

Modeling Imprecision in Perception, Valuation and Choice*

Michael Woodford
Columbia University

September 15, 2019

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.

*Prepared for the *Annual Review of Economics*, volume 12. I 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.

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.¹

But how should imprecision in people's recognition of their true situations be introduced into economic models? This review will argue 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 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.²

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 1 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 discusses the applicability of similar models to the problem of stochastic choice. Section 2 then considers more complex experimental designs in which the magnitude of a single stimulus is estimated, 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. It is argued that the

¹See Luce (1956) and Rubinstein (1988) for early discussions of this issue.

²For alternative discussions of possible lessons for economics, see Weber (2004) or Caplin (2012).

estimation biases that are observed in sensory contexts provide a potential explanation for patterns of choice behavior that are instead often attributed to non-standard preferences. And finally, section 3 discusses ways in which judgments about a given sensory magnitude can be influenced by the context in which it appears; it is proposed that similar mechanisms can explain choice behavior that appears inconsistent with the existence of any coherent preference ordering.

1 The Stochasticity of Comparative Judgments

A first important lesson of the psychophysics literature is that not only are people (or other organisms) 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 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 (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 (“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.”³

Figure 1 provides an example of such plots (from Krueger, 1984), when the stimulus feature is *numerosity*, 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,⁴ it has the advantage (for purposes of this review) of shedding light on the

³See, e.g., Gescheider (1997, chap. 3), Glimcher (2011, chap. 4), or Kingdom and Prins (2010, chap. 4).

⁴In 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




Figure 1: Psychometric functions for comparisons of numerosity. The number of items in the reference array is 25, 100, or 400 in panels (a), (b), and (c) respectively. (Reproduced from Krueger, 1984.)

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, 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 function of the true difference in numerosity $n_2 - n_1$,⁵ 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 (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 values, so that one can meaningfully speak of response probabilities as varying continuously with the true stimulus magnitude.

⁵This difference is reported in “steps.” A “one step” increase means one more X in the top panel, three more in the middle panel, and nine more in the bottom panel.

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 that 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 due to Thurstone (1927) assumes that the internal representation can be summarized by a single real number, and that it is drawn from a normal distribution $N(m(x), \nu^2)$, where $m(x)$ is an increasing function of the true magnitude, and the standard deviation $\nu > 0$ is independent of the true magnitude.⁶ 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 $r_2 > r_1$.⁷ 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

$$\text{Prob}[\text{“}x_2 \text{ greater”} | x_1, x_2] = \Phi\left(\frac{m(x_2) - m(x_1)}{\sqrt{2}\nu}\right), \quad (1.1)$$

where $\Phi(z)$ is the CDF 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 $m(x)$.⁸

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

⁶This is Thurstone’s celebrated “Case V.”

⁷Here we assume a “two-alternative forced choice” experimental design, in which the subject must select one of the 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 ((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.

⁸Data on the frequency with which different comparative judgments are made, such as that plotted in Figure 1, can allow the identification of this scale up to an affine transformation.

response probabilities are z -transformed,⁹ equation (1.1) implies that one should have

$$z(\text{Prob}) = \frac{m(x_2) - m(x_1)}{\sqrt{2\nu}} \approx \frac{m'(x_1)}{\sqrt{2\nu}} \cdot (x_2 - x_1), \quad (1.2)$$

for values of x_2 sufficient close to the reference magnitude x_1 . Thus when the relationship is plotted as in Figure 1,¹⁰ the relationship should be approximately linear, as shown in the figure — 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.

The figure not only shows that the psychometric function in each case is roughly of the predicted form (1.2), but 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 $n_1 = 25$, but 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 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 (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 .¹¹

⁹That is, the probabilities are replaced by $z(p) \equiv \Phi^{-1}(p)$, the inverse of the function $p = \Phi(z)$.

¹⁰Note that in each of the panels of the figure, 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)

¹¹Even 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. See, for example, the plots of how discrimination thresholds vary with stimulus magnitude in a variety of sensory domains in Ganguli and Simoncelli (2016). Weber’s Law has sometimes been asserted also to hold for the perception of numerosity (e.g., Nieder and Miller, 2003; Dehaene, 2003; Cantlon and 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

1.1 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.¹²

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 benchmark) models in which the decoding is assumed to represent an *optimal* inference from the evidence provided by the internal representation.¹³

Why should one wish to complicate one’s model of imprecise comparisons in this way, rather than simply directly 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

threshold increases with the reference numerosity with an elasticity that is instead about 0.85.

¹²Dayan and Abbott (2001, chap. 3) provide a textbook discussion.

¹³Our 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 7.

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.¹⁴

But 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 giving a particular response is not a function solely 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,¹⁵ 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,¹⁶ 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 on both the decision maker's prior and on the objective (or reward) function that he/she seeks to maximize.

Signal detection theory (Green and 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 and Swets) shows a classic application of the theory. The figure plots data from a single subject, in an experiment in which on each trial one of two auditory stimuli (denoted s and n) are presented, and the subject's task is to indicate which

¹⁴See Dayan and Abbott (2001, chap. 3) for examples.

¹⁵It was assumed above, in our discussion of the normative basis for a particular rule for determining the perceptual judgment (footnote 7), that either stimulus was equally likely to be the greater one, and this is true in many experiments. But 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.

¹⁶This is a common simplifying assumption, but in fact context can influence encoding as well, as discussed in section 3.




Figure 2: Conditional response probabilities in a signal detection task. The tradeoff is shown between the “hit rate” (vertical axis) and “false alarm rate” (horizontal axis), as one varies the prior probability of occurrence of the two stimuli (left panel), or the relative rewards for correct identification of the two stimuli (right panel). In each case, the efficient frontier (“ROC curve”) is shown by the bowed solid curve. (Reproduced from Green and Swets, 1966.)

is the case (making one of the two possible responses, S or N).¹⁷ In each of several blocks of trials, the same two stimuli are used (the stimulus presented on each trial is always either s or n), 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 “Type I” as opposed to “Type II” errors. Each of the circles in the figure plots the subject’s conditional response probabilities for one block of trials.

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 figure 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

¹⁷The 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).

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 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 the figure one sees systematic effects of both aspects of the context. In the left panel, 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 the right panel, 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 (by giving response S or N). Suppose, as in the model above, that r is a single real number,¹⁸ 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 ($i = s$ 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

¹⁸In this kind of task, because there are only two possible true situations (s or n), even if the internal representation were high-dimensional, it clearly suffices to describe it using a single real number as a sufficient statistic, namely the likelihood ratio of the two hypotheses given the noisy evidence. See discussion in Green and Swets (1966).

an upper bound on the probability of type II errors, or vice versa — is necessarily a likelihood ratio test,¹⁹ 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” or 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 in response to an increased reason to prioritize an increased hit rate even at the expense of an increase in the 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, which 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 auditory stimuli before expressing a judgment); if this changes the value of ν , this should shift the location of the ROC curve along which the subject should operate (whatever the prior and incentives may be). And 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 where 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

¹⁹This follows from the Neyman-Pearson lemma of statistical decision theory; again see Green and Swets (1966).

to model the subject’s judgments as the outcome of two separate processes of encoding and decoding.²⁰

1.2 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 assume that an agent should know with certainty to which element of some partition of the state space the current state belongs.²¹ In the case of a continuous state space, such a model implies that certain states that differ only infinitesimally should nonetheless be perfectly 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 wants to capture should be analogous to 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”²² 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 non-zero). 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 which is an independent draw for each agent from some (low-

²⁰Green and Swets (1966, chap. 4) call this the “separation of sensory and decision processes.”

²¹See, e.g., Gul, Pesendorfer and Strzalecki (2017).

²²See, e.g., Morris and Shin (1998, 2003).

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 — but 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. The question whether there should be sufficient common knowledge for it to be possible for agents to coordinate on multiple equilibria is then argued to depend 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.²³ 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 and Shin (1998) should be of more general relevance than is often appreciated.²⁴

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

²³See, e.g., Angeletos and Werning (2006), or Hellwig, Mukherji and Tsyvinski (2006).

²⁴Goryunov and 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. But 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. (Based on a figure in Mosteller and Noguee, 1951.)

Standard models of rational 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 and Noguee (1951), shows a classic example. The figure plots data on the choices of a single subject, who was presented on different trials during the same experiment with multiple variants of the same kind of gamble: whether the subject would pay 5 cents in order to obtain a random outcome, equal to an amount X with probability one-half, and otherwise zero. 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 the figure, for several intermediate values of X , it is neither the

case that the subject consistently accepts the gamble nor that he consistently rejects it. The figure illustrates how they 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 the fact that the subject was exactly indifferent in this case.

A graph like Figure 3 is highly reminiscent of psychometric functions like those in Figure 1.²⁵ This suggests that the randomness depicted might fruitfully be modeled in a similar way. And indeed, a common approach within economics has been to model stochastic choice using 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 on that occasion. (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 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

²⁵Indeed, it seems likely that Mosteller and Nogee’s 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 use 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 and Prins (2010, p. 19).

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 assessment of people’s level of welfare under alternative possible regulations of market transactions.

And 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.²⁶ 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 had been reached is communicated.²⁷ 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, than under 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 3, the likelihood of choosing one good over another can be influenced by contextual factors that should have no effect on the true values of

²⁶This is implicit in an approach, like that of Mosteller and Noguee, that assumes that which of two options is more often chosen tells one which of the options is preferred, and so should allow a deterministic preference ordering to be recovered even when choice is stochastic.

²⁷Even 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 in which 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 of Mosteller and Noguee, where the data are presented in symbolic form. See further discussion in Khaw *et al.* (2019).

either item to the decision maker.

2 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, allowing 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 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 to them — for example, producing two successive taps to indicate the length of a time interval.²⁸ 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 over-estimate smaller magnitudes (on average) while under-estimating larger ones.

Figure 4 illustrates this bias, in 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 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 of a variety of true distances presented by the experimenter. The different symbols identify distinct experimental sessions, in which the range of true distances (presented in random order) was different: 10-70mm in series A of trials, 30-150mm in series B, and 70-250mm in series C. The black dots represent three sessions in which all true distances were of exactly the same length:

²⁸Both 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, discussed below.




Figure 4: Mean of the length estimates produced by an experimental subject, plotted as a function of the length (in mm) that had previously been demonstrated to the subject. The different symbols indicate sessions in which different ranges of lengths were used as stimuli. (Reproduced from Laming, 1997.)

10mm, 70mm, or 250mm.

In each of the first three sessions, the subject's estimates exhibit a clear conservative bias: the shorter distances used in that day's series are over-estimated on average, while the longer distances are under-estimated 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 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 70mm movement) may be under-estimated, over-estimated, or estimated with nearly zero bias, depending whether it is unusually long, unusually short, or about average among the stimuli used in the session. Hollingworth found the same to be true in a number of different sensory domains, and christened this




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. (Reproduced from Anobile *et al.*, 2012.)

regularity “the central tendency of judgment” (Hollingworth, 1910).²⁹

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 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 one, as in Figure 4.³⁰ However, the cross-over point, at which numerosity begins to be under-estimated rather than over-estimated, differs considerably across experiments, in a way that correlates with differences in the range of numerosities used as stimuli in the different experiments.³¹ Figure 5 shows an example of two experiments that differ only in the range of numerosities used (1-30 in the left panel, 1-100 in the

²⁹See Petzschner, Glasauer and Stephan (2015) for more recent examples from a variety of sensory domains.

³⁰See, e.g., Krueger (1984) or Kramer, De Bono and Zorzi (2011).

³¹See Izard and Dehaene (2008) for discussion.

right); note how the location of the cross-over point shifts, in a way consistent with the central tendency of judgment.

2.1 A Bayesian Model of Estimation Bias

A number of 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.³² 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.³³

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

$$r \sim N(\log x, \nu^2). \quad (2.1)$$

(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,

$$\log x \sim N(\mu, \sigma^2), \quad (2.2)$$

then the expected value of x , conditional on the representation r (i.e., the estimate given by the Bayesian posterior mean),³⁴ will equal

$$\hat{x}(r) = E[x|r] = \exp((1 - \beta) \log \bar{x} + \beta r), \quad (2.3)$$

³²See, e.g., Stocker and Simoncelli (2006), Petzschner *et al.* (2015), and Wei and Stocker (2015, 2017).

³³Of 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.

³⁴This 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 and Abbott, 2001, chap. 4), but is used by authors such as Wei and Stocker (2015).

where $\beta \equiv \sigma^2/(\sigma^2+\nu^2) < 1$ 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

$$e(x) \equiv \mathbb{E}[\hat{x}|x] = Ax^\beta, \quad \text{var}[\hat{x}|x] = Be(x)^2, \quad (2.4)$$

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

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 one 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 cross-over 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,³⁵ including estimates of numerosity.³⁶

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, but strictly concave function of x ,³⁷ 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.³⁸

³⁵Again see Petzschner *et al.* (2015).

³⁶See, e.g., Whalen *et al.* (1999), Cordes *et al.* (2001), Izard and Dehaene (2008), or Kramer *et al.* (2011).

³⁷Recall that Figure 4 is a log-log plot, so that a straight line corresponds to a power law of the kind described in (2.4).

³⁸Nor is either of these functions characterized by diminishing sensitivity in all cases. For dis-

2.2 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. And when one takes this approach, one finds that systematic behavioral tendencies that are commonly taken to reflect preferences can sometimes instead be interpreted as biases resulting from inference from noisy internal representations.

2.2.1 Application: Explaining Small-Stakes Risk Aversion

A common observation in laboratory experiments is apparently risk-averse behavior. The data from Mosteller and Nogiee (1951), plotted in Figure 3, provide an example: in this case, the gamble is a fair bet when X equals 10 cents (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. But this interpretation of the Mosteller and Nogiee 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 5.0 cents would reduce it; and while logically possible, this is actually quite an extreme degree of curvature of the utility of wealth function, and would require extraordinary risk aversion with respect to larger gambles (as explained by Rabin, 2000), of a kind that is seldom

cussion 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, see Wei and Stocker (2017).

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 probability $1/2$, but has probability $1/2$ of paying nothing; and suppose that on each of a series of trials of this kind, the values of C and X vary.³⁹ Suppose further that the amounts C and X that define the choice problem on a given trial each have a noisy internal representation r_i (for $i = C, X$) drawn from a distribution (2.1), 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, and 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 (2.2), then (2.3) 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 (2.1), the probability of acceptance is predicted to be

$$\text{Prob}[\textit{accept}] = \Phi \left(\frac{\log(X/C) - \beta^{-1} \log 2}{\sqrt{2\nu}} \right). \quad (2.5)$$

Equation (2.5) 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 Noguee should be where $X^{\textit{indiff}}/C = 2^{1/\beta} > 2$, a point to the right of the value corresponding to a fair bet, so that the subject should

³⁹Note that the data plotted in Figure 3 involve only one value of C (5 cents). However, multiple values of C were used in their experiment; the figure shows only how the value of X required for indifference is computed for one particular value of C .

appear to be risk-averse. Thus the theory provides a unified explanation for both the randomness of observed choices and 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 non-trivial 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 bias 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 (2.5), when the parameters σ and ν are fit to the data from the figure of Mosteller and Nogee.⁴⁰ 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 the Mosteller-Nogee figure 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 that indicated by the randomness of his 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 aversion). The required degree of imprecision in number representation is also relatively modest compared to that found in studies of numerosity perception.⁴¹

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

⁴⁰Note that under the assumption that C and X have the same prior distribution, the value of μ does not affect the predicted choice frequencies (2.5).

⁴¹For example, the data shown in Figure 1 imply values of ν on the order of 0.13, if the data are fit to equation (1.1) with $m(x) = \log x$. Note that it makes sense to expect greater noise in the case that the number is presented visually, by an array of Xs, rather than symbolically.

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 subjects 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 numbers in ones mind reduces the available capacity for precise internal representation of the numbers involved in the choice under risk. Finally, Frydman and 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., makes 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 ,⁴² while it would have no obvious explanation if one attributes random choice to randomness in the process of comparison of options after their expected values have been correctly computed.

2.2.2 Other Economic Applications

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

⁴²As Frydman and Jin discuss, a theory of efficient coding (section 3 below) 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 (2.5).

(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, that can be modeled as optimal inference from a noisy internal representation. Gabaix and Laibson (2017) show that optimal inference from noisy mental simulations of the future outcomes resulting from alternative choices can result in under-weighting 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 a hyperbolic function of distance in time.⁴³ 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 and 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.⁴⁴ 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 fully-informed optimizing behavior subject to 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 and Wiederholt, 2009).

⁴³The argument is in some ways similar to the proposal of Commons, Woodford and Trudeau (1991), in which hyperbolic discounting is attributed to the way in which the noise in memory traces increases with the passage of time.

⁴⁴On the relevance of strategic complementarity for the aggregate effects of cognitive imprecision, see also Morris and Shin (2006), Tirole (2015), and Angeletos and Huo (2019).

More generally, noisy internal representations can explain biases in expectations relative to those implied by the conventional hypothesis of rational expectations. Coibion and 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 and 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 over-extrapolation 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, Ma, and Thesmar, 2019).

3 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 the 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 was 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 the target good (worse on both price and quality dimensions), while not worse on both dimensions than an alternative good that many consumers prefer to the target good in the absence of the decoy (Huber, Payne and Puto, 1982; Heath and Chatterjee, 1995). Such effects are not consistent with a random-utility model in which the distribution of possible valuations for each good are 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 illusions that are observed in visual perception (Trueblood *et al.*, 2013; Summerfield

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

Valuations derived from optimal Bayesian inference from noisy internal representations, as proposed in the previous section, 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 and Russo, 1969), making choice less random for any degree of difference in the true values of the goods to the consumer.⁴⁵ 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.

And 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 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 vector \mathbf{r} that is conditioned on 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).

⁴⁵The possibility of non-zero correlation between r_1 and r_2 (conditional on the true values x_1, x_2) adds an additional parameter to the formula (1.1), derived above under the assumption of zero correlation.

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 of Khaw, Li and Woodford (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” — 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.⁴⁶ 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 and 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, Hsu and Dayan, 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* values of the different goods on offer, rather than their absolute values, so that encoding absolute values would waste representational capacity on irrelevant

⁴⁶The economics literature on “rational inattention” (e.g., Sims, 2003; Caplin and 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 of Robson (2001), Rayo and Becker (2007), Netzer (2009), and Steiner and 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 and Yang, 2019) and optimal contracting (Hébert and Woodford, 2018), among other issues.

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 theory 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 and 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, Woodford, and Ruff, 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 and Jin, 2019; Payzan-LeNestour and 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 of factors such as time pressure that play no role in rational choice theory. Rather than implying that choices are “irrational,” these phenomena would 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 the theory and test its empirical validity across varied contexts. But regardless of whether any of the particular formulations offered in the current literature prove 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.

References

- [1] Angeletos, George-Marios, and Zhen Huo, “Myopia and Anchoring,” working paper, May 2019.
- [2] Angeletos, George-Marios, and Ivan Werning, “Crises and Prices: Information Aggregation, Multiplicity, and Volatility,” *American Economic Review* 96: 1720-1736 (2006).
- [3] Anobile, Giovanni, Guido Marco Cicchini, and David C. Burr, “Linear Mapping of Numbers onto Space Requires Attention,” *Cognition* 122: 454-459 (2012).
- [4] Azeredo da Silveira, Rava, and Michael Woodford, “Noisy Memory and Over-Reaction to News,” *AEA Papers and Proceedings* 109: 557-561 (2019).
- [5] Beck, Jeffrey M., Wei Ji Ma, Xaq Pitkow, Peter E. Latham, and Alexandre Pouget, “Not Noisy, Just Wrong: The Role of Suboptimal Inference in Behavioral Variability,” *Neuron* 74: 30-39 (2012).
- [6] Beshears, John, James J. Choi, Andreas Fuster, David Laibson, and Brigitte C. Madrian, “What Goes Up Must Come Down? Experimental Evidence on Intuitive Forecasting,” *American Economic Review* 103: 570-74 (2013).
- [7] Bordalo, Pedro, Nicola Gennaioli, Yueran Ma, and Andrei Shleifer, “Over-Reaction in Macroeconomic Expectations,” NBER Working Paper no. 24932, August 2018.
- [8] Cantlon, Jessica F., and Elizabeth M. Brannon, “Shared System for Ordering Small and Large Numbers in Monkeys and Humans,” *Psychological Science* 17: 401-406 (2006).
- [9] Caplin, Andrew, “Choice Sets as Percepts,” in R. Dolan and T. Sharot, eds., *Neuroscience of Preference and Choice*, New York: Academic Press, 2012.
- [10] Caplin, Andrew, and Mark Dean, “Revealed Preference, Rational Inattention, and Costly Information Acquisition,” *American Economic Review* 105: 2183-2203 (2015).

- [11] Coibion, Olivier, and Yuriy Gorodnichenko, “Information Rigidity and the Expectations Formation Process: A Simple Framework and New Facts,” *American Economic Review* 105: 2644-2678 (2015).
- [12] Commons, Michael L., Michael Woodford, and Edward J. Trudeau, “How Each Reinforcer Contributes to Value: ‘Noise’ Must Reduce Reinforcer Value Hyperbolically,” in M.L. Commons, J.A. Nevin, and M.C. Davison, *Signal Detection: Mechanisms, Models and Applications*, Hillsdale, NJ: Lawrence Erlbaum Associates, 1991.
- [13] Cordes, Sara, Rochel Gelman, Charles R. Gallistel, and John Whalen, “Variability Signatures Distinguish Verbal from Non-Verbal Counting for Both Large and Small Numbers,” *Psychonomic Bulletin and Review* 8: 698-707 (2001).
- [14] Creelman, Carleton Douglas, “Signal Detection Theory, History of,” in J.D. Wright, ed., *International Encyclopedia of the Social & Behavioral Sciences*, Amsterdam: Elsevier Press, 2d ed., 2015.
- [15] Dayan, Peter, and L.F. Abbott, *Theoretical Neuroscience*, Cambridge, MA: MIT Press, 2001.
- [16] Dehaene, Stanislas, “The Neural Basis of the Weber-Fechner Law: A Logarithmic Mental Number Line,” *Trends in Cognitive Sciences* 7: 145-147 (2003).
- [17] Dehaene, Stanislas, *The Number Sense*, revised and updated edition, Oxford: Oxford University Press, 2011.
- [18] Drugowitsch, Jan, Valentin Wyart, Anne-Dominique Devauchelle, and Etienne Koechlin, “Computational Precision of Mental Inference as Critical Source of Human Choice Suboptimality,” *Neuron* 92: 1-14 (2016).
- [19] Fechner, Gustav T., *Elements of Psychophysics*, Leipzig: Breitkopf und Bartel, 1860. [English translation published by Holt, Rinehart and Winston, 1966.]
- [20] Frydman, Cary, and Lawrence J. Jin, “Efficient Coding and Risky Choice,” working paper, July 2019.
- [21] Gabaix, Xavier, and David Laibson, “Myopia and Discounting,” NBER Working Paper no. 23254, March 2017.

- [22] Ganguli, Deep, and Eero P. Simoncelli, “Neural and Perceptual Signatures of Efficient Sensory Coding,” working paper, February 2016. [Posted on *arXiv*: 1603.00058v1.]
- [23] Garcia, Manuel, Marcus Grüschow, Rafael Polania, Michael Woodford, and Christian C. Ruff, “Predicting Risk Attitudes from the Precision of Mental Number Representation,” poster presented at the 16th Annual Meeting of the Society for Neuroeconomics, Philadelphia, PA, USA, October 5-7, 2018.
- [24] Gerhardt, Holger, Guido P. Biele, Hauke R. Heekeren, and Harald Uhlig, “Cognitive Load Increases Risk Aversion,” SFB 649 Discussion Paper no. 2016-011, Humboldt University Berlin, March 2016.
- [25] Gescheider, George A., *Psychophysics: The Fundamentals*, 3d ed., Mahwah, NJ: Lawrence Erlbaum Associates, 1997.
- [26] Glimcher, Paul W., *Foundations of Neuroeconomic Analysis*, Oxford: Oxford University Press, 2011.
- [27] Green, David M., and John A. Swets, *Signal Detection Theory and Psychophysics*, New York: Wiley, 1966.
- [28] Gul, Faruk, Wolfgang Pesendorfer, and Tomasz Strzalecki, “Coarse Competitive Equilibrium and Extreme Prices,” *American Economic Review* 107: 109-137 (2017).
- [29] Goryunov, Maxim, and Alexandros Rigos, “Discontinuous and Continuous Stochastic Choice and Coordination in the Lab,” working paper, September 2019.
- [30] Heath, Timothy B., and Subimal Chatterjee, “Asymmetric Decoy Effects on Lower-Quality versus Higher-Quality Brands: Meta-Analytic and Experimental Evidence,” *Journal of Consumer Research* 22: 268-284 (1995).
- [31] Hébert, Benjamin, and Michael Woodford, “Information Costs and Sequential Information Sampling,” NBER Working Paper no. 25316, November 2018.

- [32] Hellwig, Christian, Arijit Mukherji, and Aleh Tsyvinski, “Self-Fulfilling Currency Crises: The Role of Interest Rates,” *American Economic Review* 96: 1769-1787 (2006).
- [33] Hollingworth, Harold L., *The Inaccuracy of Movement*, Columbia Contributions, vol. 17, no. 3, June 1909.
- [34] Hollingworth, Harold L., “The Central Tendency of Judgment,” *Journal of Philosophy* 7: 461-469 (1910).
- [35] Howes, Andrew, Paul A. Warren, George Farmer, Wael El-Deredy, and Richard L. Lewis, “Why Contextual Preference Reversals Maximize Expected Value,” *Psychological Review* 123: 368-391 (2016).
- [36] Huber, Joel, John W. Payne, and Christopher Puto, “Adding Asymmetrically Dominated Alternatives: Violations of Regularity and the Similarity Hypothesis,” *Journal of Consumer Research* 9: 90-98 (1982).
- [37] Izard, Véronique, and Stanislas Dehaene, “Calibrating the Mental Number Line,” *Cognition* 106: 1221-1247 (2008).
- [38] Kahneman, Daniel, and Amos Tversky, “Prospect Theory: An Analysis of Decision Under Risk,” *Econometrica* 47: 263-291 (1979).
- [39] Kaufman, E.L., M.W. Lord, T.W. Reese, and J. Volkman, “The Discrimination of Visual Number,” *American Journal of Psychology* 62: 498-525 (1949).
- [40] Khaw, Mel Win, Ziang Li, and Michael Woodford, “Cognitive Imprecision and Small-Stakes Risk Aversion,” NBER Working Paper no. 24978, revised January 2019.
- [41] Kingdom, Frederick A.A., and Nicolaas Prins, *Psychophysics: A Practical Introduction*, New York: Academic Press, 2010.
- [42] Kramer, Peter, Maria Grazia De Bono, and Marco Zorzi, “Numerosity Estimation in Visual Stimuli in the Absence of Luminance-Based Cues,” *PLoS ONE* 6(2): e17378 (2011).

- [43] Krueger, Lester E., “Perceived Numerosity: A Comparison of Magnitude Production, Magnitude Estimation, and Discrimination Judgments,” *Perception and Psychophysics* 35: 536-542 (1984).
- [44] Laming, Donald, *The Measurement of Sensation*, Oxford: Oxford University Press, 1997.
- [45] Laming, Donald, *Human Judgment: The Eye of the Beholder*, Andover, UK: Cengage Learning, 2011.
- [46] Landier, Augustin, Yueran Ma, and David Thesmar, “Biases in Expectations: Experimental Evidence,” working paper, May 2019.
- [47] Laughlin, Simon B., “A Simple Coding Procedure Enhances a Neuron’s Information Capacity,” *Zeitschrift für Naturforschung* 36: 910-912 (1981).
- [48] Lettvin, Jerome Y., Humberto R. Maturana, Warren S. McCulloch, and Walter H. Pitts, “What the Frog’s Eye Tells the Frog’s Brain,” *Proceedings of the Institute of Radio Engineers* 47: 1940-1951 (1959).
- [49] Luce, R. Duncan, “Semiorders and a Theory of Utility Discrimination,” *Econometrica* 24: 178-191 (1956).
- [50] Mackowiak, Bartosz, and Mirko Wiederholt, “Optimal Sticky Prices Under Rational Inattention,” *American Economic Review* 99: 769-803 (2009).
- [51] McFadden, Daniel, “Econometric Models of Probabilistic Choice,” in C. Manski and D. McFadden, eds., *Structural Analysis of Discrete Data with Economic Applications*, Cambridge, MA: MIT Press, 1981.
- [52] Morris, Stephen E., and Hyun Song Shin, “Unique Equilibrium in a Model of Self-Fulfilling Currency Attacks,” *American Economic Review* 88: 587-597 (1998).
- [53] Morris, Stephen E., and Hyun Song Shin, “Global Games: Theory and Applications,” in M. Dewatripont, L. Hansen and S. Turnovsky, eds., *Advances in Economics and Econometrics: Proceedings of the Eighth World Congress of the Econometric Society*, Cambridge, England: Cambridge University Press, 2003.

- [54] Morris, Stephen E., and Hyun Song Shin, “Inertia of Forward-Looking Expectations,” *American Economic Review* 96: 152-157 (2006).
- [55] Morris, Stephen E., and Ming Yang, “Coordination and Continuous Stochastic Choice,” working paper, February 2019.
- [56] Mosteller, Frederick, and Philip Noguee, “An Experimental Measurement of Utility,” *Journal of Political Economy* 59: 371-404 (1951).
- [57] Mullainathan, Sendhil, and Eldar Shafir, *Scarcity: Why Having Too Little Means So Much*, New York: Henry Holt & Co., 2013.
- [58] Natenzon, Paulo, “Random Choice and Learning,” *Journal of Political Economy* 127: 419-457 (2019).
- [59] Netzer, Nick, “Evolution of Time Preferences and Attitudes Toward Risk,” *American Economic Review* 99: 937-955 (2009).
- [60] Nieder, Andreas, and Earl K. Miller, “Coding of Cognitive Magnitude: Compressed Scaling of Numerical Information in the Primate Prefrontal Cortex,” *Neuron* 37: 149-157 (2003).
- [61] Payzan-LeNestour, Elise, and Michael Woodford, “‘Outlier Blindness’: Efficient Coding Generates an Inability to Represent Extreme Values,” working paper, August 2019.
- [62] Petzschner, Frederike H., Stefan Glasauer, and Klaas E. Stephan, “A Bayesian Perspective on Magnitude Estimation,” *Trends in Cognitive Sciences* 19: 285-293 (2015).
- [63] Polania, Rafael, Michael Woodford, and Christian C. Ruff, “Efficient Coding of Subjective Value,” *Nature Neuroscience* 22: 134-142 (2019).
- [64] Preston, Malcolm G., and Philip Baratta, “An Experimental Study of the Auction-Value of an Uncertain Outcome,” *American Journal of Psychology* 61: 183-193 (1948).
- [65] Rabin, Matthew, “Risk Aversion and Expected-Utility Theory: A Calibration Theorem,” *Econometrica* 68: 1281-1292 (2000).

- [66] Rayo, Luis, and Gary S. Becker, “Evolutionary Efficiency and Happiness,” *Journal of Political Economy* 115: 302-337 (2007).
- [67] Robson, Arthur J., “The Biological Basis of Economic Behavior,” *Journal of Economic Literature* 39: 11-33 (2001).
- [68] Rubinstein, Ariel, “Similarity and Decision Making Under Risk (Is There a Utility Theory Resolution to the Allais Paradox?)” *Journal of Economic Theory* 46: 145-153 (1988).
- [69] Rustichini, Aldo, Katherine E. Conen, Xinying Cai, and Camillo Padoa-Schioppa, “Optimal Coding and Neuronal Adaptation in Economic Decisions,” *Nature Communications* 8, article 1208 (2017).
- [70] Schwartz, Odelia, Anne Hsu, and Peter Dayan, “Space and Time in Visual Context,” *Nature Reviews Neuroscience* 8: 522-535 (2007).
- [71] Simoncelli, Eero P., “Vision and the Statistics of the Visual Environment,” *Current Opinion in Neurobiology* 13: 144-149 (2003).
- [72] Sims, Christopher A., “Implications of Rational Inattention,” *Journal of Monetary Economics* 50: 665-690 (2003).
- [73] Steiner, Jakub, and Colin Stewart, “Perceiving Prospects Properly,” *American Economic Review* 106: 1601-1631 (2016).
- [74] Stocker, Alan A., and Eero P. Simoncelli, “Noise Characteristics and Prior Expectations in Human Visual Speed Perception,” *Nature Neuroscience* 9: 578-585 (2006).
- [75] Summerfield, Christopher, and Kostantinos Tsetsos, “Do Humans Make Good Decisions?” *Trends in Cognitive Sciences* 19: 27-34 (2015).
- [76] Thurstone, Leon L., “A Law of Comparative Judgment,” *Psychological Review* 34: 273-286 (1927).
- [77] Tirole, Jean, “Cognitive Games and Cognitive Traps,” working paper, July 2015.

- [78] Trueblood, Jennifer S., Scott D. Brown, Andrew Heathcote, and Jerome R. Busemeyer, “Not Just for Consumers: Context Effects are Fundamental to Decision Making,” *Psychological Science* 24: 901908 (2013).
- [79] Tversky, Amos, and J. Edward Russo, “Substitutability and Similarity in Binary Choices,” *Journal of Mathematical Psychology* 6: 112 (1969).
- [80] Weber, Elke U., “Perception Matters: Psychophysics for Economists,” in I. Brocas and J.D. Carrillo, eds., *The Psychology of Economic Decisions, Volume 2: Reasons and Choices*, Oxford: Oxford University Press, 2004.
- [81] Wei, Xue-Xin, and Alan A. Stocker, “A Bayesian Observer Model Constrained by Efficient Coding Can Explain ‘Anti-Bayesian’ Percepts,” *Nature Neuroscience* 18: 1509-1517 (2015).
- [82] Wei, Xue-Xin, and Alan A. Stocker, “Lawful Relation Between Perceptual Bias and Discriminability,” *Proceedings of the National Academy of Sciences USA* 114: 10244-10249 (2017).
- [83] Whalen, J., Charles R. Gallistel, and Rochel Gelman, “Non-Verbal Counting in Humans: The Psychophysics of Number Representation,” *Psychological Science* 10: 130-137 (1999).
- [84] Woodford, Michael, “Imperfect Common Knowledge and the Effects of Monetary Policy,” in P. Aghion *et al.*, eds., *Knowledge, Information and Expectations in Modern Macroeconomics*, Princeton: Princeton University Press, 2003.
- [85] Woodford, Michael, “Prospect Theory as Efficient Perceptual Distortion,” *American Economic Review* 102(3): 1-8 (2012).