

Choice Architecture for Healthier Insurance Decisions: Ordering and Partitioning Together Can Improve Consumer Choice

Benedict G.C. Dellaert , Eric J. Johnson , Shannon Duncan, and Tom Baker

Journal of Marketing
 2024, Vol. 88(1) 15-30
 © The Author(s) 2022



Article reuse guidelines:
sagepub.com/journals-permissions
 DOI: 10.1177/00222429221119086
journals.sagepub.com/home/jmx



Abstract

Making good health insurance decisions is important for health outcomes and longevity, but consumers' errors are well documented. The authors examine whether targeted choice architecture interventions can reduce these mistakes. The article examines the interaction of two choice architecture tools on improved consumer insurance decisions in online health care exchanges: (1) ordering the options from best to worst based on a high-quality user model and (2) partitioning the total set of options. Although ordering and partitioning do not always improve choices separately, the authors use one field study and three experiments to identify the conditions that allow the combination to greatly improve health insurance decisions. Findings indicate that when options are ordered such that the best options appear at the beginning of the presented list, partitioning nudges consumers to focus on the best options. However, if the best options are not at the top of the list, partitioning discourages search and can impair consumers' discovery of the best options. Process data show that these effects are achieved by focusing consumers' limited attention on higher-quality options. These results suggest that wise choice architecture interventions need to consider the joint effect of choice architecture tools as well as the quality of the firm's user model.

Keywords

choice architecture, ordering, partitioning, health insurance choice, health care exchanges, user models, consumer decision making, digital interfaces

Online supplement: <https://doi.org/10.1177/00222429221119086>

Choice architecture affects consumers' choice behavior, and its impact is amplified by digital technology. We investigate how choice architecture affects consumers' health insurance choices. These decisions are important for consumers. They have strong financial consequences and determine access to potentially lifesaving health care. Health insurance coverage increases longevity (Woolhandler and Himmelstein 2017) and halves the probability of bankruptcy (Gotberg and Sousa 2019). Health insurance decisions are also important for society because they affect firms' labor costs and raise social justice concerns.

Health insurance choices are also complex: They involve many attributes like deductibles and copayments that people often do not understand (Loewenstein et al. 2013). Because of the economics of insurance, these attributes often require complex trade-offs involving negatively correlated attributes (Abaluck and Gruber 2011; Ericson and Starc 2012). For example, large provider networks are usually more expensive, and low deductibles generate higher premiums. As a result, it

is not surprising that consumers make mistakes that result in substantial overpayment for coverage (Johnson et al. 2013). Many consumers select options that are dominated, unnecessarily spending more for equivalent benefits. In a recent study, employees of a large U.S. firm paid 42% more than needed for equivalent coverage, essentially wasting a large part of their payments (Bhargava, Loewenstein, and Sydnor 2017). Improving these choices can increase health outcomes for consumers and broadly increase the efficiency of health care resources. In this article, we show that seemingly minor

Benedict G.C. Dellaert is Professor of Marketing, Erasmus School of Economics, Erasmus University Rotterdam, The Netherlands (email: dellaert@ese.eur.nl). Eric J. Johnson is Norman Eig Professor of Business, Columbia Business School, Columbia University, USA (email: ejj3@columbia.edu). Shannon Duncan is a doctoral student in marketing, The Wharton School, University of Pennsylvania, USA (email: smduncan@wharton.upenn.edu). Tom Baker is William Maul Measey Professor of Law, Carey Law School, University of Pennsylvania, USA (email: tombaker@law.upenn.edu).

changes in choice architecture can have a large effect on health insurance product choice and consumer health care costs.

Ideally, choice architecture improves consumers' choices (Sunstein 2018; Thaler and Sunstein 2008). For example, defaults can facilitate the choice of a good alternative while reducing the need for extensive contemplation (Brown and Krishna 2004; Donkers et al. 2020; Jachimowicz et al. 2019; Johnson and Goldstein 2003; Smith, Goldstein, and Johnson 2013). In parallel, recent rapid advances in large-scale data availability, artificial intelligence, and machine-learning-based algorithms suggest a new reality in which marketers can predict with greater accuracy which products best match a given consumer's needs (Chintagunta, Hanssens, and Hauser 2016; Ghose, Ipeirotis, and Li 2012). Combining these insights with choice architecture can help consumers choose the product that is best for them. It can also provide new business opportunities for firms while increasing consumer satisfaction with the chosen product. Domains such as health and consumer finance are especially promising for such interventions because they involve high-impact, infrequent consumer decisions. The complexity of these decisions combined with the limited opportunity to learn from experience makes it difficult for consumers to select the best available option on their own.

Health care exchanges offer promising opportunities for digital choice architectures to support consumer health insurance decisions. Like many websites, they use two ubiquitous choice architecture tools. First, health care exchanges can *order* the available options. All websites make this decision. Even an alphabetical order is a choice architecture design decision. Second, exchanges can *partition* the total choice set presented to consumers by determining whether a small number of initial options are presented on the first web page. Although ordering and partitioning do not always improve choices separately, we identify the conditions that allow the combination to greatly improve health insurance decisions. This is because the effects of ordering and partitioning are not independent: When options are ordered such that the best options appear at the beginning of the presented list, partitioning slightly and subtly nudges consumers to focus on the best options. However, if the best options are not at the top of the list, partitioning discourages search and can impair consumers' discovery of the best options.

These findings can be understood in terms of fundamental decision principles. To choose a good option, consumers must pay attention to that option and avoid paying attention to poor-quality options. High-quality ordering ensures that good options are seen by consumers. However, without partitioning, consumers sometimes search too much and revise which options they consider, also including lower-quality options. Therefore, when the best options are not presented at the beginning of a choice set (e.g., with random or low-quality ordering), partitioning can be harmful by focusing consumers' attention on options that are not truly superior.

We make three contributions in this research. First, we provide a substantive demonstration of the impact of predictive choice architecture in health care. We demonstrate that the novel combination of the existing choice architecture interventions of ordering

and partitioning can improve consumer decision outcomes in the societally relevant domain of health insurance. Second, we show that the quality of a firm's predictive model or algorithm to determine the ordering of options is an essential component of choice architecture. We call this algorithm the firm's "user model." An algorithm is a high-quality user model if there is a strong correlation between the ordering and what the consumer would like best if searching exhaustively (Alba et al. 1997). Third, we demonstrate that different choice architecture interventions interact, and the resulting effects are not simply additive. Choice architecture tools are an ensemble with complex interactive effects. This implies that choice architecture tools, like ordering and partitioning, need to be assessed jointly. Most research, in contrast, focuses on the impact of single choice architecture tools, such as defaults, on consumer choice outcomes (Cadario and Chandon 2020; Johnson et al. 2012; Szasz et al. 2018). We illustrate this type of joint effect by looking closely at how ordering and partitioning improve consumer decision outcomes. We explain that this interaction occurs because both ordering and partitioning can shift consumers' attention to different alternatives and that choice improvement happens when both tools steer attention to better outcomes.

Ordering and Partitioning Can Improve Consumer Health Insurance Decisions

Ordering

Ordering options in a choice set can help consumers because individuals are more likely to choose the options that are presented first. There can be an advantage of being first in an ordered list or a series of stimuli (Bar-Hillel 2015; Mantonakis et al. 2009). Therefore, presenting options that are better for consumers earlier in an ordered choice set can improve choice (Diehl, Kornish, and Lynch 2003; Häubl and Trifts 2000; Ursu 2018).

A common explanation for this effect is that consumers' greater attention to initial options increases the probability of these options being selected (Atalay, Bodur, and Rasolofoarison 2012). If the order is well aligned with consumers' preferences, the initially presented most attractive options will receive more attention, and later—less attractive—options will receive less attention. This effect of focusing on the best options should increase the probability that consumers choose better options, promoting higher-quality decisions (Diehl, Kornish, and Lynch 2003; Häubl and Trifts 2000; Ursu 2018).

But countervailing effects also suggest that ordering might not be as helpful to consumers as it could be. Consumers can search *too much* when they are presented with an ordered list because it is so easy for them to inspect more options (Diehl 2005). This increase in search with ordering can lead to worse choice outcomes when it is challenging for consumers to find the best option. Because consumers may consider more options, they may be more likely to end up focusing on inferior options presented later in an ordered list. Ordering can also increase the complexity of comparing options by bringing more similar options closer together, creating choice conflict

and increasing the difficulty of selecting the best option (Dellaert and Häubl 2012). Thus, ordering may not always help choice.

Partitioning

Partitioning separates a choice set into two or more sets that can be inspected independently (Cheema and Soman 2008; Dorn, Messner, and Wänke 2016; Johnson et al. 2012). This is common in online shopping. Partitioning does not remove options but influences choice while also preserving consumer autonomy. Many websites present partitions that allow consumers to easily transition from an initial smaller selection to a second larger selection of options. For example, Google Flights first presents a small “Best Flights” subset of all possible flights, followed by a set of “Other Flights.” A recent search presented 4 best flights and 216 other flights. Partitioning focuses consumers’ attention on the initial set of presented products, and a trivial action (e.g., clicking a button) reveals the remaining options. It makes examining the initial set of alternatives slightly less effortful for the consumer than examining the remaining set. We predict that the immediate effort of having to click through to the next page with more options can overwhelm the potential returns that could occur if consumers saw better options on that page (Häubl, Dellaert, and Donkers 2010; Hogarth and Einhorn 1992; Wilson et al. 2000). Partitioning can increase consumers’ focus on a smaller set of alternatives and alleviate any tendency to search “too much” in ordered sets.

User Models and Choice Architecture

Ordering and partitioning can focus consumers’ attention on a more attractive set of options, but both assume that the choice architect can identify the best choice for the consumer. We define a “user model” as the algorithm used by the choice architect to match, probabilistically, different products to a consumer on the basis of the products’ characteristics. When the user model is accurate, there is a strong correlation between the ordering and the products that are most attractive for a consumer. In that case, choice architects can encourage consumers in the right direction, supporting better decision making (Sunstein 2018; Thaler and Sunstein 2008). But if the user model is inaccurate, choice architecture can lead to worse outcomes because it directs consumers’ attention toward inferior options.

Because user models seem important, it is surprising that choice architecture research has not systematically examined their impact. User models can be inaccurate for at least two reasons. The first reason is heterogeneity in preference. The best ordering may differ dramatically for different users, and it may not be possible for firms to make a correct prediction for each consumer. New customers present a “cold start” problem because the website does not know enough about the user to make good recommendations (Padilla and Ascarza 2021). A choice architect who is naive about a consumer’s preferences cannot use preference information to order the options

or to select an appropriate subset in a partition for this consumer. The second reason is that the choice architect may be misguided in their beliefs about the consumers’ needs. Policy or political considerations may arise. In health insurance decisions, government-based choice architects may be reluctant to order products in terms of a particular user model because there might be a debate on what ordering criteria should be used. Further, private-sector choice architects in health insurance may not always have the consumer’s best interests in mind. Their goal may be to maximize their payments as brokers, for example. In those cases, a nudge can turn into a “sludge” and harm consumer decisions (Sunstein 2021). The accuracy of the user model determines whether ordering and partitioning are helpful, make no difference, or are potentially harmful.

Attention Mediates the Effect of Choice Architecture

How do ordering and partitioning affect choice? We hypothesize that ordering and partitioning focus consumers’ attention on a subset of options (Johnson 2021; Payne 1976). This focus can improve consumers’ decisions if it increases the attention they pay to the best options (see Figure 1). Ordering based on a high-quality user model can achieve this positive effect because it presents the better options first. Partitioning reduces the number of options that consumers examine. For both ordering and partitioning, focusing attention on a subset of options can help improve choice with a high-quality user model but hurt if user model quality is low and more attractive alternatives are ignored (Caplin and Dean 2015; De Los Santos, Hortaçsu, and Wildenbeest 2012).

Hypotheses

In summary, we suggest two choice architecture tools that can, in principle, improve consumer health insurance choice: (1) ordering the options in a choice set according to a high-quality user model, and (2) partitioning choice sets into two sets: a primary set with a small number of alternatives and a secondary set with all the other alternatives. Together, we predict that these interventions can improve consumer choice through a greater focus on a smaller number of higher-quality options. There may be reasons to be skeptical about the effects of ordering based on a high-quality user model. Consumers tend to search too much in these types of ordered sets because searching is easier (Diehl 2005). We predict that partitioning can overcome this downside and improve consumer choice by putting up a behavioral barrier to reduce searching. Of course, the effectiveness of both interventions depends on the accuracy of the user model. Therefore, we hypothesize:

H₁: Ordering the options in a choice set on the basis of a high-quality (vs. random) user model improves consumers’ health insurance choices.

H₂: Partitioning a choice set improves consumers’ health insurance choices when the order of options in the set is

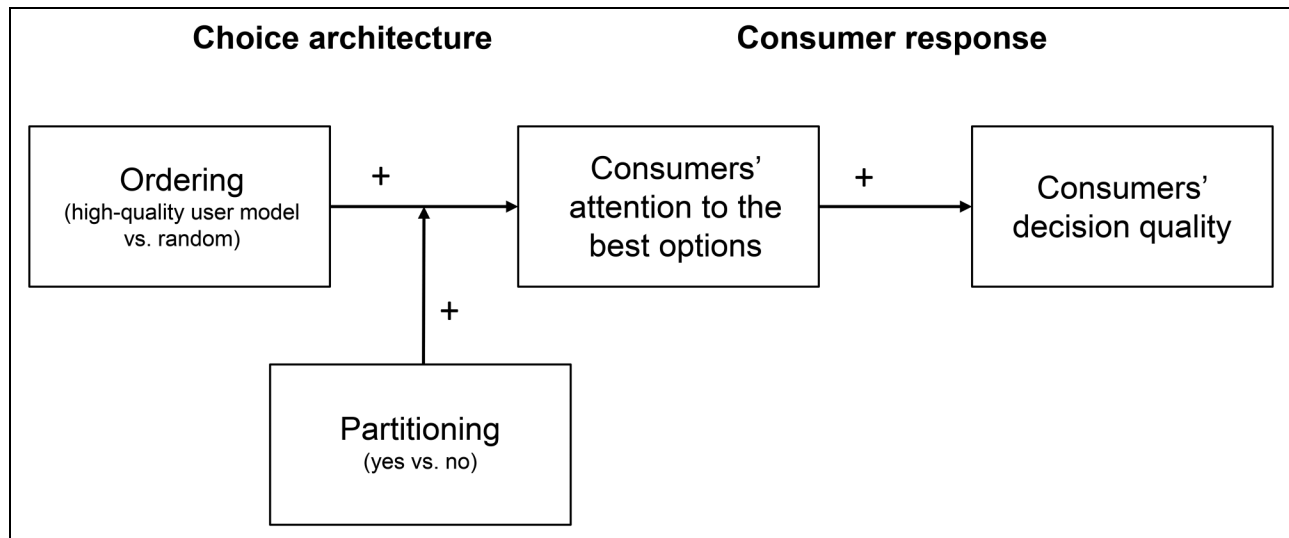


Figure 1. The Joint Impact of High-Quality Ordering and Partitioning on Consumer Decision Quality.

based on a high-quality user model but not when the order is based on a random user model.

We expect these tools to improve choice by changing how the consumer allocates attention. Ordering suggests that options presented at the top of the list will receive the most attention. Thus, we predict that, compared with random ordering, high-quality ordering will increase consumer focus on the best options. Similarly, we predict that partitioning will cause individuals to examine only a smaller initial set and allocate less attention to later options. When the two choice architecture tools are combined and the ordering is based on a high-quality user model, we hypothesize that the stronger focus on a small set of alternatives due to partitioning will promote greater attention to the best options. However, when ordering is based on a random user model, consumer attention to the initial options will mainly highlight average-quality options, and we predict that partitioning reduces attention to the best options. More formally, we hypothesize:

H₃: Ordering the options in a choice set on the basis of a high-quality (vs. random) user model increases consumers' attention to the best options.

H₄: Partitioning a choice set reduces the number of options consumers examine.

H₅: Partitioning a choice set in which the order of options is based on a high-quality user model strengthens the impact of this ordering on consumers' attention to the best options.

Finally, jointly, these three effects suggest a mediation process by which ordering and partitioning impact consumer choice outcomes. First, we predict that consumers' attention to the best alternatives mediates the impact of ordering on consumer decision outcomes. Second, we predict that the main effect of

partitioning on consumer decision outcomes is mediated by consumers' examination of a smaller number of alternatives. Third, we predict that the positive impact of partitioning on the effect of high-quality ordering on consumer decision outcomes is also mediated by consumers' attention to the best alternatives (mediated moderation).

H_{6a}: The impact of ordering the options in a choice set on consumers' health insurance choices is mediated by attention to the best options.

H_{6b}: The impact of partitioning the options in a choice set on consumers' health insurance choices is mediated by the number of options consumers examine.

H_{6c}: The synergistic impact of ordering and partitioning the options in a choice set on consumers' health insurance choices is mediated by attention to the best options.

Overview of Experiments

We present four studies that examine the joint effects of ordering and partitioning and their possible effect on consumers' health insurance choices. We first present, as a pilot study, field data showing the positive effect of these interventions in a realistic environment, when adopted by a commercial health insurance broker. Study 1 is an incentive-compatible experiment in a setting closely modeled on the U.S. health insurance exchange websites (see Wong et al. (2016) for a descriptive survey of the U.S. health exchange marketplace). The results show the positive interactive effects of ordering and partitioning (H₁ and H₂), and, using a MouselabWEB data analysis, we show that this impact is mediated by consumers' attention to the best options and the number of options examined (H₃ through H₆). Study 2 replicates the results for H₁ and H₂ and extends them to the case of a low-quality user model (i.e., generating an order that is antagonistic to the consumer's

preferences). The results show that with a low-quality user model, partitioning harms consumer choices. They also show that with a random user model, partitioning may also lead to worse consumer choices, depending on the quality of the alternatives in the partitioned set. Finally, Study 3 investigates the effects of ordering and partitioning on consumer decision quality in a setting modeled on the ordering and partitioning that are most common on U.S. health insurance exchange websites. This common setting is an ordering based on health insurance premium only (“premium-based” ordering), which is often combined with a partitioning of ten options per page. We investigate whether a higher-quality ordering based on a prediction of consumers’ total health care spending (“predicted-spending-based” ordering) combined with a partitioning of three options can improve consumer decision quality, compared with current practice. The results of Study 3 show that this is indeed the case.

We adopt a framed-field experiment approach in all three studies, whereby real consumers make choices about realistic options, in a setting that mimics the actual choice environment that consumers face on the U.S. health insurance exchange websites (Harrison and List 2004). The participants are given a well-defined goal for buying insurance and instructed to find the option that is closest to that goal. Their objective is to minimize total expected costs given specified health care needs. This approach allows us to objectively define the quality of the decision made by each participant and to quantify deviations from the best choice outcome. It also eliminates unobserved preference heterogeneity between participants (Johnson et al. 2013). The materials, data, and analysis syntax for all three studies are available at https://osf.io/xvdpn/?view_only=d40fd1c9813845b6ae51b53c744eab99.

Pilot Study: The Impact of Ordering and Partitioning on Consumer Health Insurance Choice

Our pilot study explores the effects of ordering and partitioning using field data from a commercial health insurance broker in the Netherlands. The broker perceives itself to be a strong innovator in the market, and in 2010 it changed its product comparison website introducing a higher-quality ordering with partitioning aiming to improve consumers’ choice outcomes. Health insurance plans can be purchased directly through the website, and many consumers go to the site each year to switch insurance providers. The prior choice architecture provided a complete list of health insurance products screened by a consumer’s prespecified criteria. These were ordered from low to high according to the plan’s premium, and ten options were displayed per web page. The intervention improved the user model in two ways. First, by including quality in addition to price in the ordering, the model more closely matched user preferences, based on the broker’s market research, even though the same user model was used for every visitor. Second, by introducing a partitioning that presents a subset of the top three plans, the model potentially increases the effect of ordering. There was no change in the total number of options. Consumers could simply click to see the full list of health insurance products.

We compared the impact of this choice architecture on consumers’ choices by comparing the data from the year before the website’s choice architecture redesign to the choices made after the change. Before the redesign (price-only ordering, no partitioning), 47.7% of the visitors who bought health insurance selected the first-ranked alternative (of a total of 8,581 consumer visits, 41.3% of visitors identified as female, and the average age was 36.5 years). After the redesign (price- and quality-based ordering, with partitioning), this increased to 60.7% (of a total of 34,677 consumer visits, 45.6% of visitors identified as female, and the average age was 39.2 years). To test the significance of this difference, we estimated a logistic regression model of whether or not consumers selected the first-ranked alternative (i.e., the alternative presented at the top of the list) in the years before and after the redesign was introduced. The result shows a significant positive effect of the choice architecture redesign, which suggests a strong combined impact of ordering and partitioning on consumers’ health insurance choices ($\beta_{\text{redesign}} = .53$, $SE = .02$, $p < .001$).

However, although the website managers believed that choices were improved with the new design, as more consumers chose the first-ranked alternative, we had no way of demonstrating that choice outcomes were improved objectively. We could not know whether the first-ranked alternative in either year was indeed the best option for a consumer. In addition, since the choice architecture change occurred as part of an annual update of the site, the pilot study did not control for factors such as the properties of the health insurance policies, a prelaunch media campaign, and potential changes in the demographics of consumers.

We turned to a controlled framed-field experiment to validate these findings. In the experiment, we manipulated ordering and partitioning independently, and all conditions had identical alternatives. Participants were randomly assigned to the different conditions. The results replicated the positive synergistic effect of high-quality ordering and partitioning from the pilot study (see Web Appendix A for details).

Study 1: The Impact of Ordering and Partitioning on Consumer Decision Quality and the Mediating Processes

In Study 1, we constructed an incentive-compatible experiment with a choice task modeled on a typical consumer health insurance decision task on the U.S. health insurance exchange websites used by consumers to purchase health insurance plans. We first analyzed the U.S. health exchanges in all 50 states in 2020 (of which 14 states had set up their own exchange and 36 states were using the federal health exchange, HealthCare.gov) for a person 40 years of age with medium health care use. We found that the mean number of products offered was 38.4 (with a minimum of 5 and a maximum of 95) and that exchanges most often presented consumers with ten options on the first page. Consumers could then click to see the remaining options. To mimic this choice architecture, we presented participants with 40 health insurance products drawn from

one of the U.S. health exchanges and included ten-option partitioning as one of the conditions.

Method

Experiment. Participants were given a clearly defined objective, which was to minimize total expected costs given a person's relevant health care usage. The health insurance plans differed on three key characteristics for the U.S. market: the monthly premium, the doctor visit copayment, and the annual deductible. These attributes reflect the monetary cost components of the health insurance and were used successfully in previous research (Bhargava, Loewenstein, and Sydnor 2017; Johnson et al. 2013). This information allowed participants to identify the annual expected costs of each of the health insurance products. To increase motivation, we also gave participants a monetary reward based on their performance: in addition to the \$2 completion fee, they could earn an additional bonus of up to \$4 depending on how closely their choices matched the assigned objective.

Participants were randomly assigned to six treatment conditions, based on a 2 (ordering: high-quality and random user model) by 3 (partitioning: three-option, ten-option, and no partitioning) between-subjects design. All conditions presented the same 40 health insurance plans. In the high-quality user model ordering condition, alternatives were ordered perfectly from best to worst (minimizing the participants' expected costs given the decision rule; see Web Appendix B1 for ordering details). Participants were informed about the ordering as follows: "The comparison website attempts to rank the options based on the projected cost for your family in your state. However, the ordering may not be fully accurate." The random user model ordering condition presented the same randomized order of alternatives to participants with a zero correlation between the presented ordering and the optimal ordering. Participants were informed as follows: "The comparison website does not sort the options in any particular order." The partitioned conditions showed participants either the first three or ten listed products with the possibility to click through to see the complete list of products. Conditions without partitions showed participants the complete list of 40 products. We predicted that with high-quality ordering, the three-option partitioning would produce the best decisions, followed by the ten-option partitioning, and that the no-partitioning condition would produce the worst choices.

Process data. To observe attention, we had participants make their decisions using MouselabWEB (see www.mouselabweb.org). This technique is shown to be useful in the study of individual choice, as well as games and time preferences (Costa-Gomes and Crawford 2006; Gabaix et al. 2006; Payne, Bettman, and Johnson 1993; Willemsen and Johnson 2011). MouselabWEB presented the options as a table, consistent with displays in the state-based and national exchanges. Each row presented a different health plan, and each column presented a different plan attribute (see Web Appendix B2). In the MouselabWEB interface, when participants used a mouse to move the cursor over each cell, the values of the attribute for that option were instantly revealed. When the cursor

exited the cell, the information was again hidden. This process revealed how often, for how long, and in which order participants acquired the information about the attributes of each plan. Previous research suggests a close correlation between MouselabWEB observations and eye-tracking data, an alternative way of tracing participants' decision processes (Lohse and Johnson 1996; Reisen, Hoffrage, and Mast 2008). Following previous research (Willemsen and Johnson 2011), we eliminated all information acquisitions of less than 200 milliseconds. These are too short to be seen and probably reflect travel of the cursor through a cell in transit to another. We also eliminated outlier observations of more than 50,000 milliseconds (Willemsen and Johnson 2011).

Using the MouselabWEB process data, we constructed two variables reflecting the mediators in our hypotheses. The first variable captured consumers' attention to the best options. It was measured by the information participants acquired on the best alternatives (i.e., the three options with the lowest total cost in the defined utility task) relative to the total number of inspections (i.e., number of information acquisitions for the best three options divided by the total number of information acquisitions). A value of 0 corresponds to no focus on the best three options, and a value of 1 means looking at only the three best options. The second process variable captured the number of options examined; that is, the total number of options for which a participant acquired information. This number could be observed directly from the process data.¹

Sample. The participants were recruited through a commercial panel (ROI Rocket) and restricted to U.S. participants. We predetermined the sample size to be 200 per treatment condition, and data collection was stopped when all the conditions contained at least this number. This provided a total of 1,324 completed responses. We eliminated 27 outlier participants who were not engaged in the task and had made extraordinarily few (i.e., less than five) information acquisitions during the decision (Willemsen and Johnson 2011). The average age of participants was 53.1 years, 67.1% identified as female, 42.9% had obtained a bachelor's or higher degree, and 76.1% were the primary insurance decision maker for their household.

Results

Decision quality. We first investigate the effects of ordering and partitioning on consumer decision quality (H_1 and H_2). Decision quality was measured in dollars, expressing the consumer's excess payment compared with the most cost-efficient alternative. This metric means that a smaller number reflects a higher decision quality. We tested the significance of the results using a 2×3

¹ For exploratory purposes, we also constructed a third process variable, which reflects the attention participants paid to the options conditional on them being examined. This variable did not mediate the impact of ordering and partitioning on consumer decision quality and was therefore excluded from the analyses reported in the article. We thank one of the reviewers for suggesting this additional analysis.

analysis of variance.² We find that consumers choose dramatically better health insurance plans (in other words, overpay less) with a high-quality ordering than with a random ordering ($M_{\text{High}} = \$1,185.57$ vs. $M_{\text{Rand}} = \$4,305.08$, $F(1, 1,291) = 1,709.82$, $p < .001$, $\eta_p^2 = .57$). This result supports H_1 .

As we predicted (H_2), this effect is even stronger when the choice set is partitioned, as evidenced by the significant interaction effect of ordering and partitioning ($F(2, 1,291) = 12.54$, $p < .001$, $\eta_p^2 = .02$). There is a marked increase in decision quality caused by partitioning with the high-quality user model that does not replicate when the user model is random. Figure 2 shows that with a high-quality user model, decisions improved and overpayment was further reduced by three-option partitioning ($M_{3\text{Part_High}} = \745.47) compared with ten-option partitioning ($M_{10\text{Part_High}} = \$1,265.62$) and no partitioning ($M_{\text{NoPart_High}} = \$1,627.27$). But with a random user model, partitioning was hardly helpful (three-option partitioning: $M_{3\text{Part_Rand}} = \$4,006.43$) or not helpful (ten-option partitioning: $M_{10\text{Part_Rand}} = \$4,680.82$) to improve consumer choice relative to no partitioning ($M_{\text{NoPart_Rand}} = \$4,174.09$). There is also a main effect of partitioning. This effect is nonlinear: Decision quality is highest with three-option partitioning and lower with ten-option partitioning and in the no-partitioning case ($M_{3\text{Part}} = \$2,263.88$ vs. $M_{10\text{Part}} = \$3,096.84$ vs. $M_{\text{NoPart}} = \$2,894.38$, $F(2, 1,291) = 26.32$, $p < .001$, $\eta_p^2 = .04$).

Attention to the best options and number of options examined. We hypothesized that with a high-quality user model, consumers' attention would be focused more on the best options (H_3). Analysis of variance shows that this main effect is indeed significant in the expected direction ($M_{\text{High}} = .55$ vs. $M_{\text{Rand}} = .03$, $F(1, 1,291) = 3,285.12$, $p < .001$, $\eta_p^2 = .72$; see Figure 3, Panel A). There is also a significant main effect of partitioning that increases attention to the best options ($M_{3\text{Part}} = .50$ vs. $M_{10\text{Part}} = .20$ vs. $M_{\text{NoPart}} = .16$, $F(2, 1,291) = 493.84$, $p < .001$, $\eta_p^2 = .43$).

We also hypothesized that consumers would examine fewer options when partitioning is introduced (H_4). We indeed find a significant negative effect of partitioning ($M_{3\text{Part}} = 8.10$ vs. $M_{10\text{Part}} = 14.22$ vs. $M_{\text{NoPart}} = 33.28$, $F(2, 1,291) = 592.70$, $p < .001$, $\eta_p^2 = .48$; see Figure 3, Panel B). There is a directionally negative, nonsignificant effect of high-quality ordering on the number of options examined ($M_{\text{High}} = 17.31$ vs. $M_{\text{Rand}} = 18.75$, $F(1, 1,291) = 3.74$, $p = .053$, $\eta_p^2 = .003$).

Finally, we hypothesized that partitioning combined with high-quality ordering would further strengthen the attention paid to the best options (H_5). We find that the interaction of partitioning and ordering is significant ($F(2, 1,291) = 632.46$, $p < .001$, $\eta_p^2 = .50$). Figure 3, Panel A, shows the effects of ordering and partitioning on attention to the best options. Partitioning strengthens the positive impact of high-quality ordering on attention to

the best options. With random ordering, the best options receive little attention. Without partitioning, the attention is roughly equal to chance (it is .07, which is close to the 3/40 chance level), and when partitioning is introduced, attention drops below chance (to .01).

Mediation analysis. To test the hypothesized mediated moderation structure, we use Hayes's (2013) PROCESS module (Model 8; 5,000 bootstrap samples; see Web Appendix B4). The dependent variable was the amount consumers overpaid. Ordering (high-quality user model, random user model), partitioning (three-option, ten-option, none), and their interaction were the independent variables coded as indicators. The interaction between ordering and partitioning was included, as were the two mediating process variables (i.e., attention to the best options and number of options examined).

Following the procedure outlined in Zhao, Lynch, and Chen (2010), we analyze the indirect effects of ordering and partitioning on the decision outcome via the two parallel mediating decision processes. To test H_{6a} and H_{6c} , we look at the mediation via *attention to the best options*. The results show that this variable is a significant mediator of the effect of ordering on the dependent variable for each of the three partitioning conditions (three-option, ten-option, none) ($a_1 \times b_1 = -1,012.76$, 95% confidence interval [CI] = [-1,607.47, -460.96]; $a_1 \times b_1 = -442.68$, 95% CI = [-695.10, -204.62] and $a_1 \times b_1 = -198.25$, 95% CI = [-315.91, -89.76]). This result provides support for H_{6a} . Importantly, the differences between the conditional indirect effects for the three partitioning conditions are also significant (the indices of moderated mediation between the two partitioned conditions and the nonpartitioned conditions are -814.51 , 95% CI = [-1,292.84, -371.26], and -244.44 , 95% CI = [-387.99, -112.44]). This result provides support for H_{6c} , indicating that the moderating effect of partitioning on the effect of ordering is mediated by consumers' attention to the best options.

To test H_{6b} , we look at the mediation via the *number of options examined*. The results show that this variable is indeed a significant mediator of the effect of both partitioning levels (three-option and ten-option) on the dependent variable for each of the two ordering conditions (three-option partitioning: $a_1 \times b_1 = 335.79$, 95% CI = [158.56, 534.82]; $a_1 \times b_1 = 346.40$, 95% CI = [165.71, 534.87], and ten-option partitioning: $a_1 \times b_1 = 255.95$, 95% CI = [123.99, 401.38], $a_1 \times b_1 = 262.70$, 95% CI = [125.51, 406.75]). This result provides support for H_{6b} . Thus, partitioning leads to a lower number of options being examined, which in turn leads to more overspending (i.e., lower decision quality). Note that in the case of high-quality ordering, this negative effect is more than offset by the positive moderating impact of partitioning on the effect of ordering on decision quality (via attention to the best options), as evidenced by the support for H_2 and H_{6c} .

² We also ran this analysis controlling for participants' age and education. Including these variables and their interactions with the choice architecture tools did not affect the conclusions regarding the hypothesis tests (see Web Appendix B3 for details).

Discussion

The results of Study 1 provide both outcome and process-level support for the hypothesized impact of ordering and partitioning

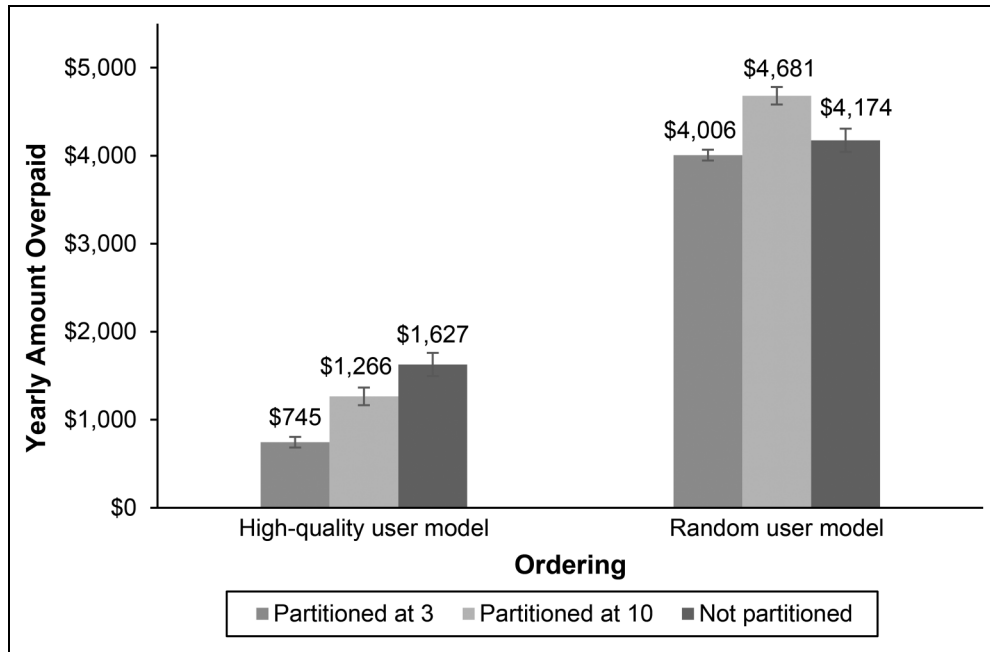


Figure 2. The Impact of Ordering and Partitioning on Consumer Decision Quality (Study 1).

Notes: Error bars = ± 1 SE.

on consumer choice quality. In this incentive-compatible study closely modeled on the U.S. health insurance exchange websites, we find that ordering based on a high-quality user model with partitioning is beneficial for consumers. Thus, the results strongly support H_1 and H_2 . Process data supported that the hypothesized change in decision process occurs, focusing attention on the better options, and that this increase in attention mediates the improvement of decisions (H_3 through H_6). Finally, the results indicate that the improvement in decision quality is a joint effect of ordering and partitioning. Ignoring this joint effect would be misleading.

Study 2: Can a Low-Quality User Model Harm Choice?

What happens when a user model does not reflect the preferences of the consumer? For example, a user model might reflect a platform's desire to sell higher-priced, more profitable options that may not be preferred by consumers. In Study 2, we extended our analysis to the more extreme case where ordering may *decrease* consumer welfare. This study also replicates Study 1 in an environment without process tracing, allowing us to examine possible reactivity to these measures.

Method

Experiment. The study was similar in structure to Study 1 with a few changes. First, we did not collect process data and presented a smaller set of health insurance products to participants. We also manipulated three levels of user model quality. The user model qualities were based on the correlation between the

order of plans presented to participants according to the goal of minimizing expected costs (high-quality, low-quality, and random user model; see Web Appendix B1 for ordering details). In the high-quality user model ordering condition, alternatives were ordered approximately from best to worst, with a positive correlation of .7 between the presented and the optimal ordering. The low-quality user model ordering condition presented an order of alternatives that was opposite, with a correlation of $-.7$ between the presented and the optimal ordering. Both random ordering conditions had a zero correlation between the presented and optimal ordering. By necessity, the correlations constrained the way the partitioned set was constructed. The best alternatives had to be present in the partitioned set for the high-quality user model and outside the partitioned set in the low-quality user model condition. To allow us to separate this effect from the effect of correlation, we used two versions of the random order: one that included the best alternative in the partitioned set and one that did not (labeled "random (in)" and "random (out)"). Participants within each ordering condition saw the same order of alternatives. Finally, as in Study 1, we also varied whether the set was partitioned or not.

Participants were randomly assigned to one of eight treatments based on a 4 (ordering: high-quality user model, low-quality user model, and the two random orders) by 2 (partitioning: yes vs. no) between-subjects experimental design. All conditions presented the same eight health insurance plans. To control for the possible effect of variation in instructions, and in contrast to Study 1, we gave participants the same information in all ordering conditions: "The comparison website attempts to rank the options in order of attractiveness for the average family in your state. However, the ordering may not be fully accurate." The range of pricing for the




Figure 3. Mediation Process Measures (Study I).

plans was similar to those of the top 50% of those found at HealthCare.gov. Participants made two choices, both in the same experimental condition, but with different alternatives.

Sample. The participants were recruited online through ROI Rocket, a U.S. market research firm. Participants received a small, fixed payment from the firm for their participation. We predetermined the sample size to be 200 per treatment condition. The assignment to treatment conditions was randomized, and the data collection was stopped after all the conditions contained at least 200 observations. The data set consisted of responses from 1,646 respondents. The average age of the participants in the sample was 50.74 years old, 62.3% identified as female, and 50.5% had obtained a bachelor’s or higher degree.

Results

We tested the effects of ordering, partitioning, and their interaction (see Figure 4, Panel A) on decision quality (excess payment) using a repeated-measures analysis of variance with the two responses per participant. There is a beneficial effect of higher-quality ordering (i.e., it lowers the amount consumers overpay; $M_{High} = \$467.52$ vs. $M_{Rand} = \$559.33$ vs. $M_{Low} = \$743.99$, $F(2, 1,640) = 18.65$, $p < .001$, $\eta_p^2 = .022$). This finding supports H_1 . The effect of partitioning also depends on the quality of ordering. There is a significant interaction effect, consistent with H_2 ($F(2, 1,640) = 13.80$, $p < .001$, $\eta_p^2 = .017$). Simple-effects analyses reveal that partitioning is helpful in the high-quality ordering condition ($M_{Part} = \$409.70$ vs. $M_{NoPart} = \$526.40$, $F(1, 404) = 3.53$, $p = .061$, $\eta_p^2 = .009$). However it is harmful in the low-quality ordering condition

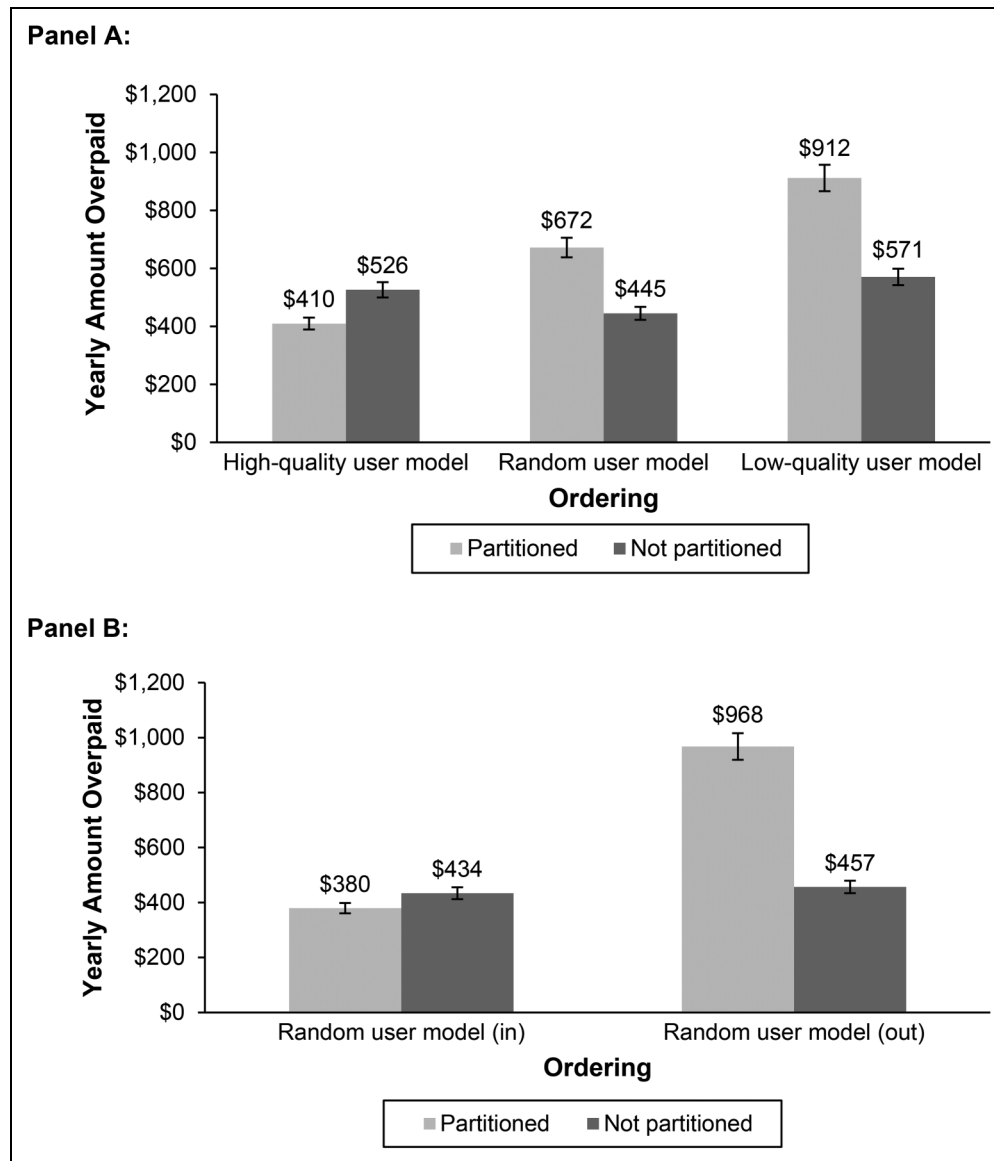


Figure 4. The Impact of Ordering and Partitioning on Consumer Decision Quality (Study 2).

Notes: Error bars = ± 1 SE. For expositional clarity, the results of the two random user model orderings are combined in Panel A.

($M_{\text{Part}} = \$912.20$ vs. $M_{\text{NoPart}} = \$570.86$, $F(1,414) = 25.29$, $p < .001$, $\eta_p^2 = .058$) and the random ordering condition ($M_{\text{Part}} = \$672.17$ vs. $M_{\text{NoPart}} = \$445.39$, $F(1,822) = 24.17$, $p < .001$, $\eta_p^2 = .029$).³

Next, we conducted a close-up analysis separating the two random ordering conditions (see Figure 4, Panel B) in simple-effects repeated measures analyses of each random ordering condition. The partitioned and nonpartitioned conditions did not differ reliably when the best alternative is included in the partitioned set ($M_{\text{Part}} = \$379.59$ vs. $M_{\text{NoPart}} = \$433.94$, $F(410) = .87$, $p = .352$, $\eta_p^2 = .002$). However, partitioning

strongly reduces decision quality and increases overpayment when the best alternative is not included in the partitioned set ($M_{\text{Part}} = \$967.59$ vs. $M_{\text{NoPart}} = \$456.73$, $F(1, 410) = 61.05$, $p < .001$, $\eta_p^2 = .130$). Thus, the negative impact of partitioning with random ordering can be attenuated depending on whether good options, through chance, end up in the initial partitioned set.

Discussion

Study 2 further supports the idea that ordering and partitioning can improve consumer decision quality (H_1 and H_2), and that these effects depend on the user model. We find that ordering based on a higher-quality user model improves consumers' choice outcomes. With a high-quality user model, partitioning is beneficial for consumers; with a low-quality user model, it is harmful.

³ We also ran this analysis controlling for participants' age and education. Including these variables and their interactions with the choice architecture tools did not affect the conclusions regarding the hypothesis tests (see Web Appendix C for details).

Study 3: Ordering Options by Premium Versus Predicted Spending

In Study 3, we investigated two kinds of ordering that are commonly used in practice on the health care exchanges: ordering based on annual premium and ordering based on predicted health care spending. Premium-based ordering sorts health insurance plans according to one component of their overall cost, the monthly premium. This is the default ordering on most states' websites (see <https://www.healthcare.gov/>). However, because health insurance prices have multiple components, this might be misleading. Other components, like the deductible and copayment, are usually negatively correlated with the monthly premium. This means that choosing a policy with a low premium implies higher deductibles and copayments. To estimate overall cost, consumers need to predict their future health care usage and do some complex calculations (Abaluck and Gruber 2011; Ericson and Starc 2012).

When the health exchanges were introduced, there was concern that consumers would overweight the premium in their health insurance decisions, resulting in choices that would be more expensive overall (Johnson et al. 2013; Wong et al. 2016). Anecdotally, some consumers bought low-premium insurance but were faced with high deductibles and copayments that, in some cases, led them to forgo needed care.

Some exchanges, such as California and Kentucky, implemented predicted-spending-based ordering in the third enrollment period of 2015 to overcome this problem (Wong et al. 2016). This ordering is based on a prediction of a consumer's health care spending. To estimate spending, they use projected health care needs, which potentially represent a more accurate reflection of costs, compared with ordering by premium alone. More states adopted this ordering in the fourth enrollment period in 2016.

In Study 3, we compared the initial standard premium-based ordering with ordering based on a prediction of total health care spending. Although this predicted-spending-based ordering may be more informative than premium-based ordering, it is not a perfect user model. Different consumers have different usage patterns. Exchanges make some attempts at customization, asking consumers which of three prototypical health care usage patterns best describes them, but this does not capture all the details of their health care needs. This study allowed us to see the difference in impact of the two different real-world ordering types on consumer decision quality. In addition, it allowed us to examine the role of partitioning using the level currently used in exchanges and one that would be considerably smaller. In essence, Study 3 examines the impact of these contemplated design changes in choices closely resembling the actual setting.

Method

Experiment. Study 3 used the same health insurance plans and participants' objective goal as in Study 1. Participants were randomly assigned to one of four treatments based on a 2 (ordering type: premium-based vs. predicted-spending-based) by 2

(partitioning size: ten-option vs. three-option) between-subjects experimental design. All conditions presented the same 40 health insurance plans as used in Study 1. To reflect the actual policy of the health insurance sites, we implemented two orderings (see Web Appendix B1 for ordering details). Premium-based ordering sorted the policies from low to high using premiums alone. This resulted in a correlation between the presented ordering and the optimal ordering of .48. Participants were told: "The comparison website ranks the options based on the plan's premium." Predicted-spending-based ordering sorted the policies from lowest to most expensive using predicted total health care expenses. In reality, these predictions are not perfect, and to reflect this uncertainty, we added some noise to the ranking, producing a correlation between the presented ordering and the optimal ordering of .70. Participants were told: "The comparison website attempts to rank the options based on the projected cost for your family in your state. However, the ordering may not be fully accurate." Participants within each ordering condition saw the same order of alternatives. We manipulated partitioning, with the initial set containing the three or ten most highly ranked products. The most common practice used on the exchanges powered by HealthCare.gov is to present ten options on the first page. We investigated whether using three options in a partition would be more helpful to consumers. In all conditions, participants could also click through to see the full list of alternatives, which was ordered the same way as the initial set. Participants received a fixed payment of \$3 for their participation, and they could earn an additional bonus of up to \$4 depending on how close the alternative they selected was to the best alternative. We also measured participants' socioeconomic status (SES) using three items: education, occupational status, and income. These items were combined into one standardized SES score following current practice reported in previous research (Mrkva et al. 2021).

Sample. The participants were recruited online through ROI Rocket, a U.S. market research firm. We predetermined the sample size to be 200 per treatment condition. The study was pre-registered on AsPredicted (<https://aspredicted.org/i8q67.pdf>). ROI Rocket delivered considerably more respondents than the number we had requested, and we included all 946 responses. They were assigned randomly to the four conditions. The average age of the participants was 53.96 years, 49.2% identified as female, 42.5% were in managerial/professional jobs (vs. 14.1% homemaker, 17.5% clerical, and 22.7% blue-collar jobs), 48.1% had a personal yearly income above \$50,000, 46.6% had obtained a bachelor's or higher degree, and 79.2% were the primary insurance decision maker for their household.

Results

To test the impact of ordering and partitioning, we conducted an analysis of variance, modeling the amount consumers overpaid as a function of ordering, partitioning, and their interaction (see Figure 5). The results show that, as expected,

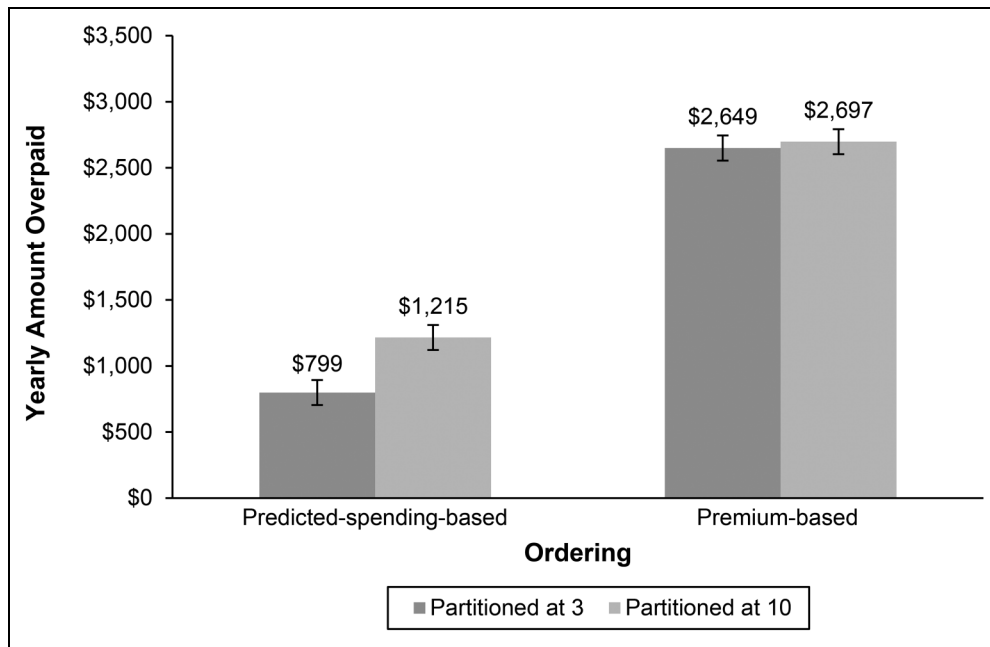


Figure 5. The Impact of Ordering and Partitioning on Consumer Decision Quality (Study 3).

Notes: Error bars = ± 1 SE.

there is a significant positive effect on decision quality of predicted-spending-based versus premium-based ordering (i.e., predicted-spending-based ordering lowers the amount consumers overpay; $M_{\text{Spending}} = \$998.26$ vs. $M_{\text{Premium}} = \$2,673.85$, $F(1, 942) = 283.56$, $p < .001$, $\eta_p^2 = .231$). Smaller partitions are beneficial ($M_{\text{Part3}} = \$1,677.21$ vs. $M_{\text{Part10}} = \$1,971.97$, $F(1, 942) = 5.51$, $p = .019$, $\eta_p^2 = .006$). The interaction effect between ordering and partitioning is marginally significant ($F(1, 941) = 3.46$, $p = .063$, $\eta_p^2 = .004$), and the positive effect of partitioning on decision quality is larger when the order is determined by predicted spending. A simple-effects analysis reveals there is a beneficial effect of partitioning in the predicted-spending-based ordering condition ($M_{\text{Part3}} = \$798.70$ vs. $M_{\text{Part10}} = \$1,215.16$, $t(478) = 2.70$, $p = .007$, Cohen's $d = .247$), but not in the premium-based ordering condition ($M_{\text{Part3}} = \$2,649.01$ vs. $M_{\text{Part10}} = \$2,697.24$, $t(464) = .391$, $p = .696$, Cohen's $d = .036$).

We conclude that premium-based ordering is of sufficiently high quality that partitioning is not harmful, but also that it is not of high enough quality that partitioning is beneficial. In summary, in this context, ordering based on predicted spending is much better: when combined with three-option partitioning, it saves consumers \$1,898 compared with premium-based ordering with ten-option partitioning. Finally, we also examined differences in the responses of participants with varying SES scores to the four treatment conditions but found no significant effects.

Discussion

Study 3 suggests that ordering and partitioning can help improve consumer decision quality in a setting similar to practice. Ordering based on predicted spending is beneficial compared with premium-based ordering, and, in that case,

partitioning is also beneficial for consumers, no matter what level of SES we examine.

General Discussion

This research shows that ordering and partitioning can have a significant and beneficial impact on consumer choice outcomes, but that achieving this effect depends critically on having a good user model that accurately predicts which options are most attractive to consumers. Consumers were more likely to choose the best alternative when presented with a set that was ordered according to a high-quality user model and partitioned into a small initial set of alternatives. The process data from Study 1 show that this impact occurs because partitioning focuses consumer attention on a small set of high-quality options when combined with high-quality ordering. Yet, as shown in Study 2, the positive effect of partitioning becomes negative when the ordering is based on a low-quality user model. When this occurs, consumers are harmed by partitioning. Finally, as an application of our findings, Study 3 shows that the most common approach in health care exchanges, premium-based ordering combined with ten-option partitioning, can be improved by presenting consumers with predicted-spending-based ordering combined with three-option partitioning.

Policy Implications

One of the advantages of using simulated decisions is that we can take a closer look at welfare effects for consumers. We should take these estimates with some caution because even though the decisions are based on realistic stimuli and made by incentivized consumers, there are other factors, such as advertising, that occur in actual marketplaces but

not in our studies. Still, our Study 3, which resembles actual marketplaces, provides initial estimates of the cost of ignoring these choice architecture interventions. Study 3 compares predicted-spending-based and premium-based ordering, both orderings that are currently used in exchanges. Consumers, on average, made larger mistakes, costing about \$2,650 annually, when the policies were ordered by premium. This ordering, used on many exchanges, represents a user model that does not correspond to their goal. Employing a higher-quality user model that ordered the options based on total predicted spending reduced that error to about \$1,000. Introducing a new choice architecture, three-option partitioning, further reduced errors to just under \$800. Together, these changes result in a yearly savings of about \$1,850 per consumer, compared with premium-based ordering with ten-option partitioning. Although these figures are specific to this study, they are similar to the size of errors reported in previous framed-field experiments (Johnson et al. 2013) and field studies (Bhargava, Loewenstein, and Sydnor 2017). Given that slightly over 200 million people are covered by private plans and 100 million by government health care plans, many of which involve choice, the stakes of getting health insurance choice architecture right could be measured in hundreds of billions of dollars in the United States alone. Not paying attention to choice architecture has huge costs. We would argue that by not adopting these changes, governments, firms, and consumers are creating a health insurance system that is less efficient, causing consumers, and, through subsidies, governments, to spend much more on health insurance than needed, using money that could be better spent on other priorities.

For consumer protection policy, asking whether choice architectures serve consumers' best interests poses new challenges and research questions. Choice architecture design could become part of regulatory frameworks. Firms might be asked to store the combination of choice architecture environments, available product alternatives, and consumer decision outcomes, for example, when offering consumers (online) financial advice (Baker and Dellaert 2018). Such data would allow for post hoc validation of the quality of the advice provided and whether consumers were supported in making better decisions. Consumer protection policy makers could also bring regulatory oversight to deceptive choice architecture shown to lead to worse consumer choices. It would be valuable to study the effectiveness of such possible interventions empirically.

It is also important to develop further insights into which choice architecture ensembles are most promising to activate consumers to consider changing their health insurance in general. For example, an approach suggested to promote comparisons by consumers without the need to provide detailed personalized information is to require that different firms offer at least the same basic version of a (financial) product (Lynch 2009).

Managerial Implications

For managers, this research highlights the potential tension between the interests of consumers and the profit-maximizing

interests of health insurance firms, but there is a twist. A classical dilemma facing insurers is adverse selection: consumers may know more about their health states than firms know, thus reducing firms' profits. However, there may be situations where firms know more about some health outcomes than consumers do and can prevent consumers from making mistakes. Health insurers may have, because of their extensive data, a more accurate idea of health risks than the individual consumer has, especially when dangerous health events are both rare and potentially catastrophic. For example, the average consumer probably has an inaccurate view of the probability and financial cost of an accident or a serious disease.

Health insurers may, in the short term, exploit these informational asymmetries, but another path may lead to greater consumer satisfaction and loyalty in the long run. Firms may build choice architectures that will help consumers make better decisions. Although the economics of this strategy is an open question, it depends on the ability of health insurance firms to develop accurate user models. A requisite step would be convincing consumers that there is value in following the firm's recommendations. Future research could address consumer trust in and reactance to firm recommendations.

Another line of research could investigate how online intermediaries' choice architecture tools can also help shape markets. For example, the new choice architecture design in the online health exchange in our pilot study not only helped consumers make more informed decisions but also led to changes in the market supply. Health insurers strengthened the consumer-satisfaction-driving aspects of their offerings in response to the greater prominence of this feature in the platform's ordering. However, the impact of choice architecture may be underappreciated in practice (Benartzi et al. 2017; Zlatev et al. 2017; but also see McKenzie, Leong, and Sher [2021] for nuances).

Theoretical Implications

Our idea of choice architecture ensembles is only a beginning. Research is needed to develop further insights into which choice architecture ensembles are the most promising to improve consumer financial decision making. Although choice architecture research has been a successful endeavor to date, better understanding the interaction of ensembles seems important. Applications that look at one tool at a time may miss important opportunities for improving choices.

Our research also makes explicit another idea long implicit in the choice architecture literature: the idea of a user model. If choice architects have access to high-quality user models (i.e., accurate prediction algorithms), they can intervene more effectively because they can identify more precisely what is the correct option for each consumer. Although we explored the impact of user model quality in the context of ordering and partitioning, it is also relevant to other choice architecture interventions, such as defaults (Goldstein et al. 2008).

An interesting question for future research is to develop a conceptual framework for helpful choice architecture depending on whether user models make strong or weak predictions.

The strength of a user model can affect the implementation of a choice architecture. Partitioning offers a good example. If a health insurance firm is certain of the quality of its user model for a given consumer, the partitioned set can be small, but as uncertainty increases, the size of the partitioned set can increase as well, reflecting that uncertainty. There is, therefore, no “one size fits all” optimal partitioning size; rather, there are different optimal sizes for different markets, different consumers, and different firm modeling and data analytics capabilities. Similarly, if user model quality is low, presenting a single strong default could be harmful, while offering a partitioned set instead allows consumers greater flexibility in selecting the best option. A deeper understanding of the decision process used by consumers with different choice architecture tools also seems important for understanding when to give recommendations and when firms could use consumer input to develop better user models. One interesting question is whether allowing consumers to tailor their choice architecture improves their choices. Many websites allow consumers to sort by an attribute, and firms usually impose a default ordering that impacts consumers’ decisions. Our results lead us to speculate that the resulting ordering, chosen by either users or firms, could lead to worse choices.

Conclusion

Rapid developments in digital technologies, artificial intelligence, and machine-learning-based algorithms may allow firms to suggest options that lead consumers and patients to more satisfying outcomes than they could determine themselves. Choice architecture, wisely applied, can be a relatively inexpensive and efficient way to use firm-level knowledge to improve social welfare. Better prediction of consumer choices and choice architecture could, in contrast, also be used to increase short-term profit. For example, a health insurer who observes consumers paying too much for a dominated policy might be pleased with the increased revenue. Regulatory oversight might be used to prevent choice architectures that lead to worse consumer health insurance choices but will depend on understanding how choice architecture ensembles affect consumers’ outcomes and choice processes. We hope that managers of health care firms will choose to use their knowledge to develop new business models that deliver longer-term value and minimize waste for customers, other stakeholders, and their firms.

Acknowledgments

The authors would like to thank Mieke van Os at Independer.nl for providing the data used in the pilot study of this research and for practical insights on the website’s structure, Nathaniel Posner for assistance in data collection, Martijn Willemsen for feedback and support in analyzing the MouselabWEB data, and Nuno Camacho and Kellen Mrkva for valuable feedback. The first author would like to thank Netspar for financial support of part of this research. The second and fourth authors would like to thank the Alfred P. Sloan Foundation for financial support of part of the research.

Special Issue Editor

Christine Moorman

Associate Editor

John Lynch

Declaration of Conflicting Interests

The author(s) declared the following potential conflicts of interest with respect to the research, authorship, and/or publication of this article: Benedict Dellaert formerly was a member of the independent board of supervisors of Independer, and Tom Baker and Eric Johnson were affiliated with Picwell.

Funding

The author(s) disclosed receipt of the following financial support for the research, authorship, and/or publication of this article: This work was supported by the Network for Studies on Pensions, Aging and Retirement and the Alfred P. Sloan Foundation.

ORCID iDs

Benedict G.C. Dellaert  <https://orcid.org/0000-0003-4637-1192>
Eric J. Johnson  <https://orcid.org/0000-0001-7797-8347>

References

- Abaluck, Jason and Jonathan Gruber (2011), “Choice Inconsistencies Among the Elderly: Evidence from Plan Choice in the Medicare Part D Program,” *American Economic Review*, 101 (4), 1180–1210.
- Alba, Joseph, John Lynch, Barton Weitz, Chris Janiszewski, Richard Lutz, Alan Sawyer et al. (1997), “Interactive Home Shopping: Consumer, Retailer, and Manufacturer Incentives to Participate in Electronic Marketplaces,” *Journal of Marketing*, 61 (3), 38–53.
- Atalay, A Selin, H. Onur Bodur, and Dina Rasolofoaarison (2012), “Shining in the Center: Central Gaze Cascade Effect on Product Choice,” *Journal of Consumer Research*, 39 (4), 848–66.
- Baker, Tom and Benedict G.C. Dellaert 2018, “Regulating Robo Advisors Across the Financial Services Industry: An Exploratory Essay,” *Iowa Law Review*, 103 (2), 713–50.
- Bar-Hillel, Maya (2015), “Position Effects in Choice from Simultaneous Displays: A Conundrum Solved,” *Perspectives on Psychological Science*, 10 (4), 419–33.
- Benartzi, Shlomo, John Beshears, Katherine L. Milkman, Cass R. Sunstein, Richard H. Thaler, Maya Shankar et al. (2017), “Should Governments Invest More in Nudging?” *Psychological Science*, 28 (8), 1041–55.
- Bhargava, Saurabh, George Loewenstein, and Justin Sydnor (2017), “Choose to Lose: Health Plan Choices from a Menu with Dominated Options,” *Quarterly Journal of Economics*, 132 (3), 1319–72.
- Brown, Christina L. and Aradhna Krishna (2004), “The Skeptical Shopper: A Metacognitive Account for the Effects of Default Options on Choice,” *Journal of Consumer Research*, 31 (3), 529–39.
- Cadario, Romain and Pierre Chandon (2020), “Which Healthy Eating Nudges Work Best? A Meta-Analysis of Field Experiments,” *Marketing Science*, 39 (3), 465–86.

- Caplin, Andrew and Mark Dean (2015), "Revealed Preference, Rational Inattention, and Costly Information Acquisition," *American Economic Review*, 105 (7), 2183–203.
- Cheema, Amar and Dilip Soman (2008), "The Effect of Partitions on Controlling Consumption," *Journal of Marketing Research*, 45 (6), 665–75.
- Chintagunta, Pradeep, Dominique M. Hanssens, and John R. Hauser (2016), "Marketing Science and Big Data," *Marketing Science*, 35 (3), 341–42.
- Costa-Gomes, Miguel A. and Vincent P. Crawford (2006), "Cognition and Behavior in Two-Person Guessing Games: An Experimental Study," *American Economic Review*, 96 (5), 1737–68.
- Dellaert, Benedict G.C. and Gerald Häubl (2012), "Searching in Choice Mode: Consumer Decision Processes in Product Search with Recommendations," *Journal of Marketing Research*, 49 (2), 277–88.
- De Los Santos, Babur, Ali Hortaçsu, and Matthijs R. Wildenbeest (2012), "Testing Models of Consumer Search Using Data on Web Browsing and Purchasing Behavior," *American Economic Review*, 102 (6), 2955–80.
- Diehl, Kristin (2005), "When Two Rights Make a Wrong: Searching Too Much in Ordered Environments," *Journal of Marketing Research*, 42 (3), 313–22.
- Diehl, Kristin, Laura J. Kornish, and John G. Lynch Jr. (2003), "Smart Agents: When Lower Search Costs for Quality Information Increase Price Sensitivity," *Journal of Consumer Research*, 30 (1), 56–71.
- Donkers, Bas, Benedict G.C. Dellaert, Rory M. Waisman, and Gerald Häubl (2020), "Preference Dynamics in Sequential Consumer Choice with Defaults," *Journal of Marketing Research*, 57 (6), 1096–112.
- Dorn, Michael, Claude Messner, and Michaela Wänke (2016), "Partitioning the Choice Task Makes Starbucks Coffee Taste Better," *Journal of Marketing Behavior*, 1 (3/4), 363–84.
- Ericson, Keith Marzilli and Amanda Starc (2012), "Heuristics and Heterogeneity in Health Insurance Exchanges: Evidence from the Massachusetts Connector," *American Economic Review*, 102 (3), 493–97.
- Gabaix, Xavier, David Laibson, Guillermo Moloche, and Stephen Weinberg (2006), "Costly Information Acquisition: Experimental Analysis of a Boundedly Rational Model," *American Economic Review*, 96 (4), 1043–68.
- Ghose, Anindya, Panagiotis G. Ipeirotis, and Beibei Li (2012), "Designing Ranking Systems for Hotels on Travel Search Engines by Mining User-Generated and Crowdsourced Content," *Marketing Science*, 31 (3), 493–520.
- Goldstein, Daniel G., Eric J. Johnson, Andreas Herrmann, and Mark Heitmann (2008), "Nudge Your Customers Toward Better Choices," *Harvard Business Review*, 86 (12), 99–105.
- Gotberg, Brook E. and Michael D. Sousa (2019), "Moving Beyond Medical Debt," *American Bankruptcy Institute Law Review*, 27 (2), 93–137.
- Harrison, Glenn W. and John A. List (2004), "Field Experiments," *Journal of Economic Literature*, 42 (4), 1009–55.
- Häubl, Gerald, Benedict G.C. Dellaert, and Bas Donkers (2010), "Tunnel Vision: Local Behavioral Influences on Consumer Decisions in Product Search," *Marketing Science*, 29 (3), 438–55.
- Häubl, Gerald and Valerie Trifts (2000), "Consumer Decision Making in Online Shopping Environments: The Effects of Interactive Decision Aids," *Marketing Science*, 19 (1), 4–21.
- Hayes, Andrew F. (2013), *Introduction to Mediation, Moderation, and Conditional Process Analysis: A Regression-Based Approach*. New York: Guilford Press.
- Hogarth, Robin M. and Hillel J. Einhorn (1992), "Order Effects in Belief Updating: The Belief-Adjustment Model," *Cognitive Psychology*, 24 (1), 1–55.
- Jachimowicz, Jon M., Shannon Duncan, Elke U. Weber, and Eric J. Johnson (2019), "When and Why Defaults Influence Decisions: A Meta-Analysis of Default Effects," *Behavioural Public Policy*, 3 (2), 159–86.
- Johnson, Eric J. (2021), *The Elements of Choice: Why the Way We Decide Matters*. New York: Riverhead Books.
- Johnson, Eric J. and Daniel Goldstein (2003), "Do Defaults Save Lives?" *Science*, 302 (2003), 1338–39.
- Johnson, Eric J., Ran Hassin, Tom Baker, Allison T. Bajger, and Galen Treuer (2013), "Can Consumers Make Affordable Care Affordable? The Value of Choice Architecture," *PLoS ONE*, 8 (12), e81521.
- Johnson, Eric J., Suzanne B. Shu, Benedict G.C. Dellaert, Craig Fox, Daniel G. Goldstein, Gerald Häubl, et al. (2012), "Beyond Nudges: Tools of a Choice Architecture," *Marketing Letters*, 23 (2), 487–504.
- Loewenstein, George, Joelle Y. Friedman, Barbara McGill, Sarah Ahmad, Suzanne Linck, Stacey Sinkula, et al. (2013), "Consumers' Misunderstanding of Health Insurance," *Journal of Health Economics*, 32 (5), 850–62.
- Lohse, Gerald L. and Eric J. Johnson (1996), "A Comparison of Two Process Tracing Methods for Choice Tasks," *Organizational Behavior and Human Decision Processes*, 68 (1), 28–43.
- Lynch, John G. (2009), "Information Remedies, Choice Architecture, and Plain Vanilla Financial Products," working paper, Russell Sage Foundation.
- Mantonakis, Antonia, Pauline Rodero, Isabelle Lesschaeve, and Reid Hastie (2009), "Order in Choice: Effects of Serial Position on Preferences," *Psychological Science*, 20 (11), 1309–12.
- McKenzie, Craig R.M., Lim M. Leong, and Shlomi Sher (2021), "Default Sensitivity in Attempts at Social Influence," *Psychonomic Bulletin & Review*, 28 (2), 695–702.
- Mrkva, Kellen, Nathaniel A. Posner, Crystal Reeck, and Eric J. Johnson (2021), "Do Nudges Reduce Disparities? Choice Architecture Compensates for Low Consumer Knowledge," *Journal of Marketing*, 85 (4), 67–84.
- Padilla, Nicolas and Eva Ascarza (2021), "Overcoming the Cold Start Problem of Customer Relationship Management Using a Probabilistic Machine Learning Approach," *Journal of Marketing Research*, 58 (5), 981–1006.
- Payne, John W. (1976), "Task Complexity and Contingent Processing in Decision Making: An Information Search and Protocol Analysis," *Organizational Behavior and Human Performance*, 16 (2), 366–87.
- Payne, John W., James R. Bettman, and Eric J. Johnson (1993), *The Adaptive Decision Maker*. Cambridge, UK: Cambridge University Press.
- Reisen, Nils, Ulrich Hoffrage, and Fred W. Mast (2008), "Identifying Decision Strategies in a Consumer Choice Situation," *Judgment and Decision Making*, 3 (8), 641–58.

- Smith, N. Craig, Daniel G. Goldstein, and Eric J. Johnson (2013), "Choice Without Awareness: Ethical and Policy Implications of Defaults," *Journal of Public Policy & Marketing*, 32 (2), 159–72.
- Sunstein, Cass R. (2018), "Better Off, as Judged by Themselves: A Comment on Evaluating Nudges," *International Review of Economics*, 65 (1), 1–8.
- Sunstein, Cass R. (2021), *Sludge: What Stops Us from Getting Things Done and What to Do About It*. Cambridge, MA: MIT Press.
- Szaszi, Barnabas, Anna Palinkas, Bence Palfi, Aba Szollosi, and Balazs Aczel (2018), "A Systematic Scoping Review of the Choice Architecture Movement: Toward Understanding When and Why Nudges Work," *Journal of Behavioral Decision Making*, 31 (3), 355–66.
- Thaler, Richard H. and Cass R. Sunstein (2008), *Nudge*. New Haven, CT: Yale University Press.
- Ursu, Raluca M. (2018), "The Power of Rankings: Quantifying the Effect of Rankings on Online Consumer Search and Purchase Decisions," *Marketing Science*, 37 (4), 530–52.
- Willemsen, Martijn C. and Eric J. Johnson (2011), "Visiting the Decision Factory: Observing Cognition with MouselabWEB and Other Information Acquisition Methods," in *A Handbook of Process Tracing Methods for Decision Making*, Michael Schulte-Mecklenbeck, Anton Kühberger, and Rob Ranyard, eds. New York: Taylor & Francis, 21–42.
- Wilson, Timothy D., Thalia Wheatley, Jonathan M. Meyers, Daniel T. Gilbert, and Danny Axsom (2000), "Focalism: A Source of Durability Bias in Affective Forecasting," *Journal of Personality and Social Psychology*, 78 (5), 821–36.
- Wong, Charlene A., Arthur T. Jones, Janet Weiner, Robert J. Town, and Tom Baker (2016), "For Third Enrollment Period, Marketplaces Expand Decision Support Tools to Assist Consumers," *Health Affairs*, 35 (4), 680–87.
- Woolhandler, Steffie and David U. Himmelstein (2017), "The Relationship of Health Insurance and Mortality: Is Lack of Insurance Deadly?" *Annals of Internal Medicine*, 167 (6), 424–31.
- Zhao, Xinshu, John G. Lynch, and Qimei Chen (2010), "Reconsidering Baron and Kenny: Myths and Truths About Mediation Analysis," *Journal of Consumer Research*, 37 (2), 197–206.
- Zlatev, Julian J., David P. Daniels, Hajin Kim, and Margaret A. Neale (2017), "Default Neglect in Attempts at Social Influence," *Proceedings of the National Academy of Sciences*, 114 (52), 13643–48.