Computational mechanisms for context-based behavioral interventions: A large-scale analysis

Wenjia Joyce Zhao^{1,2}, Aoife Coady², and Sudeep Bhatia^{2,3}

January 31, 2022

Abstract

Choice context influences decision processes, and is one of the primary determinants of what people choose. This insight has been used by academics and practitioners to study decision biases and design behavioral interventions to influence and improve choices. In this paper we analyze the effects of context-based behavioral interventions on the computational mechanisms underlying decision making. We collect data from two very large laboratory studies involving nineteen prominent behavioral interventions, and model the influence of each intervention using a leading computational model of choice in psychology and neuroscience. This allows us to parametrize the biases induced by each intervention, interpret these biases in terms of underlying decision mechanisms and their properties, quantify similarities between interventions, and predict how different interventions alter key choice outcomes. In doing so, we offer researchers and practitioners a theoretically principled approach to understanding and manipulating choice context in decision making.

Keywords: decision making; behavioral interventions; context effects; computational modeling

¹ Department of Psychology, The Ohio State University

² Department of Psychology, University of Pennsylvania

³ Department of Marketing, University of Pennsylvania

Send correspondence to Wenjia Joyce Zhao. Email: joyce.wenjia.zhao@gmail.com. Funding was received from the National Science Foundation grant SES-1847794.

Choices depend on incidental contextual factors, such as defaults, social norms, and time pressure, and understanding how these contextual factors influence choice has been a major focus of research in the behavioral sciences [1,2,3]. Context dependence also forms the psychological rationale behind behavioral interventions, or “nudges” [4]. Such interventions influence behavior by altering choice context, and are important tools for policy makers as they can improve decisions without restricting people’s capacity to choose.

Although useful qualitative taxonomies of context effects and behavioral interventions have been proposed, none offer a cognitively and neurobiologically inspired model (such as in [5,6]) with which to parametrize and predict the effect of many different contextual factors on behavior. There are two major challenges to accomplishing this goal. The first is the lack of data: Large-scale experiments that collect data on multiple behavioral interventions within a single decision task are necessary to build unified models, but these experiments are very rare. The second involves the psychological complexity of context, which can alter the choice predispositions that decision makers have prior to evaluating the choice options, how desirable the options appear to the decision makers during evaluation, as well as how hard decision makers are willing to deliberate. It is not clear how theorists should incorporate these diverse effects into a single computational model of choice.

Our paper addresses these challenges using a comprehensive dataset of choices under the influence of nineteen contextual factors. These contextual factors include a diverse set of situational variables that have been shown to influence choice without altering people’s capacity to choose (for simplicity, we do not examine the role of explicit persuasion, financial incentivization, direct education, or coercion, which are also important determinants of choice). In our experiments, we assigned over 1,200 laboratory participants to different context-based behavioral interventions (adapted from existing research), in two decision scenarios involving either consumer or financial choice. Participants made 160 binary choices both with and without the interventions, resulting in over 300,000 total choices and response times (RTs; time taken to make decisions) for quantitative analysis. The interventions are summarized in Table 1, and an example

of the experimental stimuli in Exp. 1 is shown in Fig. 1A (Exp. 2 involved nearly identical stimuli presentation). See *Methods* for additional details and Tables S1 and S2 for experimental instructions.

Table 1. Summary of behavioral interventions.

Category	Intervention	Procedure and instruction
Prominence of Information	1. Attribute order	<i>Within-participant manipulation.</i> The quality attribute was positioned above or below the price attribute.
	2. Option order	<i>Within-participant manipulation.</i> The high-quality option was positioned to the left or right of the low-price option.
	3. Quality priming	<i>Implemented in Exp. 1 only.</i> Photos of appetizing and expensive food were shown on each instruction page (before the task and between blocks).
	4. Price priming	<i>Implemented in Exp. 1 only.</i> Photos of US dollars were shown on each instruction page (before the task and between blocks).
	5. Quality information	<i>Implemented in Exp. 2 only.</i> Participants read a short passage explaining health insurance deductibles, and answered multiple choice questions about why a low-deductible health insurance plan could be beneficial.
	6. Price information	<i>Implemented in Exp. 2 only.</i> Participants read a short passage explaining health insurance premiums, and answered multiple choice questions about why a low-premium health insurance plan could be beneficial.
	7. Attribute salience	One of the two attributes appeared with an orange frame, which highly contrasted with the background.
	8. Option salience	One of the two options appeared with an orange frame, which highly contrasted with the background.
Task framing	9. Default	One of the options was pre-selected, and additional key pressing was required to switch the option.
	10. Reject (vs. accept)	Participants indicated which option they preferred less instead of indicating which option they preferred more.
Social information	11. Social norm	The more popular option in each choice problem (based on a pilot study) was indicated using an orange frame.
	12. Recommendation	The option recommended by the experimenters (based on the same pilot study for condition 11) was indicated using an orange frame.
Affect	13. Positive emotion	Before the choice task, participants took 5-10 minutes to write a report of a happy event from their life. They were also instructed to reread the event during each break.

	14. Negative emotion	Before the choice task, participants took 5-10 minutes to write a report of a sad event from their life. They were also instructed to reread the event during each break.
Speed and accuracy	15. Time pressure	Participants were instructed to make choices as quickly as possible.
	16. Accuracy instruction	Participants were asked to indicate choices only after they were completely certain.
	17. Cognitive load	Participants performed an additional memory task, in which they remembered a six-digit number before each block and reported the number at the end of each block.
	18. Accountability	Participants wrote a one-paragraph justification of one of their choices (randomly selected at the end of the choice task).
	19. Font fluency	The stimuli were shown in a hard-to-read font.

Note: All interventions that are not explicitly listed as “within-participant” are “between-participant”.

We modeled our choice dataset with the diffusion decision model (DDM), which proposes that decision makers dynamically accumulate the relative evidence favoring each choice option. Choices are made when the relative evidence reaches one of two decision boundaries, and the time to reach the decision boundary corresponds to the response time. DDM is the dominant computational model of two-alternative forced choice in psychology and neuroscience [7,8,9], and has been shown to successfully describe binary responses in perception, memory, categorization, language, and value-based decision making. We fit a hierarchical DDM to all participants assigned to a given behavioral intervention condition (~ 40 participants in each condition, with each participant completing 160 baseline trials and 160 intervention trials, see *Methods* for details). For each fit, we quantified the effect of the behavioral intervention in terms of how it shifted three key DDM parameters: the starting point, the drift rate, and the decision boundary. Changes in these three parameters have systematic effects on choice probability and RT. All else equal, starting points favoring one option over the other generate more frequent choices and quicker responses for the favored option; higher drift rates favoring one option over the other generate more frequent choices for the favored option, with higher absolute drift rates causing quicker responses for both options; and higher decision

boundaries lead to slower responses and more consistent choices. Fig. 1B summarizes the diffusion decision model and illustrates these predictions in a hypothetical choice between a high-quality and a low-price option. Although the choice probability for the high-quality option increases in all three panels, RT distributions vary based on the parameter that is causing the shift. For this reason, changes to DDM parameters can be identified based on choice and RT data.

Figure 1. (A) Screenshots for Exp. 1, with an example of the baseline condition (left) and an option salience intervention (right). (B) Illustration of the diffusion decision model (DDM), with hypothetical changes to the starting point, drift rate and decision boundary parameters.

(A)

Restaurant 1		Restaurant 2
9.1	Score	8.7
25	Price	19

Restaurant 1		Restaurant 2
9.1	Score	8.7
25	Price	19

(B)

The DDM allows us to model many different behavioral interventions within a single computational framework. Importantly, it also allows us to interpret their effects in terms of underlying cognitive and neurobiological mechanisms, as well as their statistical properties [10,11,12]. The starting point parameter, for example, corresponds to an automatic pre-decisional bias (or equivalently a baseline activation bias in neural units), and can be interpreted as a statistical prior regarding the relative desirability of the available options. The drift rate governs the formation of preferences (or equivalently the rate of change of neural activation) during deliberation, and reflects evolving probabilistic beliefs as decision makers integrate evidence favoring one option over the other. Lastly, the decision boundary is a threshold-level of preference (or activation) strength necessary to initiate choice, and can be seen as specifying an acceptable error rate for the decision. By observing how different contextual factors influence these three DDM mechanisms, we are able to study the behavioral, cognitive, neural, and statistical consequences of decision context on choice.

Results

Overall Trends. The decision scenarios in Exp. 1 and 2 were taken from consumer and financial choice domains respectively. The former offered participants choices between pairs of restaurants (each with a *user rating* and a *price*), whereas the latter offered participants choices between pairs of health insurance plans (each with a *deductible* and a *premium*). These two scenarios were tested in two separate preregistered experiments (<https://osf.io/y59jk/>). Each experiment involved fifteen between-participant interventions. Thirteen between-participant interventions were implemented in both experiments. Money priming and food quality priming were implemented only in Exp. 1, and informational prompts emphasizing the importance of either low-deductibles or low-premiums were implemented only in Exp. 2. Finally, in addition to the seventeen between-participant interventions across the two experiments, we also counterbalanced the presentation order of choice options and attributes within participants in each experiment. This resulted in a total of nineteen context-based interventions for our analysis. Note that for expositional simplicity, we refer to high-rating/high-price restaurants in Exp. 1 and low-deductible/high-

premium insurance plans in Exp. 2 as “high-quality” options. Low-rating/low-price restaurants and high-deductible/low-premium insurance plans are referred to as “low-price” options.

Overall, choices in the experiments were fairly balanced, with 53.5% and 56.8% of choices favoring the high-quality options in Exp. 1 and 2 respectively. As expected, different behavioral interventions had different choice probability and RT effects (see Fig. S1 for means and 95% CIs of behavioral intervention effects). Our diffusion decision models were able to capture this variability, with correlations greater than 0.94 in predicting individual-level changes to choice probabilities and RTs, and correlations greater than 0.97 in predicting aggregate changes to choice probabilities and RTs, across interventions. These patterns are shown in Fig. 2.

Figure 2. Observed and simulated intervention effects on choice probabilities and RTs in Exp.1 and 2. Grey points correspond to changes for individual participants. Colored labels correspond to aggregate changes for interventions, averaged over participants. Displayed correlations capture the relationship between observed and simulated changes on the individual-level.

Note: Participants with behavioral shifts outside of the range of the x and y axes of this figure are shown in Fig. S9.

Our main goal was to quantify and compare the effects of the interventions on the three DDM parameters. These parameter effects are summarized in Fig. 3, which displays the shift in the starting point (panel A), the drift rate (panel B), and the decision boundary (panel C), for each intervention, and each experiment (in what follows, we analyze the effects of social norms, recommendations, defaults, and option and attribute salience separately for trials selectively targeting the high-quality and low-price options in the two experiments). Fig. S2 provides means and 95% CIs of DDM parameter shifts, Fig. S3 shows that the behavioral effects of interventions are correlated with the estimated effects of interventions on DDM parameters, Figs. S4 and S5 provide results of cross validation predictions of choice probabilities and RT percentiles, Fig. S6 shows that the DDM parameters are recoverable, and Fig. S8 shows that DDM parameters correlate strongly with those estimated using an alternate evidence accumulation model.

Figure 3. Effect of behavioral interventions on DDM parameters for Exp. 1 and 2. Positive (negative) starting point and drift rate effects correspond to biases favoring the high-quality (low-price) option. HQ and LP denote interventions selectively targeting the high-quality or low-price options in a trial. Results are based on group-level parameters in each condition.

Starting Point Effects. In Fig. 3A we see that the interventions with the strongest effect on starting points in the two experiments were the social norm [13,14] and recommendation [15] interventions. This indicates that participants developed an automatic pre-decisional response bias (equivalently, a baseline activation bias or statistical prior) in favor of the norm or recommendation, increasing its choice probability and reducing its RT. These results are consistent with prior work, which finds that social cues can alter the automatic processing of choice options [16].

Defaults [17,18] also had positive effects on starting points in both experiments, so that making an option a default, typically generated a mild pre-decisional bias for that option. We obtained similar but weaker results for option salience [19,20]. These findings validate previous claims regarding the automatic effects of defaults and salient options [21,22]. Additionally, informational prompts about the importance of low premiums (low-price) biased starting points in favor of the corresponding health insurance plans in Exp. 2. We did not observe the analogous effect for informational prompts about the importance of low deductibles (high-quality).

Overall, the starting point effects of the interventions were largely consistent across the two experiments, with a Spearman correlation of 0.46 for the interventions. That said, some interventions had stronger effects in one experiment than the other. For example, participants were more likely to display a biased starting point in favor of the low-price option when asked to reject (instead of accepting) [23] one of the options in Exp. 1. We did not find a systematic starting point effect for reject framing in Exp. 2, in which both attributes were monetary. This could reflect a tendency to prioritize monetary attributes in rejection decisions. Conversely negative emotion [24,25] generated a starting point bias in Exp. 2 (favoring the high-quality insurance plan), but did not influence starting points in Exp. 1. Positive emotions did not yield analogous effects, suggesting that these effects could have involved negative emotions like fear and sadness, which are known to bias choices in favor of the safer choice [26]. We also found a difference between the two domains for font fluency [27]: Disfluent fonts led to a weak bias favoring the high-quality insurance plan in Exp. 2 but had no such effect in Exp. 1. We speculate that this could reflect differences in the importance or difficulty of the two decisions. Finally note that some starting point effects could be a

product of unique features of our experimental design, such as repeated choices. Experience leads to priors (and subsequently starting point effects) favoring preferred options, and inexperienced individuals in the real-world making one-off choices may not display such biases.

Drift Rate Effects. Unsurprisingly, we found that informational prompts had an effect on the drift rates in Exp. 2 (favoring the option that was dominant on the prompted attribute; Fig. 3B). Thus, informational prompts alter the evaluation of choice options. Likewise, priming money or quality [28] led to corresponding drift rate effects in Exp. 1.

Drift rates were also very strongly influenced by social norm and experimenter recommendation interventions in both experiments. Thus, in addition to developing pre-decisional biases favoring normative and recommended options, participants also developed explicit preferences for these options. This finding is consistent with prior work that finds that social cues shift people's preferences in favor of conforming options [16]. By contrast, default and salient options did not have drift rate effects. Decision makers did begin with pre-decisional biases for these options, but did not display biased preferences for these options once they began deliberating. The null effects of defaults are consistent with prior work suggesting that defaults operate primarily through automatic and non-evaluative mechanisms [4]. The null effects of salience are consistent with research suggesting the disproportionate effect of salience and accessibility on automatic (vs. deliberative) processing [22].

Again, drift rate effects were consistent across the two experiments, with a Spearman correlation of 0.39. However, we found that cognitive load [29] and time pressure [30] biased Exp. 1 drift rates in favor of the high-quality option, but Exp. 2 drift rates in favor of the low-price option. It seems therefore that such interventions do not only alter the amount of time and effort participants are willing to put into the decision, but also how they evaluate the choice options, specifically, their preferences for price over quality, which may vary across decision scenarios. This may reflect an adaptive response to effort/accuracy tradeoffs in decision making [31,32]. Finally, as with the starting point effects discussed in the previous sections, we found that negative emotion had a stronger effect on drift rates in Exp. 2 than Exp. 1, and that the effect of positive emotion was almost completely absent in the two experiments. This indicates that the

emotion effects documented above persist beyond pre-decisional bias, and influence the evaluation of options and the formation of preferences.

Decision Boundary Effects. The final mechanism we analyzed was the decision boundary, whose effects are shown in Fig. 3C. Here we can see that increasing decision maker accountability [33] had a strong positive effect on the boundary for both experiments. Intuitively, decision makers that were accountable for their decisions deliberated longer and required a greater degree of confidence (or equivalently, activation) before choosing. This is equivalent to having a lower acceptable error probability for the decision. Our finding is consistent with considerable prior work on accountability in decision making [33]. We also found a positive effect of font fluency interventions on the decision boundary in both experiments: Making the font disfluent has also been found to increase deliberative processing [27].

Other interventions that had a strong effect on the decision boundary were reject (vs. accept) framing and defaults which increased the boundaries in both experiments. Both interventions implicitly involve rejection decision frames, which have been shown to increase RTs when options are desirable [34]. Additionally, as expected, time pressure had a strong negative effect on the decision boundary. When asked to make quicker decisions, decision makers reduced their thresholds and increased acceptable error levels.

Once again, the decision boundary effect was consistent across the two conditions, with a Spearman correlation of 0.51. However, interestingly, recommendations reduced decision boundaries in Exp. 2, but not in Exp. 1, suggesting that participants need less evidence before deciding when given social information in more important or more difficult decisions. We also found a positive effect of accuracy instructions on decision boundary for Exp. 1, but not for Exp. 2, possibly because Exp. 2 already had high boundaries.

Comparative Effects. Our analytical approach allows us to describe distinct contextual effects within a unified representational framework. This can shed light on the similarities and differences between different behavioral interventions. We illustrate this in Fig. 4. Panel A of this figure plots each intervention in each of our two experiments in a single three-dimensional *space of behavioral interventions*. The coordinates for each intervention capture its absolute, standardized effects on the starting point, drift rate, and decision boundary parameters. Here we can see that the different behavioral interventions span a wide

region of the space, with some interventions having similar effects on all three mechanisms, and others having a stronger effect on one or two of the mechanisms.

Figure 4. (A) 3-D space of behavioral interventions based their absolute, standardized effects on the starting point, drift rate, and decision boundary parameters. (B) Cognitive effect sizes of the interventions. These are based on the distance between an intervention vector and the origin of the space. (C) Hierarchical clustering of intervention vectors (averaged across Exp.1 and 2). Results are based on group-level parameters in each condition.

(A)

(B)

(C)

Fig. 4B displays the overall magnitude of the intervention vectors. Vector magnitude captures the total *cognitive effect size* of the interventions, that is, the cumulative absolute effect of the intervention on the drift rates, starting points, and decision boundaries of our DDM models. This figure shows that recommendations, social norms, accountability, time pressure and defaults have the strongest effects across decision scenarios, whereas emotions, priming, and accuracy instructions have the weakest effects across our decision scenarios.

Finally, interventions also have different degrees of similarity to each other, as revealed by the relative positions of the vectors in Fig. 4A. We show this more rigorously by clustering the interventions in Fig. 4C. Here we see that the social norm and recommendation interventions cluster together - These typically involve social cues that bias the decision maker in favor of one option over the other. We also find that default and rejection interventions (which involve the manipulation of task frames) cluster together, as do salience manipulations (which involve changes to information prominence). Finally, accountability and time pressure manipulations share a cluster as they both involve altering accuracy and speed tradeoffs.

Discussion

Computational Modeling of Large-scale Behavioral Data. The successful use of context in various scientific and practical applications requires an integrative framework for understanding how diverse contextual factors and behavioral interventions shape choice. We have taken steps towards the development of such a framework by collecting a very large experimental dataset of choice behavior under the influence of many distinct contextual factors. Our dataset is an order of magnitude larger than that used in most existing research, and offers an unparalleled insight into the computational mechanisms influenced by different behavioral interventions. We fit the parameters of the diffusion decision model [7,8], a prominent model of choice process, on decisions under the influence of each contextual factor in our dataset. The parameters of this model describe three core computational mechanisms underlying decision making: pre-decisional bias (or equivalently baseline activation or statistical prior), drift rate (change in activation/use of evidence during choice deliberation) and decision boundary (activation threshold/acceptable error) [11].

Our results show that different interventions have selective influences on these three computational mechanisms. In doing so, they reveal nuanced insights about how the interventions alter choice probability and RT. For example, we find that providing social norms increases the choice probability for the normative option by altering both starting points and drift rates, whereas making an option a default increases choice probability by only altering starting points. These differences imply that social norms have an effect on preferences, but defaults do not. Prior work that measures intervention effects using only changes to choice probability cannot disentangle these selective influences. Such prior work also cannot measure the effects of interventions such as accountability, which increase the extent of deliberation, but may not have a systematic effect on choice probability.

The power of our model relies on robust estimation of its various parameters. In the supplemental material, we demonstrate that the parameters of our model can be accurately identified, and through cross validation analysis, we show that the model makes good out-of-sample predictions of choice probabilities and response time distributions. We also show that replacing the DDM with an alternative evidence accumulation model does not alter our interpretation of the intervention effects. Our model is not too complex, but could it be too simple? Many additional cognitive mechanisms can be added to the simple diffusion decision model, such as between-trial variability in the parameters, collapsing boundaries, and non-linear transformation of attribute values. We doubt that these extensions will change our main results, though they may reveal additional nuances in the effects of behavioral interventions on cognitive processes.

Model-based Intervention Design. Our results offer academics and practitioners model-based methods for choosing between contextual influences based on the goals of the intervention. If the intervention aims to change people’s preferences, then drift rate interventions like recommendations and social norms (as well as informational prompts), are suitable nudges. By contrast, defaults, option salience manipulations, and other interventions that have a selective effect on the starting point, are better for altering choice probabilities without changing preferences. Starting point effects in the DDM diminish as decision boundaries increase, implying that the effects of starting point interventions are stronger for quick and error-prone decision makers, and weaker for decision makers that are willing to deliberate extensively [4]. Of

course, some settings may require interventions that merely encourage extensive deliberation, without systematically biasing choice in favor of either option. In such settings, our analysis would recommend decision boundary interventions, like accountability and rejection task framing.

The DDM model also offers a data-driven taxonomy of behavioral interventions in terms of their core computational mechanisms. The intervention space shown in Fig. 4 can, for example, quantify the similarities between different interventions, which in turn can be used to substitute one intervention for another or to develop novel interventions composed of two or more interventions. Consider, for example, combining an intervention that shifts the drift rate to favor a target option with an intervention that increases the decision boundary. The former would increase the choice probability of the target to above 50%. The latter would raise the threshold and thus reduce the probability of choosing the less preferred option. If we assume that intervention effects compose additively in the parameter space, the boundary intervention would amplify the effect of the drift rate intervention. More generally, choice probability depends on nuanced interactions between the three DDM parameters, and context effects that selectively influence different parameters can be combined to target the specific goals of the intervention. Although the effects of interventions on parameters may not be additive, and future experiments are necessary, our approach is unique in providing theory-driven quantitative predictions that can be used to motivate and design these experiments.

Many of the effects used to derive the above recommendations have an empirical precedent. For example, our paper replicates many prior findings regarding the effects of social cues, frames, defaults, attention, and emotion, on decision making. It also, however, provides a joint empirical investigation and integrative theoretical framework with which to compare and combine these contextual factors. Our framework is quantitative, and the insights derived from this framework come with precise predictions for choice probability and response time, allowing for a new type of rigor in the development and design of behavioral interventions.

Theory Integration Across Disciplines. Our experiments involve a laboratory paradigm with multiple choices made by each individual. This methodology is popular in psychology and neuroscience as it offers

precise experimental manipulations, data on multiple decision variables like choice and RTs, as well as controls for individual heterogeneity in model fitting. Yet behavioral interventions typically target real-world choices, which may differ from those observed in laboratory experiments. For example, most everyday decisions are spread out over time, such that people make fewer pairwise choices within a day but encounter similar problems every so often. This could lead to different intervention effects, especially for starting point parameters, which are sensitive to experience. Therefore, field experiments, ideally mega-studies that jointly test a number of different behavioral interventions in a representative population of real-world decision makers [35,36,37], should be conducted to test how our DDM-based conclusions generalize from lab experiments to field data.

It is also important to note that our approach adopts a fairly narrow conceptualization of the idea of a behavioral intervention and choice context. Behavioral interventions are typically understood to include not only features of the choice presentation (as with most of the interventions in our paper) but also crucial components of the behavior system such as people's opportunities and physical capabilities (see [38,39] for a review and discussion). Additionally, behavioral interventions can influence behavior not only through nudges [4], but also through explicit incentives, training, persuasion, and even coercion. We suspect that many of these types of interventions will selectively influence the three DDM parameters. Testing this, and by doing so, extending our approach to model and interpret prominent behavioral change frameworks in implementation science, is an important avenue for future work.

Conclusion

Our approach is unique in analyzing a large number of policy-relevant behavioral interventions using a prominent cognitive and neurocomputational theory of choice. The parameters that we analyzed in this paper have direct interpretations in terms of neural variables, and have been shown to respond to task and context manipulations across numerous domains in the cognitive and neural sciences, including perception, language, categorization, and memory. By relating prominent empirical regularities in the social and behavioral sciences to an established theoretical paradigm in the natural sciences, this paper offers a cohesive transdisciplinary approach to understanding human behavior.

Methods

Participants and Procedure. We recruited participants through various university experimental subject pools, with a combination of paid and unpaid participants to diversify our sample. All experiments took place on computers in a university behavioral laboratory, and the experiment lasted approximately one hour. Exp. 1 tested consumer choice and involved 608 participants (mean age = 23.2 years, 64.7% female) making choices between pairs of restaurants (each with a user rating and a price). Exp. 2 tested financial choice and involved 627 participants (mean age = 21.9 years, 69.6% female) making choices between pairs of health insurance plans (each with a deductible and a premium). Although we were constrained by our experimental paradigm to recruit participants that could take the experiment in the lab, our subject pools consisted of both college students and non-college students, resulting in a relatively diverse sample of participants.

In Exp. 1 we set the range of user ratings from 6 to 10, and the range of prices from 10 to 40. In Exp. 2 we set the range of deductibles from 0 to 6000, and the range of premiums from 30 to 380. We then randomly generated 80 unique choice problems within the pre-determined range of attribute values. Additionally, we ensured that only a small set of problems included a dominating option (10/80 in Exp 1 and 5/80 in Exp. 2; we excluded these problems when fitting the diffusion decision model). Before running the formal experiment in the lab, we tested the problem sets on an online experimental sample and found that choices in both experiments were fairly balanced.

Each experiment implemented fifteen between-participant interventions. Participants in each experiment were randomly assigned to one of the between-participant interventions and made multiple choices both with and without the intervention. Thirteen between-participant interventions were implemented in both experiments. The other interventions involved money priming and food quality priming in Exp. 1, and informational prompts emphasizing the importance of either low deductibles or low premiums in Exp. 2. This change was necessary as money and quality priming would not have different effects on the attributes in Exp. 2 (both of which were monetary), and informational prompts would be redundant in Exp. 2 (as participants have considerable prior experience with restaurant prices and user

ratings). Finally, in addition to the seventeen between-participant interventions in the two experiments, we also counterbalanced the presentation order of choice options and attributes in each experiment, allowing us to measure the effects of attribute and option ordering on choice within participants. This resulted in a total of nineteen context-based interventions for our analysis.

The between-participant intervention implemented for each participant was randomly selected at the beginning of the session. The order of the baseline condition and the intervention condition was also randomly determined for each participant. To minimize the carry-over effect of the first-presented condition, we asked participants to complete a 5-minute filler task between the two conditions. Thus, the experiment consisted of three total parts, with parts 1 and 3 corresponding to the baseline and intervention conditions (in random order) and part 2 corresponding to the filler task. The instructions for the baseline conditions for Experiments 1 and 2 can be found in Table S1. The additional procedures and instructions for the intervention conditions can be found in Table S2.

Each participant completed five practice trials in part 1 and part 3, 160 trials in the baseline condition (80 unique choice problems with option order counterbalanced), and 160 trials in the intervention condition (the same 80 unique choice problems as in the baseline condition with option and attribute order counterbalanced). The 160 trials in the two conditions were divided into 8 blocks of 20 trials each. There was a break after each block. Attribute order was held constant within a block and counterbalanced across blocks. All experimental procedures were preregistered on OSF (<https://osf.io/y59jk/>).

DDM. We grouped participants within each experiment based on which between-participant intervention condition they were assigned to, which resulted in 15 groups in each experiment. A hierarchical diffusion decision model was fit to each group using participants' choice and RTs recorded from the baseline condition, as well as the intervention condition. This hierarchical modeling approach estimates group and individual level parameters simultaneously, with group-level parameters (mean and standard deviations) forming the distributions from which individual subject estimates are sampled. We used a simple version of the diffusion decision model, with no between-trial variabilities in any of the parameters.

To analyze the two within-participant manipulations, we randomly selected 40 participants (that were assigned to any between-participant conditions) from both experiments, and used their data recorded from the baseline condition to test the order effects. The attribute order effect was studied using the parameter difference between trials in which the quality attribute was presented *above* the price attribute (which we, for simplicity, refer to as the intervention condition) and trials in which the quality attribute was presented *below* the price attribute (the baseline condition). The option order effect was studied using the parameter difference between trials in which the high-quality option was presented to the *left* of the low-price option (the intervention condition) and trials in which the high-quality option was presented to the *right* of the low-price option (the baseline condition).

In our models, the upper boundary was associated with the choice of a high-quality option (i.e., the option with a higher score or a lower deductible in Exp. 1 and 2 respectively), and the lower boundary was associated with the choice of a low-price option (i.e., the option with a lower price or a lower premium in Exp. 1 and 2 respectively). In the baseline condition, the drift rate can be written as $v = \Delta U = U_1 - U_2$, where U_1 and U_2 denote subjective utility for the high-quality option and the low-price option, respectively. We followed the common practice of the field and model subjective utilities as weighted sums of the attribute values [40]. Therefore, the utility difference between a pair of options is equal to the weighted sum of their attribute value differences. Without any context-based interventions, this is $v = \Delta U = v_{intercept} + v_{quality}\Delta Q + v_{price}\Delta P$. Here $v_{quality}$ and v_{price} denote multiplicative coefficients for the quality and price attributes respectively. The larger (the more positive) they are, the more important the attributes are to participants. The intercept, $v_{intercept}$, captures an overall stimuli-independent choice tendency. Positive $v_{intercept}$ captures a tendency to choose the high-quality option.

We captured the effects of interventions by allowing them to influence the starting point (z), the decision boundary (a), and the drift rate (v) of the DDM. For the starting point, we included a dummy variable to capture an additive shift between the intervention and the baseline condition. Moreover, for the attribute salience, option salience, default options, social norm, experimenter recommendation conditions,

we included another dummy variable to indicate which option was selectively (or asymmetrically) targeted by the manipulation (i.e., the option that dominated on the salient attribute, or the option that was salient, default, popular, or recommended). We also allowed these two types of variables (additive shift between the intervention vs. baseline condition, and asymmetric effect of one of the options) to have an effect on the drift rate. In addition to these, we allowed the weights for the two attributes (quality and price) to be different under different conditions. Finally, we permitted an additive shift in the decision boundary between the two conditions. To control for the task order effect, we generated a dummy variable for each participant, indicating whether the condition was implemented in the part 1 or part 3 of the experiment, and allowed this variable to influence the decision boundary parameter (a). See Table S4 for a summary of model specifications.

In summary, we assumed that the interventions had linear effects on all DDM components. We subsequently estimated these effects by fitting choice and RTs with the HDDMRegression function in HDDM [41], a Python package for hierarchical Bayesian estimation of diffusion decision models. The Bayesian approach permits direct inferences for parameter posterior distributions. To fit the models, 3 chains of 20,000 samples were generated, where the first 5,000 were burn-ins, and a thinning of 2 was applied. Gelman-Rubin convergence statistics for model parameters were all close to 1, suggesting that the sample size was sufficient for the chains to converge.

Additional Analysis. In many cases, we used multiple parameters to capture the difference on a specific DDM component between the intervention condition and the baseline condition, and thus had to integrate their effects to produce Fig. 3. Here we describe the procedure for generating the composite effect on the starting point first. For interventions without any asymmetric effects, this starting point effect is simply the group-level additive change of the starting point between the intervention condition and the baseline condition. For those interventions that selectively targeted (i.e., had an asymmetric influence) on one of the options (attribute salience, option salience, default options, social norm, experimenter recommendation conditions), the composite measure combines the group-level additive shift and asymmetric effect in all posterior samples. Because the asymmetric effect can either favor the high-quality

option or the low-price option, two starting point effects were computed separately for these conditions. One of the starting point effects combined the additive shift and the asymmetric effect (as if it is) favoring the high-quality option, and the other combined the additive shift and the asymmetric effect favoring the low-price option (these effects are shown separately in Fig. 3).

To derive the drift rate effect of an intervention, we first took the mean differences of both the quality and price attributes across all the choice problems, and used it to generate a typical choice problem. Based on the estimated group-level additive change, attribute weight change, and asymmetric shift between the intervention and the baseline conditions, we computed the overall drift rate change for this typical trial in all posterior samples, and use their mean to measure the drift rate effect of an intervention. As the drift rate is a linear function of the attribute values, our procedure is equivalent to first computing the drift rates for all questions and then taking the mean of all those drift rates. As in the starting point, the asymmetric effect in the drift rate can favor either the high-quality option or the low-price option, and thus two drift rate effects were computed separately for the interventions with asymmetric effects. One of the drift rate effects combined the additive shift, the weight change, and the asymmetric effect favoring the high-quality option, and the other combined the additive shift, the weight change, and the asymmetric effect favoring the low-price option. Finally, the decision boundary effect was measured as the group-level additive shifts between intervention and baseline conditions.

To generate the 3D space of behavioral interventions in Fig. 3A, we first scaled the DDM effects. This was achieved by dividing the posterior mean of the group-level composite effects by their associated standard deviations across the posterior samples. We then took these scaled effects' absolute values to reflect standardized intervention effect sizes on the DDM components. To account for interventions involving an asymmetric effect on the starting point or the drift rate, we averaged the standardized starting point or drift rate effect obtained from trials in which the highlighted option favored the high-quality or the low-price option. We used the Euclidean distance to calculate vector magnitudes, and used Ward's minimum variance method to generate the hierarchical clustering in Fig 4C.

References

1. Tversky, A., & Kahneman, D. (1981). The framing of decisions and the psychology of choice. *Science*, 211(4481), 453-458.
2. DellaVigna, S. (2009). Psychology and economics: Evidence from the field. *Journal of Economic Literature*, 47(2), 315-72.
3. Weber, E. U., & Johnson, E. J. (2009). Mindful judgment and decision making. *Annual Review of Psychology*, 60, 53-85.
4. Thaler, R. H., & Sunstein, C. R. (2009). *Nudge: Improving decisions about health, wealth, and happiness*. Penguin.
5. Glimcher, P. W., & Fehr, E. (Eds.). (2013). *Neuroeconomics: Decision making and the brain*. Academic Press.
6. Busemeyer, J. R., Gluth, S., Rieskamp, J., & Turner, B. M. (2019). Cognitive and neural bases of multi-attribute, multi-alternative, value-based decisions. *Trends in Cognitive Sciences*, 23(3), 251-263.
7. Ratcliff, R., & Smith, P. L. (2004). A comparison of sequential sampling models for two-choice reaction time. *Psychological Review*, 111(2), 333.
8. Gold, J. I., & Shadlen, M. N. (2007). The neural basis of decision making. *Annual Review of Neuroscience*, 30.
9. 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.
10. Bogacz, R., Brown, E., Moehlis, J., Holmes, P., & Cohen, J. D. (2006). The physics of optimal decision making: a formal analysis of models of performance in two-alternative forced-choice tasks. *Psychological Review*, 113(4), 700.
11. Bogacz, R. (2007). Optimal decision-making theories: linking neurobiology with behaviour. *Trends in Cognitive Sciences*, 11(3), 118-125.

12. Heekeren, H. R., Marrett, S., & Ungerleider, L. G. (2008). The neural systems that mediate human perceptual decision making. *Nature Reviews Neuroscience*, 9(6), 467-479.
13. Schultz, P. W., Nolan, J. M., Cialdini, R. B., Goldstein, N. J., & Griskevicius, V. (2007). The constructive, destructive, and reconstructive power of social norms. *Psychological Science*, 18(5), 429-434.
14. Cialdini, R. B., Goldstein, N. J. (2004). Social influence: Compliance and conformity. *Annual Review of Psychology*, 55, 591–621.
15. Harvey, N., & Fischer, I. (1997). Taking advice: Accepting help, improving judgment, and sharing responsibility. *Organizational Behavior and Human Decision Processes*, 70(2), 117-133.
16. Bohner, G., & Dickel, N. (2011). Attitudes and attitude change. *Annual Review of Psychology*, 62, 391-417.
17. Johnson, E. J., & Goldstein, D. (2003). Do defaults save lives? *Science*.
18. Madrian, B.C., & Shea, D.F. (2001). The power of suggestion: Inertia in 401(k) participation and savings behavior. *Quarterly Journal of Economics*, 116, 1149–1187
19. Mann, T., & Ward, A. (2007). Attention, self-control, and health behaviors. *Current Directions in Psychological Science*, 16, 280–283.
20. Armel, K. C., Beaumel, A., & Rangel, A. (2008). Biasing simple choices by manipulating relative visual attention. *Judgment and Decision Making*, 3(5), 396-403.
21. Baron, J., & Ritov, I. (1994). Reference points and omission bias. *Organizational Behavior and Human Decision Processes*, 59, 475-475.s
22. Kahneman, D. (2003). A perspective on judgment and choice: mapping bounded rationality. *American Psychologist*, 58(9), 697.
23. Shafir, E., Simonson, I., & Tversky, A. (1993). Reason-based choice. *Cognition*, 49(1-2), 11-36.
24. Johnson, E. J., & Tversky, A. (1983). Affect, generalization, and the perception of risk. *Journal of Personality and Social Psychology*, 45(1), 20.

25. Schwarz, N., & Clore, G. L. (1983). Mood, misattribution, and judgments of well-being: informative and directive functions of affective states. *Journal of Personality and Social Psychology*, 45(3), 513.
26. Finucane, M. L., Alhakami, A., Slovic, P., & Johnson, S. M. (2000). The affect heuristic in judgments of risks and benefits. *Journal of Behavioral Decision Making*, 13(1), 1-17.
27. Alter, A. L., & Oppenheimer, D. M. (2009). Uniting the tribes of fluency to form a metacognitive nation. *Personality and Social Psychology Review*, 13(3), 219-235.
28. Mandel, N., & Johnson, E. J. (2002). When web pages influence choice: Effects of visual primes on experts and novices. *Journal of Consumer Research*, 29(2), 235-245.
29. Deck, C., & Jahedi, S. (2015). The effect of cognitive load on economic decision making: A survey and new experiments. *European Economic Review*, 78, 97-119.
30. Guo, L., Trueblood, J. S., & Diederich, A. (2017). Thinking fast increases framing effects in risky decision making. *Psychological Science*, 28(4), 530-543.
31. Payne, J. W., Payne, J. W., Bettman, J. R., & Johnson, E. J. (1993). *The adaptive decision maker*. Cambridge university press.
32. Lieder, F., & Griffiths, T. L. (2020). Resource-rational analysis: understanding human cognition as the optimal use of limited computational resources. *Behavioral and Brain Sciences*, 43.
33. Lerner, J. S., & Tetlock, P. E. (1999). Accounting for the effects of accountability. *Psychological Bulletin*, 125(2), 255.
34. Birnbaum, M. H., & Jou, J. W. (1990). A theory of comparative response times and “difference” judgments. *Cognitive Psychology*, 22(2), 184-210.
35. Milkman, K. L., Gromet, D., Ho, H., Kay, J. S., Lee, T. W., Pandiloski, P., ... & Duckworth, A. L. (2021). Megastudies improve the impact of applied behavioural science. *Nature*, 600(7889), 478-483.
36. Muralidharan, K., & Niehaus, P. (2017). Experimentation at scale. *Journal of Economic Perspectives*, 31(4), 103-24.

37. DellaVigna, S., & Linos, E. (2020). RCTs to Scale: Comprehensive Evidence from Two Nudge Units. Working Paper.
38. Michie, S., Van Stralen, M. M., & West, R. (2011). The behaviour change wheel: a new method for characterising and designing behaviour change interventions. *Implementation Science*, 6(1), 1-12.
39. Nilsen, P. (2020). Making sense of implementation theories, models, and frameworks. In *Implementation Science* 3.0 (pp. 53-79). Springer, Cham.
40. Keeney, R. L., & Raiffa, H. (1993). *Decisions with multiple objectives: preferences and value trade-offs*. Cambridge university press.
41. Wiecki, T. V., Sofer, I., & Frank, M. J. (2013). HDDM: Hierarchical Bayesian estimation of the drift-diffusion model in Python. *Frontiers in Neuroinformatics*, 7, 14.