Risky decisions are influenced by individual attributes as a function of risk preference

Douglas G. Lee[1], Marco D'Alessandro[2], Pierpaolo Iodice[3], Cinzia Calluso[4], Aldo Rustichini[5], and Giovanni Pezzulo[2]

[1] Tel Aviv University, School of Psychological Sciences, Tel Aviv, Israel
[2] Institute of Cognitive Sciences and Technologies, National Research Council, Rome, Italy
[3] Université de Rouen, Rouen, France
[4] Luiss University, Rome, Italy
[5] University of Minnesota, Hanson Hall, Minneapolis, USA

Contact info for correspondence:

Dr. Giovanni Pezzulo
ISTC-CNR
Via S. Martino della Battaglia, 44
00185 Roma RM
Italia
giovanni.pezzulo@istc.cnr.it

Abstract

It has long been assumed in economic theory that multi-attribute decisions involving several attributes or dimensions – such as probabilities and amounts of money to be earned during risky choices – are resolved by first combining the attributes of each option to form an overall expected value and then comparing the expected values of the alternative options, using a unique evidence accumulation process. A plausible alternative would be performing independent comparisons between the individual attributes and then integrating the results of the comparisons afterwards. Here, we devise a novel method to disambiguate between these types of models, by orthogonally manipulating the expected value of choice options and the relative salience of their attributes. Our results, using behavioral measures and drift-diffusion models, provide evidence in favor of the framework where information about individual attributes independently impact deliberation. This suggests that risky decisions are resolved by running in parallel multiple comparisons between the separate attributes – possibly alongside an additional comparison of expected value. This result stands in contrast with the assumption of standard economic theory that choices require a unique comparison of expected values and suggests that at the cognitive level, decision processes might be more distributed than commonly assumed. Beyond our planned analyses, we also discovered that attribute salience affects people of different risk preference type in the opposite manner: risk-averse participants seem to focus more on probability, except when monetary amount is particularly high; risk-neutral participants, in contrast, seem to focus more on monetary amount, except when probability is particularly low.

Introduction

Most of our everyday decisions are between options defined by multiple attributes or dimensions, such as: foods that provide different levels of pleasure and nutrition; houses that vary in terms of size and location; or lotteries that offer different amounts of money available to win and probabilities of winning. The classical economic perspective is that such multi-attribute choices should be solved by first computing a unique overall *expected value* (EV) for each option, which integrates all value-relevant attributes, and then comparing the EVs of the different options. This comparison could perhaps be achieved through a process of evidence accumulation to a threshold, as in the *drift-diffusion model* (DDM; (Ratcliff & McKoon, 2008)) or similar mathematical accounts of economic choices (Rustichini & Padoa-Schioppa, 2015; Usher & McClelland, 2001).

Despite its normative character and its centrality in neuroeconomic thinking, the idea that choices require a mandatory integration of attribute values to form estimates of EV might be less compelling than it appears. This is because, from a computational perspective, a system that makes decisions by comparing EVs requires delaying the choice until all relevant attributes are integrated, which could be especially challenging when they are numerous. Indeed, previous studies have shown (using large-scale datasets) that EV alone is not the best model for choice (He et al., 2022; Peterson et al., 2021). Rather than deciding via a unique comparison of EVs (*integrate-then-compare* model), a more feasible alternative might be running in parallel multiple comparisons – each between the values along one specific attribute dimension – and then integrating the partial results of the competitions afterwards (*compare-then-integrate* model; see (Hunt et al., 2014; Roe et al., 2001; Turner et al., 2018; Usher & McClelland, 2004)). Furthermore, the two models described so far (*integrate-then-compare* and *compare-then-integrate*) are not mutually exclusive

– they could operate simultaneously. Therefore, it is possible to envisage a third model that combines the two models (*combined-comparison*; see (Stewart, 2011)), which would perform multiple comparisons – between the values of individual attributes *and* between their integrated EVs – and then combine the outcomes of all of the comparisons to determine the final choice.

The two latter models (*compare-then-integrate* and *combined-comparison*, in which attributes can compete separately) have at least two advantages: they can initiate decisions faster than the *integrate-then-compare* model, and they can operate (at least partially) in parallel. The parallelism is important to the extent that one is also interested in the brain circuits that support multi-attribute choices. Here, again, there is a dispute between classical neuroeconomic models in which decision computations are centralized in prefrontal brain regions (Padoa-Schioppa, 2011), and distributed choice models that recognize that decision computations arise as a distributed consensus that emerges across multiple brain circuits considering separate attributes and subsequently integrating them (Cisek, 2012; Pezzulo & Cisek, 2016). At the same time, the *compare-then-integrate* and *combined-comparison* models that do not only integrate attribute values into an EV would show some deviations from the purported "optimality" criteria of the classical economic model (*integrate-then-compare*) – some of which have been observed experimentally. For example, decisions in which attributes can compete separately would be particularly influenced by the attributes that are considered first or whose comparison is faster – a mechanism that is often exploited in marketing (Amasino et al., 2019; Lim et al., 2018; Maier et al., 2020; Sullivan et al., 2015). Furthermore, decision processes might be affected by options that have low EV but a single appealing attribute (Tversky & Simonson, 1993). Finally, the values of different attributes could interact during the decision, rather than act independently as assumed by classical models (Park et al., 2011).

When empirically testing whether, during multi-attribute decisions, attributes are considered individually (*compare-then-integrate model*) or integrated to form overall expected values before entering the competition (*integrate-then-compare* model), a recurrent problem is that it is difficult to disentangle the above models experimentally, as they tend to generate very similar predictions across a range of experimental conditions (Stewart, 2011). In other words, most standard paradigms used in economics and neuroeconomics are not sensitive enough to distinguish between the above models. However, some studies have attempted to do just that, with some success. For example, Noguchi and Stewart (Noguchi & Stewart, 2014) used eye-tracking to show that participants made a series of comparisons within each choice trial, with each comparison taking place within a specific attribute dimension. The authors concluded that psychological models of choice should thus be based on within-attribute comparisons (Noguchi & Stewart, 2014). Other eye-tracking and computational modeling studies investigating risky choice have found that participants seem to perform within-alternative integrations rather than within-attribute comparisons (Fiedler & Glöckner, 2012; Glickman et al., 2019). The topic has also been explored in the domain of intertemporal choice, using mouse-tracking to show that some participants seem to integrate within alternatives whereas others seem to compare individual attributes across alternatives (Reeck et al., 2017). The question about within-alternative versus within-attribute comparisons has been asked not only with respect to human decision-makers, but also rhesus macaques (Farashahi et al., 2019). A few studies have already used drift-diffusion modeling to provide evidence that participants compare individual attributes during deliberation about intertemporal choices (Dai & Busemeyer, 2014) or that they seem to both compare attributes across options and integrate values within options during value-based decisions (Lee & Hare, 2022; Yang & Krajbich, 2022).

To help resolve the inconsistencies of the above studies and to further adjudicate amongst *integrate-then-compare*, *compare-then-integrate*, and *combined-comparison* models of decision-making, we designed a "risky" choice task that parametrically varied not only the expected values of different options (as in previous studies), but also the relative *salience* of the attributes that enter the competition. Our concept of salience aligns with that set forth in Salience Theory (Bordalo et al., 2012), where an option is deemed more salient when it contains a measure of an attribute that is much higher than the attributes of the other option. We reasoned that rendering one attribute more salient than the other on each trial could be particularly diagnostic, as it might make the (putative) within-attribute competition faster, influencing the overall pattern of choice behavior. A similar effect has been demonstrated in a risky choice task featuring numerical streams of potential payoff amounts, where the authors deemed that larger numerical values were more salient to the participants (Tsetsos et al., 2012). The assumption that higher numerical values will be more salient is also supported by previous experimental and computational work based on value-based attentional capture, which assumes that more attention will be focused on options with higher expected values, and that the additional attention will impact choice behavior (Gluth et al., 2018, 2020). Assessing the relative importance of EV versus individual attributes together with the salience of attributes during choices could permit adjudicating between models that postulate integrate-then-compare, compare-then-integrate, or combined processes.

In our task, participants used a computer mouse to make a series of risky choices between two options (lotteries) characterized by two dimensions: the amount of money to be won and the probability of winning it. An example choice is between "90% probability of winning €20" and "30% probability of winning €60". From this point onward, consistent with the neuroeconomics literature, we label the option having the highest probability (e.g., "90% of probability of winning

€20”) as *safer* and the option having the lowest probability (“e.g., 30% probability of winning €60”) as *riskier*.

Importantly, in our experiment, the choice option pairs varied across two dimensions: *expected value difference* (dV) and *salience* (SAL). The two categories of SAL reflect the fact that for each trial, either the probability or the amount of money is particularly salient, as it has a high value that is at least 1.5 times greater than the second-highest value on display. (Note that this notion of salience is not related to bottom-up characteristics of the stimuli, such as size, color, or clarity.) We reasoned that if the attributes are processed independently, the most salient attribute might be processed differently (e.g., faster) than the less salient attribute. In contrast, if the attributes are not processed independently, the salience of individual attributes should have no impact on choice dynamics. Specifically, we assume that SAL is for the safer option if the salient number corresponds to its probability and that SAL is for the riskier option if the salient number corresponds to its amount of money. For example, the choice between “90% probability of winning €20” and “30% probability of winning €60” would have dV = equal (as both options have the same EV) and SAL = safer (as the most salient number is 90, which corresponds to the probability attribute of the safer option).

By manipulating the SAL of attributes independent of the EV of options, this experimental design generates a diverse range of trials, for which the *integrate-then-compare*, *compare-then-integrate*, and *combined-comparison* models would make different predictions. Broadly speaking, under the *integrate-then-compare* model, response probabilities, response times, and mouse trajectories should not be sensitive to the SAL of individual attributes, because the attributes would not be processed independently (before being integrated into EV). Conversely, under the *compare-then-integrate* and *combined-comparison* models, both dV and SAL should exert an influence on the decision process and the associated choice variables. However, the predictions made by the above models could be more subtle than what we described. For that reason, we base our analysis on a quantitative model comparison of a set of computational (drift-diffusion) models that formalize the alternative *integrate-then-compare*, *compare-then-integrate*, and *combined-comparison approaches.*

Finally, our experimental design allows us to examine potential differences in participant behavior according to their risk preference type (i.e., how well they tolerate risk when choosing between options with different outcome probabilities). It has long been shows that there are individual differences in risk preference (Tversky & Kahneman, 1992), with some people avoiding it (risk-averse), some indifferent to it (risk-neutral), and some drawn to it (risk-seeking). Risk preferences can be deduced by examining how people's choice probabilities deviate from classical economic theory. For instance, in the above example ("90% probability of winning €20" versus "30% probability of winning €60"), classical economic theory would predict indifference, because the expected values of the options are identical (€18). A risk-neutral person would be indifferent and thus adhere to classical economic theory. A risk-averse person would prefer the first option, because it has the higher probability (i.e., lesser risk). A risk-seeking person would prefer the second option, because it has the higher monetary amount (and thus the lower probability or greater risk). One might hypothesize, therefore, that different attributes are inherently more or less important than others for different types of people (probability > monetary amount for risk-averse people, probability < monetary amount for risk-seeking people). In terms of our salience manipulation, we might then expect a different effect to emerge for the different groups. Specifically, on trials that we classify as SAL = riskier (where the monetary amount of the riskier option is the greatest number on the screen), risk-averse people might actually be repulsed by the riskier option (because it will also have a very low probability, which is more important to them than the large monetary amount). Something analogous might occur for risk-seeking people on trials where SAL = safer. We thus tested for group differences in the behavioral data, and we also tested whether the winning computational model might differ across groups. Given that risk preference is a topic that has long been studied, it could be of interest if our modeling analysis illuminated differences in parameters or even models across risk preference groups.

To preview our results, we found that both expected value difference and salience affect the final choices as well as the choice dynamics, as indexed by response time and the curvature of the response trajectory. Furthermore, the computational model that best accounts for the data (when considering all participants together) is the one that incorporates the *compare-then-integrate* framework in addition to the *integrate-then-compare* framework, along with differential evidence accumulation rates according to salience condition. Grouping participants by risk preference reveals that both groups consider individual attributes in their choices, but they do so in addition to different overall strategies (*compare-then-integrate* for risk-averse and *integrate-then-compare* for risk-neutral).

Methods

Ethics approval

The experimental procedure was approved by the Ethics Board of the Institute of Cognitive Sciences and Technologies, National Research Council of Italy.

Participants

A sample of 43 healthy adults, all males, aged 20 to 49 years old (mean age 26.7, standard deviation 8.3), took part in the study. All participants were right-handed, with normal or corrected-to-normal vision. Due to technical issues, the data for three participants was not saved. All analyses reported below are thus for n = 40 participants.

Stimuli and procedure

The stimuli used in this study were numerical displays of two options (one towards the left and one towards the right of the computer screen). Each option represented a “risky” choice, in the sense that the chosen monetary amount of would be obtained with probability less than 100%. Each trial included a safer option (with the higher probability) and a riskier option (with the lower probability). Trials included different levels of expected value difference (dV), specifically: on each trial, the expected value (EV) of the safer option was 50% greater than, 25% greater than, equal to, 25% less than, or 50% less than the EV of the riskier option. Trials also included different categories of SAL, specifically: on each trial, either the probability or the amount of money had a value that was at least 1.5 times greater than the second-highest value on display for that trial. Each participant made 240 choices in total, as each stimulus pair in the set of 60 pairs was presented four times. The stimuli were presented in a random order, with the position of the stimuli / response buttons counterbalanced across trials.

Participants reported their choices by clicking with a computer mouse one of two response buttons, located at the top-right or top-left of their screens (Barca & Pezzulo, 2012, 2015; Calluso et al., 2015; Freeman & Ambady, 2010; Lepora & Pezzulo, 2015). To begin each trial, participants clicked on the START button located at the bottom-center of the screen (Figure 1). Then, the two response buttons appeared (one at the upper-left and one at the upper-right of the screen) and remained on the screen until a response was entered.

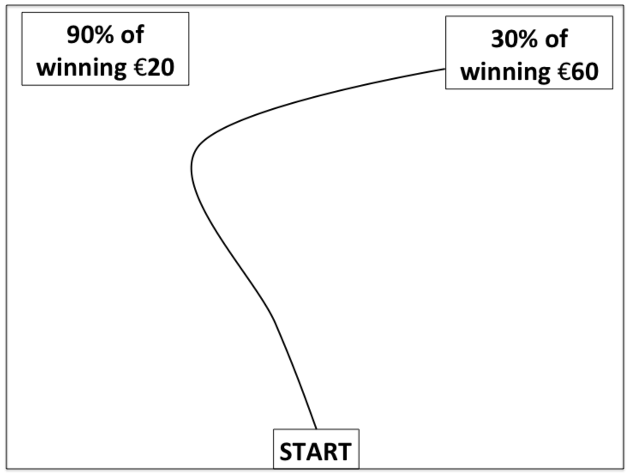

Figure 1. Experimental setup. *Each trial started when participants clicked the START button located at the bottom-center of the screen. At that point, two stimuli (a "safer" option and a "riskier" option) appeared at the top-right and top-left of the screen. Each stimulus pair appeared four times across the experiment, with the position of the safer and riskier options counterbalanced. Participants reported their choices by clicking one of the two response buttons. The figure shows a (fictive) mouse trajectory.*

Behavioral and kinematic measures

Choice and response time (RT, from when participants pressed START until they reached and pressed one of the response buttons) were recorded for each trial, as were the x and y coordinates of the mouse trajectories (with a sampling rate of approximately 70Hz). For all analyses reported below, we converted RT to seconds and ignored outlier trials (defined as having log(RT) greater than the median +/- 3x the median average deviation, within participant). With respect to mouse movement, we focused on the *maximum deviation* (MD) of the trajectories, which indexes choice uncertainty and the competition between response alternatives (Hehman et al., 2015; Spivey et al., 2005). The MD is the maximum shortest distance between each point on the observed mouse trajectory and an ideal straight line connecting the start button to the response button. The closer the mouse trajectory is to this ideal line, the smaller the MD will be, indicating a higher confidence in the choice. For each of the regression analyses that we report below, we used mixed effects models with random slopes and intercepts for each participant in addition to the regressors of interest.

Drift-diffusion model-based analysis

To further examine the issue of whether the decision process that guides risky choices is driven by calculations of expected value (i.e., the monetary amount on offer multiplied by the probability of receiving the payoff, if that option is chosen) or independently by the magnitudes of the individual attributes (monetary amount and probability), and to adjudicate between *integrate-then-compare, compare-then-integrate*, and *combined-comparison* models, we consider several variants of the drift-diffusion model (DDM) that formalize the alternative hypotheses (Ratcliff, 1978; Ratcliff & McKoon, 2008; Ratcliff & Rouder, 1998). The DDM is an evidence accumulation-to-bound model, with evidence referring to a moment-by-moment estimate of the relative value of one option compared to the other. These momentary estimates are noisy (e.g., derived from stochastic neural firing patterns), so the estimates are repeated across time and accumulated to cancel out the noise. The process terminates when a sufficient amount of evidence has been accrued, at which point a choice is made.

We followed the approach used by Lee and Hare (Lee & Hare, 2022) and considered three main models. Model 1 (*expected value DDM* or *evDDM*) is based on comparisons of expected value (EV). Model 2 (*multi-attribute DDM* or *maDDM*) is based on independent comparisons of the monetary amount (X) and probability (P) attributes. Model 3 (*multi-attribute DDM plus expected value* or *ev+maDDM*) is based on comparisons of both individual attributes X and P *and* EV. We also included variants of Models 1-3 in which the rates of evidence accumulation can differ as a function of salience condition (Model 4: *multi-attribute DDM plus salience* or *maDDM+sal*, Model 5: *expected value DDM plus salience* or *evDDM+sal*, Model 6: *expected value DDM plus salience* or *ev+maDDM+sal*). These latter models reflect a key assumption of our design, that the salience of attributes influences evidence accumulation. The mathematical details of each model are provided below.

Model 1: expected value DDM (evDDM)

The first model is a basic DDM in which only the statistical expected values (EV = X * P) of the two options influence the evidence accumulation on each trial. Specifically, at each time step, the incremental evidence equals the EV of the riskier option (EV_R) minus the EV of safer option (EV_S), scaled by an efficiency parameter (the drift rate, *d*), plus Gaussian noise (ε) with mean 0 and variance σ^2. The cumulative evidence (*e*) evolves across deliberation time as follows:

$$DV = EV_R - EV_S$$

$$drift = d^V DV$$

$$e_t = e_{t-1} + drift + \varepsilon$$

$$\varepsilon \sim N(0, \sigma^2)$$

$$e_0 = 0$$

where EV_R and EV_S are independent variables, and d^V and σ^2 are free parameters to be estimated to capture the individual-specific rate of evidence accumulation (drift) and level of noise in the accumulation process (diffusion), respectively. Evidence accumulation terminates when *e* reaches a response boundary ∈{θ, -θ}, with the sign of the final value of *e* determining the chosen option (arbitrarily defined as positive for the riskier option, negative for the safer option) and the final value of *t* determining the response time (RT).

Choice probability (*p*, choice of the riskier option) and mean RT can be analytically derived (Alós-Ferrer, 2018) as a function of EV_R, EV_S, and σ^2, with θ being fixed (here, to θ = 1 for simplicity):

$$p \stackrel{\text{def}}{=} p(ch = riskier) = \frac{1}{1 + e^{\frac{-2*drift}{\sigma^2}}}$$

$$E[RT] = \frac{2p - 1}{drift} + NDT$$

As is common practice in drift-diffusion modeling, we included an addition free parameter *NDT* to estimate the so-called non-decision time (or alternatively, time for stimulus encoding and response implementation). These analytical formulas for choice probability and mean RT will be identical across all models, except that the drift component will be calculated differently.

Model 2: multi-attribute DDM (maDDM)

The second model is similar to Model 1, except that now evidence accumulation is independently driven by the individual attributes of the options (X and P). The process is otherwise identical, and evidence accumulates as follows:

$$DX = X_R - X_S$$

$$DP = P_R - P_S$$

$$drift = d^X DX + d^P DP$$

where X_R, X_S, P_R, and P_S are independent variables, d^X and d^P are free parameters to capture the individual-specific rates of evidence accumulation for X and for P, and σ^2 captures the noise in the overall evidence accumulation process.

Model 3: expected value plus multi-attribute DDM (ev+maDDM)

The third model is essentially a hybrid of the first two models, in which evidence accumulation is independently driven by individual attributes in addition to expected value. The idea here is that choices could be made according to expected value in a general sense, but that individual-specific adjustments to the relative importance of each attribute could fine-tune the decision process. The computational structure is otherwise identical to that of Models 1 and 2, and evidence accumulates as follows:

$$drift = d^V DV + d^X DX + d^P DP$$

In this study we are only concerned with comparing different models in terms of their ability to best fit the experimental data. Note that if the goal were to make inferences about the relative impact of EV, X, and P in determining the drift rate, it would have been necessary to employ orthogonalization or dimensionality reduction techniques to control for the correlations between EV and both X and P.

Model 4: expected value DDM plus salience (evDDM+sal)

The fourth model is similar to Model 1, except that now evidence accumulation rates can differ according to which option is affiliated with the most salient attribute on each trial. The idea here is that whichever option has an attribute that is salient on a given trial will capture more attention and thus adjust the rate at which information is processed about that option relative to the other option. The computational structure is otherwise identical to that of Models 1-3, and evidence accumulates as follows:

$$drift = (d^V + d_r^V * risk) * EV_R - (d^V + d_s^V * safe) * EV_S$$

where d_r^V and d_s^V are free parameters to capture the individual-specific rates of incremental evidence accumulation for the riskier option when SAL = riskier and for the safer option when SAL = safer, respectively, *risk* = 1 if salience favors the riskier option (and 0 otherwise), and *safe* = 1 if salience favors the safer option (and 0 otherwise).

Model 5: multi-attribute DDM plus salience (maDDM+sal)

The fifth model is similar to Model 2, except that now evidence accumulation rates can differ according to which attribute is salient on each trial. The idea here is that whichever attribute is salient on a given trial will capture more attention and thus adjust the rate at which information is processed about that attribute relative to when it is not salient. The computational structure is otherwise identical to that of Models 1-4, and evidence accumulates as follows:

$$drift = (d_r^X DX + d_r^P DP) * risk + (d_s^X DX + d_s^P DP) * safe$$

where d_r^X and d_r^P are free parameters to capture the individual-specific rates of evidence accumulation for X and for P when SAL = riskier, d_s^X and d_s^P are free parameters to capture the individual-specific rates of evidence accumulation for X and for P when SAL = safer, *risk* = 1 if salience favors the riskier option (and 0 otherwise), and *safe* = 1 if salience favors the safer option (and 0 otherwise).

Model 6: expected value plus multi-attribute DDM plus salience (ev+maDDM+sal)

The six model is essentially a hybrid of Modes 4 and 5. The computational structure is otherwise identical to that of Models 1-5, and evidence accumulates as follows:

$$drift = (d^V + d_r^V * risk) * EV_R - (d^V + d_s^V * safe) * EV_S + (d_r^X DX + d_r^P DP) * risk + (d_s^X DX + d_s^P DP) * safe$$

We note that our model formulations do not contain a free parameter to represent a *starting point bias* in the evidence accumulation process. This parameter is popular in drift-diffusion modeling studies, as it allows for the possibility that one of the responses may be a sort of default option for a participant before the specific options are even presented. Previous studies on risky choice have found a starting point bias related to loss aversion (Clay et al., 2017; Sheng et al., 2020; Zhao et al., 2020). Although we do not explore loss aversion in this work, we thought it might be possible that the starting point bias parameter could serve to capture the risk-aversion tendencies that we found in our data. We therefore performed quantitative comparisons of our set of models with this additional parameter versus without, using formulations provided in previous studies (Grasman et al., 2009; Lopez-Persem et al., 2016). The models performed better without the starting-point bias parameter. We therefore decided to simplify our analysis by excluding this parameter (or equivalently, fixing it to zero). We also did this (for the same reason) for an optional parameter that allows for a fixed component of the drift rate (i.e., a drift component independent of the trial-specific option values).

Using the VBA toolbox in Matlab (see below) to fit and compare models allows us to simultaneously fit other choice variables in addition to those included in the DDM (choice probability and RT). VBA simultaneously optimizes parameters over a set of analytical equations, such as those listed above for choice probability and mean RT. Importantly, we can just as easily perform our model comparison while fitting additional equations. To that end, we fit generalized predictions about the relationship between the evidence accumulation (drift) rate and maximum mouse trajectory deviation (MD) in addition to choice probability and mean RT. We estimated MD as a linear function of either expected value difference (ev), individual attribute differences (ma), both (ev+ma), or expected value differences and individual attribute differences dependent on salience condition, or both (ev+sal, ma+sal, ev+ma+sal, respectively). The formula that we used to estimate MD under each model was:

$$MD = \beta_0 + \beta_1 * drift$$

Model fitting procedure

We fit each of the candidate models to the experimental data, then performed Bayesian model comparison to determine which of the models (if any) performed significantly better than the others across the population of participants. For this model fitting and comparison exercise, we relied on the Variational Bayesian Analysis toolbox (VBA, available freely at https://mbb-team.github.io/VBA-toolbox/; (Daunizeau et al., 2014)) with Matlab R2022b. Within participant and across trials, we entered the experimental variables {monetary amount, probability, and expected value for each option; observed RT; salience condition} as input and {choice = 1 for the riskier option, 0 for the safer option; RT; MD} as output. All monetary amounts were rescaled to the range (0,1] in accord with the probability measures. We also entered the model-specific mappings from input to output as outlined in the analytical formulas above. As we fixed the threshold parameter θ to 1, the parameters to be fitted were the drift scalar (*d*), diffusion noise (σ^2), and non-decision time (NDT) terms described above in the model formulations, plus the slope and intercept parameters for MD. VBA requires prior estimates for the free parameters, for which we set the mean and variance equal to: 1.6 and 37 for the positively constrained parameters (noise and all drift scalars); 0.5 and 3.7 for NDT (constrained between 0 and 1); and 0 and 16 for the unconstrained MD parameters. The theoretical drift rate, noise, and NDT parameters are always positive; we thus constrained the search space of our model fitting algorithm to the positive domain by replacing these parameters with the following calculation: log(1 + exp(parameter)) * 2.3. VBA then recovers an approximation to both the posterior density on unknown variables and the model evidence (which is used for model comparison). We used the VBA_NLStateSpaceModel function to fit the data for each participant individually, followed by the VBA_groupBMC function to compare the results of the model fitting across models for the full group of participants.

VBA estimates parameters during model fitting using Variational Bayes: an efficient iterative algorithm that provides a free-energy approximation for the model evidence, which trades off model accuracy (goodness of fit, or log likelihood) and complexity (degrees of freedom, or KL divergence between priors and fitted parameter estimates; see (Friston et al., 2007; Penny, 2012)). This is a critical step for comparing our models, as they differ in number of parameters. The VBA algorithm starts with our relatively flat Gaussian priors for each model's free parameters and eventually provides a posterior density estimate. The log model evidence scores calculated for each participant are then fed into the group-level random-effect Bayesian model selection (BMS) procedure. A key output of the BMS is the exceedance probability, which informs about how likely it is that a given model is more frequently implemented across the population of participants (relative to all other models under consideration; (Rigoux et al., 2014; Stephan et al., 2009)). Previous studies have successfully used this approach to fitting and comparing variants of DDM (Feltgen & Daunizeau, 2021; Lee & Hare, 2022; Lee & Usher, 2021; Lopez-Persem et al., 2016).

The VBA toolbox also allows models to be formally tested across what is known as "families" based on fundamental similarities across all models within each family. We established two types of family in our set of models. The first model family category represented our primary research question: is evidence integrated then compared, compared then integrated, or some combination of both. These families were thus: ev = {Model 1, Model 4}; ma = {Model 2, Model 5}; ev+ma = {Model 3, Model 6}. The second model family category represented our secondary research question: does the salience of individual attributes impact evidence accumulation. These families were thus: DDM = {Models 1-3}; DDM+sal {Models 4-6}. We compared the aggregated model performance across families using the *options.families* input to the VBA_groupBMC function.

Model recoverability

To verify that our model-fitting procedure is suitable for this specific analysis, we performed a test of model recoverability. Specifically, we took as the model input the same monetary amounts and probabilities from the choice trials that the participants were faced with. We then simulated the set of choice probabilities, mean RTs, and MDs for each participant, separately according to each of our models, using the actual participant-specific fitted parameters for each model. Finally, we fit all simulated data (per participant) to each of our models and performed the same formal model comparison as with our real experimental data. The results of this procedure can be seen in Figure 2. This matrix shows, for each true generative model, the percentage of participants (simulated under that model) that were attributed to each of the best fit models by our model-fitting procedure. As shown in the matrix, model confusion was low and the procedure attributed the true generating model as the best-fitting model for the vast majority of participants (recovery accuracy: 82% for *evDDM*, 80% for *maDDM*, 83% for *ev+maDDM*, 99% for *evDDM+sal*, 90% for *maDDM+sal*, and 86% for *ev+maDDM+sal*). There was a small amount of confusion between the *ev+maDDM* and the *evDDM* and *maDDM*, which is to be expected since the *ev+maDDM* is essentially a combination of those other models.

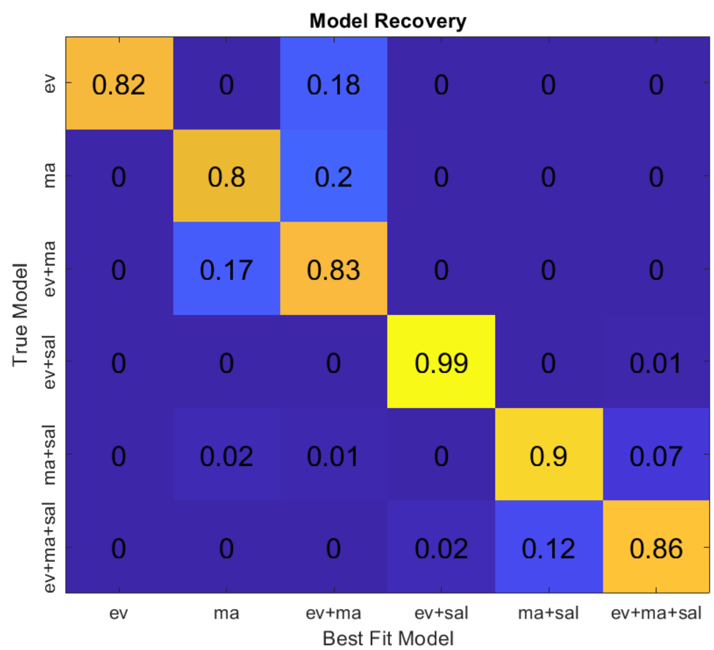

Figure 2. Model recovery analysis. *The cells in each "confusion matrix" summarize the percentage of participants (simulated under each true model) for which our model-fitting procedure attributed each of the (best fit) models.*

Results

Behavioral measures

We first examined the three decision variables of interest (choice, response time or RT, and maximum deviation or MD) in a model-free manner by observing their relationships with the relative attractiveness of the risker option compared to the safer option. Specifically, we considered the difference in the expected value (EV = monetary amount * probability) of the riskier option minus that of the safer option. The closer this difference is to zero, the more difficult the choice is presumed to be. Pooling together all trials across all participants, we separated the data into bins of equal widths covering the full range of EV difference. Within each bin, we calculated the percentage of trials where the riskier option was chosen, the mean RT, and the mean MD. Our data seem to conform to standard psychometric findings, where choices are close to chance level (50%), RTs are longest, and MDs are largest when the decision variable is closest to zero. However, we found a clear skew in our data, where this equivalence point was shifted away from zero toward the right of the EV difference scale (Figure 3A-C). This indicates that our participants (on average) were averse to risk, as has often been reported in studies of risky choice (O'Donoghue & Somerville, 2018).

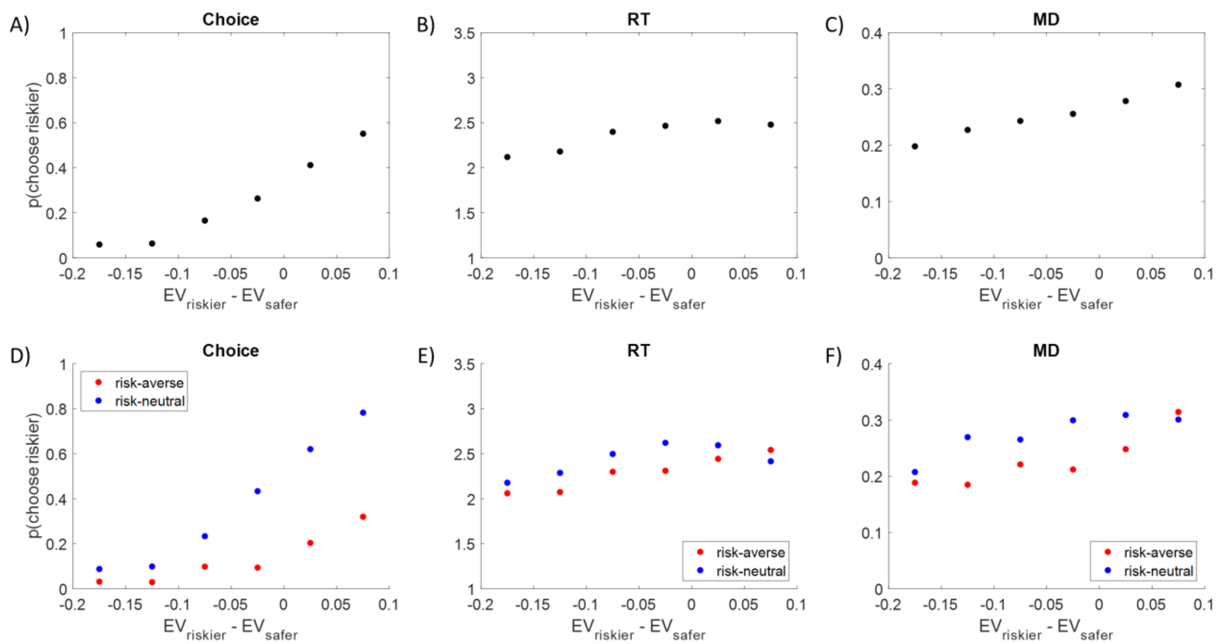

Figure 3. Behavioral results. *A-C) Choice probability (of choosing the riskier option), response time (RT), and maximum mouse trajectory deviation (MD) plotted against EV difference (riskier minus safer), pooled across all participants. D-F) Choice probability, RT, and MD plotted against EV difference, pooled across participants after performing a median split by risk preference type. Each dot represents an average across all trials within EV difference bins of equal range.*

It is well known that people vary in terms of their risk preferences. Given that our data clearly indicated that our participants (as a whole) were averse to risk, we checked if we could reasonably split our participants into two groups: risk-neutral and risk-averse. To do this, we first estimated a parameter *alpha* (α) that is commonly-used to capture risk aversion (cf, Prospect Theory (Tversky & Kahneman, 1992)), separately for each participant: $X \leftarrow X^{\alpha}$ (X = monetary amount). We used VBA to fit the basic prospect theory logistic model of choice probability for each participant (including both α and the *c* parameter that captures probability distortion; specifically, $\mathrm{EV} = \mathrm{X} * \mathrm{P}$ became $\mathrm{EU} = X^{\alpha} * \frac{P^c}{(P^c + (1-P)^c)^{\frac{1}{c}}}$ as in the original Prospect Theory formulation). A value of α close to 1 indicates no aversion to risk, whereas a value of α close to 0 indicates extreme risk aversion. We thus created groups of participants whose estimated alpha parameter was either less than or greater than the group median (0.4). We labeled these groups "risk-averse" and "risk-neutral", respectively. We did not find evidence for a "risk-seeking" group (i.e., $\alpha > 1$) within our pool of participants. We then repeated the above qualitative analysis, separately by group. Clearly this separation into groups was meaningful, as two distinct patterns in the data emerged. Risk-averse participants rarely chose the riskier option unless the EV of the riskier option was much greater than the EV of the safer option (Figure 3D in red). They also took longer to decide and deviated more from a straight trajectory as the EV of the riskier option increased (within the available range; Figure 3E-F in red). Risk-neutral participants showed the patterns expected for people who choose based on simple economic calculations unaffected by risk aversion. Their probability of choosing the riskier option also increased monotonically as a function of EV difference (riskier - safer), but their point of neutrality (p(ch) = 0.5) was when EV difference equaled zero (Figure 3D in blue). Furthermore, their RTs were longest and MDs were largest on trials where EV difference was closest to zero (Figure 3E-F in blue).

To confirm the apparent group differences, we regressed choice (of the riskier option) on EV difference (dV), log(RT) on abs(dV), and MD on abs(dV), separately for each dependent variable and risk preference group. The relative magnitudes of the regression coefficients across groups suggest that EV had a much greater influence on the choice probability and RT of the risk-neutral participants relative to the risk-averse participants (coefficients for dV: **choice**: risk-averse = 1.34, 95% CI = [0.92 1.75]; risk-neutral = 3.39, 95% CI = [2.88 3.90]; **RT**: risk-averse = -0.56, 95% CI = [-0.87 -0.26]; risk-neutral = -1.07, 95% CI = [-1.71 -0.43]), but there was no relationship between dV and MD for either group (coefficients for dV: risk-averse = 0.03, 95% CI = [-0.33 0.26]; risk-neutral = -0.34, 95% CI = [-1.02 0.35]). Given the clear differences in group behavior, we decided to also test for group differences in the model comparison analysis reported below.

Beyond the effects of EV difference, we hypothesized that choice behavior would differ according to attribute salience. Specifically, we predicted that information processing about the salient attribute on each trial would be facilitated relative to the less salient attribute. This should lead to more extreme choice probability, faster RT, and smaller MD on trials where salience aligned with EV difference (i.e., when the EV of the riskier option is higher and the salient attribute is monetary amount, or when the EV of the safer option is higher and the salient attribute is probability). We label such trials as *no conflict* trials. In the opposite situation, where salience conflicts with EV difference (*conflict* trials), choice probability should be closer to 50% (i.e., chance level), RT should be slower, and mouse trajectories should deviate more. To test for this, we split trials by conflict type (no or yes) and repeated the analysis reported above. Overall, we found no difference in choice probability (Figure 4A), but a shift in the peak of the RT and MD distributions (Figure 4B-C). Expecting to find group differences (risk-averse versus risk-neutral), we repeated this conflict analysis while splitting the data by group. With the risk-averse participants, we found the anticipated patterns: for trials where the safer EV was higher, conflict increased the probability of choosing the riskier option (towards chance level) and increased both RT and MD; for trials where the riskier EV was higher, conflict decreased the probability of choosing the riskier option (away from chance level) and decreased both RT and MD (Figure 4D-F). However, with the risk-neutral participants, we found the opposite patterns: for trials where the safer EV was higher, conflict decreased the probability of choosing the riskier option (away from chance level) and decreased both RT and MD; for trials where the riskier EV was higher, conflict increased the probability of choosing the riskier option (away from chance level) and decreased both RT and MD (Figure 4G-I). Here, the behavioral data are internally consistent, in the sense that RT and MD both decreased whenever choice probability moved away from chance level and RT and MD both increased whenever choice probability moved towards chance level. But salience seemed to affect the risk-neutral participants in a way opposite to what we predicted. It is likely that the options we labeled as salient were indeed salient even for this group, based on the behavioral effects. However, it seems like those options were salient not because of the high magnitude attributes but rather the low ones (because on average, when our options had a higher-than-average level of one attribute, they had a lower-than-average level of the other). In other words, risk-averse participants might focus more on probability, except when monetary amount is particularly high. Risk-neutral participants, in contrast, might focus more on monetary amount, except when probability is particularly low. Thus, it seems that high monetary amounts are salient for risk-averse people, and low probabilities are salient for risk-neutral people.

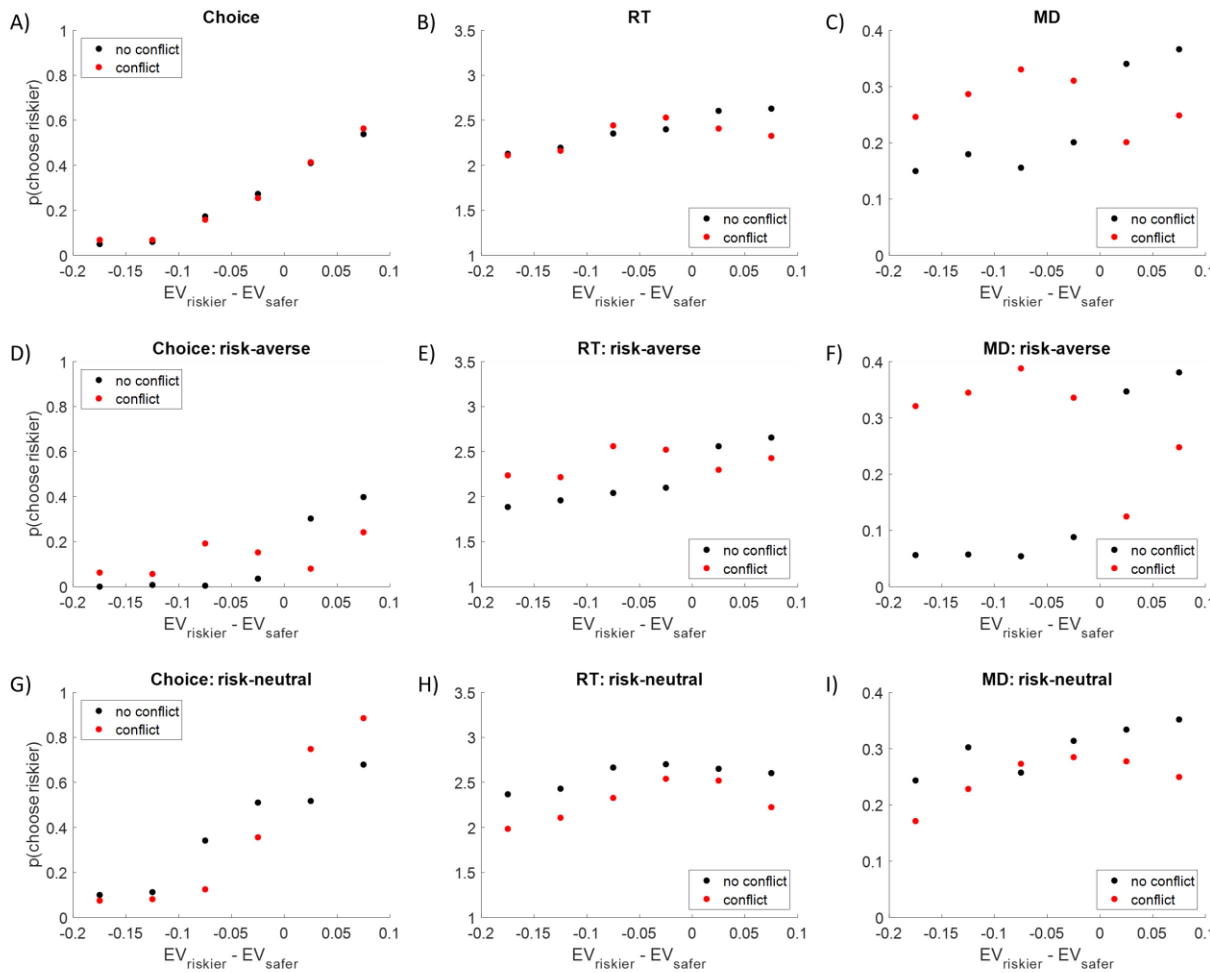

Figure 4. Impact of conflict on choice behavior. *Choice probability (of choosing the riskier option), response time (RT), and maximum mouse trajectory deviation (MD) plotted against EV difference (riskier minus safer), pooled across all participants (plots A-C), risk-averse participants (plots D-F), and risk-neutral participants (plots G-I). Each plot shows the data separated by trial type (conflict versus no conflict). Each dot represents an average across all trials within EV difference bins of equal range.*

To confirm the group differences, we regressed choice on dV + dV*conflict, log(RT) on abs(dV) + abs(dV)*conflict, and MD on abs(dV) + abs(dV)*conflict, separately for each dependent variable and risk preference group (where conflict is an indicator variable = 1 on conflict trials and = 0 on no conflict trials). The regression coefficients for dV were qualitatively similar across groups (**choice**: risk-averse = 1.95, 95% CI = [1.27 2.62]; risk-neutral = 2.73, 95% CI = [2.18 3.28]; **RT**: risk-averse = -0.88, 95% CI = [-1.27 -0.49]; risk-neutral = -0.48, 95% CI = [-1.04 0.08]; **MD**: risk-averse = -0.63, 95% CI = [-0.98 -0.27]; risk-neutral = -0.03, 95% CI = [-0.67 0.61]; Figure 5). In contrast, the regression coefficients for dV*conflict were qualitatively opposite across groups (**choice**: risk-averse = -1.30, 95% CI = [-2.16 -0.43]; risk-neutral = 1.40, 95% CI = [0.91 1.89]; **RT**: risk-averse = 0.73, 95% CI = [0.35 1.11]; risk-neutral = -1.34, 95% CI =[-2.00 -0.67]; **MD**: risk-averse = 1.35, 95% CI = [0.70 1.99]; risk-neutral = -0.69, 95% CI = [-1.57 0.19]; Figure 5).

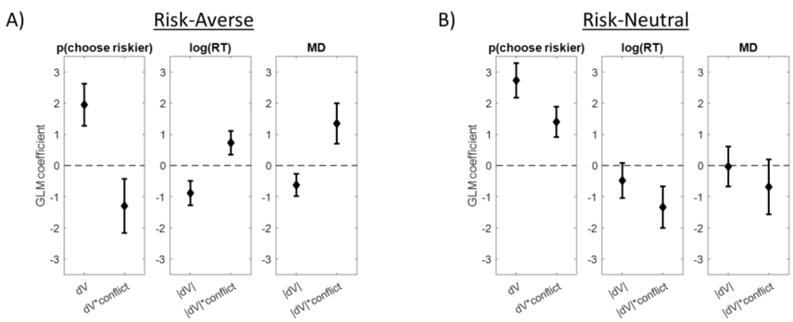

Figure 5. Impact of expected value difference and conflict across risk preference groups. *The effects of expected value difference (dV = EV of riskier option – EV of safer option) and dV on conflict trials (dV*conflict) on choice probability (for the riskier option), log(RT), and maximum deviation (MD) for risk-averse (plot A) and risk-neutral participants (plot B); dV is unsigned for RT and MD; diamonds represent fixed effects regression coefficients; error bars represent 95% confidence intervals.*

Qualitative Model Predictions

The DDM variants that we compare make distinguishable predictions regarding the impact of EV difference and conflict on choice behavior, when separately considering risk-averse and risk-neutral participants. We show in Figure 6 the qualitative predictions that each model (*evDDM*, *maDDM*, *ev+maDDM*, *evDDM+sal*, *maDDM+sal*, and *ev+maDDM+sal*, each simulated under its participant-specific best-fitting parameters) makes with respect to the effects of EV difference (dV; riskier minus safer) and dV in *conflict* trials on choice probability (for the riskier option), RT, and MD, and how this compares to the empirical data. To generate the synthetic data, we used the same choice trials that participants faced, along with each participant's best-fitting parameters for each of the four models. Next, we performed mixed effects regressions of choice (binomial) on dV and dV*conflict, and of log(RT) and MD (linear) on |dV| and |dV|*conflict, pooling together data from all simulated participants and including participants as random effects. For the risk-averse group, all the models except for the *evDDM* are able to account for the qualitative patterns of the associations between dV (with and without conflict) and choice probability, RT, and MD (Figure 6A). For the risk-neutral group, only the models with a parameter related to salience (*evDDM+sal*, *maDDM+sal*, and *ev+maDDM+sal*) are able to account for the relative magnitudes and directionality of the associations between dV (with and without conflict) and choice probability, RT, and MD (Figure 6B). Note that what is important here is the qualitative patterns of the coefficients for dV and dV*conflict *within* each data source, not the absolute magnitudes or comparisons across data sources.

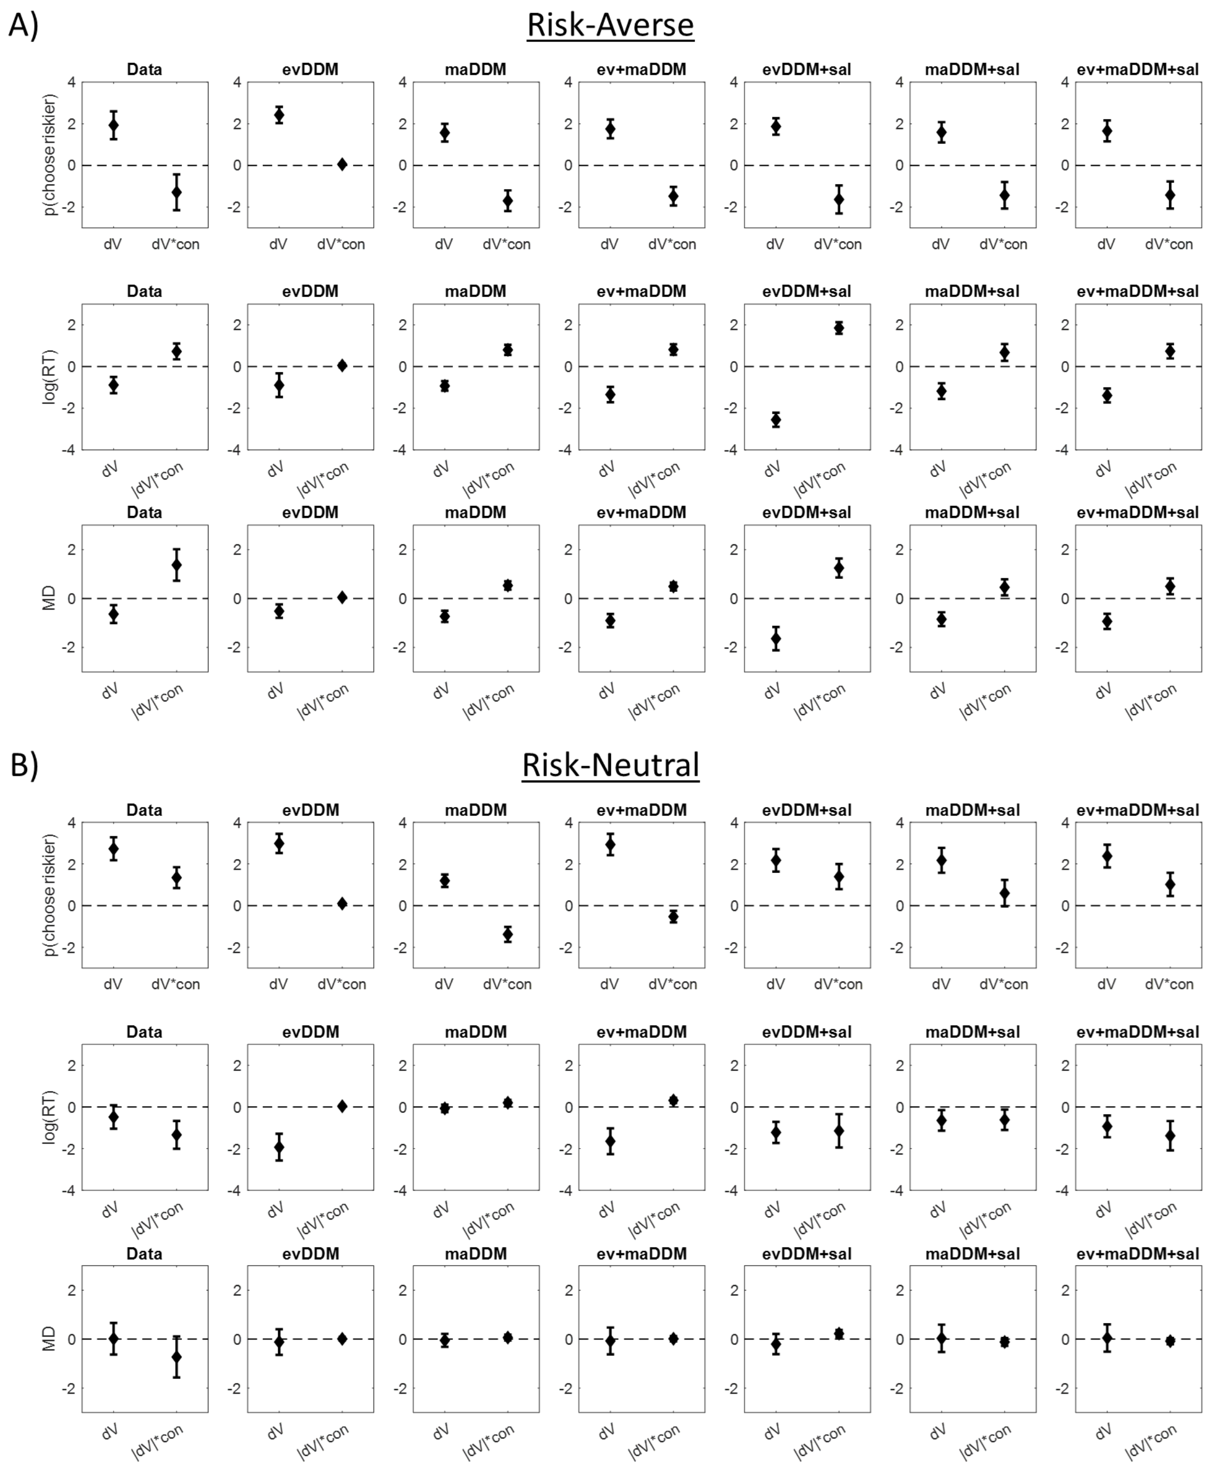

Figure 6. Model predictions. *Qualitative predictions of the effects of expected value difference (dV = EV of riskier option – EV of safer option) and dV on conflict trials (dV*conflict) on choice probability (for the riskier option), log(RT), and maximum deviation (MD) in the empirical*

(column 1) and simulated data (columns 2-7; shown for responses simulated using the best-fitting parameters for each model); dV is unsigned for RT and MD; diamonds represent fixed effects regression coefficients; error bars represent 95% confidence intervals. Panel A shows the risk-averse group, panel B shows the risk-neutral group.

Model comparison

Our primary interest when conducting a model comparison of the DDM variants was in knowing whether risky decisions were made purely based on comparisons of the expected value (EV) of the available options, on direct comparisons of each of the individual attributes (monetary amount X and probability P), or on either EV or individual attributes with an additional influence of the most salient attribute on each trial. We separated this into two core research questions: 1) if people integrate evidenced before comparing it, compare evidence before integrating it, or both; 2) if the salience of individual attributes influence evidence accumulation. We tested each of these questions separately by performing model family comparisons based on the core model features reflective of each question. In brief, the model family approach consists in pooling together model evidence for models that share core components, to test those specific components against competitors (while setting aside other differences between models within each family). This approach is more robust that testing individual models when they could be considered to be variants of subclasses of similar models. As we have two main research questions, and each individual model contains features relevant to either question, we classified our models into family sets in two different ways. The first model family classification separates the models by evidence accumulation mechanism: Models 1 and 4 both rely on comparisons of EV; Models 2 and 5 both rely on comparisons of individual attributes; and Models 3 and 6 both rely on comparisons of both EV and individual attributes. We examined this family classification with respect to research question 1 above. The second model family classification separates the models by the impact of attribute salience: Models 1-3 are not impacted by salience and Models 4-6 are. We examined this family classification with respect to research question 2 above.

For question 1, we found that the ev+ma model family that uses comparisons of both individual attributes and EV dominated (estimated family frequency = 0.68; Figure 7). Testing the risk-averse group separately, we found that both families that incorporated comparisons of individual attributes dominated the family that only incorporated comparisons of EV (estimated family frequencies: ma = 0.54, ev+ma = 0.45; Figure 7). Testing the risk-neutral group separately, we found that both families that incorporated comparisons of EV dominated the family that only incorporated comparisons of individual attributes (estimated family frequencies: ev = 0.47, ev+ma = 0.51; Figure 7). Thus, we can conclude that risk-averse participants mostly rely on a compare-then-integrate strategy and risk-neutral participants mostly rely on an integrate-then-compare strategy, but half of the participants in each group seem to use a mixed strategy by considering both individual attributes and EV.

For question 2, we found that the DDM+sal model family dominated (estimated family frequency = 0.64; Figure 7). Testing the risk-averse group separately, we found that the DDM+sal family performed best (estimated family frequency = 0.55; Figure 7). Testing the risk-neutral group separately, we found that the DDM+sal family dominated (estimated family frequency = 0.73). We can thus conclude that most participants are affected by the salience of individual attributes. Interestingly, the participants who mostly compare EV across options (the risk-neutral group) seem to be influenced more by differences in attribute salience. This provides clear evidence that even those who seem to choose based on classical economic theory (i.e., to maximize

EV; Figure 7) nevertheless pay attention to the values of the individual attributes as a function of their salience.

To summarize, both our analyses indicate that participants consider individual attributes when making their choices, possibly in combination with EV comparisons.

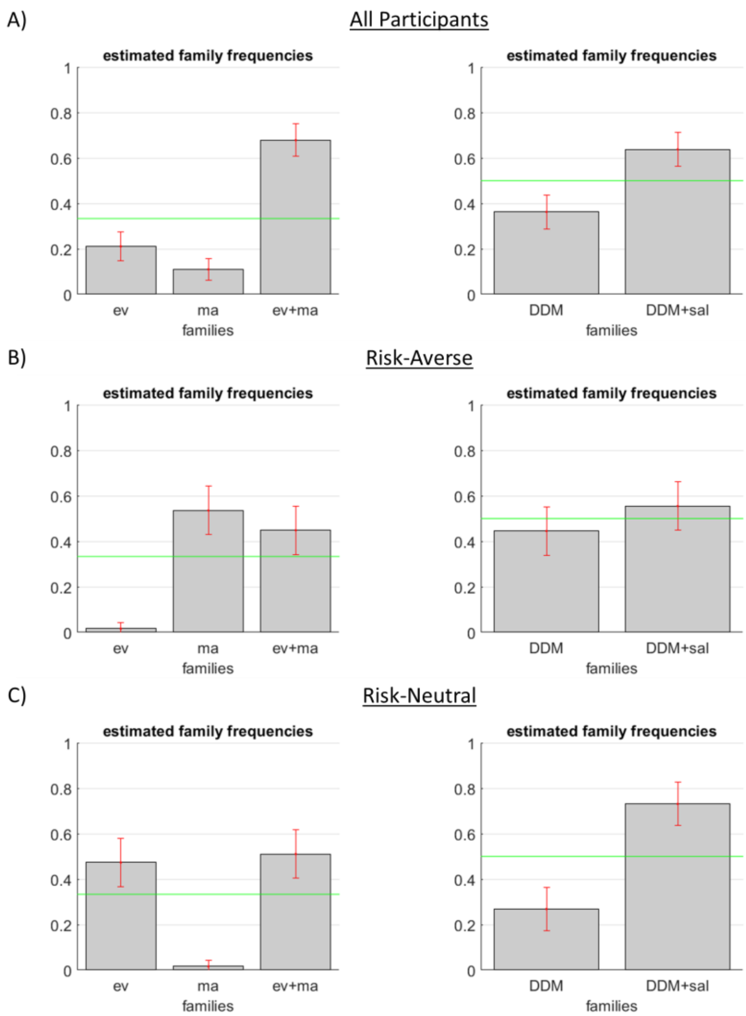

Figure 7. Model family comparison. *Estimated frequencies for model families based on evidence accumulation strategies (EV, individual attributes, or both; left column) or influence of attribute salience (not present or present; right column). Row A shows all participants, row B shows risk-averse participants only, and row C shows risk-neutral participants only. Green lines indicate chance level of family frequency if families were equally probable a priori.*

We next performed a standard quantitative model comparison that included all our individual models (Models 1-6). Across all participants, Model 6 (*ev+maDDM+sal*) dominated (estimated model frequency = 0.47, exceedance probability = 0.98; Figure 8A). Model frequency is defined as the number of participants with more evidence in favor of a particular model, and exceedance probability quantifies the likelihood that a particular model accounts for the data better than the other models (Stephan et al., 2009). Given the results of the model family comparisons reported above, it is not surprising that Model 6 won the individual model comparison, as it is the only model that is a member of both winning families. Repeating the individual model comparison for the risk-averse group only, we found that Models 2, 5, and 6 (*maDDM*, *maDDM+sal*, and *ev+maDDM+sal*) each performed above chance level, but none of these three significantly outperformed the other (estimated model frequency: M2 = 0.27, M5 = 0.27, M6 = 0.28; exceedance probability: M2 = 0.30, M5 = 0.28, M6 = 0.35; Figure 8B). Note that each of the best-performing models relies on comparisons of individual attributes. Repeating the model comparison for the risk-neutral group only, we found that Models 4 and 6 (*evDDM+sal* and *ev+maDDM+sal*) each performed above chance level, but neither significantly outperformed the other (estimated model frequency: M4 = 0.37, M6 = 0.35; exceedance probability: M4 = 0.54, M6 = 0.42; Figure 8C). Note that each of the best-performing models relies on the influence of individual attributes according to their salience. Although different models performed best for the different groups, all best-performing models required information about individual attributes in one way or another.

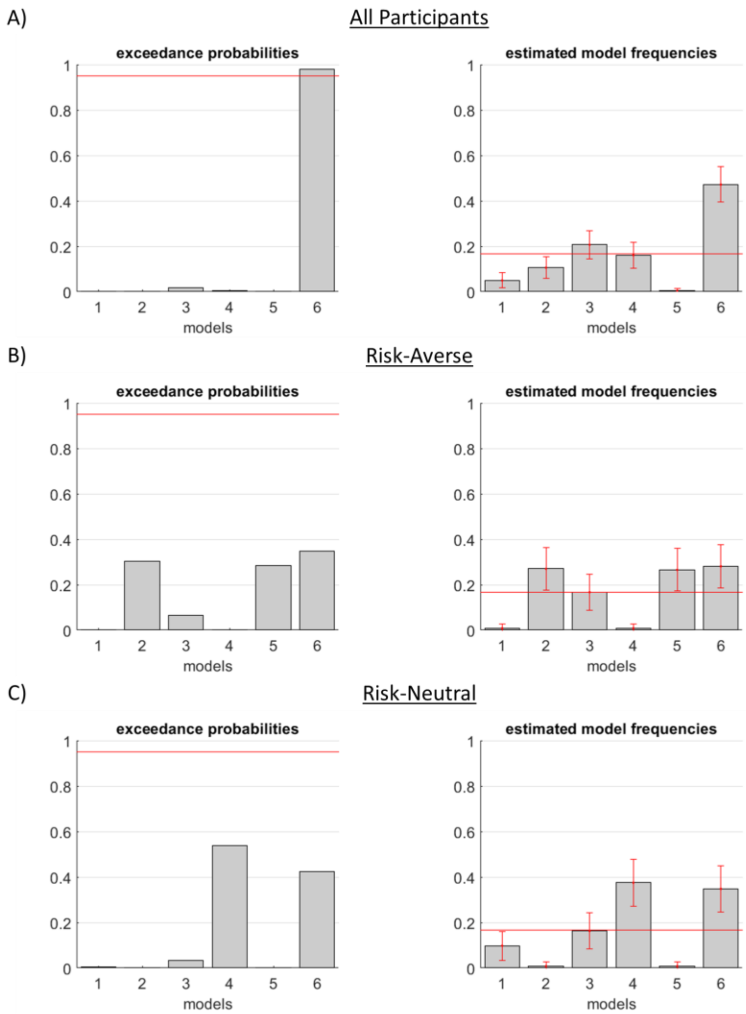

Figure 8. Model comparison. *Exceedance probabilities (left column) and estimated model frequencies across participants (right column) for each model. Row A shows all participants, row B shows risk-averse participants only, and row C shows risk-neutral participants only. Red lines indicate 95% threshold for exceedance probabilities (left plots) or chance level of model frequency if models were equally probable a priori (right plots).*

To more directly test whether both groups relied on information about both EV and individual attributes, independent of the salience effect, we performed a model comparison of Models 1-3 only (*evDDM*, *maDDM*, and *ev+maDDM*). We found that Model 3 decidedly dominated across all participants, with an estimated model frequency of 0.93 and an exceedance probability of 1. Repeating this comparison for the risk-averse group only, Model 3 won the competition (estimated model frequency = 0.64, exceedance probability = 0.91). Repeating this comparison for the risk-neutral group only, Model 3 decidedly dominated (estimated model frequency = 0.86, exceedance probability = 1). Thus, it appears that even those participants who mostly rely on comparisons of EV when deciding also rely on additional information based on comparisons of individual attributes. This aligns with previous findings (Lee & Hare, 2022).

Finally, we repeated all of the model comparisons reported above, this time excluding the equations for maximum deviation (MD). All results were nearly identical.

Discussion

In this paper, we asked whether choices between options defined by multiple attributes or dimensions, such as risky choices defined by the dimensions of "amount of money" and "probability of winning", are resolved by considering the expected value (EV) of each option that results from the preliminary integration of the different attributes (*integrate-then-compare*), the values of the individual attributes separately (*compare-then-integrate*), or both the EV and the individual attributes (*combined-comparison*). To address this question, we designed a risky choice task in which participants made decisions between pairs of options (one safer and one riskier) that varied systematically in terms of both EV and salience, here defined as one attribute having a value substantially greater than all other attribute values within a given pair of options, plausibly influencing a compare-then-integrate decision mechanism.

Our analyses of behavioral and kinematic indexes of response probabilities, response times (RTs), and curvatures of mouse trajectories during decisions indicate three main findings. First, participants performed the task properly, as shown by the fact that their response probabilities are significantly modulated by the EV of the options within each choice pair. Second, participants could be categorized into distinct groups based on risk preference type: risk-averse and risk-neutral. Third, attribute salience significantly affected response probability as well as decision dynamics (as evident in both RT and response trajectory curvature), but in an opposite direction for the two risk preference groups. When salience conflicted with the higher EV option, risk-averse participants were more likely to choose the salient option, whereas risk-neutral participants were less likely to choose the salient option. Furthermore, both groups were faster to respond when conflict (or lack of conflict) caused their choice probabilities to move away from chance level. On these trials, participants also moved the mouse with a more direct trajectory – which is an index of increased choice confidence and decreased attraction from the non-selected alternative (Hehman et al., 2015; Spivey et al., 2005). Note that the RT and trajectory curvature results are consistent with the different choice preferences of the two groups.

This pattern of results can perhaps be explained by the fact that risk-averse participants (by definition) prefer options with higher probabilities, whereas risk-neutral participants (by definition) mostly consider EV (or perhaps prefer higher monetary amounts). By design, on trials where salience was for the riskier option, such options included both an exceptionally large monetary amount and an exceptionally small probability. It seems that risk-averse participants, who usually pay more attention to probability, paid particular attention to the large monetary amount when it was salient, thus making the riskier option more attractive (hence the higher choice probability when EV favored the safer option but salience favored the riskier option) and the choice more difficult (hence the slower RT and more deviant response trajectory). The opposite applies for trials where the large probability was salient: the riskier option became less attractive (hence the lower choice probability when EV favored the riskier option but salience favored the safer option) and the choice easier (hence the faster RT and more direct response trajectory). Conversely, risk-neutral participants seem to have paid particular attention to the small probability when SAL = riskier and the small monetary amount when SAL = safer, thus making the corresponding options (with the apparently salient small-valued attributes) less attractive. Thus, for these participants, conflict between EV difference and salience (as we defined it, according to the largest attribute value) actually made the choices easier (hence the steeper psychometric choice curve, faster RT, and more direct response trajectory on conflict trials). Although this divergence of the attribute salience effect across risk preference types was not something we predicted, we believe that the overall pattern of results clearly shows that salience exerted an influence on choice behavior and decision dynamics.

Taken together, these results show that participants are sensitive to both the EV of the different options and the salience of the different attributes during their choices. The above results lend some support to conceptual models that highlight separate, within-attribute comparisons, but cannot fully distinguish whether these attribute comparisons are alternative or simultaneous to EV comparisons (as envisaged by the combined model).

To address this last question more directly, we designed and systematically compared six versions of drift-diffusion model (DDM) that incorporate different hypotheses about how multi-attribute decisions are made: an *integrate-then-compare model* (Model 1: *evDDM*) that only allows EV comparisons, a *compare-then-integrate model* (Model 2: *maDDM*) that only allows separate comparisons of individual attributes, a *combined-comparison model* (Model 3: *ev+maDDM*) that allows both EV and attribute comparisons, a variant of the *integrate-then-compare model* that allows evidence to accumulate at different rates depending on which option contains the salient attribute on a given trial (Model 4: *evDDM+sal*), a variant of the *compare-then-integrate model* that allows evidence to accumulate at different rates depending on which attribute is salient on a given trial (Model 5: *maDDM+sal*) , and a variant of the *combined-comparison model* that allows evidence to accumulate at different rates depending on which option contains the salient attribute on a given trial or on which attribute is salient on a given trial (Model 6: ev+*maDDM+sal*).

Our results show that the model family in which evidence accumulates according to comparisons of both EV and individual attributes dominates the families in which either only EV or only attributes is considered. Our results also show that the model family in which evidence accumulation is affected by attribute salience dominates the family in which it is not. At the level of individual models, the model that best accounts for the data is Model 6, which contains comparisons of both EV and individual attributes, plus additional flexibility by allowing evidence accumulation to be influenced by attribute salience. Note that the extra flexibility provided by Model 6 gives it an advantage in the model comparison, which holds despite the fact that our model comparison method balances model accuracy and complexity (hence disfavoring model flexibility when it is not accompanied by a significant increase in accuracy). Model 6 has the most free parameters of any of our models, yet still dominates in the quantitative comparison. These results suggest that both types of people (risk-averse and risk-neutral) accumulate evidence about both EV and individual attributes while deliberating, and that information processing during risky decisions is influenced by the salience of the individual option attributes.

Salience is understood to impact the cognitive processes involved in decision-making through its ability to attract and capture attention. Attention-related arousal has been shown to accompany the presentation of offers with exceptionally high monetary amounts (Leuker et al., 2019). When a decision set contains a relatively high numerical magnitude for one of the options or attributes, the attention that it captures is thought to cause the decision-relevant information about that option or attribute to be weighted more heavily (Bordalo et al., 2012; Tsetsos et al., 2012). The idea that attention leads to differential weighting has been computationally demonstrated in previous work based on the so-called *attentional drift-diffusion model* (aDDM; (Krajbich et al., 2010; Smith & Krajbich, 2019)). Under the aDDM, the evidence accumulation rate for the option that holds the attention of the decision-maker at any point in time (assumed to be the option that is visually fixated) is enhanced. This work has recently been extended to include attention towards a particular attribute in addition to attention towards a particular option (Yang & Krajbich, 2022). Our *evDDM+sal* and *maDDM+sal* models align well with the aDDM framework, in that attention (whether assumed by monitoring visual fixations or the highest numerical magnitudes) changes evidence accumulation rates. For our risk-averse group, it seems that the highest numerical magnitudes were the most salient. For our risk-neutral group, it seems that it was perhaps the *lowest* numerical magnitudes that were the most salient, rather than the highest. This would be consistent with increased evidence accumulation for the "worse" attribute (for the risk-neutral group) if these participants were choosing which option to reject rather than which to accept, which has been demonstrated in previous aDDM work (Sepulveda et al., 2020).

In sum, our results, based on both behavioral measures and drift-diffusion models, suggest that risky decisions are resolved differently for different types of people. For risk-averse people, such decisions seem to be resolved by running in parallel multiple competitions between the individual attributes – possibly in combination with an additional competition between expected values – and then integrating the partial results of the separate competitions afterwards (Hunt et al., 2014; Roe et al., 2001; Turner et al., 2018; Usher & McClelland, 2004). For risk-neutral people, risky decisions seem to be resolved by a competition between expected values but influenced by the salience of individual attributes. These results stand in contrast with the assumption of standard economic theory that choices require a unique comparison of expected values, and suggests that at the cognitive level, decision processes might be more distributed than commonly assumed. Furthermore, our results show that choice behavior can be altered when one attribute for one option is particularly salient. Finally, and intriguingly, individual differences in risk preference correspond to an opposite pattern in the effects of attribute salience, and these differences manifest differential evidence accumulation rates for either individual attributes (for risk-averse people) or options (for risk-neutral people) according to trial type and risk preference trait.

Data and Code Availability

The data and analysis code used in preparation of this manuscript are publicly available on the Open Science Framework at https://osf.io/7pu98/.

Funding

This research received funding from the European Union's Horizon 2020 Framework Programme for Research and Innovation under the Specific Grant Agreement No. 945539 (Human Brain Project SGA3) to GP and the European Research Council under the Grant Agreement No. 820213 (ThinkAhead) to GP.

References

Alós-Ferrer, C. (2018). A Dual-Process Diffusion Model. *Journal of Behavioral Decision Making*, *31*(2), 203–218. https://doi.org/10.1002/bdm.1960

Amasino, D. R., Sullivan, N. J., Kranton, R. E., & Huettel, S. A. (2019). Amount and time exert independent influences on intertemporal choice. *Nature Human Behaviour*, *3*(4), 383–392. https://doi.org/10.1038/s41562-019-0537-2

Barca, L., & Pezzulo, G. (2012). Unfolding Visual Lexical Decision in Time. *PLoS ONE*.

Barca, L., & Pezzulo, G. (2015). Tracking Second Thoughts: Continuous and Discrete Revision Processes during Visual Lexical Decision. *PLOS ONE*, *10*(2), e0116193. https://doi.org/10.1371/journal.pone.0116193

Bordalo, P., Gennaioli, N., & Shleifer, A. (2012). Salience Theory of Choice Under Risk. *The Quarterly Journal of Economics*, *127*(3), 1243–1285. https://doi.org/10.1093/qje/qjs018

Calluso, C., Committeri, G., Pezzulo, G., Lepora, N., & Tosoni, A. (2015). Analysis of hand kinematics reveals inter-individual differences in intertemporal decision dynamics. *Experimental Brain Research*, *233*(12), 3597–3611.

Cisek, P. (2012). Making decisions through a distributed consensus. *Current Opinion in Neurobiology*, *22*(6), 927–936. https://doi.org/10.1016/j.conb.2012.05.007

Clay, S. N., Clithero, J. A., Harris, A. M., & Reed, C. L. (2017). Loss Aversion Reflects Information Accumulation, Not Bias: A Drift-Diffusion Model Study. *Frontiers in Psychology*, *8*, 1708. https://doi.org/10.3389/fpsyg.2017.01708

Dai, J., & Busemeyer, J. R. (2014). A Probabilistic, Dynamic, and Attribute-wise Model of Intertemporal Choice. *Journal of Experimental Psychology. General*, *143*(4), 1489–1514. https://doi.org/10.1037/a0035976

Daunizeau, J., Adam, V., & Rigoux, L. (2014). VBA: A Probabilistic Treatment of Nonlinear Models for Neurobiological and Behavioural Data. *PLOS Computational Biology*, *10*(1), e1003441. https://doi.org/10.1371/journal.pcbi.1003441

Farashahi, S., Donahue, C. H., Hayden, B. Y., Lee, D., & Soltani, A. (2019). Flexible combination of reward information across primates. *Nature Human Behaviour*, *3*(11), Article 11. https://doi.org/10.1038/s41562-019-0714-3

Feltgen, Q., & Daunizeau, J. (2021). An Overcomplete Approach to Fitting Drift-Diffusion Decision Models to Trial-By-Trial Data. *Frontiers in Artificial Intelligence*, *0*. https://doi.org/10.3389/frai.2021.531316

Fiedler, S., & Glöckner, A. (2012). The Dynamics of Decision Making in Risky Choice: An Eye-Tracking Analysis. *Frontiers in Psychology*, *3*. https://www.frontiersin.org/article/10.3389/fpsyg.2012.00335

Freeman, J. B., & Ambady, N. (2010). MouseTracker: Software for studying real-time mental processing using a computer mouse-tracking method. *Behav Res Methods*, *42*(1), 226–241. https://doi.org/10.3758/BRM.42.1.226

Friston, K., Mattout, J., Trujillo-Barreto, N., Ashburner, J., & Penny, W. (2007). Variational free energy and the Laplace approximation. *NeuroImage*, *34*(1), 220–234. https://doi.org/10.1016/j.neuroimage.2006.08.035

Glickman, M., Sharoni, O., Levy, D. J., Niebur, E., Stuphorn, V., & Usher, M. (2019). The formation of preference in risky choice. *PLOS Computational Biology*, *15*(8), e1007201. https://doi.org/10.1371/journal.pcbi.1007201

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*, *4*(6), Article 6. https://doi.org/10.1038/s41562-020-0822-0

Gluth, S., Spektor, M. S., & Rieskamp, J. (2018). Value-Based Attentional Capture Affects Multi-Alternative Decision Making. *ELife*, *7*, e39659. https://doi.org/10.7554/eLife.39659

Grasman, R. P. P. P., Wagenmakers, E.-J., & van der Maas, H. L. J. (2009). On the mean and variance of response times under the diffusion model with an application to parameter estimation. *Journal of Mathematical Psychology*, *53*(2), 55–68. https://doi.org/10.1016/j.jmp.2009.01.006

He, L., Analytis, P. P., & Bhatia, S. (2022). The Wisdom of Model Crowds. *Management Science*, *68*(5), 3635–3659. https://doi.org/10.1287/mnsc.2021.4090

Hehman, E., Stolier, R. M., & Freeman, J. B. (2015). Advanced mouse-tracking analytic techniques for enhancing psychological science. *Group Processes & Intergroup Relations*, *18*(3), 384–401.

Hunt, L. T., Dolan, R. J., & Behrens, T. E. J. (2014). Hierarchical competitions subserving multi-attribute choice. *Nature Neuroscience*, *17*(11), 1613–1622. https://doi.org/10.1038/nn.3836

Krajbich, I., Armel, C., & Rangel, A. (2010). Visual Fixations and the Computation and Comparison of Value in Simple Choice. *Nature Neuroscience*, *13*(10), 1292–1298. https://doi.org/10.1038/nn.2635

Lee, D. G., & Hare, T. A. (2022). Evidence accumulates for individual attributes during value-based decisions. *Decision*, No Pagination Specified-No Pagination Specified. https://doi.org/10.1037/dec0000190

Lee, D. G., & Usher, M. (2021). Value certainty in drift-diffusion models of preferential choice. *Psychological Review*, No Pagination Specified-No Pagination Specified. https://doi.org/10.1037/rev0000329

Lepora, N. F., & Pezzulo, G. (2015). Embodied choice: How action influences perceptual decision making. *PLoS Comput Biol*, *11*(4), e1004110. https://doi.org/10.1371/journal.pcbi.1004110

Leuker, C., Pachur, T., Hertwig, R., & Pleskac, T. J. (2019). Too good to be true? Psychological responses to uncommon options in risk–reward environments. *Journal of Behavioral Decision Making*, *32*(3), 346–358. https://doi.org/10.1002/bdm.2116

Lim, S.-L., Penrod, M. T., Ha, O.-R., Bruce, J. M., & Bruce, A. S. (2018). Calorie Labeling Promotes Dietary Self-Control by Shifting the Temporal Dynamics of Health- and Taste-Attribute Integration in Overweight Individuals. *Psychological Science*, *29*(3), 447–462. https://doi.org/10.1177/0956797617737871

Lopez-Persem, A., Domenech, P., & Pessiglione, M. (2016). How prior preferences determine decision-making frames and biases in the human brain. *ELife*, *5*, e20317. https://doi.org/10.7554/eLife.20317

Maier, S. U., Raja Beharelle, A., Polanía, R., Ruff, C. C., & Hare, T. A. (2020). Dissociable mechanisms govern when and how strongly reward attributes affect decisions. *Nature Human Behaviour*, *4*, 949–963. https://doi.org/10.1038/s41562-020-0893-y

Noguchi, T., & Stewart, N. (2014). In the attraction, compromise, and similarity effects, alternatives are repeatedly compared in pairs on single dimensions. *Cognition*, *132*(1), 44–56. https://doi.org/10.1016/j.cognition.2014.03.006

O'Donoghue, T., & Somerville, J. (2018). Modeling risk aversion in economics. *Journal of Economic Perspectives*, *32*(2), 91–114.

Padoa-Schioppa, C. (2011). Neurobiology of Economic Choice: A Good-Based Model. *Annual Review of Neuroscience*, *34*(1), 333–359. https://doi.org/10.1146/annurev-neuro-061010-113648

Park, S. Q., Kahnt, T., Rieskamp, J., & Heekeren, H. R. (2011). Neurobiology of Value Integration: When Value Impacts Valuation. *Journal of Neuroscience*, *31*(25), 9307–9314. https://doi.org/10.1523/JNEUROSCI.4973-10.2011

Penny, W. D. (2012). Comparing Dynamic Causal Models using AIC, BIC and Free Energy. *NeuroImage*, *59*(1), 319–330. https://doi.org/10.1016/j.neuroimage.2011.07.039

Peterson, J. C., Bourgin, D. D., Agrawal, M., Reichman, D., & Griffiths, T. L. (2021). Using Large-Scale Experiments and Machine Learning to Discover Theories of Human Decision-Making. *Science*, *372*(6547), 1209–1214. https://doi.org/10.1126/science.abe2629

Pezzulo, G., & Cisek, P. (2016). Navigating the Affordance Landscape: Feedback Control as a Process Model of Behavior and Cognition. *Trends in Cognitive Sciences*, *20*(6), 414–424. https://doi.org/10.1016/j.tics.2016.03.013

Ratcliff, R. (1978). A theory of memory retrieval. *Psychol. Rev*, *85*(2), 59–108.

Ratcliff, R., & McKoon, G. (2008). The Diffusion Decision Model: Theory and Data for Two-Choice Decision Tasks. *Neural Computation*, *20*(4), 873–922. https://doi.org/10.1162/neco.2008.12-06-420

Ratcliff, R., & Rouder, J. N. (1998). Modeling Response Times for Two-Choice Decisions. *Psychological Science*, *9*(5), 347–356. https://doi.org/10.1111/1467-9280.00067

Reeck, C., Wall, D., & Johnson, E. J. (2017). Search Predicts and Changes Patience in Intertemporal Choice. *Proceedings of the National Academy of Sciences*, *114*(45), 11890–11895. https://doi.org/10.1073/pnas.1707040114

Rigoux, L., Stephan, K. E., Friston, K. J., & Daunizeau, J. (2014). Bayesian Model Selection for Group Studies—Revisited. *NeuroImage*, *84*, 971–985. https://doi.org/10.1016/j.neuroimage.2013.08.065

Roe, R. M., Busemeyer, J. R., & Townsend, J. T. (2001). Multialternative decision field theory: A dynamic connectionist model of decision making. *Psychological Review*, *108*(2), 370–392. https://doi.org/10.1037/0033-295x.108.2.370

Rustichini, A., & Padoa-Schioppa, C. (2015). A neuro-computational model of economic decisions. *Journal of Neurophysiology*, *114*(3), 1382–1398.

Sepulveda, P., Usher, M., Davies, N., Benson, A. A., Ortoleva, P., & De Martino, B. (2020). Visual attention modulates the integration of goal-relevant evidence and not value. *ELife*, *9*, e60705. https://doi.org/10.7554/eLife.60705

Sheng, F., Ramakrishnan, A., Seok, D., Zhao, W. J., Thelaus, S., Cen, P., & Platt, M. L. (2020). Decomposing loss aversion from gaze allocation and pupil dilation. *Proceedings of the National Academy of Sciences*, *117*(21), 11356–11363. https://doi.org/10.1073/pnas.1919670117

Smith, S. M., & Krajbich, I. (2019). Gaze Amplifies Value in Decision Making. *Psychological Science*, *30*(1), 116–128. https://doi.org/10.1177/0956797618810521

Spivey, M., Grosjean, M., & Knoblich, G. (2005). Continuous attraction toward phonological competitors. *Proceedings of the National Academy of Sciences of the USA*, *102*, 10393-10398.

Stephan, K. E., Penny, W. D., Daunizeau, J., Moran, R. J., & Friston, K. J. (2009). Bayesian model selection for group studies. *NeuroImage*, *46*(4), 1004–1017. https://doi.org/10.1016/j.neuroimage.2009.03.025

Stewart, N. (2011). Information Integration in Risky Choice: Identification and Stability. *Frontiers in Psychology*, *2*. https://www.frontiersin.org/article/10.3389/fpsyg.2011.00301

Sullivan, N., Hutcherson, C., Harris, A., & Rangel, A. (2015). Dietary self-control is related to the speed with which attributes of healthfulness and tastiness are processed. *Psychological Science*, *26*(2), 122–134. https://doi.org/10.1177/0956797614559543

Tsetsos, K., Chater, N., & Usher, M. (2012). Salience driven value integration explains decision biases and preference reversal. *Proceedings of the National Academy of Sciences*, *109*(24), 9659–9664. https://doi.org/10.1073/pnas.1119569109

Turner, B. M., Schley, D. R., Muller, C., & Tsetsos, K. (2018). Competing theories of multialternative, multiattribute preferential choice. *Psychological Review*, *125*(3), 329–362. https://doi.org/10.1037/rev0000089

Tversky, A., & Kahneman, D. (1992). Advances in prospect theory: Cumulative representation of uncertainty. *Journal of Risk and Uncertainty*, *5*(4), 297–323. https://doi.org/10.1007/BF00122574

Tversky, A., & Simonson, I. (1993). Context-Dependent Preferences. *Management Science*, *39*(10), 1179–1189.

Usher, M., & McClelland, J. L. (2001). On the time course of perceptual choice: The leaky, competing accumulator model. *Psychological Review*, *108(3)*, 550–592.

Usher, M., & McClelland, J. L. (2004). Loss aversion and inhibition in dynamical models of multialternative choice. *Psychological Review*, *111*(3), 757–769. https://doi.org/10.1037/0033-295X.111.3.757

Yang, X., & Krajbich, I. (2022). A dynamic computational model of gaze and choice in multi-attribute decisions. *Psychological Review*, No Pagination Specified-No Pagination Specified. https://doi.org/10.1037/rev0000350

Zhao, W. J., Walasek, L., & Bhatia, S. (2020). Psychological mechanisms of loss aversion: A drift-diffusion decomposition. *Cognitive Psychology*, *123*, 101331. https://doi.org/10.1016/j.cogpsych.2020.101331

Supplementary Materials

for

Risky decisions are influenced by individual attributes as a function of risk preference

Douglas G. Lee[1], Marco D'Alessandro[2], Pierpaolo Iodice[3], Cinzia Calluso[4], Aldo Rustichini[5], and Giovanni Pezzulo[2]

[1] Tel Aviv University, School of Psychological Sciences, Tel Aviv, Israel
[2] Institute of Cognitive Sciences and Technologies, National Research Council, Rome, Italy
[3] Université de Rouen, Rouen, France
[4] Luiss University, Rome, Italy
[5] University of Minnesota, Hanson Hall, Minneapolis, USA

Posterior predictive checks

Here we provide posterior predictive checks that demonstrate how well the models were able to fit the data. Overall, the models did a good job of predicting the data. Note, however, that the empirical MD had a highly skewed distribution, a pattern that the models (based on our simple linear formulation of MD) could not replicate. Also, recall that we fit MD using a very simple linear model in addition to the primary drift-diffusion model analysis (based on more advanced equations for choice probability and mean RT), so one might expect the fits for MD to be worse than those for the other variable.

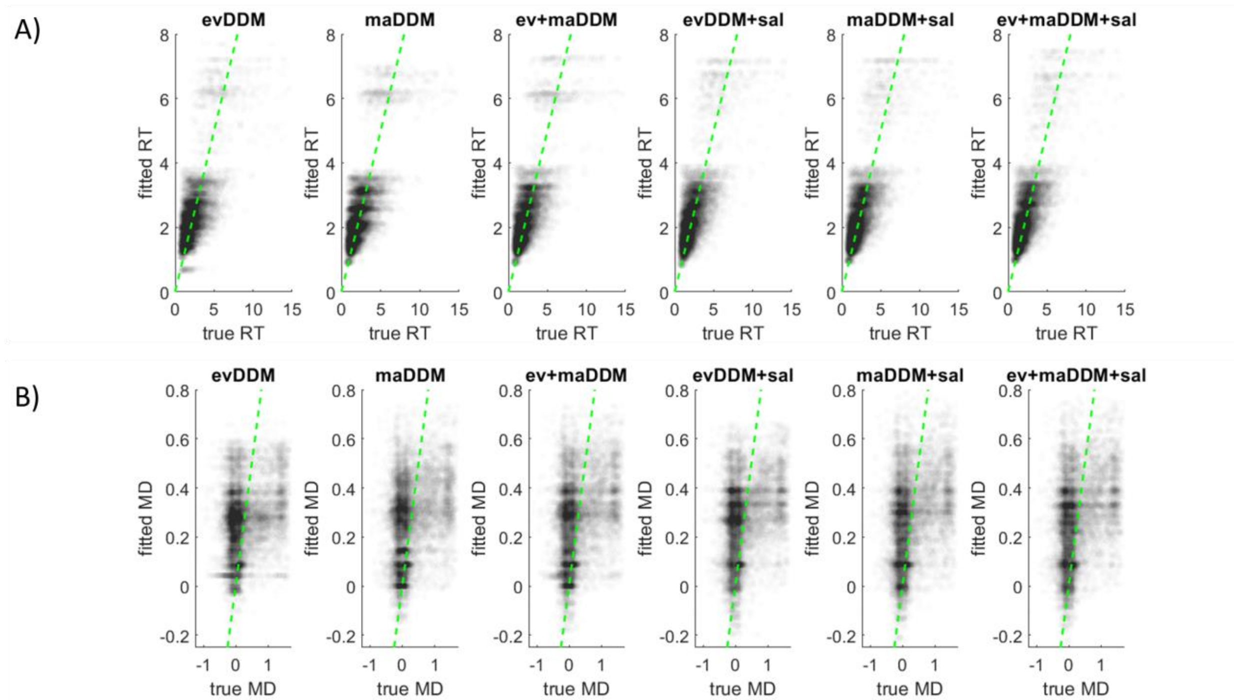

Figure S2. Posterior predictive checks. *A) Under each model, the predicted fitted RT is plotted against the empirical true RT. B) Under each model, the predicted fitted MD is plotted against the empirical true MD. Each dot represents one trial, pooled across participants. The green dashed line indicates the ideal perfect correspondence between true and fitted data.*

We also repeated the qualitative analysis as reported at the beginning of the Results section in the main text, this time for both the real data and the data predicted under each model. The models all seem to do a fair job of predicting all variables (choice, RT, and MD). However, the *ev+maDDM*, *evDDM+sal*, *maDDM+sal*, and *ev+maDDM+sal* seem to predict the data better than the *evDDM* and *maDDM*, which aligns with the fact that they performed better in the quantitative model comparison.

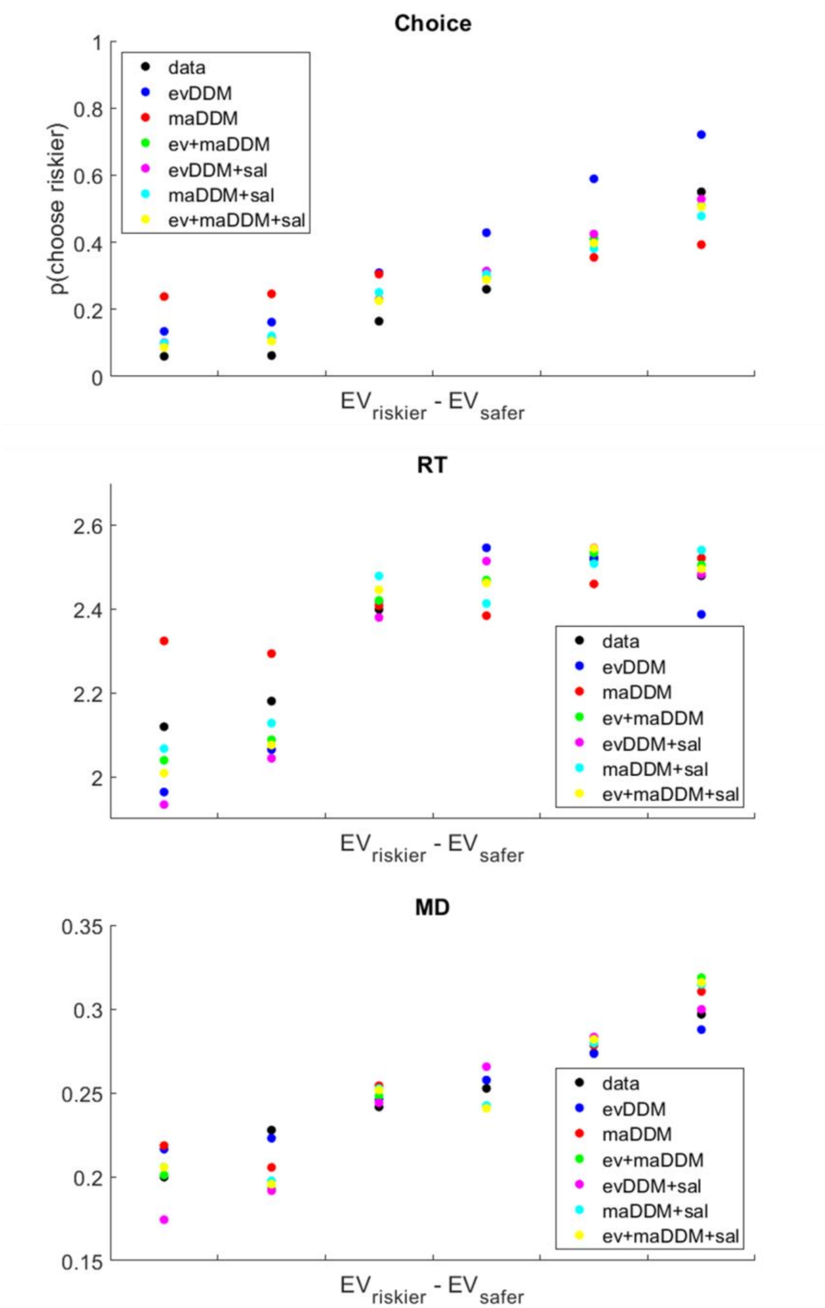

Figure S2. Posterior predictive checks. *Binned behavioral plots as in Figure 3 in the main manuscript. Here we show the true data (black dots) along with the model output data under each model (color code).*