

Annual Review of Economics

Modeling Imprecision in Perception, Valuation, and Choice

Michael Woodford

Department of Economics, Columbia University, New York, NY 10027, USA;
email: michael.woodford@columbia.edu

Annu. Rev. Econ. 2020. 12:579–601

First published as a Review in Advance on
May 1, 2020

The *Annual Review of Economics* is online at
economics.annualreviews.org

<https://doi.org/10.1146/annurev-economics-102819-040518>

Copyright © 2020 by Annual Reviews.
All rights reserved

JEL codes: C25, C91, D81, D87

ANNUAL
REVIEWS **CONNECT**

www.annualreviews.org

- Download figures
- Navigate cited references
- Keyword search
- Explore related articles
- Share via email or social media

Keywords

psychophysics, efficient coding, stochastic choice, global games, risk attitude, context effect

Abstract

Traditional decision theory assumes that people respond to the exact features of the options available to them, but observed behavior seems much less precise. This review considers ways of introducing imprecision into models of economic decision making and stresses the usefulness of analogies with the way that imprecise perceptual judgments are modeled in psychophysics—the branch of experimental psychology concerned with the quantitative relationship between objective features of an observer’s environment and elicited reports about their subjective appearance. It reviews key ideas from psychophysics, provides examples of the kinds of data that motivate them, and proposes lessons for economic modeling. Applications include stochastic choice, choice under risk, decoy effects in marketing, global game models of strategic interaction, and delayed adjustment of prices in response to monetary disturbances.

1. INTRODUCTION

Economic analysis seeks to explain human behavior in terms of the incentives that people's situations provide for taking various actions. It is common to assume that behavior responds to the objective incentives provided by the situation. However, it is evident that people can only respond to incentives (the quality of goods on offer, their prices, and so on) to the extent that they are correctly perceived; and it is not realistic to assume that people (as finite beings) are capable of perfectly precise discrimination between different objective situations. Thus, people should not be modeled as behaving differently in situations that they do not recognize as different, even if it would be better for them if they could (see Luce 1956 and Rubinstein 1988 for early discussions of this issue).

But how should imprecision in people's recognition of their true situations be introduced into economic models? This review argues that economists have much to learn from studies of imprecision in people's perception of sensory magnitudes. The branch of psychology known as psychophysics has for more than 150 years sought to carefully measure and mathematically model the relationship between the objective physical properties of a person's environment and the way these are subjectively perceived (as indicated by an experimental subject's overt responses, most often, but sometimes by physiological evidence as well). Here we review some of the key findings from this literature and suggest possible lessons for economic modeling. (For alternative discussions of possible lessons for economics, readers are referred to Weber 2004, Caplin 2012.)

While the phenomenology of sensory perception is quite rich, we stress here the power of a single modeling approach to explain many well-known findings. In this approach, imprecision in the judgments that subjects express is attributed to imprecision in the way that the objective external situation is represented by a pattern of activity within the subject's nervous system; the responses can be modeled as optimal, subject to the constraint that the response rule must be based on the imprecise internal representation. We propose that this same approach might usefully be adopted in economic modeling and provide some examples of its application.

Section 2 reviews the classic problem (first studied by Ernst Weber in the 1830s) of judgments about the comparative magnitude of two stimuli along some dimension, and it discusses the applicability of similar models to the problem of stochastic choice. Section 3 then considers more complex experimental designs in which the magnitude of a single stimulus is estimated by choosing from among a continuum of possible responses; this is of particular interest as an element in decisions that require advantages and disadvantages to be traded off against one another. The estimation biases that are observed in sensory contexts provide a potential explanation for patterns of choice behavior that are instead often attributed to nonstandard preferences. Finally, Section 4 discusses ways in which judgments about a given sensory magnitude can be influenced by the context in which the stimulus appears; it is proposed that similar mechanisms can explain choice behavior that appears inconsistent with the existence of any coherent preference ordering.

2. THE STOCHASTICITY OF COMPARATIVE JUDGMENTS

A first important lesson of the psychophysics literature is not only that people (or other organisms) are unable to make completely accurate comparisons as to which is greater of two sensory magnitudes—which of two weights is heavier, which of two lights is brighter, etc.—when the two magnitudes are not too different, but also that the responses given generally appear to be a random function of the objective properties of the two stimuli that are presented. At the same time, the random responses of experimental subjects are not pure noise—that is, completely uninformative about the truth. Thus, rather than it being the case that any two stimuli either can be told apart




Figure 1

Psychometric functions for comparisons of numerosity. The number of items (Xs) in the reference array is (a) 25, (b) 100, and (c) 400. The right scale indicates the probability of judging the second numerosity to be larger (in percent), while the left scale indicates the corresponding z score. Figure adapted with permission from Krueger (1984).

(so that the heavier of two weights is always judged to be heavier, whenever those two weights are presented for comparison) or cannot be (so that the subject is equally likely to guess that either of the weights is the heavier of the two), what one typically observes is that the probability of a given response (e.g., the second weight seems heavier) increases monotonically with increases in the actual relative magnitude of the two stimuli. Experimental procedures often focus on the estimation of this increasing functional relationship, plotted as a psychometric function (see, e.g., Gescheider 1997, chap. 3; Kingdom & Prins 2010, chap. 4; Glimcher 2011, chap. 4).

Figure 1 provides an example of such plots (from Krueger 1984) when the stimulus feature is numerosity, or the number of items in a disordered array (in this case, Xs typed on a sheet of paper); in the experiment, a subject is asked to judge whether a second array is more or less numerous than a first (reference) array on the basis of a quick impression (rather than counting). While this is not one of the most classic examples,¹ it has the advantage for the purposes of this review of shedding light on the imprecise mental representation of numerical information—and thus of offering an especially plausible analogy for judgments of economic value.

In each panel of the **Figure 1**, the number n_1 of items in the reference array is fixed, and the fraction of trials on which subjects judge the second array to be more numerous is plotted as a

¹In particular, for many standard examples, the stimulus feature that is compared—such as weight, length, brightness, speed, or direction of motion—is one that takes a continuum of possible values, so that one can meaningfully speak of response probabilities as varying continuously with the true stimulus magnitude.

function of the true difference in numerosity $n_2 - n_1$;² the value of n_1 increases from 25 to 100 to 400 as one proceeds from the top panel to the bottom. In each panel, the probability of judging n_2 to be larger than n_1 is steadily increasing as a function of the true difference.

A classic approach to modeling imprecise comparisons of this kind, dating to the work of Fechner [1966 (1860)], supposes that a true stimulus magnitude x gives rise to an internal representation r , drawn from a probability distribution $P(r|x)$ that depends on the true magnitude. This representation r can be understood to refer to a pattern of neural activation in regions of the cortex involved in processing stimuli of that kind, as a result of the person's (or other organism's) contact with a stimulus of magnitude x ; it is random because of randomness in the way that neurons fire in response to the signals they receive. A comparative judgment between two magnitudes x_1 and x_2 is made on the basis of the corresponding internal representations r_1 and r_2 ; the randomness of r_1 and r_2 makes such comparisons random, even if the rule by which responses are generated is optimal, subject to the constraint that it must be based on the noisy internal representations.

To make this more concrete, a widely used model developed by Thurstone (1927) assumes that the internal representation can be summarized by a single real number and that it is drawn from a normal distribution $N(m(x), v^2)$, where $m(x)$ is an increasing function of the true magnitude, and the standard deviation $v > 0$ is independent of the true magnitude.³ If for each of two stimuli x_1 and x_2 , the internal representation r_i is an independent draw from the corresponding distribution, then an optimal decision rule will judge that x_2 seems greater than x_1 if and only if $r_2 > r_1$.⁴ This in turn implies that the probability of such a judgment, conditional on the true magnitudes x_1 and x_2 (known to the experimenter), is predicted to be

$$\text{Prob}[\text{"}x_2 \text{ greater"}|x_1, x_2] = \Phi\left(\frac{m(x_2) - m(x_1)}{\sqrt{2}v}\right), \quad 1.$$

where $\Phi(z)$ is the cumulative distribution function of a standard normal distribution. Thus, the probability of correctly distinguishing the relative magnitudes of two stimuli depends on their distance $|m(x_2) - m(x_1)|$ from one another on the Thurstone scale established by the mapping $m(x)$.⁵

This equation predicts the shape of a psychometric function, if one plots the response probability as a function of x_2 for some fixed value of x_1 . If the measured response probabilities are z -transformed,⁶ Equation 1 implies that one should have

$$z(\text{Prob}) = \frac{m(x_2) - m(x_1)}{\sqrt{2}v} \approx \frac{m'(x_1)}{\sqrt{2}v} \cdot (x_2 - x_1) \quad 2.$$

for values of x_2 sufficiently close to the reference magnitude x_1 . Thus, when the relationship is plotted as in **Figure 1**,⁷ the relationship should be approximately linear, as shown in the figure,

²This difference is reported in steps. A one-step increase means one more X in **Figure 1a**, three more in **Figure 1b**, and nine more in **Figure 1c**.

³This is Thurstone's celebrated "Case V"

⁴Here we assume a two-alternative, forced-choice experimental design in which the subjects must select one of two possible responses, regardless of their degree of confidence in their answer. If the prior distribution from which true values (x_1, x_2) are drawn is symmetric [i.e., (x_2, x_1) has exactly the same probability of being presented as (x_1, x_2)], then this is the response rule that maximizes the probability of a correct choice.

⁵Data on the frequency with which different comparative judgments are made, such as the one plotted in **Figure 1**, can allow the identification of this scale up to an affine transformation.

⁶That is, the probabilities are replaced by $z(P) \equiv \Phi^{-1}(P)$, the inverse of the function $P = \Phi(z)$.

⁷Note that in each of the panels of **Figure 1**, the vertical axis is linear in the z score $z(P)$ rather than in the probability P (marked on the right-hand side of the panel).

with a response probability of 0.5 when $x_2 - x_1 = 0$ and a slope proportional to $m'(x)$, evaluated at the reference magnitude.

Figure 1 not only shows that the psychometric function in each case is roughly of the predicted form shown in Equation 2, but it also allows $m'(x)$ to be evaluated at three different points in the range of possible stimuli. One finds that the size of the difference in number required for a given size of effect on the response probability is not independent of the range of numerosities being compared. Suppose that one defines the discrimination threshold as the average of the increase in the number of Xs required for the probability of judging n_2 to be greater than n_1 to rise from 0.5 to 0.75 and the decrease in number required for this probability to fall from 0.5 to 0.25. Then in the data shown in **Figure 1**, this threshold is found to be 3.1 when $n_1 = 25$, 11.7 when $n_1 = 100$, and 32.3 when $n_1 = 400$.

This increase in the discrimination threshold as the reference stimulus magnitude increases is called diminishing sensitivity, and it is an ubiquitous finding in the case of extensive quantities such as length, area, weight, brightness, loudness, etc. This is consistent with the Thurstone model, under the assumption that $m(x)$ is a strictly concave function. A famous formulation, Weber's law, asserts that the discrimination threshold should increase in proportion to the reference magnitude n_1 ; as first proposed by Fechner [1966 (1860)], this would follow from the model in the case that one assumes that $m(x)$ is (some affine transformation of) the logarithm of x .⁸

2.1. Encoding and Decoding as Distinct Processes

In the early psychophysics literature, the randomness of comparisons of the kind discussed above was often modeled by simply postulating that an objective stimulus magnitude x gave rise to a perceived magnitude \hat{x} , drawn randomly from a distribution that depended on x . The probability of an incorrect comparison then depended on the degree of overlap between the distributions of possible perceived values associated with different but similar true magnitudes, as in the discussion above.

Here we have instead taken a more modern point of view in which decisions are based on a noisy internal representation r , which is not itself a perceived value of the magnitude x but only an available piece of evidence on the basis of which the brain might produce a judgment about the stimulus or a decision of some other kind. (Note that r need not be measured in the same units as x , or even have the same dimension as x .) The cognitive process through which judgments are generated is then modeled as involving (at least) two stages: encoding of the stimulus features (the process through which the internal representation is produced), followed by decoding of the internal representation to draw a conclusion about the stimulus that can be consciously experienced and reported [Dayan & Abbott (2001, chap. 3) provide a textbook discussion].

In this way of conceiving matters, perception has the structure of an inference problem—even though the decoding process is understood to occur automatically, rather than through conscious reasoning—and tools from statistical decision theory have proven useful as a source of hypotheses. In particular, once matters are conceived in this way, it is natural to consider (at least as a theoretical

⁸Even when Weber's law holds approximately, it is often only for variation in the stimulus magnitude over some range, beyond which the approximation breaks down. The plots provided by Ganguli & Simoncelli (2016) offer an example of how discrimination thresholds vary with stimulus magnitude in a variety of sensory domains. Some researchers have proposed that Weber's law also holds for the perception of numerosity (e.g., Dehaene 2003, Nieder & Miller 2003, Cantlon & Brannon 2006), but the evidence is much stronger for diminishing sensitivity at a rate that is not necessarily precisely consistent with Weber's law. In the estimates from Krueger (1984) cited above, the discrimination threshold increases with the reference numerosity with an elasticity that is instead about 0.85.

benchmark) models in which the decoding is assumed to represent an optimal inference from the evidence provided by the internal representation.⁹

Why should one adopt such a complicated model of imprecise comparisons, rather than simply postulating a distribution of perceived values about which an experimental subject might then be interrogated? One answer is that the development of constantly improving methods of measurement of neural activity has made the concept of an internal representation, distinct from the observable behavior that may be based on it, something more than just a latent variable that is postulated for convenience in explaining the logical structure of one's model of the observables; if one wishes to use such measurements to discipline models of perception, then the candidate models must include variables to which the neural measurements may be taken to correspond (for examples, see Dayan & Abbott 2001, chap. 3).

Another important answer is that such a theory allows one to understand in a parsimonious way how changes in the context in which a stimulus is presented can affect the perceptual judgments that are made about it. In a binary comparison task of the kind discussed above, the probability of a subject's giving a particular response is not only a function of the objective characteristics of the two stimuli presented; it can also depend, for example, on the frequency with which the second stimulus magnitude is greater than the first, rather than the reverse,¹⁰ or on the relative incentive for a correct response in the two possible cases. It is easy to understand how these latter considerations can influence a subject's judgments in the encoding/decoding model: Even if one supposes that in the encoding stage the distribution from which the representation r is drawn depends only on the particular stimulus feature x , and not on any other aspects of the context,¹¹ an optimal decoding rule should take other aspects of the context into account. In particular, from the standpoint of Bayesian decision theory, the optimal inference to make from any given evidence depends both on the decision maker's prior and on the objective (or reward) function that they seek to maximize.

Signal detection theory (Green & Swets 1966) applies this kind of reasoning to the analysis of perceptual judgments. **Figure 2** (which reproduces figures 4.1 and 4.2 from Green & Swets 1966) shows a classic application of the theory. The figure plots data from an experiment in which a single subject is asked to indicate on each trial which one of two auditory stimuli (denoted s and n) is presented by choosing one of two possible responses (S or N).¹² In each of several blocks of trials, the same two stimuli are used, but the prior probability of s rather than n being presented may vary across blocks, as may the financial incentive given the subject to avoid false positive as opposed to false negative responses.

The location of each circle indicates both the subject's hit rate $P(S|s)$ (the probability of correctly detecting the signal when it is present), on the vertical axis, and the subject's false alarm rate $P(S|n)$ (the probability of incorrectly reporting a signal when none is present), on the horizontal axis. The diagonal line in each panel indicates the combinations of hit rate and false alarm rate that would be possible under pure chance (that is, if the subject were entirely deaf); all points above the

⁹Our explanation above of why it makes sense to assume that the judgment " x_2 seems greater" is produced if and only if $r_2 > r_1$ is an example of such an assumption of optimal decoding (see footnote 4).

¹⁰It was assumed above, in our discussion of the normative basis for a particular rule for determining the perceptual judgment (footnote 4), that either stimulus was equally likely to be the greater one, and this is true in many experiments. However, it is possible for the frequencies to differ in a particular experiment (or block of trials) and for subjects to learn this (or be told), as illustrated below.

¹¹This is a common simplifying assumption, but in fact context can influence encoding as well, as discussed in Section 4.

¹²The stimulus s is one in which a signal (a tone) is presented amid static, while stimulus n consists only of the noise. The experiments developed out of work (originally with visual stimuli) seeking to measure the accuracy of human operators of radar equipment (Creelman 2015).

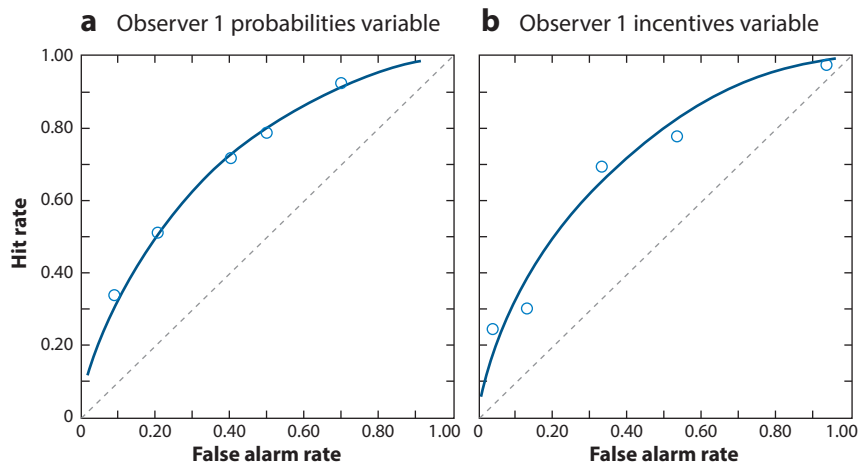


Figure 2

Conditional response probabilities in a signal detection task. Each of the circles in the figure plots the subject's conditional response probabilities for one block of trials. The trade-off is shown between the hit rate (vertical axis) and the false alarm rate (horizontal axis), as one varies the prior probability of occurrence of the two stimuli (*a*) or the relative rewards for correct identification of the two stimuli (*b*). In each case, the efficient frontier (ROC curve) is shown by the bowed solid curve. Abbreviation: ROC, receiver operating characteristic. Figure adapted from Green & Swets (1966) with permission of Peninsula Publishing.

diagonal indicate some ability to discriminate between the two stimuli, with perfect performance corresponding to the upper left corner of the figure.

If one supposes that a subject should have a perception *S* or *N* that is drawn from a probability distribution that varies depending on whether the stimulus presented is *s* or *n*—but that depends only on the stimulus presented and not on other aspects of the context—then one should expect a given subject to exhibit the same hit rate and false alarm rate in each block of trials. (Of course one could expect to see small differences owing to random sampling error, given the finite length of each block of trials, or perhaps drift over time in the subject's functioning owing to factors such as fatigue, but these differences should not be systematically related to the prior frequencies or to incentives.)

Instead, in **Figure 2** one sees systematic effects of both aspects of the context. In **Figure 2a**, the reward is the same for correct identification of either stimulus, but the probability of presenting stimulus *s* on any given trial varies across the blocks of trials (it is 0.1, 0.3, 0.5, 0.7, or 0.9, depending on the block); and one sees that as the prior probability of the state being *s* is increased, both $P(S|s)$ and $P(S|n)$ are monotonically increasing. In **Figure 2b**, both stimuli are presented with equal probability, but the relative incentive for correct identification of the two cases is varied; and one sees that as the relative reward for correct recognition of state *s* is increased, both $P(S|s)$ and $P(S|n)$ are again monotonically increasing.

Both phenomena are easily explained by a model that distinguishes between the noisy internal representation *r* upon which the subject's response on any trial must be based and the subject's classification of the situation (as indicated by response *S* or *N*). Suppose, as in the model above, that *r* is a single real number¹³ and that it is drawn from a Gaussian distribution $N(\mu_i, \nu^2)$, with a variance that is the same for both stimuli but a mean that differs depending on the stimulus

¹³In this kind of task, because there are only two possible true situations (*s* or *n*), even if the internal representation was high-dimensional, it would clearly suffice to describe it using a single real number as a sufficient

($i = s$ or n). Let us further suppose, without loss of generality, that $\mu_s > \mu_n$, so that the likelihood ratio in favor of the stimulus being s rather than n is an increasing function of r . If the subject's response on any trial must be based on r alone, an efficient response criterion—in the sense of minimizing the probability of a type I error subject to an upper bound on the probability of type II errors, or vice versa—is necessarily a likelihood ratio test,¹⁴ which in the present case means that the subject should respond S if and only if r exceeds some threshold c .

Varying the value of c (which corresponds to changing the relative weight placed on avoiding the two possible types of errors) allows one to generate a one-parameter family of efficient response rules, each of which implies a particular set of conditional response probabilities (and hence corresponds to a point in the kind of plots shown in **Figure 2**). In the case of a particular value of the ratio $(\mu_s - \mu_n)/\nu$, which determines the degree of discriminability of the two stimuli given the subject's noisy encoding of them, the points corresponding to this family of efficient response rules can be plotted as a curve, known as the receiver operating characteristic (ROC) curve.

This is shown as a concave, upward-sloping curve in each of the panels of **Figure 2**; here the ROC curve is plotted under the assumption that μ_s and μ_n differ by 0.85 standard deviations. We see that under this parameterization, the subject's pattern of responses falls close to the efficient frontier in all cases. Moreover, the change in the subject's response frequencies is in both cases consistent with movement along the efficient frontier in the direction that would be desirable if the subject had an increased reason to prioritize an increased hit rate even at the expense of an increased false alarm rate.

Hence the subject's response probabilities are easily interpreted as reflecting a two-stage process, in which encoding and decoding are influenced by separate factors that can be experimentally manipulated independently of one another. On the one hand, an experimenter can change aspects of the stimuli, unrelated to the feature on the basis of which they are to be classified, that can affect the precision of encoding (for example, varying the length of time that a subject is able to listen to the stimuli before expressing a judgment); if this changes the value of ν , it should shift the location of the ROC curve along which the subject should operate (whatever the prior and incentives may be). On the other hand, an experimenter can change the prior and/or incentives, which should affect the relative priority assigned to the two possible types of error, and hence the location on any given ROC curve in which the subject would ideally operate. The fact that encoding and decoding are determined by distinct sets of parameters that can be independently manipulated makes it useful to model the subject's judgments as the outcome of two separate processes of encoding and decoding.¹⁵

2.2. Implications for Economic Models

The above review of the way in which imprecise comparisons have been successfully modeled in sensory domains suggests a number of implications for models of imprecise economic decisions. For example, some authors propose to introduce cognitive imprecision into economic models by assuming that agents can respond only to a coarse classification of the current state of the world, but they should know with certainty to which element of some partition of the state space the current state belongs (see, e.g., Gul et al. 2017). In the case of a continuous state space, such a model implies that certain states that differ only infinitesimally should nonetheless be perfectly

statistic, namely the likelihood ratio of the two hypotheses given the noisy evidence (see discussion in Green & Swets 1966).

¹⁴This follows from the Neyman-Pearson lemma of statistical decision theory (see Green & Swets 1966).

¹⁵Green & Swets (1966, p. 86) call this the "separation of sensory and decision processes."

distinguished, because they happen to fall on opposite sides of a category boundary; but nothing of the sort is ever observed in the case of perceptual judgments. Instead, if one supposes that the imprecision that the model wants to capture should be analogous to the imprecision in the way that our brains recognize physical properties of the world, then one should model internal representations as probabilistically related to the external state, but not necessarily as discrete.

The literature on global games (see, e.g., Morris & Shin 1998, 2003) models imprecise awareness of the state of the world in a way that is more consistent with what we know about perception. In models of this kind, the true state is assumed to be a continuous variable, but agents are not assumed to be able to observe it precisely. This is shown to have important consequences for the nature of equilibrium, even when the imprecision in individual agents' private observations of the state is infinitesimally small (but nonzero). For example, even in models of bank runs or currency attacks that allow for multiple equilibria when all agents observe the state with perfect precision, there can be a unique equilibrium (and hence predictable timing of the run or attack) if decisions must be based on slightly imprecise private observations.

Here the imprecision in private observations is modeled by assuming that each agent has access to the value of a signal r , equal to the true value of the state x plus a random error term that is an independent draw for each agent from some (low-variance) distribution; for example, r might be a draw from $N(x, \nu^2)$ for some small (but positive) value of ν , just as in a Thurstonian model of imprecise perception. Indeed, it is important for the conclusions of the literature that the imprecision is modeled in this way: It is the overlap in the distributions of possible values of r associated with different nearby true states x that results in the failure of common knowledge that implies a unique equilibrium.

However, understanding the noise in private observations of the state as reflecting inevitable cognitive imprecision would change the interpretation of global games models in some respects. Many discussions assume (in line with conventional models of asymmetric information) that the imprecise private observations represent opportunities that different individuals have to observe different facts about the world, owing to their different situations; however, they also assume that some facts (such as government data releases or market prices) should be publicly visible, so that everyone should observe them with perfect precision and this should also be common knowledge. According to this perspective, the question of whether there should be sufficient common knowledge for the agents to be able to coordinate on multiple equilibria depends on how informative the publicly observable signals (about which there should be common knowledge) are about the relevant state variable; a number of authors have proposed reasons why there should be public signals that should overturn the classic uniqueness result of the global games analysis (see, e.g., Angeletos & Werning 2006, Hellwig et al. 2006). But if one regards at least a small amount of randomness in the internal representation of quantities observed in the world as inevitable—as both psychophysical and neurophysiological evidence would indicate—then there should be no truly “public signals” in the sense assumed in this literature; and since only a small amount of idiosyncratic noise in private observations of the state is needed to obtain the global games result, the case emphasized in the classic result of Morris & Shin (1998) should be of more general relevance than is often appreciated.¹⁶

The stochasticity of comparative judgments in perceptual domains is perhaps most obviously relevant as a model of randomness in observed choice behavior. While standard models of rational

¹⁶Goryunov & Rigos (2019) show that the global games result obtains in an experiment in which both players are shown the same value for the state, but the value is displayed visually (by the location of a dot). It is possible that this results only from the ambiguity of visual rather than symbolic presentation of information. However, there is good evidence for imprecise semantic internal representations even of quantity information that is presented symbolically (see discussion below and in Khaw et al. 2019).




Figure 3

The fraction of trials in which a simple gamble was observed to be accepted, as a function of the amount (on the horizontal axis) that could be won. The indifference point identifies the terms under which it is inferred that the subject would be equally likely to accept or reject. Data from Mosteller & Noguee (1951).

choice imply that people’s choices should be a deterministic function of the characteristics of the options presented to them (assuming that these are described sufficiently completely), choices observed in laboratory experiments typically appear random, in the sense that the same subject does not always make the same choice when presented on multiple occasions with the same set of options.

Figure 3, based on a similar figure in Mosteller & Noguee (1951), shows a classic example. The figure plots data on the choices of a single subject presented on different trials during the same experiment with multiple variants of the same kind of gamble: whether to pay 5 cents in order to obtain, with 50% probability, a random amount X . The amount X differed from trial to trial; the figure shows the fraction of times that the subject accepted the gamble as a function of X (plotted in cents on the horizontal axis).

The experimenters’ goal was to elicit preferences with regard to gambles that could be compared with the predictions of expected utility theory (EUT). A problem that they faced (and the reason for showing the figure) is that their subjects’ choices were random; note that in **Figure 3**, for several intermediate values of X , it is neither the case that the subject consistently accepts the gamble nor that they consistently reject it. The figure illustrates how the researchers dealt with this issue: The observed choice frequencies were interpolated in order to infer the value of X for which the subject would accept the gamble exactly half the time, and this was labeled a case of indifference. The prediction that was required to be consistent with EUT (in the case of some nonlinear utility function, inferred from the subject’s choices) was that the subject should be exactly indifferent in this case.

A graph like the one in **Figure 3** is highly reminiscent of psychometric functions like those in **Figure 1**.¹⁷ This suggests that the randomness depicted might fruitfully be modeled in a similar way. Indeed, a common approach within economics has been to model stochastic choice using

¹⁷Indeed, it seems likely that Mosteller & Noguee’s (1951) experimental method—repeating the same questions many times on randomly ordered trials and tabulating response frequencies—reflected a familiarity with psychophysics. The method that they used to identify the indifference point is one commonly used with psychometric functions to identify a “point of subjective equality” as a measure of bias in comparative judgments (see, e.g., Kingdom & Prins 2010, p. 19).

an additive random utility model (McFadden 1981). It is assumed that on any given occasion of choice, each choice option i is assigned a valuation $v_i = u(x_i) + \epsilon_i$, where $u(x_i)$ is a deterministic function of the vector of characteristics x_i of that option, and ϵ_i is an independent draw from some distribution $F(\epsilon)$, assumed not to depend on the characteristics of the option. The option that is chosen on that occasion should then be the one with the highest value of v_i . (Because the $\{v_i\}$ are random variables, choice will be stochastic.) If the function $u(x)$ is linear in its arguments, and F is either a normal distribution or an extreme-value (Gumbel) distribution, this leads to a familiar econometric model of either the probit or logit form.

Such a model (especially if applied to binary choice, and if the random terms are assumed to be Gaussian) has many similarities with the Thurstone model of random perceptual judgments. However, economists often interpret random utility models as if the valuation v_i assigned to an option represents the true value of the option to the consumer on that occasion—that is, the model is interpreted as a model of rational choice with random fluctuation in tastes. In the case of perceptual judgments, instead, it is clear that the randomness of the judgments must be interpreted as random error in the recognition of the situation (since an objective truth exists as to which of two physical magnitudes is greater); and one wonders if much of the randomness observed in choice should not be interpreted the same way. Even if it requires no change in the mathematical form of the model of choice, the alternative interpretation matters for an assessment of people's level of welfare under alternative possible regulations of market transactions.

Even if one thinks of the random term ϵ_i as representing error in the process of evaluating the subject's degree of liking for the options, it is common to assume that a precise valuation for each option is computed, with a random term added only at the end of such a calculation; in this way, the relative likelihood of choice between two options depends on the relative magnitudes of the two deterministic components of their valuations, so that the core of the theory is still a deterministic preference ordering.¹⁸ Yet once one admits that the cognitive process involves random error, it is not obvious why it should be assumed to occur only at the end, adding a random term to an otherwise correctly computed quantity, rather than introducing error into the way that different pieces of information (the different elements of x_i) are assessed and integrated to produce an estimate of the value of the option.

It should be recalled that in many modern models of random perceptual judgment, noise is assumed to enter at earlier stages of processing: Noise in the nervous system corrupts the evidence that must subsequently be decoded to produce a judgment, rather than corrupting only the accuracy with which an answer that has been reached is communicated.¹⁹ The same idea can be used to model the way in which valuations of economic options are derived; but the predictions are different, in general, compared to a model in which the random valuation assigned to an option is assumed to equal its true value to the agent plus an independent error term. For example, as discussed in Section 4, the likelihood of choosing one good over another can be influenced by contextual factors that should have no effect on the true value of either item to the decision maker.

¹⁸This is implicit in an approach, like that of Mosteller & Noguee (1951), that assumes that which of two options is more often chosen indicates which of the options is preferred, and this should allow a deterministic preference ordering to be recovered even when choice is stochastic.

¹⁹Even in the case of purely sensory judgments, the relevant noise often occurs at later stages of processing (though at stages earlier than that of action choice), rather than simply representing noise in sensory receptors (Beck et al. 2012, Drugowitsch et al. 2016). Such later-processing noise—noise in the way quantities are stored and subsequently retrieved for use in further computations, rather than noise in initial perceptions of the data—is almost certainly the more important factor in situations like the experiment by Mosteller & Noguee (1951), where the data are presented in symbolic form (see further discussion in Khaw et al. 2019).




Figure 4

Mean of the length estimates produced by an experimental subject, plotted as a function of the length (in millimeters) that had previously been demonstrated to the subject. The different symbols identify distinct experimental sessions, in which the range of true distances (presented in random order) was different: 10–70 mm in series A of trials, 30–150 mm in series B, and 70–250 mm in series C. The red dots represent three sessions in which all true distances were of exactly the same length: 10 mm, 70 mm, or 250 mm. Both axes use a log scale. Figure adapted with permission from Laming (1997).

3. IMPRECISION AND BIAS

We have thus far considered only a classic form of experiment in which a subject is asked to compare the magnitudes of two stimuli along some dimension. Another kind of experiment requires the subject to estimate the magnitude of a single stimulus within a (possibly continuous) range of responses. This allows one not only to observe the randomness of the responses elicited by a given stimulus, but also to measure whether the responses are biased, in the sense that the subject's estimates are not even correct on average. In fact, bias is commonplace in perceptual judgments; and there is reason to think that both its nature and magnitude are closely connected to the noise in the internal representations on which judgments are based.

There are a variety of ways in which subjects can be asked to estimate the magnitude of a stimulus presented to them. They might be asked to choose from among a set of possibilities the new stimulus that is most similar in magnitude to one previously presented; or they might be asked to produce themselves a stimulus of equal magnitude to the one presented—for example, producing two successive taps to indicate the length of a time interval.²⁰ A common finding with respect to estimates of extensive magnitudes (such as distance, area, angular distance, or length of a time interval) is a conservative bias in subjects' estimates: Subjects tend to overestimate smaller magnitudes (on average) while underestimating larger ones.

Figure 4 illustrates this bias, using data from a classic study by Hollingworth (1909). In this experiment, a subject is asked to reproduce a particular spatial distance by moving their arm, after

²⁰Both of these methods avoid relying on any ability of the subject to verbally describe their subjective estimate of a sensory magnitude. Symbolic expression of estimates is instead common in experiments testing people's ability to estimate numerosity, as discussed below.

having had the distance shown to them by the experimenter also through a movement of their arm; the figure plots the mean distance estimate produced by the subject for each true distance presented by the experimenter.

In each of the sessions with variable lengths, the subject's estimates exhibit a clear conservative bias: The shorter distances used in that day's series are overestimated on average, while the longer distances are underestimated on average. The figure also illustrates another important finding: that the average estimate produced in response to a given stimulus depends not only on the objective magnitude of that stimulus in isolation, but also on how it compares to the range of stimuli used in that particular session. The mapping from true distance to mean estimated distance is similar in sessions A, B, and C (in each case, an increasing function, roughly linear when presented as a log-log plot), but the function shifts from day to day as the range of stimuli used is changed. The same stimulus (a 70-mm movement) may be underestimated, overestimated, or estimated with nearly zero bias, depending on whether it is unusually long, unusually short, or about average among the stimuli used in the session. Hollingworth (1910) found the same to be true in a number of different sensory domains, and christened this regularity "the central tendency of judgment." Petzschner et al. (2015) offer more recent examples from a variety of sensory domains.

Among the domains in which this kind of bias is observed is that of judgments of numerosity, already discussed above; this is one of the several respects in which the imprecision in judgments about numerical magnitudes resembles the imprecision in judgments about physical magnitudes like distance, leading Dehaene (2011) to speak of the existence of a "number sense." Conservative bias in estimates of numerosity has been documented many times, following the classic study by Kaufman et al. (1949). Often the relationship between true numerosity and average estimated numerosity is found to fall on a roughly linear log-log plot, but with a slope slightly less than 1, as in **Figure 4** (see, e.g., Krueger 1984, Kramer et al. 2011). However, the crossover point at which numerosity begins to be underestimated rather than overestimated differs considerably across experiments, in a way that correlates with differences in the range of numerosities used as stimuli in the different experiments (see Izard & Dehaene 2008 for discussion). **Figure 5** shows an example of two experiments that differ only in the range of numerosities used (1–30 in **Figure 5a**, 1–100 in **Figure 5b**); note that the location of the crossover point shifts, in a way consistent with the central tendency of judgment.




Figure 5

The mapping from true numerosity to the distribution of numerosity estimates in two experiments that differ only in the range of true numerosities used: (a) 1–30 and (b) 1–100. Figure adapted with permission from Anobile et al. (2012).

3.1. A Bayesian Model of Estimation Bias

Several authors have noted that estimation biases in these and other sensory domains are consistent with a model of optimal decoding of the stimulus magnitude implied by a noisy internal representation (see, e.g., Stocker & Simoncelli 2006; Petzschner et al. 2015; Wei & Stocker 2015, 2017). An optimal inference from noisy evidence will depend on the prior distribution from which the true state is expected to be drawn. The observed dependence of the mapping from objective magnitudes to average estimated magnitudes on the range of objective magnitudes used in a given experiment can then be interpreted as a natural consequence of inference using a prior that is appropriate to the particular context.²¹

As an example, suppose that a true magnitude x (say, one of the distances on the horizontal axis in **Figure 4**) has an internal representation r drawn from the distribution

$$r \sim N(\log x, v^2). \quad 3.$$

[The assumption that $m(x)$ is logarithmic is consistent with Fechner's explanation for Weber's law, a regularity that is observed in the case of distance comparisons.] If the true distance is assumed to be drawn from a log-normal prior distribution,

$$\log x \sim N(\mu, \sigma^2), \quad 4.$$

then the expected value of x , conditional on the representation r (i.e., the estimate given by the Bayesian posterior mean),²² will equal

$$\hat{x}(r) = E[x|r] = \exp[(1 - \beta) \log \bar{x} + \beta r], \quad 5.$$

where $\beta \equiv \sigma^2 / (\sigma^2 + v^2) < 1$, and $\bar{x} \equiv \exp[\mu + (1/2)\sigma^2]$ is the prior mean. Conditional on the true x , this implies that the estimate \hat{x} will be a log-normally distributed random variable, with mean and variance

$$e(x) \equiv E[\hat{x}|x] = Ax^\beta, \quad \text{var}[\hat{x}|x] = B e(x)^2, \quad 6.$$

where $A \equiv \exp(\beta^2 v^2 / 2) \cdot \bar{x}^{1-\beta}$ and $B \equiv \exp(\beta^2 v^2) - 1 > 0$.

This simple model (based on Petzschner et al. 2015) implies that a plot of the mean estimate as a function of the true magnitude should yield a linear log-log plot, as in **Figure 4**, with a slope equal to $\beta < 1$; the slope less than 1 implies a conservative bias. Moreover, if in different contexts, the degree of prior uncertainty is similar (in percentage terms), so that σ remains the same across contexts but μ is different, then optimal Bayesian estimation (with learning about the statistics of each context) would imply a different function $e(x)$ in each context. The elasticity (slope of the log-log plot) should remain the same across contexts, but the crossover point should increase in proportion to the prior mean \bar{x} in each context, in accordance with the central tendency of judgment. The model also implies that estimates should be more variable, the larger is x ; specifically, the standard deviation of \hat{x} should grow in proportion to the mean estimate $e(x)$. This latter property of scalar variability is also observed for many types of magnitude estimates (see Petzschner

²¹Of course, the appropriate prior has to be learned; one should therefore expect the mapping from objective magnitudes to estimates to shift over the course of an experimental session, especially at the beginning. Such learning effects can explain the often observed difference in estimation bias depending on the sequence in which different magnitudes are presented (see Petzschner et al. 2015 for discussion).

²²This estimate of x will be optimal in the sense of minimizing the mean squared error of the estimate under the prior. It is not the only possible rule that might be used in a Bayesian model of decoding (see, e.g., Dayan & Abbott 2001, chap. 4), but it is used by authors such as Wei & Stocker (2015).

et al. 2015), including estimates of numerosity (see, e.g., Whalen et al. 1999, Cordes et al. 2001, Izard & Dehaene 2008, Kramer et al. 2011).

Experimental results of the kind shown in **Figures 4** and **5** again exhibit diminishing sensitivity to increases in the stimulus magnitude, but in a different sense than the classic one discovered by Weber; in these figures, $e(x)$ is an increasing, strictly concave function of x ,²³ but this is not equivalent to the claim that $m(x)$ is a concave function of x . The functions $m(x)$ and $e(x)$ are measured using different experimental procedures, and in a model based on optimal Bayesian decoding, they should not generally coincide. For example, in the model just presented, $m(x)$ is logarithmic, while $e(x)$ is a power law. Both are increasing, strictly concave functions, but they exhibit diminishing sensitivity at different rates.²⁴

3.2. Biased Economic Valuations and Errors in Choice

Bayesian models of perceptual bias of the kind just illustrated provide a possible interpretation of some otherwise puzzling features of choice behavior. If choices are based on imprecise internal representations of the characteristics of the available options, and subjective valuations of economic options are imprecise in a similar way as perceptual judgments, then we should expect such valuations to be not only noisy (subject to random variability from one occasion of choice to another, even over short periods of time), but also biased on average.

However, it is worth noting that in such models, bias exists only to the extent that estimates are also noisy. Hence it is essential, under this program, that choice biases and randomness of choice be modeled together, rather than treating the specification of biases and the specification of random errors in choice as two completely independent aspects of a statistical model of the data controlled by different sets of parameters, as is often the case. Based on this approach, systematic behavioral tendencies that are commonly taken to reflect preferences can sometimes be interpreted instead as biases resulting from inference from noisy internal representations.

3.2.1. Application: explaining small-stakes risk aversion. A common observation in laboratory experiments is apparently risk-averse behavior. The data from Mosteller & Noguee (1951), plotted in **Figure 3**, provide an example: In this case, the gamble is a fair bet when X equals 10 cents (i.e., 5 cents/0.5), but the indifference point appears to be around 10.7 cents. A standard interpretation of risk aversion, of course, notes that it is implied by expected utility maximization in the case of diminishing marginal utility of wealth. However, this interpretation of Mosteller & Noguee's (1951) data would require one to suppose that an increase in wealth of 5.7 cents raises the subject's utility by no more than a loss of 5 cents would reduce it; and while logically possible, this is actually quite an extreme degree of curvature of the utility of wealth function, and it would require extraordinary risk aversion with respect to larger gambles (as explained in Rabin 2000) of a kind that is seldom observed.

There is instead no puzzle if we suppose that a subject's decision whether to accept the gamble must be based on a noisy internal representation of the payoffs associated with the alternative choices (Khaw et al. 2019). Consider the case of a choice between having an amount of money $C > 0$ with certainty and a gamble that promises an amount $X > 0$ with 50% probability but has 50% probability of paying nothing; and suppose that on each of a series of trials of this kind,

²³Recall that **Figure 4** is a log-log plot, so that a straight line corresponds to a power law of the kind described in Equation 6.

²⁴Neither of these functions is characterized by diminishing sensitivity in all cases. Wei & Stocker (2017) offer a discussion of how the acuity of discrimination between nearby stimuli [which depends on $m'(x)$] and estimation bias [measured by $e(x) - x$] vary over the stimulus space in a variety of sensory domains.

the values of C and X vary.²⁵ Suppose further that each of the amounts C and X that define the choice problem on a given trial has a noisy internal representation r_i (for $i = C, X$) drawn from the distribution shown in Equation 3, the mean of which is given by the log of the corresponding true value in each case, and that the decision whether to accept the gamble on any given trial must be based on $\mathbf{r} = (r_C, r_X)$.

Assuming that the amounts C and X are small enough for the subject's marginal utility of wealth to be essentially the same regardless of the choice made or the outcome of the gamble, an optimal decision criterion will be one that maximizes the mathematical expectation of the monetary payoff from the experimental trial. Thus, under the hypothesis that judgments with regard to such gambles are optimal (given the imprecision of the internal representation of the problem), the gamble should be accepted if and only if $E[X|\mathbf{r}] > 2E[C|\mathbf{r}]$. If these expectations are computed for a prior under which the values of C and X are independent draws from a log-normal distribution (as shown in Equation 4), then Equation 5 implies that the gamble should be accepted if and only if $\beta r_X > \beta r_C + \log 2$. Under the assumption that r_C and r_X are drawn from distributions of the form given in Equation 3, the probability of acceptance is predicted to be

$$\text{Prob}[\text{accept}] = \Phi\left(\frac{\log(X/C) - \beta^{-1} \log 2}{\sqrt{2}\nu}\right). \quad 7.$$

Equation 7 predicts that if the acceptance probability is plotted as a function of X (for fixed C), as in **Figure 3**, one should obtain an increasing sigmoid function like the one shown in the figure. Moreover, the indifference point identified using the method of Mosteller and Nogee should be where $X^{\text{indiff}}/C = 2^{1/\beta} > 2$, a point to the right of the value corresponding to a fair bet, so that the subject should appear to be risk averse. Thus the theory provides a unified explanation for both the randomness of observed choices and the apparent risk aversion even when stakes are quite small. Because the predicted ratio X^{indiff}/C is independent of the value of C , the model predicts nontrivial risk aversion even when stakes are arbitrarily small; at the same time, it does so without predicting an extraordinary degree of risk aversion in the case of large bets.

Under this explanation, the apparent risk aversion results from a bias in subjective estimates of the values of the different monetary payoffs, which depends on noise in the internal representation. One can ask whether the degree of noise needed to explain the observed degree of risk aversion is plausible. The curve in **Figure 3** represents the prediction shown in Equation 7 when the parameters σ and ν are fit to the data from Mosteller & Nogee (1951).²⁶ The measure of prior uncertainty σ required to fit the choice data is equal to 0.26.

Note that the difference between the largest and smallest values of $\log X$ in **Figure 3** is 1.16; this corresponds to a range of 4.5 standard deviations if $\sigma = 0.26$. Thus the assumed degree of prior uncertainty about the value of X on each trial is consistent with the range of values actually used in the experiment. If we take σ to be determined by the range of values used (known apart from the subject's behavior), then the curve in **Figure 3** shows that the value of ν needed to account for the subject's degree of apparent risk aversion is approximately the same as the one indicated by the randomness of their decisions: The same value ($\nu = 0.07$) fits both the slope of the choice curve (measuring the degree of randomness) and its horizontal location (measuring apparent risk

²⁵Note that the data plotted in **Figure 3** involve only one value of C (5 cents). However, Mosteller & Nogee (1951) used multiple values of C in their experiment; the figure shows only how the value of X required for indifference is computed for one particular value of C .

²⁶Note that under the assumption that C and X have the same prior distribution, the value of μ does not affect the predicted choice frequencies in Equation 7.

aversion). The required degree of imprecision in number representation is also relatively modest compared to that found in studies of numerosity perception.²⁷

There are additional reasons to think that apparent risk aversion may result from imprecise internal representations. Khaw et al. (2019) find that when choice curves of the kind shown in **Figure 3** are fit to the data of individual subjects, there is a strong positive correlation between subjects' apparent degree of risk aversion (measured by the horizontal intercept) and the degree of randomness of their responses (measured by the slope). Garcia et al. (2018) show furthermore that both the randomness and the apparent risk aversion in choice under risk can be predicted by the subject's degree of randomness in an independent numerosity comparison task (like the one shown in **Figure 1**).

The model also provides an explanation for the observation that the apparent risk aversion of subjects in laboratory experiments can be increased by increasing their "cognitive load" (e.g., Gerhardt et al. 2016), if we assume that having simultaneously to hold other information in one's mind reduces the available capacity for precise internal representation of the numbers involved in the choice under risk. Finally, Frydman & Jin (2019) find that reducing the degree of dispersion of the distributions of values from which X and C are drawn increases the precision of choice (i.e., it makes the subjects' choice functions more steeply increasing as a function of X/C for values of X/C around the value required for indifference). This has a straightforward explanation if the randomness of choice is attributed to noise in the internal representations of the values of X and C ,²⁸ while it would have no obvious explanation if one attributes random choice to randomness in the process of comparing options after their expected values have been correctly computed.

3.2.2. Other economic applications. The idea that gambles are valued on the basis of a noisy internal representation of their features can explain other aspects of experimentally observed behavior as well. For example, both Steiner & Stewart (2016) and Khaw et al. (2019) show how biases in the perceived probability of different outcomes of the kind postulated by Kahneman & Tversky (1979) can result from Bayesian decoding of noisy internal representations of the probabilities presented to the experimental subject. Essentially, the overestimation of small probabilities and underestimation of larger ones is another example of the kind of conservative bias found in many perceptual domains (discussed above), as first suggested by Preston & Baratta (1948).

The discounting of future payments relative to ones that can be received sooner is another area in which valuation biases that are commonly interpreted as indicating an aspect of subjects' preferences may instead be due to a perceptual bias, which can be modeled as optimal inference from a noisy internal representation. Gabaix & Laibson (2017) show that optimal inference from noisy mental simulations of the future outcomes resulting from alternative choices can result in underweighing outcomes farther in the future, even when the objective that decisions maximize is the expected value of the undiscounted stream of payments. This is again an example of conservative bias, with the bias greater in the case of payments that are represented with less precision; the authors show that this source of discounting can easily make the apparent time discount factor

²⁷For example, the data shown in **Figure 1** imply values of ν on the order of 0.13, if the data are fit to Equation 1 with $m(x) = \log x$. Note that it makes sense to expect greater noise if the number is presented visually, by an array of Xs, rather than symbolically.

²⁸As Frydman & Jin (2019) discuss, a theory of efficient coding (see Section 4) implies that when the dispersion of values occurring in the experiment is higher, less precise discriminations should be made between values that occur relatively often in the lower-variance context. This amounts effectively to an increase in the parameter ν in the above model, while the ratio β remains the same, reducing the slope of the curve predicted by Equation 7.

a hyperbolic function of distance in time.²⁹ This alternative interpretation has the advantage of helping to make sense of a variety of ways in which apparent time preference varies across contexts and can be affected by factors such as stress or cognitive load (e.g., Mullainathan & Shafir 2013).

The hypothesis that people's decisions are based on noisy internal representations of external conditions also provides an explanation for the failure of firms' prices to fully respond to changing macroeconomic conditions, such as changes in monetary policy (Woodford 2003). This is again an example of a conservative bias, the aggregate effects of which are greater in the case of strategic complementarity between the pricing decisions of different firms.³⁰ Such an explanation for the delayed adjustment of prices (and hence for the effects of monetary disturbances on aggregate economic activity) has an advantage over explanations that posit a fully informed optimizing behavior subject to the objective costs of price adjustment, in that it can more easily explain the fact that some kinds of market developments are much more rapidly reflected in prices than others. This is what one should expect if price setters pay closer attention to some aspects of their environment than others (and hence represent them more precisely), in accordance with a model of optimal allocation of scarce cognitive resources (Mackowiak & Wiederholt 2009).

More generally, noisy internal representations can explain biases in expectations relative to those implied by the conventional hypothesis of rational expectations. Coibion & Gorodnichenko (2015) document biases in the average forecasts of professional forecasters of the kind that should result if forecasts are based on a noisy internal representation of external conditions. Azeredo da Silveira & Woodford (2019) show that the further hypothesis of noisy memory of past cognitive states predicts a more complex pattern of biases, including the kinds of overextrapolation from recent observations that are often observed in the forecasts of individual forecasters (Bordalo et al. 2018) and in the laboratory (Beshears et al. 2013, Landier et al. 2019).

4. CONTEXT-DEPENDENT VALUATIONS

One of the features of observed choice behavior that is most problematic for normative models of rational choice is the fact that choices sometimes appear to reflect valuations of choice options that depend on what other options are available, and not simply on the consequences that follow from selecting a particular option. For example, an extensive literature in marketing studies the way in which adding a decoy good to the set of options available to consumers can sometimes increase purchases of one of the goods that were already available.

A well-known example is the asymmetric dominance effect, in which purchases of a target good are increased by introducing a decoy that is dominated by (i.e., worse than) the target good on both price and quality dimensions, but it is not worse on both dimensions than an alternative good that many consumers prefer to the target good in the absence of the decoy (Huber et al. 1982, Heath & Chatterjee 1995). Such effects are not consistent with a random utility model in which the distribution of possible valuations for each good is the same regardless of what other goods are available in the choice set.

Context dependence is, however, an ubiquitous feature of perceptual judgments (Laming 2011). How long a line or bar appears to be, how tilted it appears to be, how fast it appears to be moving, and so on, depend on the features of other items that appear next to the line or bar or that have been seen just before it. Moreover, the context dependence observed in choice behavior is often directly analogous to the illusions observed in visual perception (Trueblood et al. 2013, Summerfield &

²⁹The argument is in some ways similar to the one proposed by Commons et al. (1991), in which hyperbolic discounting is attributed to the way in which the noise in memory traces increases with the passage of time.

³⁰On the relevance of strategic complementarity for the aggregate effects of cognitive imprecision, readers are referred to Morris & Shin (2006), Tirole (2015), and Angeletos & Huo (2019).

Tsetsos 2015), suggesting that similar mechanisms may be responsible in both cases. A variety of mechanisms have been proposed for context effects in perceptual domains, and some of these are consistent with the Bayesian view of perception presented above (e.g., Schwartz et al. 2007).

As proposed in the previous section, valuations derived from optimal Bayesian inference from noisy internal representations might depend on the other items in a choice set for either of two reasons. First, even supposing that the internal representation r_i of the features of good i has a marginal distribution that is independent of the choice set, the noise in the internal representations of different goods might be correlated. As Natenzon (2019) shows, in this case the optimal value estimate $\hat{x}_i = E[x_i|\mathbf{r}]$ will generally depend on elements of the vector of internal representations other than just the element r_i that encodes the value of x_i . Natenzon proposes that differing degrees of noise correlation in the case of different pairs of goods provides a way of capturing the fact that goods that are more similar to each other are more easily compared (Tversky & Russo 1969), making choice less random for any degree of difference in the true values of the goods to the consumer.³¹ In this theory an effective decoy is a good that is easily comparable with the target good and inferior to it; adding a noisy representation of the value of the decoy to the vector \mathbf{r} can then increase the estimated value of the target good, with less of an effect on the estimated values of other goods that are less similar to the decoy.

Second, one need not assume that the elements of the internal representation each encode the value of some feature of only one good, considered in isolation; they might instead encode the relative value of some attribute of a given good, compared with one or more goods in a comparison set. This is a well-documented feature of neural representations of sensory information: What is encoded is often information about changes either in space or in time, as in the case of the contrast detectors, edge detectors, convexity detectors, and dimming detectors found in the frog retina (Lettvin et al. 1959). If one supposes that the internal representation \mathbf{r} includes imprecise measures of relative attribute values between different goods in the choice set, then a Bayesian model can easily predict decoy effects of the kind that are observed (Howes et al. 2016).

This raises a question: How freely should a modeler be allowed to specify the joint distribution of posited internal representations and the observable features of a choice situation? One way of developing a more parsimonious theory is by looking for similarities across domains in the structure of imprecise internal representations, as with the analogy in the model by Khaw et al. (2019), between the way that monetary payoffs are assumed to be encoded and the way that other numerical magnitudes, such as the numerosity of visual arrays, are encoded.

An alternative approach would seek to derive the form of the imprecise internal representation in any given case from a general theory of efficient coding—i.e., the proposition that internal representations have a form that maximizes the average accuracy of classifications, given the prior probability of being in different possible situations and subject to a constraint on the feasible complexity of representations, understood to follow from a limit on available cognitive resources.³² Theories of this kind are often used to explain aspects of the neural coding of particular stimulus features in the early stages of sensory processing (e.g., Laughlin 1981, Simoncelli 2003, Ganguli

³¹The possibility of nonzero correlation between r_1 and r_2 (conditional on the true values x_1, x_2) adds an additional parameter to the formula in Equation 1, derived above under the assumption of zero correlation.

³²The economic literature on rational inattention (e.g., Sims 2003, Caplin & Dean 2015) has a similar goal. Even more closely related are economic models that consider the optimal use of a finite range of possible classifications, such as those by Robson (2001), Rayo & Becker (2007), Netzer (2009), and Steiner & Stewart (2016). The latter kind of models, which are more consistent with evidence from psychophysics and neurophysiology, imply that nearby states must be difficult to distinguish from one another; this has important implications for the analysis of coordination games (Morris & Yang 2019) and optimal contracting (Hébert & Woodford 2018), among other issues.

& Simoncelli 2016) and have also been used to explain neural representations of the values of individual options in decision problems (Rustichini et al. 2017).

In the literature on perception, the encoding of changes or contrasts rather than absolute stimulus magnitudes is often argued to be efficient because it reduces the redundancy of the neural code, given the degree to which absolute magnitudes (such as light intensity) are correlated in space and time (Schwartz et al. 2007). There is a further reason for it to be efficient to encode comparisons rather than the absolute values of the attributes of individual goods in the case of consumer choice: The accuracy of a person's choices depends only on having an accurate view of the relative, rather than absolute, values of the different goods on offer, so that encoding absolute values would waste representational capacity on irrelevant information. At the same time, if limited capacity means that the information that is encoded about relative values must be imprecise, optimal Bayesian decoding can lead to preference reversals, as shown by Howes et al. (2016).

A theory of efficient coding together with the hypothesis of Bayesian decoding provides a parsimonious model of how the distribution of errors in perceptual judgments should depend on the statistics of a particular perceptual domain. Models of this kind have been used to explain patterns of perceptual bias in a range of different sensory domains (Wei & Stocker 2015, 2017) and have been proposed as explanations for biases in economic valuations as well. Among other applications, they have been used to explain how biases in economic valuations vary with time pressure (Polania et al. 2019) or framing (Woodford 2012) and how the sensitivity of subjective valuations to the objective situation depends on the distribution of values used in a given experiment (Frydman & Jin 2019, Payzan-LeNestour & Woodford 2019).

Such models offer the prospect of a unified theory of a range of different types of behavioral biases, context effects, and effects on economic decisions produced by factors, such as time pressure, that play no role in rational choice theory. Rather than implying that choices are irrational, this approach posits that these phenomena should be understood as consequences of patterns of mental processing that serve people well, in the sense of maximizing their rewards on average, subject to the constraints imposed by the finiteness of cognitive resources. Much work remains to be done to flesh out the details of this theory and test its empirical validity across varied contexts. But regardless of whether any of the particular formulations offered in the current literature proves generally valid, it is likely that study of the mechanisms responsible for biases in perceptual domains will be an important source of insights into the nature of biases in economic valuations as well.

DISCLOSURE STATEMENT

The author is not aware of any affiliations, memberships, funding, or financial holdings that might be perceived as affecting the objectivity of this review.

ACKNOWLEDGMENTS

The author would like to thank Rava Azeredo da Silveira, Andrew Caplin, Paul Glimcher, Mel Win Khaw, Ziang Li, Stephen Morris, Rafael Polania, Arthur Prat-Carrabin, Antonio Rangel, and Christian Ruff for helpful discussions.

LITERATURE CITED

Angeletos G-M, Huo Z. 2019. *Myopia and anchoring*. Work. Pap., Mass. Inst. Technol., Cambridge, MA. <https://economics.mit.edu/files/14946>

- Angeletos G-M, Werning I. 2006. Crises and prices: information aggregation, multiplicity, and volatility. *Am. Econ. Rev.* 96:1720–36
- Anobile G, Cicchini GM, Burr DC. 2012. Linear mapping of numbers onto space requires attention. *Cognition* 122:454–59
- Azaredo da Silveira R, Woodford M. 2019. Noisy memory and over-reaction to news. *AEA Pap. Proc.* 109:557–61
- Beck JM, Ma WJ, Pitkow X, Latham PE, Pouget A. 2012. Not noisy, just wrong: the role of suboptimal inference in behavioral variability. *Neuron* 74:30–39
- Beshears J, Choi JJ, Fuster A, Laibson D, Madrian BC. 2013. What goes up must come down? Experimental evidence on intuitive forecasting. *Am. Econ. Rev.* 103:570–74
- Bordalo P, Gennaioli N, Ma Y, Shleifer A. 2018. *Over-reaction in macroeconomic expectations*. NBER Work. Pap. 24932
- Cantlon JF, Brannon EM. 2006. Shared system for ordering small and large numbers in monkeys and humans. *Psychol. Sci.* 17:401–6
- Caplin A. 2012. Choice sets as percepts. In *Neuroscience of Preference and Choice*, ed. R Dolan, T Sharot, pp. 295–304. New York: Academic
- Caplin A, Dean M. 2015. Revealed preference, rational inattention, and costly information acquisition. *Am. Econ. Rev.* 105:2183–203
- Coibion O, Gorodnichenko Y. 2015. Information rigidity and the expectations formation process: a simple framework and new facts. *Am. Econ. Rev.* 105:2644–78
- Commons ML, Woodford M, Trudeau EJ. 1991. How each reinforcer contributes to value: “Noise” must reduce reinforcer value hyperbolically. In *Signal Detection: Mechanisms, Models and Applications*, ed. ML Commons, JA Nevin, MC Davison, pp. 139–68. Hillsdale, NJ: Lawrence Erlbaum Assoc.
- Cordes S, Gelman R, Gallistel CR, Whalen J. 2001. Variability signatures distinguish verbal from non-verbal counting for both large and small numbers. *Psychon. Bull. Rev.* 8:698–707
- Creelman CD. 2015. Signal detection theory, history of. In *International Encyclopedia of the Social & Behavioral Sciences*, ed. JD Wright, pp. 952–56. Amsterdam: Elsevier Press. 2nd ed.
- Dayan P, Abbott LF. 2001. *Theoretical Neuroscience*. Cambridge, MA: MIT Press
- Dehaene S. 2003. The neural basis of the Weber-Fechner law: a logarithmic mental number line. *Trends Cogn. Sci.* 7:145–47
- Dehaene S. 2011. *The Number Sense*. Oxford, UK: Oxford Univ. Press. Rev. updat. ed.
- Drugowitsch J, Wyart V, Devauchelle A-D, Kochlin E. 2016. Computational precision of mental inference as critical source of human choice suboptimality. *Neuron* 92:1398–411
- Fechner GT. 1966 (1860). *Elements of Psychophysics*. New York: Holt, Rinehart & Winston
- Frydman C, Jin LJ. 2019. *Efficient coding and risky choice*. Work. Pap., Univ. South. Calif., Los Angeles
- Gabaix X, Laibson D. 2017. *Myopia and discounting*. NBER Work. Pap. 23254
- Ganguli D, Simoncelli EP. 2016. Neural and perceptual signatures of efficient sensory coding. arXiv:1603.00058v1 [q-bio.NC]
- Garcia M, Grüschow M, Polania R, Woodford M, Ruff CC. 2018. *Predicting risk attitudes from the precision of mental number representation*. Poster presented at the 16th Annual Meeting of the Society for Neuroeconomics, Philadelphia
- Gerhardt H, Biele GP, Heekeren HR, Uhlir H. 2016. *Cognitive load increases risk aversion*. SFB 649 Discuss. Pap. 2016-011, Humboldt Univ., Berlin
- Gescheider GA. 1997. *Psychophysics: The Fundamentals*. Mahwah, NJ: Lawrence Erlbaum Assoc. 3rd ed.
- Glimcher PW. 2011. *Foundations of Neuroeconomic Analysis*. Oxford, UK: Oxford Univ. Press
- Goryunov M, Rigos A. 2019. *Discontinuous and continuous stochastic choice and coordination in the lab*. Work. Pap., Lund Univ., Lund, Swed.
- Green DM, Swets JA. 1966. *Signal Detection Theory and Psychophysics*. New York: Wiley. Repr., 1988. Los Altos, CA: Peninsula Publ.
- Gul F, Pesendorfer W, Strzalecki T. 2017. Coarse competitive equilibrium and extreme prices. *Am. Econ. Rev.* 107:109–37

- Heath TB, Chatterjee S. 1995. Asymmetric decoy effects on lower-quality versus higher-quality brands: meta-analytic and experimental evidence. *J. Consum. Res.* 22:268–84
- Hébert B, Woodford M. 2018. *Information costs and sequential information sampling*. NBER Work. Pap. 25316
- Hellwig C, Mukherji A, Tsyvinski A. 2006. Self-fulfilling currency crises: the role of interest rates. *Am. Econ. Rev.* 96:1769–87
- Hollingworth HL. 1909. *The Inaccuracy of Movement*. New York: Sci. Press
- Hollingworth HL. 1910. The central tendency of judgment. *J. Philos. Psychol. Sci. Methods* 7:461–69
- Howes A, Warren PA, Farmer G, El-Deredey W, Lewis RL. 2016. Why contextual preference reversals maximize expected value. *Psychol. Rev.* 123:368–91
- Huber J, Payne JW, Puto C. 1982. Adding asymmetrically dominated alternatives: violations of regularity and the similarity hypothesis. *J. Consum. Res.* 9:90–98
- Izard V, Dehaene S. 2008. Calibrating the mental number line. *Cognition* 106:1221–47
- Kahneman D, Tversky A. 1979. Prospect theory: an analysis of decision under risk. *Econometrica* 47:263–91
- Kaufman EL, Lord MW, Reese TW, Volkman J. 1949. The discrimination of visual number. *Am. J. Psychol.* 62:498–525
- Khaw MW, Li Z, Woodford M. 2019. *Cognitive imprecision and small-stakes risk aversion*. NBER Work. Pap. 24978
- Kingdom FAA, Prins N. 2010. *Psychophysics: A Practical Introduction*. New York: Academic
- Kramer P, De Bono MG, Zorzi M. 2011. Numerosity estimation in visual stimuli in the absence of luminance-based cues. *PLOS ONE* 6:e17378
- Krueger LE. 1984. Perceived numerosity: a comparison of magnitude production, magnitude estimation, and discrimination judgments. *Percept. Psychophys.* 35:536–42
- Laming D. 1997. *The Measurement of Sensation*. Oxford, UK: Oxford Univ. Press
- Laming D. 2011. *Human Judgment: The Eye of the Beholder*. Andover, UK: Cengage Learn.
- Landier A, Ma Y, Thesmar D. 2019. *Biases in expectations: experimental evidence*. Work. Pap., HEC, Paris
- Laughlin SB. 1981. A simple coding procedure enhances a neuron's information capacity. *Z. Naturforschung* 36:910–12
- Lettvin JY, Maturana HR, McCulloch WS, Pitts WH. 1959. What the frog's eye tells the frog's brain. *Proc. Inst. Radio Eng.* 47:1940–51
- Luce RD. 1956. Semiorders and a theory of utility discrimination. *Econometrica* 24:178–91
- Mackowiak B, Wiederholt M. 2009. Optimal sticky prices under rational inattention. *Am. Econ. Rev.* 99:769–803
- McFadden D. 1981. Econometric models of probabilistic choice. In *Structural Analysis of Discrete Data with Economic Applications*, ed. C Manski, D McFadden, pp. 198–272. Cambridge, MA: MIT Press
- Morris SE, Shin HS. 1998. Unique equilibrium in a model of self-fulfilling currency attacks. *Am. Econ. Rev.* 88:587–97
- Morris SE, Shin HS. 2003. Global games: theory and applications. In *Advances in Economics and Econometrics: Proceedings of the Eighth World Congress of the Econometric Society*, ed. M Dewatripont, L Hansen, S Turnovsky, pp. 56–114. Cambridge, UK: Cambridge Univ. Press
- Morris SE, Shin HS. 2006. Inertia of forward-looking expectations. *Am. Econ. Rev.* 96:152–57
- Morris SE, Yang M. 2019. *Coordination and continuous stochastic choice*. Work. Pap., Mass. Inst. Technol., Cambridge, MA
- Mosteller F, Nogee P. 1951. An experimental measurement of utility. *J. Political Econ.* 59:371–404
- Mullainathan S, Shafir E. 2013. *Scarcity: Why Having Too Little Means So Much*. New York: Henry Holt & Co.
- Natenzon P. 2019. Random choice and learning. *J. Political Econ.* 127:419–57
- Netzer N. 2009. Evolution of time preferences and attitudes toward risk. *Am. Econ. Rev.* 99:937–55
- Nieder A, Miller EK. 2003. Coding of cognitive magnitude: compressed scaling of numerical information in the primate prefrontal cortex. *Neuron* 37:149–57
- Payzan-LeNestour E, Woodford M. 2019. “Outlier blindness”: *Efficient coding generates an inability to represent extreme values*. Work. Pap., Univ. N. S. W., Sydney, Aust.
- Petzschner FH, Glasauer S, Stephan KE. 2015. A Bayesian perspective on magnitude estimation. *Trends Cogn. Sci.* 19:285–93

- Polania R, Woodford M, Ruff CC. 2019. Efficient coding of subjective value. *Nat. Neurosci.* 22:134–42
- Preston MG, Baratta P. 1948. An experimental study of the auction-value of an uncertain outcome. *Am. J. Psychol.* 61:183–93
- Rabin M. 2000. Risk aversion and expected-utility theory: a calibration theorem. *Econometrica* 68:1281–92
- Rayo L, Becker GS. 2007. Evolutionary efficiency and happiness. *J. Political Econ.* 115:302–37
- Robson AJ. 2001. The biological basis of economic behavior. *J. Econ. Lit.* 39:11–33
- Rubinstein A. 1988. Similarity and decision making under risk (is there a utility theory resolution to the Allais paradox?). *J. Econ. Theory* 46:145–53
- Rustichini A, Conen KE, Cai X, Padoa-Schioppa C. 2017. Optimal coding and neuronal adaptation in economic decisions. *Nat. Commun.* 8:1208
- Schwartz O, Hsu A, Dayan P. 2007. Space and time in visual context. *Nat. Rev. Neurosci.* 8:522–35
- Simoncelli EP. 2003. Vision and the statistics of the visual environment. *Curr. Opin. Neurobiol.* 13:144–49
- Sims CA. 2003. Implications of rational inattention. *J. Monet. Econ.* 50:665–90
- Steiner J, Stewart C. 2016. Perceiving prospects properly. *Am. Econ. Rev.* 106:1601–31
- Stocker AA, Simoncelli EP. 2006. Noise characteristics and prior expectations in human visual speed perception. *Nat. Neurosci.* 9:578–85
- Summerfield C, Tsetso K. 2015. Do humans make good decisions? *Trends Cogn. Sci.* 19:27–34
- Thurstone LL. 1927. A law of comparative judgment. *Psychol. Rev.* 34:273–86
- Tirole J. 2015. *Cognitive games and cognitive traps*. Work. Pap., Toulouse Sch. Econ., Toulouse, Fr. https://www.tse-fr.eu/sites/default/files/TSE/documents/doc/by/tirole/cognitive_games_210715.pdf
- Trueblood JS, Brown SD, Heathcote A, Busemeyer JR. 2013. Not just for consumers: Context effects are fundamental to decision making. *Psychol. Sci.* 24:901–8
- Tversky A, Russo JE. 1969. Substitutability and similarity in binary choices. *J. Math. Psychol.* 6:1–12
- Weber EU. 2004. Perception matters: psychophysics for economists. In *The Psychology of Economic Decisions*, Vol. 2: *Reasons and Choices*, ed. I Brocas, JD Carrillo, pp. 163–76. Oxford, UK: Oxford Univ. Press
- Wei X-X, Stocker AA. 2015. A Bayesian observer model constrained by efficient coding can explain “anti-Bayesian” percepts. *Nat. Neurosci.* 18:1509–17
- Wei X-X, Stocker AA. 2017. Lawful relation between perceptual bias and discriminability. *PNAS* 114:10244–49
- Whalen J, Gallistel CR, Gelman R. 1999. Non-verbal counting in humans: the psychophysics of number representation. *Psychol. Sci.* 10:130–37
- Woodford M. 2003. Imperfect common knowledge and the effects of monetary policy. In *Knowledge, Information, and Expectations in Modern Macroeconomics*, ed. P Aghion, R Frydman, J Stiglitz, M Woodford, pp. 25–58. Princeton, NJ: Princeton Univ. Press
- Woodford M. 2012. Prospect theory as efficient perceptual distortion. *Am. Econ. Rev.* 102:41–46

Contents

Economics with a Moral Compass? Welfare Economics: Past, Present, and Future <i>Amartya Sen, Angus Deaton, and Timothy Besley</i>	1
Trade Policy in American Economic History <i>Douglas A. Irwin</i>	23
An Econometric Perspective on Algorithmic Subsampling <i>Sokbae Lee and Serena Ng</i>	45
Behavioral Implications of Causal Misperceptions <i>Ran Spiegler</i>	81
Poverty and the Labor Market: Today and Yesterday <i>Robert C. Allen</i>	107
The Econometrics of Static Games <i>Andrés Aradillas-López</i>	135
On Measuring Global Poverty <i>Martin Ravallion</i>	167
Taxation and the Superrich <i>Florian Scheuer and Joel Slemrod</i>	189
How Distortions Alter the Impacts of International Trade in Developing Countries <i>David Atkin and Amit K. Khandelwal</i>	213
Robust Decision Theory and Econometrics <i>Gary Chamberlain</i>	239
Cities in the Developing World <i>Gharad Bryan, Edward Glaeser, and Nick Tsivanidis</i>	273
New Developments in Revealed Preference Theory: Decisions Under Risk, Uncertainty, and Intertemporal Choice <i>Federico Echenique</i>	299
Computing Economic Equilibria Using Projection Methods <i>Alena Miftakbova, Karl Schmedders, and Malte Schumacher</i>	317

Social Identity and Economic Policy	
<i>Moses Shayo</i>	355
Empirical Models of Lobbying	
<i>Matilde Bombardini and Francesco Trebbi</i>	391
Political Effects of the Internet and Social Media	
<i>Ekaterina Zhuravskaya, Maria Petrova, and Ruben Enikolopov</i>	415
Nash Equilibrium in Discontinuous Games	
<i>Philip J. Reny</i>	439
Revealed Preference Analysis of School Choice Models	
<i>Nikhil Agarwal and Paulo Somaini</i>	471
Social Networks and Migration	
<i>Kaivan Munshi</i>	503
Informality: Causes and Consequences for Development	
<i>Gabriel Ulyssea</i>	525
The Theory and Empirics of the Marriage Market	
<i>Pierre-André Chiappori</i>	547
Modeling Imprecision in Perception, Valuation, and Choice	
<i>Michael Woodford</i>	579
Peer Effects in Networks: A Survey	
<i>Yann Bramoullé, Habiba Djebbari, and Bernard Fortin</i>	603
Alternative Work Arrangements	
<i>Alexandre Mas and Amanda Pallais</i>	631
Shotgun Wedding: Fiscal and Monetary Policy	
<i>Marco Bassetto and Thomas J. Sargent</i>	659
Social Identity, Group Behavior, and Teams	
<i>Gary Charness and Yan Chen</i>	691
Aspirations and Economic Behavior	
<i>Garance Genicot and Debraj Ray</i>	715
The Search Theory of Over-the-Counter Markets	
<i>Pierre-Olivier Weill</i>	747
Econometric Models of Network Formation	
<i>Aureo de Paula</i>	775
Dynamic Taxation	
<i>Stefanie Stantcheva</i>	801
Capital Flows and Leverage	
<i>Şebnem Kalemli-Özcan and Jun Hee Kwak</i>	833