No Such Thing as Too Much Chocolate: Evidence Against Choice Overload in E-Commerce

Carol Moser[1], **Chanda Phelan**[1], **Paul Resnick**[1], **Sarita Y. Schoenebeck**[1], **and Katharina Reinecke**[2]

[1]University of Michigan
Ann Arbor, MI, US
{moserc, cdphelan, presnick, yardi}@umich.edu

[2]University of Washington
Seattle, WA, US
reinecke@cs.washington.edu

ABSTRACT
E-commerce designers must decide how many products to display at one time. Choice overload research has demonstrated the surprising finding that more choice is not necessarily better—selecting from larger choice sets can be more cognitively demanding and can result in lower levels of choice satisfaction. This research tests the choice overload effect in an e-commerce context and explores how the choice overload effect is influenced by an individual's tendency to maximize or satisfice decisions. We conducted an online experiment with 611 participants randomly assigned to select a gourmet chocolate bar from either 12, 24, 40, 50, 60, or 72 different options. Consistent with prior work, we find that maximizers are less satisfied with their product choice than satisficers. However, using Bayesian analysis, we find that it's unlikely that choice set size affects choice satisfaction by much, if at all. We discuss why the decision-making process may be different in e-commerce contexts than the physical settings used in previous choice overload experiments.

Author Keywords
Choice overload; e-commerce; experiment; maximizing

ACM Classification Keywords
H.5.m. Information interfaces and presentation (e.g., HCI): Miscellaneous;

INTRODUCTION
E-commerce is a large and still-growing industry in the United States. Business-to-consumer online spending was predicted to increase from \$1.3 trillion in 2014 to \$2.1 trillion in 2018 [25]. Designers who craft e-commerce experiences for users are faced with a simple but critical design decision: how many products to display at once [23]. On one hand, some research shows that greater choice is better: greater choice can increase the chance of preference matching, enhance decision-making certainty, and enhance perceptions of freedom of choice and control (see [5] for a review). On the other hand, choice overload research has shown that selecting from large choice sets can be more cognitively demanding and feel more overwhelming [40], can lead to less satisfaction with the final choice [11], and can deter actual purchases [21]. Further, how users respond to different choice set sizes may depend on an individual's tendency to maximize or satisfice decisions. Maximizers, who search for the best option (in contrast to satisficers, who opt for the first satisfactory option), tend to be less satisfied with their choices [22], even when they expend time and effort to select from larger choice sets [10].

This research tests the choice overload effect in an e-commerce context, a domain that has received little attention in the choice overload literature. We conducted an online experiment with 611 participants who were recruited to complete an online survey for a chance to win gourmet chocolate. After completing the survey, participants were redirected to an online chocolate shop to select their desired chocolate bar; participants were randomly assigned to select from either 12, 24, 40, 50, 60, or 72 different options. We find that maximizers are less satisfied with their product choice than satisficers, a finding that is consistent with prior literature. However, we find that in an e-commerce context, choice set size does *not* affect choice satisfaction—at least not more than a negligible amount—regardless of a participant's maximizing or satisficing tendencies.

We discuss why the decision-making process may be different in an e-commerce context and present implications for designers and directions for future work.

RELATED WORK AND HYPOTHESES
Choice overload is grounded in the theory of bounded rationality, the argument that humans do not possess the cognitive processing ability to carefully consider all possible options and their potential consequences [44]. Schwartz characterized choice overload as the "paradox of choice": the unexpected finding that more choice is not necessarily better [40]. More specifically, choice overload has been described as a mental state of the decision maker—in response to a large, cognitively demanding choice set—that manifests as either dissatisfaction, frustration, or regret [7].

CHI 2017, May 06 - 11, 2017, Denver, CO, USA

ACM 978-1-4503-4655-9/17/05⋯\$15.00
DOI: http://dx.doi.org/10.1145/3025453.3025778

Iyengar and Lepper [21] demonstrated that when selecting from large versus small choice sets, consumers may find the decision-making process more difficult with large sets and forego making a purchase. Through a series of field and in-lab experiments, the authors found that a greater percentage of shoppers (30% versus 3%) purchased jam in a grocery store when presented with 6 versus 24 jam options and that a significantly higher percentage of students (74% versus 60%) completed an extra credit essay assignment when offered 6 versus 30 topic options. Further, they demonstrated that even when consumers commit to a final choice, they are less satisfied with that final choice when selecting from large choice sets: subjects who chose 6 gourmet chocolate (versus 30) reported more satisfaction with their choice than those selecting from the larger choice set. The choice overload effect has been widely studied and demonstrated with a variety of different product types including food products, electronics, office supplies, magazines, mutual funds, and music (see [7,39]).

Choice overload as an HCI problem

HCI research in e-commerce has investigated how to visually present product choices [16], how photography can affect e-commerce website credibility [36,37,46], and how e-commerce web aesthetics can influence online purchase behavior [4]. Other work has argued that the online customer experience encompasses not only usability and UI considerations, but also "the overall experience and satisfaction a customer has when purchasing or using a product or service" [31]. However, despite evidence that choice set size can affect both the chooser's experience and resulting choice satisfaction [11, 40], little prior work has studied choice overload as an e-commerce design problem.

In other online contexts, prior work has shown that participants using search engines to complete fact finding tasks feel more satisfied and more confident in their answers when selecting from a short (6 results) versus a long list of search engine results (24 results) [33], especially when selecting under time constraints [8]. In an online dating pool, users reported lower choice satisfaction when presented with a large pool of potential partners versus a small one (24 versus 6 matches) [9].

While the research reviewed above demonstrates negative consequences of large choice set sizes, other research suggests that greater choice is better for the user experience, as it can increase the chance of preference matching [23], enhance decision-making certainty [3], and enhance perceptions of freedom of choice and control [5]. Recent work has also shown that online shoppers spend more than catalog-only shoppers, despite larger choice set sizes online [28]. Given the conflicting research, web designers are left with little concrete guidance beyond advice to provide the "fewest product options that still allow for meaningful product comparison" [23]. The current study empirically tests whether providing fewer product options is indeed sound design practice.

Maximizing and satisficing decisions

Choice overload theory has not gone unchallenged: meta-analyses have both confirmed [7] and called into question [39] the negative effect of too many options. These conflicting results suggest perhaps that choice set size does not have a universal effect but can influence certain people in certain contexts. Several studies have found that an individual's tendency to maximize or satisfice is especially relevant to the decision-making process [10,40,41]. A maximizer is someone who typically tries to consider all options and attempts to select the best possible choice; a satisficer is someone who tends to select an option that is satisfactory, but not necessarily the best [40,42]. The nascent literature has shown that maximizers tend to score lower in overall life satisfaction, higher in perfectionism and regret [42], and lower in outcome satisfaction despite securing objectively better outcomes than satisficers [22]. In terms of purchasing behavior, maximizers report a greater tendency to conduct product comparisons, heightened post-decision consumer regret, and diminished positive feelings toward purchases [41]. These findings inform the current study's first hypothesis.

Hypothesis 1: In an e-commerce context, maximizers are less satisfied with their product choices than satisficers.

The choice overload effect has been largely demonstrated with experimental designs comparing two choice set sizes, low versus high [7,39]. However, there is some evidence to suggest that as the number of options increases, satisfaction initially increases and then declines. Reutskaja and Hogarth [35] measured choice satisfaction for participants selecting a gift box from a set of either 5, 10, 15, or 30 alternatives. The authors found an inverted U-shaped relationship with the highest levels of choice satisfaction occurring at intermediate levels of choice (10 and 15 options) versus the smallest choice set (5 options) or largest choice set (30 options). Similarly, Shah and Wolford found that the greatest proportion of participants who purchased a pen were those presented with 8, 10, 12, or 14 options versus a small choice set size (2, 4, or 6 options) or large choice set size (16, 18, or 20 options) [43]. These findings inform this study's second hypothesis.

Hypothesis 2: In an e-commerce context, choice satisfaction declines (or increases and then declines to form an inverted U-shape) as the number of product options increases.

Finally, limited prior work has investigated the relationship between maximizer/satisficer tendencies and choice overload. One study that tested the choice overload effect for maximizers versus satisficers (n=25) did not include any high scorers on the maximizer scale [33]. Work that asked participants to select and listen to a classical music album for two minutes failed to find an effect of maximizing tendencies on satisfaction but this choice scenario may not have evoked maximizing and satisficing behavior [38].

In contrast, other research demonstrated that maximizers place a greater dollar value on their choice of soda when selecting from a set of 6 options versus a set of 24 options [1]. Although this work does not address choice satisfaction directly, it does suggest that maximizers react differently (i.e., negatively) when tasked with handling larger choice sets. Maximizers have also been shown to be more likely than satisficers to sacrifice resources (i.e., time and/or effort) to attain larger choice sets but ultimately report less satisfaction with their choice in comparison to satisficers [10]. Taken together, the literature reviewed here suggests maximizers not only tend to experience less choice satisfaction (H1) but that they respond more negatively to large choice sets than satisficers. These findings inform the current research's third hypothesis.

Hypothesis 3: In an e-commerce context, as the number of product options increases, the rate of decline in choice satisfaction is steeper for maximizers.

EXPERIMENT

We designed this experiment with two main goals in mind: first, to test whether the number of product choices in an e-commerce shop affects choice satisfaction, and second, to test whether this finding depends on a person's decision-making style. Our online experiment followed a between-subjects design with participants randomly assigned to one of six choice set sizes.

Method

Experimental conditions

There were six conditions for the number of products displayed on the simulated e-commerce webpage (choice set sizes: 12, 24, 40, 50, 60, 72). Choice set sizes were derived from an examination of thirteen top-performing e-commerce sites, as ranked by multiple sources [45,47]. The lead author visited each website on the same day, using the same browser and same size window (1024x768 pixels). An incognito browser session was also used to ensure any display settings were cleared. The lead author drilled down to several different product types per website to confirm each site's default maximum number of products displayed per page. The four most commonly occurring set sizes were 24, 40, 50, and 60; the smallest and largest set sizes were 12 and 96 respectively (see supplementary materials for additional details). Due to the limited number of chocolate flavors available on the market at the time, this study was only able to test a maximum of 72 instead of 96 products.

Materials

An experimental chocolate e-commerce shop was designed to simulate an e-commerce shopping website with a home page, product description pages, a shopping cart, and a checkout procedure (see Figure 1). The three-column website design was informed by an examination of 13 top-performing e-commerce sites [45,47] where the average default number of product columns displayed was approximately three columns (M=2.8, *median*=3.0).

Chocolate was used as the experimental product in order to remain consistent with the design of several prior choice overload experiments, including Iyengar and Lepper's [21] influential study. In addition, previous research has demonstrated that the number of product attributes between choices can influence the consumer decision-making process, especially when leading to information overload [23]. Utilizing a simple product such as chocolate bars limited the number of product attributes that participants needed to consider during product comparisons. Chocolate used in this study differed in flavor and color (e.g. dark chocolate and milk chocolate), with small variations in size. Chocolate bars did not vary in other attributes such as price, brand, or packaging.

Procedure

Participants completed the study online at a location of their choice. Participants were invited to participate in the study through online recruitment ads on Facebook, Craigslist, Google AdWords, and Microsoft Bing. The ads indicated that, in return for participating, 1 in 10 participants would receive their choice of chocolate shipped to them for free. This lottery design, similar to those used in prior choice overload research [17], was implemented to create an authentic, as opposed to hypothetical, choice scenario for participants, while keeping costs down. After completing an informed consent form, subjects were directed to the maximizer survey questions and a short demographics survey. Next, participants were directed to the chocolate e-commerce shop to select the chocolate bar they wished to receive.

When directing participants to the chocolate shop, the system randomly assigned subjects to one of the six choice set size conditions and displayed the appropriate number of chocolate flavor options on the participant's web page. Consistent with a real-world e-commerce experience, participants made their chocolate selection, proceeded to the shopping cart, and completed the checkout procedure by entering their shipping address and answering "customer satisfaction" questions including the main choice satisfaction measure (described below) and a measure of how frequently they shop online (on a 6-point scale from *never* to *every day*). Finally, participants were entered into the online lottery and informed as to whether they were winners or not. The maximizer scale and demographics survey were intentionally administered before the choice exercise to remain consistent with the procedure implemented in prior work [1,10] and to enhance believability that the chocolate choice exercise was conducted simply to compensate for participation in the survey. This study was approved by the research team's Institutional Review Board. We reminded participants twice that we were collecting their addresses for shipping purposes only.

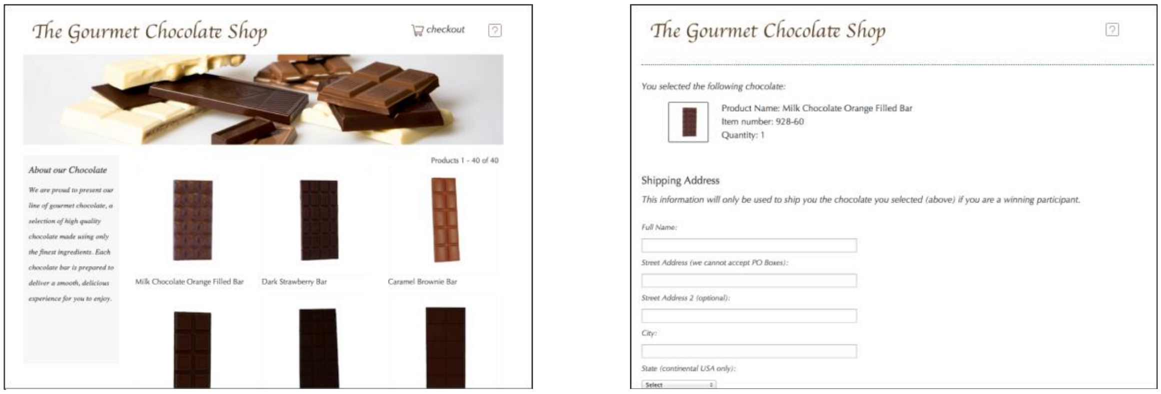

Figure 1. Simulated e-commerce website homepage (left) and checkout page (right).

Measures

Maximizing tendency was measured using a 13-item maximizer scale [41]. Questions included items such as, "Whenever I'm faced with a choice, I try to imagine what all the other possibilities are, even ones that aren't present at that moment." Responses were provided on a 7-point scale from 1 (*completely disagree*) to 7 (*completely agree*). Following [40], responses were summed to generate a maximizer score (*min* = 13, *max* = 91), with higher scores indicating a tendency to maximize and lower scores indicating a tendency to satisfice.

Choice satisfaction was measured using a survey question modeled after those used in previous work [21, 35]. After confirming their chocolate selection, subjects were asked, "How satisfied are you with the chocolate you decided to pick?" with responses provided on a 7-point Likert scale, ranging from 1 (*not at all*) to 7 (*extremely)*.

Participants

A total of 621 participants completed the study, which was available online for a total of seven weeks. Ten participants (1.6%) were removed from analysis due to data quality concerns; these participants spent less than 2.5 seconds per survey question and provided the same answer for all questions. (*Note:* These outliers had a misleading, disproportionate effect when analyzed with the full sample. See supplementary material for results that include the excluded data points.) The remaining 611 participants (164 men, 442 women, and 5 other) were on average 37.61 years old (*SD* = 15.31, *range* 18-86,) and all reported living in the continental United States. Participants were frequent online shoppers: 61% reported making an online purchase at least a few times per month. The average maximizer score was 57.23 (*SD* = 11.54). Participants considered the options presented to them, indicated by 66.61% selecting chocolates located below the first two rows of products ("below the fold").

Note: A second sample of participants (*N* = 113) came from an undergraduate course from a large university who participated in exchange for extra course credit. Because this study aimed to simulate an online shopping experience, we choose to only include participants recruited through online advertising. The results including the student participants are reported in the supplementary material and are consistent with the results reported here.

Analysis

In our analysis, we estimated several alternative linear models predicting choice satisfaction from independent variables choice set size and maximizer score (see Table 1 for reference). After estimating the models, we applied information criteria to assess which of the alternative functional forms was the best fit for the data.

Variable	*Description*	*Measurement*
Choice set size	Number of product options (chocolate bars) presented	12, 24, 40, 50, 60, or 72 choices
Maximizer score	Tendency to look for better options beyond satisfaction	13 items, 7-pt Likert scale
Choice satisfaction	Satisfaction with chocolate choice	Single-item, 7-pt Likert scale

Table 1: Variables included in analyses.

All models assume that choice satisfaction was conditionally normally distributed. That is, choice satisfaction y_i is predicted to be a value μ_i that depends on the choice set size and maximizer score, plus some normally distributed "error term" with mean 0 and standard deviation σ. The way that μ_i depends on the other variables is what differs between models. In summary:

$$y_i \sim \text{Normal}(\mu_i, \sigma)$$

For each model, we conducted both frequentist and Bayesian analysis. The frequentist analysis will be familiar for many readers: ordinary least squares (OLS) regression and null hypothesis significance testing (NHST).

The Bayesian analysis may be less familiar, but has two major advantages. First, Bayesian analysis overcomes some weakness in NHST (e.g. [12,14,26,30]). In particular, NHST cannot support the null hypothesis; it can only fail to reject it [14]. In contrast, Bayesian analysis yields a probability distribution for the parameters. If a particular parameter is close to the null hypothetical value with high probability, that provides evidence *in favor* of the null hypothesis [19].

Second, interpretation of Bayesian results is more intuitive than the results produced by frequentist methods [13,24]. For example, a Bayesian analysis produces a *credible interval*. Frequentist confidence intervals are often misinterpreted, even by experienced researchers [18]. A 95 percent confidence interval is often interpreted as meaning that there is a 95 percent probability that the true parameter value lies within the interval, which is incorrect. Conditional on the model's priors, Bayesian credible intervals *are* appropriately interpreted that way.

A Bayesian model requires three components: parameters, a likelihood function, and priors for each of the parameters [30]. The parameters, as in OLS, are the variables in a model. The model determines a likelihood function, which computes the probability of any potential observation, given any set of values for the parameters. The priors are the initial sets of plausibilities (probability distributions) for each of the model's parameters. Given these three elements and a set of observed data points, an estimation process yields a posterior distribution of plausible values for the parameters [30]. We estimated posterior beliefs in R using the R package *rethinking* [29].

Setting priors carefully is an important part of Bayesian analysis [12,30]; they can constrain a parameter to reasonable ranges, telling the estimation process to be skeptical of implausible parameter values even if they fit the observed data. We based the priors for some parameters on the results of a previous experiment [21]. When previous literature did not directly inform the priors for a parameter, we set weakly informative priors that gave very low probabilities to very implausible parameter values (e.g., those that would lead to changes in satisfaction of more than 3.0 on a 7-point scale) but otherwise constrained the parameter values as little as possible.

The prior for the error term σ of the models was not based on previous literature. As is common in Bayesian analysis for model parameters that are not the main focus of interest, we set a very weakly informative prior for this error term: σ is drawn from a uniform distribution between 0 and 3. This prevents implausibly high values: for example, if σ were 4, 5% of the time the model would predict a satisfaction score more than eight away from the predicted mean, which is hardly plausible on scale of 1-7.

Model 1: Constant
The intercept-only constant model is the simplest. There is no effect of choice set size or maximizer score on satisfaction, so satisfaction scores are random draws from the same distribution for all experimental conditions:

$\mu_i = \alpha$
$\alpha \sim \text{Normal}(5.86, 0.60)$

The Bayesian priors for α are based on the results from a previous study in which participants also chose between different chocolate bars, and their choice satisfaction was also measured on a 7-point Likert scale [21]. The mean satisfaction scores were pooled across conditions in that study to create the mean 5.86 for our prior. The standard deviation 0.6 was set so that 5.86 ± 2 standard deviations includes the 95% confidence intervals of the means for both experimental conditions from the previous study.

To provide intuitions, ~~the figure~~Figure 2a shows μ_i is 5.86 regardless of choice set size. Clicking on the graph yields an animation showing other random draws from the prior probability distribution for α.[1] The animation visually clarifies the meaning of the prior: the most plausible value for mean satisfaction α is 5.86, but a mean of 6.2 or 5.2 would not be very surprising. The graph does not illustrate σ, the model's prediction of how much individual satisfaction scores will vary from the predicted mean.

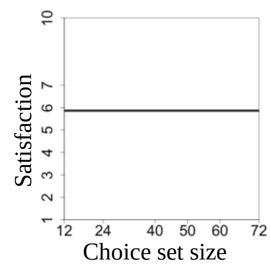

Fig. 2a, Model 1. Click to see animated draws illustrating the priors for α.

H1: Maximizers vs. Satisficers

Model 2: Linear in maximizer score
The second model predicts that mean satisfaction will vary linearly with the subject's maximizer score, but still is not affected by choice set size:

$\mu_i = \alpha + \beta_{max} * maxim\_score$
$\alpha \sim \text{Normal}(7.20, 2.0)$
$\beta_{max} \sim \text{Normal}(-0.026, 0.03)$

In the absence of guidance from prior studies, we centered the prior on a value that would yield about a two-point average difference in satisfaction between complete maximizers and complete satisficers (see Figure 2b). Because we have great uncertainty about the actual effect, we chose a standard deviation that assigns a probability of approximately 20% to parameter values > 0. For different values of β_{max}, α will need to be correspondingly lower or higher in order to maintain a mean value for μ_i centered at 5.86. Thus, we adjust the mean of α and increase its standard deviation.

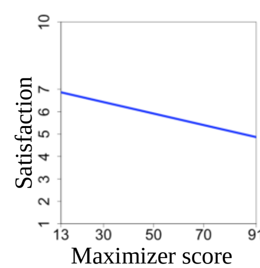

Fig. 2b, Model 2. Click to see animated priors.

[1] Animations of priors are also available at https://chocolate-animations.appspot.com/appendix/.

H2: Effect of Choice Set Size

We estimated a series of models to analyze the effect of choice set size. The constant model (Model 1) assumes mean satisfaction was the same in all conditions. A linear model allows mean satisfaction to increase or decrease linearly with the number of items. A quadratic model allows for a U-shape. Finally, a fixed-effects model with dummy variables for the different set sizes provides the most degrees of freedom. These models do not control for maximizer score.

Model 3: Linear in choice set size

$\mu_i = \alpha + \beta_{lin} * set\_size$
$\alpha \sim \text{Normal}(7.30, 2.0)$
$\beta_{lin} \sim \text{Normal}(-0.033, 0.01)$

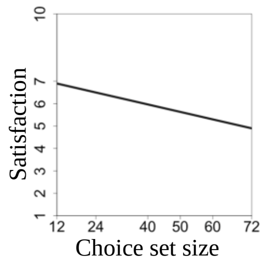

Fig. 2c, **Model 3. Click to see animated priors.**

The prior for β_{lin} says that a linear effect of -0.03 is the most plausible value for the coefficient (see Figure 2c). This comes from the previous study [21], where a difference in choice set size of 24 led to a decline in mean satisfaction of 0.82 (0.82/24 ≈ 0.033). If this were the true parameter value, an increase in set size from 12 to 72 would yield a decrease in satisfaction of 2.0 points. The prior's standard deviation is the standard error of the previous study's data (0.17/24 ≈ 0. 01). This is a strong prior in that it assigns a probability of just 0.13% to positive values. As with Model 2, the mean of the prior for α was adjusted so that the overall mean for μ_i would still be 5.86, with a high standard deviation.

Model 4: Quadratic

The quadratic model adds an additional degree of freedom, allowing for a parabolic or inverted parabolic relationship between choice set size and satisfaction, as H2 predicts:

$\mu_i = \alpha + \beta_{lin} * set\_size + \beta_{quad} * set\_size^2$
$\alpha \sim \text{Normal}(3.90, 2)$
$\beta_{lin} \sim \text{Normal}(0.15, 0.3)$
$\beta_{quad} \sim \text{Normal}(-0.002, 0.004)$

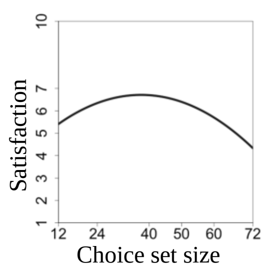

Fig. 2d, **Model 4. Click to see animated priors.**

There was no previous literature to draw from for priors for this model. Because the theory predicts an inverted-U shape, we centered the priors on a positive value for β_{lin} and a negative value for β_{quad} (see Figure 2d). The exact values were chosen so that predicted satisfaction peaked at a set size of 40, with satisfaction 1.3 less at 12 items and 2.4 less at 72 items. The standard deviations were set high relative to the means of the priors, to reflect great uncertainty about these values. The mean of the prior for α was adjusted so that the overall mean for μ_i would still be 5.86.

Model 5: Fixed effects

A fixed-effects model adds the most degrees of freedom; it has a series of dummy variables corresponding to choice set sizes other than 12:

$\mu_i = \alpha + \beta_{24}x_{24} + \beta_{40}x_{40} + \beta_{50}x_{50} + \beta_{60}x_{60} + \beta_{72}x_{72}$
$\alpha \sim \text{Normal}(5.86,$
$0.60)\beta_{24} \sim \text{Normal}(0, 1.5)$
$\beta_{40} \sim \text{Normal}(0, 1.5)$
$\beta_{50} \sim \text{Normal}(0, 1.5)$
$\beta_{60} \sim \text{Normal}(0, 1.5)$
$\beta_{72} \sim \text{Normal}(0, 1.5)$

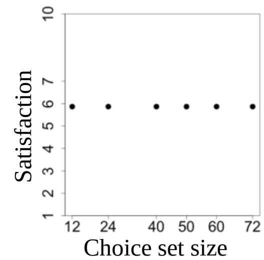

Fig. 2e, **Model 5. Click to see animated priors.**

Again in the absence of guidance from previous literature, priors for these parameters were chosen to be very weakly informative (see Figure 2e). They are centered around 0 but with high standard deviations: about 5% of the probability is assigned to values that yield increases or decreases in mean satisfaction of more than 3.0, as compared to the experiment condition with 12 items to choose from .

H3: Combined Effects

Two models included both choice set size and maximizer score.

Model 6: No Interaction Term

$\mu_i = \alpha +$
$\quad \beta_{lin} * set\_size +$
$\quad \beta_{max} * maxim\_score$

$\alpha \sim \text{Normal}(8.63, 2.0)$
$\beta_{lin} \sim \text{Normal}(-0.033, 0.01)$
$\beta_{max} \sim \text{Normal}(-0.026, 0.03)$

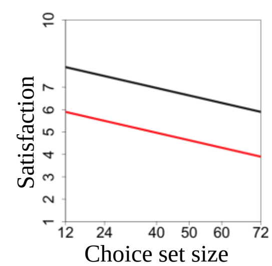

Fig. 2f, **Model 6. Click to see animated priors. Black line (upper): extreme satisficers. Red line (lower): extreme maximizers.**

The priors for β_{lin} and β_{max} were the same as those used in previous models (see Figure 2f). The mean of the prior for α was adjusted so that the mean value for μ_i would still be 5.86.

Model 7: With Interaction Term

$\mu_i = \alpha +$
$\quad \beta_{lin} * set\_size +$
$\quad \beta_{max} * maxim\_score +$
$\quad \beta_{int} * set\_size * maxim\_score$
$\alpha \sim \text{Normal}(7.43, 2.0)$
$\beta_{lin} \sim \text{Normal}(0, 0.01)$
$\beta_{max} \sim \text{Normal}(0, 0.03)$
$\beta_{int} \sim \text{Normal}(-0.0007, 0.0007)$

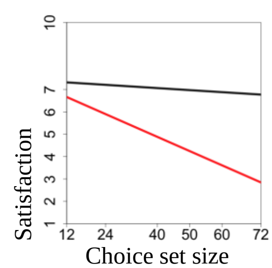

Fig. 2g, **Model 7. Click to see animated priors. Black line (upper): extreme satisficers. Red line (lower): extreme maximizers.**

Because H3 predicts a negative interaction term, we centered the prior on a negative value (see Figure 2g). In the absence of guidance from previous literature, we centered it at -0.0007 with a large enough standard deviation to treat any reasonable effect size as plausible. At -0.0007, the decline in satisfaction from 12 to 72 items is 3.27 greater for extreme maximizers than extreme satisficers. With the interaction term, we no longer have guidance from prior literature on the main effects, so the priors for β_{lin} and β_{max} were centered at 0, with the same standard deviations as in prior models. The mean of the prior for α was adjusted so that the mean value for μ_i would still be 5.86.

Model Comparison

For H2 and H3, there were multiple alternative models with different functional forms. With extra degrees of freedom (more parameters) it is always possible to improve predictions in-sample (i.e., reduce the residual errors, and thus the R^2 value in frequentist regressions). However, this can lead to overfitting, producing worse predictions out of sample (e.g., on unseen data from additional subjects).

Information criteria provide a principled way to choose among models with different parameters that predict the same outcome. We applied the Widely Applicable Information Criterion (WAIC), a measure of uncertainty for each independent observation that is calculated by taking averages of log-likelihoods over the posterior distribution [30]. The models were then compared using Akaike weights, which are rescaled WAIC scores. These weights estimate the probability that a given model will be best for predicting new out of sample data; higher weights are better [30]. WAIC and Akaike weights were also calculated using the *rethinking* R package.

Another common method of model comparison is the Bayes factor, which is the ratio of two models' average likelihoods [12]. However, there are issues with using the Bayes factor in model comparison; for example, priors can have a major effect even when they are weak [30]. We therefore prefer information criteria over the Bayes factor as a comparison measure.

RESULTS

We combine results from frequentist OLS regressions and Bayesian analyses in our result tables 2-4. We report the point estimate from OLS for each parameter. NHST significance tests for those parameters are summarized by asterisks next to the point estimates. Below the OLS point estimates are Bayesian 95% credible intervals, in square brackets. These are the smallest intervals that cover 95% of the weight of the posterior probability distribution. At the bottom of the table are the Akaike weights.

In the bottom row of ~~each table~~Tables 2-4 are graphs showing the modeled relationship between independent variables and satisfaction scores. These are analogous to Figures 2~~-3~~a-g, but based on the inferred posterior distributions for the parameters rather than the priors. Each static graph is a maximum a posteriori probability (MAP) plot, reflecting the most plausible parameter values.

	M1	**M2**
intercept	6.13*** [6.05, 6.22]	6.68*** [6.28, 7.09]
maximizer score	—	-0.01 [-0.02,-0.003]
Akaike weight	0.06	0.94

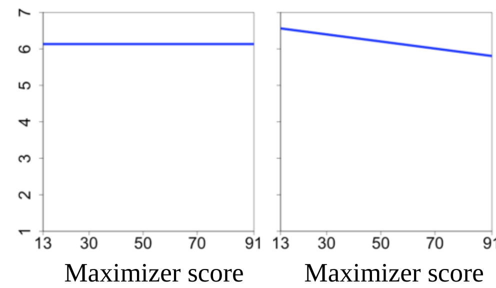

Table 2: Results for H1 (n=611). Click on any graph to visualize the distribution of plausible relationships
Note: *p < 0.05. **p < 0.01. ***p < 0.001.
95% credible intervals are presented in [square brackets].

Clicking on any graph yields an animation of alternative draws (a variant of Hypothetical Outcome Plots [20]), providing visual intuitions about the uncertainty of parameter estimates.

Details of all analyses, and the dataset itself, can be found in the supplementary materials.

H1: Maximizers vs. Satisficers

Our results support H1: maximizers were overall less satisfied with their choices than satisficers. The linear regression predicting choice satisfaction by maximizer score reveals that for every one-point increase in maximizer score, participants are on average 0.01 points less satisfied with their choice (see Table 2, Model 2).

A comparison of Model 2 against the constant model (Model 1) assigns an Akaike weight of 0.94 to Model 2: Model 2 is very likely to have a higher predictive accuracy than the constant model on out of sample data (e.g., from another experiment).

H2: Effect of Choice Set Size

Our results do not support H2. Overall choice satisfaction was high (M = 6.13, SD = 1.03), but did not vary depending on the number of options.

In the linear regression model (Table 3, Model 3), the coefficient for choice set size was not significant. The Bayesian credible interval is nearly centered around 0 and tight enough to imply that the effect of choice set size is very unlikely to be large enough to be meaningful in an e-commerce setting.

In the quadratic regression model (Table 3, Model 4), none of the predictors were statistically significant. Further, the point estimates of the quadratic model show an effect opposite of expectations: it shows a very slightly parabolic shape, instead of the inverted parabola predicted in H2.

	M1	M3	M4	M5
intercept	6.13*** [6.05, 6.22]	6.17*** [6.04, 6.40]	6.18*** [5.67, 6.34]	6.19*** [6.00, 6.37]
choice set size	—	-8×10^{-4} [-0.006,+0.001]	-0.002 [-0.01,+0.03]	—
choice set size^2	—	—	1×10^{-5} [-3×10^{-4}, $+1\times10^{-4}$]	—
choice set: 24	—	—	—	-0.18 [-0.47, +0.10]
choice set: 40	—	—	—	0.12 [-0.14, +0.40]
choice set: 50	—	—	—	-0.10 [-0.36, +0.16]
choice set: 60	—	—	—	-0.17 [-0.44, +0.24]
choice set: 72	—	—	—	-0.02 [-0.28, +0.24]
Akaike weight	0.54	0.22	0.08	0.16

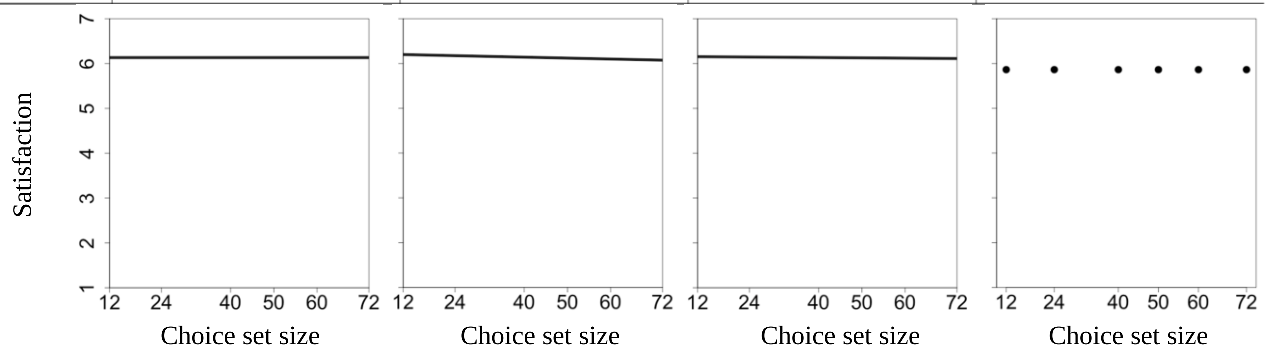

Table 3: Results for H2 (M1, M3-M5 (n=611)). Click on any graph to visualize the distribution of plausible relationships given the evidence from the experiment. Note: *p < 0.05. **p < 0.01. ***p < 0.001. 95% credible intervals are presented in [square brackets].

Last was the fixed-effects model, a regression model with dummy variables representing each choice set size (Table 3, Model 5). None of the choice set size variables showed a mean satisfaction significantly different from that for 12 choices. Moreover, there is no clear trend in mean satisfaction as choice set sizes increase.

Models 3-5 are compared to the constant model (Model 1) in Table 3, in which satisfaction scores are drawn randomly from the same distribution, regardless of choice set size. Model 1 has the highest Akaike weight, 0.59. The other models have non-zero Akaike weights, so we cannot reject these models entirely; however, they are much less plausible than the constant model. This is evidence in favor of the null hypothesis: there is no effect of choice set size on satisfaction that is large enough to be meaningful.

H3: Interaction of Set Size and Maximizer Score

H3 stated that as the number of product options increases, the rate of decline in choice satisfaction is steeper for maximizers. Our results do not support this hypothesis.

In Model 6, only the main effect of maximizer score is significant (Table 4), consistent with Model 2. Adding an interaction term for maximizer scores and choice set size in Model 7 did not reveal a statistically significant predictor of choice satisfaction (Table 4). Clicking on the graph to see the animation yields pictures showing that both positive and negative interaction terms are plausible, though negative interaction terms are somewhat more frequent.

Comparing Models 6 and 7 to Model 1, Model 6 claims the largest Akaike weight of 0.62. As in H2, the other models have non-zero weights so cannot be dismissed entirely, but are much less plausible than Model 6.

DISCUSSION

The goal of this research was to test the relationship between choice set size, maximizing/satisficing tendencies, and choice satisfaction in an e-commerce context. We find that:

- Maximizers are overall less satisfied with their choices than satisficers, but
- the number of product options online does not have much, if any, impact on choice satisfaction and accordingly,
- maximizers' choice satisfaction does not decline at a steeper rate than for satisficers.

Prior research in choice overload theory suggests that as choice set size increases, satisfaction either decreases or increases and then declines in a downward curvilinear

	M1	M6	M7
intercept	6.13*** [6.05, 6.22]	6.72*** [6.35, 7.23]	6.64*** [5.86, 7.26]
choice set size	—	-9×10^{-4} [-0.006, +0.002]	0.001 [-0.01, +0.01]
maximizer score	—	-0.01** [-0.02, -0.003]	-0.008 [-0.02, +0.005]
maximizer score * choice set size	—	—	-3×10^{-5} [-3×10^{-4}, $+2\times10^{-4}$]
Akaike weight	0.11	0.53	0.36

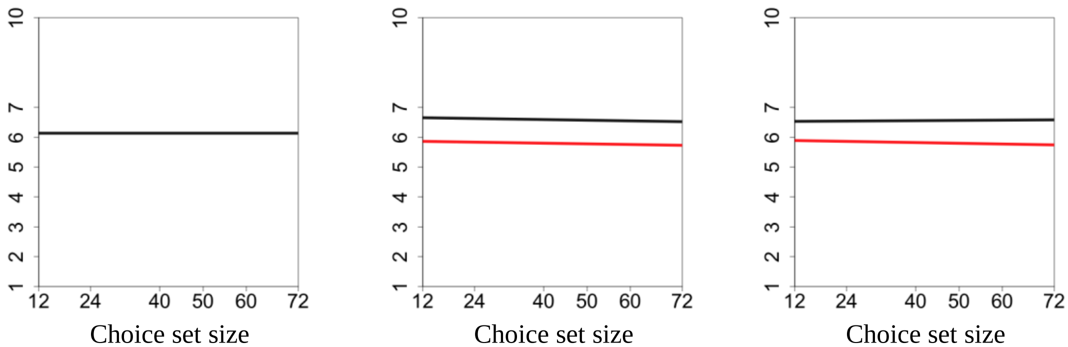

Table 4: Results for H3 (M1, M6-7 (n=611)).
Black lines in M6-M7 represent predictions at the lowest maximizer score of 13 and red lines represent predictions at the highest maximizer score of 91. Click on any graph to visualize the distribution of plausible relationships given the evidence from the experiment. Note: ****p* **< 0.05. *****p*** **< 0.01. ******p*** **< 0.001. 95% credible intervals are presented in [square brackets].**

fashion [21,35,43]. Even when these findings are taken into account, via the priors of the Bayesian analysis, the results of this study not only fail to provide support to extend these prior findings to an e-commerce context, but provide evidence against a choice overload effect in e-commerce contexts. Similarly, prior work suggests that an individual's tendency to maximize or satisfice choices can moderate the effect of choice set size on choice satisfaction. While the current study confirms the general tendency of maximizers to be less satisfied than satisficers, the current study does not support the hypothesis that maximizing tendencies moderate choice satisfaction in the e-commerce context. These surprising results suggest that the decision-making process may be different in an e-commerce context versus in face-to-face contexts—choice overload research has largely been studied in retail spaces or in-lab [7,39].

Large choice sets are normal and expected online

One potential explanation is that shoppers are accustomed to large choice sets online. For years, e-retailers such as Amazon.com, BestBuy.com, and Walmart.com have offered online shoppers thousands, if not millions, of products to select from. Moreover, online shoppers may be more accustomed to processing large choice sets that are presented *all at one time*. Our preliminary study of top-performing e-commerce sites (see Methods section) revealed that sites such as Target.com and eBay.com offer consumers the option to view as many as 100 or 200 products per page. The majority of participants in this study were frequent online shoppers (61% of participants reported making online purchases at least a few times per month in the last twelve months). Thus, it is likely that participants were already comfortable using sites that present large choice sets all at one once—choice sets even larger than presented in this research.

There is even evidence to suggest that shoppers *expect* large choice sets online; industry research shows that one reason consumers shop online versus offline is to gain access to a large variety of product offerings [34]. Recent work has also shown that online shoppers spend more than catalog-only shoppers, despite larger choice sets online [28]. This may be due to lower "search costs" online, where the marginal cost of searching for and considering each additional product is less online than through catalogs. Search costs may also help explain why choice overload may manifest in-store but not online. Considering each additional product in-store can involve physical inspection, locating product information, and product comparison. The online context reduces this workload by providing a systematic and predictable display of product information.

Large choice sets mitigate risks online

Consumers perceive a number of risks when shopping online. Consumers shopping online are concerned about not being able to physically inspect the product before purchasing, not being able to successfully return unwanted products, the potential difficulties in reaching customer service, and generally feeling disappointed with their choice when it arrives [15,27,32]. Because of these perceived risks, consumers are more likely to purchase high risk products offline, in traditional retail venues, than online [27]. One method for mitigating these perceived risks is to provide large and varied product assortments. Prior work has demonstrated that consumers prefer to select from large-variety assortments when the purchase is perceived as more risky. In a study where participants were asked to indicate which assortment of chocolates they would prefer to choose from (high variety or low variety), participants were more likely to select the low-variety assortment when they were told that the store allows product returns for any reason (low risk condition) [3]. Participants were more likely to select the high-variety assortment when they were told the store only allows returns with a receipt and with the product unopened (high risk condition) [3].

If consumers perceive greater risk in shopping online than offline, it is likely that consumers would prefer to select from more varied choice sets online in order to mitigate those perceived risks. Because larger choice sets tends to be perceived as having more variety [6], a larger choice set size online may be perceived as more varied and therefore beneficial for mitigating the risks of online shopping. Therefore, while larger choice set sizes are more cognitively demanding, they may also mitigate the risks associated with online shopping. Inversely, while small choice set sizes are less cognitively demanding, they do less to help mitigate risk. The net result of choice set size online: no effect.

Further, maximizers, who were hypothesized to react even more negatively to large choice sets, may also perceive a risk-mitigating benefit from large choice sets online. Maximizers, who tend to experience higher post-purchase regret [41], may be especially attuned to the potential risks associated with online shopping. So while large choice sets may prompt more agonizing product comparisons for maximizers, they may also feel the benefits of reducing the perceived risks with shopping online—explaining the absence of a heightened choice overload effect for these consumers.

LIMITATIONS AND FUTURE WORK

The results of this research should be considered in light of certain limitations. First, although great care was taken to replicate an authentic e-commerce experience, the lottery compensation design may have created lower-stakes consequences for participants' choices. Following [21], this study also focused on chocolate purchases, a low-value, low-stakes product choice. More expensive and consequential product choices made online may yield different findings. Second, unlike true online shopping, this research investigated online decision-making where 100% of consumers ultimately made a product choice. Future work should consider exploring choice deferral or shopping cart abandonment in relation to choice set size in e-commerce contexts.

The choice set sizes tested in this study were derived by examining the design practices of current, top-performing e-commerce websites. As a result, the smallest choice set size tested was 12 options. Prior work demonstrated differences in choice satisfaction between very small choice set sizes (e.g., 4 or 6 options) and larger sets of 12 or more options; the current study's design was not able to test these comparisons. Future work may wish to test the choice overload effect with even smaller choice sets than are typically displayed by e-retailers today. Further, considering consumers' increasing reliance on mobile shopping technologies [34], it would be valuable to investigate the implications of selecting from different choice set sizes on smaller screens. Finally, future work could improve on our one-item scale for choice satisfaction, to capture a more nuanced measure for the subjective and complex nature of choice satisfaction.

CONCLUSION

The Internet provides e-retailers the ability to present more product options than is feasible in brick-and-mortar stores. Prior choice overload research suggests that this abundance of choice may have detrimental consequences for the user experience and levels of choice satisfaction. However, by directly testing the choice overload effect in an e-commerce context we demonstrate that the number of options presented online does *not* affect choice satisfaction. Where web designers may have been previously concerned about offering too many product options, we provide evidence that more is not necessarily less.

SUPPLEMENTARY MATERIAL

Supplemental material can be accessed at https://chocolate-animations.appspot.com/appendix/ or in the ACM Digital Library.

ACKNOWLEDGMENTS

We thank Matt Kay for his insights and the participants for their time. This research was funded in part by the Rackham Graduate School at the University of Michigan and by the National Science Foundation under Grant No. 1318143.

REFERENCES

1. Bharath Arunachalam, Shida R. Henneberry, Jayson L. Lusk, and F. Bailey Norwood. 2009. An Empirical Investigation into the Excessive-Choice Effect. *American Journal of Agricultural Economics* 91, 3: 810–825. http://doi.org/10.1111/j.1467-8276.2009.01260.x
2. Dirk Bollen, Bart P. Knijnenburg, Martijn C. Willemsen, and Mark Graus. 2010. Understanding Choice Overload in Recommender Systems. In *Proceedings of the Fourth ACM Conference on Recommender Systems* (RecSys '10), 63–70. http://doi.org/10.1145/1864708.1864724

3. D. Eric Boyd and Kenneth D. Bahn. 2009. When Do Large Product Assortments Benefit Consumers? An Information-Processing Perspective. *Journal of Retailing* 85, 3: 288–297. http://doi.org/10.1016/j.jretai.2009.05.008
4. Shun Cai, Yunjie Xu, and Jie Yu. 2008. The Effects of Web Site Aesthetics and Shopping Task on Consumer Online Purchasing Behavior. In *CHI '08 Extended Abstracts on Human Factors in Computing Systems* (CHI EA '08), 3477–3482. http://doi.org/10.1145/1358628.1358877
5. Alexander Chernev. 2011. Product Assortment and Consumer Choice: An Interdisciplinary Review. *Foundations and Trends in Marketing* 6, 1: 1–61. http://doi.org/10.1561/1700000030
6. Alexander Chernev. 2011. Product Assortment and Consumer Choice: An Interdisciplinary Review. *Foundations and Trends® in Marketing* 6, 1: 1–61. http://doi.org/10.1561/1700000030
7. Alexander Chernev, Böckenholt, and Goodman. 2014. Choice overload: A conceptual review and meta-analysis. *Journal of Consumer Psychology*. http://doi.org/10.1016/j.jcps.2014.08.002
8. Pawitra Chiravirakul and Stephen J. Payne. 2014. Choice Overload in Search Engine Use? In *Proceedings of the 32Nd Annual ACM Conference on Human Factors in Computing Systems* (CHI '14), 1285–1294. http://doi.org/10.1145/2556288.2557149
9. Jonathan D. D'Angelo and Catalina L. Toma. 2016. There Are Plenty of Fish in the Sea: The Effects of Choice Overload and Reversibility on Online Daters' Satisfaction With Selected Partners. *Media Psychology*: 1–27. http://doi.org/10.1080/15213269.2015.1121827
10. Ilan Dar-Nimrod, Catherine D. Rawn, Darrin R. Lehman, and Barry Schwartz. 2009. The Maximization Paradox: The costs of seeking alternatives. *Personality and Individual Differences* 46, 5–6: 631–635. http://doi.org/10.1016/j.paid.2009.01.007
11. Kristin Diehl and Cait Poynor. 2010. Great Expectations?! Assortment Size, Expectations, and Satisfaction. *Journal of Marketing Research* 47, 2: 312–322. http://doi.org/10.1509/jmkr.47.2.312
12. Zoltan Dienes. 2011. Bayesian Versus Orthodox Statistics: Which Side Are You On? *Perspectives on Psychological Science* 6, 3: 274–290. http://doi.org/10.1177/1745691611406920
13. Pierre Dragicevic. 2016. Fair Statistical Communication in HCI. In *Modern Statistical Methods for HCI*, Judy Robertson and Maurits Kaptein (eds.). Springer International Publishing, 291–330. Retrieved September 2, 2016 from http://link.springer.com.proxy.lib.umich.edu/chapter/10.1007/978-3-319-26633-6_13
14. C.R. Gallistel. 2009. The importance of proving the null. *Psychological Review* 116, 2: 439–453. http://doi.org/10.1037/a0015251
15. Anthony Griffin and Dennis Viehland. 2012. Demographic Factors in Assessing Perceived Risk in Online Shopping. In *Proceedings of the 13th International Conference on Electronic Commerce* (ICEC '11), 9:1–9:6. http://doi.org/10.1145/2378104.2378113
16. Gerald Häubl and Pablo Figueroa. 2002. Ineractive 3D Presentations and Buyer Behavior. In *CHI '02 Extended Abstracts on Human Factors in Computing Systems* (CHI EA '02), 744–745. http://doi.org/10.1145/506443.506576
17. Graeme A. Haynes. 2009. Testing the boundaries of the choice overload phenomenon: The effect of number of options and time pressure on decision difficulty and satisfaction. *Psychology and Marketing* 26, 3: 204–212. http://doi.org/10.1002/mar.20269
18. Rink Hoekstra, Richard D. Morey, Jeffrey N. Rouder, and Eric-Jan Wagenmakers. 2014. Robust misinterpretation of confidence intervals. *Psychonomic Bulletin & Review* 21, 5: 1157–1164. http://doi.org/10.3758/s13423-013-0572-3
19. George S. Howard, Scott E. Maxwell, and Kevin J. Fleming. 2000. The proof of the pudding: An illustration of the relative strengths of null hypothesis, meta-analysis, and Bayesian analysis. *Psychological Methods* 5, 3: 315–332. http://doi.org/10.1037/1082-989X.5.3.315
20. Jessica Hullman, Paul Resnick, and Eytan Adar. 2015. Hypothetical Outcome Plots Outperform Error Bars and Violin Plots for Inferences about Reliability of Variable Ordering. *PloS one* 10, 11: e0142444.
21. Iyengar and Mark R. Lepper. 2000. When choice is demotivating: Can one desire too much of a good thing? *Journal of Personality and Social Psychology* 79, 6: 995–1006. http://doi.org/10.1037/0022-3514.79.6.995
22. Sheena S. Iyengar, Rachael E. Wells, and Barry Schwartz. 2006. Doing Better but Feeling Worse: Looking for the "Best" Job Undermines Satisfaction. *Psychological Science* 17, 2: 143–150.
23. Eric J. Johnson, Suzanne B. Shu, Benedict G. C. Dellaert, et al. 2012. Beyond nudges: Tools of a choice architecture. *Marketing Letters* 23, 2: 487–504. http://doi.org/10.1007/s11002-012-9186-1
24. Matthew Kay, Gregory L. Nelson, and Eric B. Hekler. 2016. Researcher-Centered Design of Statistics: Why Bayesian Statistics Better Fit the Culture and Incentives of HCI. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems* (CHI '16), 4521–4532. http://doi.org/10.1145/2858036.2858465
25. Scott Kessler. 2014. Industry Surveys, Computers: Consumer Services & the Internet. *S&P Capital IQ*.

26. John K. Kruschke. 2010. Bayesian data analysis. *Wiley Interdisciplinary Reviews: Cognitive Science* 1, 5: 658–676. http://doi.org/10.1002/wcs.72
27. Khai Sheang Lee and Soo Jiuan Tan. 2003. E-retailing versus physical retailing: A theoretical model and empirical test of consumer choice. *Journal of Business Research* 56, 11: 877–885. http://doi.org/10.1016/S0148-2963(01)00274-0
28. Junzhao Ma. 2016. Does Greater Online Assortment Pay? An Empirical Study Using Matched Online and Catalog Shoppers. *Journal of Retailing*.
29. Richard McElreath. 2012. *Statistical rethinking R package*. R package version 1.
30. Richard McElreath. 2016. *Statistical Rethinking: A Bayesian Course with Examples in R and Stan*. CRC Press.
31. Shailey Minocha, Marian Petre, Ekaterini Tzanidou, et al. 2006. Evaluating e-Commerce Environments: Approaches to Cross-disciplinary Investigation. In *CHI '06 Extended Abstracts on Human Factors in Computing Systems* (CHI EA '06), 1121–1126. http://doi.org/10.1145/1125451.1125663
32. Anthony D. Miyazaki and Ana Fernandez. 2001. Consumer Perceptions of Privacy and Security Risks for Online Shopping. *Journal of Consumer Affairs* 35, 1: 27–44. http://doi.org/10.1111/j.1745-6606.2001.tb00101.x
33. Antti Oulasvirta, Janne P. Hukkinen, and Barry Schwartz. 2009. When More is Less: The Paradox of Choice in Search Engine Use. In *Proceedings of the 32Nd International ACM SIGIR Conference on Research and Development in Information Retrieval* (SIGIR '09), 516–523. http://doi.org/10.1145/1571941.1572030
34. PWC. 2015. Total Retail 2015: Retailers and the Age of Disruption. *PricewaterhouseCoopers*.
35. Elena Reutskaja and Robin M. Hogarth. 2009. Satisfaction in choice as a function of the number of alternatives: When "goods satiate." *Psychology and Marketing* 26, 3: 197–203. http://doi.org/10.1002/mar.20268
36. Jens Riegelsberger and M. Angela Sasse. 2002. Face It - Photos Don't Make a Web Site Trustworthy. In *CHI '02 Extended Abstracts on Human Factors in Computing Systems* (CHI EA '02), 742–743. http://doi.org/10.1145/506443.506575
37. Jens Riegelsberger, M. Angela Sasse, and John D. McCarthy. 2003. Shiny Happy People Building Trust?: Photos on e-Commerce Websites and Consumer Trust. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (CHI '03), 121–128. http://doi.org/10.1145/642611.642634
38. Benjamin Scheibehenne, Rainer Greifeneder, and Peter M. Todd. 2009. What moderates the too-much-choice effect? *Psychology and Marketing* 26, 3: 229–253. http://doi.org/10.1002/mar.20271
39. Benjamin Scheibehenne, Rainer Greifeneder, and Peter M. Todd. 2010. Can There Ever Be Too Many Options? A Meta-Analytic Review of Choice Overload. *Journal of Consumer Research* 37, 3: 409–425. http://doi.org/10.1086/651235
40. Barry Schwartz. 2004. *The paradox of choice: why more is less*. ECCO, New York.
41. Barry Schwartz, Andrew Ward, John Monterosso, Sonja Lyubomirsky, Katherine White, and Darrin R. Lehman. 2002. Maximizing versus satisficing: Happiness is a matter of choice. *Journal of Personality and Social Psychology* 83, 5: 1178–1197. http://doi.org/10.1037/0022-3514.83.5.1178
42. Barry Schwartz, Andrew Ward, John Monterosso, Sonja Lyubomirsky, Katherine White, and Darrin R. Lehman. 2002. Maximizing versus satisficing: Happiness is a matter of choice. *Journal of Personality and Social Psychology* 83, 5: 1178–1197. http://doi.org/10.1037/0022-3514.83.5.1178
43. Avni M. Shah and George Wolford. 2007. Buying Behavior as a Function of Parametric Variation of Number of Choices. *Psychological Science* 18, 5: 369–370. http://doi.org/10.1111/j.1467-9280.2007.01906.x
44. Herbert A. Simon. 1996. *The Sciences of the Artificial*. MIT Press.
45. Statista. 2015. E-commerce in the United States. *Statista Dossier*. Retrieved from http://www.statista.com/study/28028/e-commerce-in-the-united-states-statista-dossier/
46. Ulrike Steinbrück, Heike Schaumburg, Sabrina Duda, and Thomas Krüger. 2002. A Picture Says More Than a Thousand Words: Photographs As Trust Builders in e-Commerce Websites. In *CHI '02 Extended Abstracts on Human Factors in Computing Systems* (CHI EA '02), 748–749. http://doi.org/10.1145/506443.506578
47. Vertical Web Media LLC. 2015. Top 500 Guide. *Internet Retailer*.