity between discrimination threshold and bias. It is conceivable that some of the theoretical assumptions we made in deriving the law may prove incorrect (see Fig. S1) despite the fact that the law itself is empirically well supported. It is difficult to imagine, however, how a lawful relation between perceptual bias and discrimination threshold could emerge without a functional

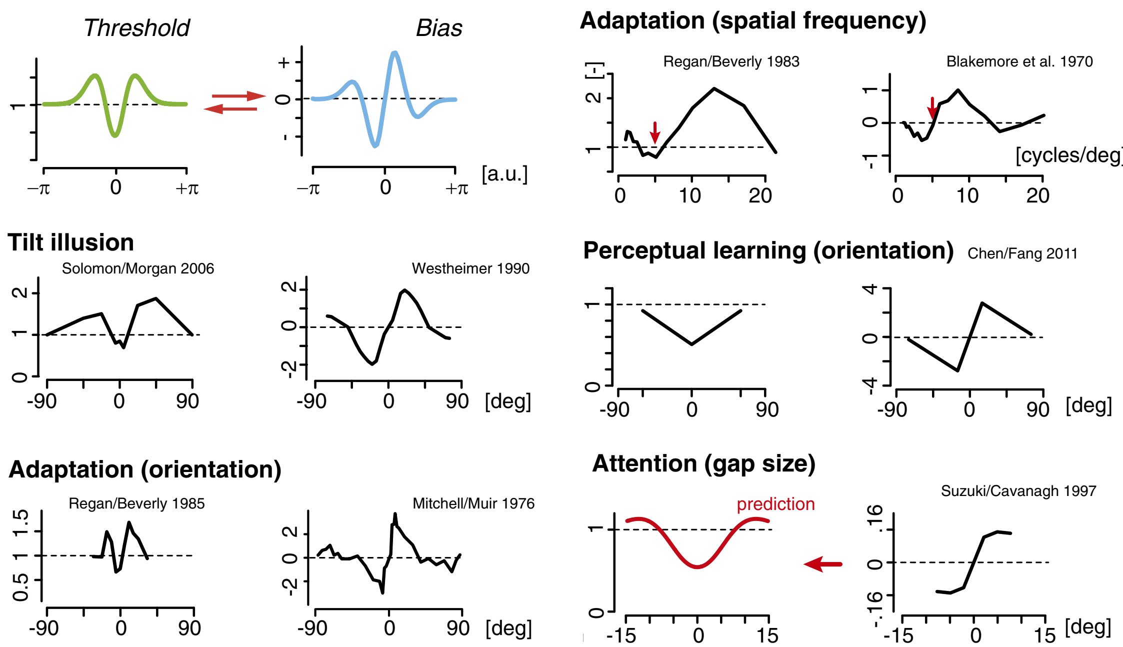

Fig. 4. Predicted and measured bias–threshold patterns induced by various forms of contextual modulation. Spatial context as in the tilt illusion (29, 30) [or using motion stimuli (31)] and temporal context during adaptation experiments [orientation (32, 33), relative to adaptor orientation; spatial frequency (34, 35), red arrows indicating the value of the adaptor stimulus] induce similar biases and changes in thresholds that are qualitatively well predicted (relative to the unadapted condition). The predicted relation seems to also hold for perceptual changes induced by perceptual learning (36) and spatial attention. Spatial attention has been tied to repulsive biases (37) at the locus of attention and is also known for improvements in discriminability, although little is known about how discriminability changes with stimulus value. All data curves are replotted from the corresponding publications. All data represent biases and changes in threshold (ratio) relative to the corresponding control conditions (no surround; preadapt; prelearn; nonattended).

constraint that tightly links the encoding and decoding processes of perception (Fig. 2*A*).

The law allows us to predict either perceptual bias based on measured data for discrimination threshold or vice versa. One general prediction is that stimulus variables that follow Weber's law should exhibit perceptual biases that are linearly proportional to the stimulus value as demonstrated with examples in Fig. 3*D*. Furthermore, because perceptual illusions are often examples of a strong form of perceptual bias induced by changes in context, we predict that these illusions should be accompanied with substantial threshold changes according to our law.

Perceptual biases can arise for different reasons, not all of which are aligned with the assumptions we made in our derivations. In particular, because we assumed that the uncertainty in the inference process is predominantly due to internal (neural) noise of the observer, we do not expect the proposed law to hold under conditions where stimulus ambiguity/noise is the dominant source of uncertainty. In this case, we expect discrimination threshold to be mainly determined by the stimulus uncertainty and not the prior expectations as we have assumed (Eq. **2**).

It is worth noting that the law can also be expressed in terms of perceptual variance rather than discrimination threshold. This can be useful because some psychophysical experiments designed for measuring perceptual bias (e.g., by a method of adjustment) often record variance in subjects' estimates as well. Using the Cramer–Rao bound on the variance $\hat{\sigma}^2(\theta)$ of a biased estimator, we can rewrite Eq. **4** as

$$b(\theta) \propto (\hat{\sigma}^2(\theta)/(1+b'(\theta))^2)'. \quad \textbf{[5]}$$

For relatively small and smoothly changing biases, the predictions for variance and discrimination threshold are similar (see *Supporting Information* for details).

Last but not least, perhaps the most surprising finding is that the law seems to hold for bias and discriminability patterns induced by contextual modulation (Fig. 4). This implies not only that changes in encoding and decoding can happen immediately (e.g., spatial context), or at least on short time scales, but also that these changes are matched between encoding and decoding by relying on identical assumptions (i.e., prior expectations) about the structure of the sensory environment. This fundamentally contrasts with existing theories that assume mismatches between encoding and decoding (i.e., the "coding catastrophe") to be responsible for many of the known contextually modulated bias effects (46). It also contrasts with findings that put the locus of perceptual learning either at the encoding (47) or at the decoding (48) level; we predict learning to occur at both levels. Whether these contextual priors actually match the stimulus distributions within these contexts or not is unclear and remains a subject for future studies. Data from spatial attention experiments (Fig. 4) at least suggest that they may reflect subjective rather than objective expectations. This would imply that the distinction is not relevant in the context of the observer model considered here (Fig. 2*A*) and that efficient encoding and Bayesian decoding are both optimized for identical prior expectations, irrespective of whether these expectations are subjective or objective. We believe that the proposed law and its underlying theoretical assumptions have profound implications for the computations and neural mechanisms governing perception, which we have just started to explore.

ACKNOWLEDGMENTS. We thank V. De Gardelle, S. Kouider, and J. Sackur for sharing their data. We thank Josh Gold and Michael Eisenstein for helpful comments and suggestions on earlier versions of the manuscript. This research was performed while both authors were at the University of Pennsylvania. We thank the Office of Naval Research for supporting the work (Grant N000141110744).

1. Helmholtz Hv (1867) *Handbuch der Physiologischen Optik, Allgemeine Enzyklopädie der Physik* (Voss, Leipzig, Germany), Vol 9.
2. Fechner GT (1876) *Vorschule der Ästhetik* (Breitkopf and Härtel, Leipzig, Germany), Vol 1.
3. Green D, Swets J (1966) *Signal Detection Theory and Psychophysics* (Wiley, New York).
4. Wei XX, Stocker AA (2015) A Bayesian observer model constrained by efficient coding can explain "anti-Bayesian" percepts. *Nat Neurosci* 18:1509–1517.
5. Seung H, Sompolinsky H (1993) Simple models for reading neuronal population codes. *Proc Natl Acad Sci USA* 90:10749–10753.
6. Seriès P, Stocker AA, Simoncelli EP (2009) Is the homunculus 'aware' of sensory adaptation? *Neural Comput* 21:3271–3304.
7. Wei XX, Stocker AA (2012) Efficient coding provides a direct link between prior and likelihood in perceptual Bayesian inference. *Advances in Neural Information Processing Systems NIPS 25* (MIT Press, Cambridge, MA), pp 1313–1321.
8. Attneave F (1954) Some informational aspects of visual perception. *Psychol Rev* 61:183–193.
9. Barlow HB (1961) Possible principles underlying the transformation of sensory messages. *Sensory Communication*, ed Rosenblith W (MIT Press, Cambridge, MA), pp 217–234.
10. Linsker R (1988) Self-organization in a perceptual network. *Computer* 21:105–117.
11. Brunel N, Nadal JP (1998) Mutual information, Fisher information, and population coding. *Neural Comput* 10:1731–1757.
12. Ganguli D, Simoncelli EP (2014) Efficient sensory encoding and Bayesian inference with heterogeneous neural populations. *Neural Comput* 26:2103–2134.
13. Wei XX, Stocker AA (2016) Mutual information, Fisher information and efficient coding. *Neural Comput* 28:305–326.
14. Caelli T, Brettel H, Rentschler I, Hilz R (1983) Discrimination thresholds in the two-dimensional spatial frequency domain. *Vision Res* 23:129–133.
15. De Gardelle V, Kouider S, Sackur J (2010) An oblique illusion modulated by visibility: Non-monotonic sensory integration in orientation processing. *J Vis* 10:6.
16. Crane BT (2012) Direction specific biases in human visual and vestibular heading perception. *PLoS One* 7:e51383.
17. Gros BL, Blake R, Hiris E (1998) Anisotropies in visual motion perception: A fresh look. *J Opt Soc Am A Opt Image Sci Vis* 15:2003–2011.
18. Rauber HJ, Treue S (1998) Reference repulsion when judging the direction of visual motion. *Perception* 27:393–402.
19. Krukowski AE, Stone LS (2005) Expansion of direction space around the cardinal axes revealed by smooth pursuit eye movements. *Neuron* 45:315–323.
20. Smyrnis N, Mantas A, Evdokimidis I (2014) Two independent sources of anisotropy in the visual representation of direction in 2-D space. *Exp Brain Res* 232:2317–2324.
21. Portfors-Yeomans C, Regan D (1996) Cyclopean discrimination thresholds for the direction and speed of motion in depth. *Vision Res* 36:3265–3279.
22. Welchman AE, Lam JM, Bülthoff HH (2008) Bayesian motion estimation accounts for a surprising bias in 3D vision. *Proc Natl Acad Sci USA* 105:12087–12092.
23. Sweeny TD, Haroz S, Whitney D (2012) Reference repulsion in the categorical perception of biological motion. *Vision Res* 64:26–34.
24. Weber EH (1978) *The Sense of Touch*, trans Ross HE, Murray DJ (Academic, San Diego).
25. Georgeson M, Ruddock K (1980) Spatial frequency analysis in early visual processing. *Philos Trans R Soc Lond B Biol Sci* 290:11–22.
26. Waugh S, Hess R (1994) Suprathreshold temporal-frequency discrimination in the fovea and the periphery. *J Opt Soc Am A Opt Image Sci Vis* 11:1199–1212.
27. Vintch B, Gardner JL (2014) Cortical correlates of human motion perception biases. *J Neurosci* 34:2592–2604.
28. Stocker AA, Simoncelli EP (2006) Noise characteristics and prior expectations in human visual speed perception. *Nat Neurosci* 9:578–585.
29. Solomon JA, Morgan MJ (2006) Stochastic re-calibration: Contextual effects on perceived tilt. *Proc Biol Sci* 273:2681–2686.
30. Westheimer G (1990) Simultaneous orientation contrast for lines in the human fovea. *Vision Res* 30:1913–1921.
31. Tzvetanov T, Womelsdorf T (2008) Predicting human perceptual decisions by decoding neuronal information profiles. *Biol Cybern* 98:397–411.
32. Regan D, Beverley K (1985) Postadaptation orientation discrimination. *J Opt Soc Am A* 2:147–155.
33. Mitchell DE, Muir DW (1976) Does the tilt after-effect occur in the oblique meridian? *Vision Res* 16:609–613.
34. Regan D, Beverley K (1983) Spatial-frequency discrimination and detection: Comparison of postadaptation thresholds. *J Opt Soc Am A* 73:1684–1690.
35. Blakemore C, Nachmias J, Sutton P (1970) The perceived spatial frequency shift: Evidence for frequency-selective neurones in the human brain. *J Physiol* 210:727–750.
36. Chen N, Fang F (2011) Tilt aftereffect from orientation discrimination learning. *Exp Brain Res* 215:227–234.
37. Suzuki S, Cavanagh P (1997) Focused attention distorts visual space: An attentional repulsion effect. *J Exp Psychol Hum Percept Perform* 23:443–463.
38. Appelle S (1972) Perception and discrimination as function of stimulus orientation: The "oblique effect" in man and animals. *Psychol Bull* 78:266–278.
39. Campbell FW, Nachmias J, Jukes J (1970) Spatial-frequency discrimination in human vision. *J Opt Soc Am* 60:555–559.
40. Mandler M (1984) Temporal frequency discrimination above threshold. *Vision Res* 24:1873–1880.
41. Ross HE (1964) Constant errors in weight judgements as a function of the size of the differential threshold. *Br J Psychol* 55:133–141.
42. Szpiro S, Spering M, Carrasco M (2014) Perceptual learning modifies untrained pursuit eye movements. *J Vis* 14:8.

43. Carrasco M, McElree B (2001) Covert attention accelerates the rate of visual information processing. *Proc Natl Acad Sci USA* 98:5363–5367.
44. Cutrone E (2017) Attention, appearance, and neural computation. PhD thesis (New York University, New York).
45. Stevens SS (1957) On the psychophysical law. *Psychol Rev* 64:153–181.
46. Schwartz O, Hsu A, Dayan P (2007) Space and time in visual context. *Nat Rev Neurosci* 8:522–535.
47. Schoups A, Vogels R, Qian N, Orban G (2001) Practising orientation identification improves orientation coding in V1 neurons. *Nature* 412:549–553.
48. Law CT, Gold JI (2008) Neural correlates of perceptual learning in a sensory-motor, but not a sensory, cortical area. *Nat Neurosci* 11:505–513.
49. Kullback S, Leibler RA (1951) On information and sufficiency. *Ann Math Stat* 22:79–86.
50. McDonnell MD, Stocks NG (2008) Maximally informative stimuli and tuning curves for sigmoidal rate-coding neurons and populations. *Phys Rev Lett* 101:058103.