Origins of noise in both improving and degrading decision making

Bo Shen[1]*, Jailyn Wilson[3], Duc Nguyen[2], Paul W. Glimcher[1,2†], Kenway Louie[2,1†]
[1] New York University, Grossman School of Medicine, New York, NY, 10016
[2] New York University, Center for Neural Science, New York, NY, 10003
[3] Department of Psychology, Cornell University, Ithaca, NY 14853

† These authors contributed equally.
* Corresponding at bs3667@nyu.edu

Abstract

Unlike the predictions of normative choice theories, the real-world accuracy of human and animal decision-making depends on context. Observed impairments of choice performance by context can be explained by computations like divisive normalization. However, recent behavioral evidence suggests that context can sometimes enhance choice performance, a finding that challenges existing explanations. Here, we propose and test an extension of existing frameworks for context-dependent choice that incorporates a richer notion of noise. Model simulations show that noise arising before (early noise) and after (late noise) normalization predict opposing effects: early noise can cause contextual information to enhance choice performance while late noise can cause context to degrade choice performance. In experiments with human subjects, we confirm key elements of this model. Manipulating early and late noise - by inducing uncertainty in option values and controlling time pressure – produces dissociable positive and negative effects of context. Together, these findings suggest a new unifying mechanism for contextual modulation, highlighting the significant role of noise source in neural computation and behavior.

Keywords: Noise, context-dependence, decision-making, choice, normalization

Introduction

Classical rational decision theory requires that choices and choice accuracy should be context-independent. For instance, a preference for an apple over an orange should not be influenced by the presence of a banana - a property known as *independence of irrelevant alternatives*, or IIA (Luce, 1959; Von Neumann & Morgenstern, 1944). Nevertheless, a wealth of empirical evidence gathered in species ranging from humans to insects has shown that real-world decision-making is inherently context-dependent. The presence of an additional option, even when this option is never chosen, will change choice behavior between other options in the choice set (Bateson et al., 2003; Chau et al., 2014, 2020; Glimcher & Tymula, 2023; Holper et al., 2017; Huber et al., 1982; Louie et al., 2013; Shafir et al., 2002; Soltani et al., 2012; Tversky & Simonson, 1993; Webb et al., 2021). Context-dependent choices are widely documented in choice tasks involving multiple attributes, where the options differ in multiple feature dimensions (see review in Trueblood, 2022). Beyond that, recent findings show context-dependent preferences can arise even when option values are unidimensional (D. Cohen et al., 2019; Louie et al., 2011, 2013). Such value-driven effects suggest that contextual processing is a general feature of biological decision-making, but the specific neural mechanisms behind context-dependent choice remains unclear.

One promising explanation for context-dependent choice is the divisive normalization computation, a widely documented nonlinear neural computation seen in diverse species, brain regions, and cognitive processes, that has been proven to be optimal for many environments (Bucher & Brandenburger, 2022). The ubiquity of this computation suggests it is canonical and pervasive in neural systems (Heeger, 1992; Heeger et al., 1996; Louie et al., 2011, 2013, 2015; Schwartz & Simoncelli, 2001). Functionally, divisive normalization adjusts the representation of input-driven activity in a neuronal pool by dividing it by the activity from other inputs, a general feature required of all optimal systems that have limited capacity (Steverson et al., 2019). This operation leads to a form of contextual suppression. Additional contextual inputs inhibit the activities driven by the pre-existing inputs and can thus lead to reduced precision in the internal representation. In sensory brain areas, normalization-driven contextual suppression leads to phenomena like surround suppression (Carandini & Heeger, 2012; Heeger, 1992); in decision-related brain areas, normalization produces relative value coding (Glimcher, 2022; Louie et al., 2011; Pastor-Bernier & Cisek, 2011) . In choice, contextual suppression reduces the neural representation of option values when additional contextual options are present. This suppression predicts that choice accuracy between two fixed options will decrease as a function of the value of a third distracter alternative in the choice set. This has been referred to as a *negative distracter effect* and has been observed experimentally in animals and humans (Chau et al., 2014; D. Cohen et al., 2019; Holper et al., 2017; Louie et al., 2013; Noonan et al., 2017; Otto et al., 2022; Webb et al., 2021). Analogous extended normalization models explain multi-attribute context effects (Daviet & Webb, 2023; Dumbalska et al., 2020; Webb et al., 2022), suggesting that divisive normalization plays a central role in context-dependent preferences.

However, recent empirical studies have also begun to document contextual effects on choice accuracy at odds with the predictions of the standard normalization model. Some studies have reported no overall effect of distracter value on choice accuracy (Gluth et al., 2020), while other studies have identified a *positive distracter effect*: at least in certain conditions, the value of a contextual option can enhance the accuracy of choosing between other options (Chau et al., 2014, 2020; Shevlin et al., 2022; Webb et al., 2020). Empirically, both positive and negative distracter effects can coexist in the same experiment, leading to the proposal that the two phenomena reflect different circuit mechanisms, perhaps operating in different brain areas (Kohl et al., 2023). These observations challenge the standard normalization-based understanding of context-dependent preferences and underscore the need for a unifying framework for all context effects.

Here, we extend the standard framework, adding a richer definition of neural noise. We find that with this richer definition of noise, a richer set of context effects should be observable. Standard normalization models of stochastic choice behavior (e.g., Glimcher, 2022; Glimcher & Tymula, 2023; Louie et al., 2013; Webb et al., 2021) represent noise as being added to the internal representation after the normalization process is complete, implying that the value inputs into the normalization process must be noiseless and thus deterministic. However, this assumption overlooks the fact that neural value representation is inherently uncertain, with uncertainty inherited from a noisy sensory environment (Drugowitsch et al., 2016) or originating from internal processes such as inference or memory retrieval of reward-related information (Findling et al., 2019; Shadlen & Shohamy, 2016; Shohamy & Daw, 2015). It has been previously noted that noise arising early or late in the decision-making process can lead to distinct impacts on choice under specific neural computations (Tsetsos et al., 2016). Here we extend that notion, showing that the stage at which noise enters the valuation process — early in the process of representing values prior to normalization or later during the comparison process post-normalization — can critically influence the direction of observed context effects.

We find theoretically that a context-setting option (for example a low-valued distracter) can actually improve choice accuracy by suppressing early noise. This stands in contrast to the detrimental effect of the same context-setting option in the presence of more traditional post-normalization noise. To empirically validate our predictions, we devised an experimental design allowing us to tease apart the impacts of early and late noise on the contextual effects of decision-making. Our theoretical and empirical findings reveal that when early noise dominates, a contextual item enhances choice accuracy; in contrast, when late noise dominates, a contextual item impairs choice accuracy. These results reconcile a broad and conflicting literature (Chau et al., 2014, 2020; Gluth et al., 2020; Louie et al., 2013; Webb et al., 2020), supporting a canonical role for normalization in neural computation and emphasizing the critical importance of different types of uncertainty in neural coding.

Results

Divisive normalization under different noise sources

While it is known that divisive normalization shapes the noise structure of neural response variability (Beck et al., 2011; Coen-Cagli & Solomon, 2019; Weiss et al., 2023), how neural noise and normalization work together to affect cognitive processes and behavior is unknown. Here, we examine how neural noise arising from different stages of value normalization affects context effects in decision-making: *early noise*, which arises prior to the normalization process, and *late noise*, which arises after normalization (**Figure 1a)**. We conceptualize early noise as the variability inherent in the input values feeding into the normalization process. This variability captures uncertainties in a subject's ability to assign a valuation to an individual item, which can result from the imprecision of sensory inputs, the vagueness of reward-related associations, and the fluctuation of the decision-maker's motivational states. In contrast, we conceptualize late noise as the stochasticity of neural processing after normalization, for example, in post-normalization value coding or when performing a choice comparison amongst multiple options that have already been efficiently represented.

To theoretically examine the differential impact of early and late noise in context effects, we extended a previous algorithmic model of divisive normalization for stochastic choice behavior (Louie et al., 2013). In this model, the represented value of each option (FR_i) is conceptualized as neural activity, akin to the firing rates of a neuronal pool, driven under the direct input of the option (V_i) and normalized by the integration of the input from all options (Eq. 1). Unlike the previous model, which assumes no noise on the input values (V_i), we included an early noise term regulating the variability of each input. This early noise is defined by a zero-mean *Gaussian* distribution additive to its mean value, $\varepsilon \sim N(0, \sigma_\varepsilon^2)$. Late noise is then added onto the product of divisive normalization as drawn from another zero-mean *Gaussian* distribution, $\eta \sim N(0, \sigma_\eta^2)$. A decision is subsequently made by choosing the option with the largest represented value under the two sources of noise; see *Methods for* other details of implementation.

$$FR_i = R_{max} \frac{V_i + \varepsilon_i}{\sigma_H + w \sum_j (V_j + \varepsilon_j)} + \eta_i \qquad \text{Equation 1}$$

where R_{max} is the maximum neural response, σ_H is a semi-saturation constant determining the baseline normalization, w is the weight of the normalization, ε_i is the early noise, and η_i is the late noise.

Normalization under early and late noise predicts opposite contextual effects

To examine how early and late noise contribute to context-dependent choice, we simulated the classic trinary choice decision-making experiment of Louie and colleagues (Gluth et al., 2020; Louie et al., 2013). In this widely used approach to

quantifying contextual choice effects, two higher-value options (V_1 and V_2), *the targets*, are used to probe the impact of a context-setting option (V_3), *the distracter*. In this framework, target-coding neural activities consist of two distributions, and stochastic choice between targets arises from the overlap between those distributions. Context effects are evident when the presence and value of the distracter option changes the overlap between target distributions (**Figures 1b,c**).

In previous work, stochastic choice and the contextual effect of the distracter is presumed to derive from late noise (e.g., Chau et al., 2020; Louie et al., 2013; Webb et al., 2021). With such post-normalization late noise, a contextual input suppresses the mean activities of other options in the choice set, leading to diminished discriminability between target activities and, consequently, impaired choice accuracy. Larger distracter values lead to more impaired choice, therefore producing a negative context effect, which has been shown in previous computational and empirical data (Glimcher & Tymula, 2023; Louie et al., 2013; Webb et al., 2021). Here, we reproduced the original visualization of this process in **Figure 1c**. This reduction in choice accuracy is evident in the signal-to-noise ratio (SNR, measured as the inverse of the Coefficient of Variance) of individual option representations: when the distracter value increases, the model representation of each target becomes relatively noisier (**Figure 1d**, right panel). Further, given the divisive nature of normalization, increasing contextual value reduces the space between the two target distributions and, therefore, enlarges their amount of overlap (**Figure 1e**, right panel). At the behavioral level, these changes lead to a decrease in conditional choice accuracy (relative choice accuracy between the two target options, excluding distracter choices) (**Figure 1f,** right panel). Note that when the distracter value increases to the point at which it becomes roughly comparable to that of the targets, a 'hook' in the accuracy curve occurs because the distracter begins to compete selectively with the second-best target; decreasing the fraction of trials on which subjects choose the second-best target naturally increases the relative ratio of choosing the best target over the second-best target. This phenomenon of *conditional choice accuracy* has been well explained in previous research (Louie et al., 2013) and is not specific to divisive normalization.

In contrast to these well-described effects, we report here that when noise is introduced prior to normalization, distracter values can enhance choice accuracy. Under early noise, a contextual input suppresses the noise associated with the targets during the normalization process. This suppression leads to an improved representation of individual targets by decreasing their noise relative to the mean activity, evident as sharpened activity distributions (**Figure 1b**) and increasing SNR with larger distracter values (**Figure 1d**, left panel). The improved representation leads to decreased overlap between the target distributions even when the distance between their mean values is reduced (**Figure 1e**, left panel). Overall, that improves conditional choice accuracy (**Figure 1f**, left panel). Thus, the interaction of early noise and normalization predicts an increase in choice accuracy with larger distracter values, a positive distracter effect. Note that, as discussed above with regard to late noise, a steeper inflection in the accuracy curve occurs when the distracter value approaches that of the targets, due to competition between the distracter and the second-best target.

When applying a (biological) constraint that all values below zero must be represented as zero, i.e., the *non-negative constraint for neural activities*, we observed subtle differences in model predictions (solid lines in **Figures 1d-f**) compared to the implementation allowing negative neural activities (dashed lines). Since distracter values are lower than target values by design, imposing a non-negative constraint primarily cuts the variance of the distracter option. This, in turn, improves target representation SNR and increases conditional choice accuracy when the distracter value approaches zero. Such an effect exclusively arises under early noise; applying non-negative constraints under late noise does not appreciably change the effects of context on representations or choice (**Figures 1d-f**, right panels). Although it is not the main focus of the current study, the impact of non-negativity constraints on contextual processing highlights the impact of biological constraints during the computational operations of neural noise (Averbeck et al., 2006).

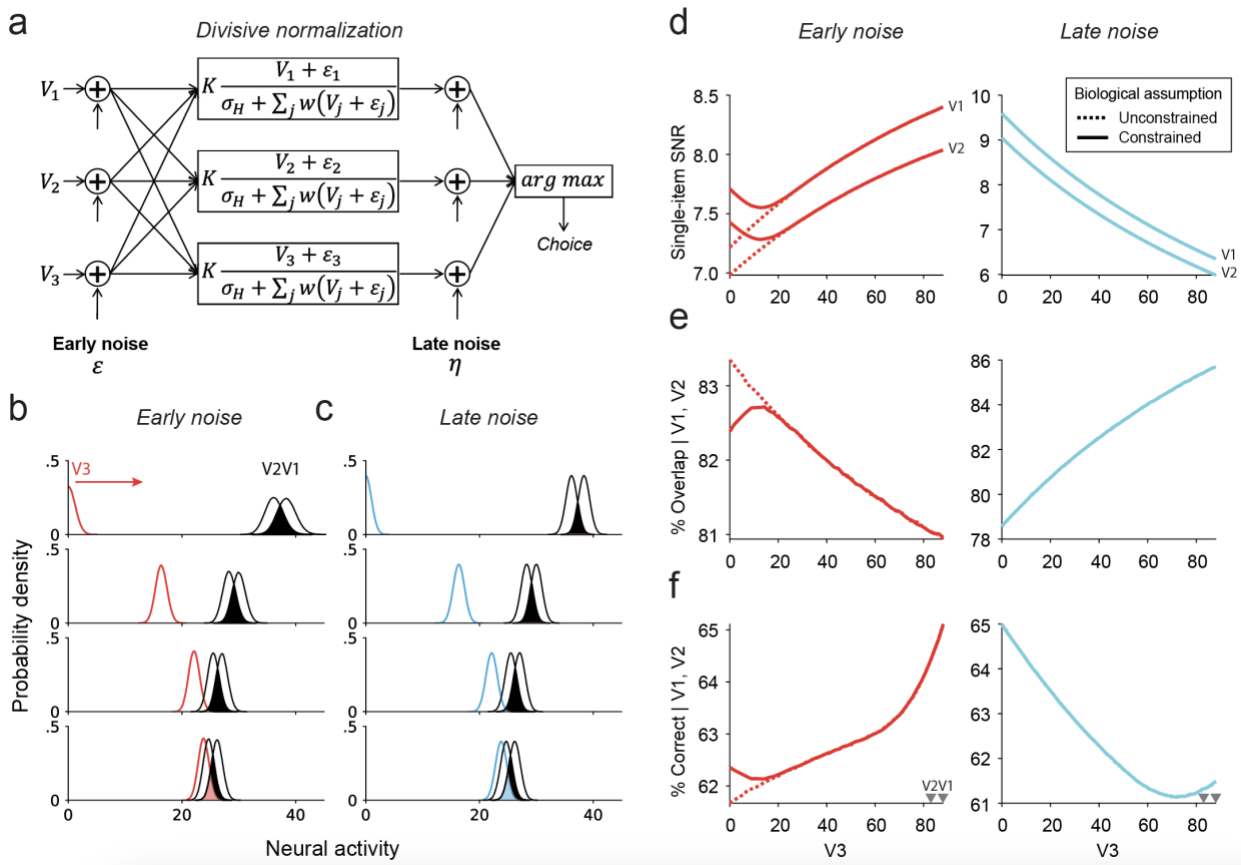

Figure 1. Divisive normalization leads to opposite contextual effects under early and late noise. (**a**) An illustration of the divisive normalization model operating with two sources of noise. (**b**) Under early noise, increased contextual value induced by a distracter (V_3) reduces the mean activities of the targets (V_1 and V_2) by divisive normalization, meanwhile sharpening the target activity distributions. (**c**) Under late noise, increased V_3 reduces the mean activities of the targets; the sharpness of the target activity distributions remains unchanged. (**d**) The signal-to-noise ratio (measured as the ratio of mean over standard deviation) of the target activities can be improved by

V_3 when under early noise (left panel) but impaired when under late noise (right panel). A non-negative threshold was applied for the biologically constrained simulations (solid curves), whereas activities were allowed to be negative in the unconstrained assumption (dashed curves). (**e**) The percentage of overlap between the neural activities of the two target distributions decreases and increases respectively under early (left panel) and late (right panel) noise. (**f**) The divisive normalization model under early and late noise predicts opposite contextual effects. V_3 under early noise facilitates the target choice accuracy; in contrast, V_3 under late noise impairs the target choice accuracy. Under both early and late noise, the choice accuracy curves exhibit a 'hook' of increase in both cases when V_3 approaches V_1 and V_2 due to a different non-normalization mechanism (see text).

Divisive normalization with mixed noise can replicate all context effects

We focused above on the distinct effects of early and late noise separately, whereas, in the brain both types of variability likely co-exist. We thus examined predictions under mixed noise by integrating both sources of noise into the normalization model. We found that opposing contextual effects compete and progressively change from improving to impairing choice accuracy, when early and late noise tradeoff in the model (**Figure 2a**). When early noise predominates, distracter value facilitates choice, evident from the upward trend of the curves (red-to-orange). Conversely, increasing late noise progressively shifts the context effect to impairing choice accuracy, as indicated in the downward slope of the curves (green-to-blue). Therefore, a divisive normalization model with two sources of noise can replicate the range of context effects reported in the literature, either positive, negative, or non-effects (Chau et al., 2014, 2020; Gluth et al., 2020; Kohl et al., 2023; Louie et al., 2013; Webb et al., 2020).

To quantify the context effects that were consistently different between early and late noise conditions, we reasoned that analyses should exclude the ranges of contextual values near zero or near the target values, where simulations show complicated curvatures that are non-specific to normalization and similar effects under both early and late noise. We thus focused on quantifying context effects as the slope of the curves in the middle range of V_3 (between 0.2 and 0.8 when V_3 is scaled to the minimum value of the targets). This measure shows a clear pattern of progressive change in context effects as the degree of noise shifts from early-dominant to late-dominant (inset bar graph, **Figure 2a**).

How robust are the predicted context effects trading off early and late noise? Using the same quantification approach as above, we visualized predicted context effects across a wide range of early and late noise magnitudes (**Figure 2b**). These results support the general conclusion that larger early noise drives positive context effects (red) while larger late noise drives negative context effects (blue). In addition to this general pattern, ceiling and floor effects occur when both sources of noise are small (bottom-left corner) or large (upper-right corner); in these regions, accuracy is either very low (near random choice) or saturates to 100%, minimizing the potential to observe context

effects. By controlling noise magnitudes, we find a clear pattern of tradeoffs between early and late noise effects along the lines of a constant Euclidean norm (lines in **Figure 2c**) of the two sources of noise. In other words, given a constant sum of the two sources of noise, decreasing late noise and increasing early noise always increases the context effect, and vice versa. Thus, the observed variability in contextual effects, contingent on the balance of early and late noise, offers a testable empirical approach to validate our proposed framework.

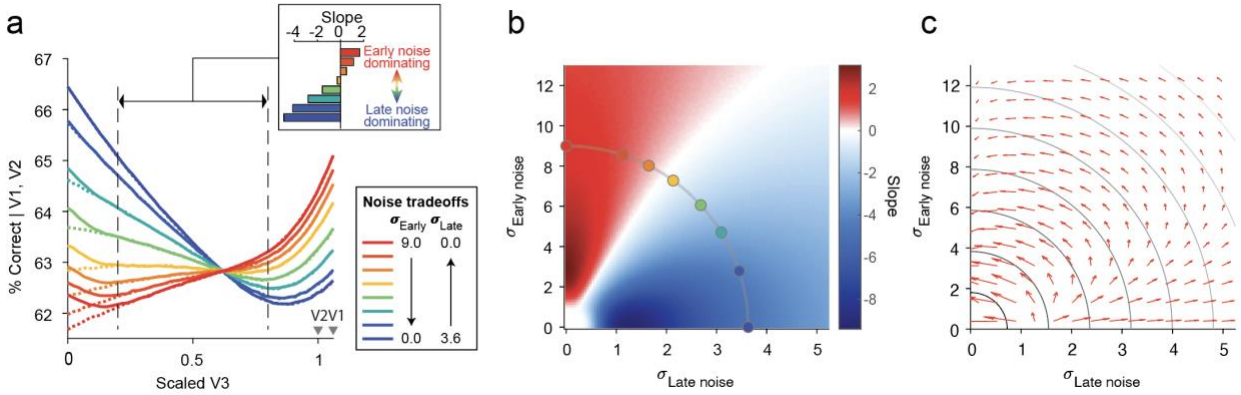

Figure 2. Context effects under mixed noise. (**a**) Early and late noise compete in driving opposite trends of context effects. In the simulations, early noise was reduced from red to blue curves; meanwhile, late noise was increased from blue to red curves. Inset illustrates the regression slope between the middle range of scaled V_3, from 0.2 to 0.8, to eliminate the impacts when distractors approach zero or the targets. (**b**) Early and late noise tradeoffs in a wide range of noise magnitudes. The quantified slopes from the middle range of V_3. Spots along the line have a constant Euclidean norm of the two sources of noise, which were the set parameters visualized in panel **a**. (**c**) The vector field indicates the changing trend of context effects at every combination of early and late noise. The vectors consistently point towards the directions where early noise increases and late noise decreases, along the lines of constant Euclidean norm (lines in grayscale).

Model-driven experimental design

To test the predictions of our model, we designed an experiment to systematically manipulate and quantify the influences of early and late noise on context effects. A 2-by-2 factorial design was employed to create four conditions in a within-subjects design that combined varying degrees of early and late noise (**Figure 3a**). We adapted a two-stage task used in previous studies (Gluth et al., 2020; Louie et al., 2013). In our version of this task, subjects (N = 55) were first presented with different consumer items (e.g., drone, tea kettle, or camera; **Figure 3b**) and provided their subjective monetary valuation for each item (**Figure 3c**) in an incentive-compatible procedure. Subsequently, participants were presented with trinary choices pairing two high-valued items and a

distracter of varying value (**Figure 3d**). They were asked to choose their preferred item from these sets of three items (**Figure 3e**).

The degree of early noise was modulated by presenting items to subjects with varying levels of vagueness. To elicit valuations that were as precise as possible (and thus to reduce early noise), half of the consumer items were physically present in the laboratory and the participants were required to handle them. We hoped in this way to reduce early noise by providing enough information for subjects to form precise evaluations of their desirability. To elicit variable valuations (and thus to increase early noise), the other half of the consumer items were not physically present and were depicted only as cartoon images. This had the effect of rendering their specific characteristics and quality uncertain (**Figure 3b**). Following interaction with the items, participants' valuations were assessed through a Becker-DeGroot-Marschak (BDM) auction task, as depicted in **Figure 3c**. Participants were asked to bid three times for each item, allowing us to characterize the variability of valuations as a proxy for early noise (see *Methods* for details). Bids for all items demonstrated a variance that depended on the mean bid. As expected, the mean-variance relationship was steeper for the vaguely described items than for the items handled by the subjects (difference of slopes: β = 1.21, s.e. = .46, t = 2.62, p = .009; **Figure 3f**). This affirmed the effectiveness of our early-noise manipulation.

The three-item choice sets were constructed based on the early noise structure for our context experiment (**Figure 3d**). We picked distracter items having a range of values, but sorted into those handled by the subjects (precise) and those represented only by cartoons (vague). In contrast, all target options (the higher-valued pair in our design) had been seen and handled by the subjects (members of the precise pool).

To manipulate the degree of late or decisional noise, we imposed a decisional time pressure, requiring subjects to make their choices either quickly or slowly (in a block-wise fashion). This is an approach widely used to increase decisional stochasticity (Cisek et al., 2009). In the low time-pressure block, the participants had 10 s to make a decision (subjects took 1.17 ± .05 s, with .06 ± .04% time out); in the high time-pressure block, the participants had only 1.5 s to choose (subjects took .78 ± .02 s, with 3.72 ± .66% time out; responded significantly faster than under low-pressure, t_{54} = -9.91, p < .001). This is indicated by a color code in **Figure 3e**. Consistent with an effect on late noise, time pressure significantly reduced participants' overall decision accuracy across conditions ($F_{1, 162}$ = 39.17; p < .001; **Figure 3g**).

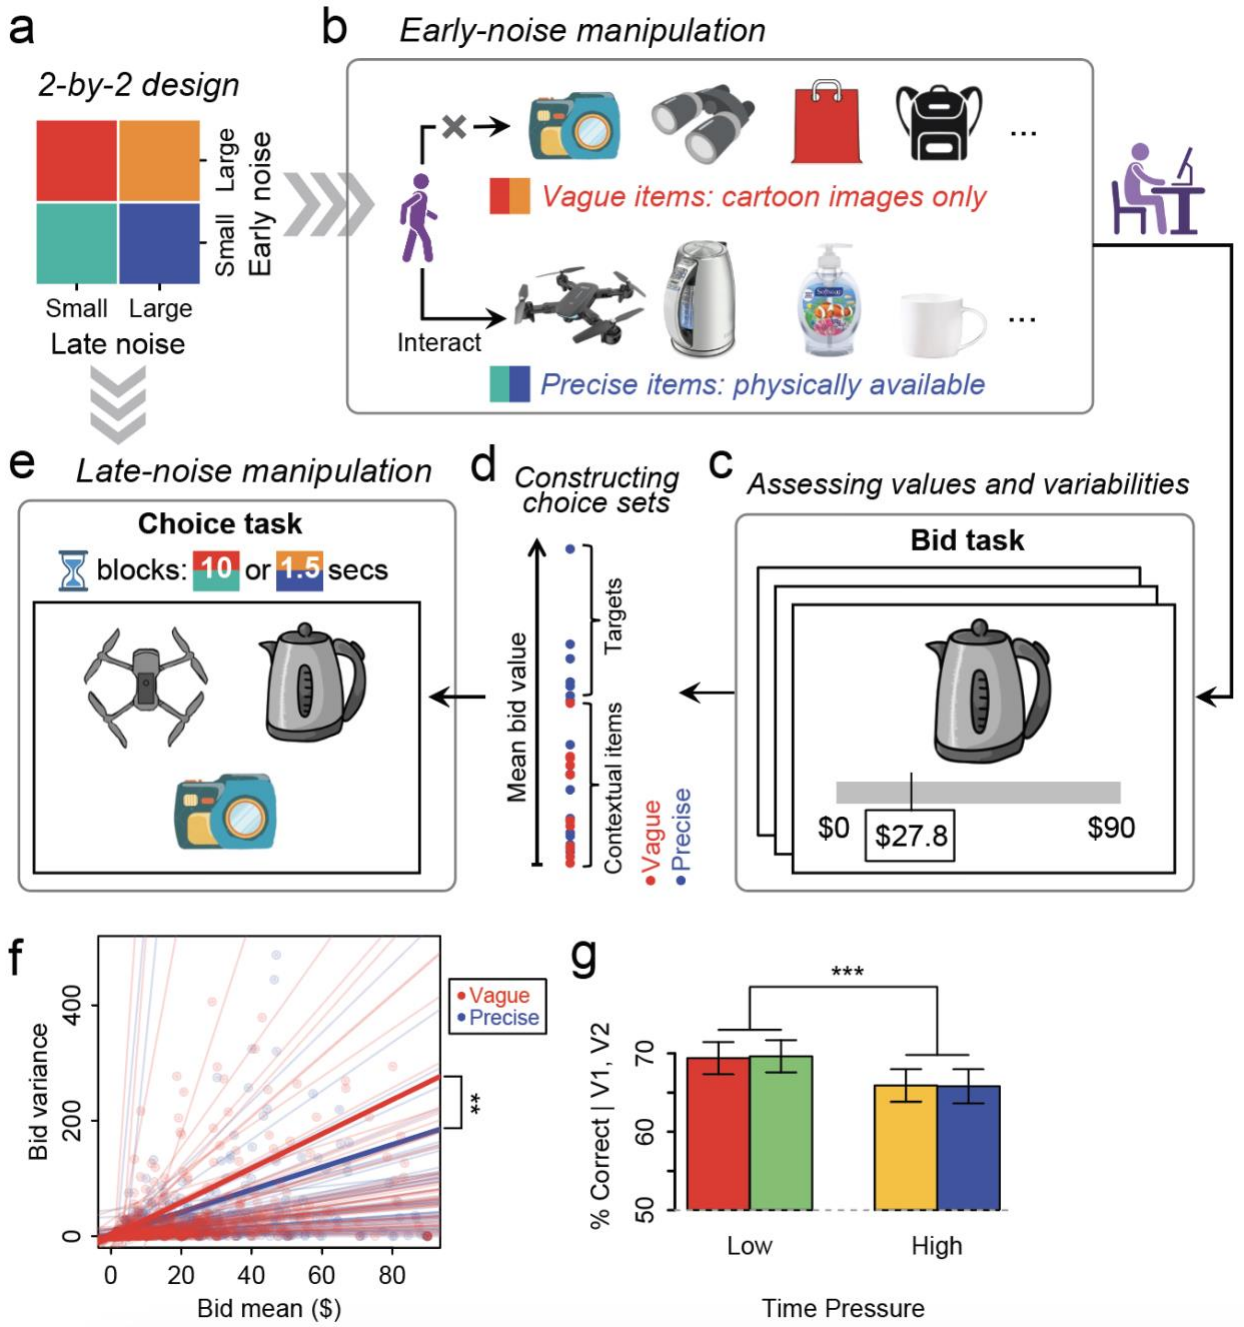

Figure 3. Experimental design and procedure in testing context effects under different sources of noise. (**a**) A 2-by-2 design matrix was used to independently vary the levels of early and late noise. (**b**) Early noise was manipulated by presenting the consumer goods either physically (handled and viewed) or only as cartoon images. This induced two different degrees of precision in object valuations. (**c**) To reveal the subjective value and the noise associated with the representation of each item, participants were asked to BDM-bid on each item three times. (**d**) To build the choice sets for the subsequent task, a fixed set of the targets (V_1 and V_2) was constructed from the six highest-bid precisely presented items, resulting in $C_6^2 = 15$ combinations of target pairs. In each trial, one target pair was combined with one distracter (V_3) selected from

twenty-four lower-bid items, where half of them were precisely presented items. (**e**) In the choice task, the participants were asked to choose their preferred item from a triplet under low- (10 s) and high- (1.5 s) time-pressure blocks. All images were presented as cartoons even when they were precise items to control for visual confounds. The positions of the options were randomized across trials. (**f**) The effect of the manipulation on early noise confirmed that the bid variance increased faster with the bid mean in the vague items than in the precise items (**, p = .009). Each dot indicates an item; thin lines indicate fitting to each individual, and thick lines indicate group averages. (**g**). The effect of the manipulation on late noise confirmed that the participants' choice accuracy decreased when time pressure increased (***, p < .001).

Empirical results closely track model predictions

We first simulated model predictions under conditions matched to the experimental design. By adjusting the noise on the distracters only but fixing the targets, the two opposite effects from early and late noise remain. These simulations are similar to the ones shown in the modeling studies of mixed-noise above: larger early noise drives more positive context-modulation slopes and larger late noise drives more negative context-modulation slopes. The predicted effect of early noise can be seen in **Figure 4a** by comparing the left and the right panels. Simulated choice accuracy under large early noise (left panel) shows more positive context effects than under small early noise (right panel), given the controlled late noise between the panels across paired lines. The effect of late noise can be seen within each panel. Increasing late noise progressively shifts context-effect slopes to be more negative within each panel under fixed early noise (lines within each panel). Overall, these results lead to two main effects on the context-modulation slopes, quantified from the middle range of distracter values (0.2-0.8) and visualized in **Figure 4b**. We test this clear two-fold hypothesis in empirical participant choice behavior below.

In analyzing the conditional choice accuracy for our empirical data, we scaled V_3 to the minimum target value for each individual subject. This normalized the scale of distracter values across subjects allowing us to collapse data across subjects. To extract the trends of the context effects over the entire range of V_3, we pooled the normalized data from all subjects and performed sliding windows along V_3 to reduce the nuisance variability in the data. The full range of data from the four experimental conditions is shown in **Figure 4c**. As noted in the theory section of this report, to focus on dissociable early and late noise effects (see above), we quantified data from the middle range of the scaled V_3 (0.2 to 0.8), illustrated in **Figure 4d**. We found clear context effects across the noise conditions that were consistent with the model predictions. Linear regressions on the factorial design revealed both significant main effects of early and late noise manipulations, with the vagueness of the distracter items driving positive context-modulation slopes and time pressure driving negative context-modulation slopes (Vagueness: β = 11.03, s.e. = 1.54, t_{152} = -7.17, p < .001; time pressure: β = -4.51, s.e. = 1.54, t_{152} = -2.93, p = .004; interaction: β = -4.27, s.e. = 2.15, t_{152} = -1.99, p = .049).

More specifically, vague distracters under low time pressure (red curve in the left panel) – a condition with large early noise and small late noise – led to the most positive context effect (linear regression: β = 11.82, s.e. = .62, t_{38} = 19.05, p < .001). Conversely, precise distracters under high time pressure (blue curve in the right panel) – a condition with small early noise and large late noise – led to the most negative context effect (linear regression: β = -3.72, s.e. = 1.26, t_{39} = -2.96, p < .005). The other two conditions showed intermediate slopes between the two extreme cases (see slopes contrasted with bars in **Figure 4e**). Under higher time pressure, the vague distracters led to a less positive effect than under low time pressure (the orange curve: β = 3.03, s.e. = .82, t_{38} = 3.69, p < .001; slope change from red to orange: β = -8.78, s.e. = 1.03, t_{76} = -8.52, p < .001). The same effect of time pressure arose for the precise distracters, where the blue curve under high time pressure showed a steeper negative slope than the cyan curve under low time pressure (cyan curve: β = .79, s.e. = 1.44, t_{37} = .55, p = .58; slope change from blue to cyan: β = -4.51, s.e. = 1.90, t_{76} = -2.37, p = .020). Comparing the two levels of vagueness given fixed time pressure showed consistent patterns: vague distracters drove a more positive context-modulation slope than the precise distracters (slope change from cyan to red: β = 11.03, s.e. = 1.53, t_{75} = 7.19, p < .001; slope change from blue to orange: β = 6.75, s.e. = 1.51, t_{77} = 4.49, p < .001). Overall, our empirical data closely aligned with the predictions of the new divisive normalization model that allowed for two sources of noise. Positive and negative contextual effects co-existed in our empirical observations as they did in the model. The vagueness of distracter values, associated with early noise, drove a positive context effect; in contrast, time pressure during the choice, associated with late noise, drove a negative context effect.

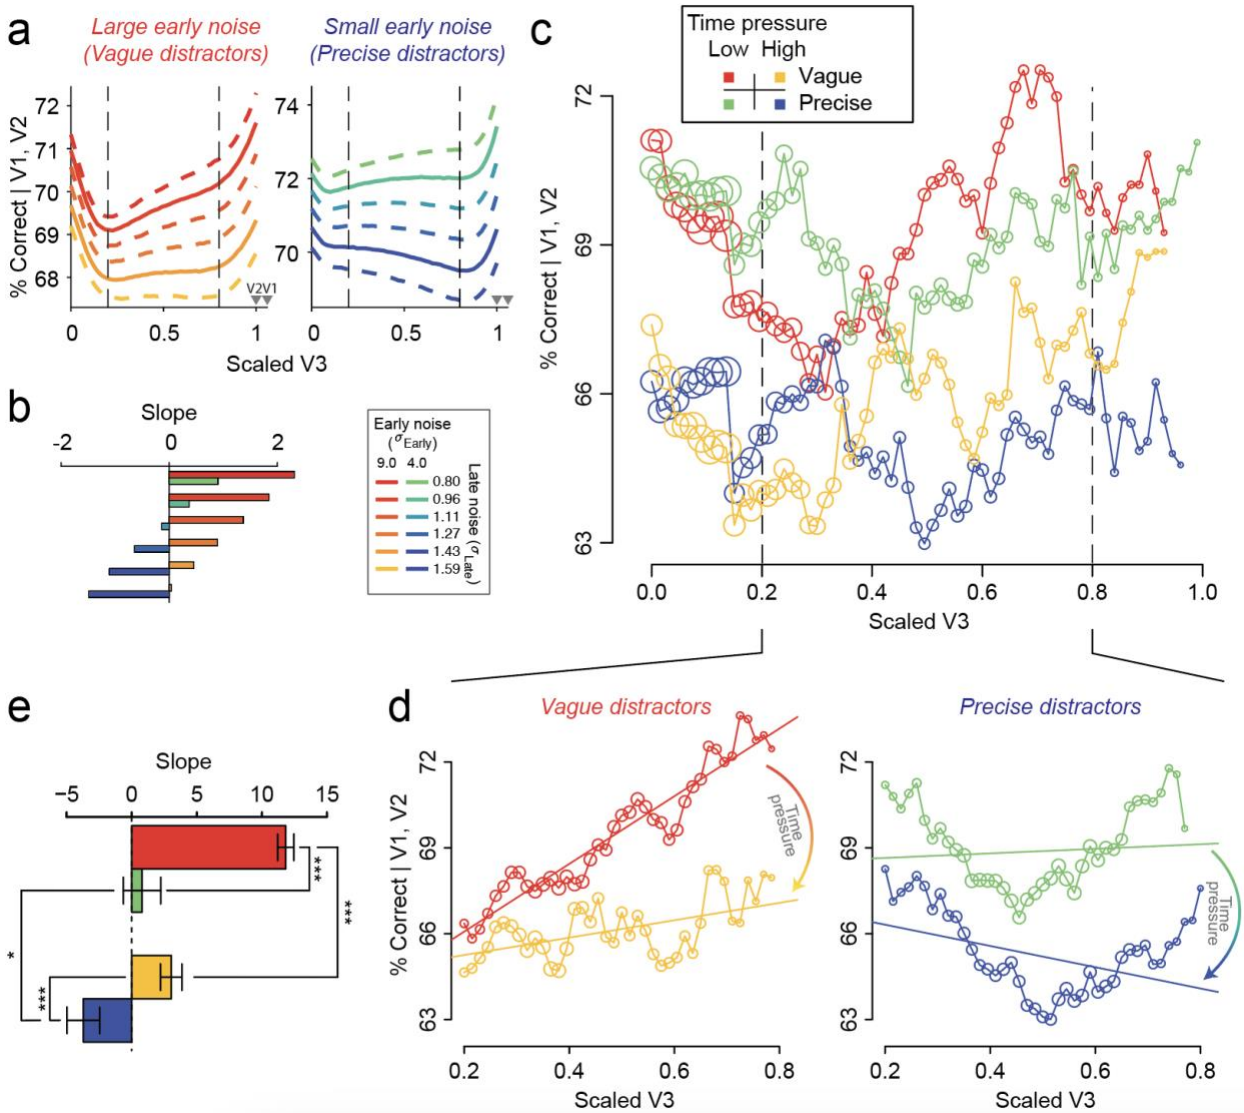

Figure 4 Human choice behavior shows opposite noise-dependent context effects consistent with model predictions. (**a**) Model predictions on conditional choice accuracy that matched the specific experimental design, with fixed pairs of targets and varying distracters with larger (left panel) or smaller (right panel) early noise. Smaller and larger late noise states were employed in the two solid lines and were extended to more magnitudes that progressively changed in the dashed lines (indicated by color) to reveal a general pattern driven by the late noise. (**b**) The quantified slopes of the context effects in the middle range of the distractor values (0.2 to 0.8) uncovered two main effects of the early and late noise. (**c**) In the participants' choices, the context effects over the range of V_3 that were scaled to the minimum target value were disclosed by implementing sliding windows over V_3 (window span = 0.3, step size = .03). Circle size indicates the number of trials in each sliding window. (**d**) Sliding windows on choice accuracy over the constrained range of scaled V_3 (0.2 to 0.8) were isolated from the full span. Lines indicate linear fit to the curves, weighted by the number of trials on each data point. (**e**) The fitted slopes to the curves were visualized as bars and compared (* $p < .05$; *** $p < .001$).

Discussion

Variability is a central feature of sensory information, neural coding, and behavior, but how cognitive noise and neural computation interact to control behavior is unknown. Here, we examined how different types of noise interact with divisive normalization-style models. We found that divisive normalization models that incorporate noise in their inputs can replicate the full range of observed context effects. With early noise arising in the inputs to the normalization process, a contextual option can improve the representation of pre-existing options, thus enhancing choice accuracy. In contrast, noise added to the output of the normalization process can react to a contextual option in a way that impairs the representation and thus reduces choice accuracy. These novel model predictions were validated in empirical human choice behavior, in a task designed to dissociate the effects of early and late noise. Our behavioral findings illuminate the role of noise on choice and suggest that an enhanced divisive normalization model can account for empirically observed effects of early and late noise.

In contrast to the assumptions of classical normative choice theories, the preferences of human and animal choosers change in response to context. Such violations of traditional rational choice models have been widely studied in the economics and psychology literature and are of increasing interest in neuroscience. However, the scientific dialogue on context-dependent effects has been recently marked by controversy, especially regarding how the presence of a contextual option influences the accuracy of decision-making. The classical divisive normalization model requires that a contextual option impair choice accuracy by diminishing the magnitude of the neural representation of the target options—a view substantiated by a breadth of studies across various species and cognitive tasks (Chang et al., 2019; D. Cohen et al., 2019; Holper et al., 2017; Louie et al., 2011, 2013; Soltani et al., 2012; Webb et al., 2014, 2020). However, recent findings have introduced an alternative narrative, with some evidence suggesting that contextual options can sometimes enhance or have no impact on decision-making (Chau et al., 2014, 2020; Gluth et al., 2020). This controversy points to a critical gap in our understanding of the underlying neural mechanisms of decision-making.

The significance of the current study lies in its enhancement of the representation of noise in divisive normalization class models as a tool for reconciling these two narratives. By differentiating the sources of noise in the model, our work offers a comprehensive framework that reconciles the conflicting evidence in the field. We demonstrate that the timing of noise—whether it occurs before or after the normalization process—can significantly alter the influence of contextual options. Our findings suggest that a contextual input can enhance decision accuracy under early noise, whereas impairs it when under late noise. This systematic understanding not only clarifies the discrepancies in the literature but also broadens our understanding of neural computations, emphasizing the temporal structure of noise in neural processing and decision-making. The current study, therefore, provides a coherent and physiologically

informed explanation of how decisions are shaped by context. While our results show that different patterns of observed context effects (Chau et al., 2014, 2020; Gluth et al., 2020; Kohl et al., 2023; Louie et al., 2013; Webb et al., 2020) can be driven by the balance of early and late noise, we note this does not rule out other coexisting mechanisms for varying context effects (Chau et al., 2020; Kohl et al., 2023).

Before addressing its neurobiological implications, it is important to recognize the limitations of the current experimental design. The current design does not allow us to measure early and late noise during choices directly. Hence, it is worth clarifying our assumptions regarding the underpinning neural computations. The early noise associated with single items was manipulated by controlling the uncertainty of information provided to subjects about choice objects. The impact of this manipulation was measured by assessing the variability of bidding behavior. This measured variability was incorporated into the normalization computation as distributions on inputs to the model. This aligns with a neurobiological mechanism recently uncovered in the study of value-based learning, where individual neurons encode a spectrum of values, shaping the overall signal distributions (Dabney et al., 2020). A second limitation which must be acknowledged is our assumption that time pressure solely affects the late noise arising from the output of the normalization process. While longer decision times have been suggested to reduce the uncertainty of value representation (Callaway et al., 2022; Jang et al., 2021), our empirical findings and model predictions did not provide marked evidence that early noise co-varies with late noise under the two time-pressure conditions.

From an evolutionary perspective, uncertainty in the representation of value poses significant challenges for humans and animals in making accurate decisions. The uncertainty of an option's value is an inevitable aspect of decision-making and is due to various factors like the stochasticity inherited from the environment (Drugowitsch et al., 2016), uncertain reward associations (Daw et al., 2005; Schlichting & Preston, 2016; Shohamy & Wagner, 2008; Wimmer & Shohamy, 2012), stochastic memory retrieval (Bakkour et al., 2019; Shadlen & Shohamy, 2016), imprecise inferential reasoning (Barron et al., 2020; Xue et al., 2024), and/or fluctuating motivational states (Priestley et al., 2024; Shen, Chen, et al., 2023). Our findings regarding early noise illustrate a special case where noise can be effectively managed by the brain circuit. Divisive normalization emerges as a crucial mechanism that enhances value representations by squeezing the uncertainty of input values with contextual information. In contrast to a more traditional perspective, which focuses on the detrimental effects of context under late noise, the early noise scenario reveals another biological benefit of divisive normalization in real-world environments. Together with other well-documented benefits of divisive normalization, such as range adaptation (Carandini & Heeger, 2012; Louie et al., 2015; Zimmermann et al., 2018) and redundancy elimination (Bucher & Brandenburger, 2022; Schwartz & Simoncelli, 2001), the current work highlights the biological value of divisive normalization under noise.

Another implication of the current work is the potential importance of noise correlation. To predict the positive contextual effect observed under early noise, a certain degree of

independence between the denominators of different options is required. Otherwise, integrating the same set of input values into the denominators would result in a uniform denominator across all options. No contextual effects would be predicted in this case since the rank of input values would always be preserved when scaled by the same constant. This implies a biological substrate where inhibitory neurons, a biological counterpart of the divisive gain-control from the denominator (Louie et al., 2014; Shen, Louie, et al., 2023; Zimmermann et al., 2018), are distinct pools for different options. This motif is supported by recent discoveries of inhibitory neurons with choice-selective properties (Najafi et al., 2020) and implemented in models like the local disinhibition decision model (LDDM) (Shen, Louie, et al., 2023). Similarly, in the late noise scenario, when the late noise covariance is fully correlated across options, the choice would lose its stochastic nature. Adding a constant value to all inputs would always preserve their rank during choice, inevitably resulting in 100% choice accuracy, which obviously contradicts empirical observations. Simultaneous recording of multiple units from the visual cortex shows that the trial-by-trial covariance across individual neurons decreases when animals are more engaged in a task and the performance in discriminating between stimuli is enhanced (M. R. Cohen & Maunsell, 2009). This implies that the noise structure imposed on the task-relevant signal is inherently decorrelated, aligning with the independent noise in our model.

In conclusion, our extension of the divisive normalization model to incorporate input stochasticity provides new insights into contextual effects in decision-making. This extended model, validated through our behavioral data, reconciles previously conflicting findings that challenged the traditional divisive normalization model. More broadly, this study underscores the importance of considering the temporal structure of noise in cognitive modeling, paving the way for more precise and predictive frameworks in neuroscience.

Methods

Participants

Sixty adult participants were recruited through an online platform in the New York area (https://newyork.craigslist.org/). Among these, five participants consistently bid zero for most items, not allowing effective construction of their choice set, thus excluded from the analysis. Fifty-five participants remained after exclusion (29 female, age = 37.9±12.5). The sample size was determined by power analysis, which showed that for a medium effect size (Cohen's D = .5), more than fifty-four subjects would be required to detect the difference of context-modulation slopes between conditions within-subject (paired t-test, two-tailed) at a power level of 95% and significance level of 0.05 (G*Power 3.1) (Faul et al., 2007). The study protocol was approved by the Institutional Review Board (IRB) of New York University Grossman School of Medicine. All participants were informed of the potential risk of the task and provided written consent. All participants were compensated with $90 or equivalent depending on their task choices.

Procedures

Overall task structure

Subjects performed a variant of the two-stage valuation and trinary choice task previously used to examine context-dependent preferences (Gluth et al., 2020; Louie et al., 2013). In this task, subjects first provide their subjective valuations of different consumer items in a bidding-based valuation task. Subsequently, they select their preferred item from trinary choice sets; these choice sets were constructed based on their valuations to quantify the degree of context-dependent choice behavior. Experimental details are provided below and in *Supplementary Materials*.

Early noise manipulation

Upon the participants' arrival, we introduced two levels of representational vagueness (precise or vague) of the consumer goods by controlling the information provided to the participants. The participants were provided with a selection of 36 goods exclusively obtained from a prominent online retail store with a wide range of market values (from $0.47 to $99.99) based in the New York area, United States (see the full list in *Supplementary Materials*). Half of the items were designated as "precise items" and presented as physical products arranged on a table. Participants were instructed to examine and interact with these items thoroughly for a minimum of 10 minutes and informed that the subsequent task would involve the appearance of a cartoon image corresponding to each observed item. The other half of the items were designated as "vague items", and were only presented as cartoon images in the task; participants were required to make their selections based on their best estimations derived from their daily

life experiences. We expected the participants to have a relatively more precise representation of the precise items than the vague items. The assignment of which half of the items were precise was counterbalanced across the participants. The market values and the categories of the goods between the two halves were carefully matched (*Supplementary Materials*).

Bid task to access subjective values and variability

Subsequently, the participants played an incentive-compatible Becker-Degroot-Marschak (BDM) auction for each good (Becker et al., 1964), which was designed to reveal the subjective value for each item as well as the variability of that value. Each participant was provided an initial endowment of $90. In each bidding trial, a cartoon image of an item was presented in the center of the screen. The participants were free to bid on the item continuously from $0 to $90 by moving their computer mouse (**Figure 3c**). They were informed that a random trial would be drawn at the end of the experiment. If they bid higher in that trial than a random number from $0 to $90 (uniformly distributed), they will get the good at that random price; otherwise, they will keep the full amount of endowment but have no chance to get the good. The goods were presented in randomized order, and each item was presented three times in total (108 trials). To quantify the variability of subjective valuations, we calculated the variance across the three bids for each item. In addition, at the end of each trial, participants were asked to rate "how certain you are about the value you bid" on a Likert scale from 0 to 10; certainty ratings showed a high correlation with the bidding variance but are not further analyzed here. Most of the participants completed the bid task in about 10 minutes.

Choice task and late noise manipulation

Following the previous tasks, the participants completed an incentive-compatible choice task to measure their choice accuracy between two target options under the influence of a distracter option. In each trial, three goods were presented in counterbalanced locations on the screen (**Figure 3e**). Among the three, two target options were selected from the six top-ranked precise items based on their bid mean values, resulting in fifteen different combinations of target pairs. A third option was selected from the lower-ranked goods, consisting of twelve precise items and another twelve vague items (**Figure 3d**). To manipulate late noise, two levels of time pressure (10 and 1.5 seconds) were implemented in the choice task, alternating in a blockwise manner (block order counterbalanced across participants). Participants were informed at the start of each block about the time limit; failure to make a choice within the time limit would lose the chance of receiving the item. The mean bid values of the distracters were matched between the two time-pressure conditions.

These conditions produced a 2-by-2 design with two levels of early noise (precise and vague items) and two levels of late noise (low and high time pressure), with six distracter items under each condition. Each distracter was combined with the same set of fifteen target pairs across randomized trials, resulting in a total of 90 trials for each

condition (360 trials in total). Participants were informed that the choice-task trials would be pooled with the bid-task trials at the end of the experiment, and a random trial from the mixed pool would be drawn for realization. If they received a choice-task trial, they would receive their selected item in that trial; otherwise, if they received a bid-task trial, they would participate in the BDM auction as per the specified rules outlined above.

Numerical simulations

Simulations were conducted in Matlab R2023a (The Mathworks Inc., 2023). To determine the choice accuracy for each set of option values, 4,096,000 trials were simulated. In each trial, the noise term for each option was randomly drawn from an independently and identically distributed (i.i.d.) *Gaussian* distribution and added either to (1) the option value before divisive normalization (early noise), (2) the product of normalization (late noise), or (3) both (mixed noise), depending on the tests specified in the main text. Choice accuracy between the two targets was computed as the ratio of the number of trials when V_1 appeared as the maximum value versus the total number of trials when V_1 or V_2 appeared as the maximum; trials when V_3 appeared as the maximum were excluded.

The non-negative biological constraint on each option will lead to different predictions from the situation when allowing the option value with noise to be negative (see **Figures 1d-f** and **2a**). To implement the constraint, we forced each option's value from both the intermediate stage of normalization with early noise and the final product added with late noise to be non-negative.

The target values in the simulations were set arbitrarily as $V_1 = 88$, $V_2 = 83$; and V_3 varied from 0 to V_1 given the specific setting of each test required. The maximum response value (R_{max}) was set arbitrarily at 75 Hz. Changing the scale of the inputs and the maximum response will not impact the results as long as their magnitudes are balanced with those of the early noise (ε) and late noise (η). Baseline normalization was assumed at a low level of $\sigma_H = 1$. The weight of the normalization was assumed at $w = 1$. Parameters that were specific to each figure are listed below.

- **Figures 1b,c**: Four examples of V_3 were set in panels valued at 0, 47.4, 71.1, and 79. In the early noise example, ε was drawn from $N(0, 2.8^2)$ for every option; in the late noise example, η was drawn from $N(0, 1^2)$ for every option.
- **Figures 1d-f**: V_3 was sampled from 0 to V_1 for 50 equally spaced values. In the early noise example, $\varepsilon \sim N(0, 9^2)$ for every option, and η was set as zero; in the late noise example, ε was set as zero, and $\eta \sim N(0, 4^2)$ for every option.
- **Figure 2a**: V_3 was sampled from 0 and V_1 for 50 equally spaced values. The standard deviation of ε increased from 0 to 9; meanwhile, the standard deviation of η decreased from 3.63 to 0. The amplitudes of early and late noise were balanced at a constant Euclidean norm, equally spaced in eight values.
- **Figures 2b,c**: A wider range of ε and η were tested. The standard deviation of ε ranged between 0 and 13, with 101 equally spaced values; the standard deviation of η

ranged between 0 and 5.24, with 100 equally spaced values. Other parameters kept the same as in Figure 2a.

- **Figure 4a:** V_3 was sampled between 0 and V_1 for 50 equally spaced values. For the small and large early noise, the standard deviations of ε were set as 4 and 9 for V_3, and fixed at 9 for V_1 and V_2. The standard deviation of η for all options varied from 0.8 to 1.59, with six values equally spaced within each panel.

Acknowledgments

The authors thank Dr. Ryan Webb and Miss Aliya Koishina for their constructive comments and feedback. This study is supported by NIH U19 NS107616 (to K.L.).

References

Averbeck, B. B., Latham, P. E., & Pouget, A. (2006). Neural correlations, population coding and computation. *Nature Reviews Neuroscience*, *7*(5), Article 5. https://doi.org/10.1038/nrn1888

Bakkour, A., Palombo, D. J., Zylberberg, A., Kang, Y. H., Reid, A., Verfaellie, M., Shadlen, M. N., & Shohamy, D. (2019). The hippocampus supports deliberation during value-based decisions. *eLife*, *8*, e46080. https://doi.org/10.7554/eLife.46080

Barron, H. C., Reeve, H. M., Koolschijn, R. S., Perestenko, P. V., Shpektor, A., Nili, H., Rothaermel, R., Campo-Urriza, N., O'Reilly, J. X., Bannerman, D. M., Behrens, T. E. J., & Dupret, D. (2020). Neuronal Computation Underlying Inferential Reasoning in Humans and Mice. *Cell*, *183*(1), 228-243.e21. https://doi.org/10.1016/j.cell.2020.08.035

Bateson, M., Healy, S. D., & Hurly, T. A. (2003). Context–dependent foraging decisions in rufous hummingbirds. *Proceedings of the Royal Society of London. Series B: Biological Sciences*, *270*(1521), 1271–1276. https://doi.org/10.1098/rspb.2003.2365

Beck, J. M., Latham, P. E., & Pouget, A. (2011). Marginalization in Neural Circuits with Divisive Normalization. *Journal of Neuroscience*, *31*(43), 15310–15319. https://doi.org/10.1523/JNEUROSCI.1706-11.2011

Becker, G. M., Degroot, M. H., & Marschak, J. (1964). Measuring utility by a single-response sequential method. *Behavioral Science*, *9*(3), 226–232. https://doi.org/10.1002/bs.3830090304

Bucher, S. F., & Brandenburger, A. M. (2022). Divisive normalization is an efficient code for multivariate Pareto-distributed environments. *Proceedings of the National Academy of Sciences*, *119*(40), e2120581119. https://doi.org/10.1073/pnas.2120581119

Callaway, F., van Opheusden, B., Gul, S., Das, P., Krueger, P. M., Griffiths, T. L., & Lieder, F. (2022). Rational use of cognitive resources in human planning. *Nature Human Behaviour*, *6*(8), Article 8. https://doi.org/10.1038/s41562-022-01332-8

Carandini, M., & Heeger, D. J. (2012). Normalization as a canonical neural computation. *Nature Reviews Neuroscience*, *13*(1), Article 1. https://doi.org/10.1038/nrn3136

Chang, L. W., Gershman, S. J., & Cikara, M. (2019). Comparing value coding models of context-dependence in social choice. *Journal of Experimental Social Psychology*, *85*, 103847. https://doi.org/10.1016/j.jesp.2019.103847

Chau, Kolling, N., Hunt, L. T., Walton, M. E., & Rushworth, M. F. S. (2014). A neural mechanism underlying failure of optimal choice with multiple alternatives. *Nature Neuroscience*, *17*(3), Article 3. https://doi.org/10.1038/nn.3649

Chau, Law, C.-K., Lopez-Persem, A., Klein-Flügge, M. C., & Rushworth, M. F. (2020). Consistent patterns of distractor effects during decision making. *eLife*, *9*, e53850. https://doi.org/10.7554/eLife.53850

Cisek, P., Puskas, G. A., & El-Murr, S. (2009). Decisions in Changing Conditions: The Urgency-Gating Model. *Journal of Neuroscience*, *29*(37), 11560–11571. https://doi.org/10.1523/JNEUROSCI.1844-09.2009

Coen-Cagli, R., & Solomon, S. S. (2019). Relating Divisive Normalization to Neuronal Response Variability. *Journal of Neuroscience*, *39*(37), 7344–7356. https://doi.org/10.1523/JNEUROSCI.0126-19.2019

Cohen, D., Teichman, G., Volovich, M., Zeevi, Y., Elbaum, L., Madar, A., Louie, K., Levy, D. J., & Rechavi, O. (2019). Bounded rationality in C. elegans is explained by circuit-specific normalization in chemosensory pathways. *Nature Communications*, *10*(1), Article 1. https://doi.org/10.1038/s41467-019-11715-7

Cohen, M. R., & Maunsell, J. H. R. (2009). Attention improves performance primarily by reducing interneuronal correlations. *Nature Neuroscience*, *12*(12), Article 12. https://doi.org/10.1038/nn.2439

Dabney, W., Kurth-Nelson, Z., Uchida, N., Starkweather, C. K., Hassabis, D., Munos, R., & Botvinick, M. (2020). A distributional code for value in dopamine-based reinforcement learning. *Nature*, *577*(7792), 671–675. https://doi.org/10.1038/s41586-019-1924-6

Daviet, R., & Webb, R. (2023). A test of attribute normalization via a double decoy effect. *Journal of Mathematical Psychology*, *113*, 102741. https://doi.org/10.1016/j.jmp.2022.102741

Daw, N. D., Niv, Y., & Dayan, P. (2005). Uncertainty-based competition between prefrontal and dorsolateral striatal systems for behavioral control. *Nature Neuroscience*, *8*(12), 1704–1711. https://doi.org/10.1038/nn1560

Drugowitsch, J., Wyart, V., Devauchelle, A.-D., & Koechlin, E. (2016). Computational Precision of Mental Inference as Critical Source of Human Choice Suboptimality. *Neuron*, *92*(6), 1398–1411. https://doi.org/10.1016/j.neuron.2016.11.005

Dumbalska, T., Li, V., Tsetsos, K., & Summerfield, C. (2020). A map of decoy influence in human multialternative choice. *Proceedings of the National Academy of Sciences*, *117*(40), 25169–25178. https://doi.org/10.1073/pnas.2005058117

Faul, F., Erdfelder, E., Lang, A.-G., & Buchner, A. (2007). G*Power 3: A flexible statistical power analysis program for the social, behavioral, and biomedical sciences. *Behavior Research Methods*, *39*(2), 175–191. https://doi.org/10.3758/BF03193146

Findling, C., Skvortsova, V., Dromnelle, R., Palminteri, S., & Wyart, V. (2019). Computational noise in reward-guided learning drives behavioral variability in volatile environments. *Nature Neuroscience*, *22*(12), Article 12. https://doi.org/10.1038/s41593-019-0518-9

Glimcher, P. W. (2022). Efficiently irrational: Deciphering the riddle of human choice. *Trends in Cognitive Sciences*. https://doi.org/10.1016/j.tics.2022.04.007

Glimcher, P. W., & Tymula, A. A. (2023). Expected subjective value theory (ESVT): A representation of decision under risk and certainty. *Journal of Economic Behavior & Organization*, *207*, 110–128. https://doi.org/10.1016/j.jebo.2022.12.013

Gluth, S., Kern, N., Kortmann, M., & Vitali, C. L. (2020). Value-based attention but not divisive normalization influences decisions with multiple alternatives. *Nature Human Behaviour*, 1–12. https://doi.org/10.1038/s41562-020-0822-0

Heeger, D. J. (1992). Normalization of cell responses in cat striate cortex. *Visual Neuroscience*, *9*(2), 181–197. https://doi.org/10.1017/S0952523800009640

Heeger, D. J., Simoncelli, E. P., & Movshon, J. A. (1996). Computational models of cortical visual processing. *Proceedings of the National Academy of Sciences*, *93*(2), 623–627. https://doi.org/10.1073/pnas.93.2.623

Holper, L., Van Brussel, L. D., Schmidt, L., Schulthess, S., Burke, C. J., Louie, K., Seifritz, E., & Tobler, P. N. (2017). Adaptive Value Normalization in the Prefrontal Cortex Is Reduced by Memory Load. *eNeuro*, *4*(2). https://doi.org/10.1523/ENEURO.0365-17.2017

Huber, J., Payne, J. W., & Puto, C. (1982). Adding Asymmetrically Dominated Alternatives: Violations of Regularity and the Similarity Hypothesis. *Journal of Consumer Research*, *9*(1), 90–98. https://doi.org/10.1086/208899

Jang, A. I., Sharma, R., & Drugowitsch, J. (2021). Optimal policy for attention-modulated decisions explains human fixation behavior. *eLife*, *10*, e63436. https://doi.org/10.7554/eLife.63436

Kohl, C., Wong, M. X., Wong, J. J., Rushworth, M. F., & Chau, B. K. (2023). Intraparietal stimulation disrupts negative distractor effects in human multi-alternative decision-making. *eLife*, *12*, e75007. https://doi.org/10.7554/eLife.75007

Louie, K., Glimcher, P. W., & Webb, R. (2015). Adaptive neural coding: From biological to behavioral decision-making. *Current Opinion in Behavioral Sciences*, *5*, 91–99. https://doi.org/10.1016/j.cobeha.2015.08.008

Louie, K., Grattan, L. E., & Glimcher, P. W. (2011). Reward Value-Based Gain Control: Divisive Normalization in Parietal Cortex. *Journal of Neuroscience*, *31*(29), 10627–10639. https://doi.org/10.1523/JNEUROSCI.1237-11.2011

Louie, K., Khaw, M. W., & Glimcher, P. W. (2013). Normalization is a general neural mechanism for context-dependent decision making. *Proceedings of the National Academy of Sciences*, *110*(15), 6139–6144. https://doi.org/10.1073/pnas.1217854110

Louie, K., LoFaro, T., Webb, R., & Glimcher, P. W. (2014). Dynamic Divisive Normalization Predicts Time-Varying Value Coding in Decision-Related Circuits. *Journal of Neuroscience*, *34*(48), 16046–16057. https://doi.org/10.1523/JNEUROSCI.2851-14.2014

Luce, R. D. (1959). *Individual Choice Behavior: A Theoretical Analysis*. Wiley.

Najafi, F., Elsayed, G. F., Cao, R., Pnevmatikakis, E., Latham, P. E., Cunningham, J. P., & Churchland, A. K. (2020). Excitatory and Inhibitory Subnetworks Are Equally Selective during Decision-Making and Emerge Simultaneously during Learning. *Neuron*, *105*(1), 165-179.e8. https://doi.org/10.1016/j.neuron.2019.09.045

Noonan, M. P., Chau, B. K. H., Rushworth, M. F. S., & Fellows, L. K. (2017). Contrasting Effects of Medial and Lateral Orbitofrontal Cortex Lesions on Credit Assignment and Decision-Making in Humans. *Journal of Neuroscience*, *37*(29), 7023–7035. https://doi.org/10.1523/JNEUROSCI.0692-17.2017

Otto, A. R., Devine, S., Schulz, E., Bornstein, A. M., & Louie, K. (2022). Context-dependent choice and evaluation in real-world consumer behavior. *Scientific Reports*, *12*(1), Article 1. https://doi.org/10.1038/s41598-022-22416-5

Pastor-Bernier, A., & Cisek, P. (2011). Neural Correlates of Biased Competition in Premotor Cortex. *Journal of Neuroscience*, *31*(19), 7083–7088. https://doi.org/10.1523/JNEUROSCI.5681-10.2011

Priestley, L., Chiew, M., Shahdloo, M., Mahmoodi, A., Cheng, X., Cleveland, R., Rushworth, M., & Khalighinejad, N. (2024). *DORSAL RAPHE NUCLEUS CONTROLS MOTIVATIONAL STATE TRANSITIONS IN MONKEYS* [Preprint]. Neuroscience. https://doi.org/10.1101/2024.02.13.580224

Schlichting, M. L., & Preston, A. R. (2016). Hippocampal–medial prefrontal circuit supports memory updating during learning and post-encoding rest. *Neurobiology of Learning and Memory*, *134*, 91–106. https://doi.org/10.1016/j.nlm.2015.11.005

Schwartz, O., & Simoncelli, E. P. (2001). Natural signal statistics and sensory gain control. *Nature Neuroscience*, *4*(8), 819–825. https://doi.org/10.1038/90526

Shadlen, M. N., & Shohamy, D. (2016). Decision Making and Sequential Sampling from Memory. *Neuron*, *90*(5), 927–939. https://doi.org/10.1016/j.neuron.2016.04.036

Shafir, S., Waite, T. A., & Smith, B. H. (2002). Context-dependent violations of rational choice in honeybees (Apis mellifera) and gray jays (Perisoreus canadensis). *Behavioral Ecology and Sociobiology*, *51*(2), 180–187. https://doi.org/10.1007/s00265-001-0420-8

Shen, B., Chen, Y., He, Z., Li, W., Yu, H., & Zhou, X. (2023). The competition dynamics of approach and avoidance motivations following interpersonal transgression. *Proceedings of the National Academy of Sciences*, *120*(40), e2302484120. https://doi.org/10.1073/pnas.2302484120

Shen, B., Louie, K., & Glimcher, P. (2023). Flexible control of representational dynamics in a disinhibition-based model of decision-making. *eLife*, *12*, e82426. https://doi.org/10.7554/eLife.82426

Shevlin, B. R. K., Smith, S. M., Hausfeld, J., & Krajbich, I. (2022). High-value decisions are fast and accurate, inconsistent with diminishing value sensitivity. *Proceedings of the National Academy of Sciences*, *119*(6), e2101508119. https://doi.org/10.1073/pnas.2101508119

Shohamy, D., & Daw, N. D. (2015). Integrating memories to guide decisions. *Current Opinion in Behavioral Sciences*, *5*, 85–90. https://doi.org/10.1016/j.cobeha.2015.08.010

Shohamy, D., & Wagner, A. D. (2008). Integrating Memories in the Human Brain: Hippocampal-Midbrain Encoding of Overlapping Events. *Neuron*, *60*(2), 378–389. https://doi.org/10.1016/j.neuron.2008.09.023

Soltani, A., De Martino, B., & Camerer, C. (2012). A Range-Normalization Model of Context-Dependent Choice: A New Model and Evidence. *PLoS Computational Biology*, *8*(7), e1002607. https://doi.org/10.1371/journal.pcbi.1002607

Steverson, K., Brandenburger, A., & Glimcher, P. (2019). Choice-theoretic foundations of the divisive normalization model. *Journal of Economic Behavior & Organization*, *164*, 148–165. https://doi.org/10.1016/j.jebo.2019.05.026

Trueblood, J. S. (2022). Theories of Context Effects in Multialternative, Multiattribute Choice. *Current Directions in Psychological Science*, *31*(5), 428–435. https://doi.org/10.1177/09637214221109587

Tsetsos, K., Moran, R., Moreland, J., Chater, N., Usher, M., & Summerfield, C. (2016). Economic irrationality is optimal during noisy decision making. *Proceedings of the National Academy of Sciences*, *113*(11), 3102–3107. https://doi.org/10.1073/pnas.1519157113

Tversky, A., & Simonson, I. (1993). Context-Dependent Preferences. *Management Science*, *39*(10), 1179–1189. https://doi.org/10.1287/mnsc.39.10.1179

Von Neumann, J., & Morgenstern, O. (1944). *Theory of games and economic behavior* (pp. xviii, 625). Princeton University Press.

Webb, R., Glimcher, P. W., & Louie, K. (2014). Rationalizing Context-Dependent Preferences: Divisive Normalization and Neurobiological Constraints on Choice. *SSRN Electronic Journal*. https://doi.org/10.2139/ssrn.2462895

Webb, R., Glimcher, P. W., & Louie, K. (2020). Divisive normalization does influence decisions with multiple alternatives. *Nature Human Behaviour*, 1–3. https://doi.org/10.1038/s41562-020-00941-5

Webb, R., Glimcher, P. W., & Louie, K. (2021). The Normalization of Consumer Valuations: Context-Dependent Preferences from Neurobiological Constraints. *Management Science*, *67*(1), 93–125. https://doi.org/10.1287/mnsc.2019.3536

Webb, R., Zhao, C., Osborne, M., Landry, P., & Camerer, C. (2022). *A Neuro-Autopilot Theory of Habit: Evidence from Canned Tuna* (SSRN Scholarly Paper 4130496). https://doi.org/10.2139/ssrn.4130496

Weiss, O., Bounds, H. A., Adesnik, H., & Coen-Cagli, R. (2023). Modeling the diverse effects of divisive normalization on noise correlations. *PLOS Computational Biology*, *19*(11), e1011667. https://doi.org/10.1371/journal.pcbi.1011667

Wimmer, G. E., & Shohamy, D. (2012). Preference by Association: How Memory Mechanisms in the Hippocampus Bias Decisions. *Science*, *338*(6104), 270–273. https://doi.org/10.1126/science.1223252

Xue, C., Markman, S. K., Chen, R., Kramer, L. E., & Cohen, M. R. (2024). *Task interference as a neuronal basis for the cost of cognitive flexibility* [Preprint]. Neuroscience. https://doi.org/10.1101/2024.03.04.583375

Zimmermann, J., Glimcher, P. W., & Louie, K. (2018). Multiple timescales of normalized value coding underlie adaptive choice behavior. *Nature Communications*, *9*(1), 3206. https://doi.org/10.1038/s41467-018-05507-8