ECONtribute
MARKETS & PUBLIC POLICY

ECONtribute
Discussion Paper No. 442

Errors in Aggregation and Adaptive Choice

Malte Kornemann

Philipp Schirmer

November 2026

www.econtribute.de

UNIVERSITÄT BONN

UNIVERSITÄT ZU KÖLN

Funding by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy – EXC 2126/2 – 390838866 – is gratefully acknowledged.

Errors in Aggregation and Adaptive Choice*

Malte Kornemann — University of Bonn
Philipp Schirmer — SAFE

October 29, 2026

Abstract

Most economic decisions require aggregating multiple attributes into an overall evaluation. We develop and experimentally test a theory where aggregation is error-prone and, in response, attributes are weighted optimally given an environment of likely options. We characterize these decision weights, which generically differ from attributes' impact on utility, and the bias they induce in choice. When attributes are independently distributed among likely options, higher-variance attributes are over-weighted, and this bias increases in the level of aggregation errors. We test the theory in an experiment by first manipulating between subjects which attribute varies more in choice options, and then comparing choices between options that are identical across subjects. Consistent with the model's predictions, choice probabilities shift toward the option favored by the high-variance attribute, in our baseline treatments by about 10 percentage points. We also test the effect of aggregation complexity by lowering the number of attributes. Compared to the baseline, this reduces the shift in choice probabilities by about 72%. We apply our model to purchases where base and add-on prices may be displayed separately. A consumer reacts to a short-term change in the display as if it were a price change, while her reaction to a long-run change depends on the details of the environment. Consumer welfare is generally maximized by a fully aggregated price.

Keywords: aggregation errors, cognitive noise, multi-attribute choice, decision weights, complexity, adaptation

JEL Codes: C91, D01, D81, D91

*__Contact__: Malte Kornemann: University of Bonn, malte.kornemann@uni-bonn.de; Philipp Schirmer: Leibniz Institute for Financial Research SAFE, philipp.schirmer@safe-frankfurt.de. **Acknowledgments**: We greatly thank Sarah Auster, Thomas Graeber, Botond Kőszegi, and Florian Zimmermann for valuable feedback and guidance. Funding by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy – EXC 2126/2 – 390838866 and by the Joachim Herz Foundation is gratefully acknowledged. **Ethics approval**: The project received ethics approval from the German Association for Experimental Economic Research (EzmSjnzv). **Preregistration**: The data collection was preregistered at AsPredicted.org (`https://aspredicted.org/333rn8.pdf`).

1 Introduction

Most economic choices are between options whose overall value is an aggregate of multiple attributes. For example, the net utility of consumption goods combines several quality features and a price, often split into a base price and add-ons; the effective return of savings products aggregates gross returns, fees, and tax treatment; and a job offer's attractiveness depends on base pay, bonuses, amenities, and location.

A variety of evidence shows that such aggregation is error-prone, leading to mistakes in choice. Work in cognitive science finds that much of choice variability stems from aggregation errors (e.g., Drugowitsch et al., 2016), cognitive economics classifies them as a key source of noise (Enke, 2026), and measures related to them predict mistakes, for example in lottery choice (Enke and Shubatt, 2026; Oprea, 2024).

Building on work on other cognitive errors (see Oprea and Woodford, 2026, for a review), in this paper we develop and experimentally test a theory of adaptation in response to unavoidable aggregation errors. In our model, the agent uses decision weights on attributes that maximize expected utility given errors and the joint distribution of attributes characterizing likely available options. We prove that optimal decision weights differ from utility weights and thus lead to biased choices. Intuitively, evaluations made with optimal decision weights set likely options further apart, which leaves choices less often decided by errors at the cost of reducing the accuracy of average evaluations. The former consideration matters more the greater the aggregation errors. In the benchmark case with independent attributes, this trade-off results in over-weighting attributes that vary more, and this over-weighting increases in the level of aggregation errors. We provide a characterization for any covariance structure, where the trade-off can lead to more complicated biases.

To test these predictions, we conduct a choice experiment in which product attributes are transparently displayed as integers but have to be aggregated to determine a product's value. In a design inspired by experiments on efficient coding (Frydman and Jin, 2022), we randomize between participants which attribute varies more in initial choices, and compare behavior on identical choices made afterward. Consistent with our theory's predictions, choice probabilities shift significantly toward the option which is better on the attribute that varied more. In addition, we test the effect of a reduction in aggregation errors by removing an attribute that is irrelevant for the identical choices. Compared to

the baseline condition, the difference in choice probabilities is significantly smaller.

We then apply our model to goods whose prices are displayed as several components that need to be aggregated, such as a base price plus a tax or a fee. Adaptation predicts that taxes and small fees receive no weight in decisions. Moving these into the posted price therefore acts like a price increase when decision weights do not adapt to the new display. When decision weights do adapt in response to display changes, they affect behavior more fundamentally. Our model suggests that, even when every component is clearly displayed and correctly understood, consumer welfare is generally highest when fully aggregated prices are shown.

In our model, which we introduce in Section 2, a decision maker (DM) chooses between two multi-attribute goods. The evaluation of a good is a weighted sum of attributes—where the decision weights may differ from utility weights—and is subject to an additive aggregation error that is independent of weights and attributes. The DM chooses the good with the higher evaluation. Adaptation assumes that the decision weights used during choice maximize true expected utility given an expectation of goods to choose between and errors to be made.[1] For tractability, we assume that the joint distribution of attributes characterizing goods and the distribution of aggregation errors are Gaussian.

Section 3 presents the main theoretical results. We prove that adapted weights optimize the trade-off between (i) decision ease, which is the expected differentiation between evaluations that decision weights induce, and (ii) distortion, which is the expected mismatch between the induced evaluation and true utility. The former is valuable since a choice between evaluations that differ more is less often decided by errors. The greater the level of errors, the more valuable is decision ease and the more strongly weights are distorted. This formalizes the logic of a heuristic, in which the optimal decision rule partially departs from utility maximization because it is less susceptible to errors. To our knowledge, ours is the first formalization of this logic from primitives in which errors are not directly influenced by the decision rule.

This trade-off generically results in a bias. When attributes are independently distributed in the expected choice environment, adaptation leads to higher decision weights on attributes that vary more. Intuitively, an attribute that varies more distinguishes the

[1] Adapted weights are optimal given unavoidable errors, but we do not assume that DMs necessarily choose them consciously. Weights may adjust unconsciously to experience without the DM recognizing that her decision rule has come apart from her utility.

goods more in expectation, so over-weighting it affords decision ease. The bias therefore rises in the level of aggregation errors, and for sufficiently large errors, low-variance attributes receive no weight.

We characterize adapted weights for any covariance structure in Section 3.3, which readers primarily interested in the experiment may skip. For any support of attributes, we provide a closed-form solution as a linear combination of two weight vectors. The first captures the motive of increasing decision ease and scales up weights while keeping distortion constant. The second vector shifts weights along the least distorting direction to enable the scaling that increases decision ease. When the level of errors is sufficiently small, every attribute receives weight, and for large errors, only the attribute that covaries most with utility receives weight. For intermediate levels of errors, which attributes receive weight can be more complicated. The proof establishes an equivalence to a well-understood problem, known, for instance, from maximizing a Sharpe ratio under a short-sale constraint.

Section 4 introduces the experiment we design to test the theory's predictions. Participants face a series of one hundred binary multi-attribute choices that consists of two phases. In the initial 'adaptation' phase, consisting of forty choices, we randomly assign participants to one of two alternative choice environments that differ in which of two attributes, A and B, has a high and which a low variance. According to our theory, this induces different adapted weights across the treatment groups. To identify the effects of adaptation, in the second phase of the experiment, we insert 'probe' choices that are identical across treatment groups and exhibit a conflict between attributes A and B. In order to maintain adaptation while eliciting probe choices, we mix these probes with maintenance choices that also reflect the initial environment. To study how adaptation depends on the magnitude of aggregation errors, we separately manipulate aggregation complexity. In our baseline treatment, there is an additional attribute C that varies little in non-probe choices and is (virtually) constant in probe choices. In a cross-randomized treatment, we omit this third attribute, while retaining the relative variance of A and B across treatment groups.

The experimental results, discussed in Section 5, are consistent with our theory of adaptation. We find significant differences in choice probabilities for the identical probes between adaptation treatments. In our baseline, participants for whom A varied more

are about 10 percentage points (pp) more likely to choose the option favored by A than participants for whom B had the higher variance ($p < 0.001$). Relative to the three-attribute baseline, participants in the lower-complexity two-attribute arm perceive aggregating product values as significantly easier and make significantly fewer incorrect choices in the initial phase. Within the lower-complexity treatment, the difference in choice probabilities between adaptation treatments is about 3 pp ($p = 0.017$), and the difference-in-differences of 7.6 pp ($p < 0.001$) suggests that aggregation complexity is a key driver of how previous environments shape on choice probabilities. This is consistent with our theory, which predicts that the level of errors increases the bias.

We then estimate participant-level probit models that directly implement the structural form implied by the theory, allowing us to quantify how adaptation to the choice environment shifts the decision weights on attributes. Even though both attributes are incentivized equally for all participants, the estimated rates of substitution differ significantly between treatments. In the baseline three-attribute setting, the average participant for whom attribute A varied more places 24% more weight on A relative to B, compared to the average participant for whom B varied more ($p < 0.001$). In the lower-complexity treatment, this difference is about 75% smaller ($p = 0.003$ for the difference in logs).

Section 6 applies our model to a specific three-attribute setting capturing add-on pricing. Goods consist of a quality, a base price, and an add-on price such as a junk fee, which is independent of quality and base price, or a tax, which covaries strongly with both. We prove that under adaptation, add-ons like taxes and small junk fees receive no weight. Short-lived changes in the inclusion of add-ons in the posted price, to which weights do not readapt, therefore act like a price change (consistent, for example, with findings by Blake et al., 2021; Chetty et al., 2009). We also consider the long-term effect of display policies that differ in the proportion of the total price that add-ons make up, once weights have adapted to them. Price sensitivity typically falls in the share add-ons make up, but can rise under certain circumstances. Consumer welfare, however, generally falls in that share and is highest when all-in prices are displayed.

We conclude in Section 7 by discussing avenues for further research. All proofs are in Appendix D.

Related literature. We contribute to a growing literature in economics that considers the cognitive foundations of behavior, in particular cognitive noise and errors, as well as adaptation to them (for reviews, see Enke, 2026; Oprea and Woodford, 2026; Woodford, 2020). To the best of our knowledge, we are the first to analyze adaptation in response to errors modeled directly at the aggregation stage, to identify the endogenous trade-off between distortion and susceptibility to errors, and to provide experimental evidence for adaptation in multi-attribute choice.[2]

The most closely related papers are ones that formalize adaptation in response to perception rather than aggregation errors (see Oprea and Woodford, 2026, for a review). In contrast to our trade-off between distortion and decision ease, these analyses follow the broad logic of Bayesian information processing.[3] Moreover, most applications focus on risky choice and probability weighting (Enke and Graeber, 2023; Frydman and Jin, 2022, 2023; Khaw et al., 2021, 2022; Netzer et al., 2025; Steiner and Stewart, 2016), intertemporal choice (Enke et al., 2025), and perception biases such as the illusion of control and over-precision (Gossner and Steiner, 2018), whereas we study multi-attribute choice. Nevertheless, some basic implications, such as a bias toward higher-variance attributes similar to ours, were noted by Woodford (2012). A natural question is how the implications of the two error types differ in our setting in general, which we analyze in Appendix A.2. We show that under some knife-edge conditions, a model of perception errors can match the predictions of a fixed parameterization of ours, although it cannot immediately generate our comparative statics.

The phenomenon of behavioral attenuation documented by Enke et al. (2026) is a response to cognitive uncertainty (Enke and Graeber, 2023), which acts like high-level noise that incorporates errors in aggregation. In models of cognitive uncertainty, an improvement in the true utility (fundamental) of an option increases its choice probability,

[2]In cognitive science, a few papers consider adaptation to aggregation-type errors, but they use other methods, study other kinds of decisions, and provide distinct experimental evidence. Most closely related are biases in averaging tasks, where individuals discount outliers or locally weaker numbers, which can make decisions more accurate in the presence of late-stage errors (Li et al., 2017; Tsetsos et al., 2016); Prat-Carrabin and Woodford (2022, 2026) explain behavior in similar number tasks by efficient coding of the numbers instead. In these tasks individuals do not aggregate different attributes but average identically distributed instances of the same attribute. If attributes in our setting were identically distributed, no bias would arise.

[3]This is true both of models of efficient coding (e.g., Frydman and Jin, 2022, 2023; Woodford, 2012) and of noisy representation with optimal decoding (e.g., Khaw et al., 2021, 2022; Vieider, 2024). Some other models reflect responses to perceptual limitations that account for a winner's-curse-like effect (e.g., Steiner and Stewart, 2016).

albeit in an attenuated way. In our model, that is not the case in general, i.e., reversals rather than attenuation can occur.

Rational inattention (Sims, 2003; see Maćkowiak et al., 2023, for a review) is closely related to models of adaptation, the main difference being the type of information constraint or cost. Flexible information acquisition with the standard Shannon cost about a higher-level object like utility or directly about choice (e.g., Caplin et al., 2022; Matějka and McKay, 2015) would not imply the bias that adaptation to aggregation errors implies. We are not aware of any cost function that captures our mechanism, that aggregation in itself is error-prone, so what individuals aggregate deviates from utility and induces a bias.

In models of sparsity (Gabaix, 2014), a decision maker optimizes attention to attributes given an exogenously given cost of attention. Attention weights act similarly to decision weights; however, this approach has not been applied to our setting of discrete multi-attribute choice. Furthermore, in our model, the cost of increasing the weight on an attribute is an endogenously arising opportunity cost directly linked to aggregation. Without additional assumptions, sparsity would not predict the difference-in-differences effect we experimentally document in our complexity variation.

Besides the evidence for aggregation errors in cognitive science (e.g., Drugowitsch et al., 2016; see also Findling and Wyart, 2021, for a review)[4] and on the computation in lottery choice (Enke and Shubatt, 2026; Oprea, 2024) we mentioned above, several other works in economics analyze complexity of and errors in processes similar to aggregation. Closest to our multi-attribute choice setting, Shubatt and Yang (2026a,b) consider comparison complexity. While these papers, and others on complexity, analyze the form of errors and determinants of complexity within and across choices, ours is a complementary approach, taking errors as given and considering adaptation and the resulting bias.

Finally, we contribute to a large literature on behavioral anomalies in multi-attribute choice. Most theories and evidence point to effects of the choice set (e.g., Bordalo et al., 2013; Bushong et al., 2021; Kőszegi and Szeidl, 2013; Landry and Webb, 2021), whereas our adapted decision weights are independent of the choice at hand. Closest to our model is focusing (Kőszegi and Szeidl, 2013), which predicts that attributes with a larger range in

[4]In cognitive science, such errors are also called inference or computation noise, or late noise, since they arise later in processing than perceptual noise (e.g., Summerfield and Parpart, 2022). Woodford (2020) likewise notes that the relevant noise often arises at later stages of processing.

the present choice set are weighted more strongly. Extending such a model to experiences or memory would share predictions with our independent case.[5] Our theory gives such a model a cognitive foundation in aggregation errors and a connection to stochastic choice. Conversely, adaptation might take place to some degree during choice, with the choice set informing the expected alternatives. Choice-set effects like focusing could then arise from adaptation to aggregation errors.

2 Model

A decision maker (DM) faces a binary[6] choice between two multi-attribute goods and is subject to errors in aggregating attributes into an overall decision variable. Each good has I attributes and is represented by a vector $x \in \mathbb{R}^I$. The DM's utility is linear and separable in attributes, so that $U(x) = \sum_{\mathrm{i}=1}^{\mathrm{I}} x_i$.[7] During choice, the attributes are scaled by decision weights d_i, which are set in an adaptation phase prior to choice. We first describe the choice phase for given weights, then the adaptation phase that determines them.

Choice phase. The DM aims to choose according to the *decision variable* $D(x) = \sum_{\mathrm{i}=1}^{\mathrm{I}} d_i x_i$. During the mental computation of this quantity she makes unavoidable random errors, so her choice is governed by a *choice variable*, the decision variable plus a realized error. Facing goods x and y she chooses according to[8]

$$\sum_{\mathrm{i}=1}^{\mathrm{I}} d_i x_i + \epsilon_x \lesseqgtr \sum_{\mathrm{i}=1}^{\mathrm{I}} d_i y_i + \epsilon_y, \tag{1}$$

where $\epsilon_x, \epsilon_y \overset{iid}{\sim} \mathcal{N}(0, \nu^2)$ are the realized errors. We take the extent of errors as unavoidable and exogenous to the problem we analyze. In particular, we assume errors are independent of the decision weights and of the attributes, $\epsilon_x, \epsilon_y \perp\!\!\!\perp (x, y, d)$.

[5] Bordalo et al. (2020), for example, extend a model of salience as a choice-set effect (Bordalo et al., 2013) this way.

[6] We prove in Appendix B that the theory and the optimal weights directly extend to choice among any number of goods. Nevertheless, we state the main model as a binary choice, since this is the case we implement experimentally, it keeps choice probabilities tractable, and it eases exposition.

[7] Attributes may be any transformation of physical characteristics into utils, so linearity is not itself restrictive. Separability remains a substantial assumption.

[8] Since both goods are ex-ante identically distributed, as we assume below, the good with the higher choice variable also has the higher decision variable in expectation. The rule therefore maximizes the probability of choosing the good with the higher decision variable—the DM's intent in this phase—whether or not she accounts for the errors.

Adaptation phase. Decision weights are adapted prior to choice, that is, they maximize expected utility given an *expected choice environment*. The environment consists of a distribution of goods expected to choose between, the *expected alternatives*, and the *expected level of aggregation errors*, ν^2, during the choice phase.

We assume that limited cognitive resources make error-free aggregation impossible. The level of errors matters only relative to the sum of decision weights and our assumption thus requires the sum to be bounded, which we normalize to 1.[9] Additionally, we restrict weights to be non-negative, so that d lies in the unit simplex Δ^I and $d_i = 1/I$ is the equal-weight benchmark. Optimally adapted weights solve

$$d^* = \arg\max_{d \in \Delta^I} \mathbb{E}_{x,y} \left[U(x)\, \mathbb{P}[\text{Choosing } x \mid d] + U(y)\, \mathbb{P}[\text{Choosing } y \mid d] \right], \tag{2}$$

where the expectation is over alternatives and choice is probabilistic given by (1).

For tractability we assume that goods are drawn independently and the expected alternatives are normally distributed. Hence, they are characterized by their mean μ and a positive definite covariance matrix Σ. It is immediate that the mean is irrelevant to the problem, because both goods are drawn from the same distribution and the mean thus cancels out in any choice.[10]

We consider the case in which expectations about the choice environment are correct. This need not be the case, however, and the weights may be adapted to an environment other than the one actually faced. Relaxing the assumption of correct expectations is straightforward: under a misspecified environment the weights in (2) are unchanged in form but only subjectively optimal. The assumption of correct expectations matters only for our results on choice probabilities in Section 3.4.

A few comments are in order. First, we describe adaptation in terms of expectations, but use the term in its mathematical sense only. While adapted decision weights are optimal from an ex-ante perspective, the DM need not be consciously aware of her decision weights or the environment they adapt to. In some cases, she may knowingly over-weight some attributes because doing so makes the decision easier. In other cases, adaptation

[9] Assuming that the decision weights sum to 1 is a normalization rather than a restriction: for any fixed total K, weights summing to K give the same behavior as weights summing to 1 with ν replaced by ν/K. Were the DM free to scale weights up without limit, she could shrink the effective level of aggregation errors to zero and would always weight equally. Our assumption of limited cognitive resources rules this out but the exact finite bound is irrelevant.

[10] Formally, we can consider $x' = x - \mu$ and $y' = y - \mu$, and for any weights d the choice rule (1) is equivalent when using x' and y' instead of x and y. What the weights adapt to is how much attributes vary, not how large they are.

may arise from a subconscious process that aims at optimizing behavior. For example, decision weights might adjust to experience without the DM necessarily recognizing that her decision rule has come apart from her utility. The environment which weights adapt to can be built from experience with similar choices, from researching the options ahead of time, or from advice or other outside information.[11]

Second, we assume aggregation errors to be independent and exogenous. Hence, our results on adaptation do not rest on assumptions about what makes choices simpler but are a result of the basic fact that aggregation is error-prone in a symmetric, unchanging, and unavoidable way. While the severity of errors may change between choices, abstracting from it keeps errors disciplined and the model tractable. Similarly, the severity of errors also changes across environments. On this we make no assumptions in the theoretical part and characterize decision weights within an environment taking the level of errors as exogenous. In the experimental part, we use a change in the environment, in the number of attributes, as a proxy manipulation of the level of aggregation errors.

Third, we have interpreted the error as imprecision in aggregating each good's attributes separately. Given the simple form of the errors, two other choice rules with appropriately adjusted errors generate the same choice probabilities, and hence the same optimally adapted weights. The DM might instead first consider the differences of weighted attributes across the two goods and then aggregate those differences, with the error entering at that aggregation step.[12] This yields the choice rule

$$\sum_{\mathrm{i}=1}^{\mathrm{I}} d_i(x_i - y_i) + \epsilon \lesseqgtr 0. \tag{3}$$

With the error rescaled to $\epsilon \sim \mathcal{N}(0, 2\nu^2)$, this is equivalent to (1). Choice could also be an imprecise ordinal comparison, where the difficulty of the comparison depends on how far apart the decision variables are. Formally, whenever $\sum_{\mathrm{i}=1}^{\mathrm{I}} d_i x_i \neq \sum_{\mathrm{i}=1}^{\mathrm{I}} d_i y_i$, (1) is equivalent to

$$\operatorname{sgn}\left(\sum_{\mathrm{i}=1}^{\mathrm{I}} d_i x_i - \sum_{\mathrm{i}=1}^{\mathrm{I}} d_i y_i\right) + \frac{\sqrt{2}\nu}{\left|\sum_{\mathrm{i}=1}^{\mathrm{I}} d_i x_i - \sum_{\mathrm{i}=1}^{\mathrm{I}} d_i y_i\right|}\epsilon \lesseqgtr 0 \tag{4}$$

[11] Decision weights could also arise from bottom-up cognition not necessarily designed to be optimal. Our theory shows that, even though not designed to, certain weights and the associated biases can still perform better than utility weights would under aggregation errors.

[12] It is also possible that errors occur when taking the differences, itself a form of aggregation, or both when taking the differences and when aggregating them—with the total of the errors then being the ϵ below. Mezzadri and Woodford (2026) find evidence for such attribute-wise processing when participants judge which of two number sequences has the larger sum. In their models the noise may enter in each attribute-wise difference or only when the differences are aggregated.

for $\epsilon \sim \mathcal{N}(0,1)$. Here the DM compares the two goods on their weighted attributes and picks the better one. The comparison is noisy, and the noise scales inversely with the gap between decision variables. This is the greatly simplified version of the model by Shubatt and Yang (2026a) that we mentioned earlier.

As long as errors take such simple forms, the adaptation we characterize and our results therefore hold across models of good-wise aggregation noise, noisy aggregation of attribute-wise differences, and noisy comparison.

3 Adapted Weights

We show that optimally adapted weights trade off two forces, decision ease and distortion. To build intuition for both and illustrate why decision weights differ from utility weights, we first consider the following example. To keep it as simple as possible, we work in a minimalistic discrete setting.

Suppose errors take values in $\{-1, 0, 1\}$, so that the errors of the two goods combine to favor one of them strongly, by 2, or weakly, by 1, or cancel out. Consider two attributes, the first taking the values 0 or 3 and the second 0 or 1.5. A decision rule following utility places weight $1/2$ on each attribute, so the first attribute moves the decision variable by 0 or 1.5 and the second by 0 or 0.75. When the attributes favor different goods, the decision variables differ by 0.75, and the choice is dictated by the errors unless they cancel out. Now suppose instead that the DM places all weight on the first attribute, so that it moves the decision variable by 0 or 3. Whenever the attributes favor different goods, the good with the larger first attribute is chosen regardless of the errors. This force is decision ease. In the first case the DM decides between an option evaluated at 1.5 and one evaluated at 0.75; in the second, between one evaluated at 3 and one at 0. In arriving at these evaluations the DM makes the same potential errors, but the decision in the latter case is easier because the evaluations differ more, and thus choice is less often overturned. Importantly, decision weights that increase decision ease do not necessarily favor the utility-maximizing good more, even though they do here. Whether decisions are aligned with the true utility-maximizing good is captured by a second, in principle independent, force: distortion. When all weight is placed on the first attribute, the decision variable no longer matches utility. In particular, when both goods have the

same first attribute but differ in the second, this rule chooses at random. Even when the errors cancel out, it can therefore select the inferior good. Assuming that all error and attribute realizations are equally likely, the trade-off favors this extreme decision rule over following the utility weights in this example.

3.1 Decision Ease and Distortion

We now show that formalizations of these two forces are exactly what optimally adapted weights trade off in general. How well decision weights track true utility in expectation is measured by the correlation between the induced decision variable $D(x)$ and utility $U(x)$. We thus define the *distortion* of a decision rule as the shortfall of this correlation from one, $1 - \mathrm{Corr}(U(x), D(x))$.

Decision ease is determined by how far apart the goods' evaluations are in expectation, since this is what makes choice less susceptible to errors. It is not a measure of whether the decision variables favor the utility-maximizing good, only of how far apart the evaluations are and thus how little errors influence choice. We define it to be $\mathrm{Var}(D(x))$, which is a monotone transformation of the expectation $\mathbb{E}|\sum_{\mathrm{i}=1}^{\mathrm{I}} d_i x_i - \sum_{\mathrm{i}=1}^{\mathrm{I}} d_i y_i|$.[13]

Optimizing decision weights reduces to trading off these two forces:

Proposition 1 (Reduction to Decision Ease and Distortion)**.** *The objective* (2) *of optimizing decision weights is an increasing affine function of*

$$\sqrt{\mathrm{Var}(U(x))} \times \mathrm{Corr}\left(U(x),\ D(x)\right) \times \sqrt{\frac{\mathrm{Var}\left(D(x)\right)}{\mathrm{Var}\left(D(x)\right) + \nu^2}}. \tag{5}$$

The first factor is the stakes of the expected choice environment, which depends on the expected alternatives alone. The second factor is one minus distortion, and the third is decision ease relative to the level of aggregation errors. Decision weights matter only through these two factors.

Optimally adapted weights therefore lie on the efficient frontier of minimizing distortion and maximizing decision ease. Utility weights induce no distortion and are one endpoint of the frontier, but other decision weights may have more decision ease and could be optimal, as illustrated by the example. The frontier itself is determined by the expected alternatives.

[13] Because the goods are drawn independently, $\mathrm{Var}(\sum_{\mathrm{i}=1}^{\mathrm{I}} d_i x_i - \sum_{\mathrm{i}=1}^{\mathrm{I}} d_i y_i) = 2\,\mathrm{Var}(\sum_{\mathrm{i}=1}^{\mathrm{I}} d_i x_i)$, and under normality $\mathbb{E}|\sum_{\mathrm{i}=1}^{\mathrm{I}} d_i x_i - \sum_{\mathrm{i}=1}^{\mathrm{I}} d_i y_i| = 2\sqrt{\mathrm{Var}(\sum_{\mathrm{i}=1}^{\mathrm{I}} d_i x_i)/\pi}$, so the variance of the decision variable is, up to a monotone transformation, the expected gap between the two decision variables.

How the two forces are traded off depends on the level of aggregation errors. Decision ease has no value without errors but its value rises with the level of errors. The larger the errors, therefore, the more adaptation favors decision ease at the cost of distortion. Formally, we have the following.

Proposition 2 (Decision Ease and Distortion Trade-off)**.** *For any expected alternatives, both the decision ease* $\mathrm{Var}(D(x))$ *and the distortion* $1 - \mathrm{Corr}(U(x), D(x))$ *of optimally adapted weights are weakly increasing in* ν*.*

3.2 Independent Attributes

This trade-off leads to over-weighting higher-variance attributes when the attributes of the expected alternatives are independent. To state the result formally, we assume distinct variances ordered $\sigma_1^2 > \cdots > \sigma_I^2$.[14]

Proposition 3 (Optimal Weights for Independent Attributes)**.** *The optimally adapted weights are*

$$d_i^* = \frac{1}{J} + \frac{\nu^2}{J} \sum_{j=1}^{J} \left(\sigma_j^{-2} - \sigma_i^{-2}\right) \tag{6}$$

for $i \leq J$*, and* $d_i^* = 0$ *for* $i > J$ *where* $J \in \{1, \ldots, I\}$ *is the largest integer for which the weights given by* (6) *are non-negative.*

The weights decompose into the equal-weight benchmark $1/J$ over the attributes that receive weight and a term proportional to ν^2 that governs over- and under-weighting. The term is positive for the attributes with sufficiently high variances (implying low σ_i^{-2}) and negative for low-variance ones. On the intensive margin, the weights of any two attributes that receive weight differ by $d_i^* - d_j^* = \nu^2(\sigma_j^{-2} - \sigma_i^{-2})$, so they move apart as the errors grow, and they are further apart the more the variances differ. On the extensive margin, attributes drop out in order of their variance, lowest first, until only the highest-variance attribute receives weight.

The intuition is that a higher-variance attribute differs more between goods on average, so weighting it more heavily widens the expected gap between decision variables and

[14]Distinct variances are assumed only to keep the statement clean. With ties the closed form below has to account for ties in the index J, since tied attributes at the boundary enter or leave together. Apart from this adjustment of the index, the proposition and the closed form extend to that case. Given distinct variances, assuming this order is without loss of generality.

increases decision ease. This departure from utility weights distorts the decision variable, but for any positive level of aggregation errors some over-weighting is worth it. How much depends on the level of errors. Larger errors mean that, all else equal, choice is governed more by the errors and less by the decision variables. The cost of distorting the decision variable therefore falls relative to the gain from decision ease, and adaptation moves further toward ease, as in Proposition 2. In the independent case this results in more uneven weights the larger the errors.

3.3 The General Case

We now characterize adapted weights for any covariance structure of attributes. The trade-off between decision ease and distortion still determines the optimally adapted weights, but how it resolves is more complicated. The predictions for the experiment, in Section 3.4, rest on the independent case alone, and readers mainly interested in the experiment may turn to them directly.

We state the characterization in vector notation to condense formulas. Let $\iota \in \mathbb{R}^I$ be the vector of ones and e_i the i-th standard basis vector. Attribute i is in the support of a weight vector d if $d_i > 0$. A generic support is denoted $\mathcal{S} \subseteq \{1, \ldots, I\}$, and a subvector or submatrix carrying $\mathcal{S}$ as a subscript is the one with elements indexed by the attributes in $\mathcal{S}$.

Theorem 1 (Characterization of Optimal Weights)**.** ***A.****(Closed Form within a Support) Let $\mathcal{S}$ be the support of the optimally adapted weights at ν. Then $d_i^* = 0$ for $i \notin \mathcal{S}$ and*

$$d_{\mathcal{S}}^* = \alpha_{\mathcal{S}} + \nu^2 \gamma_{\mathcal{S}} (\alpha_{\mathcal{S}} - \beta_{\mathcal{S}}), \tag{7}$$

$$\textit{where} \quad \alpha_{\mathcal{S}} = \frac{\Sigma_{\mathcal{S}\mathcal{S}}^{-1}(\Sigma\iota)_{\mathcal{S}}}{\iota_{\mathcal{S}}^\top \Sigma_{\mathcal{S}\mathcal{S}}^{-1}(\Sigma\iota)_{\mathcal{S}}} \in \mathbb{R}^{|\mathcal{S}|}, \quad \beta_{\mathcal{S}} = \frac{\Sigma_{\mathcal{S}\mathcal{S}}^{-1}\iota_{\mathcal{S}}}{\iota_{\mathcal{S}}^\top \Sigma_{\mathcal{S}\mathcal{S}}^{-1}\iota_{\mathcal{S}}} \in \mathbb{R}^{|\mathcal{S}|}, \quad \gamma_{\mathcal{S}} = \iota_{\mathcal{S}}^\top \Sigma_{\mathcal{S}\mathcal{S}}^{-1}\iota_{\mathcal{S}} \in \mathbb{R}\,.$$

B.*(Global Behavior) The support of optimally adapted weights is unique. As a function of ν^2, the support changes a finite number of times, optimal weights are continuous at the points of changes, and smooth elsewhere. There is a threshold $\underline{\nu} > 0$ below which the support is full, and a threshold $\overline{\nu} < \infty$ above which the weights are constant and supported within* $\arg\max_i \iota^\top \Sigma e_i$.

To unpack the theorem, first consider Part A when all attributes receive weight. Then $\alpha_{\mathcal{S}}$ is the vector of equal weights, as $\Sigma^{-1}\Sigma\iota = \iota$, and the first term of (7) is the equal-weight benchmark. As in the independent case, the second term shifts weight relative

to this benchmark. Scaling all weights up along $\alpha_{\mathcal{S}}$ widens the expected gap between decision variables while all weights remain equal. This increases decision ease and keeps distortion minimal. Scaling up alone is infeasible, since over-weighting some attributes requires under-weighting others. The vector $\beta_{\mathcal{S}}$ captures the optimal way of resolving this conflict. It is the direction along which the decision variable varies least across realizations.[15] Thus subtracting weights in this proportion keeps the decision variable closest to the one induced by $\alpha_{\mathcal{S}}$ and pays for the increase in decision ease with the least distortion. For independent attributes, $\beta_{\mathcal{S}}$ weights attributes in proportion to their inverse variances, and subtracting it indeed simplifies (7) to (6).

This reweighting is scaled by $\nu^2\gamma_{\mathcal{S}}$. It rises in ν, as before, since larger aggregation errors make decision ease more valuable. It also rises in $\gamma_{\mathcal{S}}$, which captures how much weight is freed up by moving along $\beta_{\mathcal{S}}$ per unit of change in decision variables. The greater $\gamma_{\mathcal{S}}$ is, the less distortion needs to be paid for decision ease, and the more attractive reweighting is.

When some attributes receive no weight, the interpretation of (7) is unchanged but now conditional on the support. The weighting that best aligns the decision variable with utility is given by $\alpha_{\mathcal{S}}$, and $\beta_{\mathcal{S}}$ is again the direction along which reducing weights induces the least distortion.

Part B describes how changes in the support shape the path of the weights in ν. The support is unique and, by Part A, so are the weights. There are finitely many support changes, and weights are continuous at these points. For sufficiently small aggregation errors every attribute receives weight. For sufficiently large ones, only attributes whose covariance with utility is largest, generically a single one, receive weight. Beyond that point the weights no longer move. Since weights are continuous, an attribute exits the support when its weight in (7) hits 0. Yet, unlike in the independent case, attributes can also re-enter, and the order in which they would drop out changes from one support to the next. Tracking the support in the intermediate cases is therefore involved.

The key technical step of the proof is the decomposition of Proposition 1, which can be rewritten to turn the problem into a well-understood form that is shared for example with the maximization of a Sharpe ratio (Sharpe, 1966).[16]

[15] Formally, $\beta_{\mathcal{S}}$ solves $\min\{d_{\mathcal{S}}^\top \Sigma_{\mathcal{S}\mathcal{S}} d_{\mathcal{S}} : \iota_{\mathcal{S}}^\top d_{\mathcal{S}} = 1\}$, and $\gamma_{\mathcal{S}}^{-1}$ is the minimum. Only the adding-up constraint is imposed; $\beta_{\mathcal{S}}$ is a direction along which weights are adjusted, not a weighting the DM uses, and it may have negative coordinates.

[16] In the Sharpe-ratio problem, Part A is the two-fund separation result (Merton, 1972; Tobin, 1958),

The bias induced by adaptation follows a simple rule as long as aggregation errors are small enough that every attribute receives weight. Then, as in the independent case, a fixed ranking of the attributes holds and the bias weakly increases in the level of aggregation errors.

Corollary 1 (Bias in General Environments)**.** *For any expected alternatives there is a* $\underline{\nu} > 0$ *and an order of attributes, wlog the same as the index, such that for all* $\nu \in [0, \underline{\nu})$*,* $d_1^* \geq d_2^* \geq \ldots \geq d_I^*$*. Moreover, for any* $i < j$*,* $d_i^* - d_j^*$ *is weakly increasing in* $\nu \in [0, \underline{\nu})$*.*

Beyond this range the bias need not follow such a simple rule. As long as the same attributes receive weight, the weights move in one direction as ν grows, but when an attribute drops out of or re-enters the support, the direction can change with it, and larger aggregation errors can leave the weights strictly less uneven than smaller ones.

3.4 Implications for Choice

Given the adapted weights, the choice between two goods x and y remains random, because the aggregation errors are realized only at the choice phase. By (1),

$$\mathbb{P}[\text{choosing } x] = \Phi\left(\sum_{\mathrm{i}=1}^{\mathrm{I}} \frac{d_i^*(x_i - y_i)}{\sqrt{2}\nu}\right) \quad \text{and} \quad \mathbb{P}[\text{choosing } y] = 1 - \mathbb{P}[\text{choosing } x]. \tag{8}$$

Here, choice inherits a bias from uneven decision weights. Choice probabilities respond more to attribute i than to attribute j if and only if $d_i^* > d_j^*$, whereas utility responds to both equally. The bias is quantified by the rate of substitution d_i^*/d_j^* at which the utility from attribute j substitutes for utility from i while leaving choice probabilities unchanged. Importantly, as a feature of adaptation, the bias is determined by the expected choice environment, not by the choice set. The same two goods are chosen with different probabilities by DMs whose weights are adapted to different environments. Moreover, since the bias is induced by adaptation to aggregation errors, it occurs even when attributes are clearly displayed and correctly understood.

We note that adaptation and the induced bias mitigate larger aggregation errors but cannot undo them. Holding the expected alternatives fixed, both expected utility and the probability of a correct, that is utility-maximizing, choice strictly fall in the level of aggregation errors.

with $\alpha_{\mathcal{S}}$ the tangency and $\beta_{\mathcal{S}}$ the minimum-variance portfolio, and the non-negativity of weights is a short-sale constraint. The constrained problem is well studied (Dybvig, 1984; Markowitz, 1956), and the difficulty of tracking the support is similar, with no closed form for the sequence of supports known.

Proposition 4 (Aggregation Errors and Utility)**.** *For any expected alternatives, expected utility and the probability of choosing according to utility are strictly decreasing in* ν*.*

By contrast, the probability that the DM chooses according to her aim at the choice phase, the decision variable, may rise.

Remark 1. The probability that a DM chooses according to her decision variable can be increasing in the level of aggregation errors ν.

Intuitively, adaptation to larger errors can widen the gap between decision variables by enough to more than offset the larger errors, so greater errors can leave choices more in line with the decision variable. This requires specific circumstances, but is already possible in the independent case with more than two attributes and in the two-attribute case with negative correlation.[17]

Our experiment compares choice probabilities and we estimate decision weights on the same choices, displayed the same way, across different environments. To derive the formal predictions, we consider two conditions that both feature independent attributes. In one condition, denoted *A-HighVar*, attribute A varies the most and B the least while in the other condition, denoted *B-HighVar*, variances of A and B are swapped. Adaptation thus induces a bias toward A in *A-HighVar* and toward B in *B-HighVar*, which increases in the level of aggregation errors.

Corollary 2 (Predictions in Decision Weights)**.** *Fix a choice between two goods and assume a positive level of aggregation errors, the same across conditions and in a range where both attributes receive weight.*

A.*(Cross-Environment) The bias toward attribute A is higher in A-HighVar than in B-HighVar:*

$$(d_A^*/d_B^*)_{A\text{-}HighVar} > (d_A^*/d_B^*)_{B\text{-}HighVar}.$$

[17]This bears on the cognitive uncertainty measure proposed by Enke and Graeber (2023). Consider the subconscious interpretation of adaptation, and interpret it (very) strictly: at the choice phase the DM believes that choosing by her decision variable is choosing well. Asking her how confident she is in making a correct choice is then asking for the probability that she chose according to her decision variable—which can rise as her cognitive noise rises. The possibility requires a DM entirely unaware of her own objective at the choice phase, and a particular expected choice environment besides. It is therefore not a general critique of the measure, and whether such concerns bite is an empirical question for future work. What the remark shows is a general theoretical possibility: a decision rule can optimally adapt to cognitive errors by simplifying enough that following the simplified rule is associated with a lower cognitive uncertainty.

B.*(Cross-Error) The bias toward attribute A strictly increases in the level of aggregation errors in A-HighVar and strictly decreases in B-HighVar. Hence,*

$$(d_A^*/d_B^*)_{A\text{-}HighVar} - (d_A^*/d_B^*)_{B\text{-}HighVar}$$

is strictly increasing in the level of aggregation errors.

To test these predictions more directly from behavior, we consider the choice probabilities predicted by our model. This requires choice problems in which the options differ only in their values for attributes A and B and there is a conflict, in that one good has a higher A and the other a higher B value. We call the one with a higher A value A-favored. In this case:

Proposition 5 (Predictions in Choice Probabilities)**.** *Fix a choice between two such goods and assume a positive level of aggregation errors, the same across conditions and in a range where both attributes receive weight.*
A.*(Cross-Environment) The A-favored good is chosen more often in A-HighVar than in B-HighVar:*

$$\mathbb{P}[\textit{choose A-favored} \mid \textit{A-HighVar}] > \mathbb{P}[\textit{choose A-favored} \mid \textit{B-HighVar}].$$

B.*(Cross-Error) The difference in choice probabilities of the A-favored option*

$$\mathbb{P}[\textit{choose A-favored} \mid \textit{A-HighVar}] - \mathbb{P}[\textit{choose A-favored} \mid \textit{B-HighVar}]$$

is strictly increasing in the level of aggregation errors.

4 Experimental Design

In this section, we present an experiment designed to study whether decision weights adapt to the distribution of attribute values and whether this adaptation is stronger when aggregation is more complex. Directly testing the theory requires several features that are challenging to achieve in field data. In particular, testing the theory requires (i) exogenous (and observable) variation in the choice environment to induce adaptation, (ii) observing choice behavior of individuals on identical decisions to cleanly distinguish adaptation from context effects, (iii) ensuring the identical decisions are observed within an otherwise maintained choice environment, and (iv) an exogenous shift in the level of aggregation errors.

Our experiment is designed to match these requirements and provides a clean, proof-of-concept test for the theory's predictions. Screenshots of the full experimental instructions are in Appendix F.

4.1 Setting and Timeline

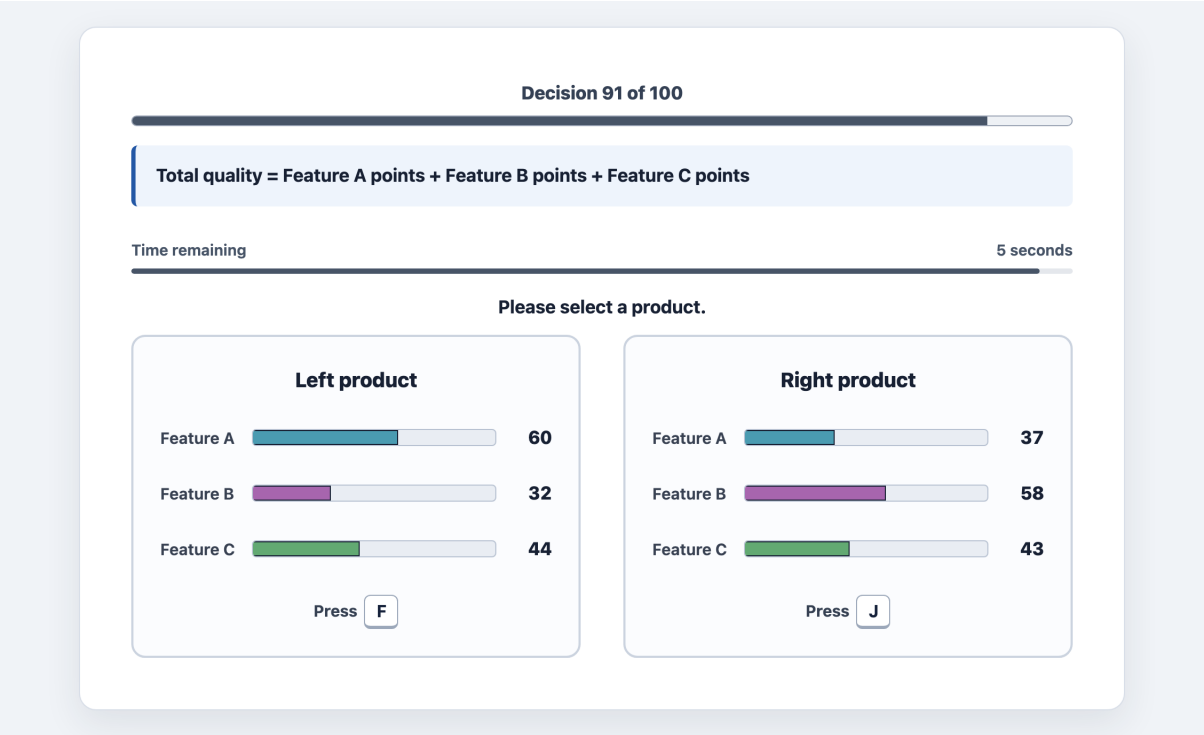
The main experiment consists of 100 binary choices between two multi-attribute products shown side by side. Figure 1 shows a screenshot for one of the choices. Each of a product's features is an integer from 0 to 100, presented both as a number and as a colored bar. The total value of a product is the equally weighted sum of its features. For one in ten participants, one choice is drawn at random and paid at \$0.10 per point of the chosen product's total value. Hence the induced utility over products is given by the equally weighted sum of the features. To reduce motor noise from misclicks with the mouse, participants choose the left product by pressing 'F' (conveniently located for the left index finger) and the right product by pressing 'J' (conveniently located for the right index finger). We discuss the design of these choices in detail below.

Participants first provide informed consent and pass a short video-based attention screener before receiving detailed instructions about the main task.

Instructions and comprehension. The instructions introduce participants in detail to the choices of the main experiment. Participants are taught about the products' features, learn that a product's total value (or 'quality') is the equally weighted sum of all its features, and see examples of the choices they will face. They also learn how choices are incentivized, that they have five seconds to respond to each choice, and that they select products by pressing 'F' or 'J'. Participants have to press both keys before proceeding to the comprehension questions, which ensures that they have access to a keyboard. In four comprehension questions, we ensure that participants understand these instructions. Participants have two attempts to answer all four comprehension questions correctly; otherwise they cannot proceed to the main part of the experiment.

Main experiment. After answering the comprehension questions correctly, participants enter the main part of the experiment, in which they make 100 binary multi-attribute choices. Figure 2 shows its timeline. We distinguish two phases: an initial

Figure 1: Example of a main-task decision screen ($K = 3$)



Notes: Screenshot of a decision as encountered during the experiment. Participants choose the product with the higher total quality by pressing F (left) or J (right) within five seconds. The depicted decision is a common probe, featuring a typical contrast between A and B features across products. The probe decision's total gap is $G = \Delta A + \Delta B + \Delta C = 23 - 26 + 1 = -2$ ($G = -3$ for $K = 2$) and width $W = |\Delta A| + |\Delta B| + |\Delta C| = 23 + 26 + 1 = 50$ ($W = 49$ for $K = 2$). The right product has a higher total value at $U_R = 37 + 58 + 43 = 138$ vs $U_L = 60 + 32 + 44 = 136$.

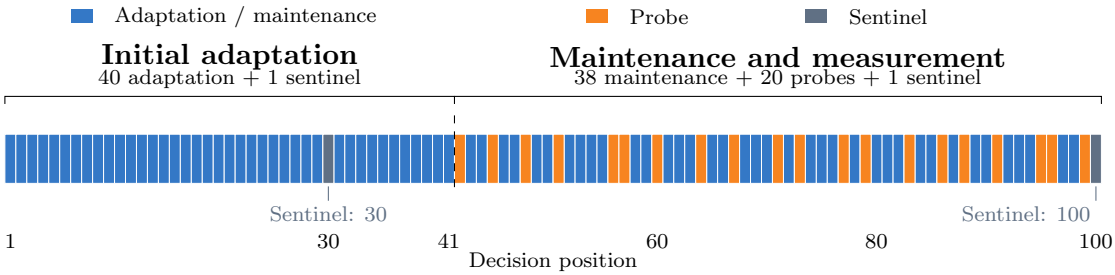
adaptation phase and a subsequent measurement and maintenance phase.

In the initial adaptation phase, participants make 40 *adaptation choices*. These choices differ across the environment treatments and expose participants to the distribution of product attributes in their assigned environment before measurement begins.

In the maintenance and measurement phase, participants make 20 *probe choices*, which are the same across the environment treatments. Our analysis focuses on these choices, allowing us to compare how participants choose between the same products after experiencing different environments. To avoid inducing a strong and discrete shift in the choice environment, we mix the probes with 38 further adaptation choices, the *maintenance choices*. Two additional sentinel choices, at positions 30 and 100, bring the total to 100 decisions. We describe these below.

Treatment conditions. We employ a 2×2 between-subjects design (Table 1). In the *Baseline* condition (*K3*), each product has three features (A, B, C) while the reduced-

Figure 2: Timeline of the 100 main-experiment decisions



Notes: Each tile represents one decision. Blue tiles indicate choices drawn from the assigned distribution depending on the treatment condition; orange tiles indicate probe decisions, which are identical across treatment conditions. Gray tiles indicate seeded sentinel decisions. The sequence of adaptation, probe, and sentinel decisions is held constant across participants. Which adaptation or probe decision is assigned to which location is randomized per participant.

complexity treatment (*K2*) omits feature C and each product has two features (A, B). In addition, we vary the variance of features A and B in the adaptation choices to induce different decision weights across treatments. Either A has a high variance and B a low one (treatment group *A-HighVar*), or the reverse (*B-HighVar*). In the baseline, feature C always has the low variance. Participants are not told about the distributions they encounter, instead learning about them by experience as they make decisions.[18]

Table 1: The 2×2 experimental design

	Aggregation complexity	
Environment	Features A, B, C (*K3*)	Features A, B (*K2*)
A-HighVar	*K3*, *A-HighVar*	*K2*, *A-HighVar*
B-HighVar	*K3*, *B-HighVar*	*K2*, *B-HighVar*

Notes: This table illustrates the factorial treatment design, described in detail in the main text.

Pre-registered exclusion criteria. As pre-registered, participants who fail the video-based attention screener or fail to answer the comprehension questions correctly within two attempts cannot complete the survey. We exclude from the main analyses choices in which participants time out. In the rare cases in which a participant times out on five or more choices, we drop the participant from the analysis altogether. In addition, we exclude participants whose average response time (excluding timeouts) is below 500 ms.

[18] Describing the distribution would likely have only a minor effect on adaptation (cf. Frydman and Jin, 2025).

4.2 Choice Design

We design the 100 choices to be fully symmetric across the two environment treatments *A-HighVar* and *B-HighVar*, with the set of adaptation choices differing only in that the A and B values are swapped. Across the complexity treatments *K3* and *K2*, the choices are the same except that feature C is omitted.

Adaptation choices. The adaptation choices were generated quasi-randomly by a computer ahead of the experiment to represent the distribution to which the weights adapt. In the *A-HighVar* baseline condition, the features are drawn from truncated independent normal distributions with means $(50, 50, 50)$ and standard deviations $(\sigma_A, \sigma_B, \sigma_C) \approx (20, 10, 10)$. In *B-HighVar*, obtained by relabeling, the means are the same but the standard deviations are $(\sigma_A, \sigma_B, \sigma_C) \approx (10, 20, 10)$. In the reduced-complexity treatments, the distributions are the same, restricted to features A and B.[19]

Probe choices. The choices we use to compare treatments, the probes, are the same in *A-HighVar* and *B-HighVar* and are designed to measure adaptation effectively and cleanly. Each probe exhibits a conflict between the two features A and B, in that one product has the higher value of A and the other the higher value of B. An example of a probe decision is shown in Figure 1. In addition, the probes have little to no variation in C in the baseline treatment ($|x_C - y_C| \leq 1$), which keeps them comparable with the reduced-complexity treatment.[20]

Across the 20 probes, we systematically vary the 'width' of the contrast, $W := |x_A - y_A| + |x_B - y_B|$, and the total value 'gap' between the products, $G := x_A - y_A + x_B - y_B$. Probes are designed factorially to reflect all pairings of a gap in $\{\pm 3, \pm 9, \pm 15\}$ and width in $\{41, 49, 57\}$ with the instances ± 9 and 49 appearing twice. For each probe, we also add its mirror pair, in which the A and B values are swapped (and hence the sign of the gap inverted).

[19] Appendix Figures 7 and 8 show the realized marginal distributions and between-feature correlations, respectively. The adaptation decisions closely match the targeted normal distributions in their marginal distribution and uncorrelatedness of features. None of the 30 distinct off-diagonal correlations is significant in nominal two-sided Pearson tests (all $p \geq 0.105$).

[20] The difference in C is never pivotal in any probe. We do not set it to zero in all probes so that the probes do not overly stand out.

Sentinel choices. We include two 'sentinel' choices, at positions 30 and 100. Both are dominated choices, in which one product has a higher value for each feature A, B, and C. The sentinel at position 30 differs across the environment treatments, with larger dominance in the targeted feature, $(\Delta A, \Delta B, \Delta C) = (20, 10, 10)$ in *A-HighVar* and $(10, 20, 10)$ in *B-HighVar*. The sentinel at position 100 is symmetric, with $(\Delta A, \Delta B, \Delta C) = (10, 10, 10)$ in both groups. Choosing the dominated product for these choices constitutes an obvious and symmetric error, allowing us to identify potentially inattentive participants.[21]

Cross-randomizations. In addition to the central 2×2 treatment randomization, we cross-randomize which color (blue, teal, or violet) each feature is assigned and the exact placement of the adaptation and probe choices within their respective categories.[22] Each of the slots 1 to 100 in Figure 2 has a fixed orientation that determines whether the product with the higher total value is placed on the left or the right; orientations are balanced within the adaptation and probe choices and within both halves of the probes. Because choices are assigned to slots at random, whether the higher-value product of a given choice appears on the left or the right is also random.

Covariates and process data. After making their 100 choices in the main part of the experiment, participants complete a short questionnaire at the end of the survey. They are first asked to describe their approach to making the choices in an open-ended text question. They then report whether their approach changed during the main task (and if so, explain how in a short open-ended text). On Likert scales from 1 to 7, participants report which of features A and B they perceived to vary more across choices, and how difficult they perceived it to be to determine the total value of a product. Finally, participants report their subjective numeracy on a short three-item scale taken from McNaughton et al. (2015).

[21] We opted against pre-registering strict filters based on sentinel choices because even generally attentive participants might make mistakes on these choices.

[22] To place the probes throughout the experiment, we split the last 60 choices into 20 "packets" of three choices, each containing one probe and two maintenance choices drawn from the same *A-HighVar* or *B-HighVar* environment as the first 40 adaptation choices. The placement of the 20 probe slots is the same for all participants. Between participants, we randomize which probe is assigned to which probe slot, subject to two constraints. If a probe is assigned to one of the first 10 slots, its A–B mirror is assigned to one of slots 11 to 20, and the mirror of the probe in slot 10 cannot be assigned to slot 11, so that mirror pairs never directly follow each other in the probe order.

5 Experimental Results

We conduct the experiment in September 2026 on the online platform Prolific with participants residing in the US. Participants take on average 15 minutes to complete the experiment and receive a fixed show-up fee of £2.25 upon completion. One out of ten participants is randomly selected for a bonus payment, which averages £9.97 among recipients.[23]

We collect 654 complete responses in total. After applying the pre-registered exclusions, we retain $n = 635$ valid respondents who constitute our main analysis sample.[24] As shown in Appendix Table 3, demographic characteristics are overall well balanced within both *A-HighVar* and *B-HighVar*, as well as between the two complexity conditions. Unless noted otherwise, we do not include demographic controls in our main regression specifications.

We begin by analyzing treatment effects in reduced form, first for adaptation in the *K3* baseline setting and then for differences in adaptation effect between *K3* and *K2*. We then present structural evidence by estimating participants' relative decision weights on the products' features.

5.1 Reduced-Form Evidence

We first investigate participants' behavior in the main task of the baseline condition *K3*, where each product has three features A, B, and C. The average (median) baseline participant spends ca. 4.60 (4.49) minutes on the main task and takes around 1.96 (1.92) seconds for each decision.

5.1.1 Adaptation to Aggregation Errors

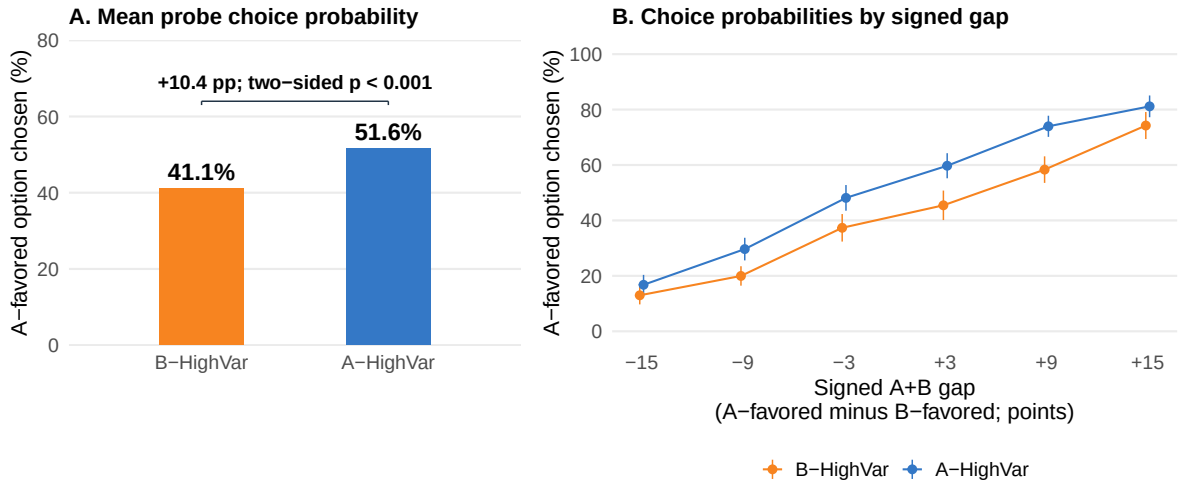
Adaptation choices. The initial adaptation choices allow participants to experience the distribution of the features A, B, and C in their assigned environment. We observe that participants generally select the higher-total-value option, doing so in 84.6% of

[23] The UK-based platform Prolific pays its participants in British pounds.

[24] We exclude 14 participants who time out on five or more decisions as well as five participants whose average response time is below 500 ms on non-timeout responses. Attrition after starting the main task is low, with 98.2% of participants completing it. We find no statistically significant differences across the four treatment conditions in retention in the final analysis sample, including the preregistered exclusions (Pearson's $\chi^2(3) = 1.76$, $p = 0.624$).

decisions in the *A-HighVar* condition and 83.0% in the *B-HighVar* condition ($p = 0.073$ for the difference).

Figure 3: Adaptation in the three-variable baseline condition *K3*



Notes: Panel A reports the mean share of choosing the A-favored option, defined as the option with the higher value of attribute A, by treatment within the *K3* baseline condition. The bracket reports the *A-HighVar* minus *B-HighVar* difference in percentage points and a two-sided p-value. Panel B reports choice shares by signed A+B gap, defined as the A+B total of the A-favored option minus that of the B-favored option, excluding C. Positive gaps favor the A-favored option on A+B value. Probe widths are pooled within each participant and gap before averaging across participants. Error bars indicate 95% confidence intervals.

Probe choices. To test for adaptation, we turn to participants' choices on the 20 common probes, which are identical across *A-HighVar* and *B-HighVar*. Recall that each probe involves a trade-off between features A and B: the option with the higher value of A has the lower value of B, and vice versa. We refer to these options as *A-favored* and *B-favored*, respectively.

Let $A\text{-}favored_{i,p}$ be an indicator equal to one if participant i selects the A-favored option on probe p, and zero otherwise. For the *K3* baseline, we estimate

$$A\text{-}favored_{i,p} = \alpha + \beta_1 A\text{-}HighVar_i + \varepsilon_{i,p}, \tag{9}$$

where $A\text{-}HighVar_i$ indicates assignment to *A-HighVar*, with *B-HighVar* as the reference group. The coefficient β_1 measures the difference between the two treatment groups in the probability of choosing the A-favored option. Proposition 5, part A, predicts $\beta_1 > 0$: participants whose weights are adapted to higher variation in A should select the A-favored option more often on otherwise identical choice problems.

Figure 3 reports the conditional choice probabilities between *A-HighVar* and *B-HighVar* on the 20 probe choices. We find strong evidence for adaptation: the *A-HighVar* group selects the A-favored option for 51.6% of the probes, while the *B-HighVar* group selects this option only 41.1% of the time and hence about 10.4 percentage points (pp) less often (two-sided $p < 0.001$, $t = 6.21$).[25] Hence, adaptation leads *A-HighVar* participants to exhibit a (slight) tendency to favor options where their targeted feature A is higher, while *B-HighVar* participants exhibit a (stronger) tendency to favor options where feature B is higher.[26]

By design, both treatment groups face the same 20 probes, with cross-randomized probe placement and color schemes. The central adaptation treatment effect remains similar after jointly controlling for color scheme, probe identity, and probe-slot effects (Appendix Table 5, columns (1) and (2)).

The positive *K3* treatment effect persists when restricting the sample to participants with no timeouts across the 100 decisions, those who answer all four comprehension questions correctly on their first attempt, or those who answer both sentinel decisions correctly (Appendix Table 7, columns (3), (4), and (7), respectively). The estimated treatment effect is also similar in the first and second halves of the probe block ($p = 0.839$ for the difference; Appendix Table 4, column (3)).

We also investigate how the treatment effect varies by characteristics of the probes. As expected, participants are more likely to select the A-favored option as the signed value gap $\Delta A + \Delta B$ increases. Consistent with the theory, the treatment effect is stronger for smaller absolute gaps: a ten-point increase in $|\Delta A + \Delta B|$ is associated with a 5.9 percentage-point smaller treatment effect ($p = 0.010$; Appendix Table 4, column (1)).[27]

Result 1 (Adaptation to the Choice Environment)**.** *Participants in the A-HighVar treatment condition are significantly more likely to select the A-favored option among identical probe choices than participants in the B-HighVar treatment condition.*

[25] We pre-registered directional treatment effects, in line with our directional hypotheses. For simplicity, we report p-values for the stricter two-sided tests here and elsewhere, unless stated otherwise.

[26] Despite the probes being balanced, we observe that participants in *K3* are overall more likely to select the B-favored option on average. This difference in levels across attributes is not explained by our theory, but could stem for example from a different salience of the attributes that appear top, center and bottom (Bordalo et al., 2013). As we hold this order constant across treatment groups, it is orthogonal to our estimated treatment effect.

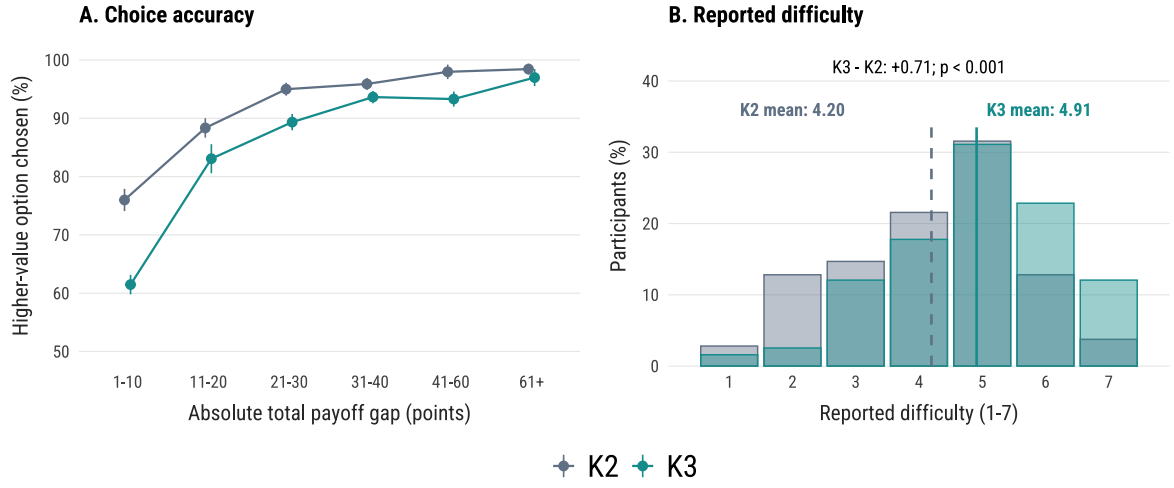
[27] The estimated interaction between treatment and probe width is positive but not statistically significant (Appendix Table 4, columns (1) and (2)).

5.1.2 Aggregation Complexity and Adaptation

We next investigate the role of aggregation complexity in the bias induced by adaptation. To do so, we compare the difference in choice probabilities between *A-HighVar* and *B-HighVar* across the *K3* baseline and the *K2* condition described in Section 4. Response times average around two seconds per decision in both conditions, differing by only 0.01 seconds ($p = 0.823$).

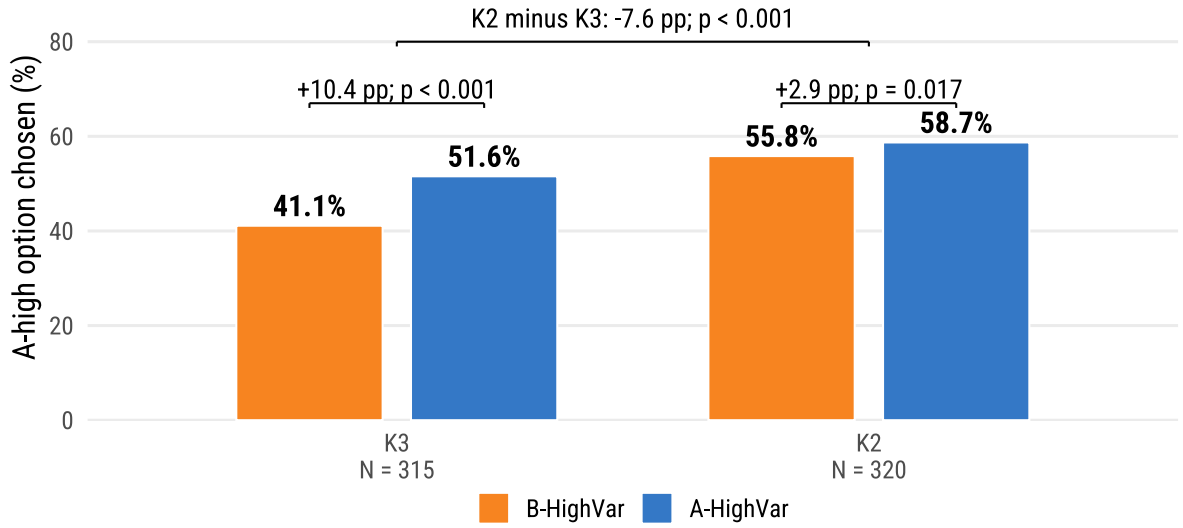
Number of features shifts aggregation complexity. As a manipulation check, we assess the difference in aggregation complexity using choice accuracy and participants' perceptions. First, participants select the higher-value option in 84% of the initial adaptation decisions in *K3*, compared with 90% in *K2* ($p < 0.001$). As shown in Figure 4, Panel A, this difference is largest for small value gaps, with accuracy of 62% and 77% in *K3* and *K2*, respectively, but narrows as value gaps increase and choices become more accurate. Second, participants report greater difficulty in aggregating a product's total value in *K3*. On a Likert scale from 1 ('not at all difficult') to 7 ('extremely difficult'), the average rating is 4.9 in *K3* compared with 4.2 in *K2*, a difference of about 0.51 pooled standard deviations ($p < 0.001$; Figure 4, Panel B).

Figure 4: Higher error rates and reported complexity in three-feature setting



Notes: This figure compares measures of complexity across the *K2* and *K3* treatment arms, pooling across *A-HighVar* and *B-HighVar* participants. Panel A reports the probability of participants selecting the higher-value option in the first 40 adaptation decisions by total payoff gap between both options for the *K2* arm (in gray) and the *K3* arm (in turquoise). Panel B shows participants' perceived difficulty in aggregating an option's total value between both treatment arms.

Figure 5: Reduced aggregation complexity reduces adaptation bias



Notes: This figure compares the reduced-form evidence on adaptation between treatment arms *K2* and *K3* as the mean share of choosing the A-favored option, defined as the option with the higher value of attribute A. The left two bars present the evidence from Figure 3, Panel A for *K3*, compared to the choice probabilities in *K2* on the right side. Smaller brackets indicate differences between *A-HighVar* and *B-HighVar* participants within *K3* and *K2*, with the large bracket reporting the difference-in-differences between treatment groups across *K2* and *K3*. All p-values are from two-sided t-tests, with standard errors clustered on the participant level.

Reduced bias in the less complex condition. We next test whether the adaptation-induced shift in choice probabilities is stronger when aggregation is more complex. We estimate the adaptation effect within *K2* using the same reduced-form specification as above. To compare this effect between the two complexity conditions, we pool participants from all four treatment groups and estimate

$$A\text{-}favored_{i,p} = \alpha + \beta_1 A\text{-}HighVar_i + \gamma_0 K3_i + \gamma_1 (A\text{-}HighVar_i \times K3_i) + \varepsilon_{i,p},$$

where $K3_i$ indicates assignment to the three-feature condition. In this pooled specification, β_1 captures the adaptation effect in *K2*, while $\beta_1 + \gamma_1$ captures the adaptation effect in *K3*. The interaction coefficient γ_1 therefore measures how much more strongly the environment shifts choices in *K3* than in *K2*.[28]

As shown in Figure 5, we find that the adaptation treatment effect is significantly smaller, by 7.6pp or about 72%, than in the more complex three-feature baseline ($p <$

[28]For the comparison across complexity conditions, we normalize weights to sum to 1 in *K3* and 2/3 in *K2*. This matches equal-weight choice probabilities at a common level of aggregation errors on probes with equal C values. Under the interior-solution assumptions of Proposition 9, the adaptation-induced difference in choice probabilities on such probes is smaller in *K3* at the same level of aggregation errors. Under these conditions, testing for a positive interaction $\gamma_1 > 0$ therefore provides a conservative test of the aggregation-error mechanism. The formal result of this, and the same argument for our structural estimation below, is in Appendix A.1.

0.001 for the difference-in-differences comparison): On the same 20 probe choices, participants in *A-HighVar* select the A-favored choice about 2.9pp more often than participants in *B-HighVar*, at 58.7% vs 55.8%, respectively ($p = 0.017$ for the difference).[29] Appendix Table 7 shows that the *K3–K2* difference remains positive and statistically significant across alternative sample restrictions. The *K3–K2* difference also remains positive and statistically significant when restricting both conditions to the four probes with equal C values across products in the *K3* design or to the 16 probes with unequal C values (Appendix Table 8). The adaptation effect is larger in *K3* than in *K2* at each absolute A+B gap. This difference is larger at smaller gaps, although not statistically significantly so ($p = 0.272$).

Result 2 (Stronger Adaptation Bias in More Complex Environment)**.** *Participants in the more complex K3 treatment condition exhibit significantly stronger shifts in their choice probabilities than participants in K2. Within K3, the average A-HighVar participant is about 10pp more likely to select the A-favored option among identical probe choices than participants in the B-HighVar treatment condition. Within the K2 condition, this treatment effect is reduced by about 72% to 3pp.*

5.2 Estimating Decision Weights

A key advantage of the experiment is that its design closely matches the theory laid out in Section 2. This allows us to go beyond the reduced-form treatment effects and estimate participants' relative decision weights on the product features.

Following the model's assumption of normally distributed aggregation errors (see Equation 8), we estimate the probit model as

$$\mathbb{P}(A\text{-}favored_{i,p} = 1) = \Phi\left(\frac{\sum_{j \in J_i} d_{i,j}\Delta_{p,j}}{\sqrt{2}\nu_i}\right), \tag{10}$$

where $\Delta_{p,j}$ is the value of feature j in the A-favored option minus its value in the B-favored option. The set J_i contains the displayed features: $\{A, B, C\}$ in *K3* and $\{A, B\}$ in *K2*. The parameter $d_{i,j}$ denotes participant i's decision weight on feature j, while ν_i

[29] Compared to the *K3* condition, we find a shift in levels toward participants in both *A-HighVar* and *B-HighVar* preferring the A-favored option. As for *K3* (see the corresponding footnote above) this shift is orthogonal to our treatment effect, and could reflect differential salience on the first vs the second feature in a two-vs.-three variable setting.

is the standard deviation of the participant's aggregation error for each product, and Φ is the standard normal cumulative distribution function.

We estimate the model separately for each participant on their non-timeout probe choices. Our main object of interest is the ratio of decision weights, or bias, $d_{i,A}/d_{i,B}$. Because the probit coefficients identify decision weights divided by the same participant-specific level of aggregation errors, the ratio of the coefficients on A and B identifies this relative weight without requiring separate identification of ν_i.

We compare biases between *A-HighVar* and *B-HighVar*, and assess whether this difference is larger in *K3* than in *K2*. The within-complexity comparison tests the cross-environment prediction in Corollary 2, part A. The comparison between *K3* and *K2* assesses whether greater aggregation complexity increases the bias in relative decision weights, as suggested by the cross-error prediction in Corollary 2, part B.

Table 2: Decision weights from individual-level probit estimation

	$K=3$			$K=2$			
	A-HighVar	B-HighVar	A-HighVar vs. B-HighVar	A-HighVar	B-HighVar	A-HighVar vs. B-HighVar	Diff.-in-diff.
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Individual probits							
Geometric mean d_A/d_B	0.999	0.805	$\times 1.241^{***}$	1.179	1.114	$\times 1.058^{**}$	$\times 1.172^{***}$
	(0.029)	(0.030)	(0.059)	(0.024)	(0.017)	(0.027)	(0.063)
Reduced form							
A-favored option chosen (%)	51.55	41.11	10.45^{***}	58.71	55.83	2.88^{**}	7.57^{***}
	(1.17)	(1.20)	(1.68)	(0.91)	(0.78)	(1.20)	(2.06)
Participants	161	154		155	165		

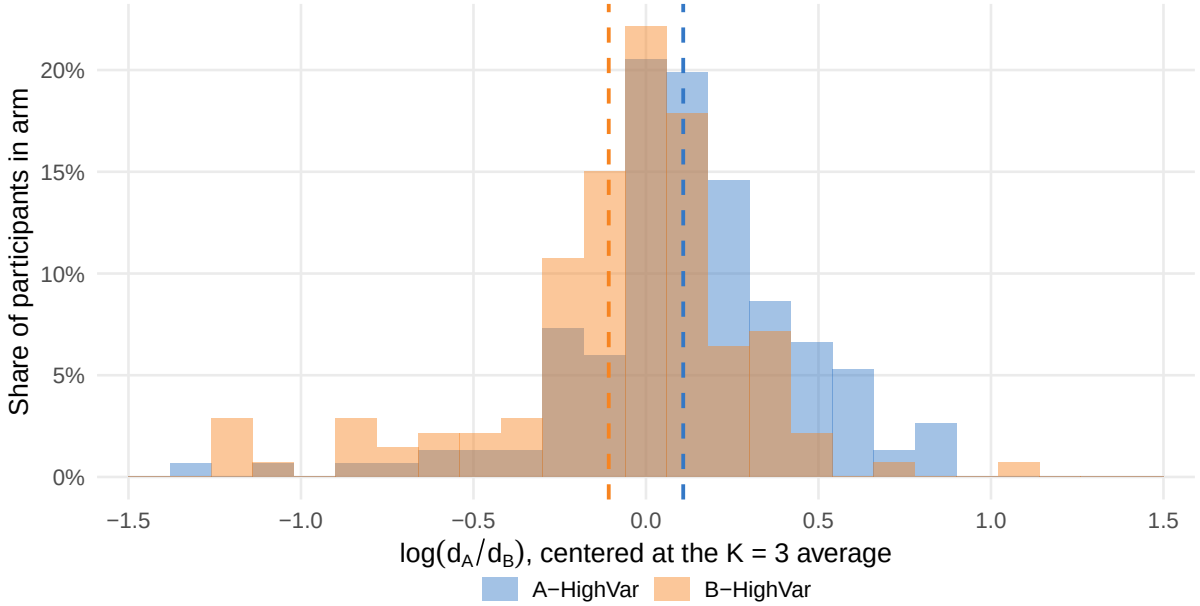
Notes: Columns (1), (2), (4) and (5) report geometric means of d_A/d_B across participants with positive weights on both A and B, excluding the top and bottom 1% of individual $\log(d_A/d_B)$ within each K, with standard errors in parentheses. Columns (3) and (6) report the ratio of the *A-HighVar* to the *B-HighVar* value, and column (7) the ratio of the $K=3$ to the $K=2$ factor (delta-method standard errors; tests on $\log(d_A/d_B)$). Reduced-form: percentage of probe choices for the A-favored option, averaged within and then across participants; columns (3) and (6) report *A-HighVar* minus *B-HighVar*, and column (7) the $K=3$ minus the $K=2$ difference. Reported tests are Welch t-tests. Timeout choices are excluded. $^{*}p<0.1$, $^{**}p<0.05$, $^{***}p<0.01$.

Table 2 reports a summary of this exercise. In line with the reduced-form results discussed above, we find that decision weights on average differ strongly between the treatment groups in the *K3* baseline. When comparing the relative weights between *A-HighVar* and *B-HighVar* in *K3*, we observe that the average *A-HighVar* participant assigns about 24% higher weight to A relative to B than the average *B-HighVar* participant ($p<0.001$), with the average participant valuing their targeted feature about 11% more than the other feature ($p<0.001$). As reported in column (1), the average

A-HighVar participant is estimated to assign close to equal weight to features A and B, while the average *B-HighVar* participant assigns about 19% lower decision weight to feature A than to feature B (Column 2, Row 1). At the margin, this amounts to *B-HighVar* participants valuing a point in A only as much as 0.81 points in B. Estimated relative decision weights therefore shift significantly in the direction predicted by the theory, although the point estimates indicate less reweighting than the model's quantitative benchmark.[30] This incomplete adaptation is not surprising, but rather suggests that forty choices might be too few for participants to become fully adapted to the environment.

The treatment differences in estimated relative decision weights are about 75% smaller for the *K2* setting, where the average *A-HighVar* participant places about 6% more weight on A relative to B than does the average *B-HighVar* participant ($p = 0.025$ for *K2*, Table 2 Column 6, Row 1, $p = 0.003$ for the difference in logs, Column 7, Row 1).

Figure 6: Distribution of relative decision weights



Notes: Decision weights estimated separately for each participant from a probit on their choices in the 20 common probes. Timeout choices are excluded. $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$.

In Figure 6, we consider the heterogeneity in the estimated bias, plotting the logarithm of the weight ratio d_A/d_B, centered by subtracting the logarithm of the geometric mean. The shift in decision weights between *A-HighVar* and *B-HighVar* is broad rather

[30] As a calibration exercise, we calculate optimally adapted weights separately for each participant using their individual probe-noise estimate, and then take the ratio of the treatment-specific geometric means of d_A^*/d_B^*. This yields an optimal treatment factor of 1.917, compared with an observed factor of 1.241 (Table 2, column 3). For the calibration, we additionally exclude 14 participants whose estimated noise implies zero optimal weight on the low-variance attributes ($\hat{\nu}_i^2 \geq 400/3$). The mean individual probe-noise estimate in this calibration sample is $\overline{\hat{\nu}} = 3.124$.

than driven by a few participants, and there is substantial heterogeneity in estimated relative decision weights. Relative to the average participant in *K3*, 72% of *A-HighVar* participants place more weight on A relative to B, compared with 49% of *B-HighVar* participants ($p < 0.001$).

Result 3 (Shifted Decision Weights)**.** ***A.*** *Participants' structurally estimated decision weights adapt in the direction predicted by the theory. In the baseline K3 condition, participants in the A-HighVar treatment group place an about 24% higher relative decision weight d_A/d_B on attribute A than participants in B-HighVar.*

B. *Reducing the aggregation complexity also reduces relative differences in decision weights. In the less complex K2 condition, participants in A-HighVar place about 6% higher relative decision weight on attribute A than participants in B-HighVar.*

6 Application: Add-On Pricing

We now turn to a natural application of our model in add-on pricing. Prices are often displayed as several components rather than as a single all-in amount, be it a sticker price plus sales tax or tip, a ticket price plus a service fee, or an item price plus shipping (see, e.g., Ellison, 2005; Morwitz et al., 1998). While only the total price is relevant for consumers, when it is displayed as multiple components, aggregation is necessary and adaptation results in decision weights that determine consumers' price sensitivity. We first show how we bring the model to the application, and then discuss the implications.

6.1 Products and Displays

Products have three attributes, a quality q and a total price p which is displayed as a base price b and an add-on a, with $p = b + a$, and all-in display is the edge case $a = 0$. Prices are measured in utils, so that a higher value of p, a, or b, represents a lower price and utility is $U = q + p$. We fix a distribution of products throughout, and thus the behavior of consumers not subject to aggregation errors ($\nu = 0$) is independent of the display. The distribution of products is characterized by the standard deviations $\sigma_q, \sigma_p > 0$ of quality and price respectively, and their correlation $\rho \in (-1, 0)$.[31] The price add-on is modeled as

[31] Since the price is measured in utils, a higher value p indicates a lower price and the negative ρ thus captures that higher-quality products, on average, come with higher prices or a lower value p.

a share α of the total price and thus of price variation. We distinguish between two types of add-ons. First, add-ons that are strongly correlated with a product's base price and thus its quality like taxes or proportional fees, and second, add-ons that apply relatively independently of a product's other attributes like junk fees. We refer to the first type as a 'tax' and the second as a 'junk fee.' For simplicity, we model them as perfectly correlated and uncorrelated with base price and quality, respectively. The base price is then given by $b = p - a$. For a junk fee, this means that the larger the share of the add-on, the stronger the correlation of base price and quality. This bounds the share strictly by $\alpha_{\max} = 1 - \rho^2$. For the tax, we assume that the base price is strictly higher than the tax, so that $\alpha \in [0, \frac{1}{2})$.[32]

6.2 Decision Weights on Price Components

For any display, decision weights on quality and base price generically induce a bias in choice and thus influence consumer behavior, as we showed in the general case. For add-ons specifically, adaptation predicts that taxes and proportional fees receive no weight. Junk fees, on the other hand, receive no weight only when they are a relatively small share of the total price. Formally, we have the following.

Proposition 6 (Decision Weights on Add-Ons)**.** *Let $\nu > 0$.* ***Tax.*** *For any $\alpha \in [0, \frac{1}{2})$ the decision weight on a tax is zero, $d_a^* = 0$.*

Junk Fee. *There is $\bar{\alpha}(\nu) > 0$ non-decreasing in ν such that $d_a^* = 0$ if and only if $\alpha \leq \bar{\alpha}(\nu)$. When ν is small enough that both price and quality receive weight under all-in display, then $\bar{\alpha}(\nu) \in (0, \alpha_{\max})$.*

Intuitively, increasing decision weight on the base price is effectively accounting for a tax but does so in a more efficient way. For junk fees, this is not the case because the base price, and also quality, carry no information about the size of such add-ons. When junk

[32]The expected alternatives are given by the covariance matrices

$$\Sigma_T(\alpha) = \begin{pmatrix} \sigma_q^2 & (1-\alpha)\rho\sigma_q\sigma_p & \alpha\rho\sigma_q\sigma_p \\ (1-\alpha)\rho\sigma_q\sigma_p & (1-\alpha)^2\sigma_p^2 & \alpha(1-\alpha)\sigma_p^2 \\ \alpha\rho\sigma_q\sigma_p & \alpha(1-\alpha)\sigma_p^2 & \alpha^2\sigma_p^2 \end{pmatrix}, \Sigma_F(\alpha) = \begin{pmatrix} \sigma_q^2 & \rho\sigma_q\sigma_p & 0 \\ \rho\sigma_q\sigma_p & (1-\alpha)\sigma_p^2 & 0 \\ 0 & 0 & \alpha\sigma_p^2 \end{pmatrix}$$

of (q, b, a) for taxes and junk fees respectively. These matrices are singular, $\Sigma_T(\alpha)$ at every share and $\Sigma_F(\alpha)$ at $\alpha = 0$. For $\nu > 0$, however, $\Sigma + \nu^2 \iota\iota^\top$ is positive definite at every share, which, together with $\iota^\top \Sigma \iota > 0$, is all our proofs require. The optimally adapted weights exist and are unique. Under all-in display the add-on receives no weight, and the other two weights are those of the model with the two attributes q and p.

fees are relatively small then they are not important in deciding which product is better and receive no weight. When they are sufficiently large, they will influence decisions.

These decision weights may also connect to the literature on hidden or shrouded prices, in which firms shift part of the price into add-ons that consumers do not fully take into account (Brown et al., 2010; Ellison, 2005; Gabaix and Laibson, 2006). In our model a small fee receives no weight even if it is displayed and understood, and so acts as a hidden price—even though nothing is hidden. Moreover, which part of the price is effectively hidden is not a fixed feature of the display but depends on the environment the weights are adapted to. This may lead to equilibrium and dynamic effects. Under adaptation to the true environment, firms choose the split of the price across components, and the under-reaction to a component is endogenous. When the weights lag behind changes in the environment, a fee that has received no weight may receive weight only over time, while firms can make use of that hidden price in the meantime.

6.3 Price Display Policies

We now consider the effects of different display policies characterized by the fraction α add-ons take up. For instance, a policy that enforces a display of a single price is captured by $\alpha = 0$. Policies that limit the size of fees or regulate which fees may be excluded from a total price can be thought of as a reduction of α relative to a less regulated setting. We distinguish between short-term changes where the policy changes but decision weights do not adapt to the new environment and long-term changes where the policy also influences adaptation.

Short-term display changes. When temporarily including taxes or proportional fees in the posted price (or excluding them from it) while decision weights stay stable, consumers react as though the display change is a change only in realizations of the base price and add-ons. Since add-ons receive no weight, consumers respond to the change in base price and treat the display change like an increase (or decrease) of the total price by the add-on's amount (for evidence, see Blake et al., 2021; Chetty et al., 2009).[33] In our

[33]The findings of Blake et al. (2021) also suggest that the display that matters is the one at the stage at which options are compared. Since our model is about choice between options, it seems reasonable to assume that decision weights are applied at this stage, and the decrease in price sensitivity would thus be compatible with a full price being shown later.

model this is a result of the need for (error-prone) aggregation and holds even if people are aware of and correctly understand tax rates (for evidence that they do, see Chetty et al., 2009; Taubinsky and Rees-Jones, 2018), and even when these add-ons are a relatively large proportion of total prices (15% in Blake et al., 2021).

Long-term display changes. When decision weights adapt to new display policies, these policies affect price sensitivity and consumer welfare through changes in weights.

The effective weight on the total price is $e^* = (1-\alpha)d_b^* + \alpha d_a^*$.[34] We thus define price sensitivity relative to quality as e^*/d_q^*, with the convention that $e^*/0 = \infty$. Display policies affect taxes and small fees as follows.

Proposition 7 (Price Sensitivity)**.** *Let $\nu > 0$.* ***Tax.*** *If $\mathrm{Cov}(p,U) \leq 2\,\mathrm{Cov}(q,U)$, then for a tax relative price sensitivity is non-increasing in α. Otherwise there exists a ν' for which it rises at some α'.*

Junk Fee. *If $\mathrm{Cov}(q,U) \geq 0$, then for a junk fee relative price sensitivity is non-increasing in $\alpha < \bar{\alpha}(\nu)$. Otherwise, there exists a ν' such that relative price sensitivity rises at some α'.*

Relative price sensitivity falls in the share of the add-on unless the DM cares sufficiently more about prices than about quality, which is expressed through the covariance with utility. Intuitively, when a greater share is displayed as add-ons, capturing the total price in decisions is harder. This increase in difficulty is asymmetric and does not affect quality. This often leads to relative price sensitivity falling. However, when consumers care sufficiently more about price than about quality, relative price sensitivity increases, similar to the logic for a Giffen good.

For consumer welfare, greater dispersion in display is generally harmful, and expected utility always falls in the share of add-ons for taxes and small fees.

Proposition 8 (Welfare)**.** *Let $\nu > 0$.* ***Tax.*** *For a tax expected utility is non-increasing in α.* ***Junk Fee.*** *For a junk fee expected utility is non-increasing in $\alpha < \bar{\alpha}(\nu)$.*

One can show that it decreases even when junk fees receive weight as long as they receive less than the base price, which seems a plausible assumption in most realistic settings.[35]

[34] It is the response of the decision variable to a change in the total price that is split between base and add-on in proportion to their shares. For a tax this is the exact weight in any case, while for a junk fee it is the weight in expectation, $\mathbb{E}[a \mid p] = \alpha p$.

[35] For sufficiently large fees this may overturn, for two reasons. Without correlation between quality

If in addition the number of add-ons increases aggregation errors directly, similar to our experiment, then this is another source of harm. From the perspective of adaptation to aggregation errors, expected utility is generally highest when all-in prices are displayed, and ensuring that attributes are clearly displayed and understood is not sufficient.

7 Conclusion

We developed a model of aggregation errors and adaptation of the decision variable that is aggregated. We showed how it induces a bias in any multi-attribute choice setting, provided experimental evidence for our predictions in a stylized setting close to theory, and highlighted one application to add-on pricing. As we are, to the best of our knowledge, the first following an approach of directly modeling errors at the aggregation step, there remain many open avenues for further research.

As a first step in analyzing adaptation in response to errors directly modeled at the aggregation step, we assumed a simple form of aggregation errors for tractability and to make our theoretical results reflect the baseline implications of adaptation. However, the specific form of these errors would refine the predictions, and existing insights on this form could be built in. For example, an attribute dropping out of consideration might have the consequence of making aggregation simpler. Integrating this as an endogenous drop in the level of aggregation errors implies that the extensive margin of decision weights would be more pronounced, and attributes would drop out of the support discontinuously. On the intensive margin, however, our results and characterization would survive. Other features of aggregation errors, such as a more direct modeling of trade-off complexity, would amount to a more substantial extension of our model and analytical tools.

We focused on a general theory and provided experimental evidence in a stylized setting close to theory. As multi-attribute choice and errors in aggregation are ubiquitous, our insights may be theoretically and empirically expanded and applied in many settings. A notable difficulty is that predictions depend on the environment the weights are adapted to, which in less stylized settings may be less clearly defined and open to interpretation. For example, one could consider an experiment in which attributes are not induced by the

and price, base price and fee are symmetric, so a fee with share $\alpha > \frac{1}{2}$ is just a relabeled base price with share $1 - \alpha$, and expected utility is mirrored at $\alpha = \frac{1}{2}$. With correlation, in addition, base price and quality correlate more strongly the larger the fee, to a degree that decision weights on one inform about both and (in the limit) erase the need for one dimension. The two reasons can also combine.

experimenter but arise more naturally as features of 'real' products. Then, predictions rely on assumptions about what the attributes are and how utility varies in them. This is ultimately an empirical question, and it connects to the more general question of how people categorize and organize product characteristics.

Finally, aggregation errors and adaptation to these errors is likely highly relevant also beyond the multi-attribute choices we study: to update beliefs, one has to combine priors and signals, to decide between politicians or policy proposals one has to aggregate one's views on several issues, and forecasts necessary for intertemporal decisions require the aggregation of many contemporaneous inputs. Even for settings that our model does not immediately apply to, the approach of explicitly modeling errors at such later-stage aggregation steps put forward in this paper might provide new explanations and valuable insights.

References

Blake, Thomas, Sarah Moshary, Kane Sweeney, and Steven Tadelis, "Price Salience and Product Choice," *Marketing Science*, 2021, *40* (4), 619–636.

Bordalo, Pedro, Nicola Gennaioli, and Andrei Shleifer, "Salience and Consumer Choice," *Journal of Political Economy*, 2013, *121* (5), 803–843.

_ , _ , and _ , "Memory, Attention, and Choice," *The Quarterly Journal of Economics*, 2020, *135* (3), 1399–1442.

Brown, Jennifer, Tanjim Hossain, and John Morgan, "Shrouded Attributes and Information Suppression: Evidence from the Field," *The Quarterly Journal of Economics*, 2010, *125* (2), 859–876.

Bushong, Benjamin, Matthew Rabin, and Joshua Schwartzstein, "A Model of Relative Thinking," *The Review of Economic Studies*, 2021, *88* (1), 162–191.

Caplin, Andrew, Mark Dean, and John Leahy, "Rationally Inattentive Behavior: Characterizing and Generalizing Shannon Entropy," *Journal of Political Economy*, 2022, *130* (6), 1676–1715.

Chetty, Raj, Adam Looney, and Kory Kroft, "Salience and Taxation: Theory and Evidence," *American Economic Review*, 2009, *99* (4), 1145–1177.

Drugowitsch, Jan, Valentin Wyart, Anne-Dominique Devauchelle, and Etienne Koechlin, "Computational Precision of Mental Inference as Critical Source of Human Choice Suboptimality," *Neuron*, 2016, *92* (6), 1398–1411.

Dybvig, Philip H, "Short sales restrictions and kinks on the mean variance frontier," *The Journal of Finance*, 1984, *39* (1), 239–244.

Ellison, Glenn, "A Model of Add-on Pricing," *The Quarterly Journal of Economics*, 2005, *120* (2), 585–637.

Enke, Benjamin, "The Cognitive Turn in Behavioral Economics," 2026. Working paper.

_ **and Cassidy Shubatt**, "Quantifying Lottery Choice Complexity," 2026. Working paper; NBER Working Paper 31677 (2023).

_ **and Thomas Graeber**, "Cognitive uncertainty," *The Quarterly Journal of Economics*, 2023, *138* (4), 2021–2067.

_ **,** _ **, and Ryan Oprea**, "Complexity and Time," *Journal of the European Economic Association*, 2025, *23* (5), 1838–1867.

_ **,** _ **,** _ **, and Jeffrey Yang**, "Behavioral Attenuation," *Journal of Political Economy*, 2026. Forthcoming.

Findling, Charles and Valentin Wyart, "Computation noise in human learning and decision-making: origin, impact, function," *Current Opinion in Behavioral Sciences*, 2021, *38*, 124–132.

Frydman, Cary and Lawrence J Jin, "Efficient coding and risky choice," *The Quarterly Journal of Economics*, 2022, *137* (1), 161–213.

_ **and** _ , "On the source and instability of probability weighting," Technical Report, National Bureau of Economic Research 2023.

_ **and** _ , "Adaptation Through Experience or Description? Evidence from an Experiment on Efficient Coding," Technical Report, National Bureau of Economic Research 2025.

Gabaix, Xavier, "A sparsity-based model of bounded rationality," *The Quarterly Journal of Economics*, 2014, *129* (4), 1661–1710.

_ **and David Laibson**, "Shrouded Attributes, Consumer Myopia, and Information Suppression in Competitive Markets," *The Quarterly Journal of Economics*, 2006, *121* (2), 505–540.

Gossner, Olivier and Jakub Steiner, "On the cost of misperception: General results and behavioral applications," *Journal of Economic Theory*, 2018, *177*, 816–847.

Khaw, Mel Win, Ziang Li, and Michael Woodford, "Cognitive imprecision and small-stakes risk aversion," *The Review of Economic Studies*, 2021, *88* (4), 1979–2013.

_ **,** _ **, and** _ , "Cognitive imprecision and stake-dependent risk attitudes," Technical Report, National Bureau of Economic Research 2022.

Kőszegi, Botond and Adam Szeidl, "A Model of Focusing in Economic Choice," *The Quarterly Journal of Economics*, 2013, *128* (1), 53–104.

Landry, Peter and Ryan Webb, "Pairwise normalization: A neuroeconomic theory of multi-attribute choice," *Journal of Economic Theory*, 2021, *193*, 105221.

Li, Vickie, Santiago Herce Castañón, Joshua A. Solomon, Hildward Vandormael, and Christopher Summerfield, "Robust averaging protects decisions from noise in neural computations," *PLOS Computational Biology*, 2017, *13* (8), e1005723.

Maćkowiak, Bartosz, Filip Matějka, and Mirko Wiederholt, "Rational Inattention: A Review," *Journal of Economic Literature*, 2023, *61* (1), 226–273.

Markowitz, Harry, "The optimization of a quadratic function subject to linear constraints," *Naval Research Logistics Quarterly*, 1956, *3* (1-2), 111–133.

Matějka, Filip and Alisdair McKay, "Rational Inattention to Discrete Choices: A New Foundation for the Multinomial Logit Model," *American Economic Review*, 2015, *105* (1), 272–298.

McNaughton, Candace D, Kerri L Cavanaugh, Sunil Kripalani, Russell L Rothman, and Kenneth A Wallston, "Validation of a short, 3-item version of the subjective numeracy scale," *Medical Decision Making*, 2015, *35* (8), 932–936.

Merton, Robert C, "An analytic derivation of the efficient portfolio frontier," *Journal of Financial and Quantitative Analysis*, 1972, *7* (4), 1851–1872.

Mezzadri, Giulia and Michael Woodford, "Sources of imprecision in integrated value comparisons," *Cognition*, 2026, *275*, 106619.

Morwitz, Vicki G., Eric A. Greenleaf, and Eric J. Johnson, "Divide and Prosper: Consumers' Reactions to Partitioned Prices," *Journal of Marketing Research*, 1998, *35* (4), 453–463.

Netzer, Nick, Arthur J. Robson, Jakub Steiner, and Pavel Kocourek, "Risk Perception: Measurement and Aggregation," *Journal of the European Economic Association*, 2025, *23* (4), 1309–1349.

Oprea, Ryan, "Decisions under Risk Are Decisions under Complexity," *American Economic Review*, 2024, *114* (12), 3789–3811.

_ **and Michael Woodford**, "Behavioral Signatures of Cognitive Noise," 2026. Annual Review of Economics, forthcoming; NBER Working Paper 35714.

Prat-Carrabin, Arthur and Michael Woodford, "Efficient coding of numbers explains decision bias and noise," *Nature Human Behaviour*, 2022, *6* (8), 1142–1152.

_ **and** _ , "Endogenous precision of the number sense," *eLife*, 2026, *13*, RP101277.

Rosenbaum, S., "Moments of a Truncated Bivariate Normal Distribution," *Journal of the Royal Statistical Society. Series B*, 1961, *23* (2), 405–408.

Sharpe, William F, "Mutual fund performance," *The Journal of Business*, 1966, *39* (1), 119–138.

Shubatt, Cassidy and Jeffrey Yang, "Tradeoffs and Comparison Complexity," 2026.

_ **and** _ , "Valuations under Tradeoff Complexity," *Available at SSRN 6810701*, 2026.

Sims, Christopher A., "Implications of Rational Inattention," *Journal of Monetary Economics*, 2003, *50* (3), 665–690.

Steiner, Jakub and Colin Stewart, "Perceiving Prospects Properly," *American Economic Review*, 2016, *106* (7), 1601–1631.

Summerfield, Christopher and Paula Parpart, "Normative Principles for Decision-Making in Natural Environments," *Annual Review of Psychology*, 2022, *73*, 53–77.

Taubinsky, Dmitry and Alex Rees-Jones, "Attention Variation and Welfare: Theory and Evidence from a Tax Salience Experiment," *The Review of Economic Studies*, 2018, *85* (4), 2462–2496.

Tobin, James, "Liquidity preference as behavior towards risk," *The Review of Economic Studies*, 1958, *25* (2), 65–86.

Tsetsos, Konstantinos, Rani Moran, James Moreland, Nick Chater, Marius Usher, and Christopher Summerfield, "Economic irrationality is optimal during noisy decision making," *Proceedings of the National Academy of Sciences*, 2016, *113* (11), 3102–3107.

Vieider, Ferdinand M., "Decisions Under Uncertainty as Bayesian Inference on Choice Options," *Management Science*, 2024, *70* (12), 9014–9030.

Woodford, Michael, "Prospect theory as efficient perceptual distortion," *American Economic Review*, 2012, *102* (3), 41–46.

_ , "Modeling imprecision in perception, valuation, and choice," *Annual Review of Economics*, 2020, *12* (1), 579–601.

A Additional Theoretical Results

We now provide the additional theoretical results on which several discussions in the main text are based: the comparison between *K2* and *K3* (A.1) and the relation to perception errors. As the main text discusses these results, we keep the discussion here brief and focus on the formal statements.

A.1 Comparing *K2* and *K3*

The formal statement that for the experiment the comparison of *K2* and *K3* under the normalization of Section 5 is a conservative test is as follows.

Proposition 9 (Comparing *K2* and *K3*)**.** *Let attributes be independent, let A and B have variances $\sigma_h^2 > \sigma_l^2$ in A-HighVar and swapped in B-HighVar, and let K3 add an attribute C with $\sigma_C^2 = \sigma_l^2$. Let the weights sum to $2/3$ in K2 and to 1 in K3, and fix the same $\nu > 0$ in both, with all weights interior. Then the difference between A-HighVar and B-HighVar in the bias d_A^*/d_B^*, in $\log(d_A^*/d_B^*)$, and in the probability of choosing x in any choice with $x_A > y_A$, $x_B < y_B$ and $x_C = y_C$ is smaller in K3 than in K2.*

A.2 Relation to Models based on Perception Errors

We base our discussion on the following model of adaptation to perception errors, with Bayesian decoding under normal distributions.

The DM does not observe goods exactly: instead of x she receives a signal $r_x \mid x \sim \mathcal{N}(x, \Psi)$, where Ψ is a symmetric positive definite covariance matrix of perception errors, and the signals for the two goods are conditionally independent of each other. We allow diagonal entries of Ψ to be infinite, for attributes perceived with so much noise that they are ignored, and then require the submatrix of the remaining attributes to be positive definite. The expected choice environment is given by μ and Σ, as before. The DM decodes her signals optimally, that is, she maximizes the ex-ante fit between the true attributes and what she uses downstream, given her noisy perception.[36] This reduces to

[36]Formally, $\mathbb{E}[x \mid r_x]$ minimizes $\mathbb{E}\|x - \hat{x}\|^2$ over all functions of the signal, and with Gaussian signals the posterior is a sufficient statistic. Such a criterion for the optimal response to perception errors matches what several other studies assume. Yet, there are also models in which the optimal response incorporates the downstream choice problem. Only absent any aggregation errors do these two criteria agree in our model.

her using the Bayesian posterior

$$\hat{x} = \mathbb{E}[x \mid r_x] = \mu + \Sigma(\Sigma + \Psi)^{-1}(r_x - \mu), \tag{11}$$

understood as a limit if some entries of Ψ are infinite. She chooses the good whose representation has the larger total value, without aggregation errors, and hence weights her signals by $\iota^\top \Sigma(\Sigma + \Psi)^{-1}$.

We call this the perceptual model with noise Ψ, and our main model with level of aggregation errors ν the aggregation model. The two are equivalent if they have the same choice probabilities for every pair of goods.

Theorem 2 (Aggregation Errors as Perception Errors)**.** ***A.*** *Fix any $\nu > 0$ and let d^* be the optimally adapted weights. The perceptual model with diagonal Ψ and entries ν^2/d_i^*, with the convention $\nu^2/0 = \infty$, is equivalent to the aggregation model.* ***B.*** *Let $\mathbb{E}[U(x) \mid r_x]$ be increasing in the signal of every attribute perceived with finite errors. The perceptual model is equivalent to an aggregation model for some $\nu > 0$ if and only if, after perception, all attributes perceived with finite errors have the same covariance with utility, $\mathrm{Cov}(x_i, U(x) \mid r_x)$, and no attribute perceived with infinite noise has a higher one.*

Part A shows that any aggregation model is equivalent to a perceptual one, obtained mechanically by choosing the perception errors that reproduce the adapted weights. Part B shows that only a particular class of perceptual models is equivalent to an aggregation model. Concretely, perception has to be efficient in the sense that a marginal shift of precision from one attribute to another does not raise expected utility.

B Extensions and Model Variants

We now consider two natural extensions: choice among an arbitrary number of goods, and choice between a random good and a known outside option. Optimally adapted weights do not depend on the number of goods the DM chooses between. The choice against a known outside option retains the incentives we identified, but a new one arises that makes the problem in general intractable. In Appendix B.3 we also consider, as a more technical robustness check, weights that may take either sign.

B.1 Multiple Goods

Suppose the DM at the choice phase faces $K \geq 2$ goods, drawn independently from the same normal distribution. We still assume that each is evaluated with its own aggregation error, that the one with the highest choice variable is chosen, and that the weights adapt to this environment.[37] The optimally adapted weights are invariant to the number of goods:

Proposition 10 (Choice among K Goods)**.** *There exists a factor c_K, depending only on K and increasing in K, such that*

$$\mathbb{E}[U(\textit{chosen good})] = \iota^\top \mu + c_K \, \iota^\top \Sigma d / \sqrt{d^\top \Sigma d + \nu^2}. \tag{12}$$

The optimally adapted weights therefore do not depend on K, and expected utility under them is strictly increasing in K.

The number of alternatives scales what is at stake, so more options raise expected utility. The effect of the decision weights, however, is only scaled by this stakes effect, and the same weights remain optimal.

B.2 Choice with an Outside Option

Suppose the DM chooses between one good drawn from the expected alternatives and an outside option she knows perfectly, e.g., a default product or not buying. We assume that aggregating the outside option is not subject to errors.[38] Let x^0 be the vector of attributes characterizing the outside option and write $\delta = \mu - x^0$ for the expected advantage of the good over the outside option.

Proposition 11 (Choice with an Outside Option)**.** *The DM's expected utility gain over the outside option is*

$$\iota^\top \delta \; \Phi\left(\frac{d^\top \delta}{\sqrt{d^\top \Sigma d + \nu^2}}\right) + \phi\left(\frac{d^\top \delta}{\sqrt{d^\top \Sigma d + \nu^2}}\right) \iota^\top \Sigma d / \sqrt{d^\top \Sigma d + \nu^2}\,. \tag{13}$$

The objective splits into two parts. The first is the expected utility advantage of the good, $\iota^\top \delta$, weighted by the probability that the good is chosen. This is the new incentive: when the good is better in expectation, adaptation favors weights that make it chosen

[37] The reformulations of the choice rule in Section 2 do not carry over to more than two goods.

[38] Assuming instead that it too is aggregated with an error as before is equivalent to replacing ν by $\sqrt{2}\nu$ in the next result.

often, which pulls weight toward the attributes on which it beats the outside option, large δ_i, regardless of how they covary with utility; when it is worse in expectation, the pull is reversed. The second part carries the incentives of the main model: one factor, $\iota^\top \Sigma d / \sqrt{d^\top \Sigma d + \nu^2}$, is exactly the main model's objective and captures the same force — how utility co-moves with the decision variable that governs which option is picked. Between two ex-ante symmetric goods, as in the main model, this co-movement is the only relevant force. With an outside option its importance depends on how close the comparison is expected to be: the co-movement can only tip choices near the margin, and it is therefore scaled by $\phi(\cdot)$, which is largest when the comparison is expected to be close and small when the choice is a foregone conclusion.

If the outside option equals the mean of the expected alternatives ($x^0 = \mu$), the model collapses to the main one. If, for example, the outside option is a default product whose characteristics match the mean expectation, our previous analysis applies as is.

The individual incentives remain interpretable, but their interaction can make the problem intractable:

Remark 2. The objective (13) need not be quasi-concave in the weights.

Hence the first-order conditions need not identify the optimal weights, and no analytical characterization is available.

B.3 Model Variant: Decision Weights of Either Sign

We now consider a model variant in which we allow decision weights to be negative. This changes neither the incentives nor the core mechanism. The reformulation of the objective is unchanged, so optimally adapted weights still trade off decision ease and distortion, and Theorem 1 largely extends. The key difference is that the general closed form depends on the support (which weights are non-zero) and also on the signs of the weights. Describing the optimally adapted weights in general therefore becomes even more complex.

As in the main model, the weights require a bound. We take it as an intuitive benchmark that weighting an attribute negatively does not free cognitive resources, and hence impose a (normalized) bound of 1 on the sum of absolute weights. A *signed support* is a pair $(\mathcal{S}, s)$, where $\mathcal{S}$ is the set of attributes receiving non-zero weight and s the vector of signs of the non-zero weights.

Theorem 3 (Weights of Either Sign)**.** *Let $|d|$ be the vector with entries $|d_i|$ and let weights be restricted only by $\iota^\top |d| = 1$. The optimally adapted weights d^* exist and are unique.* ***A.*** *(Closed Form within a Signed Support) Let $(\mathcal{S}, s)$ be the signed support of d^* at ν. Then (7) holds with s in place of $\iota_\mathcal{S}$ throughout $\alpha_\mathcal{S}$, $\beta_\mathcal{S}$ and $\gamma_\mathcal{S}$.* ***B.*** *(Global Behavior) The signed support is unique and changes a finite number of times. As a function of ν^2 the optimal weights are continuous, and smooth except at the finitely many points where the signed support changes. There is a threshold $\overline{\nu}' < \infty$ above which the weights are constant and supported within $\arg\max_i |\iota^\top \Sigma e_i|$, each weight carrying the sign of $\iota^\top \Sigma e_i$.*

A negative weight can be optimal when it, for example, reduces distortion (equivalently, increases the correlation with utility) when a high value of an attribute predicts lower utility. This is the case when the attribute covaries sufficiently negatively with the other attributes. Moreover, even when an attribute predicts higher utility, a negative weight on it can reduce distortion by counteracting the distortion from over-weighted attributes. This, too, relies on sufficiently strong covariances.

In our main applications, when only the intensive margin is active or when attributes are independent, weights are never negative.

Proposition 12 (Agreement with the Main Model)**.** *The optimally adapted weights of Theorem 3 coincide with those of the main model when the latter have full support, and hence for every $\nu < \underline{\nu}$, with $\underline{\nu}$ the threshold of Theorem 1B; and at every $\nu \geq 0$ when attributes are independent.*

C Notation and Supporting Results

Before we prove the theoretical results, we introduce the following notation that we will use throughout Appendix D and state some supporting results.

Notation. Recall that Σ is positive definite, ι is the vector of ones, and for a nonempty $\mathcal{S} \subseteq \{1, \ldots, I\}$ we write $\Sigma_{\mathcal{S}\mathcal{S}}$ for the principal submatrix on $\mathcal{S}$, which is positive definite as a principal submatrix of a positive definite matrix. We use the shorthand $m = \Sigma\iota$ and, for the closed form, $\gamma_\mathcal{S} = \iota_\mathcal{S}^\top \Sigma_{\mathcal{S}\mathcal{S}}^{-1} \iota_\mathcal{S}$, $\alpha_\mathcal{S} = \Sigma_{\mathcal{S}\mathcal{S}}^{-1} m_\mathcal{S} / (\iota_\mathcal{S}^\top \Sigma_{\mathcal{S}\mathcal{S}}^{-1} m_\mathcal{S})$, and $\beta_\mathcal{S} = \Sigma_{\mathcal{S}\mathcal{S}}^{-1} \iota_\mathcal{S} / \gamma_\mathcal{S}$. We write $R(d;\nu) = \iota^\top \Sigma d / \sqrt{d^\top \Sigma d + \nu^2}$ (R for ratio) for the reduced objective. Note that

$d^\top \Sigma d > 0$ for every $d \in \Delta^I$, since Σ is positive definite and $d \neq 0$; in particular $R(\cdot;\nu)$ is well defined on Δ^I for every $\nu \geq 0$.

Supporting results. We will make use of the following formula for the expected value of a truncated bivariate normal, which follows directly, for example, from Rosenbaum (1961).

Fact 1. Let (V, W) be jointly normal with $\mathrm{Var}(W) > 0$. Then

$$\mathbb{E}[V\,\mathbf{1}\{W \geq 0\}] = \mathbb{E}[V]\,\Phi\Big(\frac{\mathbb{E}[W]}{\sqrt{\mathrm{Var}(W)}}\Big) + \frac{\mathrm{Cov}(V,W)}{\sqrt{\mathrm{Var}(W)}}\,\phi\Big(\frac{\mathbb{E}[W]}{\sqrt{\mathrm{Var}(W)}}\Big),$$

where ϕ is the standard normal density.

Next, a key technical step is the reduction of the adaptation problem (2) to the ratio maximization (14), together with the identities it delivers.

$$d^* = \arg\max_{d \in \Delta^I} \frac{\iota^\top \Sigma d}{\sqrt{d^\top \Sigma d + \nu^2}}. \tag{14}$$

Lemma 1 (Reduction to a Ratio)**.** *Problem* (2) *is equivalent to maximizing* $R(d;\nu)$ *over* Δ^I*, and the two share the same maximizers. Moreover, at any* d*,*

$$\mathbb{E}[U(\textit{chosen good})] = \iota^\top \mu + \frac{R(d;\nu)}{\sqrt{\pi}}, \tag{15}$$

$$\mathbb{P}[\textit{choice maximizes utility}] = \frac{1}{2} + \frac{1}{\pi}\arcsin(R(d;\nu)/\sqrt{\iota^\top \Sigma \iota}), \tag{16}$$

$$\mathbb{P}[\textit{choice maximizes } D(x)] = \frac{1}{2} + \frac{1}{\pi}\arcsin\Big(\sqrt{d^\top \Sigma d/(d^\top \Sigma d + \nu^2)}\Big), \tag{17}$$

the argument of the last being the decision ease relative to the errors.

It will be helpful to analyze this as the following equivalent convex program, which also directly delivers existence and uniqueness of the optimally adapted weights.

Lemma 2 (Equivalent Program)**.** *Fix* $\nu \geq 0$*. Maximizing* $R(\cdot;\nu)$ *over* Δ^I *is equivalent to*

$$\max_{a \geq 0}\ m^\top a \quad s.t. \quad a^\top \Omega a \leq 1, \qquad \Omega = \Sigma + \nu^2 \iota\iota^\top, \tag{18}$$

with $d^* = a^*/(\iota^\top a^*)$*. A maximizer exists, is unique, and satisfies the constraint with equality.*

The first-order conditions that characterize the optimum are as follows.

Lemma 3 (KKT Conditions)**.** *$d \in \Delta^I$ maximizes $R(\cdot;\nu)$ iff there is $\xi > 0$ with*

$$m_i = \xi(\Sigma d)_i + \xi\nu^2 \quad (d_i > 0), \qquad m_i \leq \xi(\Sigma d)_i + \xi\nu^2 \quad (d_i = 0).$$

Finally, we state two elementary facts about how the ratio objective responds to scaling weights and errors; they formalize the normalization discussed in Section 2 and are used in the proofs of Appendix D. Here $R(d;\nu)$ is read as a function on all of $\mathbb{R}^I$ (with $d^\top\Sigma d+\nu^2 > 0$).

Lemma 4 (Scaling of Weights and Errors)**.** *Let $d \in \mathbb{R}^I$, $\lambda > 0$ and $\nu \geq 0$. Then (i) $R(\lambda d;\nu) = R(d;\nu/\lambda)$; and (ii) if $\iota^\top\Sigma d > 0$, the map $\lambda \mapsto R(\lambda d;\nu)$ is non-decreasing on $(0,\infty)$, and strictly increasing if $\nu > 0$.*

D Proofs

D.1 Proofs: Main Text Results

Proof of Proposition 1. For $x \sim \mathcal{N}(\mu,\Sigma)$, $\mathrm{Var}(U(x)) = \iota^\top\Sigma\iota$, $\mathrm{Var}(D(x)) = d^\top\Sigma d > 0$ and $\mathrm{Cov}(U(x), D(x)) = \iota^\top\Sigma d$, so for $d \in \Delta^I$ the product (5) equals

$$\sqrt{\iota^\top\Sigma\iota} \times \frac{\iota^\top\Sigma d}{\sqrt{\iota^\top\Sigma\iota}\,\sqrt{d^\top\Sigma d}} \times \sqrt{\frac{d^\top\Sigma d}{d^\top\Sigma d+\nu^2}} = \frac{\iota^\top\Sigma d}{\sqrt{d^\top\Sigma d+\nu^2}} = R(d;\nu).$$

By (15) of Lemma 1, the objective in (2) is an affine function of $R(d;\nu)$, which finishes the proof. □

Proof of Proposition 2. Fix $0 \leq \nu_1 < \nu_2$ and let d_j be the optimally adapted weights at ν_j. By Lemma 1, d_j maximizes $R(\cdot;\nu_j)$ and thus $R(d_j;\nu_j) \geq R(d_{-j};\nu_j)$. Written out, we have

$$\frac{\iota^\top\Sigma d_1}{\sqrt{\mathrm{Var}(D_1(x))+\nu_1^2}} \geq \frac{\iota^\top\Sigma d_2}{\sqrt{\mathrm{Var}(D_2(x))+\nu_1^2}} \text{ and } \frac{\iota^\top\Sigma d_2}{\sqrt{\mathrm{Var}(D_2(x))+\nu_2^2}} \geq \frac{\iota^\top\Sigma d_1}{\sqrt{\mathrm{Var}(D_1(x))+\nu_2^2}}$$

and multiplying the two inequalities and canceling $(\iota^\top\Sigma d_1)(\iota^\top\Sigma d_2) > 0$ gives

$$(\mathrm{Var}(D_2(x))+\nu_1^2)(\mathrm{Var}(D_1(x))+\nu_2^2) \geq (\mathrm{Var}(D_1(x))+\nu_1^2)(\mathrm{Var}(D_2(x))+\nu_2^2)$$
$$\iff (\nu_2^2-\nu_1^2)(\mathrm{Var}(D_2(x)) - \mathrm{Var}(D_1(x))) \geq 0.$$

Hence, decision ease weakly increases in ν, and from Proposition 1 it follows immediately that distortion then has to weakly increase as well. □

Proof of Proposition 3. For $\Sigma = \operatorname{diag}(\sigma_1^2, \dots, \sigma_I^2)$ and any support $\mathcal{S}$, $\Sigma_{\mathcal{S}\mathcal{S}}^{-1}(\Sigma\iota)_{\mathcal{S}} = \iota_{\mathcal{S}}$ and $\Sigma_{\mathcal{S}\mathcal{S}}^{-1}\iota_{\mathcal{S}} = (\sigma_i^{-2})_{i\in\mathcal{S}}$, so $\alpha_{\mathcal{S}} = \iota_{\mathcal{S}}/|\mathcal{S}|$ and $\gamma_{\mathcal{S}}\beta_{\mathcal{S}} = (\sigma_i^{-2})_{i\in\mathcal{S}}$. By Theorem 1A the optimally adapted weights on their support $\mathcal{S}$ are thus

$$d_i^* = \frac{1}{|\mathcal{S}|} + \frac{\nu^2}{|\mathcal{S}|}\sum_{j\in\mathcal{S}}\left(\sigma_j^{-2} - \sigma_i^{-2}\right), \qquad i \in \mathcal{S},$$

and it remains to show that $\mathcal{S} = \{1, \dots, J\}$. By Lemma 3 there is $\xi > 0$ with $\sigma_i^2 = \xi(\sigma_i^2 d_i^* + \nu^2) > \xi\nu^2$ for $i \in \mathcal{S}$ and $\sigma_i^2 \le \xi\nu^2$ for $i \notin \mathcal{S}$. Every attribute outside the support thus has a lower variance than every attribute in it, so $\mathcal{S} = \{1, \dots, k\}$ for some k, and the weights above are (6) with $J = k$, which are non-negative. If $k < I$, summing $\xi^{-1} = d_i^* + \nu^2\sigma_i^{-2}$ over $i \le k$ turns the inequality at $k+1$ into

$$1 + \nu^2\sum_{j=1}^{k+1}\left(\sigma_j^{-2} - \sigma_{k+1}^{-2}\right) \le 0,$$

that is, (6) with $J = k+1$ gives attribute $k+1$ a non-positive weight. Since each term $\sigma_j^{-2} - \sigma_{J'}^{-2}$ decreases in J', the weight that (6) gives attribute J' with $J = J'$ is then negative for every $J' > k+1$. Hence $J = k$, or $J = k+1$ with a zero weight on attribute $k+1$, and one directly sees that both give the same weights. □

Proof of Theorem 1. By Lemma 2, the optimally adapted weights d^* exist and are unique at every $\nu \ge 0$, and hence so is their support.

Part A. Let $\mathcal{S}$ be the support of d^* and let $\xi > 0$ be as in Lemma 3. Since d^* vanishes off $\mathcal{S}$, the equalities of Lemma 3 read $m_{\mathcal{S}} = \xi(\Sigma_{\mathcal{S}\mathcal{S}}d_{\mathcal{S}}^* + \nu^2\iota_{\mathcal{S}})$. With $\Sigma_{\mathcal{S}\mathcal{S}}^{-1}\iota_{\mathcal{S}} = \gamma_{\mathcal{S}}\beta_{\mathcal{S}}$ this gives $d_{\mathcal{S}}^* = \xi^{-1}\Sigma_{\mathcal{S}\mathcal{S}}^{-1}m_{\mathcal{S}} - \nu^2\gamma_{\mathcal{S}}\beta_{\mathcal{S}}$. Multiplying by $\iota_{\mathcal{S}}^\top$ and using $\iota_{\mathcal{S}}^\top d_{\mathcal{S}}^* = \iota_{\mathcal{S}}^\top\beta_{\mathcal{S}} = 1$ yields $\xi^{-1}\iota_{\mathcal{S}}^\top\Sigma_{\mathcal{S}\mathcal{S}}^{-1}m_{\mathcal{S}} = 1 + \nu^2\gamma_{\mathcal{S}}$. Since $\Sigma_{\mathcal{S}\mathcal{S}}^{-1}$ is positive definite, $\gamma_{\mathcal{S}} > 0$, and with $\xi > 0$ this gives $\iota_{\mathcal{S}}^\top\Sigma_{\mathcal{S}\mathcal{S}}^{-1}m_{\mathcal{S}} > 0$. Hence $\alpha_{\mathcal{S}}$ is well defined, and

$$\xi = \frac{\iota_{\mathcal{S}}^\top\Sigma_{\mathcal{S}\mathcal{S}}^{-1}m_{\mathcal{S}}}{1 + \nu^2\gamma_{\mathcal{S}}} \quad \text{and} \quad d_{\mathcal{S}}^* = (1 + \nu^2\gamma_{\mathcal{S}})\alpha_{\mathcal{S}} - \nu^2\gamma_{\mathcal{S}}\beta_{\mathcal{S}},$$

which is (7) rearranged.

Part B. Since there are finitely many supports, it suffices to show that each support occurs on a single interval of ν^2, on which the weights are affine in ν^2. Fix $\mathcal{S}$ with $\iota_{\mathcal{S}}^\top\Sigma_{\mathcal{S}\mathcal{S}}^{-1}m_{\mathcal{S}} > 0$ and let $d^{\mathcal{S}}$ equal the right-hand side of (7) on $\mathcal{S}$ and 0 off $\mathcal{S}$. It sums to one and, since $\Sigma_{\mathcal{S}\mathcal{S}}\alpha_{\mathcal{S}}$ is proportional to $m_{\mathcal{S}}$ and $\Sigma_{\mathcal{S}\mathcal{S}}\beta_{\mathcal{S}}$ to $\iota_{\mathcal{S}}$, it satisfies the equalities of Lemma 3 with ξ as in Part A. By Part A and Lemma 3, $\mathcal{S}$ is thus the support at ν if and only if

$$d_i^{\mathcal{S}} > 0 \quad (i \in \mathcal{S}) \qquad \text{and} \qquad (1 + \nu^2\gamma_{\mathcal{S}})\, m_i \le \iota_{\mathcal{S}}^\top\Sigma_{\mathcal{S}\mathcal{S}}^{-1}m_{\mathcal{S}}((\Sigma d^{\mathcal{S}})_i + \nu^2) \quad (i \notin \mathcal{S}),$$

where the second condition is the inequality of Lemma 3 multiplied by $1 + \nu^2\gamma_{\mathcal{S}}$. All conditions are affine in ν^2 and thus hold on an interval, on which $d^* = d^{\mathcal{S}}$ is affine in ν^2. Hence the support changes finitely many times, and the weights are smooth in ν^2 except at these changes. By Berge's maximum theorem they are also continuous, since R is jointly continuous in (d, ν), Δ^I is compact and the maximizer is unique.

It remains to find the two thresholds. At $\nu = 0$, $d = \iota/I$ satisfies the conditions of Lemma 3 with $\xi = I$, since $\Sigma\iota/I = m/I$. Hence $d^* = \iota/I$ at $\nu = 0$, and by continuity the support is full on $[0, \underline{\nu})$ for some $\underline{\nu} > 0$. Since finitely many intervals partition $[0, \infty)$, exactly one of them is unbounded. Let $\mathcal{S}$ be the support on it and choose $\overline{\nu}$ such that it contains all $\nu^2 > \overline{\nu}^2$. There $d^* = d^{\mathcal{S}}$ is affine in ν^2 but also bounded because it is in Δ^I. Its slope $\gamma_{\mathcal{S}}(\alpha_{\mathcal{S}} - \beta_{\mathcal{S}})$ therefore must equal 0. So d^* is constant above $\overline{\nu}$ and $\alpha_{\mathcal{S}} = \beta_{\mathcal{S}}$, that is, $m_i = \iota_{\mathcal{S}}^\top \Sigma_{\mathcal{S}\mathcal{S}}^{-1} m_{\mathcal{S}}/\gamma_{\mathcal{S}}$ for all $i \in \mathcal{S}$. For $i \notin \mathcal{S}$, dividing the inequality above by ν^2 and letting $\nu \to \infty$ gives $m_i \leq \iota_{\mathcal{S}}^\top \Sigma_{\mathcal{S}\mathcal{S}}^{-1} m_{\mathcal{S}}/\gamma_{\mathcal{S}}$. Hence $\mathcal{S} \subseteq \arg\max_i m_i = \arg\max_i \iota^\top \Sigma e_i$. □

Proof of Corollary 1. By Part B of Theorem 1, there is $\underline{\nu} > 0$ such that the support is full for $\nu \in [0, \underline{\nu})$. For $\mathcal{S} = \{1, \ldots, I\}$, $m_{\mathcal{S}} = m = \Sigma\iota$ gives $\alpha_{\mathcal{S}} = \iota/I$ and $\gamma_{\mathcal{S}}\beta_{\mathcal{S}} = \Sigma^{-1}\iota$. Hence (7) reads

$$d^* = \frac{\iota}{I} + \nu^2\Big(\frac{\gamma_{\mathcal{S}}}{I}\,\iota - \Sigma^{-1}\iota\Big) \qquad \text{on } [0, \underline{\nu}),$$

where the vector in parentheses does not depend on ν. Relabel the attributes so that $(\Sigma^{-1}\iota)_1 \leq (\Sigma^{-1}\iota)_2 \leq \ldots \leq (\Sigma^{-1}\iota)_I$. Then for $i < j$, $d_i^* - d_j^* = \nu^2((\Sigma^{-1}\iota)_j - (\Sigma^{-1}\iota)_i) \geq 0$, which is weakly increasing in ν on $[0, \underline{\nu})$. □

Proof of Proposition 4. By Lemma 1, the optimally adapted weights $d^*(\nu)$ maximize $R(\cdot; \nu)$ over Δ^I, and by (15) and (16) both measures are strictly increasing functions of the maximal value $R(d^*(\nu); \nu)$. It therefore suffices to show that this value is strictly decreasing in ν. Fix $0 \leq \nu_1 < \nu_2$ and let $d_j = d^*(\nu_j)$. Since $R(d_2; \nu_2) \geq R(\iota/I; \nu_2) > 0$, $\iota^\top \Sigma d_2 > 0$, so $R(d_2; \nu)$ is strictly decreasing in ν. Hence

$$R(d_2; \nu_2) < R(d_2; \nu_1) \leq R(d_1; \nu_1).$$

□

Proof of Corollary 2. Since attributes are independent, the weights are given by Proposition 3. One can directly see that swapping σ_A^2 and σ_B^2 results in weights where d_A^* and d_B^* are swapped. Now let d^* denote the weights in *A-HighVar*. There A has

the highest variance, so $d_A^* \geq 1/J > 0$, and B has the lowest. If $d_B^* > 0$, (6) gives $d_A^* - d_B^* = \nu^2(\sigma_B^{-2} - \sigma_A^{-2}) > 0$, and otherwise $d_A^* > 0 = d_B^*$. **A.** Hence, the bias toward A is $d_A^*/d_B^* > 1$ in *A-HighVar* and by symmetry $d_B^*/d_A^* < 1$ in *B-HighVar*. **B.** From (6) it follows directly that $d_A^* - d_B^* = \nu^2(\sigma_B^{-2} - \sigma_A^{-2})$ and thus

$$\frac{d_A^*}{d_B^*} = 1 + \frac{\nu^2(\sigma_B^{-2} - \sigma_A^{-2})}{d_B^*} = 1 + \frac{J\nu^2(\sigma_B^{-2} - \sigma_A^{-2})}{1 + \nu^2 \sum_{j=1}^{J} \left(\sigma_j^{-2} - \sigma_B^{-2}\right)} = 1 + \frac{J(\sigma_B^{-2} - \sigma_A^{-2})}{\nu^{-2} + \sum_{j=1}^{J} \left(\sigma_j^{-2} - \sigma_B^{-2}\right)}.$$

In the last fraction the numerator is a positive constant, and the denominator is a constant plus ν^{-2}, which decreases in ν. Hence the bias strictly increases in ν. Again by symmetry, the same bias in *B-HighVar* is strictly decreasing. □

Proof of Proposition 5. As in the proof of Corollary 2, the weights in *B-HighVar* are those in *A-HighVar*, denoted d^*, with d_A^* and d_B^* swapped, and $d_A^* > d_B^*$. Let x be the A-favored good, y the other, and $z = x - y$, so that $z_A > 0 > z_B$ and $z_i = 0$ for $i \notin \{A, B\}$. By (8), x is chosen with probability $\Phi(\omega_+)$ in *A-HighVar* and $\Phi(\omega_-)$ in *B-HighVar*, where $\omega_+ = (d_A^* z_A + d_B^* z_B)/(\sqrt{2}\nu)$ and $\omega_- = (d_B^* z_A + d_A^* z_B)/(\sqrt{2}\nu)$.

Part A. Since $\omega_+ - \omega_- = (d_A^* - d_B^*)(z_A - z_B)/(\sqrt{2}\nu) > 0$ and Φ is strictly increasing, $\Phi(\omega_+) > \Phi(\omega_-)$.

Part B. Since B has the lowest variance and receives weight, all attributes receive weight, and substituting in (6) with $J = I$ for d_A^* and d_B^* into ω_+ and ω_- gives

$$\omega_\pm = (\sqrt{2}\nu)^{-1} \left(\frac{z_A + z_B}{I} + c_\pm \nu^2 \right) = (\sqrt{2})^{-1} \left(\frac{z_A + z_B}{I\nu} + c_\pm \nu \right),$$

where $c_+ = \frac{1}{I} \sum_{\mathrm{i}=1}^{\mathrm{I}} [(\sigma_i^{-2} - \sigma_A^{-2}) z_A + (\sigma_i^{-2} - \sigma_B^{-2}) z_B]$ and c_- is the same with σ_A and σ_B swapped. Since $\sigma_A^{-2} \leq \sigma_i^{-2} \leq \sigma_B^{-2}$ and $z_A > 0 > z_B$, $c_+ > 0 > c_-$. The derivative of $\Phi(\omega_+) - \Phi(\omega_-)$ in ν is thus $1/\sqrt{2}$ times

$$c_+ \phi(\omega_+) - c_- \phi(\omega_-) - \frac{z_A + z_B}{I\nu^2} (\phi(\omega_+) - \phi(\omega_-)).$$

The first two terms are positive. The last is non-negative, since $\omega_+ > \omega_-$ and $\omega_+ + \omega_- = (d_A^* + d_B^*)(z_A + z_B)/(\sqrt{2}\nu)$ has the sign of $z_A + z_B$, so that $\phi(\omega_+) \leq \phi(\omega_-)$ if and only if $z_A + z_B \geq 0$. □

Proof of Proposition 6. Tax. Write $e = (1 - \alpha) d_b + \alpha d_a$ for the absolute price sensitivity induced by weights d_b and d_a. Since $b = (1 - \alpha)p$ and $a = \alpha p$, the decision variable is $D = d_q q + ep$, so $R(d; \nu)$ is the same objective as $R((d_q, e)^\top; \nu)$ for the two-attribute environment (q, p). Note that (d_q, e) do not sum to 1, but $d_q + e/(1 - \alpha) = 1 - d_a(1 - 2\alpha)/(1 - \alpha) \leq 1$, with equality if and only if $d_a = 0$ (even then e is scaled

up by $(1-\alpha)^{-1}$). If $d_a^* > 0$, let $\lambda = 1/(d_q^* + e^*/(1-\alpha)) > 1$ and consider the weights λd_q^* on quality, $\lambda e^*/(1-\alpha)$ on the base price and 0 on the tax. They sum to one and their absolute price sensitivity is λe^*, so their decision variable is λ times that of d^* and their objective is $R(\lambda d^*; \nu)$. By Lemma 4(ii) this is strictly larger than $R(d^*; \nu)$, since $\iota^\top \Sigma d^* > 0$ because $R(d^*; \nu) \geq R(\iota/3; \nu) > 0$. Hence $d_a^* = 0$.

Fee. Consider first the case with $d_a = 0$. Since the fee is uncorrelated with quality and the base price, $(\Sigma d)_a = 0$ for any weights on base price and quality. Hence the objective reduces to that of the two-attribute environment $\mathcal{S} = \{q, b\}$ with covariance matrix $\Sigma_{\mathcal{S}\mathcal{S}}$ and $m_{\mathcal{S}} = \Sigma_{\mathcal{S}\mathcal{S}}\iota_{\mathcal{S}}$. The best weights with $d_a = 0$ are thus the optimally adapted weights (d_q, d_b) of this environment. When both receive weight, the equalities of Lemma 3 are $\Sigma_{\mathcal{S}\mathcal{S}}\iota_{\mathcal{S}} = \xi(\Sigma_{\mathcal{S}\mathcal{S}} d_{\mathcal{S}} + \nu^2 \iota_{\mathcal{S}})$, so $d_{\mathcal{S}} = \xi^{-1}\iota_{\mathcal{S}} - \nu^2 \Sigma_{\mathcal{S}\mathcal{S}}^{-1}\iota_{\mathcal{S}}$. Summing gives $\xi = 2/(1 + \nu^2 \gamma_{\mathcal{S}})$, and the two weights differ by

$$d_b - d_q = \nu^2[(\Sigma_{\mathcal{S}\mathcal{S}}^{-1}\iota_{\mathcal{S}})_q - (\Sigma_{\mathcal{S}\mathcal{S}}^{-1}\iota_{\mathcal{S}})_b] = \frac{\nu^2((1-\alpha)\sigma_p^2 - \sigma_q^2)}{(1-\rho^2-\alpha)\sigma_q^2\sigma_p^2},$$

which we denote by ζ. The base price alone receives weight if and only if $\zeta \geq 1$, and quality alone receives weight if and only if $\zeta \leq -1$. Hence $d_b = \min\{1, \max\{0, (1+\zeta)/2\}\}$ and $d_q = 1 - d_b$.

Extended by $d_a = 0$, these weights also satisfy the conditions of Lemma 3 at q and b for the three attributes, with the same ξ. Since $m_a = \alpha\sigma_p^2$ and $(\Sigma d)_a = 0$, the condition at a is $\alpha\sigma_p^2 \leq \nu^2\xi$, so the fee receives no weight if and only if $\alpha\sigma_p^2 \leq \nu^2\xi$. The left side is zero at $\alpha = 0$ and strictly increasing in α. The right side is positive and continuous in α and ν, since $\xi = m^\top d/(d^\top \Sigma d + \nu^2)$ and the weights are continuous. It is also non-increasing in α and increasing in ν. To see this, note that by the equalities of Lemma 3 it equals

$$\frac{2}{\nu^{-2} + \gamma_{\mathcal{S}}}, \qquad \frac{\nu^2 \operatorname{Cov}(q, U)}{\sigma_q^2 + \nu^2} \qquad \text{or} \qquad \frac{\nu^2 \operatorname{Cov}(b, U)}{(1-\alpha)\sigma_p^2 + \nu^2}$$

when both attributes, quality alone or the base price alone receive weight. Each increases in ν, the last two because their covariances are positive where they apply. In α, the first does not rise, because $\gamma_{\mathcal{S}}^{-1}$ is the smallest variance of $d_q q + d_b b$ with $d_q + d_b = 1$, which does not rise when the variance $(1-\alpha)\sigma_p^2$ of the base price falls. The second is constant, and the third falls with $(1-\alpha)\sigma_p^2$, as $\operatorname{Cov}(b, U) = \rho\sigma_q\sigma_p + (1-\alpha)\sigma_p^2$ and $\rho\sigma_q\sigma_p \leq 0 < \nu^2$. Since ζ is a ratio of two affine functions of α and proportional to ν^2, each case holds on an interval of shares and on an interval of error levels, so by continuity $\nu^2\xi$ has these properties throughout. Hence the shares at which the fee receives no weight form an initial

interval, which contains a neighborhood of 0 and grows with ν. Since $\nu^2\xi$ is continuous, this interval is closed in $[0, \alpha_{\max})$, so the fee receives no weight exactly at the shares $\alpha \le \bar{\alpha}(\nu)$ for some $\bar{\alpha}(\nu) > 0$ non-decreasing in ν, with $\bar{\alpha}(\nu) = \alpha_{\max}$ if it receives weight at no share.

Finally, let both price and quality receive weight under all-in display, that is, $|\zeta| < 1$ at $\alpha = 0$. Written out, the condition is $\nu^2|\sigma_p^2 - \sigma_q^2| < (1-\rho^2)\sigma_q^2\sigma_p^2$. Then the first formula applies at $\alpha = 0$, and since $\nu^2\xi$ is non-increasing, at every share it is at most

$$\frac{2\nu^2(1-\rho^2)\sigma_q^2\sigma_p^2}{(1-\rho^2)\sigma_q^2\sigma_p^2 + \nu^2(\sigma_q^2 + \sigma_p^2 - 2\rho\sigma_q\sigma_p)} < (1-\rho^2)\sigma_p^2 = \alpha_{\max}\sigma_p^2.$$

The inequality rearranges to $\nu^2(\sigma_q^2 - \sigma_p^2 + 2\rho\sigma_q\sigma_p) < (1-\rho^2)\sigma_q^2\sigma_p^2$, which follows from the hypothesis since $\rho \le 0$. Hence the fee receives weight at every share at which $\alpha\sigma_p^2$ exceeds the left-hand side, and $\bar{\alpha}(\nu) < \alpha_{\max}$. □

Proof of Proposition 7. Tax. By the proof of Proposition 6, the tax receives no weight, so the weight on the base price is $e^*/(1-\alpha)$ and $d_q^* + e^*/(1-\alpha) = 1$. Applying Lemma 3 results in the appropriate KKT conditions for this case, but since $b = (1-\alpha)p$, the condition at b needs to be divided by $1-\alpha$. They are

$$\operatorname{Cov}(q, U) \le \xi(\sigma_q^2 d_q^* + \rho\sigma_q\sigma_p e^* + \nu^2) \text{ and } \operatorname{Cov}(p, U) \le \xi\Big(\rho\sigma_q\sigma_p d_q^* + \sigma_p^2 e^* + \frac{\nu^2}{1-\alpha}\Big),$$

with equality where the weight is positive. When both hold with equality, eliminating ξ and using $\sigma_q^2 \operatorname{Cov}(p, U) - \rho\sigma_q\sigma_p \operatorname{Cov}(q, U) = (1-\rho^2)\sigma_q^2\sigma_p^2 = \sigma_p^2 \operatorname{Cov}(q, U) - \rho\sigma_q\sigma_p \operatorname{Cov}(p, U)$ shows that $e^* - d_q^* = \frac{\nu^2\left(\operatorname{Cov}(p,U) - \operatorname{Cov}(q,U)/(1-\alpha)\right)}{(1-\rho^2)\sigma_q^2\sigma_p^2}$, which we denote by ω. Together with $d_q^* + e^*/(1-\alpha) = 1$ this gives $e^*/d_q^* = (1+\omega)/(1-\omega/(1-\alpha))$. By additionally substituting ξ similarly from the equality into the inequality, it follows that quality alone receives weight if and only if $\omega \le -1$ and the price alone if and only if $\omega \ge 1 - \alpha$.

If $\operatorname{Cov}(p, U) \le 2\operatorname{Cov}(q, U)$, then $\operatorname{Cov}(q, U) > 0$, as $\operatorname{Cov}(q, U) + \operatorname{Cov}(p, U) = \operatorname{Var}(U) > 0$. Hence ω strictly decreases in α, while $\omega/(1-\alpha)$ does not increase, since its derivative is $\nu^2((1-\alpha)\operatorname{Cov}(p, U) - 2\operatorname{Cov}(q, U))/((1-\alpha)^3(1-\rho^2)\sigma_q^2\sigma_p^2)$ and $(1-\alpha)\operatorname{Cov}(p, U) \le \max\{\operatorname{Cov}(p, U), 0\} \le 2\operatorname{Cov}(q, U)$. Relative price sensitivity is thus ∞ on an initial interval of shares, where $\omega/(1-\alpha) \ge 1$, and 0 on a final one, where $\omega \le -1$. In between it is $(1+\omega)/(1-\omega/(1-\alpha))$, whose numerator is positive and strictly decreasing and whose denominator is positive and non-decreasing, so it is non-increasing in α.

Otherwise $\operatorname{Cov}(p, U) > 2\operatorname{Cov}(q, U)$, and we will show that the derivative of relative price sensitivity at $\alpha = 0$ is positive for some ν. First, since $\operatorname{Var}(U) > 0$, $\operatorname{Cov}(p, U) >$

$2\,\mathrm{Cov}(q,U)$ implies $\mathrm{Cov}(p,U) > \mathrm{Cov}(q,U)$. Therefore $\omega > 0$ at $\alpha = 0$, and the derivative of $\omega/(1-\alpha)$ above is positive there. At $\alpha = 0$ both attributes thus receive weight if and only if $\omega < 1$, that is, $\nu^2 < (1-\rho^2)\sigma_q^2\sigma_p^2/(\mathrm{Cov}(p,U) - \mathrm{Cov}(q,U))$. For simplicity we now consider $\log(e^*/d_q^*)$. It equals $\log(1+\omega) - \log(1-\omega/(1-\alpha))$, and thus its derivative is

$$\frac{\partial\omega/\partial\alpha}{1+\omega} + \frac{\partial(\omega/(1-\alpha))/\partial\alpha}{1-\omega/(1-\alpha)}.$$

As ν^2 approaches the above bound from below, the denominator $1-\omega$ of the second term tends to zero at $\alpha = 0$, while the first term stays bounded and the numerator of the second tends to a positive limit. Hence for ν' with ν'^2 slightly below the bound, relative price sensitivity rises at $\alpha = 0$.

Fee. For $\alpha < \bar{\alpha}(\nu)$ the fee receives no weight, so by the proof of Proposition 6, $d_b^* = \min\{1, \max\{0, (1+\zeta)/2\}\}$, $d_q^* = 1 - d_b^*$ and $e^* = (1-\alpha)d_b^*$. Since the derivative of ζ in α has the sign of $\rho^2\sigma_p^2 - \sigma_q^2$, ζ is non-increasing in α if and only if $\sigma_q \geq |\rho|\sigma_p$. This condition is exactly $\mathrm{Cov}(q,U) = \sigma_q(\sigma_q + \rho\sigma_p) \geq 0$, and in this case d_b^* is non-increasing. Hence, $e^* = (1-\alpha)d_b^*$ is non-increasing, and similarly e^*/d_q^*, because its denominator $d_q^* = 1 - d_b^*$ is non-decreasing. Moreover, it is ∞ on the initial interval of shares where $d_q^* = 0$.

Otherwise $\mathrm{Cov}(q,U) < 0$, and as for the tax we will show that relative price sensitivity rises at $\alpha = 0$ for ν just below the level at which quality stops receiving weight under all-in display. First, ζ now strictly increases in α, and at $\alpha = 0$ it equals $\nu^2(\sigma_p^2 - \sigma_q^2)/((1-\rho^2)\sigma_q^2\sigma_p^2) > 0$, since $\sigma_q < |\rho|\sigma_p \leq \sigma_p$. At $\alpha = 0$ both quality and the base price thus receive weight if and only if $\zeta < 1$, that is, $\nu^2 < (1-\rho^2)\sigma_q^2\sigma_p^2/(\sigma_p^2 - \sigma_q^2)$, and then $e^*/d_q^* = (1-\alpha)(1+\zeta)/(1-\zeta)$. The derivative of its logarithm is

$$-\frac{1}{1-\alpha} + \frac{2\,\partial\zeta/\partial\alpha}{1-\zeta^2}.$$

As ν^2 approaches the above bound from below, $1-\zeta^2$ tends to zero at $\alpha = 0$, while $\partial\zeta/\partial\alpha$ tends to a positive limit. Hence for ν' with ν'^2 slightly below the bound relative price sensitivity rises at $\alpha = 0$. □

Proof of Proposition 8. By (15), expected utility at the optimally adapted weights is $\iota^\top\mu + R(d^*;\nu)/\sqrt{\pi}$, and $\iota^\top\mu$ does not depend on α. **Tax.** By the proof of Proposition 6, $R(d^*;\nu)$ is the maximum of the ratio objective of the two-attribute environment (q,p), which does not depend on α, over the weights $(d_q, e) \geq 0$ with $d_q + e/(1-\alpha) = 1$, and by Lemma 4(ii) also over those with $d_q + e/(1-\alpha) \leq 1$. This set shrinks as α rises, so

$R(d^*;\nu)$ is non-increasing in α. **Fee.** For any $d \in \Delta^3$, $m^\top d$ and $d^\top \Sigma d$ are affine in α by (32), with slopes $\sigma_p^2(d_a - d_b)$ and $\sigma_p^2(d_a^2 - d_b^2)$. Since the maximizer d^* is unique, the envelope theorem gives, for $\alpha < \bar{\alpha}(\nu)$ where $d_a^* = 0$,

$$\frac{dR(d^*;\nu)}{d\alpha} = -\frac{\sigma_p^2 d_b^*}{\sqrt{d^{*\top}\Sigma d^* + \nu^2}}\Big(1 - \frac{\xi d_b^*}{2}\Big) \leq 0,$$

with $\xi = m^\top d^*/(d^{*\top}\Sigma d^* + \nu^2)$ the multiplier of Lemma 3. The sign holds because $\xi d_b^* < 2$ wherever $d_b^* > 0$. To see this, note that $d_b^* \leq 1$ and that, by the proof of Proposition 6, ξ is then either $2/(1 + \nu^2\gamma_{\{q,b\}}) < 2$, when quality also receives weight, or $\mathrm{Cov}(b,U)/((1-\alpha)\sigma_p^2 + \nu^2)$, when the base price alone does. The latter is below 1, since $\mathrm{Cov}(b,U) = \rho\sigma_q\sigma_p + (1-\alpha)\sigma_p^2 \leq (1-\alpha)\sigma_p^2$. □

D.2 Proofs: Additional Theoretical Results

D.2.1 Proofs: Comparing *K2* and *K3*

Proof of Proposition 9. In *K2* the weights sum to 2/3, which by Lemma 4 is equivalent to scaling the errors. Hence the weights in *K2* are $\frac{2}{3}$ times the optimally adapted weights d^* at level $\frac{3}{2}\nu$, and those in *K3* are d^* at level ν. As in the proof of Corollary 2, the weights in *B-HighVar* are those in *A-HighVar* with the weights on A and B swapped. For *K2* ($K = 2$) and *K3* ($K = 3$), let $\bar{d}_K$ be the average and δ_K half the difference of the weights on A and B in *A-HighVar*, so that these weights are $\bar{d}_K + \delta_K$ and $\bar{d}_K - \delta_K$, and swapped in *B-HighVar*. Since attributes are independent and all weights are interior, (6) holds with $J = I$. With A of variance σ_h^2 and B and C of variance σ_l^2 in *A-HighVar*, it gives

$$\delta_2 = \frac{2}{3}\cdot\frac{1}{2}\Big(\frac{3}{2}\nu\Big)^2(\sigma_l^{-2} - \sigma_h^{-2}) = \frac{3}{4}\nu^2(\sigma_l^{-2} - \sigma_h^{-2}) > \delta_3 = \frac{1}{2}\nu^2(\sigma_l^{-2} - \sigma_h^{-2}) > 0$$

and $\bar{d}_2 = \frac{1}{3} < \bar{d}_3 = \frac{1}{2}(1 - d_C^*) = \frac{1}{3} + \frac{\nu^2}{6}(\sigma_l^{-2} - \sigma_h^{-2})$. Moreover, $\delta_K < \bar{d}_K$, since the weight on B is positive.

In *A-HighVar* the bias is $(\bar{d}_K + \delta_K)/(\bar{d}_K - \delta_K)$ and in *B-HighVar* its reciprocal, so the difference between the conditions is

$$\frac{\bar{d}_K + \delta_K}{\bar{d}_K - \delta_K} - \frac{\bar{d}_K - \delta_K}{\bar{d}_K + \delta_K} = \frac{4\delta_K/\bar{d}_K}{1 - (\delta_K/\bar{d}_K)^2}$$

for the bias, which is strictly increasing in $\delta_K/\bar{d}_K \in (0,1)$. Since $\bar{d}_2 < \bar{d}_3$ and $\delta_2 > \delta_3$, this proves the first part.

Now fix a choice with $z_A = x_A - y_A > 0$, $z_B = x_B - y_B < 0$ and, in *K3*, $z_C = 0$,

and set $\bar{z} = (z_A + z_B)/(\sqrt{2}\nu)$ and $\tilde{z} = (z_A - z_B)/(\sqrt{2}\nu) > 0$. By (8), x is chosen with probability $\Phi(\bar{d}_K\bar{z}+\delta_K\tilde{z})$ in *A-HighVar* and $\Phi(\bar{d}_K\bar{z}-\delta_K\tilde{z})$ in *B-HighVar*. Their difference is the integral of ϕ over an interval with midpoint $\bar{d}_K\bar{z}$ and length $2\delta_K\tilde{z}$, which is strictly increasing in the length and non-increasing in the absolute value of the midpoint, since ϕ is symmetric and decreasing on $[0,\infty)$. With $\delta_3 < \delta_2$ and $|\bar{d}_3\bar{z}| \geq |\bar{d}_2\bar{z}|$ this gives

$$\Phi(\bar{d}_3\bar{z}+\delta_3\tilde{z}) - \Phi(\bar{d}_3\bar{z}-\delta_3\tilde{z}) \leq \Phi(\bar{d}_2\bar{z}+\delta_3\tilde{z}) - \Phi(\bar{d}_2\bar{z}-\delta_3\tilde{z}) < \Phi(\bar{d}_2\bar{z}+\delta_2\tilde{z}) - \Phi(\bar{d}_2\bar{z}-\delta_2\tilde{z}).$$

□

D.2.2 Proofs: Relation to Models based on Perception Errors

Proof of Theorem 2. Let F be the set of attributes perceived with finite errors and $z = x - y$. If F is not all attributes, (11) is the limit in which the diagonal entries of Ψ off F grow without bound, and by block inversion $(\Sigma + \Psi)^{-1}$ then tends to $(\Sigma_{FF} + \Psi_{FF})^{-1}$ on $F \times F$ and to zero elsewhere. The DM thus compares $\iota^\top \hat{x} = \iota^\top \mu + w^\top(r_x - \mu)$ across goods, with signal weights $w_F = (\Sigma_{FF} + \Psi_{FF})^{-1} m_F$ on F and $w_i = 0$ off F. Since the posterior covariance is $\Sigma - \Sigma(\Sigma+\Psi)^{-1}\Sigma$, the covariances with utility after perception are $h_i = \mathrm{Cov}(x_i, U(x) \mid r_x) = m_i - (\Sigma w)_i$, and by the definition of w_F they are $h_F = \Psi_{FF} w_F$ on F.

Writing $\eta_x = r_x - x$ and $\eta_y = r_y - y$ for the perceptual errors, the DM chooses x if $w^\top(r_x - r_y) = w^\top z + w_F^\top(\eta_{x,F} - \eta_{y,F})$ is positive, and given the goods the last term is $\mathcal{N}(0, 2w_F^\top \Psi_{FF} w_F) = \mathcal{N}(0, 2w^\top h)$. Hence x is chosen with probability $\Phi(w^\top z/\sqrt{2w^\top h})$ if $w \neq 0$ and $\frac{1}{2}$ if $w = 0$, while in the aggregation model at level $\nu > 0$ it is $\Phi(d^{*\top} z/(\sqrt{2}\nu))$ by (8), which is not constant in z. Since z ranges over all of $\mathbb{R}^I$, the two models are equivalent at ν if and only if $w \neq 0$ and $w/\sqrt{w^\top h} = d^*/\nu$.

This holds if and only if $w = \xi d^*$, where ξ is the multiplier of Lemma 3. If $w = \xi d^*$, then $h_i = m_i - \xi(\Sigma d^*)_i = \xi\nu^2$ for i in the support $\mathcal{S}$ of d^* by Lemma 3, so $w^\top h = \xi^2\nu^2$ and $w/\sqrt{w^\top h} = d^*/\nu$. Conversely, $w/\sqrt{w^\top h} = d^*/\nu$ gives $w = (\iota^\top w)d^*$ with $\iota^\top w = \sqrt{w^\top h}/\nu$, and then $\nu^2(\iota^\top w)^2 = w^\top h = \iota^\top w\,(m^\top d^* - \iota^\top w\, d^{*\top}\Sigma d^*)$ shows $\iota^\top w = m^\top d^*/(d^{*\top}\Sigma d^* + \nu^2)$, which is ξ by Lemma 3.

Part A. Here Ψ is diagonal with entries ν^2/d_i^*, so $F = \mathcal{S}$. For $i \in \mathcal{S}$, $((\Sigma_{FF}+\Psi_{FF})\,\xi d_F^*)_i = \xi(\Sigma d^*)_i + \xi\nu^2 = m_i$, where the first equality uses that d^* vanishes off $\mathcal{S}$ and that $\Psi_{ii} d_i^* = \nu^2$, and the second is Lemma 3. Since $\Sigma_{FF} + \Psi_{FF}$ is invertible, $w_F = \xi d_F^*$, and $w = \xi d^*$ because both vanish off $\mathcal{S}$. By the criterion above, the two models are equivalent.

Part B. The statement presupposes that at least one attribute is perceived with finite errors, so $F \neq \emptyset$. By (11), $\mathbb{E}[U(x) \mid r_x] = \iota^\top \hat{x}$ has slope w_i in the signal of attribute i, so the hypothesis says that $w_i > 0$ for every $i \in F$. Hence $\{i : w_i > 0\} = F$ and $\iota^\top w > 0$.

(⇒) Let the perceptual model be equivalent to the aggregation model at some level $\nu > 0$, and let d^*, $\mathcal{S}$ and ξ refer to this level. By the criterion above, $w = \xi d^*$, so $\mathcal{S} = \{i : w_i > 0\} = F$. Hence $h_i = \xi\nu^2$ for every $i \in F$ and $h_i \leq \xi\nu^2$ for every $i \notin F$.

(⇐) Let $h_i = \bar{h}$ for all $i \in F$ and $h_i \leq \bar{h}$ for all $i \notin F$. Then $\bar{h}\,\iota^\top w = \sum_{i \in F} h_i w_i = w^\top h > 0$, so $\bar{h} > 0$. Set $\zeta = \iota^\top w$, $d = w/\zeta \in \Delta^I$ and $\nu = \sqrt{\bar{h}/\zeta} > 0$. Since $h = m - \Sigma w = m - \zeta \Sigma d$ and $d_i > 0$ exactly for $i \in F$, the assumption reads $m_i = \zeta(\Sigma d)_i + \zeta\nu^2$ for $d_i > 0$ and $m_i \leq \zeta(\Sigma d)_i + \zeta\nu^2$ for $d_i = 0$. These are the conditions of Lemma 3 with multiplier ζ, so d maximizes $R(\cdot;\nu)$ and, by the uniqueness in Lemma 2, $d = d^*$ at this level. Since the multiplier is unique, $\zeta = \xi$, so $w = \xi d^*$, and by the criterion above the perceptual model is equivalent to the aggregation model at this level. □

D.3 Proofs: Extensions and Model Variant Results

Proof of Proposition 10. Let $x^1, \ldots, x^K$ be independent draws from $\mathcal{N}(\mu, \Sigma)$ and $\epsilon_1, \ldots, \epsilon_K$ independent $\mathcal{N}(0, \nu^2)$ errors, independent of the goods. Write $U_k = \iota^\top x^k$ for the utility of good k, $W_k = d^\top x^k + \epsilon_k$ for its choice variable, and $W = (W_1, \ldots, W_K)$. The DM chooses $k^* = \arg\max_k W_k$. Each W_k has variance $\sigma_d^2 = d^\top \Sigma d + \nu^2 > 0$, so ties occur with probability zero. The pairs (U_k, W_k) are independent across k and jointly normal with $\mathrm{Cov}(U_k, W_k) = \iota^\top \Sigma d$, so

$$\mathbb{E}[U_k \mid W] = \mathbb{E}[U_k \mid W_k] = \iota^\top \mu + \frac{\iota^\top \Sigma d}{\sigma_d^2}\,(W_k - d^\top \mu).$$

Since k^* is a function of W,

$$\begin{aligned}\mathbb{E}[U_{k^*}] &= \mathbb{E}\Big[\sum_k \mathbf{1}\{k^* = k\}\ \mathbb{E}[U_k \mid W]\Big] = \iota^\top \mu + \frac{\iota^\top \Sigma d}{\sigma_d^2}\ \mathbb{E}\,[\max_k (W_k - d^\top \mu)] \\ &= \iota^\top \mu + c_K\, R(d;\nu),\end{aligned}$$

where $c_K = \mathbb{E}[\max\{Z_1, \ldots, Z_K\}]$ for independent standard normal $Z_1, \ldots, Z_K$, because the variables $(W_k - d^\top \mu)/\sigma_d$ are independent and standard normal. This is (12), and c_K depends on K alone and is strictly increasing in it. Hence for every $K \geq 2$ the maximizers of (12) over Δ^I are those of $R(\cdot;\nu)$, and the optimally adapted weights do not depend on K. Since $R(d^*;\nu) \geq R(\iota/I;\nu) > 0$, expected utility at the optimally adapted weights

is strictly increasing in K. □

Proof of Proposition 11. Let $\epsilon \sim \mathcal{N}(0, \nu^2)$ be the aggregation error, independent of x. The DM chooses the good if and only if $W = d^\top(x - x^0) + \epsilon \geq 0$, so her expected utility gain over the outside option is $\mathbb{E}[\iota^\top(x - x^0)\,\mathbf{1}\{W \geq 0\}]$. Since $\iota^\top(x - x^0)$ and W are jointly normal with means $\iota^\top\delta$ and $d^\top\delta$, covariance $\iota^\top\Sigma d$ and $\mathrm{Var}(W) = d^\top\Sigma d + \nu^2 > 0$, Fact 1 gives (13). □

Proof of Theorem 3. By (15) of Lemma 1, which holds at every $d \neq 0$, the optimally adapted weights maximize $R(\cdot;\nu)$ over the d with $\iota^\top|d| = 1$. For these d, $d^\top\Sigma d + \nu^2 = d^\top\Sigma d + \nu^2(\iota^\top|d|)^2$, so the proof of Lemma 2 carries over with $a^\top\Sigma a + \nu^2(\iota^\top|a|)^2$ in place of $a^\top\Omega a = a^\top\Sigma a + \nu^2(\iota^\top a)^2$, with $\iota^\top|a|$ in place of $\iota^\top a$ and without $a \geq 0$. The reason for uniqueness changes. In this case, the midpoint of two maximizers would satisfy the constraint with strict inequality, since $a^\top\Sigma a$ is strictly convex and $(\iota^\top|a|)^2$ is convex. If there were a maximizer with strict inequality, then the same weight proportions scaled up would be admissible and raise the value $m^\top a$, a contradiction. Hence d^* exists and is unique, and so is its signed support. Moreover, d is optimal if and only if $a = d/\sqrt{d^\top\Sigma d + \nu^2}$ maximizes $m^\top a$ subject to $a^\top\Sigma a + \nu^2(\iota^\top|a|)^2 \leq 1$.

This program is convex and satisfies Slater's condition at 0, and the constraint binds at a. As in the proof of Lemma 3, a is thus a maximizer if and only if m is a non-negative multiple of a subgradient of the constraint function at a. These subgradients are $2\Sigma a + 2\nu^2(\iota^\top|a|)g$, where $g_i = \mathrm{sgn}(d_i)$ if $d_i \neq 0$ and $|g_i| \leq 1$ if $d_i = 0$. Note that they are proportional to $\Sigma d + \nu^2 g$, and therefore d is optimal if and only if $m = \xi(\Sigma d + \nu^2 g)$ for some $\xi > 0$ and such a g. Finally, the signed KKT conditions are given by

$$m_i = \xi(\Sigma d)_i + \xi\nu^2\,\mathrm{sgn}(d_i) \quad (d_i \neq 0), \qquad |m_i - \xi(\Sigma d)_i| \leq \xi\nu^2 \quad (d_i = 0). \tag{19}$$

Part A. Let $(\mathcal{S}, s)$ be the signed support of d^* and ξ as in (19). Since d^* vanishes off $\mathcal{S}$, the equalities in (19) are $m_\mathcal{S} = \xi(\Sigma_{\mathcal{S}\mathcal{S}}d^*_\mathcal{S} + \nu^2 s)$, and $s^\top d^*_\mathcal{S} = \iota^\top|d^*| = 1$. With s in place of $\iota_\mathcal{S}$, these are the equalities and the sum used in the proof of Theorem 1A, and since $s^\top\Sigma^{-1}_{\mathcal{S}\mathcal{S}}s > 0$ the same steps give $s^\top\Sigma^{-1}_{\mathcal{S}\mathcal{S}}m_\mathcal{S} > 0$,

$$\xi = \frac{s^\top\Sigma^{-1}_{\mathcal{S}\mathcal{S}}m_\mathcal{S}}{1 + \nu^2\gamma_\mathcal{S}} \quad \text{and} \quad d^*_\mathcal{S} = \alpha_\mathcal{S} + \nu^2\gamma_\mathcal{S}(\alpha_\mathcal{S} - \beta_\mathcal{S}),$$

where $\alpha_\mathcal{S}$, $\beta_\mathcal{S}$ and $\gamma_\mathcal{S}$ are defined with s in place of $\iota_\mathcal{S}$. In the remainder of this proof we will use these alternative vectors.

Part B. As in the proof of Theorem 1B, there are finitely many signed supports, and

it suffices to show that each occurs on a single interval of ν^2, on which the weights are affine in ν^2. Fix $(\mathcal{S}, s)$ with $s^\top \Sigma_{\mathcal{S}\mathcal{S}}^{-1} m_{\mathcal{S}} > 0$ and let $d^{(\mathcal{S},s)}$ equal $\alpha_{\mathcal{S}} + \nu^2 \gamma_{\mathcal{S}}(\alpha_{\mathcal{S}} - \beta_{\mathcal{S}})$ on $\mathcal{S}$ and 0 off $\mathcal{S}$. As in the proof of Theorem 1B, it satisfies $s^\top d_{\mathcal{S}}^{(\mathcal{S},s)} = 1$ and the equalities in (19) with ξ as in Part A. By Part A and (19), $(\mathcal{S}, s)$ is thus the signed support at ν if and only if

$$s_i d_i^{(\mathcal{S},s)} > 0 \quad (i \in \mathcal{S}) \text{ and } |(1 + \nu^2 \gamma_{\mathcal{S}})\, m_i - s^\top \Sigma_{\mathcal{S}\mathcal{S}}^{-1} m_{\mathcal{S}}\, (\Sigma d^{(\mathcal{S},s)})_i| \leq s^\top \Sigma_{\mathcal{S}\mathcal{S}}^{-1} m_{\mathcal{S}}\, \nu^2 \quad (i \notin \mathcal{S}),$$

where the second condition is the inequality in (19) multiplied by $1+\nu^2\gamma_{\mathcal{S}}$. All conditions are affine in ν^2; for the second, this applies after splitting the absolute value into two inequalities. Hence, all conditions hold on an interval, on which $d^* = d^{(\mathcal{S},s)}$ is affine in ν^2. Therefore, the signed support changes finitely many times, and the weights are smooth in ν^2 except at these changes. By Berge's maximum theorem they are also continuous, since R is jointly continuous in (d, ν) for $d \neq 0$, the set of d with $\iota^\top |d| = 1$ is compact and the maximizer is unique.

It remains to find the threshold. Since finitely many intervals partition $[0, \infty)$, exactly one of them is unbounded. Let $(\mathcal{S}, s)$ be the signed support on it and choose $\overline{\nu}'$ such that it contains all $\nu^2 > \overline{\nu}'^2$. There $d^* = d^{(\mathcal{S},s)}$ is affine in ν^2 but also bounded, because $\iota^\top |d^*| = 1$. Its slope $\gamma_{\mathcal{S}}(\alpha_{\mathcal{S}} - \beta_{\mathcal{S}})$ therefore must equal 0. So d^* is constant above $\overline{\nu}'$ and $\alpha_{\mathcal{S}} = \beta_{\mathcal{S}}$, that is, $m_i = s_i\, s^\top \Sigma_{\mathcal{S}\mathcal{S}}^{-1} m_{\mathcal{S}} / \gamma_{\mathcal{S}}$ for all $i \in \mathcal{S}$. Hence $|m_i|$ is the same for all $i \in \mathcal{S}$, and the weight on i carries the sign s_i of $m_i = \iota^\top \Sigma e_i$. For $i \notin \mathcal{S}$, dividing the inequality above by ν^2 and letting $\nu \to \infty$ gives $|m_i| \leq s^\top \Sigma_{\mathcal{S}\mathcal{S}}^{-1} m_{\mathcal{S}} / \gamma_{\mathcal{S}}$. Hence $\mathcal{S} \subseteq \arg\max_i |\iota^\top \Sigma e_i|$. □

Proof of Proposition 12. Let d^* be the optimally adapted weights of the main model and ξ as in Lemma 3. Since $d^* \geq 0$, it is feasible in the model variant, and the conditions of Lemma 3 give all conditions in (19) except the lower bound $m_i - \xi(\Sigma d^*)_i \geq -\xi\nu^2$ for $d_i^* = 0$. If d^* has full support, there is no such i. If attributes are independent, then $(\Sigma d^*)_i = \sigma_i^2 d_i^* = 0$ for such i, so $m_i - \xi(\Sigma d^*)_i = \sigma_i^2 > 0$. In both cases d^* thus satisfies (19) and, by the uniqueness in Theorem 3, equals the optimally adapted weights of the model variant. For every $\nu < \underline{\nu}$ the support is full by Theorem 1B. □

D.4 Proofs: Supporting Results

Proof of Lemma 1. Let $z = x - y$ and $\epsilon_z = \epsilon_x - \epsilon_y$. Since $x, y \overset{iid}{\sim} \mathcal{N}(\mu, \Sigma)$ and $\epsilon_x, \epsilon_y \overset{iid}{\sim} \mathcal{N}(0, \nu^2)$ are mutually independent, $z \sim \mathcal{N}(0, 2\Sigma)$ and $\epsilon_z \sim \mathcal{N}(0, 2\nu^2)$ are independent,

and $U(x) - U(y) = \iota^\top z$. Since ties have probability zero, x is chosen if and only if $d^\top z + \epsilon_z \geq 0$. For $w \in \mathbb{R}^I \setminus \{0\}$, the pair $(w^\top z,\, d^\top z + \epsilon_z)$ is centered jointly normal with covariance $2w^\top \Sigma d$ and variances $2w^\top \Sigma w$ and $2(d^\top \Sigma d + \nu^2)$, so its correlation is

$$\frac{w^\top \Sigma d}{\sqrt{w^\top \Sigma w}\,\sqrt{d^\top \Sigma d + \nu^2}}.$$

Now $\mathbf{1}\{\text{choose } y\} = 1 - \mathbf{1}\{\text{choose } x\}$, and therefore the objective in (2) equals

$$\mathbb{E}[U(y)] + \mathbb{E}\left[\iota^\top z\, \mathbf{1}\{d^\top z + \epsilon_z \geq 0\}\right] = \iota^\top \mu + \frac{2\iota^\top \Sigma d}{\sqrt{2(d^\top \Sigma d + \nu^2)}}\,\phi(0) = \iota^\top \mu + \frac{R(d;\nu)}{\sqrt{\pi}},$$

where the equalities follow from Fact 1 with $V = \iota^\top z$ and $W = d^\top z + \epsilon_z$, which have zero means, and from $\sqrt{2}\,\phi(0) = 1/\sqrt{\pi}$. This is (15), and since $\iota^\top \mu$ does not depend on d, the two problems have the same maximizers.

The choice maximizes utility if and only if $\iota^\top z$ and $d^\top z + \epsilon_z$ have the same sign, and it maximizes $D(x)$ if and only if $d^\top z$ and $d^\top z + \epsilon_z$ have the same sign. By the classical orthant-probability formula, two centered jointly normal variables have the same sign with probability $\frac{1}{2} + \frac{1}{\pi}\arcsin r$, where r is their correlation. With $w = \iota$ the correlation above is $R(d;\nu)/\sqrt{\iota^\top \Sigma \iota}$, and with $w = d$ it is $\sqrt{d^\top \Sigma d/(d^\top \Sigma d + \nu^2)}$, which gives (16) and (17). □

Proof of Lemma 2. Since $\iota^\top d = 1$ on Δ^I, $d^\top \Omega d = d^\top \Sigma d + \nu^2$ and hence $R(d;\nu) = m^\top d/\sqrt{d^\top \Omega d}$ there, a ratio that is homogeneous of degree zero in d. Every feasible $a \neq 0$ of (18) has $\iota^\top a > 0$, so the maps $d \mapsto d/\sqrt{d^\top \Omega d}$ and $a \mapsto a/\iota^\top a$ are mutually inverse bijections between Δ^I and the feasible a with $a^\top \Omega a = 1$, and $m^\top a = R(a/\iota^\top a;\nu)$ at these a. It therefore suffices to show that (18) has a unique maximizer a^* and that the constraint binds at it. In this case $d^* = a^*/(\iota^\top a^*)$ is the unique maximizer of $R(\cdot;\nu)$ over Δ^I.

Since $\Omega = \Sigma + \nu^2 \iota\iota^\top$ is positive definite, the feasible set is compact, so a maximizer exists, and its value is positive, since $m^\top \iota = \iota^\top \Sigma \iota > 0$. The constraint binds at every maximizer, since otherwise scaling it up would be feasible and raise the value. If $a_1 \neq a_2$ were two maximizers, their midpoint $\bar{a}$ would be feasible and attain the same value. By the parallelogram law, however, $\bar{a}^\top \Omega \bar{a} = 1 - (a_1 - a_2)^\top \Omega (a_1 - a_2)/4 < 1$, so $\bar{a}$ would be a maximizer at which the constraint does not bind, a contradiction. □

Proof of Lemma 3. Fix $d \in \Delta^I$ and let $a = d/\sqrt{d^\top \Omega d}$. By the proof of Lemma 2, d maximizes $R(\cdot;\nu)$ over Δ^I if and only if a solves (18). This is a convex program

that satisfies Slater's condition, for example at $\iota/(2\sqrt{\iota^\top \Omega \iota})$. Moreover, its quadratic constraint binds at a, so complementary slackness holds for every multiplier. Since $\Omega a = (\Sigma d + \nu^2 \iota)/\sqrt{d^\top \Omega d}$, a thus solves (18) if and only if there is $\xi \geq 0$ with $m_i = \xi(\Sigma d)_i + \xi\nu^2$ whenever $a_i > 0$ and $m_i \leq \xi(\Sigma d)_i + \xi\nu^2$ whenever $a_i = 0$, where $\xi\sqrt{d^\top \Omega d}/2$ is the multiplier of the quadratic constraint. Since $a_i > 0$ if and only if $d_i > 0$, these are the conditions of the lemma. Finally, $\xi > 0$, since $\xi = 0$ would give $m \leq 0$, contradicting $\iota^\top m = \iota^\top \Sigma \iota > 0$. □

Proof of Lemma 4. Dividing numerator and denominator by λ gives

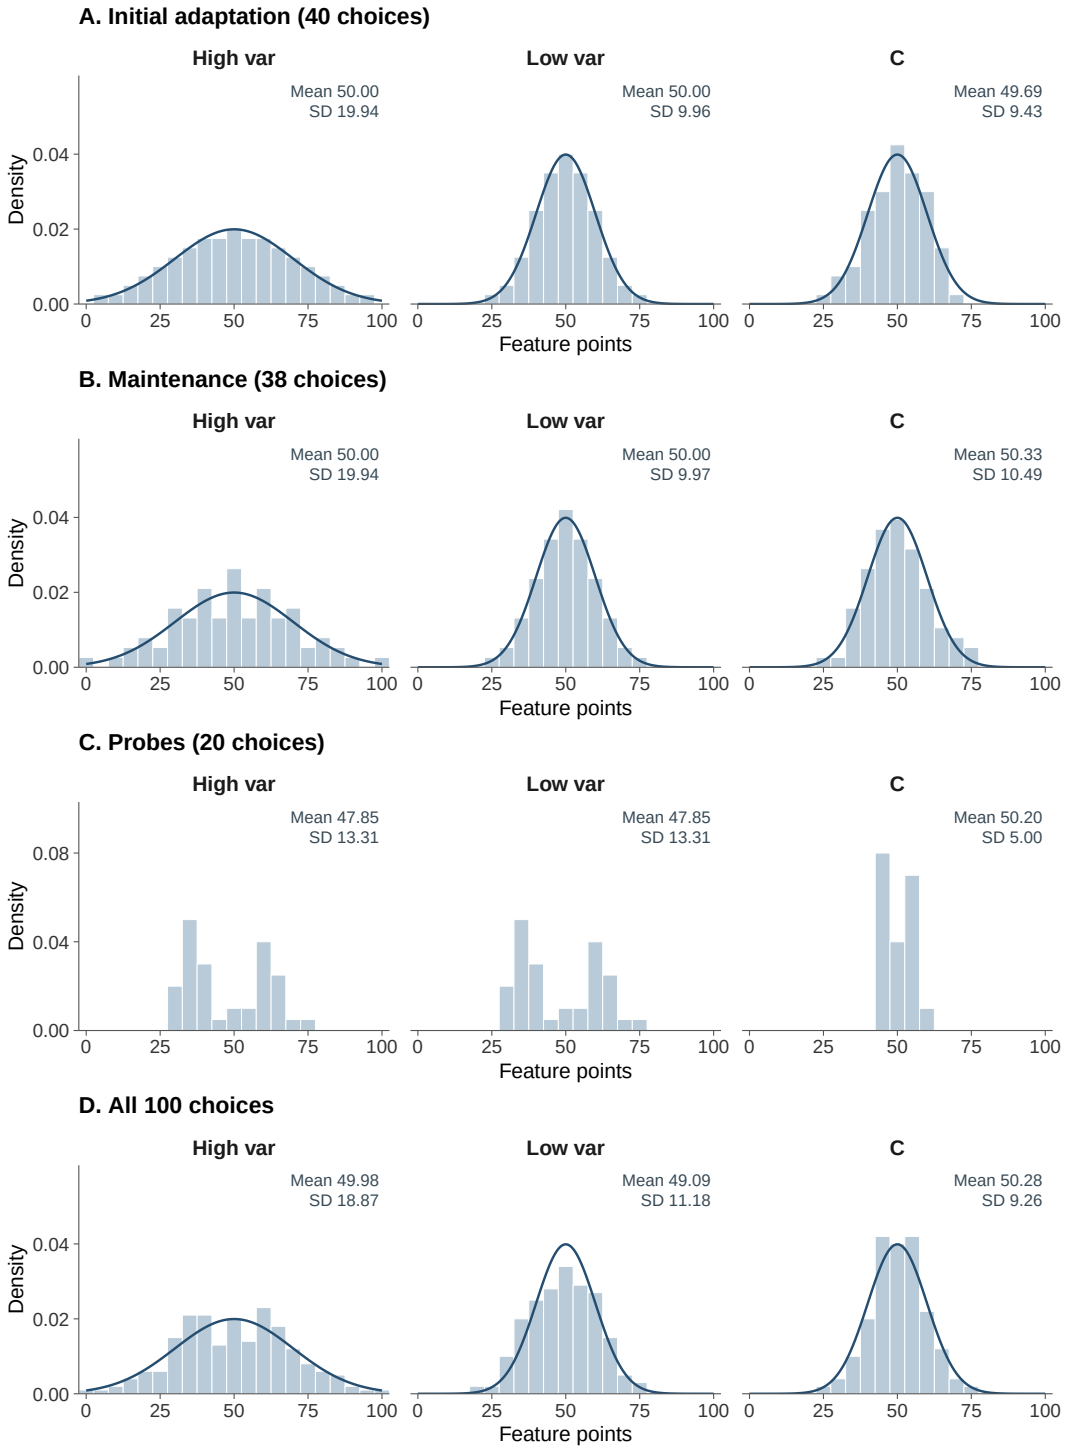
$$R(\lambda d; \nu) = \frac{\lambda\, \iota^\top \Sigma d}{\sqrt{\lambda^2 d^\top \Sigma d + \nu^2}} = \frac{\iota^\top \Sigma d}{\sqrt{d^\top \Sigma d + \nu^2/\lambda^2}},$$

which is (i). For (ii), let $\iota^\top \Sigma d > 0$. Then the numerator on the right-hand side is a positive constant, and the denominator is non-increasing in λ and strictly decreasing if $\nu > 0$. □

E Additional Figures and Tables

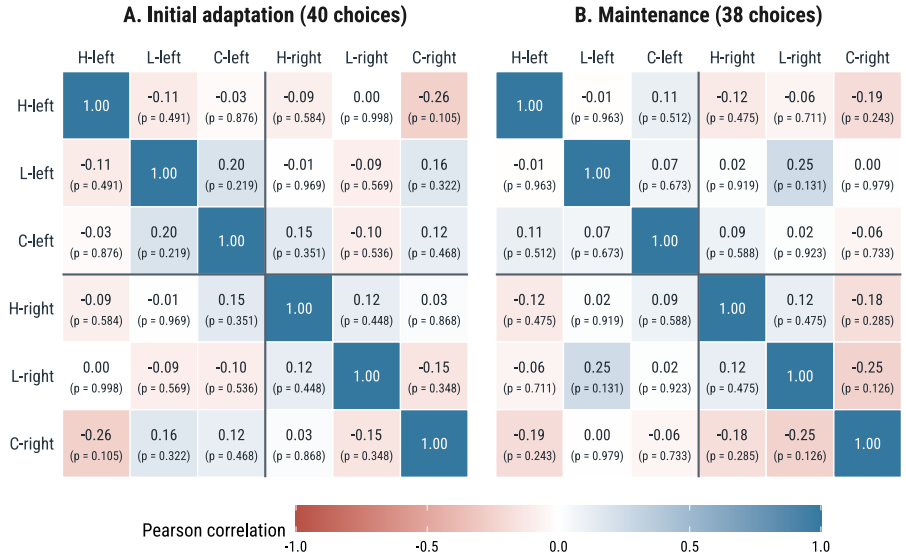
E.1 Additional Figures

Figure 7: Distribution of designed decisions



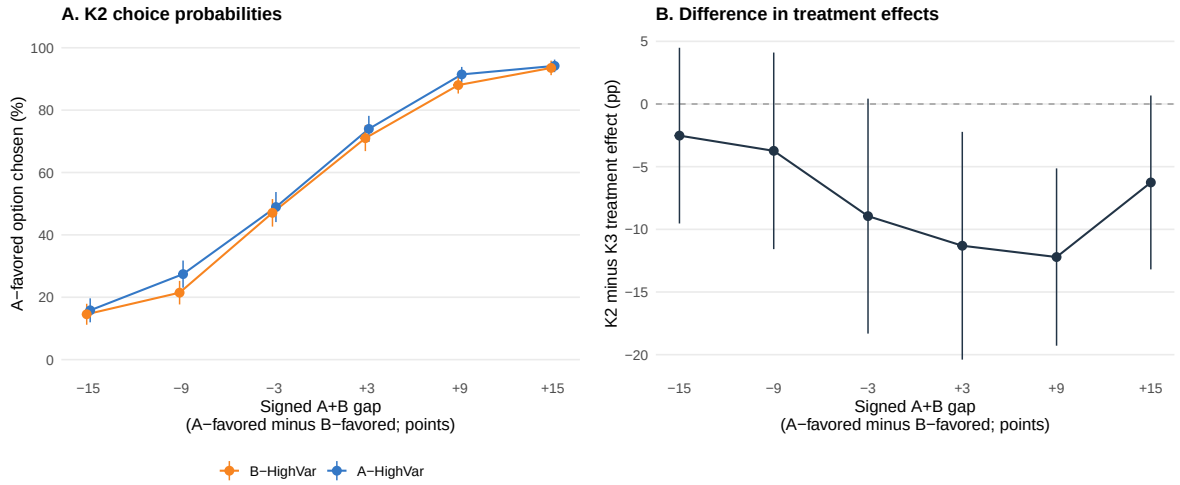
Notes: This figure summarizes the realized distributions of the displayed feature values by aggregation-complexity condition and environment treatment. *A-HighVar* assigns greater variation to feature A, while *B-HighVar* exchanges the roles of A and B. Feature C is displayed only in *K3*. Blue lines correspond to normal distributions.

Figure 8: Within- and between-product feature correlations



Notes: This figure shows a heatmap of Pearson correlations across 40 adaptation choices (Panel A) and 38 maintenance choices (Panel B), counting each fixed choice once. H/L denote A/B in *A-HighVar* and B/A in *B-HighVar*, respectively, while left/right identify the two products. Both treatments yield identical matrices after relabeling; deleting C gives the K2 matrices. Probes and sentinels are excluded. Parentheses report unadjusted, two-sided Pearson-test p-values with 38/36 degrees of freedom.

Figure 9: Adaptation by signed gap in the two- and three-variable conditions



Notes: Panel A reports the share of choosing the A-favored option, defined as the option with the higher value of attribute A, by treatment and signed A+B gap within the two-variable condition (320 participants; 6,388 common-probe choices). Panel B reports the two-variable minus *K3* difference in treatment effects, using 635 participants and 12,670 common-probe choices. Each treatment effect is the *A-HighVar* minus *B-HighVar* difference in percentage points; negative values indicate a larger effect in the *K3* baseline condition. The signed A+B gap is the A+B total of the A-favored option minus that of the B-favored option, excluding C. Positive gaps favor the A-favored option on A+B value. Probe widths are pooled within each participant and gap before averaging across participants with equal weights. Error bars indicate pointwise 95% confidence intervals, using participant-level t intervals in Panel A and Welch–Satterthwaite intervals for the four-arm contrast in Panel B. The retained sample applies the participant exclusions and removes individual timeout choices.

E.2 Additional Tables

Table 3: Demographic characteristics and treatment balance

Variable	ACS/SCF (2022)	All	K2			K3			Pooled by K		
			A-wide	B-wide	p	A-wide	B-wide	p	K2	K3	p
Gender											
Female	50%	49%	50%	45%	0.461	49%	54%	0.366	48%	51%	0.321
Age											
18–34	29%	37%	41%	38%	0.570	35%	35%	0.975	40%	35%	0.244
35–54	32%	45%	42%	47%	0.332	42%	51%	0.124	45%	46%	0.710
55+	38%	17%	17%	15%	0.572	23%	14%	0.054	16%	19%	0.314
Household income											
Below 50k	34%	27%	30%	22%	0.092	30%	26%	0.416	26%	28%	0.499
50k–100k	29%	38%	39%	36%	0.551	36%	42%	0.221	38%	39%	0.824
Above 100k	37%	35%	30%	42%	0.032	34%	32%	0.622	36%	33%	0.394
Education											
Bachelor's degree or more	33%	64%	59%	70%	0.044	61%	67%	0.342	64%	64%	0.907
Region											
Northeast	17%	18%	19%	18%	0.767	17%	19%	0.653	18%	18%	0.919
Midwest	21%	19%	23%	19%	0.373	21%	14%	0.082	21%	17%	0.191
South	39%	40%	39%	36%	0.586	39%	47%	0.130	37%	43%	0.126
West	24%	23%	19%	27%	0.074	23%	20%	0.529	23%	22%	0.642
Joint balance test (p)					0.306			0.163			0.676
Sample size		635	155	165		161	154		320	315	

Notes: This table presents summary statistics for participants' demographics in the *K2* and *K3* conditions and compares them with the ACS/SCF 2022 benchmarks. Columns 5 and 8 report p-values for equality between the *A-HighVar* and *B-HighVar* groups within *K2* and *K3*, respectively; column 11 compares *K2* and *K3* with treatment groups pooled. Tests use session fixed effects and participant-identity-clustered HC1 standard errors. The joint balance row reports p-values from a joint Wald test of nine nonredundant demographic indicators (χ^2, 9 degrees of freedom), allowing for cross-equation covariance. All p-values are nominal. Household income refers to 2025 income in US dollars before taxes and deductions. Education is a dummy variable indicating the participant holds a bachelor's degree or higher degree, and the regions follow the four US Census groups.

Table 4: Adaptation effects by probe characteristics in the Baseline (K3)

	Dependent variable:				
	A-favored choice (0 or 100)				
	Absolute gap	Signed gap	Probe location	Combined abs.	Combined signed
	(1)	(2)	(3)	(4)	(5)
A-HighVar	10.45***	10.44***	10.20***	10.21***	10.44***
	(1.68)	(1.68)	(2.11)	(2.12)	(2.06)
Absolute AB gap (per 10, centered)	1.82			1.82	
	(1.63)			(1.63)	
Signed AB gap (per 10)		20.51***			20.51***
		(0.93)			(0.93)
Width (per 10, centered)	0.41	0.42		0.42	0.42
	(1.17)	(1.17)		(1.17)	(1.17)
Last ten probes			1.05	1.05	1.21
			(1.62)	(1.62)	(1.52)
A-HighVar × absolute gap	−5.88**			−5.88**	
	(2.27)			(2.27)	
A-HighVar × signed gap		1.87			1.87
		(1.24)			(1.24)
A-HighVar × width	0.99	0.97		0.99	0.97
	(1.67)	(1.68)		(1.67)	(1.68)
A-HighVar × last ten			0.49	0.48	0.003
			(2.36)	(2.36)	(2.20)
Constant	41.11***	41.13***	40.58***	40.58***	40.52***
	(1.20)	(1.20)	(1.51)	(1.51)	(1.47)
Outcome mean	46.45	46.45	46.45	46.45	46.45
Participants	315	315	315	315	315
Observations	6,282	6,282	6,282	6,282	6,282

Notes: All columns use the *K3* participants and 6,282 human common-probe choices. The outcome is a dummy variable taking value 100 if the participant selects the *A-favored* option. The weighted linear probability regressions give each participant total weight one over retained probes (excluding timeout choices). Signed AB gap is the A-favored option's A+B total minus the B-favored option's A+B total, with the absolute AB gap being the unsigned equivalent. Width is the sum of the absolute A and B differences. All gap and width terms are scaled per ten points; absolute gap and width are centered at 9 and 49 points. We exclude C for both measures. Last ten probes' identifies the probes at location 11–20. All columns include the treatment main effect and all constituent main effects. Participant-clustered HC1 standard errors are reported in parentheses. $^{*}p < 0.05$, $^{**}p < 0.01$, $^{***}p < 0.001$.

Table 5: Robustness to color scheme, orientation, and probe order in the Baseline (K3)

	Dependent variable:			
	Choosing the A-favored option (%)			
	Baseline	Joint controls	Color only	Linear order
	(1)	(2)	(3)	(4)
A-HighVar	10.447	10.281	10.275	10.278
	(1.680)	(1.673)	(1.666)	(1.670)
Colors: green / teal / purple		4.105	4.229	4.116
		(3.030)	(3.012)	(3.026)
Colors: purple / green / teal		3.971	4.061	3.979
		(2.709)	(2.697)	(2.705)
Colors: purple / teal / green		7.852	7.868	7.857
		(2.986)	(2.973)	(2.982)
Colors: teal / green / purple		5.919	5.949	5.926
		(2.656)	(2.648)	(2.653)
Colors: teal / purple / green		3.459	3.579	3.465
		(3.028)	(3.025)	(3.024)
A-favored option on right				17.822
				(1.392)
Probe rank (centered)				0.081
				(0.095)
Constant	41.108	35.183	36.954	35.613
	(1.202)	(3.984)	(2.243)	(3.433)
A-HighVar 95% CI	[7.15, 13.74]	[7.00, 13.56]	[7.01, 13.54]	[7.00, 13.55]
A-HighVar p	<0.001	<0.001	<0.001	<0.001
Participants	315	315	315	315
Probe identity FE	No	Yes	No	Yes
Probe-slot FE	No	Yes	No	No
Observations	6,282	6,282	6,282	6,282

Notes: All columns use 315 retained participants and 6,282 common-probe choices. Weighted linear probability regressions model choosing the higher-A option, coded 0 or 100. A-HighVar equals one; B-HighVar is the reference. Each participant receives total weight one across non-timeout probes. Parentheses report participant-clustered HC1 standard errors; confidence intervals and p-values use a normal reference. Color rows list A/B/C; the reference is green/purple/teal. Probe and slot fixed-effect coefficients are suppressed. Slot effects absorb orientation and order because side is fixed by slot. Column (4) estimates side and linear rank separately under its restricted order specification. Rank is determined before timeout removal and centered at 10.5. The baseline constant is the B-HighVar mean. Adjusted constants use the reference colors, G01a, and slot 1 where included; column (4) uses A-favored-left at centered rank zero. Inference is exploratory.

Table 6: Robustness to color scheme, orientation, and probe order in K2

	Dependent variable:			
	Choosing the A-favored option (%)			
	Baseline	Joint controls	Color only	Linear order
	(1)	(2)	(3)	(4)
A-HighVar	2.878	2.880	2.885	2.879
	(1.195)	(1.194)	(1.192)	(1.193)
Colors: green / teal		1.373	1.373	1.373
		(2.173)	(2.167)	(2.170)
Colors: purple / green		2.341	2.382	2.340
		(2.342)	(2.335)	(2.339)
Colors: purple / teal		2.236	2.206	2.235
		(2.268)	(2.259)	(2.265)
Colors: teal / green		1.005	1.032	1.010
		(2.504)	(2.507)	(2.502)
Colors: teal / purple		1.279	1.320	1.278
		(2.404)	(2.399)	(2.401)
A-favored option on right				4.890
				(1.090)
Probe rank (centered)				−0.012
				(0.085)
Constant	55.830	77.454	54.403	75.250
	(0.778)	(3.506)	(1.929)	(2.961)
A-HighVar 95% CI	[0.54, 5.22]	[0.54, 5.22]	[0.55, 5.22]	[0.54, 5.22]
A-HighVar p	0.016	0.016	0.016	0.016
Participants	320	320	320	320
Probe identity FE	No	Yes	No	Yes
Probe-slot FE	No	Yes	No	No
Observations	6,388	6,388	6,388	6,388

Notes: All columns use 320 retained participants and 6,388 common-probe choices. Weighted linear probability regressions model choosing the higher-A option, coded 0 or 100. A-HighVar equals one; B-HighVar is the reference. Each participant receives total weight one across non-timeout probes. Parentheses report participant-clustered HC1 standard errors; confidence intervals and p-values use a normal reference. Color rows list A/B; the reference is green/purple. Probe and slot fixed-effect coefficients are suppressed. Slot effects absorb orientation and order because side is fixed by slot. Column (4) estimates side and linear rank separately under its restricted order specification. Rank is determined before timeout removal and centered at 10.5. The baseline constant is the B-HighVar mean. Adjusted constants use the reference colors, G01a, and slot 1 where included; column (4) uses A-favored-left at centered rank zero. Inference is exploratory.

Table 7: Reduced-form adaptation: sample restrictions and assignment controls

	(1) Prereg. sample	(2) All completers	(3) No timeouts	(4) First-attempt pass	(5) Joint slot controls	(6) Side + linear order	(7) Both sentinels correct
K3: effect (pp)	10.45	9.78	10.31	9.07	10.28	10.28	9.58
Standard error	(1.68)	(1.67)	(1.96)	(1.88)	(1.67)	(1.67)	(1.80)
95% CI	[7.15, 13.74]	[6.51, 13.04]	[6.47, 14.14]	[5.38, 12.76]	[7.00, 13.56]	[7.00, 13.55]	[6.06, 13.10]
Two-sided p	<0.001	<0.001	<0.001	<0.001	<0.001	<0.001	<0.001
K2: effect (pp)	2.88	2.90	2.86	3.08	2.88	2.88	2.30
Standard error	(1.20)	(1.19)	(1.32)	(1.40)	(1.19)	(1.19)	(1.22)
95% CI	[0.54, 5.22]	[0.58, 5.23]	[0.26, 5.45]	[0.34, 5.82]	[0.54, 5.22]	[0.54, 5.22]	[-0.09, 4.68]
Two-sided p	0.016	0.014	0.031	0.027	0.016	0.016	0.059
K3 minus K2: effect (pp)	7.57	6.87	7.45	5.99	7.40	7.40	7.28
Standard error	(2.06)	(2.04)	(2.36)	(2.34)	(2.06)	(2.05)	(2.17)
95% CI	[3.53, 11.61]	[2.87, 10.88]	[2.82, 12.08]	[1.39, 10.58]	[3.37, 11.43]	[3.38, 11.42]	[3.03, 11.54]
Two-sided p	<0.001	<0.001	0.002	0.011	<0.001	<0.001	<0.001
K3 participants	315	326	246	235	315	315	266
K3 probe decisions	6282	6484	4920	4686	6282	6282	5306
K2 participants	320	328	260	235	320	320	296
K2 probe decisions	6388	6531	5200	4692	6388	6388	5909
Color and probe FE	No	No	No	No	Yes	Yes	No
Probe-slot FE	No	No	No	No	Yes	No	No
Side and linear rank	No	No	No	No	No	Yes	No

Notes: This table reports robustness analyses for the reduced-form results, reporting the *A-HighVar* minus *B-HighVar* difference in the probability of choosing the A-favored option in percentage points. The third block reports the *K3* effect minus the *K2* effect; positive values indicate stronger adaptation in K3.

Column (1) uses the preregistered analysis sample: full-outro completers with fewer than five timeouts across 100 decisions and mean non-timeout response time of at least 500 ms. Column (2) includes all 654 completers without these participant-level exclusions. Columns (3) and (4) separately restrict column (1) to participants with no timeouts and those answering all four comprehension questions correctly on their first attempt. Columns (5) and (6) use the same sample as column (1). Column (7) restricts column (1) to participants answering both sentinels correctly, excluding 73 participants and retaining 562. There are no missing or timed-out sentinel responses in the preregistered sample.

All columns exclude computer-generated timeout choices. Weighted linear probability regressions give each participant total weight one across their non-timeout probe decisions. Models are estimated separately within K, allowing all control coefficients to differ. Columns (1)–(4) and (7) include only treatment and a constant. Column (5) adds color-scheme, probe-identity, and probe-rank fixed effects; rank effects absorb orientation because side is fixed by rank. Column (6) includes color and probe fixed effects, an indicator for the A-favored option appearing on the right, and linear probe rank. Ranks are assigned before timeout removal. No demographic or session controls are included. Confidence intervals and two-sided p-values use a standard normal reference. For K3 minus K2, the variances of the independent within-K estimates are added. Parentheses report participant-clustered HC1 standard errors.

Table 8: Robustness to equal and unequal values of feature C

	(1) All probes	(2) Equal-C probes	(3) Unequal-C probes
Adaptation effect in $K = 3$ (pp)	10.45	13.55	9.66
	(1.68)	(2.56)	(1.77)
Adaptation effect in $K = 2$ (pp)	2.88	4.11	2.58
	(1.20)	(2.02)	(1.32)
$K = 3$ **minus** $K = 2$ **(pp)**	**7.57**	**9.44**	**7.08**
	(2.06)	(3.26)	(2.20)
95% confidence interval	[3.53, 11.61]	[3.05, 15.83]	[2.76, 11.40]
Two-sided p-value	< 0.001	0.004	0.001
Distinct probes	20	4	16
Participants	635	635	635
Valid probe decisions	12,670	2,536	10,134

Notes: Effects are A-HighVar minus B-HighVar in the probability of choosing the higher-A product, in percentage points. The difference-in-differences subtracts the $K = 2$ effect from the $K = 3$ effect. All columns retain the same 635 participants (315 in $K = 3$, 320 in $K = 2$) and the main analysis's weighting, timeout handling, and inference. Weighted linear probability models give each participant total weight one over valid probes within each column; timeout choices are excluded. Parentheses contain participant-clustered HC1 standard errors with separate finite-sample corrections within K. Independent within-K variances are added for the difference-in-differences; confidence intervals and two-sided p-values use a standard normal reference.

Probe subsets use the original $K = 3$ definitions, retaining identical probe IDs in $K = 2$ and $K = 3$. Equal-C probes are G05a/b and G10a/b: two mirror pairs, all with width 49 and absolute A+B gap 9. Width is $|\Delta A|+|\Delta B|$; the A+B gap excludes C. The other probes have $|\Delta C|= 1$. The equal-minus-unequal difference-in-differences is 2.36 pp (95% CI $[-4.13, 8.85]$; $p = .477$), accounting for within-participant covariance across subsets. Different probe characteristics prevent interpreting this comparison as the causal effect of unequal C alone.

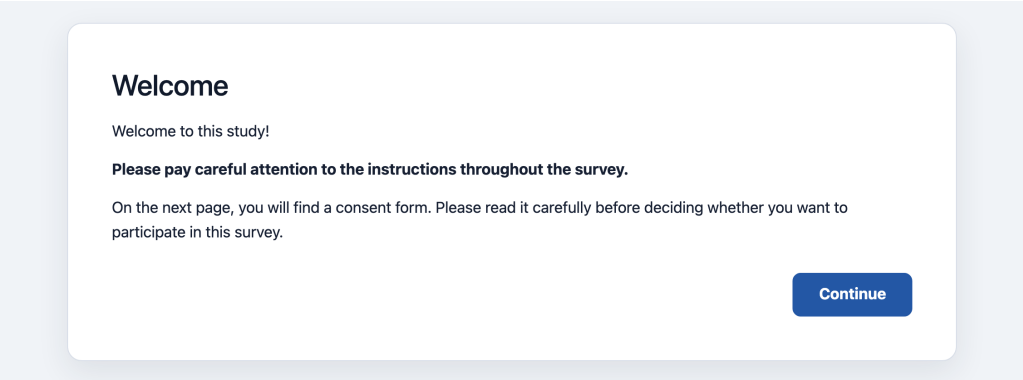
F Screenshots of the Experiment

This appendix shows the screens participants saw, in order. Bracketed labels indicate whether a screen appeared on a new page or was added to the page shown before it. Participants were randomly assigned to one of four arms, crossing the number of product features (two or three) with the feature whose ratings vary more across choices (Feature A or Feature B). All instruction screens are identical in the A-wide and B-wide arms. Screens are shown for the three-feature arm; screens whose content differs in the two-feature arm are shown separately. Bar colors were randomized across participants. The example decisions show the same stimulus in the three- and two-feature arms.

Before the task

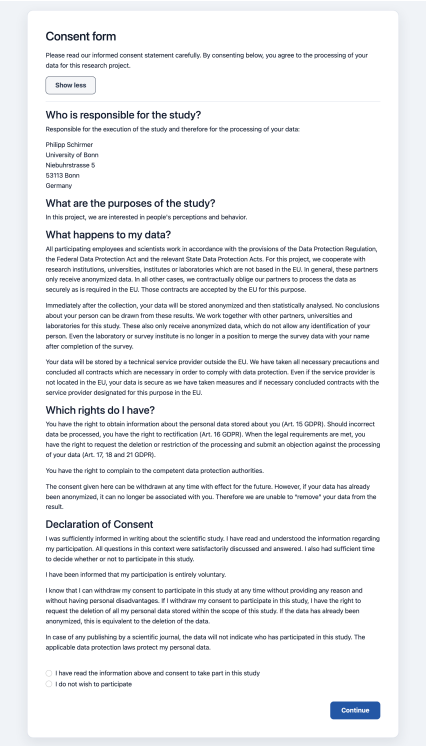
Welcome

[First page]



Consent

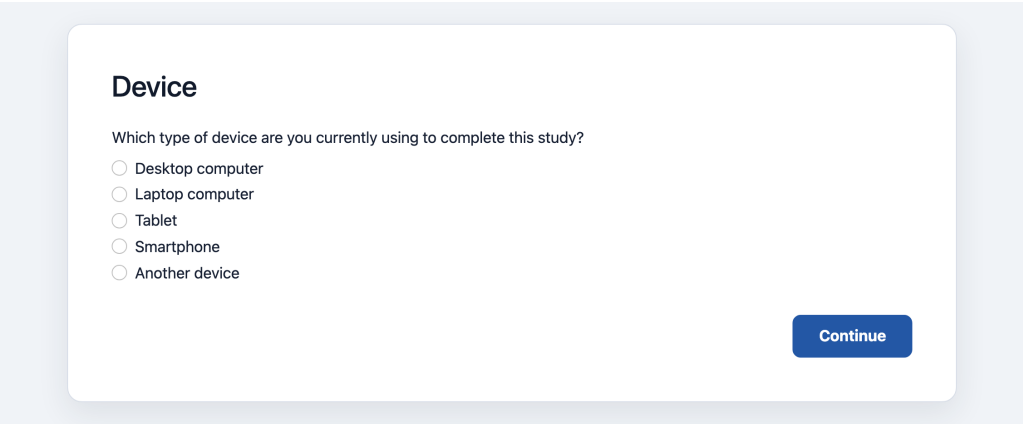
[Next page]



Shown expanded; by default the full consent text is collapsed behind a “Show more” button.

Device question

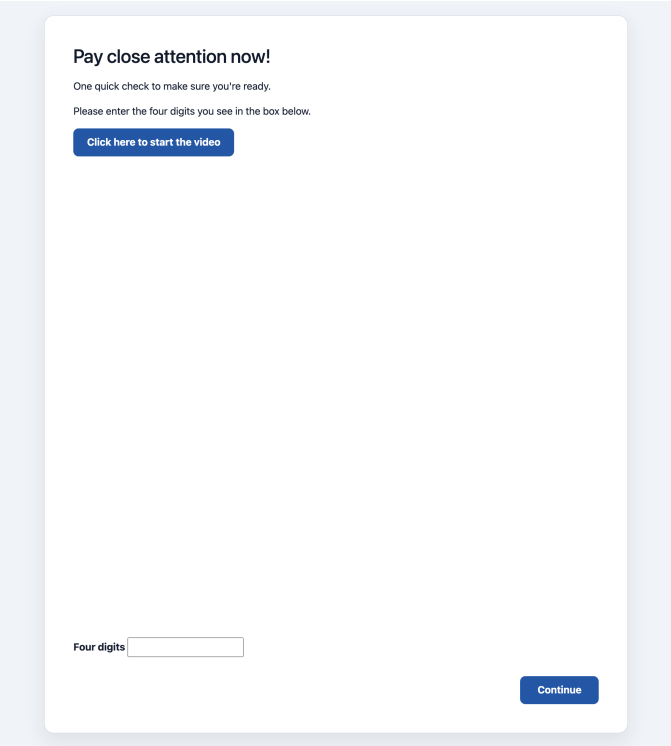
[Next page]



Only “Desktop computer” or “Laptop computer” continue; other answers lead to a page asking the participant to return the study.

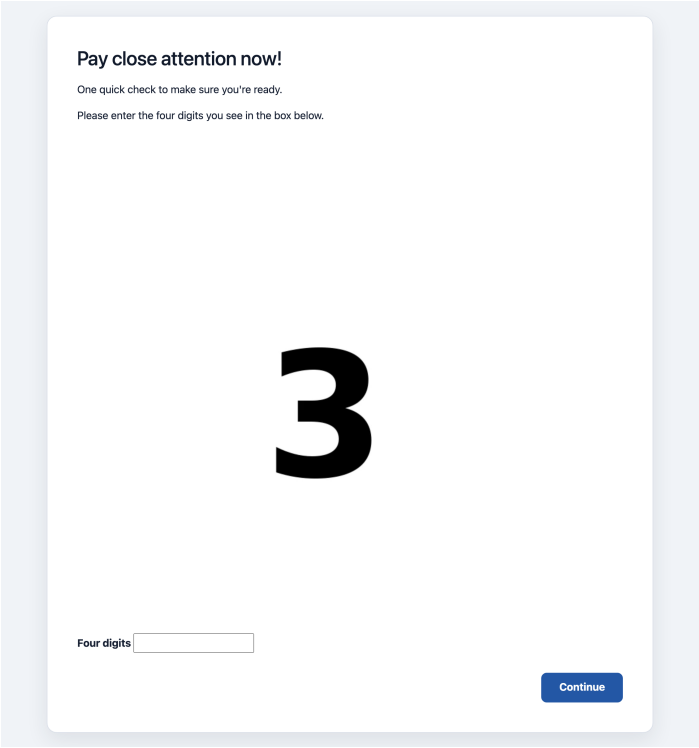
Attention check

[Next page]



Attention check: video playing

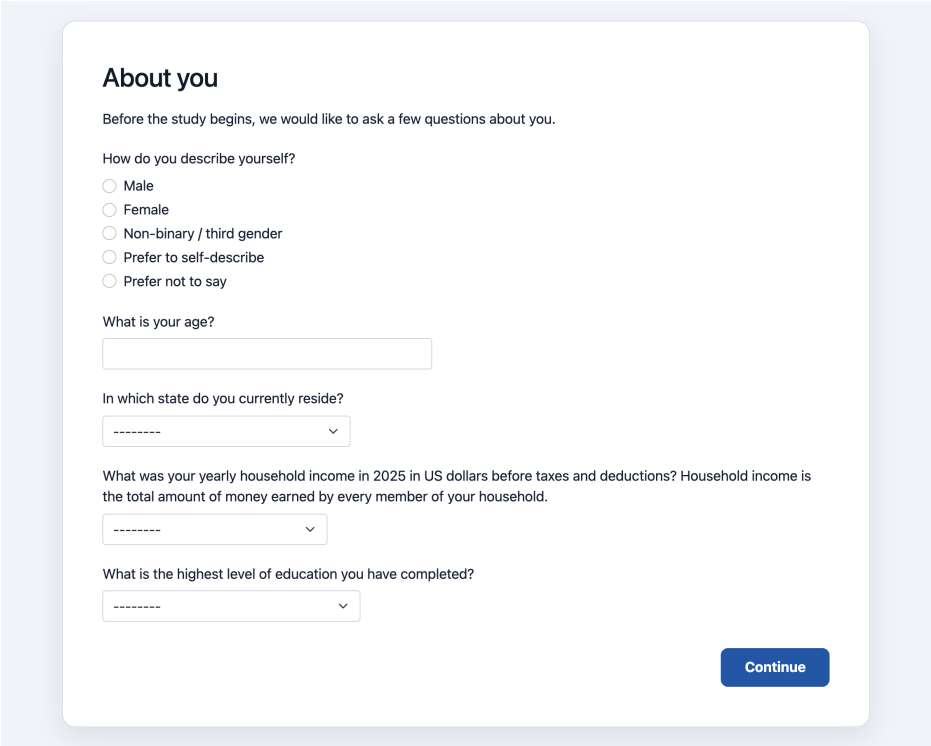
[Same page: after clicking the start button]



The video can be played once and shows four digits; only the correct four digits continue.

Demographics

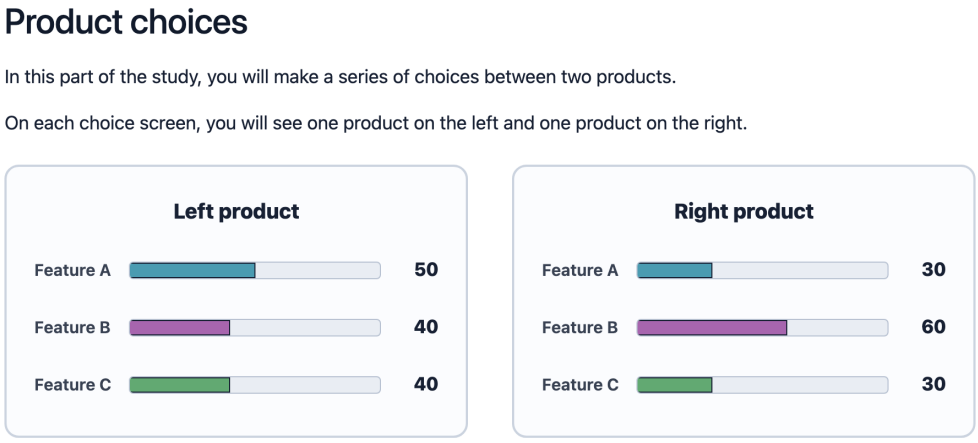
[Next page]



Instructions (three-feature arms)

Instructions, part 1: Product display

[Next page]



Instructions, part 2: Comparing total quality

[Same page: added below part 1 after clicking "Continue"]

Comparing total quality

A product's **total quality** is the sum of all its features' ratings:

Total quality = Feature A points + Feature B points + Feature C points

Your task is to **choose the product with the higher total quality.**

Consider the same two products from the example:



The left product has:

50 + 40 + 40 = 130 total quality points

The right product has:

30 + 60 + 30 = 120 total quality points

Because **130 is greater than 120**, you should choose the **left product** in this example.

It does not matter how the points are distributed across features—only the sum of all feature ratings determines total quality.

Instructions, part 3: Bonus payment

[Same page: added below part 2 after clicking "Continue"]

Bonus payment

In addition to the fixed payment, you may receive a bonus based on your choices.

If you complete the study, you have a **1 in 10 chance** of being selected for a bonus. If you are selected, the computer will randomly choose one of your main-task choices. Every main-task choice is equally likely to be selected.

The bonus depends on the total quality of the chosen product for that choice:

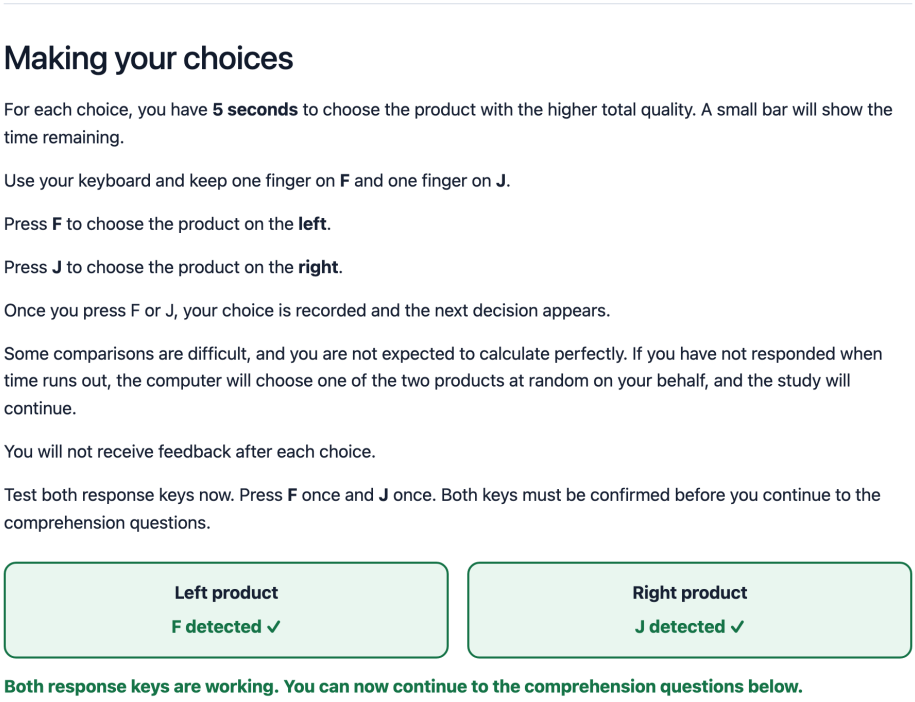
Bonus = Total quality points of the chosen product × $0.10

Your best strategy for earning the highest possible bonus is therefore to **choose the product with the higher total quality in every choice.**

Any bonus will be paid separately through Prolific.

Instructions, part 4: Making your choices

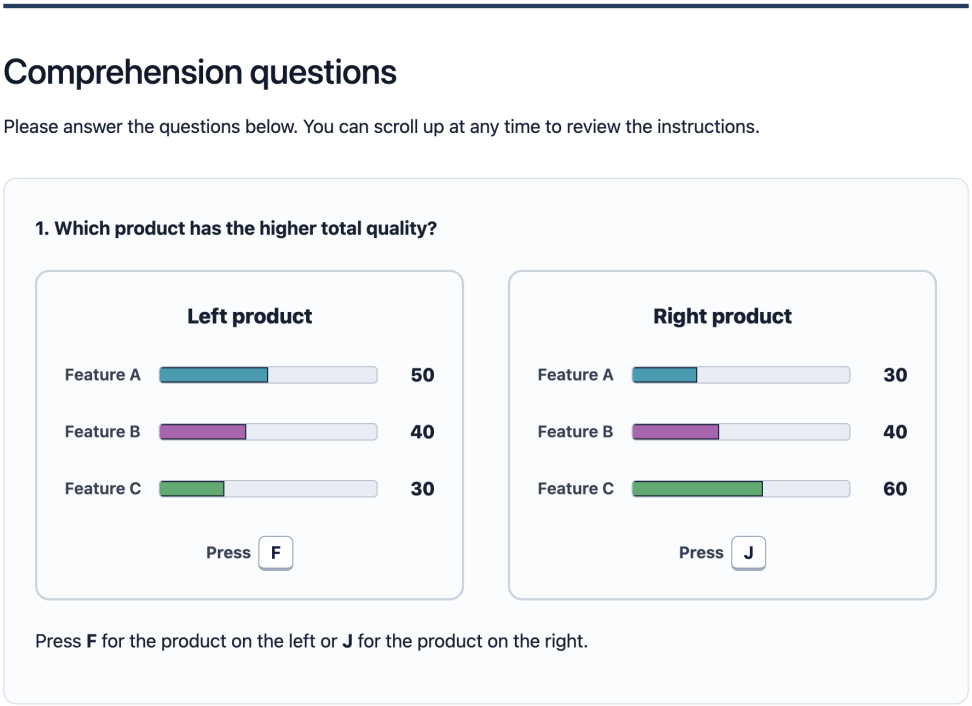
[Same page: added below part 3 after clicking "Continue"]



The response keys must register both F and J before "Next" is enabled.

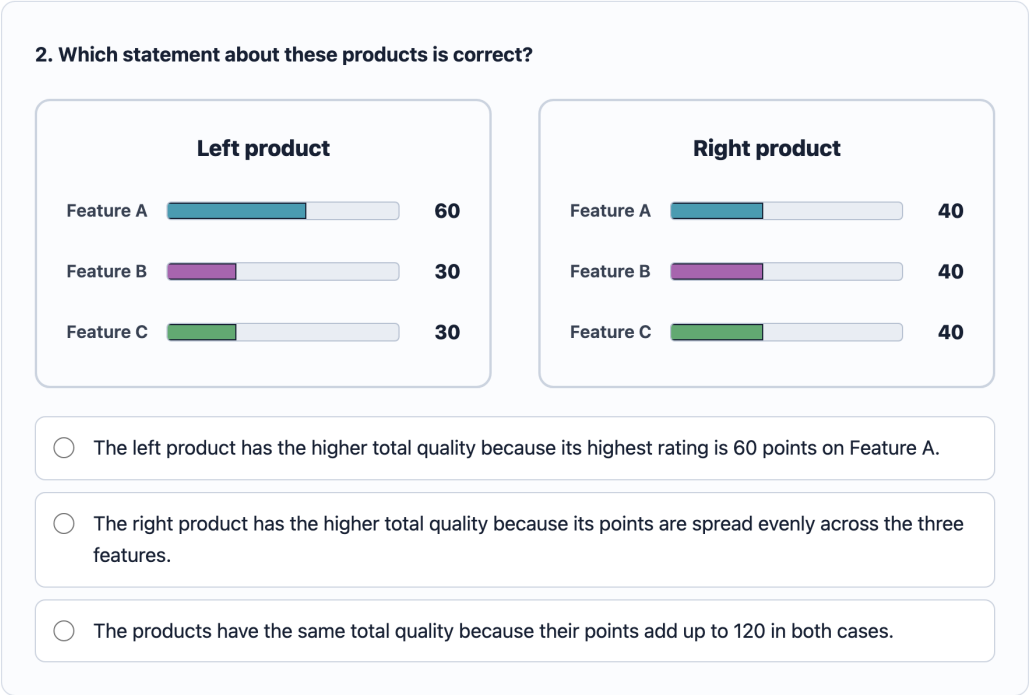
Comprehension questions: introduction and question 1

[Same page: shown below the complete instructions after clicking "Next"]



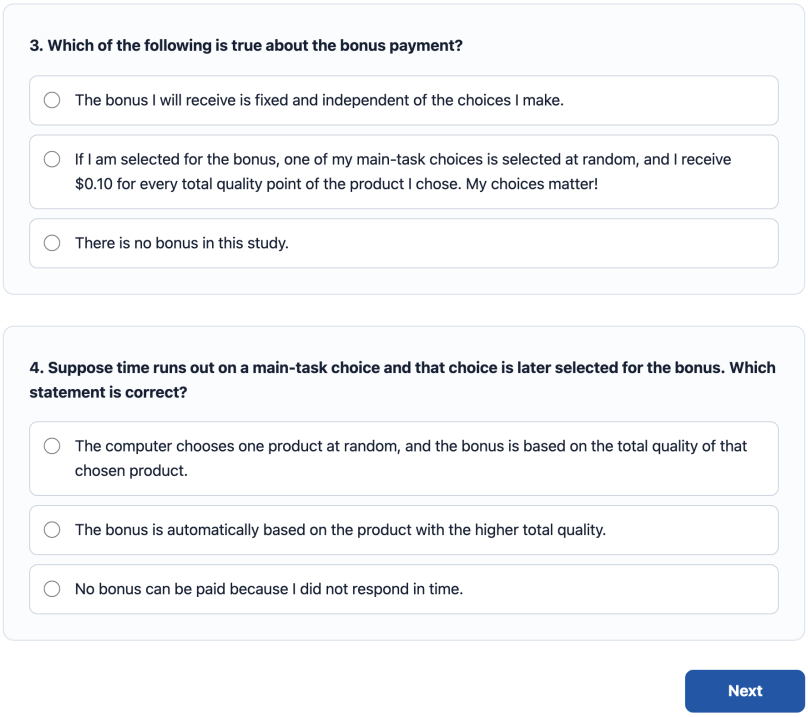
Comprehension questions: question 2

[Same page: continued below]



Comprehension questions: questions 3 and 4

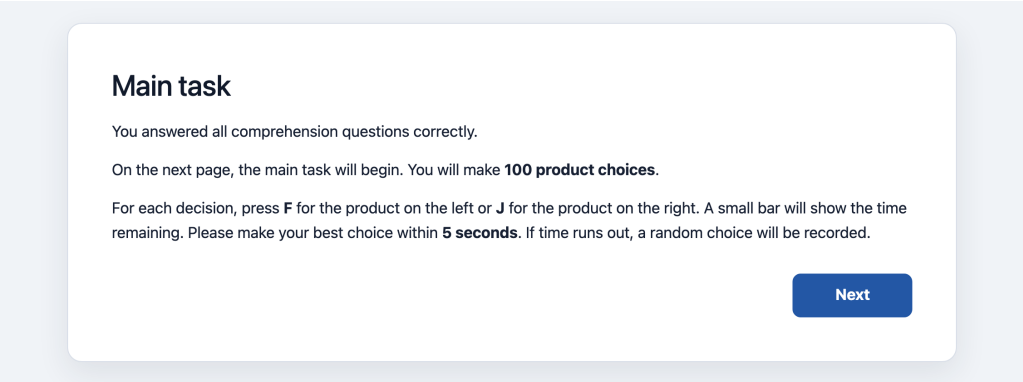
[Same page: continued below]



Participants answer all four questions and then click "Next".

Start of the main task

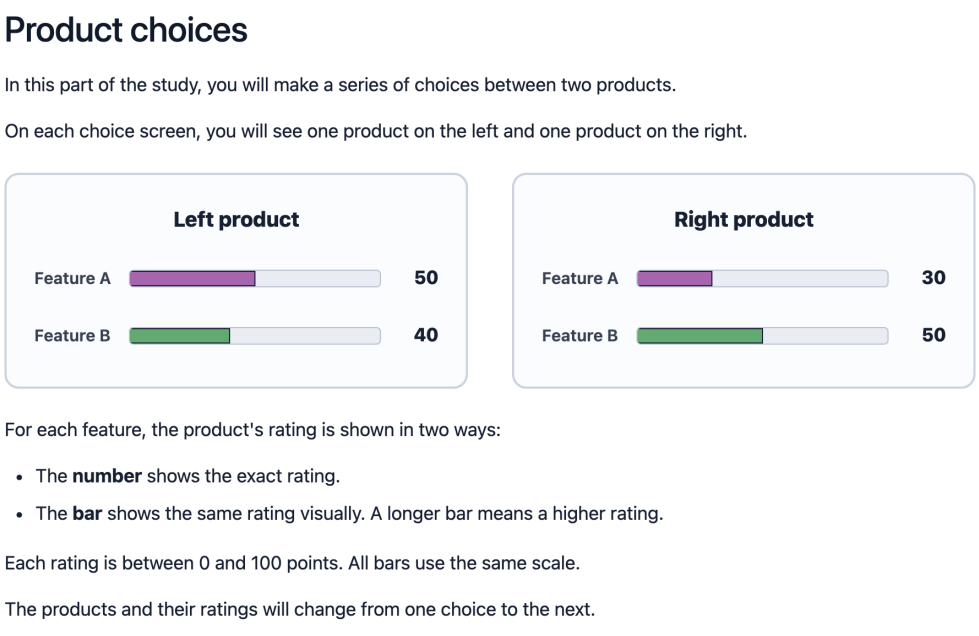
[Next page]



Instructions in the two-feature arms (screens whose content differs)

Instructions, part 1: Product display

[Next page]



Instructions, part 2: Comparing total quality

[Same page: added below part 1 after clicking "Continue"]

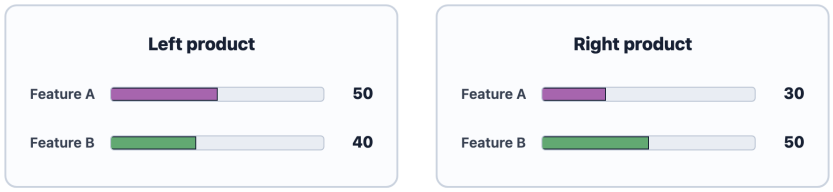
Comparing total quality

A product's **total quality** is the sum of all its features' ratings:

Total quality = Feature A points + Feature B points

Your task is to **choose the product with the higher total quality.**

Consider the same two products from the example:



The left product has:

50 + 40 = 90 total quality points

The right product has:

30 + 50 = 80 total quality points

Because **90 is greater than 80**, you should choose the **left product** in this example.

It does not matter how the points are distributed across features—only the sum of all feature ratings determines total quality.

Comprehension questions: introduction and question 1

[Same page: shown below the complete instructions after clicking "Next"]

Comprehension questions

Please answer the questions below. You can scroll up at any time to review the instructions.

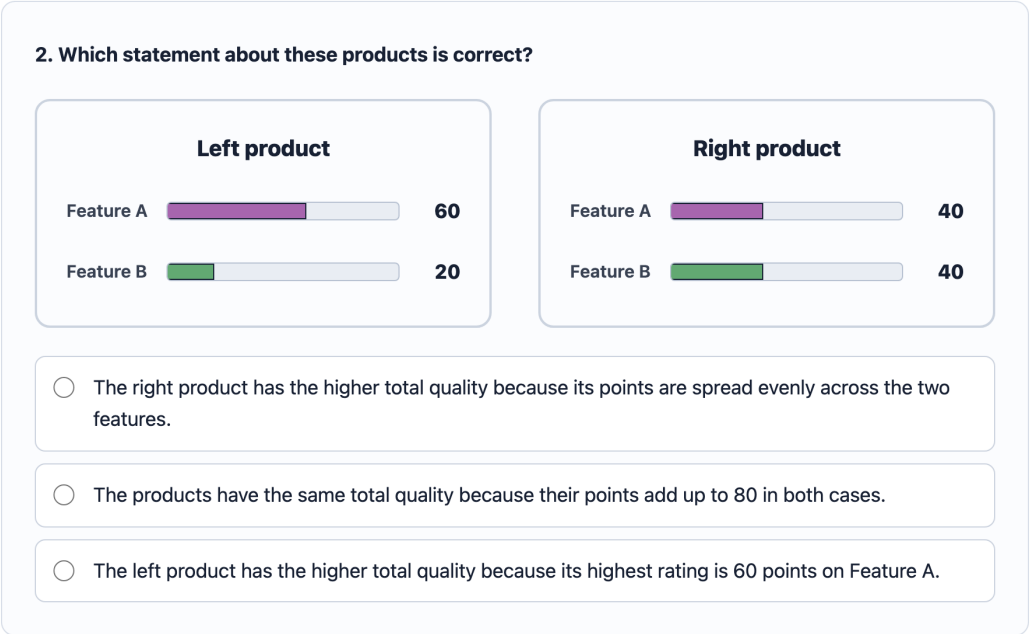
1. Which product has the higher total quality?



Press **F** for the product on the left or **J** for the product on the right.

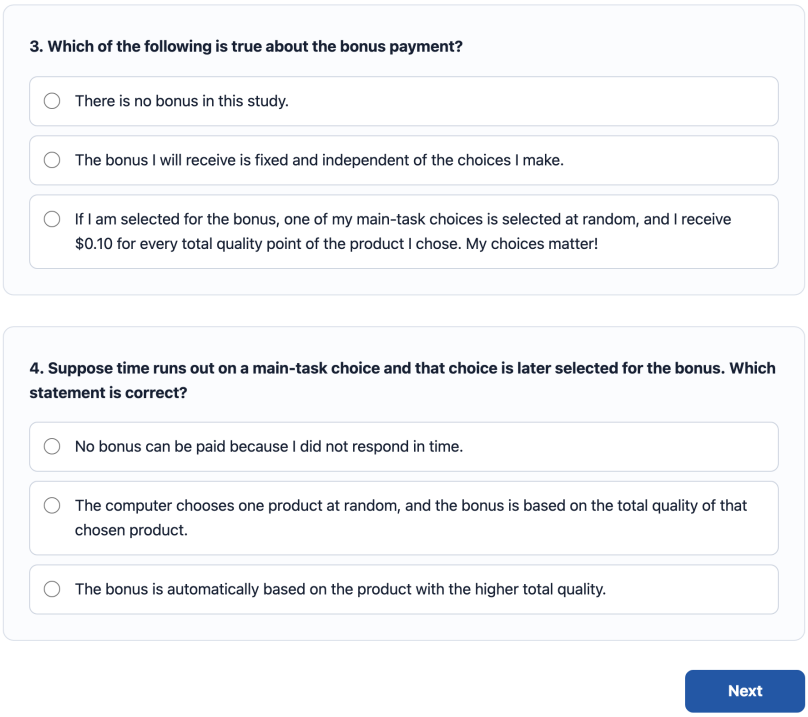
Comprehension questions: question 2

[Same page: continued below]



Comprehension questions: questions 3 and 4

[Same page: continued below]

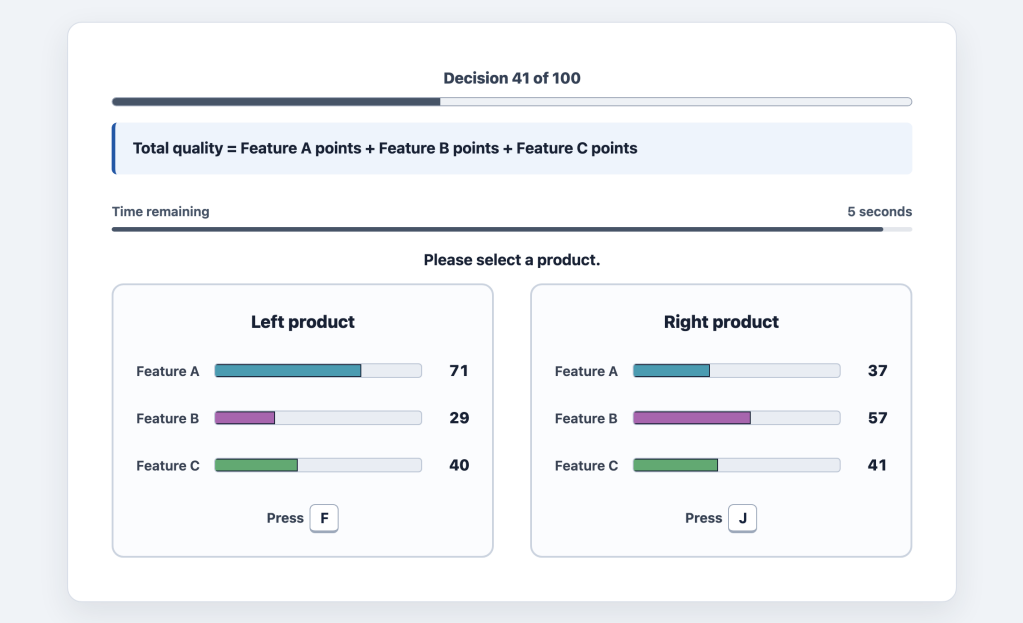


Participants answer all four questions and then click "Next".

Main task: example decisions

Adaptation choice, 3 features (decision 41 of 100)

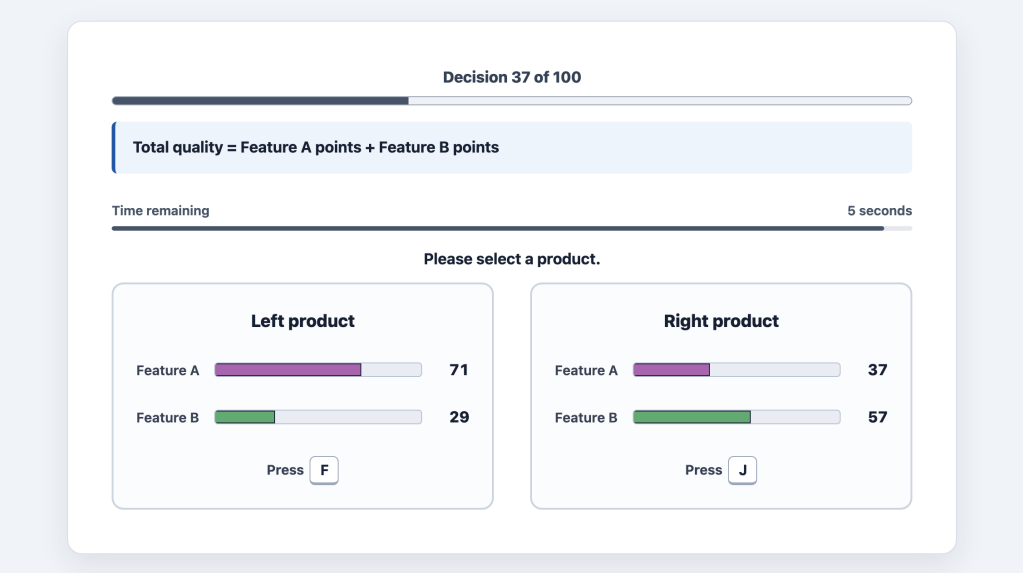
[One screen per decision]



Left 71/29/40 (total 140) vs. right 37/57/41 (total 135).

Adaptation choice, 2 features (decision 37 of 100)

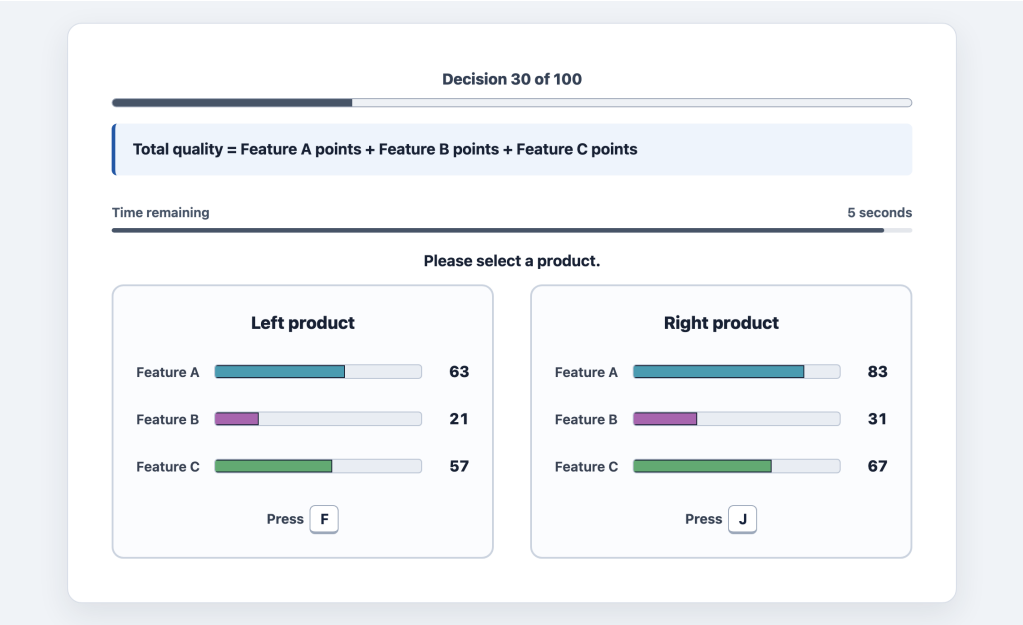
[One screen per decision]



Left 71/29 (total 100) vs. right 37/57 (total 94).

Dominance check, 3 features (decision 30 of 100)

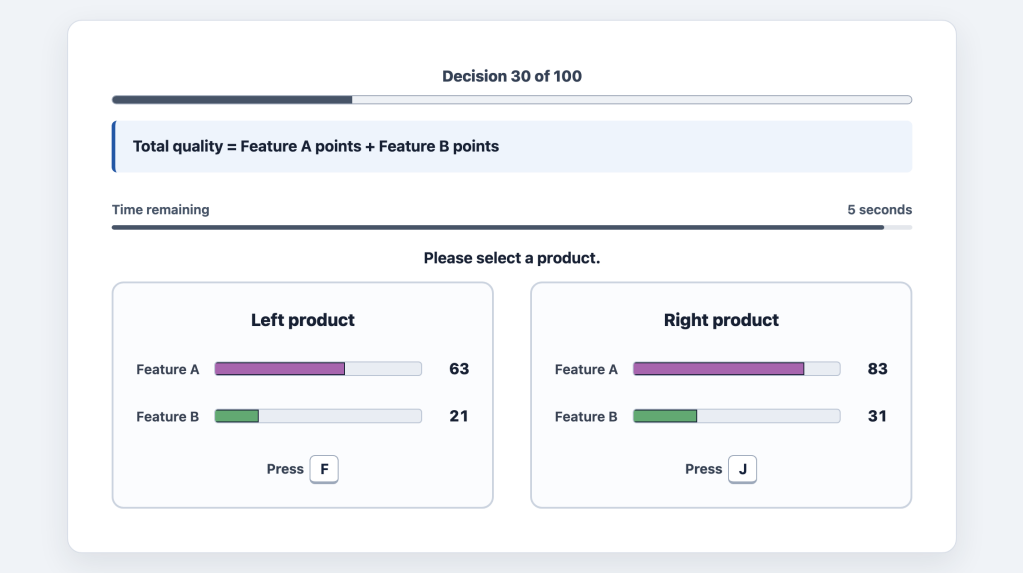
[One screen per decision]



Left 63/21/57 (total 141) vs. right 83/31/67 (total 181).

Dominance check, 2 features (decision 30 of 100)

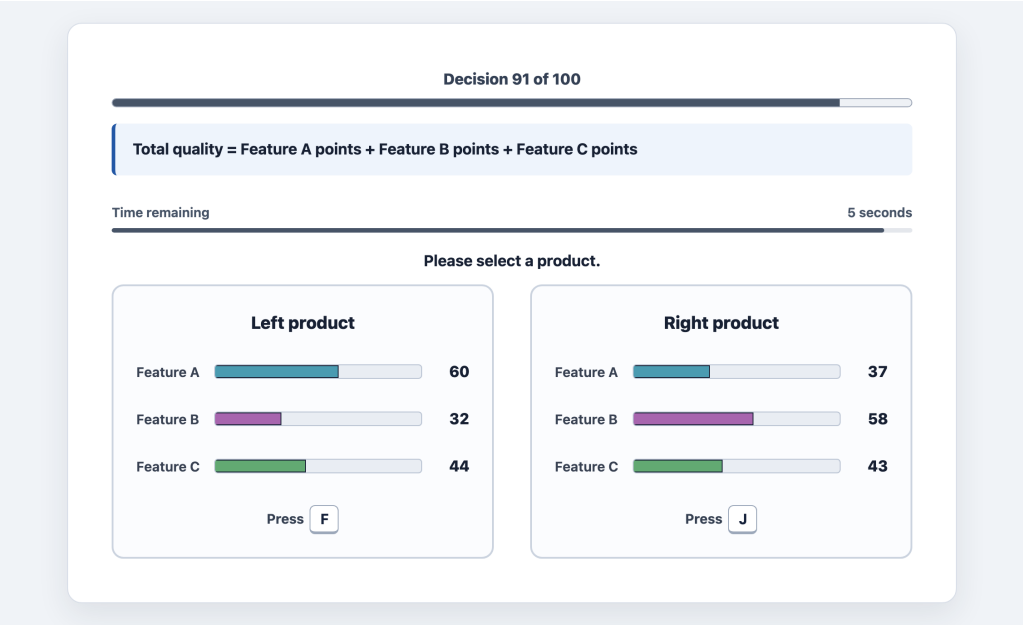
[One screen per decision]



Left 63/21 (total 84) vs. right 83/31 (total 114).

Common probe, 3 features (decision 91 of 100)

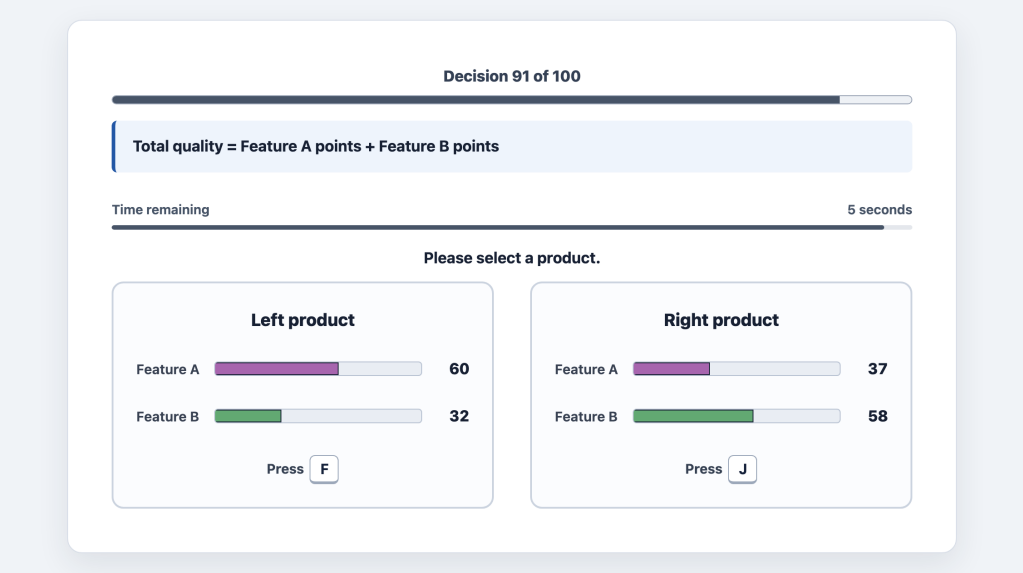
[One screen per decision]



Left 60/32/44 (total 136) vs. right 37/58/43 (total 138).

Common probe, 2 features (decision 91 of 100)

[One screen per decision]

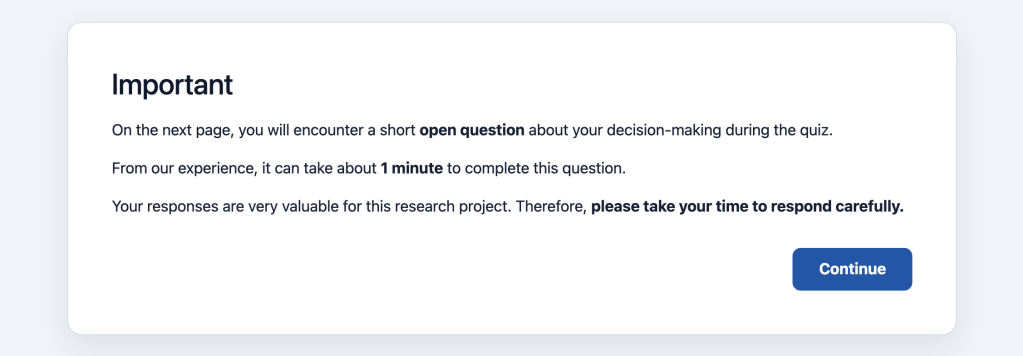


Left 60/32 (total 92) vs. right 37/58 (total 95).

After the task

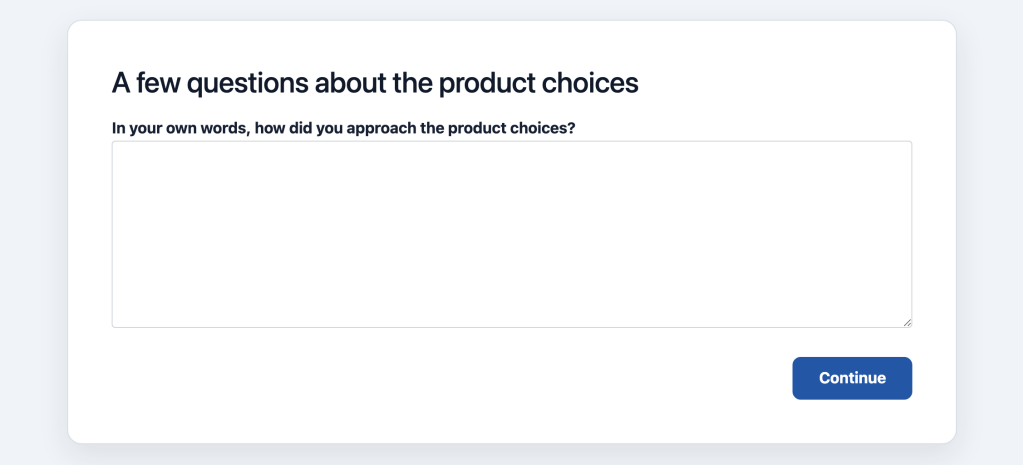
Transition to the final questions

[Next page]



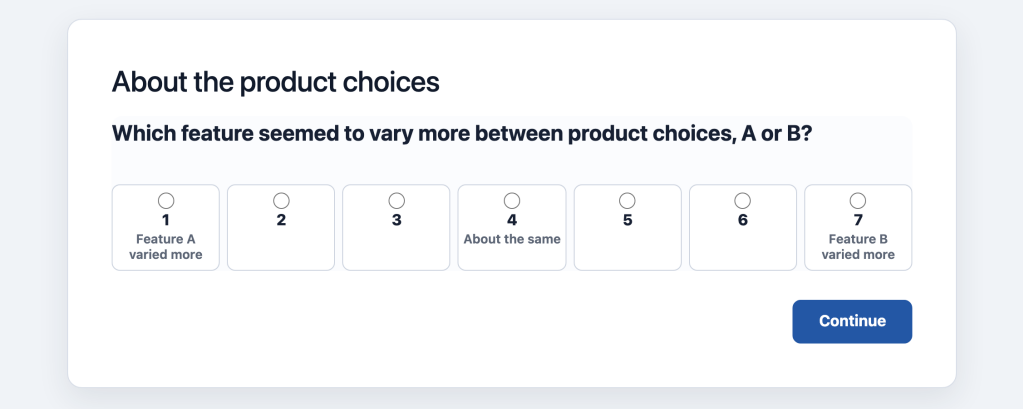
Strategy (open question)

[Next page]



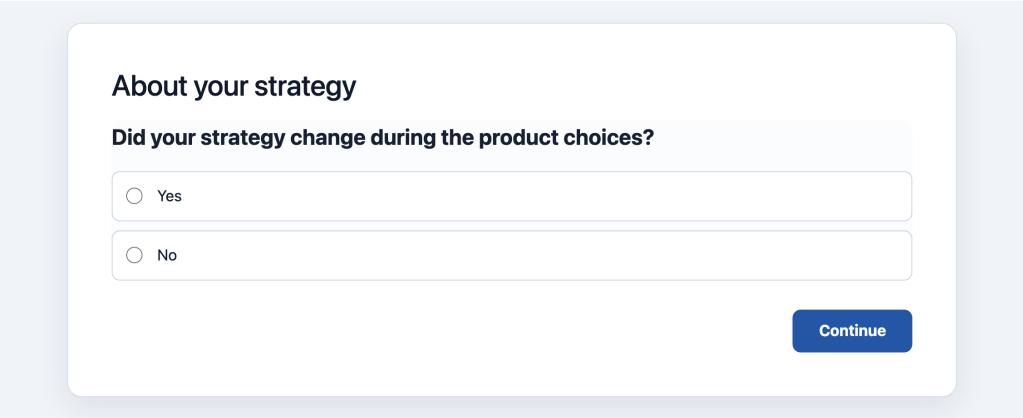
Perceived feature variability

[Next page]



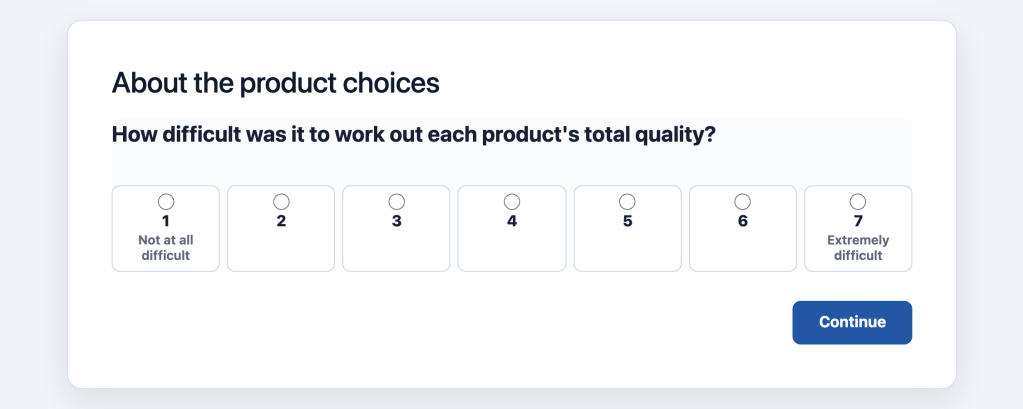
Strategy change

[Next page]



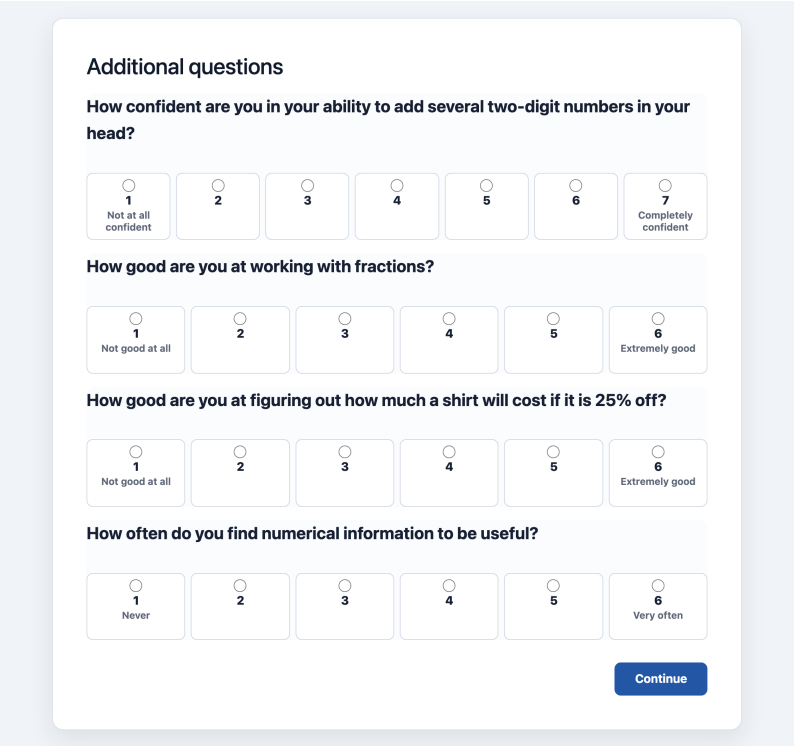
Difficulty

[Next page]



Math confidence and subjective numeracy

[Next page]



End of study

[Next page]

