Research Article

A tale of two environments: divisive normalization and the (in)flexibility of choice

Short title: Divisive Normalization and Flexibility of Choice

Vered Kurtz-David[1,2]*, Shreya Sinha[1], Vinayak Alladi[3,4], Stefan Bucher[5,6], Adam Brandenburger[7,8,9], Kenway Louie[1,10], Paul Glimcher[1,10] ¶ and Agnieszka Tymula[3]¶.

[1] Neuroscience Institute, NYU Langone, New York, NY, USA.

[2] The Bogen Family Department of Economics, Hebrew University of Jerusalem, Jerusalem, Israel.

[3] School of Economics, University of Sydney, Sydney, NSW, Australia.

[4] The World Bank Group, Washington, DC, USA.

[5] Faculty of Economics, University of Cambridge, Cambridge, United Kingdom.

[6] Max Planck Institute for Biological Cybernetics, Tübingen, Germany.

[7] Stern School of Business, New York University, New York, NY, USA.

[8] Tandon School of Engineering, New York University, New York, NY, USA.

[9] NYU Shanghai, Shanghai, China.

[10] Center for Neural Science, New York University, New York, NY, USA.

*Corresponding author

Email: vered.kurtz@mail.huji.ac.il

¶ equally contributing senior authors.

Keywords: decision-making, efficient coding, divisive normalization, risky choice

Abstract

The Divisive Normalization (DN) function has been described as a "canonical neural computation" in the brain that achieves efficient representations of sensory and choice stimuli. Recent work shows that it efficiently encodes a specific class of Pareto-distributed stimuli. Does the brain shift to different encoding functions or is there evidence for DN encoding in other types of environments? In this paper, using a within-subject choice experiment, we show evidence of the latter. Subjects made decisions in two distinct choice environments with choice sets either drawn from a Pareto distribution or from a uniform distribution. Our results indicate that subjects' choices are better described by a divisive coding strategy in both environments. Moreover, subjects appeared to calibrate a DN function to match, as closely as possible, the actual statistical properties of each environment. These results suggest that divisive representations of encoded stimuli may be inherent to the nervous system.

Introduction

We make some decisions more often than others – in dozens of instances during our life, we choose between having two regular dishes for dinner, but rarely have to indicate which of two acclaimed restaurants we prefer. An often overlooked fact is that these encounter frequencies play a critical role in defining efficient encoding strategies – given constraints on neural coding, more accurate encoding must generally be allocated to more frequently encountered stimuli [1,2]. Indeed, experimental studies confirm this theoretical insight, showing a dependency of preference orderings, choice patterns [3–7], and choice efficiency [3,8] on the frequency with which subjects encounter different rewards.

This work has led to the conclusion that during the decision process, the brain adheres to principles of *efficient coding*, allocating resources to optimize decision outcomes [3,8–13]. A canonical example of a well-studied efficient code [13,14] is Divisive Normalization (henceforth DN) [15], which has been related to neuronal firing rates across many sensory modalities [16–19] and across various cognitive domains as well [20]. The DN function enables a system with limited information capacity to employ a flexible encoding of naturally occurring stimuli that is sensitive to encounter frequency [17,21,22]. Ample evidence has supported the notion that DN is also highly predictive of reward value encoding in the human and animal choice mechanism [6,7,23–25], although alternative value encoding mechanisms, some of which include division, have also been suggested [26,27].

At least one form of DN has been analytically shown to be an efficient code for stimuli with a probability of occurrence that is described by the asymmetrical heavy-tailed Pareto Type III distribution (see eq. (iv-vi) in Materials and Methods) [28]. This prompts the empirical question of whether the brain employs non-DN encoding functions when the statistical properties of the input stimuli (in our case, *choice environments*) are not Pareto-distributed. Would we expect to find evidence of divisive normalization across dimensions [28], or divisive encoding mechanisms in general [13], only in Pareto-distributed environments? The latter might imply that previous documentation of DN encoding mechanisms may say more about the stimulus distributions used in experiments than about constraints on encoding mechanisms. An alternative hypothesis, however, is that our brains are constrained to employ DN-like encoding mechanisms [13]. Such a constraint might reflect an adaptation of the nervous system to Pareto-distributed real-world natural stimuli, such as the sensory [14,18,29] and even ecological [30,31] environments we typically encounter. Our primary aim here is therefore to assess whether the encoding function itself is sensitive to the structure of the environment – specifically, to differences in the shape of the distribution of valuations.

In this study, our subjects complete a two-stage task design. The first stage recovers subjects' mapping of dollar amounts from objective to subjective values. These mappings are then used in the second stage of the study, in which subjects face a binary-choice task where lotteries are drawn from two individually-tailored environments characterized by different distributions of subjective lottery valuations. In one environment, lottery valuations are Pareto-distributed, while in the other, lottery valuations are uniformly distributed (Fig. 1A). Our novel task design controls for individual heterogeneity in subjects' risk preferences, thus ensuring that the second stage solely tests for contextual effects induced by the two environments.

We test hypotheses about our subjects' value encoding functions by fitting the patterns of errors in their choices with two random expected utility models (henceforth, RUM) [32,33]. The first one is a form of DN function designed for representation in risky choice [34]. As a non-DN benchmark, we use RUM with a power utility function that nests within its parameterization a linear, a concave, and convex encoder (henceforth, power utility; Fig 1B). Power utility is a common model in economic research for describing risky choice and has been applied across many subfields, including experimental settings [35–38], psychophysics [39], study of life-cycle consumption [40], and health [41]. (See [42] for a concise theoretical discussion of why power utility is so widely adopted and often fits choice data better than other functional forms.) Power utility is a natural comparator against which to evaluate DN.

We use a form of DN to examine if subjects are better described as obligate-DN choosers who use DN in both environments, or alternatively, if subjects' choices are better described with our DN function in one environment and with a power utility function in the other (Fig 1C).

The form of DN we test incorporates information about the environment by tracking expected valuations, allowing for context dependency in subjects' value encoding. Because the median valuation in the uniform environment is higher than in the Pareto environment, we further hypothesized that subjects would anticipate higher valuations in the uniform case, evidenced by a calibration of the DN function.

We find that in a Pareto-distributed environment, subjects employ an encoding of values that is well modeled by DN. However, we also find that the DN model better captures subjects' choices in the uniformly distributed environment than does a standard power utility. This suggests that subjects' choices are more accurately described by divisive encoders, like those found in DN models, than by standard power utility functions. We find further evidence for context dependency in subjects' choices, as, within the constraints of DN encoding, they adapt their reward expectations according to changes in the specific statistical properties of the choice environment.

Taken together, our results suggest that divisive mechanisms may be an inherent component of the encoding mechanism used during the choice process. Future work could generalize these findings to other types of statistical environments. Finally, the current study focuses on decision-making processes, but, given the dominance of DN representations across cortical systems, our findings may be of general interest to the study of encoding mechanisms in sensory and other cognitive domains.

[Insert Fig 1 here]

Fig 1. Research Question. (A) Choice environments are determined by the distribution of valuations. We compare a long-tailed bivariate Pareto Type III environment with a uniformly distributed environment for which DN is not an efficient code. Figures show 2D histograms of simulated choice trials with valuations in the range $u_k \in [0, u_i^{max} = 60]$ for every lottery $k \in \{1,2\}$. Each reward's value was drawn from 40 bins. Insets show their corresponding marginal distributions. We simulate 100k valuations per environment. See Materials and Methods for further details. (B) Value encoding choice functions. We test two different RUM models: classic power utility (left) and DN (right). The figure shows the probability of choosing a lottery with valuation u_1 over a lottery with valuation u_2 for various parameter values in each model. Insets show the subjective representation of u_1 in power utility (R), and in DN (S). For every combination of u_1 and u_2, we simulate 1k binary choice sets. We allow stochasticity in choice by incorporating additive noise, drawn from $\eta \sim N(0, 0.05 * R^{max}(u1)$, such that $R^{max}(u_1)$ denotes the maximal subjective value of u_1 in the power utility model (and $\zeta \sim N(0, 0.05 * S^{max}(u_1)$ in the DN model, respectively). We cast 10K noisy draws per simulated trial and reported average choice probabilities across simulated sets. (C) Contour plots indicate the mass of occurrences of (u_1, u_2) choice trial combinations in each environment. Contours were laid over a representative DN model with $\alpha = 4$, $M = 30$ (middle right panel in (B)).

Results

Two-stage task design

Seventy-six subjects completed a two-stage choice task. In STAGE I (Fig 2A left panel), subjects reported their valuations (willingness to pay) for 33 50/50 lotteries that pay either y_1 or y_2 dollars with a probability of 0.5 each (see S1 Table for a complete lottery list). These valuations were used to estimate, for every subject *i*, the curvature parameter of the expected power utility specification: $E[u_i(y)] = 0.5y_1^{\rho_i} + 0.5y_2^{\rho_i}$, using a standard non-linear least squares (NLS) estimation. We note that in STAGE I, our goal was not to test normalization, but to flexibly capture heterogeneity in subjective valuations of lotteries (risk preferences). We therefore employed the standard expected power utility function, which is the most widely used and well-understood functional form for eliciting risk preferences [42].

The subjective value function curvature (ρ_i) varied substantially from subject to subject (Fig 2E). Using individual ρ_i estimates, we generated subject-specific distributions of rewards in terms of their subjective – rather than dollar – values for the STAGE II task (Fig 2B). This first step was critical. It allowed us to perform all our analyses in the domain of subjective value, removing simple utility curvature from our primary analyses and allowing us to create individualized choice sets with specific distributional properties that were essential for our design. Without this transformation, small subject-specific differences in utility curvature (risk attitudes) would have made the construction of probative choice sets required for the experiment impossible.

In STAGE II, on two separate days, subjects made binary choices between 50/50 lotteries (Fig 2A, right panel), with 320 decisions on each day. We created two choice environments: on one day, subjects were choosing between lotteries with subjective values drawn from a Pareto Type III distribution (henceforth, Pareto), and on the other day between lotteries with subjective values drawn from a uniform distribution. Subjects encountered each distributional environment on a different day (counter-balanced across subjects) to avoid contextual spillovers.

Using these risky-choice lotteries, rather than choices over consumer goods, enabled us to generate *continuous* distributions of valuations for STAGE II and to fully control their *distributional shape*. Our decision to generate the distributions of STAGE II lotteries in subjective value space, rather than in dollar space, ensured that any observed environmental effects were not confounded by the heterogeneity in subjects' subjective valuations of lotteries, that is, their risk attitudes (Fig 2E and S2 Table). Consider three subjects exposed to the same set of 50/50 lotteries with uniformly distributed dollar payoffs. Subject 1 is risk-seeking: subjective value grows slowly (convex subjective value function); Subject 2 is risk-averse: subjective value grows more than proportionally (concave); Subject 3 is risk-neutral: subjective value is linear in objective value. Now, imagine we created a choice environment with uniformly distributed 50/50 lottery payoffs (in dollar amount). The same uniform distribution would induce a left-skewed distribution of subjective values in Subject 1, a right-skewed distribution in Subject 2, and only for Subject 3 does the subjective value distribution remain uniform. Since what we wish to study is the subjective value distribution rather than the expected value distribution, we must first factor out this heterogeneity. Our two-stage procedure thus ensured that the *shapes* of the individually tailored distributions in STAGE II are controlled and comparable across subjects. This design is therefore crucial for valid between-subject comparisons of environment-induced effects and for cleanly addressing our central questions: how statistical environments shape the value-encoding function, and how well subjects adapt to these environments.

Across subjects, we fixed the first moment (mean) of valuations and the range of monetary payoffs in both environments. Of course, this also fixed the second moment (standard deviation) of the uniform distribution across subjects. The second moment of the Pareto distribution (as measured in dollars) varied by subjects' subjective valuations of money as assessed in STAGE I (risk attitudes) (Fig 2C, S1 Fig). Accordingly, this heterogeneity also varied the distributions of the high and low monetary payoffs in each lottery (S1 Fig). As a result, the median expected monetary payoff in each environment was fully determined by subjects' risk attitudes, so that the difference in expected payoffs between environments was smallest for risk-averse subjects ($\rho_i < 1$, S4 Fig). To ensure that we fully captured each distributional environment, we matched the mean and standard deviation of the choice sets with those of larger sets of 100k draws (Fig 2D). See Materials and Methods for further details on our sampling design.

Overall, subjects appeared to pay careful attention during the study – only six subjects in the uniform environment, and nineteen subjects in the Pareto environment failed to choose the higher subjective value lottery in more than 20% of trials (S2 Fig). On average, subjects violated first-order stochastic dominance in 0.97% of trials in the uniform treatment and in 1.08% of trials in the Pareto treatment, respectively (S2 Fig). Note that a higher incidence of mistakes in the Pareto environment is expected. The correlational structure across lotteries made the value difference between lotteries (on average) smaller, and thus choices were harder in this case [43]. Finally, even though the experiment was quite demanding (320 trials in each of the two sessions), subjects' performance was not affected by fatigue. The propensity to choose the lottery with the higher subjective value did not vary between the first and second halves of each experimental session (Pareto sessions: p=0.2791, uniform sessions: p=0.5109, paired t-test (df=75), S2 Fig).

[Insert Fig 2 here]

Fig 2. Experimental Design. (A) Timeline. In STAGE I, subjects reported their valuations for 33 lotteries. Valuations were used to recover the curvature of the subjective value function for each subject using NLS estimation. Based on those estimates, we generated subject-specific bi-dimensional uniform and Pareto Type III distributions of valuations for STAGE II of the study. In STAGE II, subjects completed two sets of 320 binary choices between 50/50 lotteries (640 choices in total). (B) Bi-dimensional Pareto and Uniform distributions. In the uniform distribution, we created 40 bins of subjective values between 0 and the maximal payoff in the study (\$60, $u_i^{max} = 60^{\rho_i}$) with eight lotteries in each bin. We then picked pairs of lotteries from this set to create binary choice sets. In the Pareto distribution, we used a Gamma-weighted scale mixture of exponential random variables to capture the covariance structure of the bi-variate Pareto distribution. (C)

Choice sets in STAGE II controlled for differences in individual subjective value function (risk attitudes), modulating the second moment (std) of the Pareto distribution (see eq. (vi) in Materials and Methods). The histograms show the bi-dimensional Pareto distributions and their marginals (with 100k draws per distribution) from three representative subjects: a risk averse subject (*left*), a risk neutral subject (*middle*), and a risk seeking subject (*right*). (D) Experimental sets with 320 trials were prone to under-sampling (see top, unmatched distribution). We matched experimental sets to the distributional shape of a larger set with 100k draws (see bottom, matched distributions). The figure shows an example corresponding to the middle panel in (C). (E) Recovered estimates of subjective value curvature (risk attitudes) from STAGE I. See Materials and Methods and S1 Fig for further details. See S2 Table for a list of the estimated subjective value function curvatures (risk parameter ρ).

Distributional properties of the choice environments influence subjects' choice behavior

Our overarching goal was to study how the distributional properties of the choice environment influenced the encoding of value, and whether subjects could flexibly switch between different types of encoding mechanisms, as evidenced by errors in their choice patterns, in different environments. We created the experimental choice environments with Pareto Type III and uniform distributions of valuations. In this section, we tackle the first part of our research question in a model-free manner, determining whether the distributional structure influenced the errors produced by our subjects in a meaningful manner.

It is useful to introduce our hypotheses using an illustration. In Fig 3A-B, we indicate the probability of choosing the higher valued lottery, given the coupling of the (u_1, u_2) valuations in a choice set. Choices along the diagonal represent trials in which the two lotteries had the same or very similar valuations, whereas trials away from the diagonal correspond to choice sets in which the valuations of the two lotteries were substantially different. A central feature of DN is the calibration of the function used to represent subjective value (the decisional variable) to the input stimuli. That is, encoding/representational resources are allocated to the range of stimuli most likely to be observed (*tuning*) [15,25]. Thus, compared with nondivisive encoders, if DN governs the choice mechanism in a Pareto environment, we would expect subjects in this case to make more mistakes in choice sets containing elements further from the high-density center of the main diagonal, since these choices are less frequent. Conversely, we would also expect subjects in the Pareto environment to make fewer mistakes in choice sets containing elements nearer the main diagonal, since these choices occur more frequently. We find both patterns in our data.

To statistically test whether the frequency of mistakes grew faster as choice sets moved away from the main diagonal in the Pareto environment, as compared with the uniform environment, we ran a probit regression with an indicator dependent variable equal to one for trials on which a subject selected the option with higher SV, and equal to zero otherwise. We controlled for the difference in difficulty across the trials by including the absolute value difference between the lottery valuations ($|u_1 - u_2|$) and for the general impact of the distribution by including a dummy for the Pareto distribution. The different rate of mistakes as a function of the distance from the diagonal in each environment is captured by a significant coefficient on the interaction of Pareto dummy and ($|u_1 - u_2|$) (Column (1) in Table 1). We found that choice accuracies increased with an increase in the subjective value distance between the two options, and that moving from the uniform distribution to the Pareto distribution reduced accuracy (see also discussion in the previous section). Importantly, in line with our hypothesis, we found a negative and significant interaction term, indicating that in the Pareto vs. the uniform case, subjects were more likely to make errors when encountering choice sets further away from the diagonal -- those sets being experienced less often in the Pareto environment. We conclude that encounter frequency, as defined by the Pareto distributional structure, did influence choice accuracy.

To examine whether subjects calibrated their encoding function to the most frequently presented choice sets, we tested if they made fewer mistakes around the high-density center of the main diagonal in the Pareto environment. We ran a complementary probit regression focusing on twenty-two valuation bins from the center of the distributions presented in Fig 3A-B (out of an equally-spaced 40-bin space), which corresponded to lotteries with $9-$42 payoffs (Column (2) in Table 1). The center (medians) of the distributions depended on subjects' subjective valuations of dollar amounts (ρ parameter). In the Pareto case, the smallest median was $11.45 and the highest was $33.58. Likewise, in the uniform case, the smallest median was $22.85 and the highest was $41.53. Thus, we set a range of $9-42 to include the center of distributions for all the subjects in our sample.

Since this regression focuses on the center of the distributions, the SV difference between the two lotteries is relatively small. Consequently, we excluded this variable from the model to avoid potential multicollinearity with the main variables of interest. In addition to the Pareto dummy regressor included in the baseline specification, we introduced a dummy variable indicating whether a lottery was located near the diagonal of the valuation space, as well as an interaction of this dummy with the Pareto dummy. Lotteries were defined as near the diagonal if the ratio between the two valuations satisfied $0.9 < u2/u1 < 1.1$.

Not surprisingly, choice accuracy was lower in choice sets around the diagonal, since these represented the most difficult choices in the experiment and exhibited the smallest SV difference. Crucially, we found a positive interaction term between the diagonal and Pareto dummies, suggesting that in the Pareto vs. uniform environment, subjects had higher accuracy in those particularly difficult trials within the highly sampled region.

A graphical illustration of this finding is depicted in Fig 3C, which plots the subject-level change in the accuracy around the center of the distributions compared with their overall accuracy in each environment. As expected, in both environments, we trace a decline in subjects' accuracy around the center of the distribution, since these are the most difficult trials in the experiment, though this decline is more moderate in the Pareto environment than in the uniform environment (one-sided paired t-test, p<0.0001).

In the supplementary materials, we present two robustness analyses supporting these results. S7 Table replicates the findings from Table 1, using a Bernoulli specification in which the mean is constrained to the interval [0.5, 1) rather than the standard probit model. This specification addresses potential discontinuities in the model below chance level (50%) as the SV difference approaches zero. In addition, S3 Table demonstrates that the results reported in Column (2) of Table 1 remain robust when applying alternative definitions of the "center of the distribution" and "around the diagonal" regions.

Together, these results suggest that in the Pareto environment, subjects adjusted their value encoding to increase choice accuracy rates at the center of the joint distribution, at the expense of the decreased choice accuracy at the margins. This is suggestive evidence for some forms of divisive value encoding, where choice discriminability is the highest near the median of the distribution.

Table 1. Results from the model-free analysis. Probit regressions with the dependent variable equal to 1 when the subject chose the lottery with the higher SV, and zero otherwise. Column (1) model was run on the full sample. The independent variables are the absolute SV difference between the two lotteries, a dummy indicating the Pareto environment, and their interaction. Column (2) model was run on data including choice sets in the center of the distributions. The model includes a dummy for the Pareto distribution, an additional dummy equal to 1 if the lottery was taken from around the diagonal (and zero otherwise, see text for definitions), and their interaction. Standard errors clustered on subject in parentheses, * p<0.05, ** p<0.01, *** p<0.001.

	(1) Full sample	(2) Center of the distributions
SV difference	0.0002*	
	(0.0001)	
Pareto	-0.3311***	-0.1587***
	(0.0405)	(0.0381)
Pareto*SV difference	-0.0001*	
	(0.0001)	
Near diagonal		-0.7376***
		(0.0683)
Pareto*Near diagonal		0.1839**
		(0.0602)
Constant	1.3260***	1.1175***
	(0.0704)	(0.0703)
N	48640	22442

pseudo R-sq	0.015	0.034

[Insert Fig 3 here]

Fig 3. Model-free evidence for DN encoding. (A-B) Probability of choosing the higher-valued lottery given (u_1, u_2) valuations in a choice set. Data is aggregated over subjects. Within subjects, valuations are divided into 40 equally spaced bins. (A) The Pareto environment. (B) The uniform environment. (C) The change in the propensity to choose the higher-valued lottery around the center of the distribution, defined as bins #5-25 along the SV space diagonal, with a band of three bins below and above the main diagonal (illustrated by the dashed rectangles at the bottom). N=76.

Evidence for DN-like value encoding across choice environments

The findings in the previous section provided initial evidence that subjects adapted to the distribution of valuations and that subjects used some form of divisive encoding in the Pareto environment. The DN model enables subjects to focus their resources on the center (median) of the distributions – those valuations that they are more likely to encounter. By contrast, in the uniform environment, subjects are less likely to encounter these valuations, making DN encoding less beneficial. Nevertheless, such encoding would still yield different choice patterns from those under standard power utility. Hence, our next goal was to evaluate whether subjects used the same or different encoding mechanisms in each of the two environments.

To answer this question, we tested which of two expected utility models – a form of DN, or power utility – better captures subjects' choices. The DN model is regarded as a canonical encoding mechanism in the brain [15–17,19], including in the choice domain [7,24,25,44], and has been considered an efficient encoder [1,9,13,45]. One variant of the DN model (cross-normalized) has been proven to efficiently encode Pareto distributed environments [28].

We estimate a form of DN model that has been used to study risky choice behavior [34]. (See [46] for modeling subjects' behavior in the same dataset with alternative DN specifications.) In this DN model, subject i's STAGE II subjective value function of a lottery $k \in \{1,2\}$ with payoffs $x_{1,k}$ or $x_{2,k}$ is given by:

$$\text{(i)} \quad S_{i,k} = 0.5\frac{(u_i(x_{1,k}))^{\alpha_i}}{(u_i(x_{1,k}))^{\alpha_i}+M_i^{\alpha_i}} + 0.5\frac{(u_i(x_{2,k}))^{\alpha_i}}{(u_i(x_{2,k}))^{\alpha_i}+M_i^{\alpha_i}} + \varepsilon_{i,k}$$

where α_i is a curvature parameter, $u_i(\cdot)$ is subject i's STAGE I's power utility function (i.e., $u_i(x_{1,k}) = x_{1,k}^{\rho_i}$), M is a reward expectation parameter, and $\varepsilon_{i,k}$ is an additive decision noise drawn in each trial from a zero-mean normal distribution, such that $\varepsilon_{i,k} \sim N(0, \theta_{DN})$. The encodings of the two marginals are then combined using the 50/50 risk probabilities to arrive at the overall formula. In line with the previous set of results, we expected some form of DN encoding in the Pareto environment.

The second model we examined was a power utility model, a common model in describing risky choice behavior [42], applied here to subjective rather than monetary values:

$$\text{(ii)} \quad R_{i,k} = 0.5\big(u_i(x_{1,k})\big)^{r_i} + 0.5\big(u_i(x_{2,k})\big)^{r_i} + \eta_{i,k}$$

The model has one free parameter (r_i), which captures the function's curvature. When $r_i = 1$, the function is linear. As in our DN model, we included an additive decision noise drawn in each trial from a zero-mean normal distribution, such that $\eta_i \sim N(0,\theta_p)$. A consequence of our design – positive-payoff 50/50 lotteries – is that Prospect Theory [47] and this model coincide.

For every subject, we estimated both models using maximum-likelihood estimation (see Materials and Methods). The subject-specific recovered parameters are reported in S4 Table, and the sample medians are in Table 2. To determine, at the population level, which model better captured subjects' choice patterns in each environment, we compared each subject's Bayesian Information Criterion (BIC) scores across the two models in each environment. Options in the uniform environment had, on average, larger value difference, and responses in this environment were more accurate (S2 Fig) and less noisy (S3 Fig). Therefore, we only compared BIC scores of the two models within the same environment, and did not compare the models across the two environments.

In line with our hypothesis, we found that in the Pareto environment, subjects' BIC scores were on average significantly lower, indicating a better model fit, for the DN model than for the power utility model (Fig 4A, one-sided Wilcoxon sign-rank test, Z=4.4603, p<0.0001). This was true for 48 subjects (out of 76). In the uniform environment, we, again, found that the BIC scores were on average significantly lower for the DN model (Fig 3D, one-sided Wilcoxon sign-rank test, Z=2.9692, p=0.0015) and this held for 42 (out of 76) subjects. Among the subjects who had a lower BIC score for the DN model in one of the environments (41 in uniform, 48 in Pareto), 31 subjects, i.e., roughly three-quarters of each group of subjects, exhibited consistent preference for the DN model across the two environments. Fewer subjects (17) had consistently lower BIC scores for the power utility model in both environments.

Moreover, for only four subjects, the curvature parameter in the power utility model was estimated as linear or as almost linear ($r_i = 1 \pm 0.05$). The mean and median *r* estimates were 0.608 and 0.366, respectively. Importantly, the asymmetrical distributions of the differences in BIC scores (see insets in Fig 4A-B) indicate that while both models do (almost) equally well for most subjects, there is a group of subjects for whom the DN model predicts their choices much better (ΔBIC>20 for 29 subjects in Pareto and 18 subjects in uniform).

For a further comparison, we also recovered the pooled (aggregate) model parameters. S5 Table presents the recovered pooled estimates from this analysis. Note that this analysis could only be done in monetary space, so as to allow comparability of lotteries across subjects, and to recover meaningful estimates of the *M* parameter in the DN model. Here, too, we find that the DN model captured subjects' choices better, evidenced by the lower BIC scores under aggregation of choices from both treatments (leftmost column), as well as within each environment (second and third columns). These results should be interpreted cautiously, since the reward distributions were not fully controlled in the monetary space (S1 Fig).

Next, we directly examined whether the DN model fully contextualizes information about the environment. We re-estimated subjects' choices using the DN specification in Eq. (i), setting M equal to the true median of the subjective valuations encountered within each statistical environment (see Materials and Methods). At the aggregate level, the DN model with fixed M provided a better fit than the power utility model, as indicated by a lower BIC (S5 Table, column 1). When analyzed by environment, this advantage persisted in the uniform case but not in the Pareto case (S5 Table, columns 2–3). Subject-level analyses mirrored this pattern (S5 Fig): In the uniform environment, BIC scores were significantly lower for the DN model (one-sided paired *t*-test, $t(75) = 1.735$, $p = 0.0434$), whereas in the Pareto environment, the two models performed comparably ($t(75) = 0.975$, $p = 0.1663$). This analysis is, nevertheless, less suited to testing our main hypothesis because fixing *M* does not allow for partial adaptation. The model's relative success in the uniform environment underscores the value of DN-like contextual choice models compared with non-contextual models.

Another way to examine the effect of the distributional environment on subjects' value-encoding – and to validate our task design – is to assess the relationship between subjects' subjective valuation estimates from STAGE I (ρ, see Fig 2E) and the STAGE II parameters: α in the DN model, and r in the power utility model. The estimated parameters in STAGE II are multiplicative combinations of the STAGE I parameters. Therefore, if participants were not adjusting to the environment, we would simply find that in STAGE II, and independent of the environment, a power utility model with the exponent equal to one (i.e. $r_i = 1$) would fit the data best. Figs 4C-4F suggest that this is not the case, further indicating that the environments influenced subjects' value encoding – specifically through capturing the residual curvature attributable to changes in the statistical environment. We also note that the two sets of parameters were fit on different datasets: ρ_i was recovered from STAGE I data, whereas α_i and r_i were recovered from STAGE II data, thereby highlighting that the STAGE II parameters can be attributed to environmentally induced effects.

Given the nature of our design, a hyperbolic relationship between STAGE I ρ and the STAGE II parameters (i.e., the function $y = 1/x$) would imply that STAGE I curvature is undone in STAGE II in both models. In the power utility model, this would also imply linear encoding of monetary payoffs (because $(x^{\rho})^{1/\rho}=x$). In contrast, in DN, it would mean that all curvature in STAGE II is associated with the DN encoding.

Fig 4C-F plots STAGE II parameters against STAGE I ρ, highlighting two key findings: (a) the scatter of points away from the hyperbolic function ($y = 1/x$) indicates that there remains a residual curvature attributable to the statistical environments, and (b) in both the Pareto and uniform environments, the relationship between ρ and the DN parameter (Figs 4C and 4E) is considerably closer to hyperbolic than that between ρ and the power utility r parameter (Figs 4D and 4F). This suggests that, although the environments further modified the curvature of the encoding function, the DN function better captured STAGE II's subjective value functions. To quantify this effect, we compared the root-mean-squared errors (RMSE) between the hyperbolic function and the parameters in both models and confirmed that across the two environments, the α parameter of the DN model was more likely to exhibit this hyperbolic relationship (α: Pareto: RMSE=0.6491, uniform: RMSE=0.6585; r: Pareto: RMSE=0.8109, uniform: RMSE=0.8588).

Taken together, all these results strengthen the notion that subjects used DN-like encoding of value in both environments.

Table 2. Median estimates.

	Power utility model			DN model			
Parameter	r	θ_p	BIC	α	M	θ_{DN}	BIC
Uniform	0.379	0.054	121.120	1.299	23.216	0.0233	104.670
Pareto	0.528	0.103	184.761	1.358	18.875	0.026	154.482

[Insert Fig 4 here]

Fig 4. Model-fitting. (A-B) Each dot is one subject's DN model BIC score (y-axis) plotted against the same subject's power utility BIC score (x-axis). A dashed 45-degree line indicates when both models are equally successful. Inset shows the difference in BIC scores ($BIC_{power} - BIC_{DN}$). (A) The Pareto environment. (B) The uniform environment. (C-D) Relationship between the STAGE I curvature of the subjective value function (ρ) and STAGE II subjective value functions in the Pareto environment. The dashed curve indicates a hyperbolic function $y = 1/x$. (C) DN model (α parameter). (D) Power utility model (r parameter). (E-F). Same as (C-D), but for the uniform environment. Dots indicate individual subjects, + indicate the sample averages. N=76.

Context-dependency: adaptation of the encoding function to the choice environment

Our next aim was to examine whether subjects calibrated their encoding functions according to the properties of the two different environments. We found that, for all subjects in our sample, the medians of the subjective valuations in the uniform environment were higher than in the Pareto environment (sample medians: 18.721 vs. 14.723 util units, respectively, Δ= 3.998, one-sided Wilcoxon sign-rank test between subject-specific medians, Z=7.572, p<0.0001). The reward expectation M in the DN model tracks the median of the reward distribution, and hence, we hypothesized it would be higher in the uniform environment. Consistent with this hypothesis, the sample median of the recovered M parameters in the uniform case was higher by 4.99 (in utility units) than in the Pareto case (Table 2, also corroborated by a one-sided Wilcoxon sign-rank test, Z=2.8907, p=0.0019). This difference between the recovered M parameters was very close to the actual difference between the distributions' medians, indicating that subjects, at least at the sample-level, quite precisely calibrated their encoding to the difference in reward expectation. On the subject level, we found that for 44 out of 76 subjects estimated M was higher in the uniform environment (Fig 5A).

We also examined how the change in recovered M across environments related to the true difference in median payoffs. To make these parameters comparable, we expressed M in dollar units (i.e., $\widetilde{\mathrm{M}}_{\mathrm{i}} = M_i^{\frac{1}{\rho_i}}$). Perhaps counterintuitively, we found a negative correlation between the two measures (r = –0.279, p = 0.015; S4 Fig), indicating that larger differences in true payoffs across environments were estimated as smaller differences in M. However, this negative correlation arises mechanically from the relationship between median payoffs and subjects' risk preferences (the ρ parameter). This result is further supported by the positive correlation between ρ and $\Delta\widetilde{\mathrm{M}}_{\mathrm{i}}$ (r = 0.311, p = 0.006; S4 Fig).

Our pooled estimation supports this conclusion with higher estimates of M in the uniform environment ($M(uniform) = 66.653$ and $M(Pareto) = 55.561$, second and third columns in S4 Table, p<0.001. As a robustness check, we estimated the DN model using the full dataset with the data from both environments and included an additive dummy variable for the Pareto environment in the estimation of the

M parameter ($M = constant + M_{pareto}\, x\, Pareto$). The output of this model split M into a constant, corresponding to the estimate of M for the uniform environment, and an additional coefficient M_{pareto} that captured the difference in M in the Pareto relative to the uniform environment. We found M_{pareto} to be negative and significant (p<0.001), indicating M was lower in the Pareto environment.

As an additional test of subjects' adaptation to the two statistical environments, we examined the stability of the M parameter. Firstly, we tested whether the statistical structure in the uniform environment may have been less informative for choosers employing the DN model. As a result, the recovered M values would be noisier than in the Pareto environment. To allow interpretability of M across subjects, we focused on the pooled estimates in dollar space. Indeed, we found that the standard error of M estimate was substantially larger in the uniform vs. the Pareto environment (15.31 vs. 9.01, S5 Table), suggesting that subjects had greater difficulty calibrating the M parameter to the uniform environment.

Secondly, we tested whether a longer exposure to a given statistical environment led to more precise estimates of M. For each environment, we estimated M for early trials (#1-160) and late trials (#161-320), separately in each session (S6 Table). We found that in the Pareto environment, the recovered M values did not change significantly (54.6 in early trials vs. 56.5 in late trials). By contrast, in the uniform environment, we observed substantial within-session variability, and the recovered M decreased from 75.2 in early trials to 59.5 in later ones. Notably, the latter estimate is closer to that of the Pareto environment, though still somewhat higher, reflecting the relative differences in the environments' actual medians. These two results provide additional evidence that subjects calibrated to the statistical structure of both environments.

In contrast to the M parameter, we had no prior hypotheses regarding the model's curvature parameter α. Nevertheless, comparing subject-specific estimates, we found that, on average, the α parameter was higher by 0.1593 in the Pareto environment (one-sided Wilcoxon sign-rank test, Z=1.9987, p=0.0228, Fig 5B). This result may indicate that higher α values in the Pareto environment allowed better discriminability between the more frequently encountered lottery options, as also indicated by our model-free analysis (Table 1). However, this result was not fully replicated in the pooled estimates. When estimating each environment separately, we found that recovered parameters were almost identical ($\alpha_{uniform} = 0.93$, $\alpha_{Pareto} = 0.92$, S4 Table, second and third columns). A full model with random effect for the Pareto environment (similarly to the one run on M) revealed that there was a tuning of the function curvature when switching between environments (S4 Table, rightmost column, p<0.001).

The power utility model is not designed to capture the dependence of the subjective value function on the distribution of valuations and hence, we did not anticipate an adaptation of the function's curvature. Indeed, when comparing estimates of r across the two environments, we obtain inconclusive results. While the pooled estimates indicated higher r values in the Pareto environment (S4 Table), the subject-level estimates pointed in the opposite direction (Fig 5C, one-sided Wilcoxon sign-rank test between subject-level estimates of r, Z=0.051, p=0.480).

To conclude, we found that subjects adapted the parameters of the DN encoding function to the two environments in line with our hypothesis, showing context dependency in choice.

[Insert Fig 5 here]

Fig 5. Cross-environment adaptation. (A-B) Adaptation of the encoding function in the DN model. (A) Best-fitting M parameter in the uniform (x-axis) vs. the Pareto (y-axis) environments. Estimates of M's are in utility space. Left inset: outliers. Right (diagonal) inset: Difference in the estimates of M across choice environments ($M_{uniform} - M_{Pareto}$). Insets do not show three additional (risk-seeking) subjects whose M's are >400 (in util units). Dots indicate individual subjects, + indicate sample average without the inset outliers, N=76. (B) same as (A) for the DN's α parameter. (C) Adaptation of the encoding function in the power utility model. Same as (B), but for the r parameter from the power utility model. (B-C) Dots indicate individual subjects, + indicate sample average, N=76.

Discussion

In this study, we tested how the distributional properties of choice environments affect value encoding. In particular, we were interested in whether the subjective value of rewards is encoded via a mechanism such as divisive normalization (DN) exclusively in the Pareto environments akin to those for which it is probably efficient [28], or whether a DN representation is also employed in environments characterized by different reward distributions. To this end, we designed an experiment in which subjects were asked to make choices in two distinct statistical environments. In one environment, rewards were drawn from a Pareto distribution of valuations, while in the other environment, valuations were uniformly distributed.

Our results indicate that subjects in our study were better described as using a DN mechanism than a power utility mechanism to encode the subjective value of rewards, regardless of which of our two distributions the rewards were drawn from. As expected, the key parameter of the model tracked the median of the distribution. A model-free analysis indicated that, as compared with the uniform environment, subjects in the Pareto environment made fewer mistakes when choice sets were drawn from the center of the distribution at the expense of the margins, in accordance with a principal property of the DN function. We then fitted our subjects' choices with two RUMs – a RUM with a DN-like utility function and the other a standard RUM with a power utility function. Our subject-level and pooled model-fitting results suggested that the DN model better captured subjects' choice patterns in both the Pareto and the uniform environments (Table 2, Fig 5C-D and S5 Table). In line with the actual statistical properties of the two environments, subjects had higher reward expectations in the uniform environment. Taken together, these findings suggest that subjects' choices were affected by the context of the choice environment, and that their choices were better described by DN-like divisive encoders than a more standard power utility model (Fig 6).

[Insert Fig 6 here]

Fig 6. Summary of main findings.

One reason to see DN encoding, even across environments, is that Pareto distributions are very common in the real world, and the human brain has evolved a mechanism that accords well with natural environments. Indeed, numerous sensory stimuli are characterized by Pareto-like statistical properties [1,14,18,29]. On a larger scale, Pareto distributions also describe various ecological quantities, such as temporal and spatial measures of biodiversity [48–51]. More relevant to value-based decisions is that certain economic and financial variables in modern societies [52,53], including consumption of several categories of consumer goods [54], have Pareto-like properties.

Another important finding is that, compared with the standard utility functions used in economics, DN provides the brain with a rather flexible tool for the representation of choice options [34,44]. Given the specific parameterization we employed for DN, our model embeds the standard concave utility function, but is also suitable for capturing preferences that follow *S-shaped* functions, similar to the one suggested by Prospect Theory [55] with expectations-based reference dependence [56]: The M parameter in the DN model tracks the median of rewards (*expectations*), which allows for scale-invariant adjustments to different environments, while ensuring a fine discrimination between stimuli that are in the center of the distribution [2,9,13,57]. These adjustments – also evident in our data – give rise to spatial and temporal context effects in choice processes [25,44,58–62], and are also the core reason for some notable perceptual illusions [63,64].

Our findings also imply that some choice patterns should not be regarded as built-in decision biases, errors, or mistakes. Rather, they reflect adjustments of the brain, as a constrained system, to its environment, thus reflecting a rational value-encoding mechanism [2,13]. Such an observation can explain the under-sampling of rare events when subjects adjust to new choice environments [65,66] since the main focus of the system is on the mass of occurrences.

Our primary aim in the current study was to assess whether the encoding function was sensitive to the distributional structure of the environment. Future work could vary both distributional shapes and their means (in a factorial design) for a targeted test of adaptation to study which environmental statistic underlies normalized encoding. Another interesting question that stems directly from our research is to what extent our results generalize beyond decision-making processes to other cognitive functions, such as sensory processing. Even though various natural sensory stimuli are described by Pareto-like properties [14,18,57], we also frequently encounter, and are required to process, non-natural non-Pareto stimuli [67,68]. Our findings, therefore, invite further investigation into the effects of DN encoding on the sensory processing of non-Pareto stimuli.

Materials and methods

Some of the data in this manuscript have been used in the conference paper in reference [46].

Experimental design

Valuation task (STAGE I)

Our goal was to establish whether the brain employs different value encoding models in environments with different reward distributions. To eliminate any additional prior heterogeneity in subjects' subjective valuations of money, we generated distributions of rewards in the subjective value (SV) space instead in dollar amounts (or expected values). To map the subject-specific SV space, we first recovered individual-specific subjective value functions over dollar amounts. To do this, in STAGE I, we used a valuation task, in which subjects reported their willingness to pay to participate in a lottery. See S1 Table for the list of 33 lotteries used in this task. On each trial, subjects were presented with a visualization of a 50-50 lottery on the computer screen and had to type in their willingness to pay to participate in it as a dollar amount (Fig 2A). For each lottery, the valuation could range between the current lottery's minimal and maximal payoff, in $0.10 increments. All subjects completed the same 33 trials in an order randomized at the subject level. At the end of the session, the realization of one randomly selected trial was implemented for payment, using a Becker–DeGroot–Marschak (BDM) [69] procedure which was designed to elicit truthful valuations.

Choice task (STAGE II). STAGE II was designed to test whether the distribution of rewards (lotteries with different subjective valuations) in a choice environment affects what value encoding model subjects use. Subjects were asked to choose the 50-50 lottery they preferred from two available options that varied from trial to trial. Lottery payoffs ranged between $0 and $60 in $0.10 increments. Overall, subjects made 640 binary choices that were divided into two blocks of 320 trials each and presented on subsequent days. Our experimental manipulation was that in each block, the valuations were drawn either from a Pareto Type III distribution or from a uniform distribution (Fig 2A-B). The order in which subjects experienced these environments was counter-balanced across subjects. One trial was randomly selected for payment at the end of each experimental session. Subjects also completed additional 640 trials with six-option choice sets with lottery valuations drawn either from a Pareto Type III or uniform distributions. Thus, in total, in each environment subjects encountered two 320 choice blocks. The six-option blocks were designed to examine another research question that is beyond the scope of the current study and will be reported in a separate paper. Blocks were presented in an order randomized across subjects but on a given day, all blocks were drawn from the same distribution. Payments for STAGE II included a realization of one choice from each of the two sessions, and could be drawn either from the two-options sets or from the six-options sets.

Subjective Value of Money. We used each subject's STAGE I single lottery valuations to estimate their subjective value function over money. We expressed each subject *i*'s subjective value of a 50-50 lottery that paid y_1 or y_2 each equally likely, using an expected power utility function as:

(iii) $E[u_i(y)] = 0.5y_1^{\rho_i} + 0.5y_2^{\rho_i}$

If the curvature parameter $\rho_i < 1$, then subject *i* is risk-averse. When $\rho_i = 1$, the subject is risk-neutral. If $\rho_i > 1$, the subject is risk-seeking. Therefore, the certainty equivalents (c) that participants stated were converted to subjective values using the same power utility function such that $c = E[u_i(y)]^{1/\rho}$. We ran an NLS regression to estimate the ρ parameter separately for each subject.

We used the subject's' estimated ρ_i, to pick different combinations of lottery dollar payoffs to create lotteries that had a specific SV to that individual. This enabled us to generate sets of lotteries whose implied SV distributions matched our target distributions (see below), regardless of individual differences in the curvature of the subjective value function.

Distributions of valuations

Uniform Distributions of SVs. For each subject *i,* we computed the upper bound of the distribution as the SV of the maximal possible monetary payoff in the study, which was \$60 (i.e., $u_i^{max} = 60^{\rho_i}$). We then divided the range [0,u_i^{max}] into 40 equally-spaced SV increments. For each of the increments, we created eight different lotteries, which would give the subject the subjective value in exactly this bracket (for a total of 320 lotteries). Since the joint distribution of a two-dimensional uniform distribution is independent, and hence determined by its marginals, we then picked pairs of lotteries from this set for generating binary choice sets.

Pareto Type III Distributions of SVs. In the Pareto treatment, subjective lottery valuations are drawn from a bivariate Pareto distribution with a joint pdf $f_{u_i}\left(u_{i,1}, u_{i,2}\right)$ given, for every subject *i* and $k \in \{1,2\}$ enumerating the choice option within the choice set (see Eq. 7 in reference [28], with μ_1=0 to match the lower bound of the uniform distribution and to avoid negative valuations), by

$$\text{(iv)}\quad f_{u_i}\left(u_{i,1}, u_{i,2}\right) = \beta^2 \frac{2\left(\prod_{k=1}^{2} \frac{1}{\sigma_{i,k}} \left(\frac{u_{i,k}}{\sigma_{i,k}}\right)^{\beta-1}\right)}{\left(1+\sum_{k=1}^{2}\left(\frac{u_{i,k}}{\sigma_{i,k}}\right)^{\beta}\right)^{3}}.$$

Its marginals are log-logistic (or Fisk) distributions, with pdf

$$\text{(v)}\quad f_{u_{i,k}}\left(u_{i,k}\right) = \beta \frac{\frac{1}{\sigma_{i,k}}\left(\frac{u_{i,k}}{\sigma_{i,k}}\right)^{\beta-1}}{\left(1+\left(\frac{u_{i,k}}{\sigma_{i,k}}\right)^{\beta}\right)^{2}}.$$

We matched, for each subject, the conditional mean to the expectation of the uniform distribution, which was \$30 (and $\bar{u}_i = 30^{\rho_i}$ in SV-space); the conditional mean is given, for β>1/2, by

$$\text{(vi)}\quad E\left(u_{i,k} \middle| u_{i,l}\right) = \sigma_{i,k}\left[1+\left(\frac{u_{i,l}}{\sigma_{i,l}}\right)^{\beta}\right]^{1/\beta} \frac{\Gamma(2-\frac{1}{\beta})\Gamma(\frac{\beta+1}{\beta})}{\Gamma(2)}, \forall k \neq l$$

where Γ denotes the Gamma function. We set $\beta = 3$, and solve for σ_i.

Following Proposition 4 in [28] and using the subject-specific parameterization, we generated the random variables following Pareto Type III distributions as

$$\text{(vii)}\quad u_{i,k} = \sigma_{i,k}\left(\frac{Y_{i,k}}{Z_i}\right)^{\frac{1}{\beta}}, \quad for\ k \in \{1,2\}$$

where $Y_{i,k} \sim Exp(\lambda = 1)$ and $Z_i \sim Exp(\lambda = 1)$ independently of all $Y_{i,k}$. Fig 2C presents three examples for such distributions with different values for ρ_i.

Note that using only 320 draws may lead to under-sampling of the distributions. Therefore, to fully capture the shape of the distribution, for each subject, we first generated joint Pareto distributions with 100K draws. We then created small 600-draw experimental distributions that matched the large 100k-draw distributions, allowing a deviation of up to 0.2 utils from the actual first and second moments (mean and standard deviation) of the large 100k-draws sets. Fig 2D compares matched and unmatched small sets, corresponding to the large 100k-draws set presented in Fig 2C (middle panel). Finally, we truncated the long tail of the Pareto Type III distributions at $u_i^{max} = 60^{\rho_i}$ (eliminating 6.5 to 23.83 percent of the distribution, depending on the ρ parameter, the curvature of the subjective value function), to match the upper bound of the uniform distribution and to avoid extreme reward amounts. We then sampled 320 SVs at random from the remaining valuations, which constituted the experimental subject-specific Pareto distributions.

Generating Binary Choice Sets from the Distributions of Valuations. The final step was to generate lottery dollar amounts from the SV distributions. For each lottery k with a valuation u_k, we first randomly drew the first monetary payoff $x_{1,k}$ from a range of possible payoffs \$0-x^{max} in \$0.10 increments. We had to restrict the maximum value of $x_{1,k}$ to make sure that including it in the lottery, does not exceed the lottery valuation (u_k), and thus to avoid negative values for the second lottery payoff. We determined the maximal value of the first payoff $x_{1,k}$ using the minimum function:

(viii) $x_{1,k}^{max} = \min\{(2u_k)^{\frac{1}{\rho}}, 60\}$.

We then solved for $x_{2,k}$ giving rise to the desired u_k, rounded to one decimal place, using the following equation:

(ix) $x_{2,k} = \left(2u_k - (x_{1,k})^{\rho}\right)^{\frac{1}{\rho}}$.

S1 Fig shows how the heterogeneity in ρ values affected the distributions of $x_{1,k}$ and $x_{2,k}$.

We restricted the share of trials with first-order stochastic dominance (FOSD) (trials on which both lottery payoffs of one lottery were higher or equal to the other lottery's payoffs) to 45 percent. For subjects with $\rho_i \rightarrow 0$, we could not generate experimental sets with only 45 percent of the trials. Thus, we fixed $\rho_i = 1$, for all subjects with $\rho_i < 0.$, (a total of 4 subjects, see S2 Table), limiting the interoperability of data from this small number of subjects. In contrast, for two subjects with very high ρ's ($\rho_i > 4$), we also had to fix $\rho_i = 1$ in STAGE II of the study, since a very large tail from their Pareto distribution of SVs exceeded $60. Respectively, the interoperability of data from these subjects is also limited. Nonetheless, we wanted to avoid any unjustified elimination of data, and therefore analyzed data from these six subjects. Importantly, our main qualitative findings do not change once we remove these subjects from our sample.

Procedures

Sessions

Experimental sessions were carried out online via Zoom while subjects completed the task on a website. Data collection took place between autumn 2021 and summer 2022. After instruction, subjects had to successfully answer a set of comprehension questions about the procedure before starting STAGE I. They could participate in STAGE II of the study only if they completed all trials in STAGE I. Subjects received all payments after completing both STAGE I and STAGE II. Subjects received a $10 participation fee and on average $24.5 in STAGE I (range $0-60) and $76.02 in STAGE II (range $7.3-120) from the decision task. All amounts are in Australian dollars. All parts of the experiment were self-paced. Both the valuation and the choice tasks were programmed in the oTree software package [70].

Ethics statements

The study was approved by the local ethics committee at the University of Sydney. All subjects gave informed written consent before participating in the study.

Participants

We recruited participants from various departments at the University of Sydney. Seventy-six subjects (44 females, mean age=21.8, std: 3.34, range: 18-30) passed the comprehension questions and completed STAGE I and the two choice tasks of STAGE II.

Model Fitting

Sample-level (pooled) estimates. We estimated subjects' aggregated choice data via a probit choice function with maximum likelihood estimation (MLE). Standard errors were clustered at the subject level. Thus, in the pooled estimation subjects were treated as one representative decision-maker. In this analysis, we used lotteries' monetary rewards (as opposed to their subjective valuations) to allow meaningful estimates of DN's M parameter, and to confine the range of lottery payoffs. For both DN and power utility models, we report the results from models estimated on the full dataset and separately on each choice environment. To test the possibility of adaptation of the encoding function to the choice environments, we further report the results from three additional models estimated on the full dataset, which also included a dummy variable indicating the Pareto environment for the reward expectation, M parameter (DN) as $M = constant + M_{Pareto}\, x\, Pareto$ and similarly for the functions' curvature parameters α (DN) and r (power utility), respectively.

Subject-level estimates. *DN.* In each choice environment, we recovered subject-specific estimates of the free parameters, restricting the search space as follows: $\alpha \in [0,1.5]$, $M \in [0, u_i^{max}]$ and $\theta > 0$ (see equation (i) in the text). We employed MLE using the Nelder-Mead algorithm with a max-iteration limit of 1,000 and a stopping criterion of 0.5 tolerance. We initialized M to the distributions' medians. θ was initialized at 0.03, matching the sample-level pooled estimate (see S5 Table), and the α parameter was initialized at 1. For calculating the likelihoods, in each of the 320 trials, we generated 10,000 samples with randomly drawn Gaussian noise. The log-likelihood function was thus given by –

(x) $log\mathcal{L}(\alpha_i, M_i, \theta_i | u_{i,t}) = y_{i,t}\log\left(Pr(y_{i,t} = 1 | u_{i,t})\right) + (1 - y_{i,t})log\left(\Pr(y_{i,t} = 0 | u_{i,t})\right),$

Where $y_{i,t} = \{0,1\}$ indicates the subject's *i* choice in trial $t = \{1, \ldots, 320\}$.

Power utility. We fitted the power utility model to recover subject-specific estimates of the r and θ parameters using a similar procedure. We restricted the search space as follows: $r \in \{0,1.5\}$, and $\theta > 0$ *(*see equation (ii) in the text). θ was initialized at 0.03, matching the sample-level pooled estimate (see S5 Table). For the r parameter, we took ten random initializations in the range $\{0.1,5\}$ with a precision of 5. All other procedures were identical to the DN model.

Acknowledgements

The authors thank the attendees of the Departmental Seminar at the Neuroscience Institute at NYU, the Annual Meeting of the Society for Neuroscience, and the Annual Meeting of the Society for Cognitive Science for their valuable input and discussions on the paper.

Funding Statement

This work was funded by the Australian Research Council (Grant No. DP190100489 to AT and PG). A.B. acknowledges financial support from NYU Stern School of Business, NYU Shanghai, and J.P. Valles.

Authors Contributions

V.K.D., V.A., S.B., A.B., K.L., P.G. and A.T. designed the study. V.A. collected the data. V.K.D, S.S. and V.A. analyzed the data. V.K.D., P.G. and A.T. wrote the manuscript.

Data Availability

All data needed to evaluate the conclusions in the paper are present in the paper and/or the Supplementary Materials. All data and code used in the paper are available on figshare at 10.6084/m9.figshare.c.8183615.

Competing Interests

The authors declare that they have no competing interests.

References

1. Simoncelli EP, Olshausen BA. Natural Image Statistics and Neural Representation. Annu. Rev. Neurosci. 2001;24:1193–216.

2. Glimcher PW. Efficiently irrational: deciphering the riddle of human choice. Trends in Cognitive Sciences 2022;26:669–87.

3. Frydman C, Jin LJ. Efficient Coding and Risky Choice. The Quarterly Journal of Economics 2022;137:161–213.

4. Stewart N, Reimers S, Harris AJL. On the Origin of Utility, Weighting, and Discounting Functions: How They Get Their Shapes and How to Change Their Shapes. Management Science 2015;61:687–705.

5. Alempaki D, Canic E, Mullett TL, Skylark WJ, Starmer C, Stewart N, et al. Reexamining How Utility and Weighting Functions Get Their Shapes: A Quasi-Adversarial Collaboration Providing a New Interpretation. Management Science 2019;65:4841–62.

6. Tymula A, Plassmann H. Context-dependency in valuation. Current Opinion in Neurobiology 2016;40:59–65.

7. Khaw MW, Glimcher PW, Louie K. Normalized value coding explains dynamic adaptation in the human valuation process. Proceedings of the National Academy of Sciences 2017;114:12696–701.

8. Heng JA, Woodford M, Polania R. Efficient sampling and noisy decisions. eLife 2020;9:e54962.

9. Louie K, Glimcher PW. Efficient coding and the neural representation of value. Annals of the New York Academy of Sciences 2012;1251:13–32.

10. Polanía R, Woodford M, Ruff CC. Efficient coding of subjective value. Nat Neurosci 2019;22:134–42.

11. Khaw MW, Li Z, Woodford M. Cognitive Imprecision and Small-Stakes Risk Aversion. The Review of Economic Studies 2021;88:1979–2013.

12. Rustichini A, Conen KE, Cai X, Padoa-Schioppa C. Optimal coding and neuronal adaptation in economic decisions. Nat Commun 2017;8:1208.

13. Steverson K, Brandenburger A, Glimcher P. Choice-theoretic foundations of the divisive normalization model. Journal of Economic Behavior & Organization 2019;164:148–65.

14. Schwartz O, Simoncelli EP. Natural signal statistics and sensory gain control. Nat Neurosci 2001;4:819–25.

15. Carandini M, Heeger DJ. Normalization as a canonical neural computation. Nat Rev Neurosci 2012;13:51–62.

16. Olsen SR, Bhandawat V, Wilson RI. Divisive Normalization in Olfactory Population Codes. Neuron 2010;66:287–99.

17. Heeger DJ. Normalization of cell responses in cat striate cortex. Visual Neuroscience 1992;9:181–97.

18. Schwartz O, Simoncelli E. Natural Sound Statistics and Divisive Normalization in the Auditory System [Internet]. In: Advances in Neural Information Processing Systems. MIT Press; 2000 [cited 2023 Feb

1]. Available from: https://proceedings.neurips.cc/paper/2000/hash/96c5c28becf18e71190460a9955aa4d8-Abstract.html

19. Carandini M, Heeger DJ, Movshon JA. Linearity and Normalization in Simple Cells of the Macaque Primary Visual Cortex. J. Neurosci. 1997;17:8621–44.

20. Kaliukhovich DA, Vogels R. Divisive Normalization Predicts Adaptation-Induced Response Changes in Macaque Inferior Temporal Cortex. J. Neurosci. 2016;36:6116–28.

21. Kello CT, Brown GDA, Ferrer-i-Cancho R, Holden JG, Linkenkaer-Hansen K, Rhodes T, et al. Scaling laws in cognitive sciences. Trends in Cognitive Sciences 2010;14:223–32.

22. Chater N, Brown GDA. Scale-invariance as a unifying psychological principle. Cognition 1999;69:B17–24.

23. Rangel A, Clithero JA. Value normalization in decision making: theory and evidence. Current Opinion in Neurobiology 2012;22:970–81.

24. Louie K, Grattan LE, Glimcher PW. Reward Value-Based Gain Control: Divisive Normalization in Parietal Cortex. J. Neurosci. 2011;31:10627–39.

25. Louie K, Khaw MW, Glimcher PW. Normalization is a general neural mechanism for context-dependent decision making. Proceedings of the National Academy of Sciences 2013;110:6139–44.

26. Gluth S, Kern N, Kortmann M, Vitali CL. Value-based attention but not divisive normalization influences decisions with multiple alternatives. Nat Hum Behav 2020;4:634–45.

27. Bavard S, Palminteri S. The functional form of value normalization in human reinforcement learning. eLife 2023;12:e83891.

28. Bucher SF, Brandenburger AM. Divisive normalization is an efficient code for multivariate Pareto-distributed environments. Proc. Natl. Acad. Sci. U.S.A. 2022;119:e2120581119.

29. Wainwright MJ, Schwartz O, Simoncelli EP. Natural Image Statistics and Divisive Normalization [Internet]. In: Rao RPN, Olshausen BA, Lewicki MS, editors. Probabilistic Models of the Brain. The MIT Press; 2002 [cited 2024 June 27]. page 203–22.Available from: https://direct.mit.edu/books/book/2733/chapter/73923/Natural-Image-Statistics-and-Divisive

30. Solé RV, Levin SA, Brown JH, Gupta VK, Li BL, Milne BT, et al. The fractal nature of nature: power laws, ecological complexity and biodiversity. Philosophical Transactions of the Royal Society of London. Series B: Biological Sciences 2002;357:619–26.

31. Newman M. Power laws, Pareto distributions and Zipf's law. Contemporary Physics 2005;46:323–51.

32. Mcfadden D. Economic Choices. The American Economic Review 2001;91.

33. Jacob Marschak. Binary Choice Constraints on Random Utility Indicators. 1959;

34. Glimcher PW, Tymula AA. Expected subjective value theory (ESVT): A representation of decision under risk and certainty. Journal of Economic Behavior & Organization 2023;207:110–28.

35. Barsky RB, Juster FT, Kimball MS, Shapiro MD. Preference Parameters and Behavioral Heterogeneity: An Experimental Approach in the Health and Retirement Study*. Q J Econ 1997;112:537–79.

36. Choi S, Fisman R, Gale D, Kariv S. Consistency and Heterogeneity of Individual Behavior under Uncertainty. American Economic Review 2007;97:1921–38.

37. Steffen Andersen, Glem W. Harirson, Morten I. Lau, Elisabet Rutström. Eliciting Risk and Time Preferences. Econometrica 2008;76:583–618.

38. Holt CA, Laury SK. Risk Aversion and Incentive Effects. American Economic Review 2002;92:1644–55.

39. Luce RD, Krumhansl CL. Measurement, scaling, and psychophysics. In: Stevens' handbook of experimental psychology: Perception and motivation; Learning and cognition, Vols. 1-2, 2nd ed. Oxford, England: John Wiley & Sons; 1988. page 3–74.

40. Gourinchas PO, Parker JA. Consumption Over the Life Cycle. Econometrica 2002;70:47–89.

41. Bleichrodt H, Quiggin J. Life-cycle preferences over consumption and health: when is cost-effectiveness analysis equivalent to cost–benefit analysis? Journal of Health Economics 1999;18:681–708.

42. Peter Wakker. Explaining the characteristics of the power (CRRA) utility family. Health Economics 2008;17:1329–44.

43. Rolls ET, Grabenhorst F, Deco G. Choice, difficulty, and confidence in the brain. NeuroImage 2010;53:694–706.

44. Webb R, Glimcher PW, Louie K. The Normalization of Consumer Valuations: Context-Dependent Preferences from Neurobiological Constraints. Management Science 2021;67:93–125.

45. Schwartz O, Hsu A, Dayan P. Space and time in visual context. Nat Rev Neurosci 2007;8:522–35.

46. Kurtz-David V, Alladi V, Bucher S, Brandenburger A, Louie K, Glimcher P, et al. Choosers Adapt Value Coding to the Environment, But Do Not Attain Efficiency [Internet]. In: Proceedings of the Annual Meeting of the Cognitive Science Society. 2023 [cited 2024 Aug 14]. Available from: https://escholarship.org/uc/item/7cp9r5hc

47. Kahneman D, Tversky A. Prospect Theory: An Analysis of Decision under Risk. Econometrica: Journal of the Econometric Society 1979;47:263–91.

48. Sophia I. Passy. Power law relationships among hierarchical taxonomic categories in algae reveal a new paradox of the plankton. A Journal of Macroecology 2006;15:528–35.

49. Enquist BJ. Universal scaling in tree and vascular plant allometry: toward a general quantitative theory linking plant form and function from cells to ecosystems. Tree Physiology 2002;22:1045–64.

50. McGill BJ, Etienne RS, Gray JS, Alonso D, Anderson MJ, Benecha HK, et al. Species abundance distributions: moving beyond single prediction theories to integration within an ecological framework. Ecology Letters 2007;10:995–1015.

51. Mark Q. Wilber, Justin Kitzes, John Harte. Scale collapse and the emergence of the power law species–area relationship. A Journal of Macroecology 2015;24:883–95.

52. Gabaix X. Power Laws in Economics and Finance. Annual Review of Economics 2009;1:255–94.

53. Gabaix X. Zipf's Law and the Growth of Cities. American Economic Review 1999;89:129–32.

54. Kim BJ, Singh V, Winer RS. The Pareto rule for frequently purchased packaged goods: an empirical generalization. Mark Lett 2017;28:491–507.

55. Kahneman D, Tversky A. Prospect Theory: An Analysis of Decision Under Risk [Internet]. In: Handbook of the Fundamentals of Financial Decision Making. WORLD SCIENTIFIC; 2012 [cited 2024 June 18]. page 99–127.Available from: https://www.worldscientific.com/doi/abs/10.1142/9789814417358_0006

56. Botond Koszegi, Matthew Rabin. A model of reference-dependent preferences. The Quarterly Journal of Economics 2006;121:1133–65.

57. Simoncelli EP, Olshausen BA. Natural Image Statistics and Neural Representation. Annual Review of Neuroscience 2001;24:1193–216.

58. Trueblood JS, Brown SD, Heathcote A, Busemeyer JR. Not Just for Consumers: Context Effects Are Fundamental to Decision Making. Psychol Sci 2013;24:901–8.

59. Tversky A, Simonson I. Context-Dependent Preferences. Management Science 1993;39:1179–89.

60. Shen B, Nguyen D, Wilson J, Glimcher PW, Louie K. Early versus late noise differentially enhances or degrades context-dependent choice. Nat Commun 2025;16:3828.

61. Louie K, Glimcher PW, Webb R. Adaptive neural coding : from biological to behavioral. Current Opinion in Behavioral Sciences 2015;5:91–9.

62. Khaw MW, Glimcher PW, Louie K. Normalized value coding explains dynamic adaptation in the human valuation process. Proceedings of the National Academy of Sciences 2017;114:12696–701.

63. Reynolds JH, Heeger DJ. The Normalization Model of Attention. Neuron 2009;61:168–85.

64. Schwartz O, Coen-Cagli R. Visual attention and flexible normalization pools. Journal of Vision 2013;13:25.

65. Hertwig R, Barron G, Weber EU, Erev I. Decisions from Experience and the Effect of Rare Events in Risky Choice. Psychol Sci 2004;15:534–9.

66. Hertwig R, Erev I. The description–experience gap in risky choice. Trends in Cognitive Sciences 2009;13:517–23.

67. Barranca VJ, Zhu XG. A computational study of the role of spatial receptive field structure in processing natural and non-natural scenes. Journal of Theoretical Biology 2018;454:268–77.

68. Heurteloup C, Merchie A, Roux S, Bonnet-Brilhault F, Escera C, Gomot M. Neural repetition suppression to vocal and non-vocal sounds. Cortex 2022;148:1–13.

69. Becker GM, DeGroot MH, Marschak J. Measuring utility by a single-response sequential method. Syst. Res. 1964;9:226–32.

70. Chen DL, Schonger M, Wickens C. oTree—An open-source platform for laboratory, online, and field experiments. Journal of Behavioral and Experimental Finance 2016;9:88–97.

Supplementary Information

S1 Fig. Representative Sets in STAGE II. Left – a risk averse subject, middle – a risk neutral subject, right – a risk seeking subject. Top to bottom: (1) Distributions of the high winning amount in Lottery 1 (in dollars); (2) Distributions of the low winning amount in lottery 1 (in dollars); (3) Distribution of the expected earnings (EV) of Lottery 1 (in dollars); (4) Distributions of the valuations (u_1) of Lottery 1 (in util units); (5) 2-dimensional histogram of the valuations of Lottery 1 and Lottery 2 (u_1 and u_2, in util units).

S2 Fig. Descriptive statistics. (A) violins show the share of trials in which subjects chose the lottery with the higher subjective value. (B) violins show the number of FOSD violations per subject. Dots indicate individual subjects. N=76. (C-D) Share of trials in which subjects chose the lottery with the higher SV, first half of the session (trials 1-160), compared with the second half of the session (trials 161-320). Each gray line indicates a subject. Colored lines are sample averages. (C) Pareto distribution sessions. (D) Uniform distribution sessions.

S3 Fig. Noise estimates. Comparing the best-fitting σ parameter (decision noise) across the distributional environments reveals noise levels were higher in the Pareto environment. *Left* - DN model (one-sided Wilcoxon sign-rank test, Z=2.2314, p=0.0257). *Right* - Power Utility model (one-sided Wilcoxon sign-rank test, Z=2.9172, p=0.0035). Scatters indicate individual subjects. N=76.

S4 Fig. Distributions' medians. (A-B) Distributions of median payoffs (in dollar amounts), (A) Pareto, (B) uniform. (C) median of monetary payoffs across environments vs subjects' risk preferences, captured by the ρ parameter from STAGE I. (D) change in the true median of monetary payoffs across environments vs subjects' ρ parameter (scatters are equivalent to the grey lines in (C)). (E) Change in the true median payoff across environments compared with the change in the recovered $\widetilde{M}_i$ parameter across the two environments. (F) Subjects' risk (ρ parameter) compared with the change in the recovered $\widetilde{M}_i$ parameter across the two environments. (D-F) Each scatter represents one subject. N=76.

S5 Fig. Model-fitting, Power utility compared with a DN model where M is fixed. *Top* - Each dot is one subject's DN model BIC score (y-axis) plotted against the same subject's power utility BIC score (x-axis).

A dashed 45-degree line indicates when both models are equally successful. *Bottom* - the difference in BIC scores (BIC_{power} – BIC_{DN}). Left panels show the uniform environment, and right shows the Pareto environment.

S1 Table. Lotteries used in STAGE I.

S2 Table. Individual-level estimates of risk preferences from subjects' bids in STAGE I. (*) For these subjects we could not generate distributions of valuations for STAGE II that would adhere to our requirement to limit the number of trials with FOSD violations (when $\rho_i \rightarrow 0$), or without having to censor a very large tail of the Pareto distribution (when $\rho_i > 4$). Instead, for these subjects we plugged-in $\rho_i = 1$ to generate the distributions for STAGE II.

S3 Table. Robustness checks for the findings presented in Column (2) in Table 1. We vary the definitions for *center of the distributions* (center) and *around the diagonal* (diagonal). Column (1) corresponds to the regression presented in the Main Text.

S4 Table. Individual-level best-fitting model parameters across environments (STAGE II). (*) Subjects who had either a STAGE I estimate of $\rho_i = 0$ or $\rho_i > 4$. For those subjects, we could not generate distributions of valuations for STAGE II that would adhere to our requirement to limit the number of trials with FOSD violations (when $\rho_i \rightarrow 0$), or without having to censor a very large tail of the Pareto distribution (when $\rho_i > 4$). Instead, for these subjects we plugged-in $\rho_i = 1$ to generate the distributions for STAGE II.

(**) Subjects who had >20 FOSD violations in at least one of the treatments.

S5 Table. Pooled estimates, dollar space. The table shows recovered parameters for the DN model (top rows), the DN model where M is fixed as the true median of the distributions (middle rows), and the Power utility model (bottom rows). In practice, to allow a better identification of the model parameters, we estimated the parameter τ, such that $\tau = M^{\alpha}$. We recovered M post-hoc by simply plugging-in τ and α into the equation. Standard errors in parentheses, + p<0.1, * p<0.05, ** p<0.01, *** p<0.001.

S6 Table. Behavioral dynamics of fitted parameters. Pooled estimates in dollar space, early vs late trials in each statistical environment. In practice, to allow a better identification of the model parameters, we estimated the parameter τ, such that $\tau = M^{\alpha}$. We recovered M post-hoc by simply plugging-in τ and α into the equation. Standard errors in parentheses, + p<0.1, * p<0.05, ** p<0.01, *** p<0.001.

S7 Table. Model-free analysis, alternative model. Same analysis as in Table 1, while controlling for potential discontinuity in the model. MLE estimation with a Bernoulli model where the inverse link cannot drop below the chance level (0.5), such that $p_i = 0.5 + 0.5\,\Phi(x_i{'}\beta)$. Variables and specifications are identical to Table 1. Standard errors clustered on subject in parentheses, + p<0.1, * p<0.05, ** p<0.01, *** p<0.001.