

Memory-based preferential choice in large option spaces

Adam Hornsby

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
of
University College London.

Department of Experimental Psychology
University College London

October 3, 2022

I, Adam Hornsby, confirm that the work presented in this thesis is my own. Where information has been derived from other sources, I confirm that this has been indicated in the work.

For Daisy-May ♡

Abstract

Whether adding songs to a playlist or groceries to a shopping basket, everyday decisions often require us to choose between an innumerable set of options. Laboratory studies of preferential choice have made considerable progress in describing how people navigate fixed sets of options. Yet, questions remain about how well this generalises to more complex, everyday choices. In this thesis, I ask how people navigate large option spaces, focusing particularly on how long-term memory supports decisions. In the first project, I explore how large option spaces are structured in the mind. A topic model trained on the purchasing patterns of consumers uncovered an intuitive set of themes that centred primarily around goals (e.g., tomatoes go well in a salad), suggesting that representations are geared to support action. In the second project, I explore how such representations are queried during memory-based decisions, where options must be retrieved from memory. Using a large dataset of over 100,000 online grocery shops, results revealed that consumers query multiple systems of associative memory when determining what choose next. Attending to certain knowledge sources, as estimated by a cognitive model, predicted important retrieval errors, such as the propensity to forget or add unwanted products. In the final project, I ask how preferences could be learned and represented in large option spaces, where most options are untried. A cognitive model of sequential decision making is proposed, which learns preferences over choice attributes, allowing for the generalisation of preferences to unseen options, by virtue of their similarity to previous choices. This model explains reduced exploration patterns behaviour observed in the supermarket and preferential

choices in more controlled laboratory settings. Overall, this suggests that consumers depend on associative systems in long-term memory when navigating large spaces of options, enabling inferences about the conceptual properties and subjective value of novel options.

Impact Statement

My research elucidates the contribution of long-term memory to preferential choices in naturalistic environments. Experimental studies of preferential choice tend to examine how participants choose between small sets of explicitly-presented options, which tends to minimise the contribution of prior knowledge about option relationships. The results presented in Chapter 2 and 3 show how conceptual representations of shoppers (and their individual differences) can be approximated through their prior choice history. I hope this will be a useful tool for cognitive scientists wishing to study the contribution of mental representations of option similarity during naturalistic preferential choice.

These findings also emphasise many of the unique characteristics of naturalistic preferential choice. As detailed throughout the thesis, long-term memory helps to guide preferential choices in everyday contexts (like the supermarket), where the option space is large and often requires retrieval of attributes and options from memory. Fewer studies have explored how people make decisions in these contexts. I hope this thesis will stimulate more studies investigating naturalistic preferential choice and prompt others to use large observational datasets of human choices to validate their theoretical claims.

The ideas proposed in my research have had a direct impact for my employer, dunnhumby, and should be of substantial relevance to those working in retail marketing. Specifically, the idea of using Latent Dirichlet Allocation (LDA) to approximate conceptual representations of products through product co-occurrences (proposed in Chapter 2) is now used widely across the

business. LDA helps retailers to discover and track the high-level goals of shoppers. Our work shows that these goals closely resemble the way that consumers think about products (e.g., “meals for tonight”) and contrast with more widely-available, highly-curated product groups created by logistics teams (e.g. “fresh food”).

Similarly the ideas proposed in Chapter 3 should be of interest to marketers and machine learning practitioners building personalised recommender systems. Retailers often wish to predict consumers’ subsequent purchases online, as it helps them to prioritise marketing, such as personalised promotions. Typically this would be achieved using machine learning models, such as recurrent neural networks. However, our results revealed that sequential choices could be predicted using a relatively simple cognitive model of memory retrieval. Cognitive models may capture new variance that machine learning models may not, helping them to improve the relevance of recommendations to customers (see Kopeinik et al., 2017, for an example). They may also be more computationally efficient than machine learning models and are often more interpretable.

Finally, the ideas proposed in Chapter 4 highlight an interesting interaction between recommender systems and human preferences. In particular, recommender systems are often designed to suggest similar products or content to those consumed by a user in the past. Our results indicate that consumers update their preferences to support their choices. This could lead to self-perpetuating cycles, in which preferences and recommendations focus on an increasingly narrow set of attributes. Machine learning practitioners could mitigate this by promoting more diverse recommendations or leveraging this interaction to – with consent – encourage more pro-social behaviour, such as healthier or more sustainable choices.

Acknowledgements

I have been very fortunate to be able to pursue a PhD and even luckier to have the support of so many great people.

Firstly, I am grateful to my amazing supervisor, Bradley Love, for lending me so much of his time, patience and brilliant ideas over the years. You've always known the right questions to ask and the right approach to take. It's always a thrill to discuss new ideas and results with you and I'll miss our long discussions in 26 Bedford Way.

I'd also like to thank my secondary advisor, Marten Speekenbrink, for his excellent guidance over the years and his exceptional statistics course, which taught me so much about model comparison.

I've loved following the progress of so many amazing researchers in the UCL Love Lab and I'm very grateful for their feedback and insightful discussions over the years. This includes (in approximate chronological order) Peter Riefer, Sebastian Bobadilla Suarez, Paula Parpaart, Christiane Ahlheim, Takao Noguchi, Kurt Braunlich, Sébastien Ohleyer, Mina Stojanovic, Brett Roads, Rob Mok, Olivia Guest, Xiaoliang Luo, Nick Sexton, Daniel Barry, Kaarina Aho, Franziska Bröker and Nikolay Dagaev, amongst others.

At dunnhumby, I had the support of many amazing Data Scientists. Thank you to Giles Pavey and Piers Stobbs for giving me the confidence to start this PhD. Thank you to Rosie Prior for the supervision over the years, the many counselling sessions and for dealing with so many gruesome aspects of industry-academic collaboration, like legal contracts, data privacy and paper approvals. I've been lucky to have the support of many great leaders during

my PhD, including Pam Campbell, Carlos Ardid, Sandra Stanley, Rex Davis and Chris Brooks; this PhD could never have happened without your positivity and championing of extended academic research. I was lucky to supervise several brilliant Machine Learning MSc interns who also shaped my thinking, including Jiayun Yu and David Errington, amongst others. There are many, many amazing dunnhumby Data Scientists that have inspired me over the years, including Gareth Brown (who checked lots of my maths), Victor Robin, Morgan Jones and countless others. Thanks also to Marco Lattanzio and his BDaSS team for maintaining dunnhumby's servers whilst I ran endless model comparisons.

I am also immensely grateful to the Commission for the Royal Exhibition of 1851 for supporting me as an Industrial Fellow of their illustrious organisation. This generous support allowed me to spend more time on my PhD each week, to visit conferences and to connect with many brilliant scientists from other fields.

There are many other great scientists whose feedback shaped the thinking in my dissertation; particularly those that listened to me present. I'm very grateful to the Bath Psychology and Economics colloquium and the DeepMind Neuroscience team for their incisive questions about my research.

I attended many great conferences and I'm particularly thankful to those at Reinforcement Learning and Decision Making 2019 for making my final in-person conference so memorable. I'm very grateful to those that stopped to give feedback on my research.

On a personal note, I'd like to thank my amazing wife and partner-in-crime Daisy-May Boutwood. For all the words in this thesis, there are none to surmise my gratitude to you. Academic research is full of ups and downs, but also full of many boring, frustrating nights and weekends in-front of a laptop; you have been patient and understanding through 5 years of this. Thank you.

I'd also like to thank my mum, Pam; thanks for bestowing in me the grit and determination that you always had and for being so supportive. And I'd

like to thank my dad, Jim, for giving me the gift of academic curiosity.

Bibliography

Kopeinik, S., Kowald, D., Hasani-Mavriqi, I., and Lex, E. (2017). Improving Collaborative Filtering Using a Cognitive Model of Human Category Learning. *The Journal of Web Science*, 2(4):45–61.

Published articles

Chapters 2, 3 and 4 are partially or fully based on the following articles:

- Chapter 2 - Hornsby, A.N., Evans, T., Riefer, P. S., Prior, R. Love, B.C. (2020). *Conceptual Organization is Revealed by Consumer Activity Patterns*. *Computational Brain and Behavior*. 3, 162–173.
- Chapter 3 - Hornsby, A.N. Love, B.C. (2022). *Sequential consumer choice as multi-cued retrieval*. *Science Advances*. 8. 8.
- Chapter 4 - Hornsby, A.N. Love, B.C. (2020). *How Decisions and the Desire for Coherency Shape Subjective Preferences Over Time*. *Cognition*. 200. 104244.

I declare that I was fully involved through the whole process of investigation for these articles which includes matters regarding experimental design, analysis, and writing of the articles themselves. Declaration forms for reusing these published works are available at the end of this thesis.

Contents

1	Memory and preferential choice	30
1.1	Introduction	30
1.2	A brief history of preferential choice research	32
1.3	A brief introduction to human memory systems	37
1.4	Retrieval of option attributes from memory	39
1.4.1	Determinants of attribute retrieval	39
1.4.2	Conceptual representation of options	41
1.5	Retrieval of options from memory	42
1.5.1	Option generation	43
1.5.2	Option selection	44
1.6	Exploiting past experience	45
1.6.1	Recognition, familiarity and fluency	46
1.6.2	Decisions from experience	47
1.6.3	Sequential sampling from memory	48
1.6.4	Reinforcement learning	49
1.7	Choice-supportive memory biases	51
1.7.1	Choice-supportive misremembering	52
1.7.2	Choice-induced preference change	53
1.8	Research question	54
1.9	Dissertation outline	55
2	Conceptual organisation of large option spaces	77
2.1	Introduction	77

2.2	Training a topic model with retail data	84
2.2.1	Method	84
2.2.2	Results and discussion	85
2.3	Evaluating topic labels with retail experts	85
2.3.1	Method	87
2.3.2	Results and discussion	88
2.4	Seasonal trends in topics	90
2.4.1	Method	90
2.4.2	Results	90
2.5	Evaluating topic coherency with typical consumers	92
2.5.1	Method	92
2.5.2	Results and discussion	94
2.6	Classifying individual consumers by their experienced topics .	95
2.6.1	Method	95
2.6.2	Results	96
2.7	General Discussion	96
3	Option retrieval in sequential, open-ended tasks	103
3.1	Introduction	103
3.2	Data	107
3.2.1	Clickstream data	107
3.2.2	In-store data	109
3.3	Past choices cue subsequent retrievals	112
3.3.1	Method	112
3.3.2	Results	113
3.4	Episodic, semantic and hierarchical knowledge explain choice .	115
3.4.1	Method	116
3.4.2	Results	118
3.5	Relying on certain knowledge formats predicts retrieval errors	120
3.5.1	Method	121
3.5.2	Results	122

3.6	Discussion	124
4	Learning preferences from past choices	134
4.1	Introduction	134
4.2	The Coherency Driven Choice (CDC) Model	140
4.3	Learning strong preferences over time	144
4.3.1	Method	144
4.3.2	Results	146
4.4	People prefer novel patterns associated with their prior choice	148
4.4.1	Method	148
4.4.2	Results	151
4.5	Political beliefs become consistent with a prior vote	153
4.5.1	Method	154
4.5.2	Results	158
4.6	General Discussion	160
5	General Conclusions	173
5.1	Preferential choice in large option spaces	174
5.2	Long-term memory and preferential choice	176
5.3	Learning from decisions in the absence of extrinsic reward	177
5.4	Bidirectional influences between choices and representations	179
5.5	Studying choices outside of the lab	180
5.6	General limitations	182
5.6.1	Asymmetric similarity	182
5.6.2	Separating memory representations from process	183
5.6.3	Causal limitations of observational studies	184
5.6.4	Beyond the supermarket	186
5.7	Future directions	188
5.7.1	A combined model of option retrieval and preferential choice	188
5.7.2	Interactions with other choice attributes	189

5.7.3	Leveraging different types of conceptual similarity . . .	190
5.7.4	Adjusting subjective preferences to support healthier choices	191
5.8	Looking ahead	192
5.9	Conclusion	193

Appendices **203**

A Probabilistic utility models **203**

B Chapter 2 Appendix **205**

B.1	Latent Dirichlet Allocation	205
B.1.1	Topic inspection and labelling	206
B.1.2	Calculating product relevancy for a topic	207
B.1.3	Initial topic labelling	209
B.2	Label confusion by retail experts	210
B.3	Topics by day of week	210

C Chapter 3 Appendix **212**

C.1	Method	212
C.1.1	Data	212
C.1.2	Representations of associative knowledge	214
C.1.3	Retrieval model	217
C.1.4	Permutation tests	218
C.1.5	Response times	219
C.1.6	Trajectory analyses	219
C.1.7	Transition clustering	219
C.1.8	Predictive modelling	221
C.2	Additional analyses of clickstream data	221
C.2.1	Correlations between representations	221
C.2.2	Similarity ripples	222
C.2.3	Timestep and IRI regression	222

- C.2.4 Similarity and timestep regression 223
- C.2.5 Representation and IRI regression 225
- C.2.6 Transitions between categories reveal hierarchical knowledge 226
- C.2.7 Correlation between attention weights 227
- C.2.8 Retrieval model parameter recovery 227
- C.2.9 Forgotten items regression 228
- C.2.10 Removed items regression 229
- C.2.11 Search arrival transitions 230
- C.3 Analyses of food fluency data 235
 - C.3.1 Method 236
 - C.3.2 Results and discussion 237

D Chapter 4 Appendix 243

- D.1 Experiment 1 243
 - D.1.1 Preferences did not change over time 243
 - D.1.2 Preference change doesn't vary as a function of political affiliation 243
- D.2 Experiment 2 245
 - D.2.1 Preference change varies within political topic 245
 - D.2.2 Responses were not biased in favour of a particular party affiliation or slider response 247

List of Figures

2.1	The input in a corpus analysis is typically item counts (i.e., word counts) within some context (e.g., a sentence or document). Analogously, products (akin to words) are organized into baskets (akin to sentences). One advantage of applying these analysis techniques to baskets is that, unlike natural language, meaning is unaffected by item order.	81
2.2	Latent Dirichlet Allocation (LDA) uncovers the higher-level product topics that can be viewed as generating the observed baskets purchased by consumers. LDA's fit is driven by the co-occurrence pattern of products within baskets. In the solution, each product has a probability of occurring within each topic (shown on the left for apple). The colours illustrate which topic each product would have been labelled with if using the maximum product topic probability. Each basket is generated by a mixture of probabilities over the topics (shown on the right for this basket).	83
2.3	Proportion correct with standard error bars for the study on label agreement involving retail experts and the intruder study involving typical consumers. All proportions were significantly different ($p < .001$) than chance levels, 25.00% (1 of 4) and 16.67% (1 of 6), respectively.	89

2.4 Topic prevalence varies by season. The proportion of baskets with a given topic label in each month of 2014, divided by the monthly mean average across all topics (i.e., index), is shown. a) Topics that should be seasonal peak at the expected time, such December for the Christmas topic. b) In contrast, topics for staple products vary less in prevalence over time. 91

3.1 Deciding what to choose next when shopping for groceries online depends on cued retrieval from multiple knowledge sources. **a.**, We used 4.3m unordered, in-store receipts to build representations of episodic, semantic and hierarchical knowledge. **b.**, To model retrieval, we collected data from 135,000 shoppers as they sequentially searched for products on the website of one of the UK’s largest supermarket retailers. **c.**, Prior choices predict future ones, by virtue of their similarity according to different representational formats. Once an item is added to their basket, shoppers use this to cue matches from long-term memory. The stronger the match with this cue, the higher the probability an item will be retrieved (this may be attenuated by increased attention towards a particular representation). Retrieved items are checked against one’s internal goals. If the retrieval is goal-relevant, the shopper adds an appropriate item from the website and uses that item to cue associations. If not, a new option is retrieved and checked for goal relevance until one is accepted. Similar heuristic strategies have been used in models of option generation for single choices (Johnson and Raab, 2003; Klein et al., 1995). Once all goals are satisfied, the user checks out. Note that the goals and the goal-checking process is not modelled here. 106

- 3.2 Consecutive purchases tend to be close episodic, semantic and hierarchical relations. **a.**, Choices are predicted by their similarity with the prior choice across each representation. Histograms show that the similarity between consecutive purchases (averaged for each visit) was higher compared to when the order of purchases was randomly permuted. (with 95% confidence intervals). **b.**, Sequential retrieval is like a ripple in semantic memory. Mean episodic similarity (with 95% confidence intervals) between the current product and those purchased recently is higher compared with products purchased later **c.**, Visitors slowed as they approached the end of their shopping trip. Mean response times (with 95% confidence intervals) as a function of timestep quantile (Small = 10 - 30 items, Medium = 31 - 49 items, Large = 50+ items)). **d.**, Consumers make more between-category transitions (i.e., taxonomy level five) towards the end of their visit. Stacked density plots denoting the proportion of switches according to each level of the taxonomy as a function of the relative timestep. **e.**, Transitions between product groups at the fourth level of the hierarchy clustered into intuitive higher-order groupings that appear similar to those in the product taxonomy, suggesting that the taxonomy closely resembles how shoppers represent products during sequential choice. The *Lift* - 1 of each transition is depicted in purple, with values less than 0 shown in grey. Boxes represent clusters identified by the optimal spectral clustering solution (more information in section 1.7. and 2.6 of the appendix). 108

3.3 Mean number of forgotten items (with 95% confidence intervals) for each model attention weight (β). Results show that relying on episodic or hierarchical knowledge predicted fewer forgotten items, whereas attending to semantic knowledge predicted more forgotten items, as measured by use of a recommender system displayed before checkout. 122

- 4.1 Many popular models of decision-making cannot easily explain how people form strong subjective preferences for free choices made in the wild. **a**, In standard decision theory, it is assumed that preferences remain stable over time (von Neumann et al., 1944; Glimcher, 2009). Indeed, for many researchers, the challenge is often learning what people’s preferences *are* (e.g., by asking them to choose between options), rather than understanding how they *became*. **b**, Reinforcement Learning (RL) models contrast in that they assume preferences change over time. Specifically, RL agents learn to prefer actions with a higher expected reward, which they learn as they monitor extrinsic feedback from their environment (Sutton and Barto, 1998). While RL has been shown to account for many aspects of human learning well (for a review, see Daw and Tobler 2013), these investigations have been largely confined to objective tasks, where there is a clear extrinsic signal steering the decision-maker. **c**, Studies of free choice — where there is no objectively correct answer — have shown that merely choosing an item increases one’s preference for it. This has often been taken to imply that free-choices are self-reinforcing, so as to increase preferences toward the chosen option. Yet, caching values in this way would arguably not scale well, as it would require people to keep track of every item they’d ever tried. **d**, In this paper, we propose a novel theory of subjective preference formation and decision making that arguably scales better to free choices made in the real world. According to this theory, people encode preferences over attributes of free choices, such as the hop content in beer. After making a decision, they then update their preferences in the direction of the attributes that defined their prior choice, thereby increasing their preference for it, as well as for other options that have similar attributes. 138

- 4.2 To illustrate how this theory can materialize in strong subjective preferences, we formalize it in a computational cognitive model, known as Coherency Driven Choice (CDC). **a**, CDC possesses a preference and attention-weight vector. When evaluating options, the model evaluates the similarity between its own attention-weighted preference and the attributes of the available options. The closer an option's attributes are to the model's attention-weighted preferences, the more likely the model is to select it. **b**, Following a choice, CDC updates its preference vector to make the past choice most likely using gradient descent, thereby moving toward the prior choice in attribute space. **c**, CDC also adjusts its attention weights to make prior choice more likely, effectively warping the preference space so that it becomes more sensitive to the attended attribute. **d**, As the model grows to prefer the options it initially chose, it tends to repeat the same choice type more over time. **e**, The model was simulated for 10,000 timesteps in a simple environment in which there were two choice types. Due to happenstance in the initial choices, the model began to prefer *choice type a*, adjusting its preferences and attention weights in favor of the attributes that make it unique. The preference history is coloured by attention weight ratio, such that blacker colors indicate a greater deal of attention paid towards attribute 2. 145

- 4.3 **a**, The trial sequence observed by participants. First, participants were introduced to a robot and asked to design aspects of it. They were then introduced to a new robot, which was designed differently. These robots then turned around, revealing random patterns on their backs. One pattern was unique to the robot participants' had previously designed (i.e., chosen-unique), one was shared across both robots (i.e., shared) and the other was unique to the robot they had previously eschewed (i.e., non-chosen). Participants were then asked to rank the patterns in order of preference. **b**, This chart depicts the proportion that each image type was chosen as a first, second and third preference (with standard error bars). These results are consistent with the theory presented here, which predicts that choice-based learning generalizes most strongly to the unique, most-discriminating feature of the original choice (i.e., chosen unique), followed by the shared feature. 151

4.4 **a**, Participants were initially asked to vote for a candidate based on trivial characteristics. Both candidates then revealed opposing beliefs on a more controversial belief. Participants were then asked to rate how much they agreed with their candidate’s newly revealed opinion. **b**, If voting for the candidate affects the extent to which participants agree with the later-revealed political belief, then one would expect participants to show different levels of agreement between groups, depending on whether their candidate revealed a traditionally left-wing view or a right-wing view. **c**, The results of the experiment are shown in kernel density plots (with rugplots below to depict the individual data points). These reveal that Republican-identifying participants were particularly prone to adjusting their preference to be in accordance with that revealed by their chosen candidate. **d**, Plotting effect sizes for each per-topic comparison shows that Republican-identifying participants were particularly prone to adjusting their preferences. 159

B.1 The relationship between item frequency within the corpus and a) item-topic probability, b) lift and c) relevancy 208

B.2 The top 50 most frequently occurring items across baskets within the corpus. Note that brand names have been removed. 209

B.3 The proportion of baskets with a given topic label on each day of the week, divided by the weekly mean average across all topics. Plot a) shows that food requiring longer preparation times (i.e. *loose fruit and veg*) or eaten specifically during weekend occasions (i.e. *afternoon tea*) are more likely to be bought on Thursday or Friday. Plot b) indicates that impulse purchases (i.e. *food for now*) or food that tends to be stored away (i.e. *branded store cupboard*) does not vary in popularity over the week. 211

- C.1 Shoppers could transition between products using different design features of the website. **a.**, Most transitions occurred through use of a search bar, which was located at the top of the page. After entering a keyword, shoppers were presented with a list of relevant products associated with the keyword. Shoppers could add products to their basket from this search results page directly or click on the product to view a dedicated page containing more information (e.g., nutritional data). **b.**, Fewer transitions occurred through use of a category dropdown, which appeared when hovering the mouse over the *Groceries* hyperlink at the top of the page. Three levels of subcategories could be revealed by hovering one’s mouse over the respective department names. In the example pictured, the shopper has hovered their mouse over the “Cupboard” department and then the “Cereals” category. The products in each subcategory were displayed in the same way as the search results. 213
- C.2 Lagged average similarity between choices indicates how choices become less similar over time 222

List of Tables

2.1	Retailer-supplied product descriptions for the 5 most relevant products within each of the 10 surveyed topics. Note that the authors had access to the full product-topic relevancy matrix (see https://osf.io/tsymx/) when they labeled the topics. Brand names have been removed from this table for publication.	86
3.1	The % BIC improvement over the random baseline and the mean attention weights (with 95% confidence intervals) for each of the candidate models. Results show that including representations of multiple knowledge formats provides the best fit to the data (shown in bold)	117
4.1	A summary of the neutral statements shown to participants .	156
4.2	A summary of the controversial statements revealed by candidates following a vote	156
B.1	The labels given to each of the 25 topics. Size is defined as the number of products that had the highest probability of belonging to the respective topic over the total number of products in the corpus. Asterisks indicate that they were surveyed in studies A and B.	206
B.2	A summary of mislabelling errors made by retail experts during the topic labelling task (***) indicates a proportion significantly different ($p < .001$) from the random baseline of 25.00% (1 of 4))	210

C.1	Spearman correlations between each of the similarity measures (95% confidence intervals shown in parentheses)	221
C.2	Standardised coefficients, significance values, upper (UB) and lower bounds (LB) for variables included in the mixed-effects linear regression predicting IRI using timestep, transition similarity and transition type	223
C.3	Standardised coefficients, significance values, upper (UB) and lower bounds (LB) for variables included in the mixed-effects linear regression predicting timestep using episodic similarity and transition type	223
C.4	Standardised coefficients, significance values, upper (UB) and lower bounds (LB) for variables included in the mixed-effects linear regression predicting timestep using semantic similarity and transition type	224
C.5	Standardised coefficients, significance values, upper (UB) and lower bounds (LB) for variables included in the mixed-effects linear regression predicting timestep using hierarchical similarity and transition type	224
C.6	Standardised coefficients, significance values, upper (UB) and lower bounds (LB) for variables included in the mixed-effects linear regression model predicting IRI using each similarity measure and confounding variables	225
C.7	Nested model comparisons (using AIC and BIC) for linear-mixed effects regressions predicting IRI. Models were chosen using backward elimination.	226
C.8	Spearman correlations between each of the attention weights .	227

C.9	Standardised coefficients, significance values, upper (UB) and lower bounds (LB) for variables included in the regression model predicting number of forgotten items using the model attention weights, total number of choices and proportion of each transition type	228
C.10	Standardised coefficients, significance values, upper (UB) and lower bounds (LB) for variables included in the regression model predicting number of products removed using the model attention weights, total number of choices and proportion of each transition type	229
C.11	Spearman correlations between each of the similarity measures in the search-arrival dataset ($N = 3086716$)	231
C.12	Standardised coefficients, significance values, upper (UB) and lower bounds (LB) for variables included in the mixed-effects linear regression predicting IRI using timestep, fit to the search-arrival dataset ($N = 2959878$)	232
C.13	Standardised coefficients, significance values, upper (UB) and lower bounds (LB) for variables included in the mixed-effects linear regression model predicting IRI using each similarity measure, fit to the search-arrival dataset	233
C.14	Nested model comparisons (using AIC and BIC) for regressions predicting IRI using each similarity measure, fit to the search-arrival dataset. Models were chosen using backward elimination.	233
C.15	The % BIC improvement over the random baseline and the mean attention weights (with 95% confidence intervals) for each of the candidate retrieval models, fit to the search-arrival dataset. Results show that including representations of all knowledge formats provides the best fit to the data (shown in bold)	234

C.16 Standardised coefficients, significance values, upper (UB) and lower bounds (LB) for variables included in a multiple linear regression predicting the number of forgotten items, fit to the search-arrival dataset	234
C.17 Standardised coefficients, significance values, upper (UB) and lower bounds (LB) for variables included in a multiple linear regression predicting the number of removed items, fit to the search-arrival dataset	234
C.18 Matching retailer products for the 40 most frequently occurring retrievals reported by Zemla et al. (2020a). Note that brand names have been redacted.	240
C.19 The % BIC improvement over the random baseline and the mean attention weights (with 95% confidence intervals) for each of the candidate retrieval models, <i>fit to each retrieval list separately</i> . Results show that including representations of all knowledge formats provides the best fit to the data (shown in bold)	241
C.20 The % BIC improvement over the random baseline and the mean attention weights (with 95% confidence intervals) for each of the candidate retrieval models, <i>fit to each participant and collapsing over multiple lists</i> . Results show that including representations of all knowledge formats provides the best fit to the data (shown in bold)	241
C.21 The % BIC improvement over the random baseline and the mean attention weights (with 95% confidence intervals) for each of the candidate retrieval models, <i>fit to the first retrieval list from each participant</i> . Results show that including representations of all knowledge formats provides the best fit to the data (shown in bold)	242

Chapter 1

Memory and preferential choice

1.1 Introduction

“Learning to choose is hard.
Learning to choose well is
harder. And learning to choose
well in a world of unlimited
possibilities is harder still”

Barry Schwartz

The Paradox of Choice

Whether deciding which restaurant to visit or which shampoo to purchase, we are often required to navigate an unfathomable amount of choice. In the supermarket, consumers often choose between tens of thousand of options and do so according to a breadth of considerations. Which products go well together? Which are interchangeable? Which are overpriced? Which ones taste bad? Despite having never tried most products on offer, we can navigate large option spaces with relative ease. In this thesis, I ask *how*, focusing particularly on how long-term memory provides the cognitive apparatus for supporting everyday decisions in-the-wild.

Since early economic models of subjective utility, psychological experiments have helped to progress our understanding of how humans learn and

use subjective preferences during decision making. For example, simple experimental manipulations — such as changing how a risky investment decision is framed (Tversky and Kahneman, 1981; Slovic et al., 1985) or introducing an additional option into a choice set (Simonson and Tversky, 1992) — have been shown to reverse preferences for options. These experimental findings contradict many of the basic assumptions of economic models, such as that preferences are stable (Savage, 1954). Accounting for these preference reversals is now considered a benchmark for models of preferential choice.

An important facet of these behavioural experiments is that they are carefully designed so as to minimise the effects of extraneous variables that are otherwise hard to account for, such as prior knowledge. For example, choice options are often presented explicitly; we refer to these as *stimulus-based choices* (Lynch and Srull, 1982). Moreover, participants are often asked to make their choices on the basis of explicit attributes, and thus make *decisions by description* (Hertwig, 2012). The set of possible alternatives also tends to be small and choices are often randomised so as to be independent.

Yet, preferential choices in real-world environments are rarely so straightforward. In a supermarket, consumers must choose between tens of thousands of products, most of which will be new to them. Attributes of options may be retrieved from memory; these are often referred to as *memory-based choices* (Bettman, 1979; Weber and Johnson, 2006) and may be integrated with explicitly presented attributes (Bettman, 1979; Lynch and Srull, 1982). Related prior experience may be exploited to guide inferences about the properties of options, such about their subjective value (known as *decisions from experience*, Hertwig, 2012). Choices may depend on those made previously, like whether a given ingredient complements the ones we already have in our shopping basket. Thus, preferential choices in lab tasks often contrast with those made in naturalistic domains in a number of important ways.

The aim of this thesis is to understand how people navigate large option spaces in naturalistic preferential choice environments. These contexts

are uniquely characterised by many of the properties described above; the option space is large and dynamic, often requiring attributes and options to be retrieved from memory and thus dependent on prior experience. I will show long-term memory facilitates preferential choice in these spaces; helping decision-makers to mentally organise the plethora of options available, infer their properties when making decisions and justify them in retrospect. To do this, I will combine laboratory experiments with analyses of large, naturally occurring datasets that capture consumer choices made in-the-wild. I will use computational models to bridge this behaviour with cognition, building on existing theories of how humans represent and retrieve long-term knowledge and update these in the context of choice.

1.2 A brief history of preferential choice research

Preferences are most typically defined as an order given by an agent to alternatives when making a choice. Early economic models of choice assumed that preferences could be described by their *utility* (Bentham, 1789). Initially, utilities were represented using real numbers over available options, allowing for options and bundles to be aggregated and compared. For example, early choice models used utilities to model the functional relationship between different products, such as the point which consumers became indifferent towards bundles of goods (a brief review of these models is provided in Appendix A and the reader is advised to consult Busemeyer and Rieskamp 2014 for further information). The terms utility and preference are nowadays used interchangeably (Warren et al., 2011), although may be measured in different ways. For example, some may equate preferences with a choice or willingness to pay (i.e., *revealed preferences*) whereas others may measure them explicitly using rating scales (i.e., *stated preferences*). *Revealed preferences* such as choices may appeal to those wanting to predict subsequent choices, given that there can

be a gap between what people *say* and *do* (LaPiere, 1934). However — as I will introduce below — preferential decisions are highly sensitive to the choice context and people may not even choose the same option consistently within the same experiment (Mosteller and Nogee, 1951; Hey, 2001).

Early work by cognitive scientists argued that — contrary to rational economic models — decision makers are limited in their capacity for processing information (Simon, 1955) and may therefore restrict their consideration of attributes and alternatives. A major focus was therefore to characterise the search strategies used when making preferential choices, particularly when making decisions-by-description. Many heuristic strategies have been proposed to explain how consumers navigate such scenarios (Tversky, 1972; Simon, 1955; Alba and Marmorstein, 1987), ranging from lexicographic strategies (which choose options based on whichever is superior on a given attribute, such as price) to weighted additive strategies (where a subjective importance weight is assigned to each available attribute). Broadly, strategies vary in the extent to which they consider available information (Bettman et al., 1998), namely the number of attributes processed, the amount of information processed for each attribute, whether information is processed by alternative (i.e., multiple attributes studied for a single option) or by attribute (i.e., multiple alternatives considered on a given attribute) and whether strategies are compensatory (i.e., a good value on one attribute can compensate for a bad value on another) or non-compensatory. Whilst some strategies make use of an invariant set of preferences (Tversky et al., 1988), a prevailing view is that people often do not have well-defined preferences; instead constructing them on the spot based on properties of the task at hand and the constraints of their cognitive systems (for a review, see Bettman et al., 1998).

A related strand of experimental research has focused on characterising the situations where decision-makers reverse their preferences. Notable early studies showed that people often violate the assumptions of subjective expected utility theory when making decisions under risk (e.g., gambles) (Savage, 1954;

von Neumann and Morgenstern, 1944). For example, suppose you are informed about the outbreak of an unusual disease, and subsequently asked whether you would save $\frac{1}{3}$ of patients for certain, or invest in a program that has a $\frac{1}{3}$ probability of saving 100% of patients. With this framing, the majority ($\sim 72\%$) choose the first option, opting to save lives with high certainty. However, when the same question is asked but framed in the context of lives lost, the majority ($\sim 78\%$) choose the second option. When framed in terms of gains, people tend to be risk-averse, whereas the prospect of losses tends to elicit risk-seeking (Tversky and Kahneman, 1981).

Preferences for options described across multiple attributes may also reverse depending on the presence or absence of irrelevant alternatives (Huber et al., 1982; Simonson, 1989). For example, suppose a decision maker is asked to choose between options A and B, where A is high in quality and price and B is low on both. When a decoy option D is introduced that is similar to A but deficient on price, the probability that people choose A increases; this is known as the *attraction effect* (Huber et al., 1982). If the decoy is instead middling on both attributes between options A and B, then decision makers may be drawn to that option by virtue of the fact that it is a compromise (Simonson, 1989). When a decoy option is added that is competitive with option A — say, slightly deficient in quality but more competitive in price — this can reduce the choice share of A but not B; this is known as the *similarity effect* (Tversky, 1972). Context effects have been demonstrated among risky prospects (Herne, 1999; Soltani et al., 2012; Farmer et al., 2017), perceptual decisions (Choplin and Hummel, 2005; Trueblood et al., 2013), consumer choices (Huber et al., 1982; Noguchi and Stewart, 2014; Simonson and Tversky, 1992) and even slime mould (Latty and Beekman, 2011). These context effects breach the assumption that choices are independent from irrelevant alternatives (IIA) (amongst other assumptions of economic models), thereby emphasising the context-sensitive nature of preferential choice. They can even drive choices towards inferior prospects (Farmer et al., 2017), thereby also

violating value-maximising models of behaviour.

Starting in the 1990s with work by Tversky and Simonson (1993), a wide array of computational models have been proposed that aim to account for context effects and the factors that amplify and diminish them (see Busemeyer et al., 2019, for a full review). Perhaps some of the best-known are evidence accumulation models, such as Multi-alternative Decision Field Theory (Roe et al., 2001) and the Multi-alternative Linear Ballistic Accumulator (Trueblood et al., 2014), amongst others. Whilst their details are beyond the scope of this introduction (for a recent review, see Trueblood, 2021), these models are unified in that they assume that evidence (i.e., valuation) for each option accumulates over time until one reaches a decision threshold. Competitors are evaluated in terms of their ability to explain both choices and their response times.

More recently, the focus has shifted towards explaining seemingly irrelevant experimental manipulations that amplify or diminish context effects, such as the spatial ordering of options on the screen (Evans et al., 2021) or the *concreteness* of the attributes (for a recent review, see Spektor et al., 2021). To date, no model can account for all of these moderating factors (Spektor et al., 2021). Some have questioned the importance of studying context effects, given that they can be amplified and diminished by seemingly irrelevant changes in experimental design (Frederick et al., 2014).

An additional challenge is that context effects may disappear or reverse in more naturalistic contexts. For example, Frederick et al. (2014) found no evidence of the attraction effect when participants were given the opportunity to taste products or when options were presented perceptually (e.g., hotels of different quality were presented pictorially). A recent follow-up similarly found no evidence for the attraction effect when participants were asked to choose between movies carefully matched for similarity (Trendl et al., 2021). In these naturalistic settings, attributes of options must be retrieved from memory, which appears to drive important differences in preferential choices.

The highly-stylised presentation of options in certain experimental tasks may therefore limit their external validity, particularly because they diminish the role of long-term prior knowledge. In this thesis, I will argue that — to understand how they are made in naturalistic settings — we must examine how long-term knowledge is exploited during preferential decision-making. This interaction has been described in several related fields. For example, studies of category learning have shown how discriminative decisions can warp semantic knowledge (i.e., general knowledge about the world) (Davis and Love, 2010; Braunlich and Love, 2019) and how selective retrieval of related exemplars can bias decisions (Nosofsky and Palmeri, 1997; Giguere and Love, 2013; Hornsby and Love, 2014). Studies of reinforcement learning have shown how semantic knowledge may be abstracted from sequential experiences, such as the map of a well-explored maze (Tolman, 1948; Eichenbaum, 2017b), and how episodic knowledge can facilitate inferences of an action’s value in large state spaces (Gershman and Daw, 2017; Mnih et al., 2015). Elsewhere, studies of judgment have shown that when people make decisions from experience they tend to underweight the probability of rare events, perhaps because they selectively retrieve recent exemplars from memory when constructing their estimates (Hertwig et al., 2004).

Summary: Psychologists have made considerable progress understanding how people make isolated, stimulus-based choices with a fixed set of options. Preferential choices are highly context sensitive, depending on the attribute similarity with competing options. Despite this, questions remain about the generalisability of well-studied effects, such as context effects. Whilst comparatively uncommon, behavioural studies investigating the role of long-term memory in preferential choice has been of growing interest, and I aim to review the key findings in the bulk of this introduction.

1.3 A brief introduction to human memory systems

I now briefly introduce some key concepts about human memory, before discussing what is known about the role of long-term memory in preferential decisions. Memory allows humans to encode, store and retrieve prior information. A primary function of memory is to help us predict future events and thus can be considered as an adaptation to allow better decision-making (Dasgupta and Gershman, 2021; Biderman et al., 2020). Memory is an information processing system thought to comprise of a sensory processor, short-term, working and long-term memory (Tulving, 1972). Sensory processors like the eye allow for information from the outside world to be perceived by the agent. Short-term memory allows for knowledge to be stored in an active, readily available state, albeit at a limited capacity (Miller, 1956). Working memory allows for the active manipulation of information and helps to determine what is encoded and what is retrieved (Baddeley, 2007). Long-term memory allows for the storage of such information over longer time periods.

Long-term memory has been subdivided into declarative and non-declarative (or implicit) memory (Squire and Zola, 1996). Storage and retrieval of declarative memory is thought to occur above the level of conscious awareness. For example, a person may store or retrieve *episodic memories* of past experiences or events and they may also store or retrieve *semantic knowledge*, such as general knowledge about concepts or facts in the world (Tulving, 1972). Non-declarative or implicit memory is said to operate below the level of conscious awareness (Graf and Schacter, 1985). For example, speaking a language, tying one's shoes may or learning a conditioned response recruits *procedural memories* that are easier to perform than to articulate (Gershman and Daw, 2017). Implicit memories may also be subconsciously activated or primed by an environmental stimulus.

Associations are fundamental building blocks of memory. During asso-

ciative learning, one learns that a cue or behaviour predicts another stimulus. Learning of associations may occur across different memory systems depending on properties of the stimulus (Eichenbaum, 2017a; Mason et al., 2021), such as whether it is rewarding (e.g., food) (Pavlov, 1927) or unrewarding (Wimmer and Shohamy, 2012). Associative memories can be organised into higher level structures, according to their consecutive, temporal and hierarchical associations (Mandler, 2011; Eichenbaum, 2017a). These organisational structures allow for memories to be associated indirectly, allowing for higher-level knowledge to be abstracted. Mental representations can therefore be constructed from associations in memory to reflect attributes, features, properties or characteristics about the world and things in it. Computational cognitive methods often use ideas from linear algebra to approximate these representations and describe how they are learned (for a recent review, see Polyn, 2022).

Insights about the organisation and interaction of memory systems can be revealed by studying memory retrieval. During list recall tasks, where lists of words are memorised and then recalled, participants tend to be more accurate at recalling items experienced first and last; these primacy and recency effects are thought to characterise episodic rehearsal and short-term memory, respectively (Murdock, 1962). During semantic fluency tasks, where items are listed from a category (e.g., animals), sequential retrievals are often close semantic associates (Bousfield and Sedgewick, 1944). Associative memories can also be highly context sensitive, with learned associations more likely to be recalled in the specific context in which they were encoded (Bouton, 1993, 2004). Memories can also be fallible, such as when the high semantic similarity between studied items causes related items to be erroneously retrieved during test (Deese, 1959; Roediger and McDermott, 1995). These characteristic patterns of human memory reach beyond simple list recall tasks. For example, in this thesis, I will demonstrate how associations across declarative memory influence retrieval of options and their attributes during preferential decision-making.

Summary: The storage and retrieval of information from memory improves our ability to make predictions and decisions about the future. Our prior experiences accumulate into long-term knowledge, such as episodic and semantic memories. This knowledge undoubtedly facilitates preferential decision-making in environments for which we have related prior experience. Yet, as discussed earlier, the interaction between declarative memory and preferential choices has received relatively little attention in the psychological literature. In the remainder of this chapter, I'll discuss what is known. In particular, I'll discuss how option attributes are retrieved from memory when they are not present, how options are retrieved from memory when they are not explicitly presented and how memory is updated following a choice. Whilst not an exhaustive review, I hope to provide a solid foundation for understanding the subsequent chapters.

1.4 Retrieval of option attributes from memory

To demonstrate context effects, attributes of choice options are often presented explicitly, in cue-attribute matrices. This affords researchers a high degree of experimental control, allowing them to model both the process (Bettman et al., 1998) and the outcome (Bröder, 2003; Rieskamp and Hoffrage, 1999) of preferential decisions. In naturalistic domains, people often make judgments on the basis of information that is not explicitly presented, such as *experience attributes* relating to the experience of sampling an option (e.g., taste) (Wright and Lynch, 1995).

1.4.1 Determinants of attribute retrieval

Information processing strategies have been a core focus in studies of stimulus-based choice (Tversky, 1972; Simon, 1956; Tversky and Simonson, 1993). During memory-based choice, it seems that people often rely on attribute-wise

processing (Lynch and Srull, 1982). Biehal and Chakravarti (1982) asked participants to study four fictitious toothpaste brands, either to choose or to prepare for questions later. When asked to freely recall information about the studied brands, participants in the choice condition tended to recall attributes of options (e.g., price), whereas those in the learning condition recalled more brand-based information. Interestingly, those in the choice condition recalled less information overall, indicating that studying options with the goal of making a choice directs attention selectively towards certain attributes of options.

The nature of the choice and the decision strategy used can also affect what is retrieved. For example, many processing strategies benefit from alternatives that can be aligned and thus used for comparison. Aligned attributes therefore tend to be better remembered and also receive comparatively larger weights than nonaligned attributes (Gentner and Markman, 1997; Lindemann and Markman, 1996; Markman and Medin, 1995). Attributes may also be retrieved to the extent to which they support a given decision. For example, Lynch et al. (1988) asked participants to choose between two previously-rated alternatives, for which prior ratings of quality and price were either diagnostic (i.e., better on both ratings) or non-diagnostic (i.e., better on one but not the other). When later tested on their memory for prior ratings and the option attributes, participants tended to recall choice attributes better for cases in which prior ratings were non-diagnostic and recalled prior ratings better for cases in which they were diagnostic during choice. Thus, attributes are more likely to be recruited when they support the decision-making process.

Of course, a great deal also depends on properties of memory itself. For example, by manipulating the depth at which participants encoded information about televisions (using a retroactive interference paradigm), Feldman and Lynch (1988) found that participants who studied option attributes with less subsequent interference were more likely to recall and subsequently use those attributes to make a subsequent choice. Thus much depends on the accessibility of knowledge at the time of decision. Memory retrieval also depends

on how knowledge is structured (Raaijmakers and Shiffrin, 1981; Howard and Kahana, 2002). Using a free recall paradigm, Biehal and Chakravarti (1982) found that participants were primarily biased towards retrieving information about brands of studied products, leading them to suggest that “consumer memory for product information is primarily brand organized”. However this likely differs depending on how options are encoded. For example, when choice options were organised by brand during encoding, people tended to recall them in an alternative-wise way, whereas when they were organised by attributes during encoding, they tended to recall options with similar attributes (Bettman and Kakkar, 1977).

1.4.2 Conceptual representation of options

So far we have discussed how extrinsically presented information is retrieved to support choices. But what about the knowledge we gain through our direct experience? Such knowledge is not directly observable, meaning that it must be inferred. One approach is to have participants explicitly organise stimuli into categories. For example, when asked to freely sort food items, emerging categories indicated that people defaulted to taxonomic (e.g., fruit and meat) and — to a lesser degree — script organisations (e.g., breakfast foods) (Ross and Murphy, 1999; Murphy and Ross, 1999). A challenge with this approach is that it may be biased by the stimuli present in the task. For example, free-sorting of larger set sizes indicated that semantic classifications tended to be more thematic as the sets became larger (Lawson et al., 2017).

Another approach is to learn approximations of mental representations through people’s behaviour. For example, distributed models of semantic memory have been used to approximate representations of objects and concepts (Landauer and Dumais, 1997; Shepard, 1974; Yee et al., 2018). Typically, these models are trained to explain the co-occurrences of words in people’s language use. Words are represented using low-dimensional vector representations and items that are nearby in vector space are considered more

similar. Words may become less similar in semantically meaningful ways as they become more distant along a vector direction (Mikolov et al., 2013). Numerous vector space models have been proposed, but they are united in their ability to successfully predict human responses, such as judgments of semantic similarity (Landauer and Dumais, 1997; Jones and Mewhort, 2007), word categorisation (Laham, 2000) and word *gist* (Griffiths et al., 2007).

Whilst initially used to approximate semantic representations of objects and concepts, vector space models can also explain conceptual representation of options in naturalistic domains. Bhatia and Stewart (2018) used a popular vector space model (known as *word2vec*) as input to several well-known cognitive models of preferential choice. This semantic space was hypothesised to capture mental representation of options in domains where attributes are not presented explicitly, such as when choosing foods or movies by images. Results showed that these computational models could accurately predict preferential choices out-of-sample, indicating that they may reflect the mental organisation of options.

Summary: When attributes of options are not presented explicitly, they must be retrieved from memory. These so-called memory-based choices can therefore be influenced by properties of memory, such as the accessibility of attributes in memory or the nature by which memory for attributes is organised. The mental organisation of option attributes has been approximated using distributed models of semantic memory, with promising results. In Chapter 2, we introduce a method for estimating conceptual representations of options directly from their choices, by observing the co-occurrences of products in shopper’s baskets.

1.5 Retrieval of options from memory

As with attributes, choice options may not be presented explicitly when making preferential choices in the real-world. For example, when shopping for

groceries online, one must retrieve ideas of what to buy from memory before searching for them. This was first identified by early marketing research on *consideration set construction*, which showed that recall of brands in the absence of explicit choice sets could be predicted by marketing variables, such as advertising and provenance (Nedungadi, 1990; Shocker et al., 1991) (for a review, see Roberts and Lattin, 1997). These studies rightly emphasised the importance of memory retrieval in cases where choice sets are large or not presented explicitly. However, I focus on the problem of option retrieval here, which is broader than recall of brands typically studied in marketing.

1.5.1 Option generation

Open-ended choices require a consideration set to be generated from memory before a selection can be made. One key determinant of what is retrieved is the strength of the association between options and the current choice context in memory (Keller and Ho, 1988)¹. For example, brands more strongly associated with a cue category during encoding (e.g., “soft drinks”) were more likely to be retrieved for consideration and indicated as a preference in a later test (e.g., “which soft drink would you most like to consume?”) (Posavac et al., 2001). Thus, options that have a strong *episodic association* with a choice context may have a high probability of being considered.

As well as episodic associations, options may also be retrieved according to their semantic association with the choice context (Bhatia, 2019). For example, when asked to list foods one would most like to consume, participants tend to recall items in semantically related clusters (Bhatia, 2019). This clustering of responses could be predicted by distributed models of semantic memory (such as those discussed in the previous section), further supporting the claim that these models capture how options are represented and retrieved in naturalistic domains. These models feature prominently within this thesis.

¹Note that other determinants of option generation have been proposed, such as cognitive control and ideation fluency (Del Missier et al., 2015). However, these proposals are less relevant to long-term memory and thus not reviewed here

The quality of generated options may depend on one’s experience with the domain. For example, Raab and Johnson (2007) found that experienced Handball players generated a greater diversity of suggested plays when watching videotapes of games and retrieved better options first, as marked by expert judges. Analyses of chess moves proposed by experienced players indicated a higher overall quality of initially generated options (Klein et al., 1995). Experts may therefore be able to perform deep searches of possibilities that are contingent on the first one being considered, in ways analogous to the tree-search algorithms used by expert game-playing AI (Silver et al., 2016).

1.5.2 Option selection

Once a consideration set has been generated, one must then decide on an option. During their analyses of suggested plays of Handball players, Johnson and Raab (2003) argued that an adequate heuristic strategy is to *take-the-first* option that comes to mind. This is because the first option generated was often the best. Analyses revealed that participants would have had a higher accuracy if they had only selected this option. This is consistent with “less-is-more” effects, which have also been demonstrated in attribute-wise processing of stimulus-based judgments (Goldstein and Gigerenzer, 2002) (though see Parpart et al., 2018). Taking the first may similarly work in preferential domains, particularly when one is experienced or under time pressure.

In preferential choice tasks, retrieved options may be evaluated according to their subjective value. For example, using the “remember-and-decide” task, in which participants learn to associate snacks with positions on a screen, Gluth and colleagues (Gluth et al., 2015a; Mechera-Ostrovsky and Gluth, 2018; Kraemer et al., 2021; Weilbacher et al., 2020) showed that participants exhibited a bias towards choosing remembered snack positions during a preferential choice task, but tended to reject remembered options if they had a very low subjective value, as indicated by initial ratings. This suggests that remembered options are compared to a reference subjective value before being selected. In

a choice-directed recall task — where participants were asked to study a list of snacks and then recall them in order of preference — Aka and Bhatia (2021) showed that the retrieval order was best explained by a model that incorporated subjective valuations with memory-based determinants of list-recall, such as recency and semantic associations. Similarly, when modelling answers to open-ended questions, such as “what is your favourite fast-food restaurant?”, choices were best predicted by a model that integrated semantic retrieval processes with subjective valuations (Zhang et al., 2021). These findings indicate that subjective valuation of generated options may interact with retrieved options to additionally influence their probability of being chosen.

Summary: In open-ended tasks, where the space of potential options is less well-defined, people must retrieve options from memory. The probability of a retrieval depends on how strongly options are associated with the cue in memory. Retrieved options may be selected according to a heuristic (e.g., “take-the-first”) or selected according to their subjective value. In Chapter 3, I discuss how long-term memory systems such as episodic and semantic knowledge guide sequential consumer choices in a real-world open-ended decision-making task; online grocery shopping.

1.6 Exploiting past experience

In laboratory decision-making tasks, the explicit presentation of novel options and their attributes deliberately minimises the contribution of prior experience and thus memory. However, decision makers in naturalistic domains like grocery stores often have prior experience with options that can help facilitate inferences about their properties and subjective value. Typically, this takes the form of episodic memories for past events (Tulving, 1972). Indeed, some have argued that the role of memory is to facilitate value-based decision making (Biderman et al., 2020) and to reduce the complexity of these computations (Dasgupta and Gershman, 2021).

1.6.1 Recognition, familiarity and fluency

Research from the heuristics and biases program (Gigerenzer and Todd, 1999) has shown that decision makers tend to choose options that they have more prior experience with. Arguably the simplest of these shortcuts is the *recognition heuristic*, which states that a higher value (on a given criterion) will be ascribed to whichever of two options is recognised. Recognition knowledge is thought to arrive early on the mental stage and thus may support choices when computational resources or prior experience are limited (Gigerenzer and Goldstein, 1996). Whilst initially thought to be a non-compensatory strategy — meaning that no other information can reverse a choice determined by recognition — studies of naturalistic consumer choice have suggested that recognition is used alongside other learned knowledge about a brand when it is available (Oeusoonthornwattana and Shanks, 2010).

When both options are recognised, choice options may be preferred to the extent that they are familiar; this is known as the *familiarity heuristic*. For example, mere exposure to choice options tends to increase liking (Zajonc, 1968) and has been demonstrated across numerous domains, including words, paintings, faces and sounds (Zajonc, 2001). However, familiarity may play a more subtle role when choice options must be retrieved from memory. For example, Gluth et al. (2015b) asked participants to memorise snack items to locations on a screen and later asked them to recall which locations contained snacks they would most like to consume. Responses were indicative of a familiarity bias, in that participants tended to recall snacks in locations that were presented more times during encoding. But not always. In particular, participants rejected the remembered option when it had a lower subjective value (as measured by an auction-task performed beforehand). Negative subjective valuations may therefore override familiarity biases when options must be retrieved from memory.

The fluency by which options are retrieved from memory can boost the subjective valuation of options. Namely, if two options are recognised but one

is more fluently retrieved from memory, then this option tends to be preferred against a given criterion; this is known as the *fluency heuristic* (Schooler and Hertwig, 2005). For example, Hertwig et al. (2008) showed that cities that were recognised more quickly accurately tracked judgments about their population size. However, fluency in domains with less ecological validity, such as judgments of company annual revenues, was a less accurate predictor of the correct answer. Fluency therefore appears sensitive to statistical regularities observed in the world, which highlights how prior experience can influence judgments in naturalistic domains.

1.6.2 Decisions from experience

Decisions may differ depending on whether they are based on explicit descriptions or prior experience; this is known as the *description-experience gap* (Hertwig et al., 2004; Wulff et al., 2018). For example, given access to explicit descriptions of the likelihood of events, such as the probability it will rain, people tend to overweight the probability of rare events, leading to a general aversion to risky choice (Kahneman and Tversky, 1979). In contrast, when people only have access to samples of prior experience, such as history of past weather, risky events tend to be underweighted relative to their actual occurrence (Hertwig et al., 2004). A proposed explanation for this gap is that decisions from experience are susceptible to the *recency effect*, or the tendency to retrieve recent exemplars from memory (Hogarth and Einhorn, 1992). In particular, rare events are less likely to be experienced recently, biasing subjective estimates of their likelihood downwards (Hertwig et al., 2004).

Episodic knowledge of one's past choices within a domain can be used to support inferences. For example, Scheibehenne et al. (2015) showed that participants could accurately estimate the market prices of wines when they had chosen similar wines during training. More recently, Jarecki and Rieskamp (2022) showed that subjective choices of novel options (pens and snacks) were better predicted by an exemplar comparison process than more tra-

ditional attribute-value comparison theories. Interestingly, these exemplar-based strategies were originally described by models of category learning (Nosofsky, 1984), which categorise based on their proximity to stored exemplars in memory. Such models have an attractive explanatory power, in that they formalise the ongoing interaction between memory and choice when making repeated decisions from experience.

Episodic knowledge from distinct events can also be integrated to guide value-based inferences. For example, having experienced that A is more valuable than B and that B is more valuable than C, A is often preferred over C, despite no prior experience of this combination (Dusek and Eichenbaum, 1997; Heckers et al., 2004). Similarly, when two stimuli are first associated in the absence of reward (e.g., through *sensory preconditioning*), the later association of one stimulus to a reward can increase preference for the unrewarded associate (Wimmer and Shohamy, 2012). Indeed, such associations may only need to be imagined for valuations to be inferred, such as imagining one’s preference for “tea-jelly” or “avocado-milkshake” (Barron et al., 2013). After choosing between two options, experiencing a reward can cause the nonchosen option to be ascribed the inverse value Biderman and Shohamy (2021), suggesting that deliberation binds choices in memory. Thus, value can propagate through episodic experience to influence the probability of choosing novel, but related options.

1.6.3 Sequential sampling from memory

During the decision making process, prior experience is thought to accumulate sequentially over time. This process can be formalised using sequential sampling models, such as drift-diffusion models (Ratcliff, 1978). Whilst initially used to describe perceptual decision-making (such as the direction of moving dots, Britten et al., 1992), adaptations have subsequently been proposed to account for preferential choice, such as risky gambles (Busemeyer and Townsend, 1993; Usher and McClelland, 2004), multi-attribute (Trueblood et al., 2014;

Bhatia, 2013) and value-based choice (Krajbich et al., 2010; Krajbich and Rangel, 2011). One proposed explanation for the close correspondence between perceptual and preferential decisions (Dutilh and Rieskamp, 2016) is that both depend on the sequential retrieval of evidence from memory (Shadlen and Shohamy, 2016). In particular, memory circuits may integrate evidence about the options (e.g., attribute values) and compute their subjective value, even when the attributes themselves are presented explicitly. As these value estimates are sequentially updated, evidence towards options grows, helping choices to become increasingly representative of one’s preferences over time.

Taken to the extreme, the theory of Decision by Sampling (Stewart et al., 2006) proposes that subjective valuations are constructed online by sequentially comparing attribute values drawn from memory and those available in the surrounding context (e.g., from other options available). By keeping a tally of the number of ordinal comparisons that favour a target option, one can generate a relative utility for each option. A consequence of this theory is that people do not possess underlying psychoeconomic scales, merely the memory of attribute values for related options, which are used during the online comparison process. Despite its simplicity, decision by sampling can account for several well-known economic phenomenon (Stewart et al., 2006), such as temporal discounting and the overestimation of rare events, in addition to multi-attribute context effects (Ronayne and Brown, 2017).

1.6.4 Reinforcement learning

Reinforcement learning (RL) is a framework that describes how agents learn to ascribe value to options through direct experience. Computational models of RL describe how knowledge about action values or the world can be extracted from many experiences to drive choice. Recent advances have driven several artificial intelligence breakthroughs (Mnih et al., 2015; Silver et al., 2016, 2017) and earlier “model-free” RL models exhibit remarkable overlap with neural firing patterns in the brain (Schultz et al., 1997; O’Doherty et al., 2003). More

recent advances have endowed these models with more deliberative, “model-based” evaluation (Daw et al., 2005; Dolan and Dayan, 2013), allowing them to adapt more easily to novel tasks. Cognitively, the processes of reward learning and knowledge abstraction map closely on to procedural and semantic memory and distinctions between model-free and model-based RL align with procedural and declarative memory systems in the brain (Eichenbaum and Cohen, 2004; Poldrack et al., 2001).

Typically, value estimates of model-free RL models are updated incrementally, leading to a decaying influence of past experiences. However, as we have discussed above, people may sample episodic knowledge in domains where they have prior experience, causing their behaviour to depart from model-free RL. For example, related episodic knowledge may be cued by the choice context, such as inferring the payoff of a gamble based on which room it was in (Bornstein and Norman, 2017). Reminders of prior choices can similarly bias our actions (Bornstein et al., 2017). These insights have inspired computational analogues, such as replay of recent experience during online training Mnih et al. (2015). Episodic knowledge can therefore help to guide value-based decisions in contexts for which we have prior experience.

When learning about an environment through repeated decisions, one may be drawn to novel options. For example, in non-stationary environments where the options are changing (like supermarkets introducing new products), it makes sense to periodically *explore* new options in addition to *exploiting* known favourites. Seeking novel options in these cases brings an opportunity to discover potentially valuable outcomes, and decision-makers appear biased towards such behaviour in certain learning environments (Wittmann et al., 2008). Attributes may be used to infer the value of novel options via function-based generalisation, helping to guide explorations towards options that are similar to previously rewarding outcomes (Stojić et al., 2020). For example, a study of takeaway purchases showed that consumers tended to explore dissimilar restaurants after a bad experience (measured through a negative rating)

or transition to similar restaurants after a positive experience (Schulz et al., 2019). Thus, exploration of new options appears to depend on both episodic memory of what was tried in the past as well as semantic knowledge to guide exploration towards options that share features with previously valuable outcomes.

Summary: Naturalistic choices rarely exist in isolation and therefore may be influenced by related past experience. This may manifest in heuristics or biases, such as the tendency to prefer options that are more familiar. In many circumstances, we must choose options to learn about their outcomes. Our success in these tasks depends on our ability to encode and retrieve related episodes from memory and generalise from those experiences. Choice may also depend on more fundamental properties of human memory, such as the bias towards retrieving recent exemplars. In Chapter 2, I discuss how knowledge can be abstracted from experience with preferential choices. In Chapter 3, I discuss how such experience may guide sequential preferential choices. In Chapter 4, I discuss how subjective preferences can be learnt over time in the absence of a clear extrinsic reward signal.

1.7 Choice-supportive memory biases

Following a choice, evidence from that past experience may be integrated into memory. For example, as discussed, prior experience of options may be stored in episodic memory to support inferences in familiar domains. Interestingly, the mere act of making a choice may skew memory in favour of that option. Such choice-supportive memory biases can be particularly pronounced in preferential domains, where choices are not followed by extrinsic feedback (Nakao et al., 2012).

1.7.1 Choice-supportive misremembering

One such bias is the tendency to retroactively ascribe positive attributes to a chosen option and/or to demote attributes of a foregone option (Lind et al., 2017). For example, a consumer may downplay the faults of a car they had purchased and emphasise the negative qualities of a car they had eschewed.

The facts of a decision — such as the attributes of chosen and nonchosen options — can be distorted to be supportive of a past choice². For example, Svenson et al. (2009) presented participants with facts about patients, such as the expected survival time, and asked them to prioritise a patient for surgery. After deciding and then performing a distractor task, participants tended to exaggerate positive characteristics of their choice and downweight negative characteristics. Memory of previously experienced attributes can therefore be distorted to favour one's prior choices.

People may also misattribute values to the wrong option when their memory is tested after a delay, so as to justify their past choice. For example, using a decision-by-description design, Mather et al. (2000) tested participants' recall of attributes belonging to prior choices of job-candidate, blind-date and roommate. Results revealed that — following a distractor task — participants tended to recall positive attributes of the foregone option as belonging to their choice, and vice versa. Such memory distortions also occurred for choices participants incorrectly believed they had made, either because they misrecalled (Henkel and Mather, 2007) or they were tricked by the experimenter (Henkel and Mather, 2007; Pärnamets et al., 2015). This suggests that misattribution errors occur during retrieval, rather than encoding.

Participants may also falsely remember attributes of their choices in ways that are supportive of their decision. For example, choice-supportive attributes

²The distortion of facts after a decision are distinct from subjective distortion of the attractiveness of options during the decision-making process. These are better characterised by processing bottlenecks during decision-making, rather than any influence of memory per se, and thus are discussed elsewhere (Svenson and Benthorn, 1992; Holyoak and Simon, 1999; Russo et al., 1996)

may be recognised, despite having never been presented before (Mather et al., 2000). In some cases, false memories of a choice may elicit evaluations and behaviour that are supportive of that false memory. For example, false memories of loving asparagus as a child can lead to higher appreciation of that food and a higher willingness to pay (Laney et al., 2008). False memories of becoming ill after eating a food can diminish their liking for it (Bernstein and Loftus, 2009). In sum, these findings show how choice-supportive biases may distort recollection of past choices and their attributes.

1.7.2 Choice-induced preference change

Choosing freely between options — in the absence of extrinsic feedback — may also shape subjective valuation of options so as to increase their favourability in subsequent ratings. In a classic demonstration of this, Brehm (1956) asked participants to rate food items by their desirability and — in a subsequent phase — choose freely between pairs, before rating them again. Results showed that ratings in the final phase increased towards chosen options and decreased towards nonchosen options, particularly when options were rated more closely in the initial phase.

The original design of the free-choice paradigm has been criticised, in that ratings may have been noisy and thus not reflective of participants' true preference (Chen and Risen, 2010; Izuma and Murayama, 2013). However, follow-up studies — such as an experiment in which participants chose between anonymous boxes only to have the contents revealed after the choice — revealed that choice-induced preference change does occur (Sharot et al., 2010; Alos-Ferrer and Shi, 2012; Koster et al., 2015; Schonberg et al., 2014; Miyagi et al., 2017; Akaishi et al., 2014; Nakao et al., 2016; Vinckier et al., 2019) and can be long-lasting (Sharot et al., 2009). Some have proposed moderating factors. For example, preference change may depend on whether one remembers the choice, suggesting a dependence on episodic memory (Salti et al., 2014, though see Lieberman et al. 2001).

Explanations for this post-choice *spread of alternatives* have varied. One explanation is that choosing between equally rated options causes cognitive dissonance (Festinger, 1957), which can be reduced by making the chosen alternative more desirable and the eschewed alternative less desirable. An alternative — self-perception theory (Brehm, 1956) — states that people retrospectively infer their preferences by observing their choices. Indeed, Brehm (1956) showed that post-choice ratings changed most for alternatives that were more comparable prior to the choice, which is consistent with both accounts. However, subsequent studies found evidence of asymmetric preference change for chosen and nonchosen options, which was particularly pronounced in depressive patients (Miyagi et al., 2017). More recent accounts have therefore proposed that option values are self-reinforced by choices (Akaishi et al., 2014; Hornsby et al., 2020; Chammat et al., 2017), which allows for asymmetric updating of option values for chosen and non-chosen options (Miyagi et al., 2017).

Summary: Following a preferential choice, memory may be updated so as to be supportive. This can distort one’s recollection for the features of that choice, such as wrongly attributing negative attributes to nonchosen options. Preferences may correspondingly change following a choice, such that a forced-choice between two equally-rated options can spread subjective value in favour of the choice. It is unclear whether preferences also increase towards novel but *related* options; this is a core focus in Chapter 4, where I propose an account of preference learning and decision making in multi-attribute space.

1.8 Research question

Laboratory studies of preferential choice have made considerable progress in describing how decision makers navigate fixed sets of options. However, there is arguably a disconnect between preferential choices made in the lab and those made in everyday life. In particular, we often have declarative knowledge,

gained through related prior experience, that helps to guide our inferences about the attributes and subjective value of different options when we make decisions in everyday life. How does prior experience — particularly represented across long-term, declarative memory — influence preferential choices? How does it help us to make sense of environments where there are innumerable options to choose from? These are the main overarching questions addressed in my thesis.

The interaction between declarative memory and preferential choice is unpacked more precisely over the following chapters. In Chapter 2, I address how abstract semantic knowledge may be learnt through direct experience of options, helping us to ascribe meaning and infer properties of options. In Chapter 3, I evaluate the contribution of episodic and semantic knowledge when making sequential, memory-based choices; here choices depend on the retrieval of options from memory and are therefore influenced by associations across declarative systems. In Chapter 4, I ask how subjective preferences for options are gained through direct experience, particularly how they are inferred and updated in option spaces that are large, multidimensional and overlapping. Throughout the thesis, I predominantly study consumer choices whilst drawing inspiration from literature on learning and memory. To do this, I use a mixture of behavioural experiments, computational models and analyses of large datasets of real-world consumer choices.

1.9 Dissertation outline

How do consumers conceptualise the countless products available to them? In Chapter 2 (Hornsby and Love, 2020), I gain insights about this semantic organisation by examining the choices of real British supermarket consumers. Using a topic model to abstract high-level topics from consumer purchases, I find that learnt topics are meaningful and centre primarily around goals (such as “stir fry” and “food for tonight”). The psychological validity of the topic

labels and the topics themselves are confirmed in two large-scale behavioural experiments conducted with retail industry experts and real consumers, respectively. The goal-oriented nature of these topics suggests that consumers abstract meaning from products through their experience to support future action. Moreover, I find that individual differences in consumers' topics predicted demographic information, such as age and gender. This suggests that choice trajectories reveal conceptual organisation and may also give rise to it.

How are declarative systems leveraged when making choices in large option spaces? In Chapter 3 (Hornsby and Love, 2022), I explore how conceptual representations influence the decision of what to choose next, when the decision depends on options first being retrieved from memory. When shopping for groceries online, consumers could feasibly search for products in any order. Yet, using a large dataset of over 100,000 online grocery shops, we find that choices can be predicted by their similarity with the prior purchase, suggesting that choices cue the retrieval of subsequent options in memory. Importantly, we develop representations of episodic, semantic and hierarchical semantic knowledge and find that a combination of all three explain choices and their response times. This suggests that consumers use their past choice to cue multiple sources of declarative knowledge when deciding what to choose next. Finally, we find that the type of knowledge people relied on predicted the type of errors they made; as hypothesised, more episodic retrievals predicted fewer items being forgotten and more semantic retrievals predicted more items being added to one's basket that they didn't otherwise need. Thus, when choices depend on the retrieval of options, much depends on how options are associated in memory.

If people represent options within large conceptual spaces, how are preferences and represented within them? In Chapter 4 (Hornsby et al., 2020), I explore how subjective preferences may be learned and used in real-world domains, where there are countless options and no extrinsic reinforcements to learn from. Building on previous findings showing that supermarket consumers

tend to explore less the more they exploit the same product, we propose that people use their past choices as the basis by which to update their subjective preferences. I propose a computational model of this process and simulate it to show how it could give rise to strong subjective preferences in the absence of extrinsic reinforcement. Because it represents preferences over attributes of choices, it is able to infer the subjective value of any option by virtue of its similarity to the current preference. In two large-scale behavioural experiments, I validate several key predictions of our model; that people prefer unseen options by virtue of their similarity with the prior choice and that this extends to domains where people have strong prior preferences (e.g., political candidates). Preferential choice is therefore cast as a learning process that seeks to align choices and preferences within attribute space in order to maintain coherency.

Bibliography

- Aka, A. and Bhatia, S. (2021). What I like is what I remember: Memory modulation and preferential choice. *Journal of Experimental Psychology: General*, 150(10):2175–2184.
- Akaishi, R., Umeda, K., Nagase, A., and Sakai, K. (2014). Autonomous mechanism of internal choice estimate underlies decision inertia. *Neuron*, 81(1):195–206.
- Alba, J. W. and Marmorstein, H. (1987). The effects of frequency knowledge on consumer decision making. *Journal of Consumer Research*, 14(1):14–25.
- Alos-Ferrer, C. and Shi, F. (2012). Choice-induced preference change and the free-choice paradigm: A clarification. *SSRN*.
- Baddeley, A. (2007). *Working memory, thought, and action*, volume 45. OuP Oxford.
- Barron, H. C., Dolan, R. J., and Behrens, T. E. J. (2013). Online evaluation of

- novel choices by simultaneous representation of multiple memories. *Nature neuroscience*, 16(10):1492–1498.
- Bentham, J. (1789). *The Principles of Morals and Legislation*. Prometheus Books.
- Bernstein, D. M. and Loftus, E. F. (2009). The Consequences of False Memories for Food Preferences and Choices. *Perspectives on Psychological Science*, 4(2):135–139.
- Bettman, J. R. (1979). Memory Factors in Consumer Choice: A Review. *Journal of Marketing*, 43(2):37–53.
- Bettman, J. R. and Kakkar, P. (1977). Effects of information presentation format on consumer information acquisition strategies. *Journal of Consumer Research*, 3(4):233–240.
- Bettman, J. R., Luce, M. F., and Payne, J. W. (1998). Constructive Consumer Choice Processes. *Journal of Consumer Research*, 25(3):187–217.
- Bhatia, S. (2013). Associations and the accumulation of preference. *Psychological Review*, 120(3):522–543.
- Bhatia, S. (2019). Semantic processes in preferential decision making. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 45(4):627–640.
- Bhatia, S. and Stewart, N. (2018). Naturalistic multiattribute choice. *Cognition*, 179:71–88.
- Biderman, N., Bakkour, A., and Shohamy, D. (2020). What Are Memories For? The Hippocampus Bridges Past Experience with Future Decisions. *Trends in Cognitive Sciences*, 24(7):542–556.
- Biderman, N. and Shohamy, D. (2021). Memory and decision making interact to shape the value of unchosen options. *Nature communications*, 12(1):1–12.

- Biehal, G. and Chakravarti, D. (1982). Information-presentation format and learning goals as determinants of consumers' memory retrieval and choice processes. *Journal of Consumer Research*, 8(4):431–441.
- Bornstein, A. M., Khaw, M. W., Shohamy, D., and Daw, N. D. (2017). Reminders of past choices bias decisions for reward in humans. *Nature Communications*, 8(1):15958.
- Bornstein, A. M. and Norman, K. A. (2017). Reinstated episodic context guides sampling-based decisions for reward. *Nature Neuroscience*, 20(7):997–1003.
- Bousfield, W. A. and Sedgewick, C. H. W. (1944). An Analysis of Sequences of Restricted Associative Responses. *The Journal of General Psychology*, 30(2):149–165.
- Bouton, M. E. (1993). Context, time, and memory retrieval in the interference paradigms of pavlovian learning. *Psychological bulletin*, 114(1):80.
- Bouton, M. E. (2004). Context and behavioral processes in extinction. *Learning & memory*, 11(5):485–494.
- Braunlich, K. and Love, B. C. (2019). Occipitotemporal Representations Reflect Individual Differences in Conceptual Knowledge. *Journal of Experimental Psychology. General*, 148(7):1192–1203.
- Brehm, J. W. (1956). Postdecision changes in the desirability of alternatives. *The Journal of Abnormal and Social Psychology*, 52(3):384–389.
- Britten, K. H., Shadlen, M. N., Newsome, W. T., and Movshon, J. A. (1992). The analysis of visual motion: A comparison of neuronal and psychophysical performance. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, 12(12):4745–4765.

- Bröder, A. (2003). Decision making with the "adaptive toolbox": Influence of environmental structure, intelligence, and working memory load. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 29(4):611–625.
- Busemeyer, J. R., Gluth, S., Rieskamp, J., and Turner, B. M. (2019). Cognitive and neural bases of multi-attribute, multi-alternative, value-based decisions. *Trends in Cognitive Sciences*, 23(3):251–263.
- Busemeyer, J. R. and Rieskamp, J. (2014). Psychological research and theories on preferential choice. In *Handbook of Choice Modelling*. Edward Elgar Publishing, Cheltenham, UK.
- Busemeyer, J. R. and Townsend, J. T. (1993). Decision field theory: A dynamic-cognitive approach to decision making in an uncertain environment. *Psychological Review*, 100(3):432–459.
- Chammat, M., Karoui, I. E., Allali, S., Hagège, J., Lehongre, K., Hasboun, D., Baulac, M., Epelbaum, S., Michon, A., Dubois, B., Navarro, V., Salti, M., and Naccache, L. (2017). Cognitive dissonance resolution depends on episodic memory. *Scientific Reports*, 7:41320.
- Chen, M. K. and Risen, J. L. (2010). How choice affects and reflects preferences: Revisiting the free-choice paradigm. *Journal of Personality and Social Psychology*, 99:573–594.
- Choplin, J. M. and Hummel, J. E. (2005). Comparison-induced decoy effects. *Memory & Cognition*, 33(2):332–343.
- Dasgupta, I. and Gershman, S. J. (2021). Memory as a Computational Resource. *Trends in Cognitive Sciences*, 25(3):240–251.
- Davis, T. and Love, B. C. (2010). Memory for category information is idealized through contrast with competing options. *Psychological Science*, 21(2):234–242.

- Daw, N. D., Niv, Y., and Dayan, P. (2005). Uncertainty-based competition between prefrontal and dorsolateral striatal systems for behavioral control. *Nature Neuroscience*, 8(12):1704–1711.
- Deese, J. (1959). On the prediction of occurrence of particular verbal intrusions in immediate recall. *Journal of Experimental Psychology*, 58(1):17–22.
- Del Missier, F., Visentini, M., and Mäntylä, T. (2015). Option generation in decision making: Ideation beyond memory retrieval. *Frontiers in psychology*, 5:1584.
- Dolan, R. J. and Dayan, P. (2013). Goals and Habits in the Brain. *Neuron*, 80(2):312–325.
- Dusek, J. A. and Eichenbaum, H. (1997). The hippocampus and memory for orderly stimulus relations. *Proceedings of the National Academy of Sciences of the United States of America*, 94(13):7109–7114.
- Dutilh, G. and Rieskamp, J. (2016). Comparing perceptual and preferential decision making. *Psychonomic Bulletin & Review*, 23(3):723–737.
- Eichenbaum, H. (2017a). Memory: organization and control. *Annual review of psychology*, 68:19.
- Eichenbaum, H. (2017b). The role of the hippocampus in navigation is memory. *Journal of Neurophysiology*, 117(4):1785–1796.
- Eichenbaum, H. and Cohen, N. J. (2004). *From Conditioning to Conscious Recollection: Memory Systems of the Brain*. Oxford University Press, USA.
- Evans, N. J., Holmes, W. R., Dasari, A., and Trueblood, J. S. (2021). The impact of presentation order on attraction and repulsion effects in decision-making. *Decision*, 8(1):36–54.
- Farmer, G. D., Warren, P. A., El-Deredy, W., and Howes, A. (2017). The Effect of Expected Value on Attraction Effect Preference Reversals. *Journal of Behavioral Decision Making*, 30(4):785–793.

- Feldman, J. M. and Lynch, J. G. (1988). Self-generated validity and other effects of measurement on belief, attitude, intention, and behavior. *Journal of Applied Psychology*, 73(3):421–435.
- Festinger, L. (1957). *A Theory of Cognitive Dissonance*. Stanford University Press.
- Frederick, S., Lee, L., and Baskin, E. (2014). The Limits of Attraction. *Journal of Marketing Research*, 51(4):487–507.
- Gentner, D. and Markman, A. B. (1997). Structure mapping in analogy and similarity. *American Psychologist*, 52(1):45–56.
- Gershman, S. J. and Daw, N. D. (2017). Reinforcement learning and episodic memory in humans and animals: An integrative framework. *Annual review of psychology*, 68:101–128.
- Gigerenzer, G. and Goldstein, D. G. (1996). Reasoning the fast and frugal way: Models of bounded rationality. *Psychological Review*, 103(4):650–669.
- Gigerenzer, G. and Todd, P. M. (1999). Fast and frugal heuristics: The adaptive toolbox. In *Simple Heuristics That Make Us Smart*, Evolution and Cognition, pages 3–34. Oxford University Press, New York, NY, US.
- Giguere, G. and Love, B. C. (2013). Limits in decision making arise from limits in memory retrieval. *Proceedings of the National Academy of Sciences*, 110(19):7613–7618.
- Gluth, S., Sommer, T., Rieskamp, J., and Büchel, C. (2015a). Effective Connectivity between Hippocampus and Ventromedial Prefrontal Cortex Controls Preferential Choices from Memory. *Neuron*, 86(4):1078–1090.
- Gluth, S., Sommer, T., Rieskamp, J., and Büchel, C. (2015b). Effective Connectivity between Hippocampus and Ventromedial Prefrontal Cortex Controls Preferential Choices from Memory. *Neuron*, 86(4):1078–1090.

- Goldstein, D. G. and Gigerenzer, G. (2002). Models of ecological rationality: The recognition heuristic. *Psychological Review*, 109(1):75–90.
- Graf, P. and Schacter, D. L. (1985). Implicit and explicit memory for new associations in normal and amnesic subjects. *Journal of Experimental Psychology: Learning, memory, and cognition*, 11(3):501.
- Griffiths, T. L., Steyvers, M., and Tenenbaum, J. B. (2007). Topics in semantic representation. *Psychological Review*, 114(2):211–244.
- Heckers, S., Zalesak, M., Weiss, A. P., Ditman, T., and Titone, D. (2004). Hippocampal activation during transitive inference in humans. *Hippocampus*, 14(2):153–162.
- Henkel, L. A. and Mather, M. (2007). Memory attributions for choices: How beliefs shape our memories. *Journal of Memory and Language*, 57(2):163–176.
- Herne, K. (1999). The Effects of Decoy Gambles on Individual Choice. *Experimental Economics*, 2(1):31–40.
- Hertwig, R. (2012). The psychology and rationality of decisions from experience. *Synthese*, 187(1):269–292.
- Hertwig, R., Barron, G., Weber, E. U., and Erev, I. (2004). Decisions from Experience and the Effect of Rare Events in Risky Choice. *Psychological Science*, 15(8):534–539.
- Hertwig, R., Herzog, S. M., Schooler, L. J., and Reimer, T. (2008). Fluency heuristic: A model of how the mind exploits a by-product of information retrieval. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 34(5):1191–1206.
- Hey, J. D. (2001). Does Repetition Improve Consistency? *Experimental Economics*, 4(1):5–54.

- Hogarth, R. M. and Einhorn, H. J. (1992). Order effects in belief updating: The belief-adjustment model. *Cognitive Psychology*, 24(1):1–55.
- Holyoak, K. J. and Simon, D. (1999). Bidirectional reasoning in decision making by constraint satisfaction. *Journal of Experimental Psychology: General*, 128(1):3–31.
- Hornsby, A. N., Evans, T., Riefer, P. S., Prior, R., and Love, B. C. (2020). Conceptual Organization is Revealed by Consumer Activity Patterns. *Computational Brain & Behavior*, 3(2):162–173.
- Hornsby, A. N. and Love, B. C. (2014). Improved classification of mammograms following idealized training. *Journal of Applied Research in Memory and Cognition*, 3(2):72–76.
- Hornsby, A. N. and Love, B. C. (2020). How decisions and the desire for coherency shape subjective preferences over time. *Cognition*, 200:104244.
- Hornsby, A. N. and Love, B. C. (2022). Sequential consumer choice as multi-cued retrieval. *Science Advances*, 8(8):eabl9754.
- Howard, M. W. and Kahana, M. J. (2002). A Distributed Representation of Temporal Context. *Journal of Mathematical Psychology*, 46(3):269–299.
- Huber, J., Payne, J. W., and Puto, C. (1982). Adding asymmetrically dominated alternatives: Violations of regularity and the similarity hypothesis. *Journal of Consumer Research*, 9(1):90–98.
- Izuma, K. and Murayama, K. (2013). Choice-Induced Preference Change in the Free-Choice Paradigm: A Critical Methodological Review. *Frontiers in Psychology*, 4.
- Jarecki, J. B. and Rieskamp, J. (2022). Comparing attribute-based and memory-based preferential choice. *DECISION*.

- Johnson, J. and Raab, M. (2003). Take the First: Option-Generation and Resulting Choices. *Organizational Behavior and Human Decision Processes*, 91:215–229.
- Jones, M. N. and Mewhort, D. J. K. (2007). Representing word meaning and order information in a composite holographic lexicon. *Psychological Review*, 114(1):1–37.
- Kahneman, D. and Tversky, A. (1979). Prospect Theory: An Analysis of Decision under Risk. *Econometrica*, 47(2):263–291.
- Keller, L. and Ho, J. (1988). Decision problem structuring: Generating options. *IEEE Transactions on Systems, Man, and Cybernetics*, 18(5):715–728.
- Klein, G., Wolf, S., Militello, L., and Zsombok, C. (1995). Characteristics of Skilled Option Generation in Chess. *Organizational Behavior and Human Decision Processes*, 62(1):63–69.
- Koster, R., Duzel, E., and Dolan, R. J. (2015). Action and valence modulate choice and choice-induced preference change. *PLOS ONE*, 10(3):1–10.
- Kraemer, P. M., Fontanesi, L., Spektor, M. S., and Gluth, S. (2021). Response time models separate single- and dual-process accounts of memory-based decisions. *Psychonomic Bulletin & Review*, 28(1):304–323.
- Krajbich, I., Armel, C., and Rangel, A. (2010). Visual fixations and the computation and comparison of value in simple choice. *Nature Neuroscience*, 13(10):1292–1298.
- Krajbich, I. and Rangel, A. (2011). Multialternative drift-diffusion model predicts the relationship between visual fixations and choice in value-based decisions. *Proceedings of the National Academy of Sciences of the United States of America*, 108(33):13852–13857.

- Laham, R. D. (2000). *Automated Content Assessment of Text Using Latent Semantic Analysis to Simulate Human Cognition*. PhD thesis, University of Colorado at Boulder, United States – Colorado.
- Landauer, T. K. and Dumais, S. T. (1997). A solution to Plato’s problem: The latent semantic analysis theory of acquisition, induction, and representation of knowledge. *Psychological Review*, 104(2):211–240.
- Laney, C., Morris, E. K., Bernstein, D. M., Wakefield, B. M., and Loftus, E. F. (2008). Asparagus, a Love Story. *Experimental Psychology*, 55(5):291–300.
- LaPiere, R. T. (1934). Attitudes vs. Actions. *Social Forces*, 13(2):230–237.
- Latty, T. and Beekman, M. (2011). Irrational decision-making in an amoeboid organism: Transitivity and context-dependent preferences. *Proceedings of the Royal Society B: Biological Sciences*, 278(1703):307–312.
- Lawson, R., Chang, F., and Wills, A. J. (2017). Free classification of large sets of everyday objects is more thematic than taxonomic. *Acta Psychologica*, 172:26–40.
- Lieberman, M. D., Ochsner, K. N., Gilbert, D. T., and Schacter, D. L. (2001). Do amnesics exhibit cognitive dissonance reduction? The role of explicit memory and attention in attitude change. *Psychological Science*, 12(2):135–140.
- Lind, M., Visentini, M., Mäntylä, T., and Del Missier, F. (2017). Choice-Supportive Misremembering: A New Taxonomy and Review. *Frontiers in Psychology*, 8.
- Lindemann, P. G. and Markman, A. B. (1996). Alignability and attribute importance in choice. In *Proceedings of the Eighteenth Annual Meeting of the Cognitive Science Society*, pages 358–363.

- Lynch, J. G., Marmorstein, H., and Weigold, M. F. (1988). Choices from Sets Including Remembered Brands: Use of Recalled Attributes and Prior Overall Evaluations. *Journal of Consumer Research*, 15(2):169–184.
- Lynch, J. G. and Srull, T. K. (1982). Memory and attentional factors in consumer choice: Concepts and research methods. *Journal of Consumer Research*, 9(1):18–37.
- Mandler, G. (2011). From association to organization. *Current directions in psychological science*, 20(4):232–235.
- Markman, A. B. and Medin, D. L. (1995). Similarity and Alignment in Choice. *Organizational Behavior and Human Decision Processes*, 63(2):117–130.
- Mason, A., Ludvig, E., and Madan, C. (2021). Conditioning and associative learning.
- Mather, M., Shafir, E., and Johnson, M. K. (2000). Misremembrance of Options Past: Source Monitoring and Choice. *Psychological Science*, 11(2):132–138.
- Mechera-Ostrovsky, T. and Gluth, S. (2018). Memory Beliefs Drive the Memory Bias on Value-based Decisions. *Scientific Reports*, 8(1):10592.
- Mikolov, T., Sutskever, I., Chen, K., Corrado, G. S., and Dean, J. (2013). Distributed Representations of Words and Phrases and their Compositionality. In *Advances in Neural Information Processing Systems*, volume 26. Curran Associates, Inc.
- Miller, G. A. (1956). The magical number seven, plus or minus two: Some limits on our capacity for processing information. *Psychological review*, 63(2):81.
- Miyagi, M., Miyatani, M., and Nakao, T. (2017). Relation between choice-induced preference change and depression. *PLoS ONE*, 12(6):1–14.

- Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A. A., Veness, J., Bellemare, M. G., Graves, A., Riedmiller, M., Fidjeland, A. K., Ostrovski, G., Petersen, S., Beattie, C., Sadik, A., Antonoglou, I., King, H., Kumaran, D., Wierstra, D., Legg, S., and Hassabis, D. (2015). Human-level control through deep reinforcement learning. *Nature*, 518(7540):529–533.
- Mosteller, F. and Noguee, P. (1951). An Experimental Measurement of Utility. *Journal of Political Economy*, 59(5):371–404.
- Murdock, B. B. (1962). The serial position effect of free recall. *Journal of experimental psychology*, 64(5):482.
- Murphy, G. L. and Ross, B. H. (1999). Induction with cross-classified categories. *Memory & Cognition*, 27(6):1024–1041.
- Nakao, T., Kanayama, N., Katahira, K., Odani, M., Ito, Y., Hirata, Y., Nasuno, R., Ozaki, H., Hiramoto, R., Miyatani, M., and Northoff, G. (2016). Post-response $\beta\gamma$ power predicts the degree of choice-based learning in internally guided decision-making. *Scientific Reports*, 6(August):1–9.
- Nakao, T., Ohira, H., and Northoff, G. (2012). Distinction between Externally vs. Internally Guided Decision-Making: Operational Differences, Meta-Analytical Comparisons and Their Theoretical Implications. *Frontiers in Neuroscience*, 6.
- Nedungadi, P. (1990). Recall and Consumer Consideration Sets: Influencing Choice without Altering Brand Evaluations. *Journal of Consumer Research*, 17(3):263–276.
- Noguchi, T. and Stewart, N. (2014). In the attraction, compromise, and similarity effects, alternatives are repeatedly compared in pairs on single dimensions. *Cognition*, 132(1):44–56.
- Nosofsky, R. M. (1984). Choice, similarity, and the context theory of clas-

- sification. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 10(1):104–114.
- Nosofsky, R. M. and Palmeri, T. J. (1997). An exemplar-based random walk model of speeded classification. *Psychological Review*, 104(2):266–300.
- O’Doherty, J. P., Dayan, P., Friston, K., Critchley, H., and Dolan, R. J. (2003). Temporal difference models and reward-related learning in the human brain. *Neuron*, 38(2):329–337.
- Oeusoonthornwattana, O. and Shanks, D. R. (2010). I like what I know: Is recognition a non-compensatory determiner of consumer choice? *Judgment and Decision Making*, 5(4):16.
- Pärnamets, P., Hall, L., and Johansson, P. (2015). Memory distortions resulting from a choice blindness task: 37th Annual Conference of the Cognitive Science Society : Mind, Technology, and Society. *Proceedings of the 37th Annual Conference of the Cognitive Science Society*, pages 1823–1828.
- Parpart, P., Jones, M., and Love, B. C. (2018). Heuristics as Bayesian inference under extreme priors. *Cognitive Psychology*, 102:127–144.
- Pavlov, P. I. (1927). Conditioned reflexes: an investigation of the physiological activity of the cerebral cortex. *Annals of neurosciences*, 17(3):136.
- Poldrack, R. A., Clark, J., Paré-Blagojev, E. J., Shohamy, D., Creso Moyano, J., Myers, C., and Gluck, M. A. (2001). Interactive memory systems in the human brain. *Nature*, 414(6863):546–550.
- Polyn, S. (2022). Attribute theories of memory. In *Oxford Handbook of Human Memory*. Oxford University Press.
- Posavac, S. S., Sanbonmatsu, D. M., Cronley, M. L., and Kardes, F. R. (2001). The Effects of Strengthening Category-Brand Associations on Consideration Set Composition and Purchase Intent in Memory-Based Choice. *ACR North American Advances*, NA-28.

- Raab, M. and Johnson, J. G. (2007). Expertise-based differences in search and option generation strategies. *Journal of Experimental Psychology: Applied*, pages 158–170.
- Raaijmakers, J. G. and Shiffrin, R. M. (1981). Search of associative memory. *Psychological Review*, 88(2):93–134.
- Ratcliff, R. (1978). A theory of memory retrieval. *Psychological Review*, 85(2):59–108.
- Rieskamp, J. and Hoffrage, U. (1999). When do people use simple heuristics, and how can we tell? In *Simple Heuristics That Make Us Smart*, Evolution and Cognition, pages 141–167. Oxford University Press, New York, NY, US.
- Roberts, J. H. and Lattin, J. M. (1997). Consideration: Review of Research and Prospects for Future Insights. *Journal of Marketing Research*, 34(3):406–410.
- Roe, R. M., Busemeyer, J. R., and Townsend, J. T. (2001). Multialternative decision field theory: A dynamic connectionist model of decision making. *Psychological Review*, 108(2):370–392.
- Roediger, H. L. and McDermott, K. B. (1995). Creating false memories: Remembering words not presented in lists. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 21(4):803–814.
- Ronayne, D. and Brown, G. D. A. (2017). Multi-attribute decision by sampling: An account of the attraction, compromise and similarity effects. *Journal of Mathematical Psychology*, 81:11–27.
- Ross, B. H. and Murphy, G. L. (1999). Food for thought: Cross-classification and category organization in a complex real-world domain. *Cognitive Psychology*, 38(4):495–553.

- Russo, J. E., Medvec, V. H., and Meloy, M. G. (1996). The Distortion of Information during Decisions. *Organizational Behavior and Human Decision Processes*, 66(1):102–110.
- Salti, M., Karoui, I. E., Maillet, M., and Naccache, L. (2014). Cognitive Dissonance Resolution Is Related to Episodic Memory. *PLOS ONE*, 9(9):e108579.
- Savage, L. J. (1954). *The Foundations of Statistics*. Wiley Publications in Statistics.
- Scheibehenne, B., von Helversen, B., and Rieskamp, J. (2015). Different strategies for evaluating consumer products: Attribute- and exemplar-based approaches compared. *Journal of Economic Psychology*, 46:39–50.
- Schonberg, T., Bakkour, A., M Hover, A., Mumford, J., Nagar, L., Perez, J., and A Poldrack, R. (2014). Changing value through cued approach: An automatic mechanism of behavior change. *Nature neuroscience*, 17.
- Schooler, L. J. and Hertwig, R. (2005). How forgetting aids heuristic inference. *Psychological Review*, 112(3):610–628.
- Schultz, W., Dayan, P., and Montague, P. R. (1997). A neural substrate of prediction and reward. *Science (New York, N.Y.)*.
- Schulz, E., Bhui, R., Love, B. C., Brier, B., Todd, M. T., and Gershman, S. J. (2019). Structured, uncertainty-driven exploration in real-world consumer choice. *Proceedings of the National Academy of Sciences*, 116(28):13903–13908.
- Shadlen, M. N. and Shohamy, D. (2016). Decision making and sequential sampling from memory. *Neuron*, 90(5):927–939.
- Sharot, T., De Martino, B., and Dolan, R. J. (2009). How choice reveals and shapes expected hedonic outcome. *Journal of Neuroscience*.

- Sharot, T., Velasquez, C. M., and Dolan, R. J. (2010). Do decisions shape preference?: Evidence from blind choice. *Psychological Science*, 21(9):1231–1235.
- Shepard, R. (1974). Representation of structure in similarity data: Problems and prospects. *Psychometrika*, 39(4):373–421.
- Shocker, A. D., Ben-Akiva, M., Boccara, B., and Nedungadi, P. (1991). Consideration Set Influences on Consumer Decision-Making and Choice: Issues, Models, and Suggestions. *Marketing Letters*, 2(3):181–197.
- Silver, D., Huang, A., Maddison, C. J., Guez, A., Sifre, L., van den Driessche, G., Schrittwieser, J., Antonoglou, I., Panneershelvam, V., Lanctot, M., Dieleman, S., Grewe, D., Nham, J., Kalchbrenner, N., Sutskever, I., Lillicrap, T., Leach, M., Kavukcuoglu, K., Graepel, T., and Hassabis, D. (2016). Mastering the game of Go with deep neural networks and tree search. *Nature*, 529(7587):484–489.
- Silver, D., Schrittwieser, J., Simonyan, K., Antonoglou, I., Huang, A., Guez, A., Hubert, T., Baker, L., Lai, M., Bolton, A., Chen, Y., Lillicrap, T., Hui, F., Sifre, L., Van Den Driessche, G., Graepel, T., and Hassabis, D. (2017). Mastering the game of Go without human knowledge. *Nature*.
- Simon, H. A. (1955). A Behavioral Model of Rational Choice. *The Quarterly Journal of Economics*, 69(1):99–118.
- Simon, H. A. (1956). Rational choice and the structure of the environment. *Psychological Review*, 63(2):129–138.
- Simonson, I. (1989). Choice Based on Reasons: The Case of Attraction and Compromise Effects. *Journal of Consumer Research*, 16(2):158–174.
- Simonson, I. and Tversky, A. (1992). Choice in context: Tradeoff contrast and extremeness aversion. *Journal of Marketing Research*, 29(3):281–295.

- Slovic, P., Fischhoff, B., and Lichtenstein, S. (1985). Characterizing perceived risk. In *Perilous Progress: Managing the Hazards of Technology*, pages 91–125. Westview.
- Soltani, A., Martino, B. D., and Camerer, C. (2012). A Range-Normalization Model of Context-Dependent Choice: A New Model and Evidence. *PLOS Computational Biology*, 8(7):e1002607.
- Spektor, M. S., Bhatia, S., and Gluth, S. (2021). The elusiveness of context effects in decision making. *Trends in Cognitive Sciences*, 25(10):843–854.
- Squire, L. R. and Zola, S. M. (1996). Structure and function of declarative and nondeclarative memory systems. *Proceedings of the National Academy of Sciences*, 93(24):13515–13522.
- Stewart, N., Chater, N., and Brown, G. D. A. (2006). Decision by sampling. *Cognitive Psychology*, 53(1):1–26.
- Stojić, H., Schulz, E., P. Analytis, P., and Speekenbrink, M. (2020). It’s new, but is it good? How generalization and uncertainty guide the exploration of novel options. *Journal of Experimental Psychology: General*, 149(10):1878–1907.
- Svenson, O. and Benthorn, L. J. (1992). Consolidation processes in decision making: Post-decision changes in attractiveness of alternatives. *Journal of Economic Psychology*, 13(2):315–327.
- Svenson, O., Salo, I., and Lindholm, T. (2009). Post-decision consolidation and distortion of facts.
- Tolman, E. C. (1948). Cognitive maps in rats and men. *Psychological Review*, 55(4):189–208.
- Trendl, A., Stewart, N., and Mullett, T. L. (2021). A zero attraction effect in naturalistic choice. *Decision*, 8(1):55.

- Trueblood, J. (2021). Theories of Context Effects in Multi-alternative, Multi-attribute Choice.
- Trueblood, J. S., Brown, S. D., and Heathcote, A. (2014). The multiattribute linear ballistic accumulator model of context effects in multialternative choice. *Psychological Review*, 121(2):179–205.
- Trueblood, J. S., Brown, S. D., Heathcote, A., and Busemeyer, J. R. (2013). Not Just for Consumers: Context Effects Are Fundamental to Decision Making. *Psychological Science*, 24(6):901–908.
- Tulving, E. (1972). Episodic and semantic memory. In *Organization of Memory*, pages xiii, 423–xiii, 423. Academic Press, Oxford, England.
- Tversky, A. (1972). Elimination by aspects: A theory of choice. *Psychological Review*, 79(4):281–299.
- Tversky, A. and Kahneman, D. (1981). The Framing of Decisions and the Psychology of Choice. *Science*, 211(4481):453–458.
- Tversky, A., Sattath, S., and Slovic, P. (1988). Contingent weighting in judgment and choice. *Psychological Review*, 95(3):371–384.
- Tversky, A. and Simonson, I. (1993). Context-Dependent Preferences. *Management Science*, 39(10):1179–1189.
- Usher, M. and McClelland, J. L. (2004). Loss Aversion and Inhibition in Dynamical Models of Multialternative Choice. *Psychological Review*, 111(3):757–769.
- Vinckier, F., Rigoux, L., Kurniawan, I. T., Hu, C., Bourgeois-Gironde, S., Daunizeau, J., and Pessiglione, M. (2019). Sour grapes and sweet victories: How actions shape preferences. *PLOS Computational Biology*, 15(1):1–24.
- von Neumann, J. and Morgenstern, O. (1944). *Theory of Games and Economic Behavior*. Princeton University Press.

- Warren, C., McGraw, A. P., and Van Boven, L. (2011). Values and preferences: Defining preference construction. *WIREs Cognitive Science*, 2(2):193–205.
- Weber, E. U. and Johnson, E. J. (2006). Constructing Preferences from Memory. SSRN Scholarly Paper ID 1301075, Social Science Research Network, Rochester, NY.
- Weilbacher, R. A., Kraemer, P. M., and Gluth, S. (2020). The Reflection Effect in Memory-Based Decisions. *Psychological Science*, 31(11):1439–1451.
- Wimmer, G. E. and Shohamy, D. (2012). Preference by Association: How Memory Mechanisms in the Hippocampus Bias Decisions. *Science*, 338(6104):270–273.
- Wittmann, B. C., Daw, N. D., Seymour, B., and Dolan, R. J. (2008). Striatal Activity Underlies Novelty-Based Choice in Humans. *Neuron*, 58(6):967–973.
- Wright, A. and Lynch, J. (1995). Communication Effects of Advertising Versus Direct Experience When Both Search and Experience Attributes are Present. *Journal of Consumer Research*, 21:708–18.
- Wulff, D. U., Mergenthaler-Canseco, M., and Hertwig, R. (2018). A meta-analytic review of two modes of learning and the description-experience gap. *Psychological Bulletin*, 144(2):140–176.
- Yee, E., Jones, M. N., and McRae, K. (2018). Semantic Memory. In *Stevens' Handbook of Experimental Psychology and Cognitive Neuroscience*, pages 1–38. John Wiley & Sons, Ltd.
- Zajonc, R. (2001). Mere Exposure: A Gateway to the Subliminal. *Current Directions in Psychological Science*, 10(6):224–228.
- Zajonc, R. B. (1968). Attitudinal effects of mere exposure. *Journal of Personality and Social Psychology*, 9(2, Pt.2):1–27.

- Zhang, Z., Wang, S., Good, M., Hristova, S., Kayser, A. S., and Hsu, M. (2021). Retrieval-constrained valuation: Toward prediction of open-ended decisions. *Proceedings of the National Academy of Sciences*, 118(20).

Chapter 2

Conceptual organisation of large option spaces

2.1 Introduction

“You shall know a word by the
company it keeps”

John Rupert Firth

A synopsis of linguistic theory

Our everyday experiences shape the way we conceptualize and act in the world. Following this intuition, previous work using text corpora has proved useful in understanding the nature of language and human concepts (Andrews et al., 2009; Deerwester et al., 1990). One appeal of this work is that text, such as from newspaper articles, reflects human behavior outside the laboratory. However, this text primarily serves a communicative role and is often scraped from curated sources, making it less reflective of real human activity.

In this chapter, I aim to build upon previous work from the text domain by analyzing real-world behavior from a broad section of the general population as they go about an everyday activity in relative anonymity, namely supermarket shopping. We apply techniques developed in computational linguistics to shopping data from nearly 1.3 million trips. Instead of words and

documents, our analyses are over products and shopping baskets. These analyses reveal that conceptualisations of products are organized around goals and interactions (e.g. tomatoes go well with vegetables in a salad), rather than their internal features (e.g. defining a tomato by the fact that it has seeds and is fleshy).

This work speaks to the relative importance of *intrinsic* and *extrinsic* features in concept representation. One way that people may reason about categories is to decompose them into intrinsic features or parts (Plato, 1973). On this view, a bird is an animal that typically has wings, feathers, a beak, and so on (Rosch and Mervis, 1975). However, extrinsic features are also critical for how humans organize concepts and come to understand the world, to the extent that some concepts may be solely defined by them (Barr and Caplan, 1987). For example, Wittgenstein (1967) asserted that the concept of game is undefinable. One might suggest that games are fun, but Russian Roulette is not fun and other activities that are fun are not games. Likewise, not all games are competitive (e.g., Ring Around the Rosie). Instead of defining game in terms of intrinsic features, one solution is to define game relationally – a game is simply something that is played (Markman and Stilwell, 2001). Human categories are therefore additionally sensitive to relationships and interactions with other concepts (Markman and Stilwell, 2001).

The importance of relations and interactions extends beyond abstract concepts. Many features of concrete concepts are extrinsic (Jones and Love, 2007). For example, whilst knowing that tomatoes are taxonomically related to fruits, people commonly associate them with other vegetables. Even for natural kinds, people commonly list extrinsic features for concepts (Jones and Love, 2007), such as noting that birds eat worms. Meanings appear to update in light of extrinsic relationships. For example, people are more likely to judge a polar bear and a dog as similar after reading vignettes in which both played the same role in a relation, such as chasing some other animal (Jones and Love, 2007). Likewise, merely sharing a thematic relationship, such as a man and a

tie (e.g., wears), makes the linked concepts more similar (Schank and Abelson, 1977; Wisniewski and Bassok, 1999; Jones and Love, 2007).

When concepts are defined in terms of other concepts, what moors or grounds our concepts to the physical world we inhabit (Harnad, 1990)? One proposed solution is that some concepts are embodied (Barsalou, 2008). For example, the action of hammering may be grounded to related motor programs and associated perceptions, linking the body, mind, and physical world. Indeed, neuroscientific evidence has shown that comprehension of language is tightly coupled with the neural regions associated with action and perception (Pickering and Garrod, 2013). A computational model developed by Mitchell et al. (2008) was able to accurately predict the neural activity elicited by a noun by considering the co-occurrence of that noun with action verbs in a large-text corpus. In effect, the action verbs, for which elicited neural activity was known, provided a grounding or bases for representing associated nouns.

These corpus models, such as *Latent Semantic Analysis*, use the co-occurrence of words within some context (e.g., a document) to learn lower dimensional, vector representations of word concepts (Deerwester et al., 1990). Like the reviewed psychological research (Jones and Love, 2007), words need not directly co-occur with one another to become more similar, but need only occur in similar contexts. Although LSA has enjoyed numerous successes, cases in which its representations diverge with those of humans has prompted further model development (Wandmacher et al., 2008).

One subsequent proposal, *Latent Dirichlet Allocation (LDA)*, is a probabilistic approach in which documents are generated according to a mixture of probabilities over latent themes or topics (Blei et al., 2003). For example, LDA may find that the words ‘Parliament’ and ‘Prime Minister’ have a high probability of belonging to the same topic (e.g., ‘politics’). A passage about the Prime Minister visiting the Houses of Parliament would make this politics topic highly probable, though other topics would also be somewhat likely, such as a topic related to tourism (Big Ben is part of the Houses of Parliament).

The representations learned by topic models appear similar to the concepts that people use (Griffiths et al., 2007; Andrews et al., 2009). For example, topic modeling can predict subsequent words in a sentence, disambiguate word meanings, and extract the gist of a sentence (Griffiths et al., 2007). Related techniques find that word meanings extracted for text corpora reflect back that society’s gender stereotypes (Bolukbasi et al., 2016). These successes emphasize the importance of extrinsic roles and relationships.

People learn thematic relations by observing co-occurrence in events or situations (Estes et al., 2011). In corpus analysis, word co-occurrence in language is assumed to be a proxy for co-occurrence in the wild. However, this assumption may not always hold. For example, words can co-occur in language without being semantically related (e.g., *iceburg* \rightarrow *lettuce*). More generally, most spoken language is concerned with effective communication of relevant information (Grice, 1975), rather than providing a faithful record of object interactions. For example, in waiting to cross the street with a companion, one would never verbalize that the passing car drives on the road. Written language also tends to be curated. For example, journalists adhere to particular guidelines and aim to report on stories of interest to their readership. Whether it’s from natural language or otherwise, data that captures co-occurrence of events in-the-wild is best suited to evaluate the structure of people’s thematic representations.

An alternative dataset that may help to further evaluate the influence of extrinsic features on people’s representations is consumer retail data. Retail data are collected from consumers as they purchase products together in the same basket, analogous to how words group together in the same document (see Figure 2.1). While a person may be conscious not to voice every item they bought in their supermarket shop, one’s grocery receipt provides a faithful record of what they purchased in a supermarket visit. Importantly, this data is traceable to an individual, which contrasts with most corpora analyses, which tends to be based on language in newspapers and books (e.g., Griffiths

The figure shows two examples of input for a corpus analysis. On the left is a news article titled 'NEWS' with the headline 'MAN BITES DOG'. Below the headline is a small graphic with a red square, a teal square, and some placeholder text. On the right is a grocery receipt titled 'RECEIPT' with a list of items: APPLE, BANANA, BEAN SPROUTS, CHILI, and LIME. Below each example is a table with two columns: 'Item' and 'Count'.

Item	Count
Dog	1
Cat	0
Man	1
Bites	1
...	...

Item	Count
Chili	2
Lime	1
Milk	0
Banana	1
...	...

Figure 2.1: The input in a corpus analysis is typically item counts (i.e., word counts) within some context (e.g., a sentence or document). Analogously, products (akin to words) are organized into baskets (akin to sentences). One advantage of applying these analysis techniques to baskets is that, unlike natural language, meaning is unaffected by item order.

et al., 2007). Large scale analyses of grocery retail data is therefore well placed to evaluate the claim that individual differences in people’s experience of the world leads them to possess different thematic representations. In particular, it may help to supplement existing research investigating how people cross-classify food (Murphy and Ross, 1999; Ross and Murphy, 1999; Lawson et al., 2017; Blake, 2008), such as elucidating how regional and generational differences affect people’s thematic representations.

An additional benefit of using consumer purchasing data is that it suits the mathematical assumptions of topic models particularly well. For example, natural language researchers typically use their domain expertise to remove function or ‘stop’ words that have little semantic meaning (such as *the*, *of*, *and*). They may also ‘stem’ words to remove prefixes and suffixes of words that have similar semantic meaning (e.g., eat vs. eating). Moreover, the order of words in sentences can also make a big difference to sentence meaning (e.g., “dog bites man” vs. “man bites dog”). However, most standard implementations of topic models (based on the original algorithm by Blei et al., 2003) typically ignore word order, instead preferring to consider language as a “bag-

of-words” (for an alternative, see Huang and Wu, 2015). In contrast, for retail data captured in-store, there is no inherent order for products within a basket, nor a need to remove stop words or perform stemming.

If people’s thematic organization of concepts arises through their interaction with the environment, then it should be possible for a topic model to recover relevant representations of these through consumer purchasing patterns, as shown in Figure 2.2. Whilst earlier research has indicated that this is possible, none (to this author’s knowledge) have explicitly measured the likeness of learned topics to consumer’s mental representations (Iwata and Sawada, 2013; Iwata et al., 2009; Hruschka, 2014). Although people have been shown to default to a taxonomic organization (e.g., tomato → fruit) when asked to freely sort food items in the lab, the presence of a goal can lead to a thematic organization (e.g., tomato → salad) during decision-making (Murphy and Ross, 1999; Ross and Murphy, 1999). Because shopping is highly goal-directed, we hypothesized that the topics recovered by a topic model would reflect thematic organization. We tested these predictions using a large, anonymized dataset of 1,252,963 shopping baskets and 5,753 unique products, supplied by one of the UK’s largest supermarket retailers. After optimizing an LDA solution using fit statistics and checking for convergence¹, we labelled the 25 topics recovered by the model.

To foreshadow, we found that LDA recovered meaningful topics that were primarily goal-directed and thematic in nature. We confirmed the psychological reality of these topics in two human studies, one with judgments from retail experts and another involving typical consumers. Further support came from analyses showing that topics tied to a season varied sensibly in their prevalence over the calendar year (e.g., the Christmas topic was most prevalent in December). Overall, these results suggest that — contrary to early research on cross-classification of food (Murphy and Ross, 1999; Ross and Murphy, 1999) — thematic relations dominate representations of food. This is in line

¹More detail about the model fitting procedure can be found in the methods section



Figure 2.2: Latent Dirichlet Allocation (LDA) uncovers the higher-level product topics that can be viewed as generating the observed baskets purchased by consumers. LDA’s fit is driven by the co-occurrence pattern of products within baskets. In the solution, each product has a probability of occurring within each topic (shown on the left for apple). The colours illustrate which topic each product would have been labelled with if using the maximum product topic probability. Each basket is generated by a mixture of probabilities over the topics (shown on the right for this basket).

with more recent claims that thematic relations may be more numerous than taxonomic associations in people’s stored semantic network and may be more easily revealed when examined at scale (Estes et al., 2011; Lawson et al., 2017). Final analyses tested whether an individual’s shopping experience shaped their conceptual organization of the products. In support of this assertion, the rate at which an individual sampled topics (based on recent shopping history) predicted the individual’s age, gender, and geographic region. This suggests that individual differences in people’s experience can lead them to possess notably different thematic representations from each other. This is important, because it suggests that food-related themes discussed in the literature are likely a function of the participants’ individual experiences and culture.

2.2 Training a topic model with retail data

2.2.1 Method

Data

The topic model was developed on a random 0.1% sample of all grocery transactions that occurred in 2014 in one of the UK's largest supermarket chains. The transactions were filtered such that only relatively popular products selling $> 50,000$ units annually were kept. Moreover, data was filtered such that only large baskets containing ≥ 20 items were kept. Filtering was performed to ensure that LDA would have enough observations to learn meaningful topics. This is typical in LDA modelling (Yan et al., 2013) and is performed by the original LDA authors (Blei et al., 2003). After filtering, the final dataset contained 1,253,183 unique baskets and 5,753 unique products.

Items were modelled at the product code level. Concretely, there is a different product code for each distinct product in the supermarket. Small variations in that product (i.e., different sizes of the same t shirt) are not given separate codes however.

Note that — unlike traditional uses of LDA in NLP — we did not remove commonly occurring items from documents (i.e., ‘stop words’). Whilst natural language may contain ‘stop words’ (i.e., common words with little semantic meaning such as ‘the’), we did not believe grocery transactions to suffer from the same problem. In the retail case, purchasing popular products, such as milk, bananas and bread, may be informative, perhaps indicating that the consumer is stocking essential items. The basket data was fully anonymised for general research purposes so as to not be personally identifiable.

Model fit

In our experiments we applied Latent Dirichlet Allocation (LDA) to the data, using the machine learning library in Spark 1.6.0 (Apache Software Founda-

tion, 2016). We conducted a range of experiments to identify the optimal set of hyperparameters (including the number of topics k) and in each case monitored the training and test log-perplexity to ensure model convergence and generalization, respectively (see supplemental for further details).

The LDA solution with the lowest log-perplexity on held out data (i.e., best generalization) had 25 topics. Models were trained for a maximum of 500 epochs, used the Online Variational Bayes optimization algorithm with an $\alpha = 0.1$. The remaining hyperparameters were set to the package defaults.

2.2.2 Results and discussion

The topics recovered by LDA were coherent and readily labeled by the authors. Table 2.1 shows the top 5 products within each of 10 randomly selected topics, according to the relevancy metric. Topics tended to be organized along activity patterns and goals, ranging from specific (e.g., *Stir Fry*) to general in scope (e.g., *Cooking from scratch*). This therefore provides early support for the hypothesis that consumers primarily recruit thematic representations when conducting their grocery shop.

When proposing labels for the topics, the authors had access to the full item-topic relevancy matrix. In some cases, products outside of the top five were instrumental in determining the topic label. For example, mince pies (a popular dessert consumed during Christmas in the UK) were the seventh most relevant item within the *Christmas* topic and chicken korma was eighth most relevant within the *food for now* topic.

2.3 Evaluating topic labels with retail experts

To evaluate the appropriateness of the topic labels, we conducted a more detailed study. Specifically, a group of industry experts were asked to look through a sample of highly-ranking products from within 10 randomly-selected topics and confirm that the proposed labels were indeed representative of the

Table 2.1: Retailer-supplied product descriptions for the 5 most relevant products within each of the 10 surveyed topics. Note that the authors had access to the full product-topic relevancy matrix (see <https://osf.io/tsymx/>) when they labeled the topics. Brand names have been removed from this table for publication.

Topic	Description
Food for now	ITALIAN BEEF LASAGNE 450G
	ITAL CHICKEN & BACON PASTA BAKE 450G
	ITALIAN MACARONI CHEESE PASTA 450G
	ITAL SPAGHETTI CARBONARA 450G
	ITAL HAM & MUSHROOM TAGLIATELLE 450G
Summer salad	BUNCHED SPRING ONIONS 100G
	ICEBERG LETTUCE EACH
	WHOLE CUCUMBER EACH
	SALAD TOMATOES 6 PACK
	GROWING SALAD CRESS EACH
Stir fry	FRESH EGG NOODLES 375G
	VEGETABLE & BEANSPROUTS FRY 333G
	CHINESE STIR FRY BOWL 300G
	UNCLE BENS EXPRESS GOLDEN VEG RICE 250G
	BEANSPROUTS 370G
Afternoon tea	2 EGG CUSTARD TARTS 2X90G
	BRS/SKIMMED MLK 1.136L/2PINTS
	DANISH SLICED WHITE BREAD 400G
	MINHUMBUGS 200G
	BANANAS LOOSE
Loose fruit and veg	CARROTS LOOSE
	BANANAS LOOSE
	PARSNIPS LOOSE
	CONFERENCE PEARS LOOSE
	BROCCOLI LOOSE
Low calorie options	LIGHFRUITS YOGUR6X175G
	BRSKIMMED MILK 2.272L/4 PINTS
	LIGHYELLOW FRUIYOGUR6X175
	LIGHTOFFEE YOGUR175G
	LIGHLIMITED EDITION YOGHURT 165G
Cheapest option	EDAY VALUEBAKED BEANS IN TOMSAUCE 420G
	EVERYDAY VALUE HAM 364G
	EDAY VALUE MILK CHOCOLATE DIGESTIVES 300G
	EDAY VALUEPENNE 500G
	EVERYDAY VALUELOW FAFRUIYOG 4X125G
Cooking from scratch	COURGETTES LOOSE
	LOOSE BROWN ONIONS
	RED ONIONS LOOSE
	CARROTS LOOSE
	GARLIC EACH
Christmas	ORIGINAL CRISPS 190G
	SOUR CREAM & ONION CRISPS 190G
	BRUSSELS SPROUTS 500G
	PARSNIPS PACK 500G
	SAL& VINEGAR CRISPS 190G
Low maintenance cooking	PREPARED BABY SPROUTS 180G
	PREPARED CARROCAULIFLOWER & BROCCOLI 370G
	PREPARED TRAD SLICED RUNNER BEANS 185G
	PREPARED BROCCOLI FLORETS 240G
	TOPSIDE OF ROASTBEEF 85G

grouped products².

2.3.1 Method

Participants

Participants were recruited internally within the UK headquarters of dunnhumby (www.dunnhumby.com), a customer marketing company with over 29 years of experience working with grocery retailers and fast moving consumer goods (FMCG) brands. Employees were asked to participate via the company intranet and were not remunerated. Fifty-one participated in the study. Participants had a wide range of roles within the business, including data analysts, category experts, company directors and client leads. Of these, 56.86% were male. Participants were surveyed in early December 2016 and were blind to the purposes of this study. The Ethics Committee at the UCL Experimental Psychology department approved the methodology and all participants consented to participation through an online consent form at the beginning of the survey.

Materials

The study was hosted on an internal company server. Participants accessed the study via their web browsers and answered questions by clicking on the appropriate radio button with their cursor. The study was 1700 x 1300 pixels within the browser.

In each trial, participants were shown 10 product images (2 rows of 5) and accompanying product descriptions from a single topic. Images were 540 x 540 pixels each. Descriptions appeared below each image in size 12 font. The displayed products were the 10 with the highest *relevance*³. Product

²This same subset of 10 topics was considered in the empirical studies of retail experts (described in section 2.3) and typical consumers (described in section 2.5). Analyses was limited to a random subset of topics to reduce the overall survey time and thereby maximize the number of people that could take the respective surveys.

³See supplementary for more information about how relevance was calculated

descriptions and images were downloaded from the retailer’s website in late November 2016.

Design

All participants were asked to label the same 10 topics in a random order. The dependent measure was the proportion of times that participants selected the topic label originally proposed during the model development phase. This proportion was then compared against a random baseline, to check whether participants were responding non-randomly.

It was not feasible to survey participants about all 25 topics in the final LDA solution given constraints on employee time. Therefore, 10 topics were chosen from the original 25 to include in the survey.

Procedure

Participants were first briefed about the purpose of the experiment. After agreeing, they were then asked to label a group of products for 10 separate topics. Four possible labels were suggested using radio buttons. The order of the presented topics was randomized. One of the four labels was the ‘target’ label proposed by the authors whereas the other three were randomly selected from the remaining nine topics. After selecting a topic label, participants then confirmed their choice with a “Continue” button, before seeing the next set of products from a randomly-selected remaining topic. At the end of the study, participants were debriefed.

2.3.2 Results and discussion

When asked to select the appropriate topic label for a group of products from a list of four possible labels, 92.8% ($SE = 0.015$) of the 51 retail industry experts selected the same topic label as was originally proposed by the authors. A two-sided binomial test showed this to be significantly above chance ($p < .001$).




Figure 2.3: Proportion correct with standard error bars for the study on label agreement involving retail experts and the intruder study involving typical consumers. All proportions were significantly different ($p < .001$) than chance levels, 25.00% (1 of 4) and 16.67% (1 of 6), respectively.

Figure 2.3 shows the proportion of times participants agreed with the originally proposed topic label for each topic.

This high-level of accuracy from a group of experts, who were naive to our research program, indicates that the topics recovered from consumer activity patterns are meaningful. Disagreement regarding topic labels was primarily driven by conceptually similar topics (further details of this are available in the supplemental). For example, the most common error in labeling the *summer salad* topic was *cooking from scratch*. These errors are reasonable and are also consistent with the notion that baskets are generated by a mix of topics, as opposed to a single topic (see Figure 2.2).

2.4 Seasonal trends in topics

The results of the expert study discussed in section 2.3 suggested that the names given to the topics were reasonable. As further confirmation of this, we attempted to evaluate the appropriateness of names pertaining to seasonal events using historic data. Specifically, we identified 4 topics that were likely to have a highly seasonal popularity (*summer fruits*, *summer salad*, *Christmas* and *low calorie options*) and 4 ‘staple-food’ topics that we believed unlikely to vary as much over the year (*loose fruit and veg*, *Northern Ireland*, *quick to prepare meals* and *food for now*).

2.4.1 Method

Data

To evaluate seasonal trends in topic prevalence, the same data used in section 2.2.1 was used.

Analyses

To calculate the monthly prevalence of each topic, we hard-assigned each basket to belong to one topic, using the maximum topic probability. We then calculated an index indicating the relative popularity of a topic in a given month by calculating the proportion of baskets belonging to a given topic in a month and dividing it by the average topic probability for a given month across all topics.

2.4.2 Results

Figure 2.4 shows the popularity of several topics in each month of 2014 in one of the UK’s largest retailers. In line with our hypotheses, *summer fruits* and *summer salad* peaked in popularity during the summer months. Contrasting, baskets labelled with the *Christmas* topic peaked in popularity during Decem-

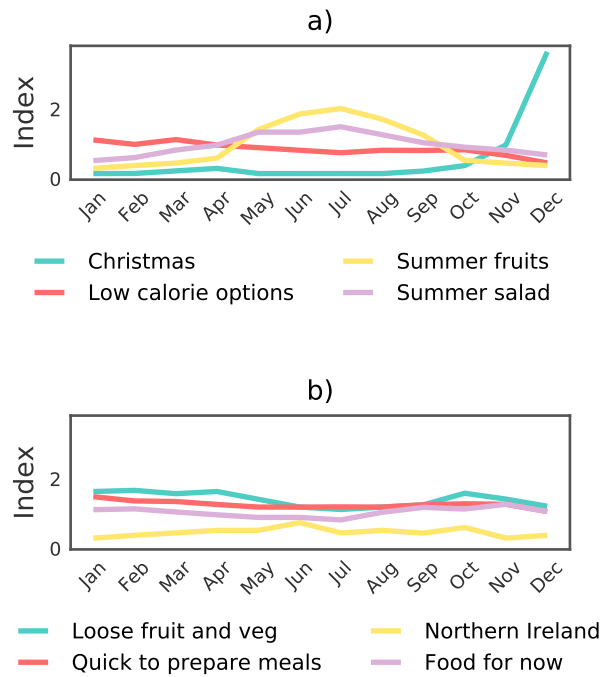


Figure 2.4: Topic prevalence varies by season. The proportion of baskets with a given topic label in each month of 2014, divided by the monthly mean average across all topics (i.e., index), is shown. a) Topics that should be seasonal peak at the expected time, such December for the Christmas topic. b) In contrast, topics for staple products vary less in prevalence over time.

ber and the surrounding winter months. *Low calorie options* appeared to peak in January and steadily decline to its lowest level of popularity in December. The ‘staple’ topics shown in Figure 2.4b appeared to vary considerably less over the year compared to the more seasonal topics. These results give further credence to the proposed topic labels and illuminate some seasonal variations in behavioural patterns that likely reflect time-dependent characteristics of people’s thematic representations. Similarly intuitive patterns were shown to occur during different days of the week, which are reported in more detail in the supplemental.

Results from the survey of retail experts (in section 2.3) and these seasonal analyses support the validity of the topic labels assigned by the experimenters. This therefore provides further credence to the hypothesis that people’s repre-

sentations are dominated by thematic categories during shopping. It is particularly exciting that this can be inferred using data collected from consumers activity patterns in-the-wild. By combining big behavioural datasets with computational modelling in this way, we are able to inferences about people’s semantic representations on a far greater scale than would be possible using standard experimental tasks (e.g., card-sorting or free association tasks) (e.g., Murphy and Ross, 1999). However, these analyses alone do not confirm the psychological reality of the representations found by this topic model alone. Specifically, there is still an open possibility that the topics recovered by LDA are only meaningful to commercial retail experts, and thus not real consumers.

2.5 Evaluating topic coherency with typical consumers

To understand whether the product representations identified by the topic model were meaningful to real consumers, we conducted a controlled experiment. Specifically, a large sample of British supermarket shoppers were shown a set of products; 5 of which were highly-ranked from within a topic and one that was an “intruder” product, randomly selected from one of the other topics. Participants were asked to select the intruding product. This paradigm is often used when evaluating the fit of computational models to human semantic memory (De Deyne et al., 2016). It was hypothesized that if topics were representative of the mental categories held by consumers, participants would be able to identify the intruding products significantly above chance levels.

2.5.1 Method

Participants

Participants were recruited using dunnhumby’s consumer survey panel; Shopper Thoughts (<https://shopperthoughts.com/>). Participants completed the

survey as part of a larger, monthly survey for 50 card loyalty points. The final sample consisted of 3,840 participants, of which 59.47% were female. The modal age group was 55-59 ($n = 501$) and 724 participants did not disclose their age. All participants were from England, Scotland or Wales with the majority of respondents based in central England ($n = 946$). Participants were surveyed during March 2017. The Ethics Committee at the UCL Experimental Psychology department approved the methodology and all participants consented to participation through an online consent form at the beginning of the survey.

Materials

The study was accessible via the web, after logging in to the survey platform. The study screen was 960 x 455 pixels. Product images and descriptions were the same as those described in section 2.3. Participants responded to the survey by clicking on a radio button next to the picture and product description of the item they believed to be the intruder.

Procedure

Participants were first informed that the purpose of the study was to help retailers group together products found in the supermarket. They were then informed about the study's procedure. After agreeing to participate, the sole trial started.

Participants were shown six images of products (two rows of three) alongside product descriptions. Five were the most relevant from within a topic and the other 'intruder' was the most relevant product from a randomly selected alternative topic⁴. Participants were asked to "spot the one that does not belong in the group" by clicking on the appropriate radio button underneath the image. Following their choice, participants were thanked, debriefed about the purpose of the research and remunerated immediately.

⁴More information about the ranking procedure used is available in the supplemental

Design

The dependent variable was the proportion of times that participants identified the intruder product. This proportion was then compared against a random baseline to assess whether participants were able to identify intruders significantly above chance levels.

Participants each completed one trial in which they were asked to identify the intruding product. Participants were randomly selected to see one of the 10 topics also used in the retail experts study. This ensured that comparisons between the two related studies were consistent.

2.5.2 Results and discussion

Of the 3840 British consumers surveyed, 74.1% ($SE = 0.007$) were able to correctly identify the intruder product. A two-sided binomial test showed this to be significantly above chance ($p < .001$). Figure 2.3 shows accuracy by topic.

One topic stands out for its below chance level of performance, *afternoon tea*. Participants were most likely (51.7% of the time) to incorrectly suggest that ‘mint humbugs’ was the intruder. One possibility for this poor classification accuracy is that participants did not have enough context to interpret them correctly. In the *afternoon tea* topic, the top 5 items were predominantly fresh and ‘staple’ foods (e.g., Milk, Bananas, Danish sliced white bread). Thus, seeing a packet of sweets (i.e., ‘humbugs’) among this fresh food may have appeared unusual. An analogous issue arises with the *low maintenance cooking* topic. Each topic is a probability distribution over thousands of products, so perhaps it is not surprising that a small sample of products could be ambiguous.

Another possibility that is more core to our theory is that individual differences in experience may help explain some of these confusions. For example, the poor classification of the *afternoon tea* topic may have been driven by the

fact that most British people no longer regularly engage in this ritualistic activity. If experience shapes people’s mental concepts, then we would expect representations of certain products to vary between demographics. Supporting this view, consumers from Northern Ireland had an average probability for the *Northern Ireland* topic 7.5× higher than the average across all regions. The fact that the model was able to recover such strong regional differences in consumers suggests that it should be sensitive to other individual differences in people’s experience of the world.

2.6 Classifying individual consumers by their experienced topics

If the topic model proposed in this paper is representative of people’s semantic categories, then it should also be able to uncover individual differences in their representations. To test this assertion, we used logistic regression to predict (5-fold cross-validation) self-reported age⁵, region⁶ and gender using consumers’ mean LDA probabilities over baskets as the predictors.

2.6.1 Method

Data

The feature set for the supervised models comprised of the training-set topic probabilities output by LDA, averaged at the customer level. These features were then used to predict customers’ self-reported age (discretized into 18-29, 30-44, 45-59 and 60+), region (binarized into England vs. regional (i.e., Scotland, Wales and Northern Ireland) and gender. These self-reported values had been gathered by the marketing panel described in section 2.5 during the last 3 years. The final modelling set contained data from 28,122 customers.

⁵Discretized into 18-29, 30-44, 45-59 and 60+

⁶Binarized into England vs. regional (i.e., Scotland, Wales and Northern Ireland)

Model

To find the best performing model, we performed a grid-search between λ values of 0.1 to 1.0. Model selection was performed using the average predictive performance over 5 cross-validation folds. Baselines were calculated by predicting the majority class in each fold.

The age model was evaluated according to the average classification accuracy across the four classes. The gender and region models were binary classification problems, and thus evaluated in terms of the Area Under the ROC curve (AUC).

2.6.2 Results

Results showed that the models were able to predict age with an accuracy of 48.51%, region with an accuracy of 58.34% and gender with an accuracy of 57.17%, considerably higher than the guessing baselines of 36.85%, 50.00%, and 50.00%, respectively.

2.7 General Discussion

Rather than being solely defined by intrinsic features (Plato, 1973), concepts gain meaning through their interaction in the real world. Support for this notion comes from laboratory studies demonstrating that object interactions alter how people conceptualize objects (Jones and Love, 2007) and from large-scale corpora analysis of text (e.g., newspaper articles) that extract meaning from word co-occurrence patterns. However, none of these previous investigations involve individuals engaging in unfiltered, goal-driven, real-world interactions with objects. Under such conditions, can meaningful conceptual organization be recovered from human activity patterns?

We tested this possibility by considering the shopping patterns of thousands of UK consumers. Using LDA, we found that the pattern of consumer purchases was highly revealing of people's conceptual organization of these

products. Topics ranged from specific and goal-driven (e.g., ingredients for a stir-fry) to very general (e.g., cooking from scratch). Interestingly, the topics tended to be goal-directed and situational, which is consistent with the notion that much of human conceptual knowledge is defined relationally and tailored to support action (Murphy and Ross, 1999; Ross and Murphy, 1999; Schank and Abelson, 1977). The situational nature of certain topics was reflected in their increasing prevalence during certain times of the year, such as the *Christmas* topic in December and the *Summer salad* topic in the Summer.

The psychological reality of the 25 LDA topics we found was confirmed by two studies, one involving retail experts and one involving everyday consumers. The experts, who were blind to the purposes of this research, agreed with our labeling of the topics. The novices were able to identify an intruder product among an array of products from the same topic. These results indicate that the topics uncovered by human activity patterns are both comprehensible and coherent.

If concepts gain meaning through the actions we take, then individual differences in experiences should be reflected in differences in conceptual organization. In support of this conjecture, topic prevalence varied across geographic regions. In our study of everyday consumers, poor performance for the topic *afternoon tea* may reflect that today's typical British consumer differs from past caricatures. Consistent with the idea that different types of people will have different topic experiences, we were able to predict basic demographic information about consumers from their topics mix (i.e., which topics best characterized their purchasing behavior). One avenue for future research is to develop, apply, and evaluate topic models in which individuals organize into higher-level groups that can vary in terms of topic prevalence or even topic composition.

Taxonomic and thematic cross-classifications of food are typically measured in free-sorting tasks, where participants must sort food items into groups (Ross and Murphy, 1999; Murphy, 2001). While originally suggesting that peo-

ple have a bias towards sorting food taxonomically, more recent, larger-scale sorting tasks have suggested that people have a thematic bias (Lawson et al., 2017), suggesting that taxonomic bias is an artifact of a small initial set size. The large-scale analyses presented here give further credence to this claim, which is notable, given that supermarkets tend to arrange food taxonomically. Another likely cause of thematic bias is that grocery retail data reflects more goal-directed behavior. For example, people may visit solely in order to purchase “food for now”, which emerged as a topic in our model. An outstanding question however is how people recruit these different representations over the course of a large shopping trip, as they complete several sub-goals. One possibility is that people recruit taxonomic and thematic representations hierarchically, using thematic representations to identify which ingredients to combine (e.g. for a salad) and taxonomic representations to identify the most suitable version of a given ingredient (e.g., best variety of tomato). Future research may wish to investigate this interaction in more detail.

What is clear is that conceptual organization is deeply tied to extrinsic relationships and that meaning can be seen as a byproduct of an element’s role within a larger system or web. Indeed, the insight behind Google’s PageRank algorithm is that web pages should be prioritized to the extent that they are central within a link graph (Page et al., 1998). Prior to PageRank, the exact same algorithm was developed in Psychology to explain why certain features of concepts are more central than others within a concept web (Love and Sisman, 1995). Whether the system is human or artificial or the domain involves natural language or shopping behavior, meaning can be inferred, and perhaps arises, from relations among elements embedded within a larger system.

Bibliography

Andrews, M., Vigliocco, G., and Vinson, D. (2009). Integrating Experiential and Distributional Data to Learn Semantic Representations. *Psychological*

- Review*, 116(3):463–498.
- Apache Software Foundation (2016). Spark.
- Barr, R. A. and Caplan, L. J. (1987). Category representations and their implications for category structure. *Memory and Cognition*, 15:397–418.
- Barsalou, L. W. (2008). Grounded Cognition. *Annual Review of Psychology*, 59(1):617–645.
- Blake, C. E. (2008). Individual differences in the conceptualization of food across eating contexts. *Food quality and preference*, 19:62–70.
- Blei, D. M., Ng, A. Y., and Jordan, M. I. (2003). Latent Dirichlet Allocation. *Journal of Machine Learning Research*, 3:993–1022.
- Bolukbasi, T., Chang, K.-W., Zou, J., Saligrama, V., and Kalai, A. (2016). Man is to computer programmer as woman is to homemaker? debiasing word embeddings. In *Proceedings of the 30th International Conference on Neural Information Processing Systems, NIPS’16*, pages 4356–4364, USA. Curran Associates Inc.
- De Deyne, S., Navarro, D., Perfors, A., and Storms, G. (2016). Structure at every scale: A semantic network account of the similarities between unrelated concepts. *Journal of Experimental Psychology: General*, 145:1228–1254.
- Deerwester, S., Dumais, S. T., Furnas, G. W., Landauer, T. K., and Harshman, R. (1990). Indexing by latent semantic analysis. *Journal of the American Society for Information Science*, 41(6):391–407.
- Estes, Z., Golonka, S., and Jones, L. L. (2011). Thematic thinking: The apprehension and consequences of thematic relations. In Ross, B. H., editor, *Advances in Research and Theory*, volume 54 of *Psychology of Learning and Motivation*, pages 249–294. Academic Press.

- Grice, H. P. (1975). Logic and conversation. In *Syntax and semantics*, pages 41–58. New York: Academic Press.
- Griffiths, T. L., Steyvers, M., and Tenenbaum, J. B. (2007). Topics in semantic representation. *Psychological Review*, 114(2):211–244.
- Harnad, S. (1990). The symbol grounding problem. *Physica D*, 42:335–346.
- Hruschka, H. (2014). Linking multi-category purchases to latent activities of shoppers: Analysing market baskets by topic models. University of Regensburg Working Papers in Business, Economics and Management Information Systems 482, University of Regensburg, Department of Economics.
- Huang, C. M. and Wu, C. Y. (2015). Effects of Word Assignment in LDA for News Topic Discovery. In *Proceedings - 2015 IEEE International Congress on Big Data, BigData Congress 2015*, pages 374–380.
- Iwata, T. and Sawada, H. (2013). Topic model for analyzing purchase data with price information. *Data Min. Knowl. Discov.*, 26(3):559–573.
- Iwata, T., Watanabe, S., Yamada, T., and Ueda, N. (2009). Topic tracking model for analyzing consumer purchase behavior. In *IJCAI-09 - Proceedings of the 21st International Joint Conference on Artificial Intelligence*, pages 1427–1432.
- Jones, M. and Love, B. C. (2007). Beyond common features: The role of roles in determining similarity. *Cognitive Psychology*, 55(3):196–231.
- Lawson, R., Chang, F., and Wills, A. J. (2017). Free classification of large sets of everyday objects is more thematic than taxonomic. *Acta Psychologica*, 172:26–40.
- Love, B. C. and Sloman, S. A. (1995). Mutability and the determinants of conceptual transformability. In *Proceedings of the 17th Annual Conference of the Cognitive Science Society*, pages 654–659, Mahwah, NJ. Lawrence Erlbaum Associates.

- Markman, A. B. and Stilwell, C. H. (2001). Role-governed categories. *Journal of Experimental and Theoretical Artificial Intelligence*, 13(4):329–358.
- Mitchell, T. M., Shinkareva, S. V., Carlson, A., Chang, K.-M., Malave, V. L., Mason, R. A., and Just, M. A. (2008). Predicting Human Brain Activity Associated with the Meanings of Nouns. *Science*, 320(5880):1191–1195.
- Murphy, G. L. (2001). Causes of taxonomic sorting by adults: a test of the thematic-to-taxonomic shift. *Psychonomic bulletin & review*, 8(4):834–839.
- Murphy, G. L. and Ross, B. H. (1999). Induction with cross-classified categories. *Memory & Cognition*, 27(6):1024–1041.
- Page, L., Brin, S., Motwani, R., and Winograd, T. (1998). The pagerank citation ranking: Bringing order to the web. In *Proceedings of the 7th International World Wide Web Conference*, pages 161–172, Brisbane, Australia.
- Pickering, M. and Garrod, S. (2013). An integrated theory of language production and comprehension. *Behavioral and Brain Sciences*, 36(4):329–347.
- Plato (1973). *Theaetetus*. Clarendon Press.
- Rosch, E. and Mervis, C. B. (1975). Family resemblances: Studies in the internal structure of categories. *Cognitive Psychology*, 7(4):573–605.
- Ross, B. H. and Murphy, G. L. (1999). Food for thought: Cross-classification and category organization in a complex real-world domain. *Cognitive Psychology*, 38(4):495–553.
- Schank, R. C. and Abelson, R. P. (1977). *Scripts, Plans, Goals and Understanding: an Inquiry into Human Knowledge Structures*. L. Erlbaum, Hillsdale, NJ.
- Wandmacher, T., Ovchinnikova, E., and Alexandrov, T. (2008). Does Latent Semantic Analysis Reflect Human Associations? In *Bridging the gap between*

- semantic theory and computational simulations: Proceedings of the ESSLLI 2008 Workshop on Distributional Semantics*, pages 63—70.
- Wisniewski, E. J. and Bassok, M. (1999). What makes a man similar to a tie? Stimulus compatibility with comparison and integration. *Cognitive Psychology*, 39:208–238.
- Wittgenstein, L. (1967). *Philosophical Investigations*, volume 17. Wiley-Blackwell.
- Yan, X., Guo, J., Lan, Y., and Cheng, X. (2013). A biterm topic model for short texts. *WWW '13 Proceedings of the 22nd international conference on World Wide Web*, pages 1445–1456.

Chapter 3

Option retrieval in sequential, open-ended tasks

“I know what I like,
And I like what I know”

Genesis, I Know What I Like

3.1 Introduction

Many studies of preferential choice have examined how people choose between a fixed menu of options (Glimcher and Rustichini, 2004; Busemeyer and Rieskamp, 2014; Rangel et al., 2008). Yet, in real-world tasks like online grocery shopping, the space of possible options is too large to be considered at once. Choices therefore depend on how options are retrieved from long-term memory (Keller and Ho, 1988; Kalis et al., 2013; Zhang et al., 2021; Bhatia, 2019). The task itself can provide the context to retrieve associated choice options. Here, we consider how a previous choice influences subsequent consumer choices in a goal-directed sequential decision task. Using computational models, we evaluate how different sources of knowledge influence choice, by decomposing associative memory into its constituent components.

Once items are retrieved from memory, they may cue subsequent re-

trievals, leading to complex sequential dynamics like semantic clustering. For example, when asked to name as many animals as possible, sequential retrievals tend to be semantically similar and faster when they are so (e.g., dog \rightarrow cat) (Bousfield and Sedgewick, 1944; Troyer et al., 1997; Gruenewald and Lockhead, 1980; Hills et al., 2012; Abbott et al., 2015). This sequential cuing of memory means that retrievals tend to cluster over time. Sequential consumer choices may also semantically cluster if they depend on the same retrieval mechanisms; we test this hypothesis here.

Retrieval is said to depend on the strength of associations in memory, although *association* is somewhat nebulous given that items can relate in different ways. For example, choosing tea could trigger childhood memories of enjoying it with cake, as it did for Proust (1913). Options that occur in the same episode could have a high probability of being retrieved; this was shown in the early experiments of memory (Ebbinghaus, 1913) and has since become a core prediction in models of memory search (Kahana, 2020).

Sequential choices could also be influenced by semantic similarity between items, such as their conceptual overlap. For example, whilst they may not be consumed in the same episode, purchasing chocolate could prompt the search for other chocolate bars, due to their shared features. Semantic space models have been shown to predict sequential retrievals in fluency tasks (Jones and Mewhort, 2007; Hills et al., 2012) and options generated to open-ended questions (Bhatia, 2019). Online shoppers may similarly retrieve products that are nearby in conceptual space when making sequential choices, given that they are not constrained by the physical layout of products in stores.

An often-cited feature of semantic memory is that people are sensitive to hierarchical relations between items. For example, responses tend to be slower when judging the correctness of statements like “apples are fruit” compared to “apples are produce” (Collins and Quillian, 1969). One might therefore expect online shoppers to retrieve items that are nearby within a structured hierarchy, such as purchasing fruit then vegetables. This seems particularly

likely during grocery shopping, given that stores tend to arrange products taxonomically in order to make them easier to locate (questions concerning whether hierarchical, semantic, and episodic knowledge are strictly separate systems from neurobiological or computational perspectives is orthogonal to our aims and claims).

We hypothesise that retrieval of options in sequential choice tasks depends on their similarity with the prior choice across different knowledge formats (visualised in Figure 3.1c). We test this by developing associative representations of these knowledge sources (using techniques inspired by those of Chapter 2), before evaluating whether sequential choices are better explained by one or a combination of these representations. We also hypothesise that individual differences may drive shoppers to attend to certain sources of knowledge more than others. For example, a shopper driven by episodic memories of breakfast might retrieve butter then bread, whereas those relying on hierarchical knowledge may retrieve butter with other dairy products, as they would in the supermarket. We operationalise these processes of associative retrieval and attention in a computational cognitive model and show that it can predict sequential consumer choices. After each retrieval, we suggest that consumers accept or reject possibilities according to their goals. For example, shoppers may consider whether retrieved options are suitable for breakfast. However, goals are not modelled or enumerated here, as we focus on the contribution of different knowledge systems during sequential option retrieval.

During online grocery shopping, consumers could feasibly search for products in any order. Yet, if options are retrieved according to their similarity with the prior choice, purchases should be non-random and predicted by their sequential similarity. We test this using a new dataset of over 5 million consumer choices. To foreshadow, results supported this hypothesis. Interestingly, representations of episodic, semantic and hierarchical knowledge explained unique variance when predicting sequential choices and their response times, supporting the idea that shoppers query multiple sources of long-term knowledge.

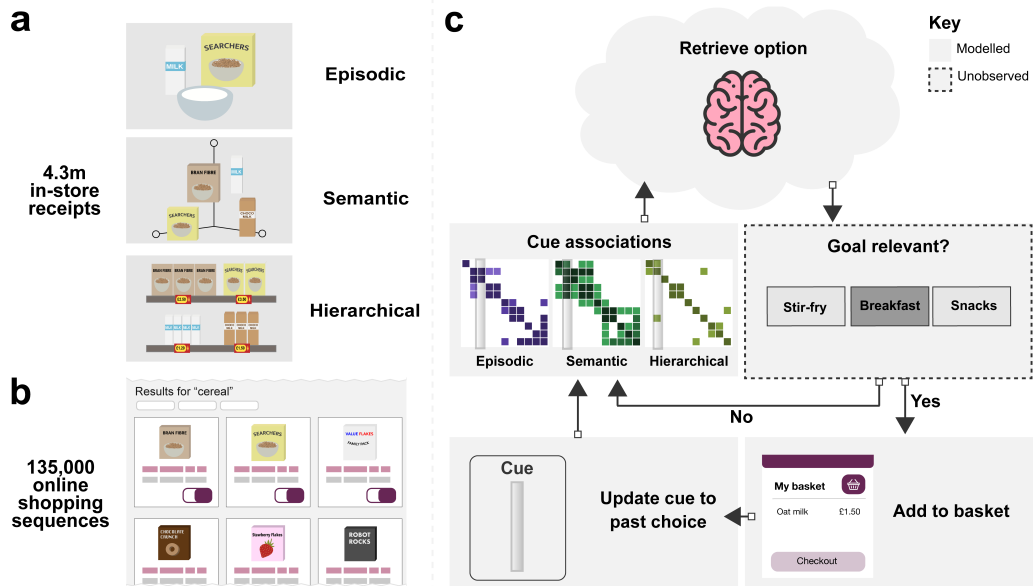


Figure 3.1: Deciding what to choose next when shopping for groceries online depends on cued retrieval from multiple knowledge sources. **a.**, We used 4.3m unordered, in-store receipts to build representations of episodic, semantic and hierarchical knowledge. **b.**, To model retrieval, we collected data from 135,000 shoppers as they sequentially searched for products on the website of one of the UK’s largest supermarket retailers. **c.**, Prior choices predict future ones, by virtue of their similarity according to different representational formats. Once an item is added to their basket, shoppers use this to cue matches from long-term memory. The stronger the match with this cue, the higher the probability an item will be retrieved (this may be attenuated by increased attention towards a particular representation). Retrieved items are checked against one’s internal goals. If the retrieval is goal-relevant, the shopper adds an appropriate item from the website and uses that item to cue associations. If not, a new option is retrieved and checked for goal relevance until one is accepted. Similar heuristic strategies have been used in models of option generation for single choices (Johnson and Raab, 2003; Klein et al., 1995). Once all goals are satisfied, the user checks out. Note that the goals and the goal-checking process is not modelled here.

Consumers retrieving options from episodic memory appeared less prone to subsequently forget products, whereas those attending to semantic knowledge were less likely to add items to their basket that they didn't otherwise need. Thus, individuals may recruit these systems to different extents, which may affect their ability to complete the task effectively.

3.2 Data

3.2.1 Clickstream data

Data capturing a sequence of clicks during a given shopping session is known as *clickstream data*. In this study, we used clickstream data collected by a large British retailer between 1st January 2015 and 31st March 2016. We used a random sample of visits resulting in a checkout during that period and only kept observations where a product was added to a shopper's basket. The data contained 5,238,469 choices from 132,146 unique visitors across 42,837 unique products (more information in Materials and Methods). By shopping online, all customers were required to participate in the loyalty scheme of the retailer and therefore consented to having their data used for research. In order to preserve user privacy, we removed all customer identifiers from the data and kept only a cryptographic hash of each visit ID. All analyses were in compliance with UCL's code of ethics

Shoppers were required to search sequentially for groceries to add to their virtual basket (the website is depicted in Figure C1). On average, they made 39.64 choices ($95\%CI = [39.49, 39.79]$). The landing page displayed a generic selection of "special offers" (e.g., discounted products), which was used relatively infrequently to purchase products ($\mu_{offers} = 1.55, 95\%CI = [1.53, 1.58]$). Shoppers tended to search for products using a search bar, which was located at the top of every page ($\mu_{searches} = 23.36, 95\%CI = [23.26, 23.46]$). They could also use a category drop-down by hovering the mouse over a hyperlink saying "Groceries" at the top of every page. This menu required users

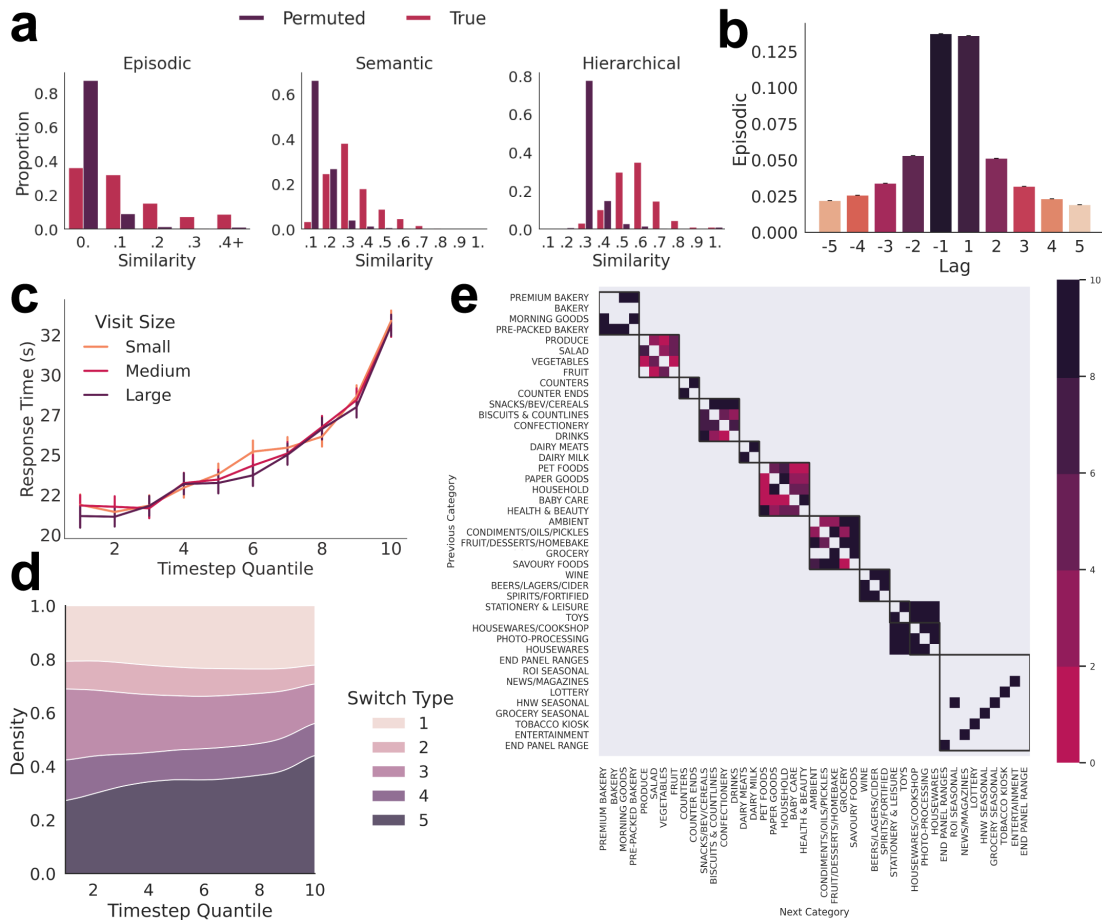


Figure 3.2: Consecutive purchases tend to be close episodic, semantic and hierarchical relations. **a.**, Choices are predicted by their similarity with the prior choice across each representation. Histograms show that the similarity between consecutive purchases (averaged for each visit) was higher compared to when the order of purchases was randomly permuted. (with 95% confidence intervals). **b.**, Sequential retrieval is like a ripple in semantic memory. Mean episodic similarity (with 95% confidence intervals) between the current product and those purchased recently is higher compared with products purchased later **c.**, Visitors slowed as they approached the end of their shopping trip. Mean response times (with 95% confidence intervals) as a function of timestep quantile (Small = 10 - 30 items, Medium = 31 - 49 items, Large = 50+ items). **d.**, Consumers make more between-category transitions (i.e., taxonomy level five) towards the end of their visit. Stacked density plots denoting the proportion of switches according to each level of the taxonomy as a function of the relative timestep. **e.**, Transitions between product groups at the fourth level of the hierarchy clustered into intuitive higher-order groupings that appear similar to those in the product taxonomy, suggesting that the taxonomy closely resembles how shoppers represent products during sequential choice. The $Lift - 1$ of each transition is depicted in purple, with values less than 0 shown in grey. Boxes represent clusters identified by the optimal spectral clustering solution (more information in section 1.7. and 2.6 of the appendix).

to navigate to the lowest level of three subcategories before viewing products (e.g., Cupboard \rightarrow Cereals \rightarrow Healthy cereals). It was used relatively infrequently ($\mu_{category} = 3.02, 95\%CI = [2.98, 3.05]$). After navigating with the search bar or the category drop-down, shoppers would be shown a list of associated products, where they could purchase products or click on products for more information. Before checkout, they could also add products from a personalised recommender system that suggested products that might have been forgotten before checkout ($\mu_{forgotten} = 0.322, 95\%CI = [0.3199, 0.327]$). Visitors removed an average of 3.23 products from their basket before checking out ($95\%CI = [3.20, 3.27]$).

3.2.2 In-store data

To prevent information leaking into our analysis of online shopping behaviour, we used a distinct dataset of in-store shopping behaviour to develop knowledge representations. In-store grocery receipts are unordered, making it particularly useful for this study. The final dataset contained purchase information from 4,336,917 distinct baskets. We followed the same procedure of encryption as with the clickstream data in order to preserve the privacy of customers.

We developed representations of episodic, semantic and hierarchical knowledge.

3.2.2.1 Episodic knowledge

The long-term episodic retrieval structure used in the Search of Associative Memory (SAM) model (Raaijmakers and Shiffrin, 1980) associates items more strongly to the extent that they co-occur during encoding. Whilst we cannot know the exact context in which consumers encoded products, we assume that products purchased together more frequently in the same basket must be stronger episodic associates. Episodic associations are therefore represented using the pairwise co-occurrence between products observed in the in-store dataset.

Episodic similarities $S(a,b)$ were determined using the probability of co-occurrence between products a and b , which is given by:

$$S(a,b) = \frac{frq(a,b)}{frq(a)}$$

Where $frq(a,b)$ is the total number of number of times that product a co-occurred with b in the same basket across the dataset.

3.2.2.2 Semantic knowledge

In addition to episodic memory, shoppers likely also rely on semantic memory to guide their retrievals. A common feature of modern semantic memory models is that they represent knowledge within a connected representational space, allowing people to generalise their knowledge to observations they haven't directly experienced (Jones et al., 2015). Unlike episodic co-occurrence, two items that have never co-occurred together may still be considered semantically similar, so long as they co-occur in similar contexts. We followed recent research (Hornsby et al., 2019) by training a 200-dimensional distributed semantic model using the in-store data.

For this project, we chose to learn 200-dimensional vector representations for products with *word2vec* Mikolov et al. (2013). This is because *word2vec* tends to scale better when trained on large datasets, because it can be trained stochastically. Rather than encoding words (e.g., as they might appear in the product descriptions), *word2vec* was trained to represent supermarket product codes, as they might appear in till receipts. Concretely, there is a different product code for each distinct product in the supermarket. Small variations in that product (i.e. different sizes of the same t shirt) are not given separate codes however. Each product code was thus represented as a one-hot encoded vector before being embedded, which resulted in a $42,837 \times 200$ matrix. During training, the model learns to associate product codes that often co-occur in baskets, which is analogous (but not the same) to how word vector models learn similarity between word tokens that often co-occur in sentences.

Associations S between any two vectors v_1 and v_2 was calculated using the cosine similarity $\text{cos}(v_1, v_2)$:

$$S(v_1, v_2) = \text{cos}(v_1, v_2) = \frac{v_1 \cdot v_2}{\|v_1\| \cdot \|v_2\|}$$

Conventionally, *word2vec* is trained using one of two network architectures. In this case, we used Continuous Bag of Words (CBOW) with negative sampling. CBOW assumes that items within baskets are un-ordered (i.e., known in natural language processing as the “bag of words” assumption). This assumption is true of in-store supermarket data, where product codes within baskets are unordered once they reach the database. Using CBOW, the training objective O is to maximise the likelihood of a target item i_j given a window of c surrounding context items:

$$O = \frac{1}{T} \sum_{j=1}^T \log p(i_j | i_{j-c}, \dots, i_{j+c})$$

Where T is the corpus size (e.g., number of baskets) and $p(i_j | i_{j-c}, \dots, i_{j+c})$ is the probability of a target item i_j given average or summation of a set of context words $(i_{j-c}, \dots, i_{j+c})$. In this case, *word2vec* used a context window size of 15.

Whilst this probability *could* be determined using a softmax, it is not practical when called over large vocabulary sizes. Negative sampling mitigates this by randomly sampling a set of k “negative” items (in our case, 20) that did not appear in the present basket. The model then learns from these negative samples by treating them as false labels in multiple binary classification tasks, evaluated using a logistic loss and then updated using gradient descent. In this project, this training process was repeated for 15 epochs.

We determined these aforementioned hyperparameters by monitoring the loss on the training set and visually inspecting the nearest products in terms of cosine similarity for a set of randomly-chosen products. During model fits, negative cosine similarities were clipped to 0.

3.2.2.3 Hierarchical knowledge

A strict organisation of products is imposed on consumers by way of a product taxonomy, which groups products from small subgroups (e.g., apples) to large departments (e.g., produce). Amongst other things, this taxonomy determines the proximity of products on the shelves and aisles of a supermarket store. The product taxonomy used here contained five levels. Each product had a unique taxonomisation.

If products a and b shared the same low-level category within the taxonomy, then they were said to be perfectly associated $S(a, b) = 1$. Products with entirely different taxonomic classifications had $S(a, b) = 0.2$. The remainder increased in increments of 0.2.

Uniquely, this hierarchy describes a taxonomy of *is-of* relations, which we used to define a measure of similarity as opposed to distance in a continuous semantic space.

3.3 Past choices cue subsequent retrievals

I begin by evaluating the first major claim of our model; that past choices cue the retrieval of subsequent options. If this is the case, then one would expect sequential purchases to be closely related in memory. One might also expect choices to become more dissimilar over time, as goals become increasingly satisfied.

3.3.1 Method

3.3.1.1 Permutation tests

To assess whether adjacent retrievals were more related than would be expected at random, we performed a permutation test. Each product within the clickstream data was encoded with each of the three representations described above. We then calculated the per-visit mean similarity between consecu-

tively added products. These were compared to the per-visit mean similarities determined by 100 random permutations of the product order, permuted within each visit. Thus, for significance tests reported; $N_{true} = 132,146$ and $N_{permuted} = 13,214,600$.

3.3.1.2 Response times

Response times (RTs) were compared for transitions of varying distances. These were capped at 60 seconds to minimise the leverage of outliers.

In the multiple linear regression comparing each similarity measure, RTs were monotonically transformed using a log function, $RT_{log} = \ln(RT + 1)$, due to positive skewness.

3.3.1.3 Trajectory analyses

Correlational analyses were conducted to assess how behaviour changed over time. Because visits contained differing numbers of products, basket adds within each visit were binned into equally-sized deciles based on their proximity to checkout. Subsequently, 13,901 (10.52%) visits were dropped because they purchased fewer than 10 products, leaving 5,174,018 choices.

3.3.2 Results

If shoppers cue retrievals from a given long-term store, then one would expect the similarity between consecutive purchases to be higher than when compared with random permutations, where the order of products has been randomly permuted within each shopping trip. For example, given that butter and bread are episodically linked (e.g., purchased in the same baskets), this should increase the probability of them being chosen consecutively. As shown in Figure 3.2a, the average trip-wise similarity between consecutively purchased items was significantly higher for the true order of purchases compared to the permuted order for episodic ($Median_{true} = 0.0756$, $IQR_{true} = 0.01465$ & $Median_{permuted} = 0.0154$, $IQR_{permuted} = 0.0193$) (Mann-Whitney U

= 2.95×10^{11} , $p < .0001$, $CLE = 0.8312$), semantic ($Median_{true} = 0.2507$, $IQR_{true} = 0.1495$ & $Median_{permuted} = 0.0770$, $IQR_{permuted} = 0.0669$) (Mann-Whitney U = 1.5×10^{11} , $p < .0001$, $CLE = 0.9141$) and hierarchical representations ($Median_{true} = 0.5169$, $IQR_{true} = 0.1433$ & $Median_{permuted} = 0.2620$, $IQR_{permuted} = 0.0504$) (Mann-Whitney U = 1.07×10^{11} , $p < .0001$, $CLE = 0.9390$). This suggests that choices were nonrandom and were cued by their similarity with the prior choice.

Sequential cued-retrieval can be viewed as a ripple through memory, in that more recently retrieved items tend to be more similar (Hills et al., 2012). As shown in Figure 3.2b, choices were most similar to the prior choice according to their episodic similarity (Figure C2 shows similar patterns for semantic and hierarchical knowledge). To confirm this, we calculated the average similarity of each choice for lags ranging from -10 to -1. Regressing lag onto the standardised average similarity for each user revealed a positive average relationship across each representation, indicating that more recent choices tend to be more similar ($\mu_{episodic} = 0.1357$, $95\%CI = [0.1331, 0.1383]$, $Z_{sign} = 56229.0$, $p < .0001$, $\mu_{semantic} = 0.1837$, $95\%CI = [0.1816, 0.1859]$, $Z_{sign} = 59170.5$, $p < .0001$, $\mu_{hierarchical} = 0.2263$, $95\%CI = [0.2247, 0.228]$, $Z_{sign} = 61651.5$, $p < .0001$).

The retrieval model described in Figure 3.1 selects options according to the similarity with the previous option. As this process repeats, the chance that a high similarity option has already been purchased increases, meaning that choices should become more dissimilar over time. Grouping choices into deciles based on their timestep (and thus adjusting for different trip sizes), we regressed each similarity measure onto timestep decile and included dummy coded representations of each transition type as confounding variables within each regression. Results showed that average similarity between sequential choices decreased over time across episodic ($b_{episodic} = -0.058$, $95\%CI[-0.061, -0.055]$, $p < .0001$), semantic ($b_{semantic} = -0.010$, $95\%CI[0.013, -0.007]$, $p < .0001$) and hierarchical representations

($b_{hierarchy} = -0.247$, 95%CI[-0.250, -0.243], $p < .0001$) (full regression equations in Table C3-C5). The increase in hierarchical similarity over time is visualised in Figure 3.2d.

One might correspondingly expect choices to become slower over time, as more dissimilar options are slower to retrieve. Regression analyses of inter-response intervals (IRIs) conformed to this expectation. Firstly, the standardised coefficients for episodic ($b_{episodic} = -0.137$, 95% CI = [-0.138 -0.136], $p < .0001$), semantic ($b_{semantic} = -0.086$, 95% CI = [-0.087 -0.085], $p < .0001$), hierarchical knowledge ($b_{hierarchy} = -0.273$, 95% CI = [-0.274 -0.271], $p < .0001$) were all negative predictors of IRI, indicating that more dissimilar options were slower to retrieve. Moreover, average IRIs appeared to slow over the duration of the trip ($b_{timestep} = 0.130$, 95%CI[0.129, 0.130], $p < .0001$). This slow-down is shown in Figure 3.2c. Including variables representing the navigation method (e.g., keyword search) and the similarity across each representation confirmed this slow-down as a general trend (full regression equation in Table C2) and this echoes similar patterns of slowing observed in category fluency tasks (Gruenewald and Lockhead, 1980; Hills et al., 2012; Abbott et al., 2015).

3.4 Episodic, semantic and hierarchical knowledge explain choice

We next evaluated whether consumers' sequential choices were best explained by one or multiple sources of knowledge, using the retrieval equation of a popular memory retrieval model, Search of Associative Memory (SAM) (Raaijmakers and Shiffrin, 1980). This equation formalizes how options may be retrieved based on their similarity with the current cue (illustrated in Figure 3.1c, with full equation in Section 1.3 of the appendix). Importantly, we evaluated its fit when including representations of episodic, semantic and hierarchical knowledge.

3.4.1 Method

3.4.1.1 Retrieval model

Broadly, our retrieval model is based on the retrieval equation from Search of Associative Memory (SAM) (Raaijmakers and Shiffrin, 1980). It assumes that retrieval and thus the decision of what to choose next is achieved by querying associative structures in memory with a memory probe. We follow previous models of semantic fluency by using the most recently chosen option O_i to probe associative memory structures (Hills et al., 2012; Abbott et al., 2015). Whilst other possibilities exist — such as a decaying influence of all previous retrievals — we focus on the role of the prior choice in order to simplify analyses (for a review of other approaches, see Kahana, 2020). The retrieval strength of the subsequently chosen option O_{i+1} is given by the product of the M associations between the present choice and itself, $S(O_i, O_{i+1})_j$. For example, in the full model, we used episodic, semantic and hierarchy-based associations between products, meaning that $M = 3$. This is then divided by the sum of that same function applied to all of the N options that remain to be added for that trip. This then gives rise to an overall probability of retrieval for each choice:

$$P(O_{i+1}|S_1, S_2, \dots, S_j, O_i) = \frac{\prod_{j=1}^M S(O_i, O_{i+1})_j^{\beta_j}}{\sum_{k=1}^N \prod_{j=1}^M S(O_i, O_k)_j^{\beta_j}} \quad (3.1)$$

We compared the inclusion of episodic, semantic and hierarchy-based associations. β values represent attention weights for each of these knowledge representations and were estimated as free parameters for each visit.

Each model was compared with a random baseline model, which predicted an equal probability of $\frac{1}{N}$ for every transition using a single representation. Thus, each of the products remaining to be purchased by each visitor is assumed to have an equal probability of being chosen at each timestep according

Table 3.1: The % BIC improvement over the random baseline and the mean attention weights (with 95% confidence intervals) for each of the candidate models. Results show that including representations of multiple knowledge formats provides the best fit to the data (shown in bold)

	Δ BIC (%)	Episodic	Semantic	Hierarchy
Episodic	9.13	0.29 (0.001)		
Semantic	4.80		0.091 (0.001)	
Hierarchy	26.28			2.217 (0.016)
Episodic & Semantic	12.30	0.258 (0.001)	0.068 (0.001)	
Semantic & Hierarchy	29.10		0.055 (0.001)	2.105 (0.017)
Episodic & Hierarchy	31.79	0.174 (0.001)		2.004 (0.017)
Multiple	33.78	0.160 (0.001)	0.044 (0.001)	1.939 (0.018)

to the baseline model.

Fit procedure: Each measure of association j was raised to its own respective attention weight β_j ; these were treated as free parameters and fit to individual visitors using maximum likelihood estimation (attention weights were forced to have a lower bound of 0, in order to prevent individual retrieval probabilities from exceeding 1). These free parameters were solved separately for each visit using the SLSQP solver within SciPy.

Model input: Models were fit to the retrieval sequences in the clickstream data. In addition to non-computable similarities, observations were dropped from the clickstream data if they occurred during or after the use of a recommender system, which prompted users about items they may have forgotten before checkout. Finally, to ensure that parameter estimates were robust, visits were dropped if they contained fewer than 10 items. This left 117,337 distinct visits.

Because of the probabilistic and multiplicative nature of the model, negative or zero similarities were replaced with a very small but positive number $1e-7$.

3.4.2 Results

Results are presented as the mean improvement in the Bayesian Information Criterion (BIC) relative to a random model, for which the probability of each transition was equal across all remaining products (see Section 1.3 of the appendix for further details about the model fitting procedure).

As shown in Table 3.1, the best fitting model contained multiple memory representations, even after penalising for multiple parameters. This suggests that online grocery shoppers query multiple knowledge formats when deciding what to choose next. A model parameter recovery study revealed that each parameter could be recovered accurately, with correlations between actual and estimated parameters > 0.6 in all cases (reported in Section 2.8 of the appendix). This indicates that parameter estimates were uniquely identifiable and could therefore be interpreted.

Inspecting the average attention weights of the best fitting model, one can gauge the relative importance of each representation. Hierarchical knowledge received the largest weight, followed by episodic, then semantic knowledge. Further analyses (presented in Section 2.6 of the appendix), revealed that transitions between product groups tended to overlap with superordinate classifications in the product taxonomy (clusters of transitions between product groups at the third taxonomic level are visualised in Figure 3.2e). Together, this suggests that shoppers rely heavily on hierarchical knowledge about how products relate, which aligns closely with the taxonomy used to arrange products in stores.

We next evaluated whether response times were best explained by one or a combination of knowledge sources. A multiple linear regression was performed, predicting the IRIs between each choice, using each of the three similarity measures as predictors. We also included the number of products remaining to be purchased and dummy coded variables representing each of the navigation methods (e.g., keyword search); these served as confounding variables (full model equation in Table C6). Model comparisons that penalised for more

variables revealed that IRIs were best explained by this full model, rather than one containing a subset of similarity measures (model comparisons in Table C7).

Importantly, these results support our key claim that sequential choice in open-ended tasks depends on retrieval of options from multiple sources of long-term memory. It is perhaps surprising that episodic and semantic knowledge explain unique variance in consumer choices, given that the latter may derive from the former (Mack et al., 2020). However, episodic knowledge provides a more direct link between experiences than semantic knowledge, which may play a unique role during goal-directed choice. Most of all, shoppers appeared to depend on hierarchical knowledge about products, which emphasises the influence of taxonomic organisations during navigation of large option spaces.

These model fits demonstrate the complementary role of different knowledge systems during everyday sequential choice tasks, but should not be limited to such settings. For example, they should extend to more well-known experimental tasks, such as semantic fluency. To test this, we fit the same retrieval model to a separate dataset of sequential food retrievals collected in a controlled experiment (originally collected by Zemla et al., 2020a, and shared via Zemla et al. 2020b). In this task, 50 participants were given three minutes to retrieve as many food words as possible. Much like keyword searches, each retrieval was typed into a text box. For each word retrieved (e.g., “hamburger”), we found a corresponding product from the retailer, allowing us to measure the episodic, semantic and hierarchical similarity between sequential retrievals as before (more details of the method and results can be found in Section 3 of the appendix). After performing the same set of model comparisons, results showed that the best fitting model contained all three representations, even after penalising for the additional parameters. Moreover — much like shoppers — participants appeared to rely most on hierarchical knowledge when sequentially retrieving food items from memory. This suggests that these knowledge systems also influence sequential retrievals in controlled experimental tasks

and that our model fits are representative of memory retrieval and not merely the design of the website.

3.5 Relying on certain knowledge formats predicts retrieval errors

If shoppers rely on certain knowledge formats during retrieval, this may increase their propensity to make certain errors such as forgetting or falsely retrieving products. Forgetting indicates the failure to retrieve a relevant item (i.e., a *miss*) whereas removing items indicates the failure to suppress irrelevant retrievals (i.e., a *false-alarm*). Indeed, forgetting is often viewed as a failure of retrieval (Shiffrin, 1970; Anderson et al., 1994) and could simply result from “searching the wrong part of memory” (Bettman, 1979, page 40).

The retrieval model used here (Raaijmakers and Shiffrin, 1980) assumes that items will be activated according to a process of spreading activation (Anderson, 1983; Collins and Loftus, 1975). When operating on an episodic representation, this would tend to chain together products found together in the same basket (e.g., purchasing a Thai pepper may cue coconut milk, bamboo shoots and other complementary ingredients). Thus, we hypothesised that shoppers relying on episodic knowledge — as measured by the attention weights from the best fitting retrieval model — would be less likely to forget products, as they would tend to co-activate items often combined in pursuit of a goal. Forgotten items were measured through the use of recommender system, which displayed products the shopper had purchased recently and frequently in prior visits before checkout.

When operating on a semantic network, a spreading activation process would tend to co-activate products that are substitutable, or conceptually similar but not necessarily purchased together (e.g., purchasing a Thai pepper may co-activate other forms of pepper). We therefore hypothesised that shoppers relying more on semantic knowledge would be more prone to remove

products from their basket, indicating that they didn't actually need them. This shares a kindred spirit with theories of confabulation in memory retrieval (Deese, 1959; Roediger and McDermott, 1995), where high semantic similarity between studied items causes related items to be erroneously retrieved.

3.5.1 Method

3.5.1.1 Forgotten items

Forgotten items were flagged through use of a personalised recommender system, which prompted users about items they may have forgotten at the end of their visit, before they checked-out. The exact products shown to each customer were determined according to the recency and frequency of purchase in previous shops (online or in-store), de-duplicated against products that had been purchased in the present visit. This page was displayed to users prior to payment.

3.5.1.2 Removed items

Shoppers could also remove products from their basket at any time during the shop. The total number of removed items were counted for each user.

3.5.1.3 Predictive modelling

We explored whether the attention weights (β) from the best-fitting retrieval model would predict the number of forgotten or removed items. These weights reflect the extent to which each visitor recruited each of the three representations to guide choice. Importantly, models were estimated using choices that preceded use of the recommender system. Outlying attention weights (three standard deviations above the mean) were clipped for this analysis for numerical stability.

Attention weights were regressed onto the number of forgotten and removed items using linear regression. We also included the proportion of times




Figure 3.3: Mean number of forgotten items (with 95% confidence intervals) for each model attention weight (β). Results show that relying on episodic or hierarchical knowledge predicted fewer forgotten items, whereas attending to semantic knowledge predicted more forgotten items, as measured by use of a recommender system displayed before checkout.

that each shopper transitioned using each transition type (e.g., search arrival) as confounding variables.

3.5.2 Results

3.5.2.1 Forgetting products

As shown in Figure 3.3, results supported the prediction that shoppers with increased episodic retrieval forgot fewer items on average ($r_s = -0.0811$, 95% CI $[-0.0868, -0.0754]$, $p < .0001$). In addition, attending to hierarchical knowledge predicted fewer forgotten items ($r_s = -0.0088$, 95%CI $[-0.0145, -0.0031]$, $p \leq .003$). In contrast, the more shoppers attended to semantic knowledge, the more likely they were to forget items ($r_s = 0.0152$, 95% CI $[0.0095, 0.0209]$, $p = 0.0001$).

To gauge their relative usefulness, we scaled each attention weight and entered them into a multiple linear regression, regressing onto the number of forgotten items. We also included the number of choices and the proportion of choices made using each search context as confounding variables. The regression was significant ($F_{8,117328} = 178.7$, $p < 0.0001$, $R^2 = .012$). Importantly, higher attendance to episodic knowledge ($b_{episodic} = -0.053$, 95% CI = $[-0.059,$

-0.048], $p < .0001$) or hierarchical knowledge ($b_{hierarchy} = -0.041$, 95% CI = [-0.046, -0.035], $p < 0.0001$) negatively predicted forgetting, whereas attending more closely to semantic knowledge ($b_{semantic} = 0.015$, 95% CI = [0.009, 0.021], $p \leq 0.0001$) positively predicted forgetting (full regression equation reported in Table C9).

Relying on episodic knowledge could indicate a greater episodic memory capacity or more experience with products and their relationships. Both possibilities could explain the general trend of forgetting fewer items as one increases attendance to episodic knowledge. The formalisation of episodic knowledge predicts that shoppers will transition to products that frequently co-occur with the past choice. Thus, another possibility is that the recommended products were less relevant to those that used episodic knowledge to guide their search, because the recommended products were ones that had been purchased by that shopper with a high frequency in the past. More work — perhaps through a cognitive battery — is required to understand the relationship between retrieval from different knowledge systems and forgetfulness.

3.5.2.2 Removing products

We next examined whether shoppers attending to certain knowledge sources removed more items from their basket. In line with our predictions, shoppers who attended more closely to the semantic similarity between items removed more items from the basket ($r_s = 0.0992$, 95% CI [0.0935, 0.1049], $p < .0001$). Conversely, those attending to episodic knowledge ($r_s = -0.014$, 95% CI [-0.0197, -0.0083], $p < .0001$) or hierarchical knowledge ($r_s = -0.084$, 95% CI [-0.0897, -0.0783], $p < .0001$) removed fewer items on average.

One explanation for these results is that attending to semantic knowledge increases one's propensity to suppress irrelevant retrievals and thus add products to one's basket that they don't otherwise need. Further analyses (presented in Section 2.10.1 of the appendix) revealed that the removed items tended to have above average similarity with the chosen options across each

knowledge representation, which is supportive of this idea. Another possibility is that shoppers attending to episodic knowledge were more experienced with products and their associations and thus were less prone to mistakes. We leave these possibilities for future analyses.

Most importantly, these results provide further support for the main claim that shoppers query multiple knowledge formats when deciding what to choose next and that individuals may differ in the extent to which they rely on these systems. Future experimental work may wish to test these findings using explicit measures and thus causally evaluate the precise relationship between attention to representations and errors. Such studies would complement our claim that knowledge formats may be recruited by individuals to different extents.

3.6 Discussion

In open-ended choice tasks like grocery shopping, how do we decide what to choose next? Many factors could influence what is chosen, but we propose that much depends on the similarity with the preceding choice across multiple knowledge formats. This view makes a number of predictions that we confirmed. First, choices and their response times were predicted by their similarity with the last choice, suggesting that choices cue the retrieval of subsequent options. Second, this behaviour was best explained by a mixture of episodic, semantic, and hierarchical knowledge, suggesting that consumers reason about associations between products in different ways by querying different sources of knowledge. Thirdly, how prone consumers were to different types of memory errors was predicted by their reliance on different types of memory, as assessed by model fits.

As our model describes, retrieved options may be cued by prior choices. This likely explains why the sequential choices of online grocery shoppers clustered over time, as they do in fluency tasks (Bousfield and Sedgewick, 1944;

Gruenewald and Lockhead, 1980; Hills et al., 2012; Abbott et al., 2015; Avery and Jones, 2018). We build upon past research (Bhatia, 2019; Zhang et al., 2021; Kaiser et al., 2013; Keller and Ho, 1988; Kalis et al., 2013) by showing how memory retrieval mechanisms influence the generation of options in sequential decision-making tasks. These results would not be predicted by many classical models of preferential choice, which consider option-retrieval to be out-of-scope (Glimcher and Rustichini, 2004; Busemeyer and Rieskamp, 2014; Rangel et al., 2008). Our results demonstrate how choice options can be dynamically constructed in the moment depending on the context supplied by the previous choice. Future work could explore the influence of other past retrievals, which have been shown to influence list recall (for a review, see Kahana, 2020).

Choices may follow different trajectories depending on which sources of knowledge are queried. Overall, choices and their response times were best explained by the sequential similarity across episodic, semantic, and hierarchical representations. In addition, individual differences in the extent to which each representation was recruited predicted how many products would be forgotten or removed. This would not be predicted by many existing models of semantic memory retrieval (Abbott et al., 2015; Hills et al., 2012) and option generation (Bhatia, 2019; Zhang et al., 2021), which rely on a single measure of association. Associative knowledge likely takes several forms (e.g., see Mirman et al., 2017), which is consistent with our modelling approach and results. Future experimental work may wish to explore the role of different associations in memory retrieval tasks and whether such systems are cognitively or neurally distinct.

Although we focused on sequential retrieval of choice options, determining whether an option is goal relevant could also be key to choice (see Figure 3.1c). Whilst modelling goals was out of the scope of this study, we hope studying the interaction between goals and retrievals will be addressed in future work. A person's subjective preferences may also affect which retrieved options are

chosen (Zhang et al., 2021; Levy and Glimcher, 2011; Hornsby and Love, 2020). Choice itself can affect preferences (Hornsby and Love, 2020), which in turn may affect memory retrieval. For example, new episodic memories could be formed after purchasing a preferred pairing of balsamic vinegar and bitter salad. One exciting direction for future research is to consider how different shopping experiences for individuals lead to different memory representations, which in turn affect future purchasing decisions.

The results of this chapter build on those of Chapter 2. Previously, I showed that the semantic organisation of options used by consumers could be approximated by their episodic co-occurrences within shopping baskets. We build on those results here to demonstrate that both episodic co-occurrences and the associated semantic embeddings are useful in predicting sequential consumer choices. The semantic embeddings used in this chapter were trained using a slightly different algorithm (known as word2vec, Mikolov et al., 2013). This algorithm scales better to large datasets such as the one used here. However, the product embeddings generated by LDA could theoretically be used as an alternative measure of semantic similarity and would likely be similar to the semantic embeddings used in this chapter, given that they are trained in similar ways. The generative algorithm of LDA should not be seen as a rival to the retrieval model of this chapter; whilst both can feasibly generate a basket, the order in which LDA draws products does not depend on previous draws, in contrast to SAM. Both chapters emphasise the contribution of episodic and semantic knowledge to preferential choice and uniquely show how it can be inferred and then used to predict the real-world choices of consumers.

One possible confound is that sequential choices were biased by the design of the website. For example, adding different brands of cola from the same page could cause retrievals to appear more hierarchical, as they belong to the same sub-category. To test this, we re-ran all analyses on a filtered dataset of product transitions that occurred through use of the search bar (detailed in Section 2.11 of the appendix). All results were consistent with those reported here, which

is reassuring given that these transitions were perhaps best representative of memory-based search. In addition, the model that best explained sequential grocery choices also provided the best fit to sequential retrievals of foods, which were observed in a controlled laboratory task (detailed in Section 3 of the appendix). Thus, whilst design features may help shoppers to retrieve certain brands (e.g., brands of cola), shoppers still seem to depend on cued retrieval from multiple knowledge formats to determine what they look for next.

Analyses of large field data such as these complement findings from the lab, allowing theories of memory and cognition to be evaluated at an unprecedented scale with high ecological validity. In this case, we've shown that the sequential purchases of grocery shoppers are well explained by a model of memory retrieval that was originally developed to explain behaviour in lab tasks (Raaijmakers and Shiffrin, 1980; Hills et al., 2012). A large driver of this model's success in this task is that it makes use of three relevant embedding spaces that relate to knowledge systems proposed in studies of memory (Tulving, 1985). We hope these findings stimulate further work in the lab, where one typically has a higher degree of control for assessing questions of cause and effect. For example, an additional explanation for choices becoming slower and more dissimilar over time is that retrieved options are increasingly rejected as they become less goal relevant (e.g., goals become increasingly satisfied). Future lab studies could assess this claim by asking participants to choose options in the presence of more or fewer goals. Others could enquire about the content of people's goals and examine how they interact with choices over time.

Our approach may make it possible to use shopping behavior to detect cognitive impairments. Longitudinal studies link performance in retrieval tasks to memory decline in pre-clinical Alzheimer's populations (Mueller et al., 2015). While many people shop, relatively few people participate in such clinical tests until they experience serious memory impairment, thereby foregoing the advantages of an early diagnosis (Rasmussen and Langerman, 2019). Although

more work would be needed to establish efficacy and suitable ethical guidelines, model fits (e.g., changes in attendance to episodic memory cues) may in the future predict the onset of cognitive impairment. Such a system operating at scale with informed consent could improve outcomes for individuals and society.

Online shoppers may be more or less responsive to certain recommendations depending on their navigational strategy. Results showed that shoppers relying on episodic memory were less likely to purchase products from a recommender system that reminded shoppers of previous purchases before checkout. This may be of practical significance to marketers designing personalised recommender systems, who could adapt recommendations to suit the retrieval strategies of shoppers as estimated by our cognitive model. For example, shoppers relying on hierarchical knowledge could benefit from recommendations promoting episodically related products (e.g., “goes well with...”) whereas those relying on episodic knowledge could benefit from seeing semantically similar products (e.g., “people also viewed...”) Such insights would complement traditional machine learning systems, which do not typically consider variations in human cognition (Griffiths, 2015).

To conclude, we find that choice sequences in an open-ended task depend strongly on the sequential cuing of long-term knowledge. Shoppers appear to use their previous choice to probe similarities in memory to determine what to purchase next. Depending on which sources of knowledge are queried, shoppers may choose products in different orders or exhibit an increased propensity to forget. Working with models and memory formats originally developed in laboratory settings, we were able to verify and extend these ideas in a real-world setting. In doing so, we strengthen the case for the complementary nature of laboratory and large-scale, real-world studies (Goldstone and Lupyan, 2016; Hornsby et al., 2019) with linkages enhanced through common modelling approaches.

Bibliography

- Abbott, J. T., Austerweil, J. L., and Griffiths, T. L. (2015). Random walks on semantic networks can resemble optimal foraging. *Psychological Review*, 122(3):558–569.
- Anderson, J. R. (1983). A spreading activation theory of memory. *Journal of Verbal Learning and Verbal Behavior*, 22(3):261–295.
- Anderson, M. C., Bjork, R. A., and Bjork, E. L. (1994). Remembering can cause forgetting: retrieval dynamics in long-term memory. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 20(5):1063.
- Avery, J. and Jones, M. N. (2018). Comparing models of semantic fluency: Do humans forage optimally, or walk randomly? In *CogSci*.
- Bettman, J. R. (1979). Memory factors in consumer choice: A review. *Journal of Marketing*, 43(2):37–53.
- Bhatia, S. (2019). Semantic processes in preferential decision making. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 45(4):627–640.
- Bousfield, W. A. and Sedgewick, C. H. W. (1944). An Analysis of Sequences of Restricted Associative Responses. *The Journal of General Psychology*, 30(2):149–165.
- Busemeyer, J. R. and Rieskamp, J. (2014). Psychological research and theories on preferential choice. In *Handbook of choice modelling*. Edward Elgar Publishing.
- Collins, A. M. and Loftus, E. F. (1975). A spreading-activation theory of semantic processing. *Psychological Review*, 82(6):407–428. Place: US Publisher: American Psychological Association.

- Collins, A. M. and Quillian, M. R. (1969). Retrieval time from semantic memory. *Journal of Verbal Learning & Verbal Behavior*, 8(2):240–247.
- Deese, J. (1959). On the prediction of occurrence of particular verbal intrusions in immediate recall. *Journal of Experimental Psychology*, 58(1):17–22.
- Ebbinghaus, H. (1913). *Memory: A contribution to experimental psychology*. Memory: A contribution to experimental psychology. Teachers College Press, New York, NY, US.
- Glimcher, P. W. and Rustichini, A. (2004). Neuroeconomics: the consilience of brain and decision. *Science (New York, N.Y.)*, 306(5695):447–452.
- Goldstone, R. L. and Lupyan, G. (2016). Discovering psychological principles by mining naturally occurring data sets. *Topics in Cognitive Science*, 8:548–568.
- Griffiths, T. L. (2015). Manifesto for a new (computational) cognitive revolution. *Cognition*, 135:21–23.
- Gruenewald, P. J. and Lockhead, G. R. (1980). The free recall of category examples. *Journal of Experimental Psychology: Human Learning and Memory*, 6(3):225.
- Hills, T. T., Jones, M. N., and Todd, P. M. (2012). Optimal foraging in semantic memory. *Psychological Review*, 119(2):431–440.
- Hornsby, A. N., Evans, T., Riefer, P. S., Prior, R., and Love, B. C. (2019). Conceptual organization is revealed by consumer activity patterns. *Computational brain & behavior*, 3:162–173.
- Hornsby, A. N. and Love, B. C. (2020). How decisions and the desire for coherency shape subjective preferences over time. *Cognition*, 200:104244.
- Johnson, J. G. and Raab, M. (2003). Take The First: Option-generation and resulting choices. *Organizational Behavior and Human Decision Processes*, 91(2):215–229.

- Jones, M. N. and Mewhort, D. J. K. (2007). Representing word meaning and order information in a composite holographic lexicon. *Psychological Review*, 114(1):1–37.
- Jones, M. N., Willits, J., and Dennis, S. (2015). Models of Semantic Memory.
- Kahana, M. J. (2020). Computational Models of Memory Search. *Annual Review of Psychology*, 71(1):107–138.
- Kaiser, S., Simon, J. J., Kalis, A., Schweizer, S., Tobler, P. N., and Mojzisch, A. (2013). The cognitive and neural basis of option generation and subsequent choice. *Cognitive, Affective & Behavioral Neuroscience*, 13(4):814–829.
- Kalis, A., Kaiser, S., and Mojzisch, A. (2013). Why we should talk about option generation in decision-making research. *Frontiers in Psychology*, 4.
- Keller, L. and Ho, J. (1988). Decision problem structuring: generating options. *IEEE Transactions on Systems, Man, and Cybernetics*, 18(5):715–728.
- Klein, G., Wolf, S., Militello, L., and Zsombok, C. (1995). Characteristics of Skilled Option Generation in Chess. *Organizational Behavior and Human Decision Processes*, 62(1):63–69.
- Levy, D. J. and Glimcher, P. W. (2011). Comparing Apples and Oranges: Using Reward-Specific and Reward-General Subjective Value Representation in the Brain. *Journal of Neuroscience*, 31(41):14693–14707.
- Mack, M. L., Preston, A. R., and Love, B. C. (2020). Ventromedial prefrontal cortex compression during concept learning. *Nature communications*, 11(1):1–11.
- Mikolov, T., Sutskever, I., Chen, K., Corrado, G., and Dean, J. (2013). Distributed representations of words and phrases and their compositionality. In *Proceedings of the 26th International Conference on Neural Information Processing Systems - Volume 2*, NIPS’13, pages 3111–3119, Red Hook, NY, USA. Curran Associates Inc.

- Mirman, D., Landrigan, J.-F., and Britt, A. E. (2017). Taxonomic and thematic semantic systems. *Psychological Bulletin*, 143(5):499–520.
- Mueller, K. D., Koscik, R. L., LaRue, A., Clark, L. R., Hermann, B., Johnson, S. C., and Sager, M. A. (2015). Verbal Fluency and Early Memory Decline: Results from the Wisconsin Registry for Alzheimer’s Prevention. *Archives of Clinical Neuropsychology*, 30(5):448–457.
- Proust, M. (1913). *In Search of Lost Time*. Modern Library, New York.
- Raaijmakers, J. G. W. and Shiffrin, R. M. (1980). SAM: A Theory of Probabilistic Search of Associative Memory. In Bower, G. H., editor, *Psychology of Learning and Motivation*, volume 14, pages 207–262. Academic Press.
- Rangel, A., Camerer, C., and Montague, P. R. (2008). A framework for studying the neurobiology of value-based decision making. *Nature Reviews Neuroscience*, 9(7):545–556.
- Rasmussen, J. and Langerman, H. (2019). Alzheimer’s Disease – Why We Need Early Diagnosis. *Degenerative Neurological and Neuromuscular Disease*, 9:123–130.
- Roediger, H. L. and McDermott, K. B. (1995). Creating false memories: Remembering words not presented in lists. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 21(4):803–814.
- Shiffrin, R. M. (1970). Forgetting: Trace Erosion or Retrieval Failure? *Science*, 168(3939):1601–1603.
- Troyer, A. K., Moscovitch, M., and Winocur, G. (1997). Clustering and switching as two components of verbal fluency: evidence from younger and older healthy adults. *neuropsychology*, 11(1):138.
- Tulving, E. (1985). How many memory systems are there. *American Psychologist*, 40:385–398.

- Zemla, J. C., Cao, K., Mueller, K. D., and Austerweil, J. L. (2020a). SNAFU: The Semantic Network and Fluency Utility. *Behavior Research Methods*, 52(4):1681–1699.
- Zemla, J. C., Cao, K., Mueller, K. D., and Austerweil, J. L. (2020b). SNAFU: the Semantic Network and Fluency Utility. [Dataset] <https://github.com/AusterweilLab/snafu-py> accessed on 13/10/2021.
- Zhang, Z., Wang, S., Good, M., Hristova, S., Kayser, A. S., and Hsu, M. (2021). Retrieval-constrained valuation: Toward prediction of open-ended decisions. *Proceedings of the National Academy of Sciences*, 118(20).

Chapter 4

Learning preferences from past choices

4.1 Introduction

“We are, in a very real sense,
characters of our own creation”

Nick Chater

The Mind is Flat

“Things just happen in life, and
pretty much after the fact, we
make up a story to make it all
seem rational”

Michael Gazzaniga

*Tales from Both Sides of the
Brain*

Every day, people are confronted with countless choices for which there is no objectively correct answer. These tend to be either preference judgments or moral decisions (Nakao et al., 2012). Rather than being guided by extrinsic feedback, people choose these options freely for themselves, using their subjective preferences. We therefore refer to these choices as *free choices*. But

how do people acquire these subjective preferences in the first place? The aim of this chapter is to understand how people learn subjective preferences over time and use these to make inferences about untried options.

We might learn about our own preferences in the same way we learn about others'; by observing and then rationalizing behavior (Bem, 1967, 1972; Cushman, 2019). This is because we tend to lack introspective access to the mechanisms driving our behavior, meaning that we have to post-rationalize in order to make sense of it (Bem, 1967, 1972; Miller and Buckhout, 1973; Mandler, 1975; Nisbett and Wilson, 1977). In a dramatic demonstration of this, people have been tricked by mischievous experimenters into justifying choices that they did not actually make (Hall et al., 2010; Sauerland et al., 2016; Strandberg et al., 2018). For example, after choosing their favourite flavour of jam in a taste test, participants were tricked into then justifying a different choice by experimenters, who covertly switched them mid-way through the experiment (Hall et al., 2010). Thus, rather than accessing the reasons for their choices directly, people seem to retrospectively infer them using evidence of their historic choices, even when that evidence is not valid.

As well as facilitating the inference of preferences, past choices also shape them. This has been demonstrated in studies of free choice, which show that after freely choosing an option, people tend to increase their subjective preference for it (Brehm, 1956; Sharot et al., 2009; Alos-Ferrer and Shi, 2012; Ariely and Norton, 2008; Schonberg et al., 2014; Koster et al., 2015; Miyagi et al., 2017; Voigt et al., 2017; Akaishi et al., 2014; Nakao et al., 2016; Vinckier et al., 2019; Chammat et al., 2017; Cockburn et al., 2014; Rieffer et al., 2017). In the original free-choice paradigm, Brehm (1956) asked participants to rate a set of items (e.g. snack products), choose between two similarly rated options and finally to rate the full set again. Results showed that after making the forced choice, they had an increased preference for the chosen item on the final rating and a decreased preference for the rejected item. This is surprising, because it suggests that merely choosing or rejecting an option causes a person to update

their subjective preference for it.

Although there has been some debate as to the validity of the free-choice paradigm in its original format, more recent studies have suggested that choice-based learning is real. For example, one particular concern about the original paradigm was that the first rating phase was noisy and therefore an imperfect measure of people's true preferences (Chen and Risen, 2010; Izuma and Murayama, 2013). However, researchers have since overcome this concern using various methodological adaptations, demonstrating choice-induced preference change does occur (Sharot et al., 2010; Alos-Ferrer and Shi, 2012; Koster et al., 2015; Schonberg et al., 2014; Miyagi et al., 2017; Akaishi et al., 2014; Nakao et al., 2016; Vinckier et al., 2019) and can be long-lasting (Sharot et al., 2009). For example, Sharot et al. (2010) asked participants to blindly choose between masked holiday destinations, which were only revealed to participants after one of two keys had been pressed. Subsequent ratings of those destinations were consistent with choice-induced preference change, even though choices had been randomly assigned to participants. More recent analyses of data collected from supermarket shoppers in-the-wild gives further credence to the claim that choices are self-reinforcing. In particular, a recent study of 283,000 British consumers found that their tendency to repeat a choice (i.e., *exploit*) strengthened as a function of the number of previous repetitions (Riefer et al., 2017). The consensus from studies inside and outside the laboratory is that free choices appear to be self-reinforcing, such that people come to prefer the options they choose.

Studies of free choice imply that people refer to past acceptances and rejections to infer what they like and dislike (Akaishi et al., 2014; Cockburn et al., 2014; Miyagi et al., 2017; Izuma et al., 2010; Chammat et al., 2017; Riefer et al., 2017). Yet, we know that people can also infer the value of things they've never tried. For example, one could infer that they would not enjoy sky diving, despite having never tried it. How do people do this? Rather than caching the value of individual options (e.g. beer varieties), people likely represent

options and preferences within a shared, continuous, multidimensional space (e.g., varieties of hop, brand and brewing style). As depicted in Figure 4.1, representing options and preferences in this way is beneficial, in that it provides a lower-dimensional learning problem and allows one to infer the relative value of any option in their environment, irrespective of whether it has been tried.

Making free choices may therefore serve a considerably broader function than first thought, helping us to learn more deeply about ourselves and the world around us. Specifically, if options are represented within the same attribute space, then free choices may help to determine where one’s preferences lie within that space. In this paper, we propose that the position of one’s preferences is determined by a general, error-driven learning process, where the error term seeks to make the last choice more likely to repeat. As well as increasing the likelihood of past choices being repeated — as has been shown in real supermarket consumers (Riefer et al., 2017) — one should also increase their preferences for other options to the extent that they are *similar* to those previously chosen. While surprising, striving for internal coherency in this way may make sense in a world where choices can be evaluated across a multitude of different criteria.

We begin by demonstrating how the intrinsic desire to *maximize coherency* between past choices and present preferences can elicit strong subjective preferences in the absence of extrinsic reinforcement. In accordance with our proposed theory, we develop a computational cognitive model that learns preferences over choice attributes and uses past choices as the basis for updating them. We call it the Coherency Driven Choice (CDC) model. CDC is similar to models in the field of human category learning, which are primarily concerned with classification of items into a set of mutually exclusive categories via their attributes (e.g., using wings, beaks and feathers to describe birds) (Kruschke, 1992; Nosofsky, 2011; Love et al., 2004). However, rather than updating based on corrective feedback, our model self-supervises using its past choices, thereby making them and similar options more likely to be sampled.

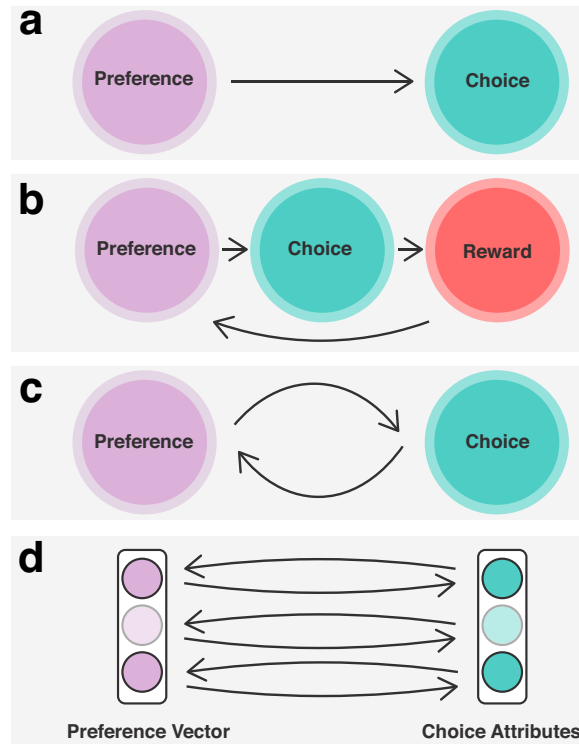


Figure 4.1: Many popular models of decision-making cannot easily explain how people form strong subjective preferences for free choices made in the wild. **a**, In standard decision theory, it is assumed that preferences remain stable over time (von Neumann et al., 1944; Glimcher, 2009). Indeed, for many researchers, the challenge is often learning what people’s preferences *are* (e.g., by asking them to choose between options), rather than understanding how they *became*. **b**, Reinforcement Learning (RL) models contrast in that they assume preferences change over time. Specifically, RL agents learn to prefer actions with a higher expected reward, which they learn as they monitor extrinsic feedback from their environment (Sutton and Barto, 1998). While RL has been shown to account for many aspects of human learning well (for a review, see Daw and Tobler 2013), these investigations have been largely confined to objective tasks, where there is a clear extrinsic signal steering the decision-maker. **c**, Studies of free choice — where there is no objectively correct answer — have shown that merely choosing an item increases one’s preference for it. This has often been taken to imply that free-choices are self-reinforcing, so as to increase preferences toward the chosen option. Yet, caching values in this way would arguably not scale well, as it would require people to keep track of every item they’d ever tried. **d**, In this paper, we propose a novel theory of subjective preference formation and decision making that arguably scales better to free choices made in the real world. According to this theory, people encode preferences over attributes of free choices, such as the hop content in beer. After making a decision, they then update their preferences in the direction of the attributes that defined their prior choice, thereby increasing their preference for it, as well as for other options that have similar attributes.

Moreover, we use these mechanisms to make decisions that do not involve any fixed set of classes; the model chooses a set of items, which is not of fixed size. Through simulation, we show how this mechanism can drive complex, multidimensional preferences from free choices alone. Thus, error-driven, self-supervision helps the agent to maximize coherency between its preferences and choices over time. As a result, CDC can achieve a sense of order in environments where there are innumerable possible options and dimensions by which to score them.

After presenting this formal demonstration of our theory, we validate its predictions using a large-scale experiment of human participants. Chiefly, the error-driven nature by which CDC learns means that it will update its preferences in order to maximize the perceived contrast between accepted and rejected options. This is analogous to contrastive learning effects documented within the field of category learning, where the experience of contrasting category exemplars causes perceived category averages to drift apart and become idealized (Davis and Love, 2010). Results from Experiment 1 demonstrated that people update their preferences in a similar way following a choice. Specifically, participants were shown to prefer never-before-seen patterns if they happened to be on the back of a toy robot they had just designed. The more discriminating the pattern was to the initial robot, the more likely they were to prefer it.

Whereas Experiment 1 concerned preference formation in a novel and well-controlled domain, Experiment 2 evaluated the model’s predictions in a domain in which people hold strong, preexisting preferences. In particular, Experiment 2 evaluated whether participants would retrospectively update their political beliefs following a vote. Results revealed that after choosing between two electoral candidates based on trivial grounds (e.g., whether they liked cats or dogs), participants were more likely to agree with a political belief later revealed by their chosen candidate, irrespective of whether that belief was traditionally left or right wing (e.g., pro-choice vs. pro-life abortion

rights). Thus, these results support the key claims of our proposed theory, in that they suggest that people retrospectively update their preferences to be coherent with their past choices. Significantly, this even occurs in domains where people possess strong prior preferences that likely have strong subjective significance.

4.2 The Coherency Driven Choice (CDC) Model

We begin by formally describing our model of subjective preference learning and decision making.

Broadly speaking, the model works by maintaining an internal set of preferences and attention weights for attributes across choices. For example, all products in a supermarket can be described in terms of nutritional attributes such as salt, sugar and saturated fat content. Individuals will possess different preferences for those attributes, and pay differing levels of attention to them. These preferences and attention weights are used to determine how favourable a choice is at a given timepoint. In particular, the higher the attention-weighted similarity, the more likely it will be to choose that option. Our model can be thought of as an agent interacting with its environment. Much like a reinforcement learning agent, the model takes an action, observes its environment, updates its internal state and then repeats the process.

Here we introduce some important notation relevant to the model’s decision making process. Note that vectors will now be denoted in bold lowercase letters and matrices in bold uppercase letters. We denote the observation of choices in the environment using the matrix \mathbf{O} , which has a shape of $N \times M$. Here, N denotes the number of choices available $\mathbf{O} = [\mathbf{o}_1, \mathbf{o}_2, \dots, \mathbf{o}_N]^T$ at a given timestep. For simplicity of notation, we assume that the model must choose between two items at any one time (i.e. $N=2$). However — in principle — the model is not constrained to this. M denotes the number of attributes for

each option. Thus, each column \mathbf{o}_i ($i \in \{1, \dots, N\}$) is a vector of M attributes $\mathbf{o}_i = [o_{i1}, o_{i2}, \dots, o_{iM}]$. Therefore, the element o_{ij} corresponds to the j th attribute ($j \in \{1, \dots, M\}$) of the i th item.

Preference similarity

In order to determine the most appropriate choice, the model first calculates a probability over the available options observed in \mathbf{O} using the preference vector $\mathbf{p} = [p_1, p_2, \dots, p_M]^T$ and the attention weight vector $\mathbf{w} = [w_1, w_2, \dots, w_M]^T$. Each element of the preference and attention weight vectors p_j and w_j maps to an attribute j in the attribute vector o_{ij} .

After sampling an option from the environment, the agent must update the preference and attention weight vectors. We now discuss the process of computing probabilities, selecting actions and updating vectors in more detail.

Choice probabilities

In order to determine the probability of an action, the model calculates an attention-weighted similarity between the preference vector and each of the $N = 2$ item vectors \mathbf{o}_i within the observation matrix \mathbf{O} . We denote the attention-weighted similarity as $a(\mathbf{o}_i)$

$$a(\mathbf{o}_i) \equiv -\gamma \left(\sum_{j=1}^M w_j (o_{ij} - p_j)^2 \right)^{1/2} \quad (4.1)$$

Where γ is a scaling hyperparameter. Note that this weighted Euclidean similarity term is very similar to the one used in the ALCOVE model of human category learning (Kruschke, 1992).

In order to determine the probability of selecting an option i , the attention-weighted similarity $a(\mathbf{o}_i)$ is then fed into a softmax function

$$\mathbf{f}(\mathbf{o}_i) \equiv P(I_i | \mathbf{O}) \equiv \text{softmax}(\mathbf{a}(\mathbf{O}))_i = \frac{\exp a(\mathbf{o}_i)}{\sum_{k=1}^N \exp a(\mathbf{o}_k)} \quad (4.2)$$

Thus, the preference $a(\mathbf{o}_i)$ for an option is a function of three things:

1. The similarity (i.e. Euclidean distance) between the preference vector \mathbf{p} and the choice attribute vector \mathbf{o}_i — The more similar the attributes and corresponding preference values, the more the model prefers that option
2. The attention weight vector w — A higher degree of attention towards a similarity leads to a greater impact on the overall preference
3. The scaling hyperparameter γ — The higher the γ , the higher the probability for selecting the preferred option (when using softmax action selection)

Similar to more traditional models, preferences are represented as ideal-points within a multidimensional space (Greenhoff and MacFie, 1994). However, unlike many of those methods, the model can have varying levels of attention to those according to the attention weights and — crucially — describes how preferences update over time as a consequence of decision making.

Action selection

Choices can be selected using one of the many popular strategies used in RL, such as ϵ -greedy, softmax action selection (Sutton and Barto, 1998) or more sophisticated directed-exploration strategies (e.g., uncertainty minimization). In each case, higher probabilities for choices (i.e. stronger preferences) increase the likelihood of exploiting that known favourite, rather than exploring disfavoured options. When using softmax selection specifically, the λ parameter can be thought of as determining the “fussiness” of the agent’s choices, such that higher λ equates to a higher likelihood of choosing the favorite. We denote the choice made by the agent as c .

Updating preference and attention-weight vectors

Following an action, the agent must then update its preference and attention weight vectors. As discussed in the main text, a battery of psychological research has shown that — in subjective choice domains where there is no

explicit feedback — preferences tend to follow choices. We therefore update the preference and attention weight vectors so as to maximize the likelihood of the previous choice. This contrasts sharply with traditional preference models, which seldom specify how preferences may change over time (Greenhoff and MacFie, 1994; DeSarbo and Kim, 2012).

The exact learning procedure used to update the preference and attention weight vectors is gradient descent on the cross-entropy loss, similar to that used during backpropagation and in the neural network literature generally (Hinton et al., 1986; Goodfellow et al., 2016). During the learning procedure, an action is determined probabilistically using the softmax choice rule. After the action, the cross-entropy loss is calculated between the preference probabilities output by the model $\mathbf{f}(\mathbf{o}_i)$ and the actual choice c that was made.

$$l(\mathbf{f}(\mathbf{O}), c) \equiv - \sum_{i=1}^N 1_{\{c=i\}} \log(f(\mathbf{o}_i)) \quad (4.3)$$

After making an action, the preference and attention weights are updated so as to minimize the cross-entropy error. Concretely, they are updated proportionally to the negative of the error gradient.

We therefore use the following calculation to find the partial derivative of the preference vector \mathbf{p} with respect to the cross-entropy loss:

$$\frac{\partial l(\mathbf{f}(\mathbf{O}), c)}{\partial p_j} \equiv \gamma^2 w_j \sum_{i=1}^N (1_{\{c=i\}} - f(\mathbf{o}_i)) \frac{1}{a(\mathbf{o}_i)} (o_{ij} - p_j) \quad (4.4)$$

And the following calculation to find the partial derivative of the attention weight vector \mathbf{w} with respect to the cross-entropy loss:

$$\frac{\partial l(\mathbf{f}(\mathbf{O}), c)}{\partial w_j} \equiv - \frac{\gamma^2}{2} \sum_{i=1}^N (1_{\{c=i\}} - f(\mathbf{o}_i)) \frac{1}{a(\mathbf{o}_i)} (o_{ij} - p_j)^2 \quad (4.5)$$

We then use these partial derivatives to update the existing preference and attention weight vectors using gradient descent. Concretely, we define the following update rules for the vectors \mathbf{p} and \mathbf{w} , respectively:

$$\mathbf{p} := \mathbf{p} - \eta_{\mathbf{p}} \frac{\partial l(\mathbf{f}(\mathbf{O}), c)}{\partial p_j} \quad (4.6)$$

$$\mathbf{w} := \mathbf{w} - \eta_{\mathbf{w}} \frac{\partial l(\mathbf{f}(\mathbf{O}), c)}{\partial w_j} \quad (4.7)$$

Where $\eta_{\mathbf{p}}$ and $\eta_{\mathbf{w}}$ represent the learning rates for the preference and attention weight vectors, respectively. As is standard during gradient descent, these learning rates scale the updated vectors and thus determine the magnitude of the update at a given time step.

4.3 Learning strong preferences over time

To illustrate how one could learn strong subjective preferences by virtue of their choice trajectory, we simulated the CDC model. In the simulated environment, there were two choice types that did not vary on the first dimension but varied significantly on the second dimension. This is analogous to choosing between two beer brands that are similar in taste but contrast in the color of branding.

4.3.1 Method

4.3.1.1 Simulation

Observations in the environment were randomly sampled from distributions of two clusters. Choice *type a* had a cluster centroid of $(0.2, 0.8)$ whilst *choice type b* had a cluster centroid of $(0.2, 0.2)$. The standard deviation of each cluster was determined apriori to be 0.05; thereby making the two choice types linearly separable. A total of 500 observations were simulated.

The agent was initialised with middling preferences and attention weights across the two attributes of $(0.5, 0.5)$. It was also set to have a learning rate of 0.01, $\varepsilon = 0.05$ and $\lambda = 1$. Choices were simulated for 10,000 timesteps. At each timestep, the agent was forced to choose between two randomly-selected options from each choice type using an ε -greedy action selection strategy.

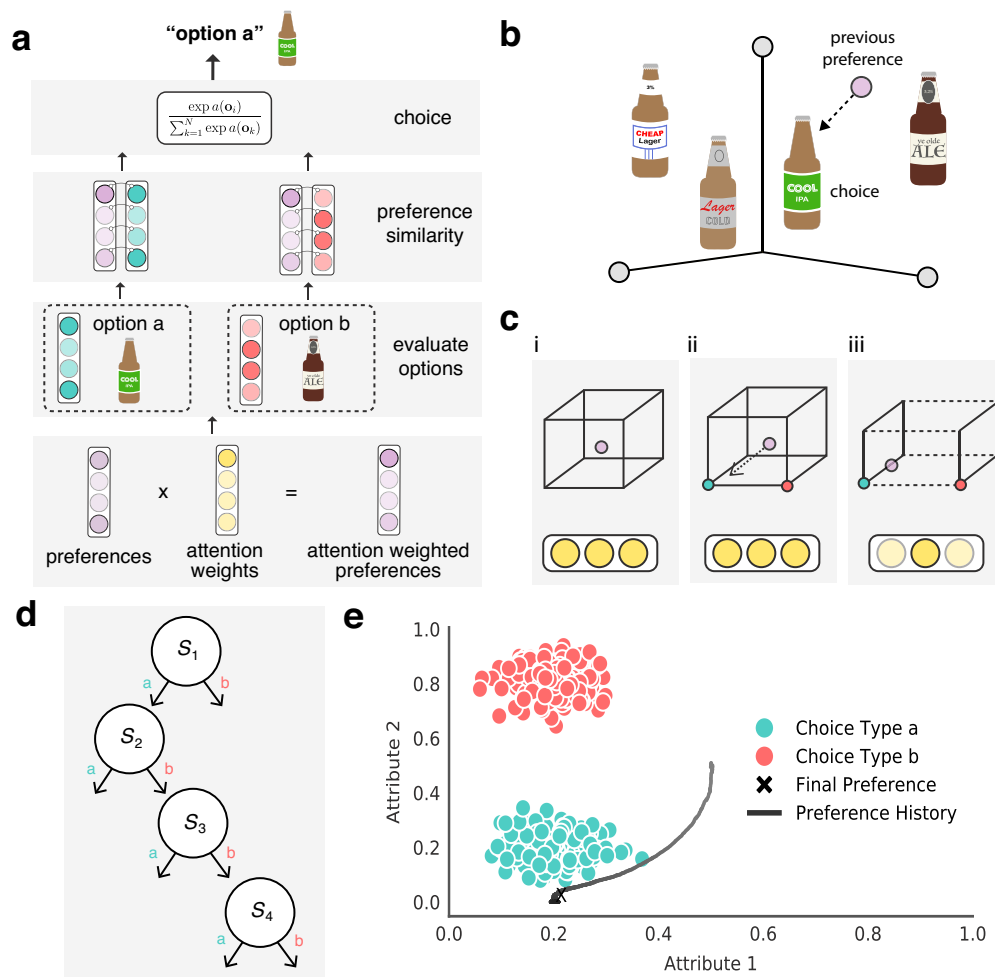


Figure 4.2: To illustrate how this theory can materialize in strong subjective preferences, we formalize it in a computational cognitive model, known as Coherency Driven Choice (CDC). **a**, CDC possesses a preference and attention-weight vector. When evaluating options, the model evaluates the similarity between its own attention-weighted preference and the attributes of the available options. The closer an option’s attributes are to the model’s attention-weighted preferences, the more likely the model is to select it. **b**, Following a choice, CDC updates its preference vector to make the past choice most likely using gradient descent, thereby moving toward the prior choice in attribute space. **c**, CDC also adjusts its attention weights to make prior choice more likely, effectively warping the preference space so that it becomes more sensitive to the attended attribute. **d**, As the model grows to prefer the options it initially chose, it tends to repeat the same choice type more over time. **e**, The model was simulated for 10,000 timesteps in a simple environment in which there were two choice types. Due to happenstance in the initial choices, the model began to prefer *choice type a*, adjusting its preferences and attention weights in favor of the attributes that make it unique. The preference history is coloured by attention weight ratio, such that blacker colors indicate a greater deal of attention paid towards attribute 2.

Choosing these parameters over a large number of timesteps helped the learning process to be smooth and stable across the agent’s lifespan. This is therefore more representative of what might happen over a timescale of several years. However, it should be noted that the same learning trajectory documented below could be found using a larger learning rate over a shorter number of timesteps, though many such situations would not involve complete movement towards a goal.

4.3.2 Results

A visualization of this model and the results of this simulation are shown in Figure 4.2.

After simulating 10,000 forced-choices, CDC eventually came to possess preferences resembling the first choice type (i.e., *type a*). After an initial sequence of random actions, CDC began to quickly develop a preference in retrospect of them, and thus select choices consistent with this new-found preference. This process was then self-reinforcing, further strengthening preferences and therefore the likelihood for choice *type a* over time. Thus, when consumers become less likely to explore new products the more they exploit the same, it may be because their preferences are being self-reinforced by their past choices, as shown in this simulation (Riefer et al., 2017).

Uniquely, as CDC chose more of *choice type a*, the preferences and attention of the model moved most in favour of the attributes of the choice that made it unique. This is a known consequence of discrimination learning, but uniquely demonstrated here within the context of subjective preference change (Davis and Love, 2010; Rescorla and Wagner, 1972; Ramsar et al., 2010). Exaggerating preferences in this way helped the model to maximize the perceived contrast between the accepted and rejected choices. This increases the likelihood of the past choice type being sampled again and reduces the likelihood of the rejected option being selected. Continuing the example introduced above, this suggests that a person will have an over-exaggerated preference toward

the unique branding of their preferred beer, helping to retrospectively justify their apparent preference.

Of course, when formalizing a cognitive theory, one must make some assumptions about the world. For example, in the case presented here, one could argue that people do not always have complete knowledge of the attributes describing each option at the time of decision. Indeed, one may need to taste a product to know how salty it is, or they may vote for a political candidate before learning of their stance on free-trade. In this case, our model would simply use a placeholder for their preference on that particular attribute. Upon revelation of the attribute value for their prior choice (e.g., discovering how salty a snack was), they would then move their preference toward the attribute point in question.

In reality, people are unlikely to develop preferences as exaggerated as the ones learned in this simulation. This is because choices in real life are often more innumerate, multidimensional and overlapping. The high degree of similarity between options in the real world would cause our agent to explore more, and thereby develop a less extreme set of preferences. The environment may also provide additional sources of noise during decision making that elicit exploration and thus movement of preferences in new directions. For example, a preferred product may be out-of-stock in a store, a person may develop a new allergy or travel to a new country. Rather than attempt to account for the multitude of ways in which preferences change as a function of choices in-the-wild, the aim of this simulation was to highlight an important consequence of the coherency maximizing mechanism proposed here; namely, that preferences are updated in order to discriminate the choice just made.

4.4 People prefer novel patterns associated with their prior choice

The new account proposed in this paper suggests that by learning preferences over attributes of choices, people can generalize their preferences to novel options that are associated with ones previously chosen. Moreover, similar to error-driven models developed in the field of category learning (Davis and Love, 2010), it predicts that people will update their preferences in the direction of the most discriminative elements of their choice, in order to maximize the likelihood it being repeated.

Experiment 1 aimed to validate these predictions. Here, participants were asked to design a robot (the trial flow is depicted in Figure 4.3a). They were then introduced to a second robot, before both turned around revealing previously unseen, randomly assigned patterns on their backs. Finally, participants were asked to choose between three patterns; one that was unique to the back of the robot they had previously designed (i.e. chosen unique), another that was shared across the backs of the two robots and a final pattern that was unique to the back of the robot they had not designed (i.e. non-chosen). It was hypothesized that — consistent with a discriminative account of learning — participants would prefer those novel patterns to the extent that they were uniquely associated with the robot they had just designed.

4.4.1 Method

4.4.1.1 Participants

One thousand and three participants were recruited from Amazon Mechanical Turk (mturk). Mturk (www.mturk.com) is generally known for being an inexpensive source of reliable human data (Crump et al. (2013), though see McDuffie (2019) for a discussion on the possible limitations). Participants were required to have completed >1000 tasks (or *HITs*) and have an acceptance rate

greater than 95%. Participants had to be based in the US or Canada. Data from 37 participants were removed due to having an average response time two standard deviations greater than the mean ($>17.64s$). The mean age of the participants was 37.0 ($SD=11.4$) and 50.2% were male. Participants were paid 50¢ for participating, which is typical for mturk (Horton and Chilton, 2010). Overall, the experiment took about 10 minutes¹.

4.4.1.2 Design

The experiment used a between-groups design with 10 trials. Participants either chose a pattern that belonged to the back of the robot they previously designed (i.e., *chosen-unique*), the one they did not design (i.e., *non-chosen*) or was shared across both. The dependent variable was therefore the sum of the preferences over each of the choice types across each of the 10 trials.

4.4.1.3 Apparatus & Stimuli

The study was designed using JavaScript and was accessed in a web browser. The task was presented in a 700 x 700 pixel screen. During the design phase of the experiment, participants responded by choosing attributes from a drop-down box. When stating their design preferences in the final phase, participants were asked to click each robot in order of their preference (from highest to lowest).

Each robot was designed using the Support Vector Graphics (SVG) format. These robots had a front — which could be designed by participants — and a back, which contained randomly-assigned patterns.

During the design phase, participants could design three aspects of the front of their robot; the stomach texture, the visor border colour and the eye colour. In each trial, participants could choose between two randomly-selected options for each design aspect.

In the final phase of each trial, participants were shown multicoloured

¹Both experiments presented here were in compliance with UCL’s code of ethics

geometric patterns. In total, there were three patterns per trial, randomly selected from one of 50 possible triplets. One pattern appeared on one random half of the back of the robot they had previously designed. Another pattern appeared on one random half of the back of the robot they had not designed. The final design was shared such that it appeared on both remaining halves of the two turned robots.

4.4.1.4 Procedure

Participants were initially briefed about the experiment in order to get their informed consent and asked to supply their age and gender. They were told that the task would take about 10 minutes and would require them to design robots and make choices. On each trial, participants were shown a robot with a randomly selected name and asked to design it. The hope was that by designing the robot, they would become more motivated about their choice, increasing ecological validity. They could choose between one of two randomly selected options for each of the three design aspects. After completing the design, participants were then introduced to a new, frowning *anti-robot*. This anti-robot was designed using all the attributes that the participant had previously eschewed. In addition, participants were warned that the anti-robot did not like the participants' design. They were then asked to reassure their own robot by clicking on it. Henceforth, we refer to this anti-robot as the *non-chosen* option. This is because it was uniquely designed using elements that had been explicitly rejected during the previous phase of the trial. After clicking on their own robot, both robots then turned around revealing randomized patterns on their backs (described above). The two robots then moved to the back of the screen. Three patterns — either shared across both or unique to one of the robots at the back of the screen — then appeared at the front of the screen and participants were asked to choose their favourite in order of preference. Participants completed 10 such trials with patterns, robot names, and other trial details randomized for each trial, they were debriefed, thanked for their

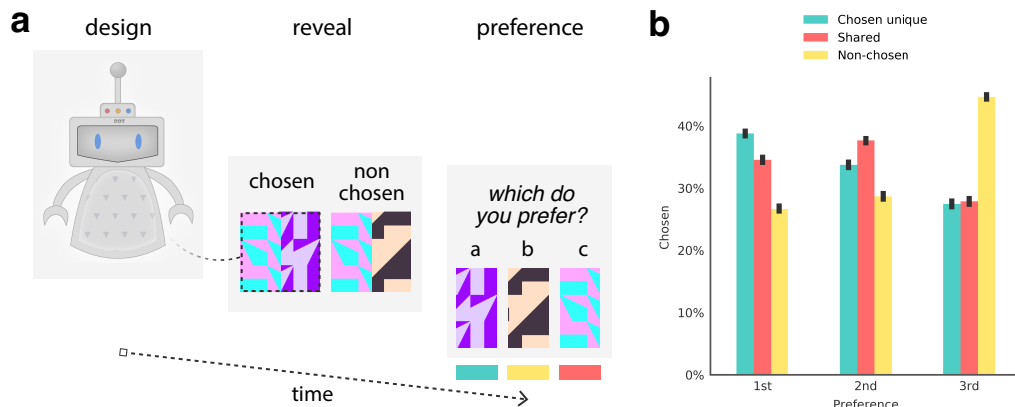


Figure 4.3: **a**, The trial sequence observed by participants. First, participants were introduced to a robot and asked to design aspects of it. They were then introduced to a new robot, which was designed differently. These robots then turned around, revealing random patterns on their backs. One pattern was unique to the robot participants had previously designed (i.e., chosen-unique), one was shared across both robots (i.e., shared) and the other was unique to the robot they had previously eschewed (i.e., non-chosen). Participants were then asked to rank the patterns in order of preference. **b**, This chart depicts the proportion that each image type was chosen as a first, second and third preference (with standard error bars). These results are consistent with the theory presented here, which predicts that choice-based learning generalizes most strongly to the unique, most-discriminating feature of the original choice (i.e., chosen unique), followed by the shared feature.

participation and paid immediately.

4.4.2 Results

A repository for all data described in this paper is available on OSF ([dataset] Hornsby and Love, 2019).

The proportion of times each image type was selected as a first, second and third preference is pictured in Figure 4.3b. As hypothesised, participants most preferred the unique chosen pattern, followed by the shared pattern and finally, the non-chosen pattern (omnibus non-parametric Friedman test of differences among repeated-measures $\chi^2 = 137.48, p < 0.001$). Specifically, summed preferences for the chosen-unique patterns (Median = 9, IQR = 4.0) were stronger than that for the shared items (Median = 9, IQR = 3.0) (Wilcoxon signed-

rank $Z = -2.91, p < 0.005, r = 0.09$) and the non-chosen items (Median = 11, $IQR = 5.0$) ($Z = -12.08, p < 0.001, r = 0.39$). A final test also revealed that preferences for the shared items were stronger than that for non-chosen items ($Z = -11.70, p < 0.001, r = 0.38$)².

These results supported our two key hypotheses and therefore the key claims of our theory. Firstly, participants exhibited an increased preference for the chosen-unique and shared patterns, demonstrating that they generalized their preference learned from the initial choice to the novel patterns by virtue of their association. Consistent with the behaviour of our model formalization, people appear to generalize their preferences to novel items that share attributes with choices just made.

Secondly, participants demonstrated an increased preference for the chosen-unique pattern over the shared pattern and a reduced preference for the pattern unique to the rejected option. This is consistent with discriminative accounts of learning, in that it suggests that people update their preference towards attributes that discriminate their prior choices. Studies of human categorization have shown that experiencing contrasting category exemplars causes their perceived difference to drift apart and become idealized. For example, because people are used to contrasting diet foods with high calorie foods, they are more likely to suggest celery as a prototypical diet food, even though it is extreme for its category (Davis and Love, 2010). Choosing freely between options appears to have similar effects during preference learning, in that preferences update to maximize the perceived contrast between the accepted and rejected options. This functions to maintain coherence between past choices and present preferences.

²All Wilcoxon-signed rank tests were evaluated against a Holm-Bonferroni corrected alpha value for multiple comparisons

4.5 Political beliefs become consistent with a prior vote

The experimental results presented so far have provided controlled, experimental support for our error-driven account of subjective preference formation and decision making. Outside the lab however, people usually have strong prior preferences for options, which likely interacts with their intrinsic tendency to coherency maximize.

The aim of Experiment 2 was therefore to explore the extent to which people retrospectively update their existing preferences following a free choice. Specifically, we evaluated whether people would modify their political beliefs after a vote. Participants from the U.S. were shown two political candidates and asked to vote for one, based on some trivial attributes (e.g., whether they liked cats or dogs)³. This experimental procedure is depicted in Figure 4.4a. Following the vote, these chosen and non-chosen political candidates revealed randomly assigned, opposing controversial beliefs on a particular topic. These beliefs were either traditionally left or right-wing. For example, if randomly assigned to the abortion issue, the chosen candidate would either show “Abortion: Pro-choice” or “Abortion: Pro-life”. Participants were then asked to state the extent to which they agreed with their chosen candidate’s newly revealed belief using a slider. Finally, they were asked to state their preferred party out of the Democrats or Republicans.

Our primary hypothesis was that people would show higher levels of agreement for the right-wing view (e.g., pro-choice abortion rights) if their chosen candidate revealed support for that view, compared to the left-wing view. This would provide support for the theory presented here and suggest that, following a vote, people are prone to adjust their political beliefs to be coherent with their chosen candidate.

³More information about the experimental design and stimuli is available in the methods section

While expecting that people would feel inclined to agree with their chosen candidate's newly revealed opinion, we also secondarily hypothesized that this would vary depending on individual differences. Specifically, we expected that different individuals would be more or less prone to updating their preferences retrospectively, due to differences in their longstanding political beliefs. Due to its simplicity, we used people's preferred political party affiliation as an indirect measure of these beliefs and subsequently predicted that Republican-identifying participants would be more likely to update their beliefs to be coherent with their initial vote. This was for two main reasons. Firstly, due to the broad and unpredictable nature of being in power, voters in support most likely have to accept more compromises and adjust their beliefs on occasion in order to remain coherent. Secondly, results from self-reported surveys suggest that conservatives value coherency comparatively more when making political decisions (e.g. they have been shown to value loyalty to the in-group and to authority comparatively more Mendez (2017); Haidt and Graham (2007); Jost et al. (2003)). While this is a somewhat secondary prediction that does not directly flow from our core theory, we used Experiment 2 as an initial exploration of group differences across domains.

4.5.1 Method

4.5.1.1 Participants

One thousand participants were recruited using mturk. All participants were required to have had at least 1000 of their previous tasks accepted on mturk, have a 95% mturk task acceptance rate and be based in the United States. All participants were paid 50¢ for participating. Of the participants, 47 were removed due to a lack of apparent concentration or understanding of the task. Specifically, 45 participants were removed for having response times more than two standard deviations from the mean for the vote response ($n = 26$, $> 38.68s$) and the slider response ($n = 19$, $> 55.38s$). A further 2 participants were

removed for clicking more than 50 times during the whole experiment. Of the final 953 participants, 50.26% were male, 49.53% were female and the remainder specified as ‘other’. The mean age was 37.44 ($SD = 11.67$). Of the participants, 64.53% said they affiliated more closely with Democrats, whereas the remainder said they affiliate more closely with Republicans. Data was collected in late July 2018.

4.5.1.2 Design

The experiment used a 2 x 2 between-groups design. The experiment involved one trial. Participants either chose a candidate that — previously unbeknownst to them — revealed a more left-wing view or a more right-wing view. Furthermore, participants either affiliated more closely with Democrats or Republicans. The dependent variable was the slider value of the participant, normalized by political direction. Specifically, these values ranged from 1 to 100, where higher values indicated more agreement with the right-wing view and lower values indicated more agreement with the more left-wing view. Note that the initial vertical location of the candidates, the allocation of neutral and controversial attributes, and the final horizontal locations of the candidates were all fully randomized.

4.5.1.3 Apparatus & Stimuli

The study was designed using JavaScript and was accessed in a web browser. The task was presented in a 1000 pixel wide screen. Participants were shown two of a possible four faces taken from the Chicago Face Database (Ma et al., 2015)⁴. To control for noise emerging from other, well-documented biases, all faces were controlled to be male, Caucasian and ranging between the ages of 37 and 43. Faces were randomly assigned one of four names. These were either James Smith, Michael Johnson, John Williams, Graham Brown. These

⁴These faces can be identified in the Chicago Face Database using the target codes WM-023, WM-221, WM-223 and WM-248

names were sourced from a list of the most popular first and second Caucasian names in the United States.

Table 4.1: A summary of the neutral statements shown to participants

Topic	Statements
Pets	"I like cats"
	"I like dogs"
Sport	"I'm a baseball fan"
	"I'm a basketball fan"
Food	"My favorite Italian food is pizza"
	"My favorite Italian food is spaghetti"

Table 4.2: A summary of the controversial statements revealed by candidates following a vote

Topic	Left-wing	Right-wing
Abortion	"Abortion: Pro-choice"	"Abortion: Pro-life"
Immigration	"I support policies that would increase immigration"	"I support policies that would decrease immigration"
Trade	"I support tariffs on imports"	"I support free trade"

The neutral opinions of participants are shown in Table 4.1. The controversial political opinions of participants are shown in Table 4.2.

4.5.1.4 Procedure

Participants were initially briefed about the purpose of the experiment in order to get their informed consent. They were told that the task would take about 5 minutes to complete and would require them to vote for political candidates. After the briefing, participants began the trial. To first ensure that the participants were alert, they were told to read a short experimental briefing. Within

this briefing was the instruction to click on the name of the university (which was displayed at the top of the screen). After clicking this, a new panel was revealed displaying some text ("It's time to vote for a candidate! Please click on a candidate to vote") and two political candidates in gray cards; one above the other. Within each card and underneath the candidate's photographs was their name and then a list of "My opinions". These lists initially displayed opposing neutral opinions selected from the same neutral topic, as shown in Table 4.1. The assignments of photographs, names, neutral topics and vertical alignment were all randomized. Participants were asked to vote for a candidate using this neutral information by clicking on their card. After voting, the cards then moved to a left or right position on the screen, aligning horizontally with each other. This horizontal allocation was also randomized. The voted and non-voted candidates then immediately revealed opposing political opinions from the same topic in their list of opinions. A blue "New" button drew attention to this newly-revealed belief. The window then moved down to reveal a new section, containing some text and a slider. The text informed participants in bold lettering that "Your candidate has revealed a new opinion above! How much do you agree with your candidate's newly revealed opinion?" Participants were prompted to use the new slider to state the extent to which they agreed with the candidates' newly revealed opinions. Small avatars of the candidate's faces were shown on either side of the slider, with a gray tick or white cross below the ones that were previously chosen, respectively. There were also prompts below each avatar reminding participants of the newly revealed controversial opinion of the candidate. Note that the slider was not initialized with any starting value. Sliding closer towards the chosen candidates indicated higher levels of agreement with that candidate. The slider had 100 possible positions. The photographs, names, neutral and political opinions and left/right direction of movement were all randomly assigned between participants.

After confirming the slider position, participants were asked to state which

party they most affiliated with (Democrat or Republican), their gender and their age. Following task completion, they were thanked for their participation and paid immediately.

4.5.2 Results

In support of the primary hypotheses, results revealed that participants' stated degree of belief was significantly influenced by the randomly-assigned belief revealed by their chosen candidate (non-parametric two-way analysis of variance (ANOVA) ($F(1, 952) = 28.89, p < 0.001, CL=0.563$). In particular, those that voted for a candidate that later revealed a right-wing opinion (Median = 56.0, IQR = 76.00) agreed an average of 43.59% more with the right-wing view compared to if the candidate later revealed a left-wing view (Median = 39.0, IQR = 69.75). This suggests that choosing a political candidate based on initially trivial characteristics made participants more likely to agree with that candidate's later-revealed controversial opinion, irrespective of whether the person self-identified as a Democrat or Republican.

In support of the secondary hypothesis, there was also a significant interaction between party affiliation and the degree of agreement with the right-wing view ($F(1, 952) = 11.77, p < 0.001$). A first Mann-Whitney U test looking at the responses of self-identifying Democrats revealed that they only slightly increased their level of agreement for the right-wing view if their candidate revealed a right-wing view (Median = 40.0, IQR = 78.00) compared to if it revealed a left-wing view (Median = 31.5, IQR = 65.75); this difference was not significant ($U = 43584.5, p=0.054$), and had a very low effect size ($CL = 0.518$). In contrast, the second test looking at Republican-identifying participants revealed a larger effect, in that — if one's candidate later revealed a more right-wing opinion — they were 65.66% more in agreement with the the right-wing view (Median = 82.0, IQR = 59.75) compared to when the candidate revealed a more left-wing view (Median = 49.50, IQR = 70.25) (U

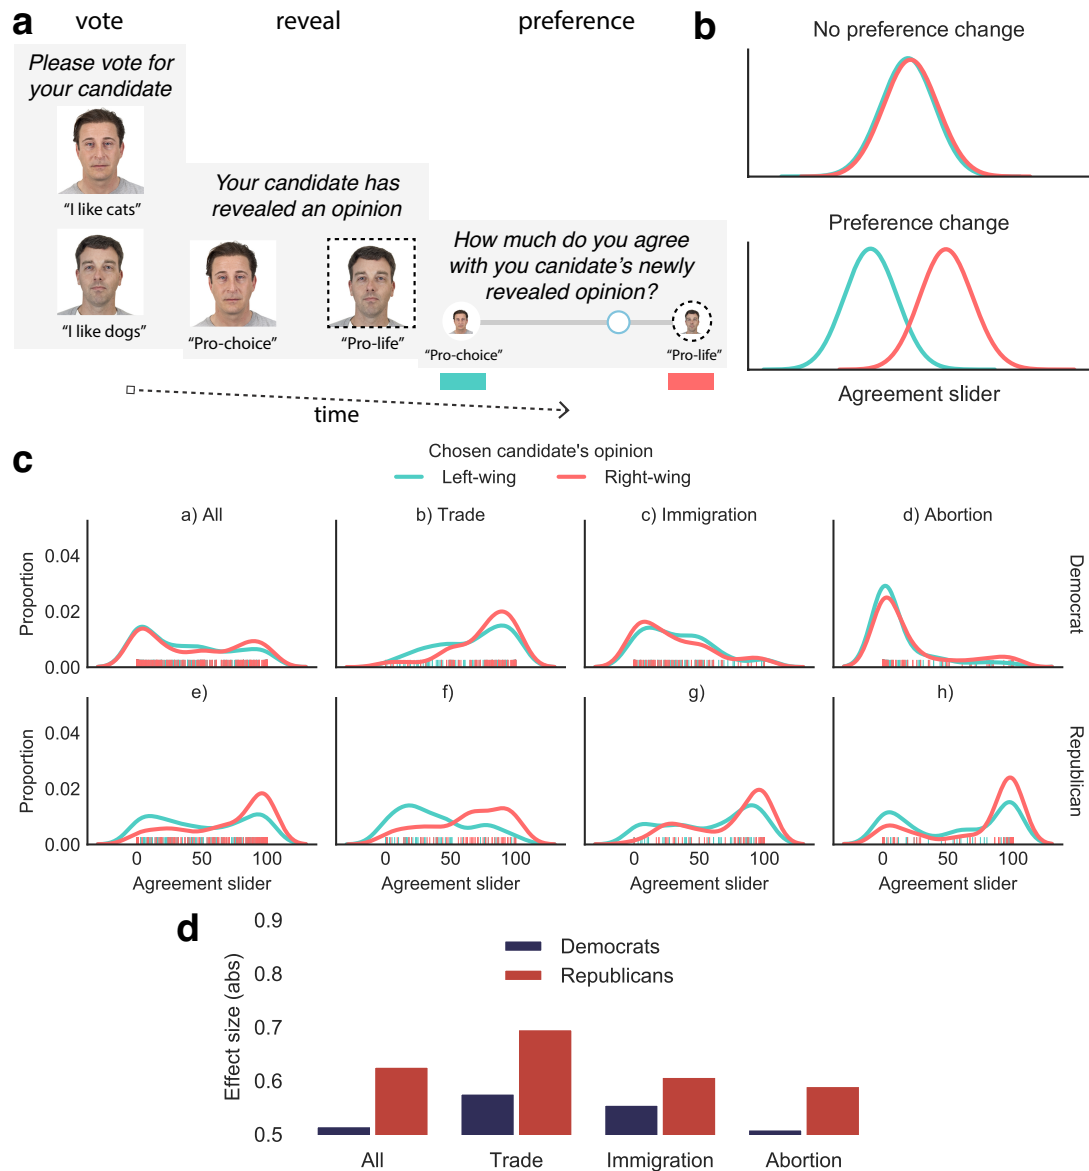


Figure 4.4: **a**, Participants were initially asked to vote for a candidate based on trivial characteristics. Both candidates then revealed opposing beliefs on a more controversial belief. Participants were then asked to rate how much they agreed with their candidate's newly revealed opinion. **b**, If voting for the candidate affects the extent to which participants agree with the later-revealed political belief, then one would expect participants to show different levels of agreement between groups, depending on whether their candidate revealed a traditionally left-wing view or a right-wing view. **c**, The results of the experiment are shown in kernel density plots (with rugplots below to depict the individual data points). These reveal that Republican-identifying participants were particularly prone to adjusting their preference to be in accordance with that revealed by their chosen candidate. **d**, Plotting effect sizes for each per-topic comparison shows that Republican-identifying participants were particularly prone to adjusting their preferences.

= 10131.5, $p < 0.001$, $CL = 0.629$)⁵. Thus, Republican-identifying participants appeared to be particularly prone to adjusting their beliefs so as to be consistent with their last choice, even when this choice was based on trivial grounds (e.g., the fact that their candidate liked cats).

It is notable that self-identified Republicans were particularly willing to adjust their preferences to be coherent with their prior choices. For example, their median support for pro-life abortion policies was 60% higher when their chosen candidate later expressed support for pro-life policies ($Median=96.0$, $IQR=61.50$) compared to pro-choice policies ($Median=60.0$, $IQR=88.00$) (the remaining within-topic comparisons are depicted in Figure 4.4c and 4.4d, and described in the appendical). One possibility is that only Republican-identifying participants adjust their preferences to be consistent with a choice, across all domains. To confirm that this was not the case, we replicated the first experiment, asking about party affiliation at the end of the study. Results revealed that both self-identifying Democrat and Republican participants exhibited preferences for the final pattern in ways consistent with the reported findings above (see the appendical for details). Thus, a more likely explanation of our results is that there are group and domain differences guiding the extent to which people learn from their past choices. For example — as discussed in the introduction to this section — differences in people’s belief systems may give rise to different propensities to adjust preferences following a choice. Clearly these topics are deserving of future investigation in light of the results presented here.

4.6 General Discussion

In this chapter, we proposed a novel account of subjective preference formation and decision making. In our model, choice options and preferences are represented in a common continuous, multi-dimensional space. When people choose

⁵For all multiple comparisons reported in this paper, significance values were compared against a Holm-Bonferroni corrected α value

options within this space, their preferences are updated to increase the likelihood of their previous choices. The objective of decision making is therefore to reduce the error between one's past choices and present preferences; we refer to this general mechanism as *coherency maximization*. Consistent with patterns of repeat-purchasing observed in supermarket consumers (Riefer et al., 2017) and studies of choice-induced preference change (Brehm, 1956; Sharot et al., 2009; Alos-Ferrer and Shi, 2012; Ariely and Norton, 2008; Schonberg et al., 2014; Koster et al., 2015; Miyagi et al., 2017; Voigt et al., 2017; Akaishi et al., 2014; Nakao et al., 2016; Vinckier et al., 2019; Chammat et al., 2017; Cockburn et al., 2014), coherency maximization boosts the likelihood of past choices being repeated by shifting preferences towards the chosen item and away from rejected alternatives. As was shown in a simple simulation, this mechanism can drive strong subjective preferences in the absence of extrinsic feedback. Thus, our model is well-suited to navigating the complex choices encountered in everyday life.

Following from this account, one would expect preferences to be exaggerated towards attributes that favored the past choice and diminished towards attributes associated with rejected choices. Results from the first experiment showed that preferences are updated in accordance with these predictions. Concretely, the more uniquely associated a novel pattern was to the back of a toy robot they had previously designed, the more likely participants were to prefer it. Given that we tend to lack direct introspective access to the mechanisms driving our behaviour, post-rationalizing choices may be a reasonable way to make sense of ourselves and the world around us (Bem, 1967, 1972; Miller and Buckhout, 1973; Mandler, 1975; Nisbett and Wilson, 1977; Cushman, 2019). Our model shows that simple learning processes can achieve similar ends by updating underlying preferences to align with a choice.

These learning rules can be seen as reflecting an internal drive for internal consistency, which is pivotal to rational models of decision making (Savage, 1954). For example, it is often assumed that subjective choices should be

stochastically transitive (Schultz et al., 1997). When preferences are not transitive, people can become liable to manipulation (or “dutch booking”) (Davidson et al., 1955). For example, if out of three beers, a person prefers beer A over B, beer B over C but would rather have C than A, their preferences are cyclical. This person could be tricked into paying for a series of costly trades in which the drinker ended with his original beer. In subjective domains, internal consistency may be at times the only rational strategy that is feasible.

The second experiment explored how the tendency to maximize coherency interacted with people’s prior preferences. The results indicated that after voting for a political candidate based on trivial criteria (e.g., the candidate likes cats), participants were more likely to agree with a controversial opinion later revealed by the candidate, such as their stance on abortion rights. This held up as a main effect, supporting our claim that people retrospectively update their preferences over attributes of their past choices to make them more likely. The fact that this was particularly pronounced for Republican-identifying participants supported our secondary hypothesis that this mechanism can vary between groups, consistent with previous studies demonstrating individual differences in choice-induced preference change (Cockburn et al., 2014). However, while individual differences may partly explain our results, it is likely that other variables outside the scope of our theory also influence decisions, such as how important affiliation is to different groups. Specifically, while Republican-identifying participants adjusted their preferences more than their Democrat-identifying counterparts in the second experiment, a replication of the first experiment (see appendix) revealed no difference between the groups when choosing between novel options. In future research, the prominence of choice-induced preference change within certain groups or domains could be estimated by fitting the CDC model to different decision tasks.

If people update their preferences to be coherent with their past choices, then why might they also feel motivated to explore? A recent analysis of consumers’ take-away purchases suggested that they were more likely to try

a different restaurant after a positive experience rather than repeat it again (Schulz et al., 2019), suggesting that people may also explore to reduce uncertainty about their environment. In the simulation, we used a stochastic, undirected exploration policy (i.e. ϵ -greedy). However, CDC could be adapted to use a more sophisticated, directed exploration strategy, such as uncertainty minimization. This is because CDC considers exploration and coherency maximization as theoretically and mechanistically independent. Understanding where and when people explore is an open question in the literature (for a review, see Hills et al. 2015), meaning that such adaptations would need to be evaluated with scrutiny. In the future, we hope to further understand how people trade-off this need to explore with the desire to coherency maximize.

In this chapter, we have assumed that preferences are adjusted following a choice. However, preference change has also been shown to occur “online” during a choice (Schonberg et al., 2014; Niv et al., 2015; Akaishi et al., 2016; Voigt et al., 2019). A recent study by Voigt et al. (2019) found that choice-induced preference changes only occurred for choices that were remembered. Both online and post-choice induced preference change mechanisms could co-exist. Though this would highlight a future area of development for the CDC model. One such modification would be to adapt the role of the attention weights during a choice. For example, their influence could be magnified in cases where the model was more familiar with or remembered the attributes of the choice. The preference vectors and attention weights would then account for both online and post-decisional effects of preference change, respectively.

A core idea formalised within CDC is that preferences should update to discriminate the past choice. Concretely, preferences should move towards attributes associated with chosen option and away from attributes associated with the rejected option. Results from Experiment 1 supported this hypothesis, with participants preferring patterns to the extent that they were uniquely associated with their previous choice. Whilst results in Experiment 2 were suggestive of this effect, one cannot be certain of the extent to which preferences

within individuals changed as a result of the choice. For example, one possibility is that participants only moved significantly towards the beliefs of their chosen candidate's when they aligned with the participants' prior ideology. Future studies could consider taking a baseline measurement of participants' affiliation within individuals to compare before and after a choice. However, one would need to be careful that initial ratings were measured reliably and did not suffer from the same confounds identified by Chen and Risen (2010).

Although our studies involved brief decision-making episodes, the basic mechanisms considered here may also apply at longer timescales. Indeed, this work was partially motivated by Riefer et al.'s 2017 discovery of self-reinforcing purchasing patterns in supermarket consumers, which extended over several months. If CDC's predictions for how preferences change held over extended periods of time, the practical consequences and possibilities for behavioural change would be substantial.

For instance, CDC predicts that purchasing the same *type* of food should increase preference for associated attributes. This can be problematic in cases where the food is unhealthy (e.g., high in sodium or saturated fat). For example, studies of nutrition have shown that repeated exposure to a particular ingredient (e.g., sodium) increases one's desire and lowers their sensitivity to it, making it difficult for them to adjust when they go on a diet (Bertino et al., 1982; Mattes, 1993; Liem and de Graaf, 2004). The advent of targeted recommender systems means that these cyclical effects may be being perpetuated further. While recommender systems often reduce diversity of alternatives in the environment (Pariser, 2011), coherency maximization causes preferences to become less diverse, thereby perpetuating the problem. Rather than blindly targeting people based on their previous consumption, marketers could incorporate external objectives to their targeting algorithms. For example, our theory predicts that recommending healthy alternatives that are similar to people's existing preferences could lead to long-term improvements in their choices.

The modeling approach presented here is readily extended to account for the richness of people's preferences. Unlike how CDC was formalized here, people are unlikely to have a single preference across each respective attribute. For example, while a foodie might prefer expensive, locally produced foods, they may also be happy to watch affordable, mass produced television. People may have different preferences for the same attribute (e.g., cost) depending on context. Fortunately, it would be straightforward to extend CDC to have multiple preference vectors to capture this context dependence, similar to how models of human category learning possess multiple clusters in which only the most contextually relevant one is updated during a learning episode (Love et al., 2004). Such a model would cast coherency maximization as a process that occurs within a domain (e.g., food, entertainment, etc.) as opposed to globally.

The model presented here may also be extended to account for some well-known decision biases. For example, CDC predicts that people will prefer novel options when they are similar to options that have repeatedly tried in the past. Thus, it may be able to account for the mere-exposure effect, where people have been shown to prefer options by virtue of their familiarity (Zajonc, 1968). Similarly — though only distantly related to the free choices described here — it is possible that CDC could be adapted to account for anchoring effects. For example, forcing CDC to update its preferences towards a given anchor would cause it to become more favorable towards related options (see attitudinal change accounts of anchoring for related arguments, e.g., Wegener et al. 2001). While such relationships are speculative at this stage, they indicate a general, coherency-driven learning mechanism may underpin several, well-known sequential choice biases.

In conclusion, preferences and choices can be characterized as existing within a common space. While we prefer options that match our preferences, we also appear to engage in error-driven learning to update our preferences to accommodate our past choices. Because preference and choice representation

lie in a shared multidimensional space, the choices we make have consequences (i.e., spillover effects) for related future choices. For example, people may be prone to agree with a controversial opinion held by a political candidate, by virtue of the fact that they voted for them. Although this behavior may appear irrational, being internally coherent may be the best we can hope for in complex, subjective domains. Being aware of these coherency maximizing dynamics may make it possible for people to ameliorate some of the potentially harmful consequences. For example, if a voter chooses a candidate based on tax policy, perhaps being mindful of that fact will make the voter less likely to reflexively adopt their candidate's positions on unrelated issues. Likewise, these same coherency maximizing principles could be incorporated into recommender systems to help consumers achieve some positive goal, such as eating more healthily.

Bibliography

- Akaishi, R., Kolling, N., Brown, J. W., and Rushworth, M. (2016). Neural mechanisms of credit assignment in a multicue environment. *Journal of Neuroscience*, 36(4):1096–1112.
- Akaishi, R., Umeda, K., Nagase, A., and Sakai, K. (2014). Autonomous mechanism of internal choice estimate underlies decision inertia. *Neuron*, 81(1):195 – 206.
- Alos-Ferrer, C. and Shi, F. (2012). Choice-induced preference change and the free-choice paradigm: A clarification. *SSRN*.
- Ariely, D. and Norton, M. I. (2008). How actions create – not just reveal – preferences. *Trends in Cognitive Sciences*, 12(1):13 – 16.
- Bem, D. J. (1967). Self-perception: An alternative interpretation of cognitive dissonance phenomena. *Psychological Review*, 74(3):183–200.

- Bem, D. J. (1972). Self-perception theory. volume 6 of *Advances in Experimental Social Psychology*, pages 1–62. Academic Press.
- Bertino, M., Beauchamp, G. K., and Engelman, K. (1982). Long-term reduction in dietary sodium alters the taste of salt. *The American Journal of Clinical Nutrition*, 36(6):1134–1144.
- Brehm, J. W. (1956). Postdecision changes in the desirability of alternatives. *Journal of Abnormal and Social Psychology*.
- Chammat, M., Karoui, I. E., Allali, S., Hagège, J., Lehongre, K., Hasboun, D., Baulac, M., Epelbaum, S., Michon, A., Dubois, B., Navarro, V., Salti, M., and Naccache, L. (2017). Cognitive dissonance resolution depends on episodic memory. *Scientific Reports*, 7:41320.
- Chen, M. K. and Risen, J. L. (2010). How choice affects and reflects preferences: Revisiting the free-choice paradigm. *Journal of Personality and Social Psychology*, 99:573–594.
- Cockburn, J., Collins, A. G. E., and Frank, M. J. (2014). A Reinforcement Learning Mechanism Responsible for the Valuation of Free Choice. *Neuron*, 83(3):551–557.
- Crump, M. J. C., McDonnell, J. V., and Gureckis, T. M. (2013). Evaluating amazon’s mechanical turk as a tool for experimental behavioral research. *PLOS ONE*, 8(3):1–18.
- Cushman, F. (2019). Rationalization is rational. *Behavioral and Brain Sciences*, page 1–69.
- [dataset] Hornsby, A. N. and Love, B. C. (2019). How decisions and the desire for coherency shape subjective preferences over time.
- Davidson, D., McKinsey, J. C. C., and Suppes, P. (1955). Outlines of a formal theory of value, i. *Philosophy of science*, 22(2):140–160.

- Davis, T. and Love, B. C. (2010). Memory for category information is idealized through contrast with competing options. *Psychological Science*, 21(2):234–242.
- Daw, N. and Tobler, P. (2013). *Value Learning through Reinforcement: The Basics of Dopamine and Reinforcement Learning*. Elsevier.
- DeSarbo, W. S. and Kim, S. (2012). *A review of the major multidimensional scaling models for the analysis of preference/dominance data in marketing*, pages 3–25.
- Glimcher, P. W. (2009). Chapter 32 - choice: Towards a standard back-pocket model. In Glimcher, P. W., Camerer, C. F., Fehr, E., and Poldrack, R. A., editors, *Neuroeconomics*, pages 503 – 521. Academic Press, London.
- Goodfellow, I., Bengio, Y., and Courville, A. (2016). *Deep Learning*. MIT Press.
- Greenhoff, K. and MacFie, H. J. H. (1994). *Preference mapping in practice*, pages 137–166. Springer US, Boston, MA.
- Haidt, J. and Graham, J. (2007). When morality opposes justice: Conservatives have moral intuitions that liberals may not recognize. *Social Justice Research*, 20(1):98–116.
- Hall, L., Johansson, P., Tärning, B., Sikström, S., and Deutgen, T. (2010). Magic at the marketplace: Choice blindness for the taste of jam and the smell of tea. *Cognition*, 117(1):54–61.
- Hills, T. T., Todd, P. M., Lazer, D., Redish, A. D., Couzin, I. D., Group, C. S. R., et al. (2015). Exploration versus exploitation in space, mind, and society. *Trends in cognitive sciences*, 19(1):46–54.
- Hinton, G. E., McClelland, J. L., and Rumelhart, D. E. (1986). Parallel distributed processing: Explorations in the microstructure of cognition, vol. 1.

- chapter Distributed Representations, pages 77–109. MIT Press, Cambridge, MA, USA.
- Horton, J. and Chilton, L. (2010). The labor economics of paid crowdsourcing. *CoRR*, abs/1001.0627.
- Izuma, K., Matsumoto, M., Murayama, K., Samejima, K., Sadato, N., and Matsumoto, K. (2010). Neural correlates of cognitive dissonance and choice-induced preference change. *Proceedings of the National Academy of Sciences*, 107(51):22014–22019.
- Izuma, K. and Murayama, K. (2013). Choice-induced preference change in the free-choice paradigm: A critical methodological review. *Frontiers in Psychology*, 4:41.
- Jost, J., Glaser, J., Kruglanski, A., and Sulloway, F. (2003). Political conservatism as motivated social cognition. *Psychological Bulletin*, 129(3):339–375.
- Koster, R., Duzel, E., and Dolan, R. J. (2015). Action and valence modulate choice and choice-induced preference change. *PLOS ONE*, 10(3):1–10.
- Kruschke, J. K. (1992). ALCOVE: An exemplar-based connectionist model of category learning. *Psychological Review*, 99:22–44.
- Liem, D. G. and de Graaf, C. (2004). Sweet and sour preferences in young children and adults: role of repeated exposure. *Physiology Behavior*, 83(3):421–429.
- Love, B. C., Medin, D. L., and Gureckis, T. M. (2004). SUSTAIN: a network model of category learning. *Psychol Rev*, 111(2):309–332.
- Ma, D. S., Correll, J., and Wittenbrink, B. (2015). The chicago face database: A free stimulus set of faces and norming data. *Behavior Research Methods*, 47(4):1122–1135.

- Mandler, G. (1975). Consciousness: Respectable, useful, and probably necessary.
- Mattes, R. D. (1993). Fat preference and adherence to a reduced-fat diet. *The American Journal of Clinical Nutrition*, 57(3):373–381.
- McDuffie, D. (2019). Using amazon’s mechanical turk: Benefits, drawbacks, and suggestions. *APS Observer*, 32(2).
- Mendez, M. F. (2017). A neurology of the conservative-liberal dimension of political ideology. *The Journal of neuropsychiatry and clinical neurosciences*, 29 2:86–94.
- Miller, G. A. and Buckhout, R. (1973). *Psychology: The science of mental life*. Harper & Row.
- Miyagi, M., Miyatani, M., and Nakao, T. (2017). Relation between choice-induced preference change and depression. *PLoS ONE*, 12(6):1–14.
- Nakao, T., Kanayama, N., Katahira, K., Odani, M., Ito, Y., Hirata, Y., Nasuno, R., Ozaki, H., Hiramoto, R., Miyatani, M., and Northoff, G. (2016). Post-response $\beta\gamma$ power predicts the degree of choice-based learning in internally guided decision-making. *Scientific Reports*, 6(August):1–9.
- Nakao, T., Ohira, H., and Northoff, G. (2012). Distinction between Externally vs. Internally Guided Decision-Making: Operational Differences, Meta-Analytical Comparisons and Their Theoretical Implications. *Frontiers in Neuroscience*, 6:31.
- Nisbett, R. E. and Wilson, T. D. (1977). Telling more than we can know: Verbal reports on mental processes. *Psychological review*, 84(3):231.
- Niv, Y., Daniel, R., Geana, A., Gershman, S. J., Leong, Y. C., Radulescu, A., and Wilson, R. C. (2015). Reinforcement learning in multidimensional environments relies on attention mechanisms. *Journal of Neuroscience*, 35(21):8145–8157.

- Nosofsky, R. M. (2011). The generalized context model: An exemplar model of classification. *Formal approaches in categorization*, pages 18–39.
- Pariser, E. (2011). *The Filter Bubble: What the Internet Is Hiding from You*. Penguin Group, The.
- Ramscar, M., Yarlett, D., Dye, M., Denny, K., and Thorpe, K. (2010). The effects of feature-label-order and their implications for symbolic learning. *Cognitive science*, 34(6):909–957.
- Rescorla, R. A. and Wagner, A. R. (1972). A theory of pavlovian conditioning: Variations in the effectiveness of reinforcement and nonreinforcement. *Classical conditioning II: Current research and theory*, 2:64–99.
- Riefer, P. S., Prior, R., Blair, N., Pavey, G., and Love, B. C. (2017). Coherency-maximizing exploration in the supermarket. *Nature Human Behaviour*.
- Sauerland, M., Sagana, A., Siegmann, K., Heiligers, D., Merckelbach, H., and Jenkins, R. (2016). These two are different. Yes, they’re the same: Choice blindness for facial identity. *Consciousness and Cognition*, 40:93–104.
- Savage, L. J. (1954). *The Foundations of Statistics*. Wiley Publications in Statistics.
- Schonberg, T., Bakkour, A., M Hover, A., Mumford, J., Nagar, L., Perez, J., and A Poldrack, R. (2014). Changing value through cued approach: An automatic mechanism of behavior change. *Nature neuroscience*, 17.
- Schultz, W., Dayan, P., and Montague, P. R. (1997). A neural substrate of prediction and reward. *Science*.
- Schulz, E., Bhui, R., Love, B. C., Brier, B., Todd, M. T., and Gershman, S. J. (2019). Structured, uncertainty-driven exploration in real-world consumer choice. *Proceedings of the National Academy of Sciences*, page 201821028.

- Sharot, T., De Martino, B., and Dolan, R. J. (2009). How Choice Reveals and Shapes Expected Hedonic Outcome. *Journal of Neuroscience*.
- Sharot, T., Velasquez, C. M., and Dolan, R. J. (2010). Do decisions shape preference?: Evidence from blind choice. *Psychological Science*, 21(9):1231–1235. PMID: 20679522.
- Strandberg, T., Sivén, D., Hall, L., Johansson, P., and Pärnamets, P. (2018). False beliefs and confabulation can lead to lasting changes in political attitudes. *Journal of Experimental Psychology: General*, 147(9):1382–1399.
- Sutton, R. and Barto, A. (1998). Reinforcement Learning: An Introduction. *IEEE Transactions on Neural Networks*, 9(5):1054–1054.
- Vinckier, F., Rigoux, L., Kurniawan, I. T., Hu, C., Bourgeois-Gironde, S., Daunizeau, J., and Pessiglione, M. (2019). Sour grapes and sweet victories: How actions shape preferences. *PLOS Computational Biology*, 15(1):1–24.
- Voigt, K., Murawski, C., and Bode, S. (2017). Endogenous formation of preferences: Choices systematically change willingness-to-pay for goods. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 43(12):1872.
- Voigt, K., Murawski, C., Speer, S., and Bode, S. (2019). Hard decisions shape the neural coding of preferences. *Journal of Neuroscience*, 39(4):718–726.
- von Neumann, J., Morgenstern, O., and Rubinstein, A. (1944). *Theory of Games and Economic Behavior (60th Anniversary Commemorative Edition)*. Princeton University Press.
- Wegener, D. T., Petty, R. E., Detweiler-Bedell, B. T., and Jarvis, W. B. G. (2001). Implications of attitude change theories for numerical anchoring: Anchor plausibility and the limits of anchor effectiveness. *Journal of Experimental Social Psychology*, 37(1):62–69.
- Zajonc, R. B. (1968). Attitudinal effects of mere exposure. *Journal of personality and social psychology*, 9(2p2):1.

Chapter 5

General Conclusions

This thesis tells a story about long-term memory before, during and after a preferential choice. Prior to a choice, our prior experience shapes perceptions of option similarity and guides inferences about the properties of novel options. I uncover this relational structure from consumer choice history in Chapter 2. Abstracting knowledge in this way allows us to achieve common tasks that precede preferential choices, such as reasoning about the complementarity of options (e.g., goes well in a salad). During decision-making, the proximity of options within these relational spaces may guide memory retrieval and thus manifestly influence open-ended decisions. In Chapter 3, I showed how the proximity of options across episodic and semantic memory influenced the order in which items were chosen in an open-ended task. Following a preferential choice, we do not benefit from a clear extrinsic reward signal, leaving them open to subjective interpretation. In Chapter 4, I discussed how an internal desire to maximise coherency between ones past choices and present preferences can drive preferences for novel options, by virtue of their similarity with the prior choice. Whilst charting the journey of a preferential choice, this thesis has discussed several important topics, such as; how we navigate large option spaces, how prior experience is leveraged during preferential choice, how preferences are learned in the absence of extrinsic reward and how choices differ between the lab and more naturalistic settings.

5.1 Preferential choice in large option spaces

A major aim of this thesis has been to understand how decision-makers navigate large option spaces. In environments like the supermarket, consumers must decide between tens of thousands of products. In these cases, it would be intractable to represent each option separately, as they might be in traditional models of stimulus-response learning (Thorndike, 1927; Schultz et al., 1997; Schultz, 2006; Rangel et al., 2008) or free-choice (Brehm, 1956; Akaishi et al., 2014; Nakao et al., 2012). A solution proposed here is to learn lower-dimensional representations of options; casting them into a connected, multidimensional attribute space. By representing options in this way, decision makers reduce the dimensionality of large option spaces, helping them to generalise properties to novel options that they've never tried.

Much is known about how explicitly presented attributes influence preferential choices (Glimcher and Rustichini, 2004; Busemeyer and Rieskamp, 2014; Rangel et al., 2008) but relatively little is known about how options are mentally represented in their absence. In Chapter 2, I proposed a method for estimating conceptual representations of consumers through observations of their preferential choices. A distributed model of semantic memory trained to predict co-occurrences of products in shopping baskets recovered representations that closely resembled those used by shoppers. Interestingly, this uncovered representations that centered around goals (e.g., stir-fry). Dimensionality reduction inevitably leads to a loss of information, so centering representations around goals ensures that they are maximally useful for making actions in large option spaces.

In Chapter 3, I showed how representations of similarity are leveraged by consumers in open-ended tasks, where the scale of the option space is so large that it would be impossible to consider all options at once. In open-ended tasks like online grocery shopping, consumers must first retrieve options from memory and our results showed that sequential purchases could be predicted by

their similarity with the prior purchase across multiple memory systems. This builds on classical findings from memory retrieval literature (see Kahana, 2020, for a review) and related findings showing that retrieval of options for open-ended questions can be predicted by their similarity across semantic memory (Aka and Bhatia, 2021; Zhao et al., 2022). The findings in this thesis uniquely describe how memory retrieval dynamics unfold when making sequential, open-ended choices in a real-world task.

In Chapter 4, I presented a theory describing how subjective preferences may be learnt and used when making decisions in large option spaces. In particular, by learning preferences over attributes of choices, people are able to generalise preferences to new options by virtue of their similarity to those tried previously. This builds on past findings showing that preferences increase towards previously chosen options (Brehm, 1956; Akaishi et al., 2014; Riefer et al., 2017; Sharot et al., 2009, 2010) and explains the novel experimental finding that this can generalise to untried options. Namely, because preferences are updated over *attributes* to make the past choice most likely, preferences spill over to options that are similar to the previous option. Results from two laboratory experiments confirmed that people indeed prefer novel options by virtue of their association with a prior choice. Making decisions and updating preferences across choice attributes helps when inferring the subjective value of untried options, which is an important skill in large option spaces.

Together, these findings demonstrate how multidimensional representations of options held in the minds of consumers facilitate naturalistic preferential choice. Representing options in a similarity space reduces the dimensionality of large option spaces, helping decision-makers to generalise properties to novel options that they've never tried. Because the space of possible options and attributes is often infinite in everyday tasks, conceptual representations play an important role in memory-based preferential choice, influencing both which options and attributes are retrieved and thus considered during decision-making. The content and use of these representations may differ between

individuals, perhaps due to differences in past choice trajectories and these differences may lead to noticeable variations in behaviour, such as increasing one's propensity to forget products.

5.2 Long-term memory and preferential choice

A related aim of this thesis was to understand how long-term knowledge is exploited when making preferential decisions in large option spaces. In Chapter 1, I reviewed how lab paradigms commonly used to study preferential choice — such as *decisions-by-description* and *stimulus-based choices* — tend to diminish the influence of prior knowledge. I then reviewed how prior knowledge may influence choices when attributes or options are not presented explicitly; namely, how attributes, options and prior experience are retrieved and how retrospective memory may be biased in favour of prior choices.

Semantic knowledge appears to guide consumers when making preferential choices in large option spaces. For example, in Chapter 2, I demonstrated how semantic properties of options could be learnt by observing their relationships with other options; casting options into a distributed semantic space. In Chapter 3, I showed how relationships within this space could influence the sequential choices of consumers when shopping for groceries online. Consumers often chose products that were nearby within this semantic space, as if they were mentally traversing it. Shoppers that relied strongly on semantic memory were more prone to add items to their basket that they didn't otherwise need, which relates to patterns of confabulation demonstrated in studies of list learning (Deese, 1959; Roediger and McDermott, 1995). These results suggest that shoppers learn semantic representations from interactions between options in the wild and exploit these in open-ended tasks.

Episodic knowledge can also influence consumer choice, which was highlighted in Chapter 3. Episodic and semantic representations explained unique variance when predicting sequential choices and their response times, indicat-

ing that multiple systems of long-term memory jointly contributed towards the retrieval of options. Options that are close episodic associates tend to be more complementary, in that they are more likely to be combined in pursuit of a goal. This may explain why shoppers relying more on episodic associations forgot fewer products on average. These findings highlight the unique but complementary nature of long-term memory systems in open-ended preferential choice tasks. It remains to be seen how many associative spaces are necessary to explain choices and how they map to regions in the brain; I hope future work will explore this question.

Models originally developed to describe properties of long-term memory and retrieval can be adapted to predict memory-based choices. In Chapter 2 and 3, we used models originally intended to estimate semantic representations from language to estimate conceptual representations of grocery products (Blei et al., 2003; Griffiths et al., 2007). In Chapter 3, we combined a semantic representation of option associations with a model of memory retrieval originally developed to explain list recall (Raaijmakers and Shiffrin, 1980); pairing these systems allowed us to predict open-ended, sequential consumer choices. In Chapter 4, we developed a model of preference learning and decision-making that was strongly influenced by models of human category learning (Kruschke, 1992; Nosofsky, 2011), with representations updating to discriminate between the chosen and rejected options, as they might during category learning (Davis and Love, 2010). This work demonstrates how models of learning and memory can be used to make explicit predictions about memory-based preferential choices and I am excited to see more of such investigations in the future.

5.3 Learning from decisions in the absence of extrinsic reward

Perhaps the best known framework for modelling memory-based decision-making is reinforcement learning (RL) (Sutton and Barto, 1998). Yet — to

date — few RL models have formalized this process in the context of subjective decisions. In Chapter 4, we proposed a new computational cognitive model which characterizes how people make subjective decisions and update their preferences over time. The probability of a choice is determined by how similar choice options (e.g., pizza) are to the agent’s preference vector, where similarity is a function of attention-weighted distance such that some attributes (e.g., taste) can be weighted more than others (e.g., calories). Preferences are updated by gradient-descent learning rules that make repeating related choices (e.g., pizza over salad) more likely in the future by adjusting attention weights and the position of the preference vector. These learning rules maximize coherency by aligning preferences to match choices, a well-documented finding within the psychological literature of free choice (Brehm, 1956; Sharot et al., 2009, 2010; Alos-Ferrer and Shi, 2012; Akaishi et al., 2014). This is validated by simulation and behavioral experiments with humans. People updated their preferences and generalized to similar choices in a manner consistent with the model.

In many ways, the CDC model proposed in Chapter 4 is comparable to an RL agent. It calculates a value function of different actions (using function approximation), uses this to determine which action to take, updates its preferences subsequent to a choice and repeats this process over multiple timesteps. In standard models of RL, the reward signal is objective and drives learning; this is known as the *reward hypothesis*. However, CDC contrasts in that it does not assume a reward is emitted by an “oracle” in the environment Sutton and Barto (1998). Whilst this may be present in computer games, rewards outside of such artificial environments are not specified. For example, no objective rewards are associated with choosing to eat a pizza. Thus, CDC’s reward function is generated intrinsically to align with whatever was last chosen. As a result, CDC updates its preferences in the absence of an environmental reward, which in turn may alter future valuations of choices.

The desire to maximise coherency between one’s preferences and past

choices may co-exist alongside a host of other intrinsic motivations. Most primitively, species are motivated to preserve homeostasis; seeking food when they're hungry and rest when they're tired. Thus, other intrinsic reward functions may exist to motivate self-regulation (Singh et al., 2010; Keramati and Gutkin, 2014). Coherency maximisation could be viewed as a form of self-regulation over time, albeit at a higher cognitive level. However, this motivation may not fully account for all subjective decision making. For example, a large part of learning in childhood is driven by play, suggesting that people are also driven by higher-level drives, such as curiosity (Oudeyer, 2018; Kidd and Hayden, 2015; Gruber et al., 2014; Schmidhuber, 2007; Burda et al., 2018). These motivations may influence the exploration policies of individuals or may be used as a basis for learning; many fields would benefit from a deeper understanding of these interactions in the future.

5.4 Bidirectional influences between choices and representations

Throughout this thesis, I have discussed the bidirectional influence between choices and prior knowledge. Preferential choices appear to shape conceptual knowledge and preferences and appear to influence how representations are queried subsequently. This can lead to complex behavioural dynamics that unfold over time, such as semantic clustering of sequential choices or alignment of preferences towards past choices.

For instance, one interesting question raised by Chapter 2 is whether shopping activity changes conceptual organization or conceptual organization drives shopping behaviour. Our results cannot definitely answer this question, but the likely answer is that the influence is bidirectional. For example, having a concept like *stir-fry* should cause certain items to be purchased together to fulfill the common goal. Likewise, ingredients in the same dish may come to be viewed as more similar over time, consistent with laboratory studies that

find that linking objects makes them more similar (Jones and Mewhort, 2007).

In Chapter 4 our model of preference learning and decision-making captured the bidirectional relationship between preferences and past choices. Because decision-makers adjust their preferences to discriminate their past choices, this can drive strong subjective preferences over time in the absence of any extrinsic feedback, particularly towards attributes that discriminate in the choices made. This was demonstrated in the model simulation, lab experiments and elsewhere in the reduced product explorations described by Riefer et al. (2017).

Relatedly, Chapter 3 highlighted a form of serial dependence in memory-based choice, where behaviour appears to drift towards the recent past. Model fits revealed that choices were best predicted by their similarity with the prior choice in memory. This aligns with disparate research demonstrating serial dependence in human cognition, such as visual processing (Fischer and Whitney, 2014; Cicchini et al., 2017), perceptual decisions (Braun et al., 2018), memory (Kiyonaga et al., 2017) and decision making (Tversky and Kahneman, 1974; Ariely et al., 2003). For example, visual perception of grating orientations tend to be biased towards those most recently experienced (Fischer and Whitney, 2014). This bias has been exploited in visual illusions where — for example — changes to a slowly ageing face tend to go unnoticed (Manassi and Whitney, 2022). The physical world is largely stable and continuous, making the recent past a good predictor of the present (Dong and Atick, 1995). A reasonable strategy in noisy domains that require repeated choices — such as the supermarket — may be to choose similarly to before.

5.5 Studying choices outside of the lab

This thesis builds on a growing trend of using large, naturally occurring datasets of human behaviour to enhance our understanding of human cognition (Paxton and Griffiths, 2017). For example, in Chapter 2, we used nearly

1.3m real shopping transactions to infer the semantic representations of supermarket shoppers. In Chapter 3, we predicted nearly 5m sequential online grocery purchases using representations trained on a separate dataset of 4.3m shopping transactions. In Chapter 4, we developed a theory of preference learning inspired by the longitudinal exploration patterns of over 280,000 real supermarket shoppers (Riefer et al., 2017). In addition to the high ecological validity afforded by these datasets, they have helped us to make important theoretical and practical advances.

For example, studying naturalistic choices has allowed us to challenge status quos established in laboratory studies. In Chapter 2, the representations of food recovered from consumer’s choices were thematic and goal-directed, which contrasted with the more taxonomic organisations recovered from lab studies that used free-sorting (Murphy and Ross, 1999; Ross and Murphy, 1999). Instead of defaulting to taxonomic organisation when representing food (as suggested by Murphy and Ross, 1999), this may instead be an artifact of tasks that use a small initial set size (Lawson et al., 2017). In Chapter 3, I discussed how choices made in the absence of an extrinsic reward signal may exhibit a reduced tendency to explore over time. This is the opposite of what would be expected of a rational agent making decisions in a nonstationary environment and contrasts with human behaviour observed in laboratory studies of reinforcement learning (Sutton and Barto, 1998; Knox et al., 2012; Blanco et al., 2013). By combining laboratory studies, computational cognitive modelling and analyses of naturalistic behaviour, we arguably form a more complete understanding of human decision-making.

Practically, we have demonstrated a new way to estimate semantic representations of options. Notably, in Chapters 2 and 3 we trained distributed representations of grocery products from the co-occurrences of products in till receipts and validated these using controlled experiments and out-of-sample predictions. This may offer a way forward for researchers wishing to use naturalistic semantic representations of choice options (Aka and Bhatia, 2021;

Bhatia and Stewart, 2018), particularly those wanting to avoid learning representations from corpora of natural language. In contrast to natural language, grocery data does not require heavy pre-processing, does not require any domain-expertise to prepare and may suit the assumptions of certain algorithms better (e.g., in-store transactional data is unordered). In addition, choices are arguably more representative of human activity compared to the curated language often captured in text datasets (e.g., newspapers and websites). Choices may be more variable between individuals, meaning that they may be better suited for evaluating questions about individual differences, such as predicting endogenous demographics or memory decline. By making this data and the associated representations available online, I hope to stimulate more of such discoveries (Hornsby and Love, 2022).

5.6 General limitations

5.6.1 Asymmetric similarity

Throughout this thesis I have argued that consumers represent options in a large, connected, multidimensional space. For example, in Chapter 3, I used `word2vec` to develop a metric space of options, where each option is embedded within a multidimensional space. A challenge with this spatial view is that it can be seen to make restrictive assumptions about how people perceive relationships between options (Tversky, 1977). For example, metric spaces assume that people think symmetrically about the similarity between options. Yet human judgments (Tversky and Gati, 1982) and free associations (Griffiths et al., 2007) routinely violate metric axioms such as symmetry. Being cued with one word in a pair may give rise to asymmetric responses (e.g., cream may be retrieved after strawberries but not vice versa), so it may be unreasonable to expect that consumers think this way.

Despite first appearances, the models used in Chapter 2 and Chapter 3 actually do not assume a symmetrical relationship between options. Firstly,

the Topic model used in Chapter 2 represents each option as a probability distribution and computes the association between options as the probability of one word given another. This allows it to make different predictions depending on the option that it is being conditioned on (Griffiths et al., 2007). In addition, the semantic memory representation used in Chapter 3 (more precisely represented as a symmetric cosine similarity matrix) is not supposed to be a model of sequential retrieval; rather, it is a memorial representation that is operated on by a retrieval model. As originally shown by Jones et al. (2018), such a model operating on a similarity space can produce asymmetries. Concretely, whilst the distance between two points is obviously equal in either direction of a metric space, the density of the landscape is not; if one ranks the similarities of products to strawberries, cream will presumably appear high in that ranking, whereas strawberries could be low in the rank of items considered most similar to cream. In the retrieval model, the similarities of options in the denominator vary depending on the ratio of similarity to other competitors. Thus a simple choice rule operating on a symmetric similarity space can give rise to asymmetric relationships when they are not encoded within the space to begin with. Whilst I do not wish to make any strong claims about the relative plausibility of topic models or metric spaces, I hope to emphasise that neither study presupposes a symmetric relationship between options.

5.6.2 Separating memory representations from process

Another well known challenge when modelling memory retrieval is that it can be difficult to separate a representation from the process that operates on it, particularly when the representation is estimated using human behaviour (Anderson, 1990). For example, the same behavioural characteristics of semantic fluency have been explained by a random walk over a network graph (Abbott et al., 2015) and a strategic switching process reminiscent of optimal foraging (Hills et al., 2012). Whilst the random walk model would seem more parsimonious, it has been argued that the free association norms used to build

the semantic network also embedded information about the retrieval process (Jones et al., 2015). The representations used in Chapter 2 were developed on a separate dataset of in-store transactions in order to be sufficiently different from online shopping (e.g., in-store receipts are unordered). Moreover, explaining sequential choices through different information sources arguably provides clearer resolution than free-associations, allowing for the prediction of individual differences, such as the propensity to forget. Nevertheless, future research may wish to explore different techniques for estimating representations, such as via choice sequences directly (e.g., Zemla et al., 2016). Whichever technique is used, these results make clear the importance of capturing different sources of knowledge when explaining sequential choice.

5.6.3 Causal limitations of observational studies

Throughout this thesis, I have drawn inspiration from large datasets of consumer choices to motivate an improved understanding of how people make decisions outside of the lab. Generally however, it can be difficult to draw causal inferences from large observational datasets due to a lack of experimental control (Shiffrin, 2016). Unless one is careful, this can be compounded by an increase risk of discovering significant, spurious correlations (Calude and Longo, 2017). One must therefore be cautious when analysing and drawing inferences from large datasets alone.

It is important to perform analyses with clear apriori hypotheses, rather than going “fishing” for results. Importantly, hypotheses in this thesis were informed using past psychological theory and our analyses sought to evaluate these predictions. Some were inspired by past lab studies, where there is usually more experimental control, such as the prediction that sequential consumer choices would semantically cluster (in Chapter 3). Other hypotheses were generated by computational cognitive models, which make explicit, testable and interpretable predictions about human behaviour. For example, the models used in Chapter 3 and 4 were used to generate predictions about

preferential choices of consumers and laboratory participants, respectively, and were subsequently confirmed.

During the analyses itself, one can account for obvious confounds or co-variates by isolating their effects. For example, in Chapter 3, an alternative explanation for shoppers purchasing similar products is that they were somehow biased by the design of the website. We therefore included variables that represented each navigation method (e.g., using the search bar or the drop down) in regressions predicting choices and their response times. We also reproduced all analyses on a filtered version of the shopping dataset capturing transitions that occurred exclusively via the search bar. By accounting for confounding variables in this way, one can isolate variance explained by the core variables of interest, such as semantic similarity between products.

Analyses of observational data can also be used to generate hypotheses for future experiments, which helps to ensure that findings are reliable and generalisable. For example, we confirmed that our retrieval model elicited a similar pattern of model fits and attention weights when estimated on a separate dataset of food fluency retrievals (in Chapter 3). Discoveries from observational data can also motivate behavioural experiments, like in Chapter 2 — where the categories discovered by a Topic model were confirmed using an odd-one-out experiment undertaken by real consumers — and Chapter 4 — where consumers' increasing propensity to exploit their favourite product was demonstrated in an analogous lab task. When follow-up experiments were not in scope, I was cautious to flag possible limitations in drawing causal inference, such as in Chapter 3, where we discussed a relationship between reliance on episodic memory and forgetting.

The approaches described above are examples of good scientific practise and are certainly not limited to the analyses of big data. Other approaches have been proposed more recently with a specific focus on minimising false discovery during investigation of big datasets. For example, Agrawal et al. (2020) proposed Scientific Regret Minimisation, where cognitive models are

adjusted so as to minimise residuals between their predictions and those of a powerful, unconstrained predictive model. This ensures that cognitive models only seek to explain cases that were known to be predictable as opposed to unexplainable noise. This seems like an appropriate approach in cases where machine learning models can be scaled to make accurate predictions about behaviour. However, it could struggle otherwise, such as in the modelling problem described in Chapter 3, where the number of available options (i.e., output classes) exceeded 40,000. More work is required to understand how best to leverage big data and “big” (i.e., highly parameterised) models for scientific discovery (Hofman et al., 2021) and I believe that cognitive science will benefit substantially from these advances.

5.6.4 Beyond the supermarket

A core aim of this thesis was to understand how people choose options in large option spaces and I believe that the supermarket is an excellent domain to study this. Shoppers must choose between tens of thousands of products, and yet navigate these environments on a regular basis with relative ease. However, supermarkets have several unique characteristics that may not be shared by other preferential choice domains. In particular, supermarkets mostly sell fast moving consumer goods (FMCG). These goods sell quickly, are often repeat-purchased and are often purchased in combination with others. This contrasts with slower moving domains such as electronics, clothing or movies, where products are purchased with a lower frequency, rarely repeat-purchased and rarely combined.

The approaches used to predict preferential choices in this thesis may be particularly well-suited to the goal of modelling choices in fast-moving domains. The topic model used in Chapter 2 learned representations of similarity from co-occurrences of products in shopping baskets. For slower moving domains, the slow rate of sale may mean leave less opportunity for meaningful representations to be abstracted from co-occurrences in shopping baskets. A

practical remedy is to use more inclusive definition of “documents”, such as modelling the items consumed by customers; this is the essence of collaborative filtering models that are used to determine novel movie recommendations (Koren, 2009). For these to be plausible cognitive models, one must assume that people also infer representations through other people’s choices; an interesting hypothesis that has recently been explored (Analytis et al., 2018) and is worthy of further investigation.

A related issue is that choices in FMCG domains are often repeat purchases (i.e., *exploitation*), whereas choices in other domains — such as movies — are typically novel (i.e., *exploration*). Predicting product discovery is a unique enterprise that requires a different class of models. For example, movie recommender systems are often censored to ensure that past choices are not recommended (Koren, 2009). As a result, these systems typically recommend options that are *similar* to those previously experienced, although it remains unclear whether this is the best approach in the long term (Qin and Zhu, 2013). As we have seen from large-scale studies of consumers in the supermarket (Riefer et al., 2017) or on fast-food delivery apps (Schulz et al., 2019), consumers go through intermittent phases of product exploration. Future work could examine these phases of exploration in more detail.

An extreme position could be that grocery shopping is entirely unique, in that it is relatively unusual to make frequent, repetitive choices. One could therefore argue that the ideas proposed in this thesis do not generalise beyond grocery shopping. I would argue that this is far from being the case. Firstly, many choices made outside of the supermarket are repetitive and frequent. For example, decisions about where and what to eat for lunch or where to order takeaway. Non-food purchases such of cosmetics or medicines may occur somewhat less regularly, but will likely be repetitive. In all cases one could imagine deriving a reasonable model of semantic memory from the choice bundles of consumers, as we do in Chapter 2.

Importantly, whilst the ideas discussed in these chapters are often framed

within the context of grocery shopping, they often speak to more general principles about memory and cognition. Thus, they could be readily extended to explain choices elsewhere. Firstly, results from Chapters 2 and 3 are supportive of the general claim that rich semantic representations can be inferred from co-occurrence statistics in the environment and may even be learnt in this way. Co-occurrence statistics can be derived from many seemingly unique domains and yet the same basic idea of deriving low dimensional representations from co-occurrences within these domains has proven to be useful in the study of semantic memory.

In addition, the retrieval model proposed in Chapter 3 can more generally be used to explain how options are sequentially retrieved from memory. Comparable models have been proposed to explain the process of choice set generation during open-ended decision making (e.g., “where would you like to go on holiday?”) (Zhang et al., 2021; Aka and Bhatia, 2021). More recently, semantic congruence has been shown to bias search behaviour in an open-ended active learning task, where participants must learn about properties of a category by sequentially searching terms in a search bar (He et al., 2022). Thus, across the many domains in which choice sets are not presented explicitly, similar ideas about associative memory retrieval manifest to explain choice behaviour.

5.7 Future directions

5.7.1 A combined model of option retrieval and preferential choice

To understand decision-making, psychologists often “divide and conquer... decompose a complex problem into simpler problems... paste these analyses together with a logical glue” (Raiffa, 1968, p. 271). My thesis has broadly followed this approach, as I have addressed the problem of option representa-

tion, generation and preference learning across separate chapters. A natural next-step could be to combine ideas from multiple chapters. For example, one could develop a computational model of option retrieval and preference learning by combining the model of option retrieval proposed in Chapter 3 with the model of preference learning proposed in Chapter 4. I have often wondered whether coherency maximisation is moderated by the extent to which choices are memory-based (e.g., rely on options being retrieved from memory) and would be excited to see this addressed as part of future work.

It would also be exciting to see how the ideas proposed in this thesis could be combined with more conventional models of preferential choice elaborated elsewhere. For example, when modelling online grocery purchases, our model ignored the fact that consumers *were* presented with choice sets after they had decided what to search for next (e.g., after deciding to look for cereals, they would enter this into the search bar and then be presented with an explicit list of possible options). Choices amongst the presented options could be predicted using models of multi-alternative, multi-attribute choice (such as process models) (Trueblood, 2021). I would be particularly interested to see the extent to which model predictions improved with the inclusion of such models and thus get a sense of their relative importance in everyday decision-making.

5.7.2 Interactions with other choice attributes

There are many variables known to influence preferential choice that have not been included here. Whilst I have primarily investigated the contiguous relationships between options, consumers are also sensitive to intrinsic properties of options. For example, consumers are sensitive to changes in price, with demand for goods typically increasing as the price decreases (Marshall, 1890). Consumers are also influenced by the perceptual properties of options, such as their nutritional properties and associated taste (Sullivan and Huettel, 2021) or their visual appearance (Milosavljevic et al., 2012; Iigaya et al., 2021). Future

work could explore how such intrinsic properties interact with the mechanisms of memory-based choice discussed here. For example, some consumers may exclusively choose the cheapest alternative; in which case the CDC model would need to be adapted correspondingly, rather than represent preferences for all attributes as ideal points.

Uncertainty in a decision can also affect what is chosen. Uncertainty can be expected, such as when a favourite product is only in stock on half of our trips to the supermarket or when products belonging to a certain brand are only good half of the time. Uncertainty can also be unexpected, such as when the burgers offered by a new local restaurant are substantially worse than those offered by your usual favourite (Yu and Dayan, 2005). Future work could investigate how such sources of uncertainty interact with the properties of memory-guided choice described here. For example, perhaps environments with high unexpected uncertainty or volatility elicit more coherency maximisation (i.e., higher learning rate in CDC), as they do in studies of RL (Behrens et al., 2007; Nassar et al., 2010).

5.7.3 Leveraging different types of conceptual similarity

I remain fascinated by the apparent flexibility in which consumers consider options to be similar. For example, when deciding what to purchase for tonight's dinner, one may consider which options *complement* each other in a recipe whilst simultaneously deciding which of a set of *substitutable* ingredients are best. Broadly, substitutability and complementarity (often used in economics, Samuelson 1974 and computer science, Han et al. 2000) align to the psychological constructs of taxonomic and thematic similarity, respectively. There has been substantial interest in understanding the relative contribution of taxonomic and thematic relations towards judgments of similarity; between tasks (Wisniewski and Bassok, 1999; Murphy and Ross, 1999), individuals (Gentner and Brem, 2020) and brain regions (Estes et al., 2011) (for a review see, Mirman et al., 2017). I presented some work in Chapter 3 discussing how these

may map on to episodic and semantic memory systems. However, I would like to see more work acknowledge and investigate their role during preferential choice. There may also be opportunities to elaborate our definitions of these constructs. For example, taxonomic relations are said to share intrinsic features, but many options considered to be substitutable bear little conceptual overlap (e.g., tofu is a vegan substitute for chicken but they share few intrinsic relations). Future research could investigate how these relationships are learnt, inferred and utilised in the service of preferential choice. This would contribute to a small existing literature on the subject (Guest et al., 2016; Estes et al., 2012; Felcher et al., 2001)

5.7.4 Adjusting subjective preferences to support healthier choices

In the absence of extrinsic feedback, memory and preferences may be adjusted to suit prior choices. For example, choosing between two options drives an increase in preference for the chosen option and decrease for the nonchosen option (Brehm, 1956; Sharot et al., 2010). In this thesis, I have shown that this is not limited to options directly experienced in the past (e.g., in Chapter 4). Instead, people generalise their preferences across attributes of choices, meaning that they also come to prefer options that are *similar* to ones they previously chose. The consequences of this choice-induced preference change may be more widespread than first thought.

For example, when decisions are facilitated by recommender systems, post-decisional choice biases may exacerbate the propensity to choose within *filter bubbles* (Pariser, 2011). As people choose more and more recommended options by recommender systems (e.g. YouTube videos, Facebook news articles, offers for grocery products), they end up consuming more and more similar content. This is because recommender systems prioritise content that is similar to content they have consumed in the past (Koren, 2009). User consumption is determined by these algorithms and their preferences corre-

spondingly adjust to suit their choices, leading to a virtuous circle. Narrowing people’s preferences is most-likely an unintended consequence of algorithms that are trained to optimise observable metrics, such as session length or the number of clicks. For people wanting to develop new consumption habits, such as those wishing to eat more healthily, they may feel unduly pressured by personalised offers encouraging them to eat more unhealthy food.

Rather than recommending options that are deliberately similar to one’s previous choices, an algorithm could recommend options to promote positive change in people’s preferences. For example, we could theoretically induce healthier preferences by determining a personalised sequence of recommendations that shifts people’s preferences towards progressively healthier options over time, whilst still being close to people’s original preferences. According to the model proposed in Chapter 4, preferences could be “garden pathed” by leading consumers to choose between options that are similar to their existing preferences but increasingly more similar to a target product. This is something we intend to study in the future. This idea builds on recent studies of choice engineering (Dan and Loewenstein, 2019; Dezfouli et al., 2020), where adversarial algorithms learn to exploit decision biases and thus shape human decisions.

5.8 Looking ahead

The next ten years will likely bring greater clarity about the role of memory in preferential choice, in much the same way that it has with reinforcement learning. For the field to make comparable progress with RL, we will need to develop a better understanding of the intrinsic motivations driving learning and decision-making, such as coherency maximisation, so that we can capture dynamics of learning and decision-making over time.

Modelling and predicting preferential choices is a lucrative industry that will undoubtedly improve as more historical data is collected. This is the

dawn of a new era for computational cognitive science, in which theory can be evaluated at an unprecedented scale. Developing recommender systems is too-often treated as a “black-box” engineering problem, meaning that they may inadvertently exploit human biases to maximize a company-defined metric (e.g., session length, click-through rate). A challenge for cognitive scientists will be to develop a deeper understanding of this relationship, allowing for the development of systems that optimise human goals, with their consent, such as becoming more healthy or sustainable.

5.9 Conclusion

Long-term memory plays a vital role in supporting everyday preferential decision-making; helping decision-makers to mentally organise the plethora of options available, infer their properties when making decisions and justify them following a choice. Whilst prior knowledge may be considered a nuisance variable in laboratory tasks, I’ve shown that it has a critical function during everyday decision-making, where the option space is large, multidimensional and overlapping and choices are sequentially dependent. Many of these findings were prompted by a desire to explain behaviour outside of lab settings. Large datasets of real-world decisions offer cognitive scientists a way to do this and test their ideas at an unprecedented scale. The growth in such datasets will present unique perspectives in the future and I hope that my work serves as inspiration for some of them.

Bibliography

- Abbott, J. T., Austerweil, J. L., and Griffiths, T. L. (2015). Random walks on semantic networks can resemble optimal foraging. *Psychological Review*, 122(3):558–569.
- Agrawal, M., Peterson, J. C., and Griffiths, T. L. (2020). Scaling up psychology

- via Scientific Regret Minimization. *Proceedings of the National Academy of Sciences*, 117(16):8825–8835.
- Aka, A. and Bhatia, S. (2021). What I like is what I remember: Memory modulation and preferential choice. *Journal of Experimental Psychology: General*.
- Akaishi, R., Umeda, K., Nagase, A., and Sakai, K. (2014). Autonomous mechanism of internal choice estimate underlies decision inertia. *Neuron*, 81(1):195–206.
- Alos-Ferrer, C. and Shi, F. (2012). Choice-induced preference change and the free-choice paradigm: A clarification. *SSRN*.
- Analytis, P. P., Barkoczi, D., and Herzog, S. M. (2018). Social learning strategies for matters of taste. *Nature Human Behaviour*, 2(6):415–424.
- Anderson, J. R. (1990). *The Adaptive Character of Thought*. Psychology Press, New York.
- Ariely, D., Loewenstein, G., and Prelec, D. (2003). “Coherent Arbitrariness”: Stable Demand Curves Without Stable Preferences. *The Quarterly Journal of Economics*, 118(1):73–106.
- Behrens, T. E. J., Woolrich, M. W., Walton, M. E., and Rushworth, M. F. S. (2007). Learning the value of information in an uncertain world. *Nature Neuroscience*, 10(9):1214–1221.
- Bhatia, S. and Stewart, N. (2018). Naturalistic multiattribute choice. *Cognition*, 179:71–88.
- Blanco, N. J., Otto, A. R., Maddox, W. T., Bevers, C. G., and Love, B. C. (2013). The influence of depression symptoms on exploratory decision-making. *Cognition*, 129(3):563–568.

- Blei, D. M., Ng, A. Y., and Jordan, M. I. (2003). Latent dirichlet allocation. *Journal of Machine Learning Research*, 3:993–1022.
- Braun, A., Urai, A. E., and Donner, T. H. (2018). Adaptive History Biases Result from Confidence-Weighted Accumulation of past Choices. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, 38(10):2418–2429.
- Brehm, J. W. (1956). Postdecision changes in the desirability of alternatives. *The Journal of Abnormal and Social Psychology*, 52(3):384–389.
- Burda, Y., Edwards, H., Pathak, D., Storkey, A. J., Darrell, T., and Efros, A. A. (2018). Large-scale study of curiosity-driven learning. *CoRR*, abs/1808.04355.
- Busemeyer, J. R. and Rieskamp, J. (2014). Psychological research and theories on preferential choice. In *Handbook of Choice Modelling*. Edward Elgar Publishing.
- Calude, C. S. and Longo, G. (2017). The Deluge of Spurious Correlations in Big Data. *Foundations of Science*, 22(3):595–612.
- Cicchini, G. M., Mikellidou, K., and Burr, D. (2017). Serial dependencies act directly on perception. *Journal of Vision*, 17(14):6.
- Dan, O. and Loewenstein, Y. (2019). From choice architecture to choice engineering. *Nature Communications*, 10(1):2808.
- Davis, T. and Love, B. C. (2010). Memory for category information is idealized through contrast with competing options. *Psychological Science*, 21(2):234–242.
- Deese, J. (1959). On the prediction of occurrence of particular verbal intrusions in immediate recall. *Journal of Experimental Psychology*, 58(1):17–22.

- Dezfouli, A., Nock, R., and Dayan, P. (2020). Adversarial vulnerabilities of human decision-making. *Proceedings of the National Academy of Sciences*, 117(46):29221–29228.
- Dong, D. W. and Atick, J. J. (1995). Statistics of natural time-varying images. *Network: Computation in Neural Systems*, 6(3):345–358.
- Estes, Z., Gibbert, M., Guest, D., and Mazursky, D. (2012). A dual-process model of brand extension: Taxonomic feature-based and thematic relation-based similarity independently drive brand extension evaluation. *Journal of Consumer Psychology*, 22(1):86–101.
- Estes, Z., Golonka, S., and Jones, L. L. (2011). Thematic thinking: The apprehension and consequences of thematic relations. In *The Psychology of Learning and Motivation: Advances in Research and Theory, Vol. 54*, The Psychology of Learning and Motivation, pages 249–294. Elsevier Academic Press, San Diego, CA, US.
- Felcher, E. M., Malaviya, P., and McGill, A. L. (2001). The role of taxonomic and goal-derived product categorization in, within, and across category judgments. *Psychology & Marketing*, 18(8):865–887.
- Fischer, J. and Whitney, D. (2014). Serial dependence in visual perception. *Nature neuroscience*, 17(5):738–743.
- Gentner, D. and Brem, S. K. (2020). *Is Snow Really Like a Shovel? Distinguishing Similarity from Thematic Relatedness*, pages 179–184. Psychology Press, New York, first edition.
- Glimcher, P. W. and Rustichini, A. (2004). Neuroeconomics: The consilience of brain and decision. *Science (New York, N.Y.)*, 306(5695):447–452.
- Griffiths, T. L., Steyvers, M., and Tenenbaum, J. B. (2007). Topics in semantic representation. *Psychological Review*, 114(2):211–244.

- Gruber, M. J., Gelman, B. D., and Ranganath, C. (2014). States of curiosity modulate hippocampus-dependent learning via the dopaminergic circuit. *Neuron*, 84(2):486–496.
- Guest, D., Gibbert, M., Estes, Z., Mazursky, D., and Lam, M. (2016). Modulation of taxonomic (versus thematic) similarity judgments and product choices by inducing local and global processing. *Journal of Cognitive Psychology*, 28(8):1013–1025.
- Han, J., Pei, J., and Yin, Y. (2000). Mining frequent patterns without candidate generation. *ACM SIGMOD Record*, 29(2):1–12.
- He, L., Richie, R., and Bhatia, S. (2022). Limitations to optimal search in naturalistic active learning. In *Proceedings of the Annual Meeting of the Cognitive Science Society*, volume 44.
- Hills, T. T., Jones, M. N., and Todd, P. M. (2012). Optimal foraging in semantic memory. *Psychological Review*, 119(2):431–440.
- Hofman, J. M., Watts, D. J., Athey, S., Garip, F., Griffiths, T. L., Kleinberg, J., Margetts, H., Mullainathan, S., Salganik, M. J., Vazire, S., Vespignani, A., and Yarkoni, T. (2021). Integrating explanation and prediction in computational social science. *Nature*, 595(7866):181–188.
- Hornsby, A. N. and Love, B. C. (2022). Sequential consumer choice as multi-cued retrieval. *Science Advances*, 8(8):eabl9754.
- Iigaya, K., Yi, S., Wahle, I. A., Tanwisuth, K., and O’Doherty, J. P. (2021). Aesthetic preference for art can be predicted from a mixture of low- and high-level visual features. *Nature human behaviour*, 5(6):743–755.
- Jones, M. N., Gruenenfelder, T. M., and Recchia, G. (2018). In defense of spatial models of semantic representation. *New Ideas in Psychology*, 50:54–60.

- Jones, M. N., Hills, T. T., and Todd, P. M. (2015). Hidden processes in structural representations: A reply to Abbott, Austerweil, and Griffiths (2015). *Psychological Review*, 122(3):570–574.
- Jones, M. N. and Mewhort, D. J. K. (2007). Representing word meaning and order information in a composite holographic lexicon. *Psychological Review*, 114(1):1–37.
- Kahana, M. J. (2020). Computational Models of Memory Search. *Annual Review of Psychology*, 71(1):107–138.
- Keramati, M. and Gutkin, B. (2014). Homeostatic reinforcement learning for integrating reward collection and physiological stability. *eLife*, 3:e04811.
- Kidd, C. and Hayden, B. Y. (2015). The psychology and neuroscience of curiosity. *Neuron*, 88(3):449–460.
- Kiyonaga, A., Scimeca, J. M., Bliss, D. P., and Whitney, D. (2017). Serial Dependence across Perception, Attention, and Memory. *Trends in Cognitive Sciences*, 21(7):493–497.
- Knox, W. B., Otto, A. R., Stone, P., and Love, B. C. (2012). The nature of belief-directed exploratory choice in human decision-making. *Frontiers in Psychology*.
- Koren, Y. (2009). The bellkor solution to the netflix grand prize. *Netflix prize documentation*, 81(2009):1–10.
- Kruschke, J. K. (1992). ALCOVE: An exemplar-based connectionist model of category learning. *Psychological Review*, 99:22–44.
- Lawson, R., Chang, F., and Wills, A. J. (2017). Free classification of large sets of everyday objects is more thematic than taxonomic. *Acta Psychologica*, 172:26–40.

- Manassi, M. and Whitney, D. (2022). Illusion of visual stability through active perceptual serial dependence. *Science Advances*, 8(2):eabk2480.
- Marshall, A. (1890). *Principles of Economics*. London : Macmillan.
- Milosavljevic, M., Navalpakkam, V., Koch, C., and Rangel, A. (2012). Relative visual saliency differences induce sizable bias in consumer choice. *Journal of Consumer Psychology*, 22(1):67–74.
- Mirman, D., Landrigan, J.-F., and Britt, A. E. (2017). Taxonomic and Thematic Semantic Systems. *Psychological bulletin*, 143(5):499–520.
- Murphy, G. L. and Ross, B. H. (1999). Induction with cross-classified categories. *Memory & Cognition*, 27(6):1024–1041.
- Nakao, T., Ohira, H., and Northoff, G. (2012). Distinction between Externally vs. Internally Guided Decision-Making: Operational Differences, Meta-Analytical Comparisons and Their Theoretical Implications. *Frontiers in Neuroscience*, 6.
- Nassar, M. R., Wilson, R. C., Heasley, B., and Gold, J. I. (2010). An Approximately Bayesian Delta-Rule Model Explains the Dynamics of Belief Updating in a Changing Environment. *Journal of Neuroscience*, 30(37):12366–12378.
- Nosofsky, R. M. (2011). The generalized context model: An exemplar model of classification. *Formal approaches in categorization*, pages 18–39.
- Oudeyer, P.-Y. (2018). Computational theories of curiosity-driven learning.
- Pariser, E. (2011). *The Filter Bubble: What the Internet Is Hiding from You*. Penguin Group , The.
- Paxton, A. and Griffiths, T. L. (2017). Finding the traces of behavioral and cognitive processes in big data and naturally occurring datasets. *Behavior Research Methods*, 49(5):1630–1638.

- Qin, L. and Zhu, X. (2013). Promoting diversity in recommendation by entropy regularizer. In *Proceedings of the Twenty-Third International Joint Conference on Artificial Intelligence, IJCAI '13*, pages 2698–2704, Beijing, China. AAAI Press.
- Raaijmakers, J. G. W. and Shiffrin, R. M. (1980). SAM: A Theory of Probabilistic Search of Associative Memory. In Bower, G. H., editor, *Psychology of Learning and Motivation*, volume 14, pages 207–262. Academic Press.
- Raiffa, H. (1968). *Decision Analysis: Introductory Lectures on Choices under Uncertainty*. Decision Analysis: Introductory Lectures on Choices under Uncertainty. Addison-Wesley, Oxford, England.
- Rangel, A., Camerer, C., and Montague, P. R. (2008). A framework for studying the neurobiology of value-based decision making. *Nature Reviews Neuroscience*, 9(7):545–556.
- Riefer, P. S., Prior, R., Blair, N., Pavey, G., and Love, B. C. (2017). Coherency-maximizing exploration in the supermarket. *Nature Human Behaviour*.
- Roediger, H. L. and McDermott, K. B. (1995). Creating false memories: Remembering words not presented in lists. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 21(4):803–814.
- Ross, B. H. and Murphy, G. L. (1999). Food for thought: Cross-classification and category organization in a complex real-world domain. *Cognitive Psychology*, 38(4):495–553.
- Samuelson, P. A. (1974). Complementarity: An Essay on The 40th Anniversary of the Hicks-Allen Revolution in Demand Theory. *Journal of Economic Literature*, 12(4):1255–1289.
- Schmidhuber, J. (2007). Simple Algorithmic Principles of Discovery, Subjective Beauty, Selective Attention, Curiosity & Creativity. In Corruble, V., Takeda,

- M., and Suzuki, E., editors, *Discovery Science*, Lecture Notes in Computer Science, pages 26–38, Berlin, Heidelberg. Springer.
- Schultz, W. (2006). Behavioral theories and the neurophysiology of reward. *Annual Review of Psychology*, 57(1):87–115.
- Schultz, W., Dayan, P., and Montague, P. R. (1997). A neural substrate of prediction and reward. *Science (New York, N.Y.)*.
- Schulz, E., Bhui, R., Love, B. C., Brier, B., Todd, M. T., and Gershman, S. J. (2019). Structured, uncertainty-driven exploration in real-world consumer choice. *Proceedings of the National Academy of Sciences*, 116(28):13903–13908.
- Sharot, T., De Martino, B., and Dolan, R. J. (2009). How choice reveals and shapes expected hedonic outcome. *Journal of Neuroscience*.
- Sharot, T., Velasquez, C. M., and Dolan, R. J. (2010). Do decisions shape preference?: Evidence from blind choice. *Psychological Science*, 21(9):1231–1235.
- Shiffrin, R. M. (2016). Drawing causal inference from Big Data. *Proceedings of the National Academy of Sciences*, 113(27):7308–7309.
- Singh, S. P., Lewis, R. L., Barto, A. G., and Sorg, J. (2010). Intrinsically motivated reinforcement learning: An evolutionary perspective. *IEEE Transactions on Autonomous Mental Development*, 2:70–82.
- Sullivan, N. J. and Huettel, S. A. (2021). Healthful choices depend on the latency and rate of information accumulation. *Nature Human Behaviour*, 5(12):1698–1706.
- Sutton, R. and Barto, A. (1998). Reinforcement learning: An introduction. *IEEE Transactions on Neural Networks*, 9(5):1054–1054.

- Thorndike, E. L. (1927). The law of effect. *The American Journal of Psychology*, 39(1/4):212–222.
- Trueblood, J. (2021). Theories of Context Effects in Multi-alternative, Multi-attribute Choice.
- Tversky, A. (1977). Features of similarity. *Psychological Review*, 84(4):327–352.
- Tversky, A. and Gati, I. (1982). Similarity, separability, and the triangle inequality. *Psychological Review*, 89(2):123–154.
- Tversky, A. and Kahneman, D. (1974). Judgment under Uncertainty: Heuristics and Biases. *Science*, 185(4157):1124–1131.
- Wisniewski, E. J. and Bassok, M. (1999). What makes a man similar to a tie? Stimulus compatibility with comparison and integration. *Cognitive Psychology*, 39(3-4):208–238.
- Yu, A. J. and Dayan, P. (2005). Uncertainty, neuromodulation, and attention. *Neuron*, 46(4):681–692.
- Zemla, J. C., Kenett, Y. N., Jun, K.-S., and Austerweil, J. L. (2016). U-invite: Estimating individual semantic networks from fluency data. In *CogSci*.
- Zhang, Z., Wang, S., Good, M., Hristova, S., Kayser, A. S., and Hsu, M. (2021). Retrieval-constrained valuation: Toward prediction of open-ended decisions. *Proceedings of the National Academy of Sciences*, 118(20).
- Zhao, W. J., Richie, R., and Bhatia, S. (2022). Process and content in decisions from memory. *Psychological Review*, 129(1):73–106.

Appendix A

Probabilistic utility models

To understand why preferential choices may be difficult to predict, it's important to first specify what we expect. Probabilistic utility models provide a normative account of preferential choices and therefore describe what we expect. I briefly highlight an example of such a model, in order to highlight its main assumptions. The reader should consult Busemeyer and Rieskamp (2014) for more detail.

Assume there exists a complete set of options $X = \{A_1, \dots, A_n\}$, of which only a subset is presented, $Y = \{B_1, \dots, B_m\} \subseteq X, m \leq n$. The probability that an individual chooses option B_i from the set Y is denoted by $p(B_i|Y)$.

According to the fixed utility model of Luce and Suppes (Luce, 1959), a utility $u(A_i)$ can be assigned to each option in X , such that — when presented with a subset of options Y — the probability of choosing B_i equals:

$$p(B_i|Y) = f(u(B_i), u(B_{i+1}), \dots, u(B_n)) \tag{A.1}$$

Options are therefore chosen stochastically according to this probability distribution. Central to these choice models is the idea that people choose stochastically, which acknowledges the noisy and unpredictable nature of choice (Mosteller and Nogee, 1951). For instance, Hey (2001) showed that, when asked to choose between the same gambles across five sessions, none of 53 participants made the same choices across each session.

In the case of this model, the function f could be — for example — given by Luce’s ratio of strength model (Luce, 1959):

$$P(B_i|Y) = \frac{e^{u(B_i)}}{\sum_{j=1}^Y e^{u(B_j)}} \quad (\text{A.2})$$

In addition to choices being taken stochastically, this model makes several more contentious simplifying assumptions. For example, utilities assigned to choice options are assumed to be *stable*, in that they do not change as a result of the choice set Y . Moreover, utilities are assumed to be *complete*, in that utilities are known for each option and integrated at the time of decision. As discussed in this thesis, extensive behavioural research has shown that preferences are neither stable nor complete, leading to several adaptations that allow utility models to account for preferential choices observed inside and outside of the laboratory.

Bibliography

- Busemeyer, J. R. and Rieskamp, J. (2014). Psychological research and theories on preferential choice. In *Handbook of choice modelling*. Edward Elgar Publishing.
- Hey, J. D. (2001). Does Repetition Improve Consistency? *Experimental Economics*, 4(1):5–54.
- Luce, R. D. (1959). *Individual choice behavior*. Individual choice behavior. John Wiley, Oxford, England. Pages: xii, 153.
- Mosteller, F. and Noguee, P. (1951). An Experimental Measurement of Utility. *Journal of Political Economy*, 59(5):371–404. Publisher: University of Chicago Press.

Appendix B

Chapter 2 Appendix

B.1 Latent Dirichlet Allocation

For this project, we used a topic model known as Latent Dirichlet Allocation (LDA) (Blei et al., 2003). LDA is a generative probabilistic model that groups data into K unobserved topics. In the case of this project, baskets are represented as random mixtures over unobservable topics. Topics are then characterized as a mixture over N distinct products and D baskets. The generative process used by LDA can be described as follows:

1. For each topic $k \in \{1, \dots, K\}$:
 - Choose a distribution over products $\phi_k \sim \text{Dir}(\beta)$.
2. For each basket $d \in \{1, \dots, D\}$ in the collection c :
 - Generate a vector of topic probabilities: $\theta_d \sim \text{Dir}(\alpha)$
 - For each product $w_{d,n}$ in basket d :
 - Generate a topic assignment: $z_{d,n} \sim \text{Multinomial}(\theta_d)$
 - Draw a product $w_{d,n} \sim \text{Multinomial}(\phi_{z_{d,n}})$

Where β and α are hyperparameters that determine the concentration of the Dirichlet prior placed on the topics' distribution over products ϕ and the baskets' distribution over topics θ , respectively. The latent variables ϕ_k and θ_d

can then be inferred using an iterative learning algorithm, such as expectation maximization (Blei et al., 2003).

B.1.1 Topic inspection and labelling

Table B.1: The labels given to each of the 25 topics. Size is defined as the number of products that had the highest probability of belonging to the respective topic over the total number of products in the corpus. Asterisks indicate that they were surveyed in studies A and B.

Topic label	Size
Loose fruit and veg *	18.8%
Young children’s shop	18.0%
Own brand shop	17.5%
Cooking from scratch *	16.9%
Snacks	16.4%
Cheapest option *	15.4%
Home baking	14.9%
Exotic cooking from scratch	14.5%
Afternoon tea *	14.5%
Quick to prepare meals	13.9%
Branded store cupboard	13.9%
Summer salad *	13.2%
Summer fruits	12.6%
Low maintenance cooking *	12.5%
Low calorie options *	11.3%
Party snacks	9.3%
Christmas *	7.3%
Northern Ireland	6.3%
Delisted Products	5.8%
Cat lover	1.3%
Stir fry *	1.1%
Own brand family party	1.0%
Food for now *	1.0%
Eating from tins	0.5%
Desserts	0.2%

For each experiment, we calculated model *perplexity* on the training set and a held out test set. The perplexity of a model on that collection is defined as:

$$\text{Perplexity}(c) = \exp\left\{\frac{-\sum_d \log p(w_d)}{\sum_d N_d}\right\}$$

Where $p(w_d)$ is the probability of a product w in a basket d and N_d is the number of products in a basket. Perplexity can be useful to monitor during model training to ensure that the algorithm is converging on the training set and generalizing to unseen data. However, some have argued that perplexity and human interpretation are uncorrelated or even negatively correlated (Chang et al., 2009).

Given that interpretability is a fundamental goal of this research, we decided to test it more directly. In particular, we tested to see whether the topics were interpretable to humans and could identify known patterns in historic purchasing data. We now discuss this in more detail.

B.1.2 Calculating product relevancy for a topic

One conventional approach to interpreting topic models is to rank items (i.e. products) within a topic and manually inspect those with the highest probabilities. One can look for similarities between items with high probabilities to understand whether there is a key theme that binds them together.

A major issue with the traditional approach to interpreting topic models is that the probability of an item given a topic ϕ_{kw} can be biased positively in favour of more frequently occurring items within the corpus (Taddy, 2012). The validity of this traditional approach was therefore significantly limited, particularly given that we did not filter high-frequency products (i.e. stop items) from our dataset.

To overcome issues with biased item probabilities, we explored two additional measures for determining the pertinence of items within a topic; *lift* and *relevancy*.

The lift measure (Taddy, 2012) is defined as:

Lift is therefore a ratio of an item w 's probability within a topic k (i.e., $p(w_k)$) to its marginal probability across the corpus p_w . Whilst this helps to

diminish the impact of overall token frequency, some have argued that the measure is noisy (Sievert and Shirley, 2014). In particular, it can give overly high rankings to items that occur very rarely within the corpus. Indeed — during our manual inspection of the topics — we found this to be the case, limiting our ability to interpret the topics’ meanings.

To overcome this problem, Sievert and Shirley (2014) proposed the *relevance* metric, which is defined as:

$$r(w,k|\lambda) = \lambda \log p(w_k) + (1 - \lambda) \log \left(\frac{p(w_k)}{p_w} \right)$$

where λ is a free parameter that determines the weight given to the item w ’s probability within a topic k relative to its lift (measured on a log scale). If one sets $\lambda = 1$ then $r(w,k|\lambda) = \log p(w_k)$. Alternatively, if $\lambda = 0$ then $r(w,k|\lambda) = \log(\text{lift}(w,k))$. Thus, the benefit of this metric over lift is that it’s possible to blend the probability of an item given a topic with lift. The metric’s authors recommend using $\lambda = 0.6$ to maximize human interpretability, which is the value that we kept throughout all analyses (Sievert and Shirley, 2014).




Figure B.1: The relationship between item frequency within the corpus and a) item-topic probability, b) lift and c) relevancy

The data displayed in Figure B.1 plot the relationship between item corpus frequency and each respective metric. As discussed, results indicate a strong correlation between item probabilities and frequency ($r = 0.8862$). The lift metric ($r = 0.2193$) and relevance metric ($r = 0.2934$) considerably reduce this correlation.

After manually inspecting the topics with each of the three measures, we

agreed that relevance provided the best measure of item salience within a topic. We therefore used this to help determine topic names.

B.1.3 Initial topic labelling

When labelling the topics, the authors inspected the relevancy scores of each item within each topic, considered the most relevant items. Table B.1 shows the topic labels along with the relative size of each topic within the corpus.




Figure B.2: The top 50 most frequently occurring items across baskets within the corpus. Note that brand names have been removed.

Figure B.2 depicts the most popular products within the corpus. As is with most retailers (and word usage in language), the data is highly right skewed. This suggests that there are a small number of products that are purchased across many baskets and a long tail of products that are less popular.

Table B.2: A summary of mislabelling errors made by retail experts during the topic labelling task (***) indicates a proportion significantly different ($p < .001$) from the random baseline of 25.00% (1 of 4))

Topic label	Proportion correct	Most confused topic	Most confused frequency
Cheapest option	0.98 ***		0
Food for now	0.961 ***	Low maintenance cooking	2
Stir fry	0.961 ***		0
Low calorie options	0.961 ***		0
Christmas	0.961 ***	Food for now	2
Afternoon tea	0.961 ***	Cheapest option	2
Loose fruit and veg	0.922 ***	Cooking from scratch	3
Summer salad	0.882 ***	Cooking from scratch	3
Cooking from scratch	0.784 ***	Loose fruit and veg	8
Low maintenance cooking	0.725 ***	Food for now	7

B.2 Label confusion by retail experts

In the retail expert study, errors in labeling appeared sensible. Namely, the most popular alternative labels tended to be related to the original topic (see Table B.2). For example, the most popular alternative label to the *cooking from scratch* table was *loose fruit and veg*; both topic labels pertain to ingredients that need to be prepared before consumption.

B.3 Topics by day of week

In addition to monthly trends, the proposed topic labels are also indicative of different weekly trends in purchasing habits. In particular, we hypothesized that topics indicative of longer preparation times (e.g. *loose fruit and veg*) or a special weekend occasion (e.g. *afternoon tea*) would be more likely to occur on or just before the weekend. Contrasting, we hypothesized that topics indicative of impulse purchasing (e.g. *food for now*) or stocking up for the long-term (e.g. *branded store cupboard*) would not vary much across the week.

Bibliography

- Blei, D. M., Ng, A. Y., and Jordan, M. I. (2003). Latent Dirichlet Allocation. *Journal of Machine Learning Research*, 3:993–1022.

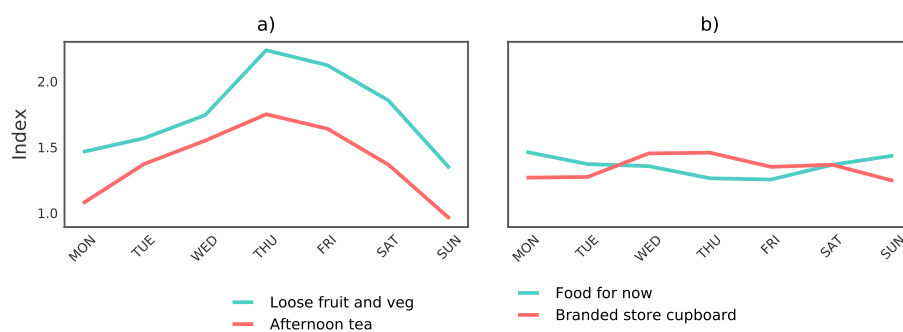


Figure B.3: The proportion of baskets with a given topic label on each day of the week, divided by the weekly mean average across all topics. Plot a) shows that food requiring longer preparation times (i.e. *loose fruit and veg*) or eaten specifically during weekend occasions (i.e. *afternoon tea*) are more likely to be bought on Thursday or Friday. Plot b) indicates that impulse purchases (i.e. *food for now*) or food that tends to be stored away (i.e. *branded store cupboard*) does not vary in popularity over the week.

Chang, J., Gerrish, S., Wang, C., and Blei, D. M. (2009). Reading Tea Leaves: How Humans Interpret Topic Models. *Advances in Neural Information Processing Systems 22*, pages 288—296.

Sievert, C. and Shirley, K. (2014). LDAvis: A method for visualizing and interpreting topics. *Proceedings of the Workshop on Interactive Language Learning, Visualization, and Interfaces*, pages 63–70.

Taddy, M. a. (2012). On Estimation and Selection for Topic Models. In *Proceedings of the Fifteenth International Conference on Artificial Intelligence and Statistics*, number 2003, pages 1184–1193.

Appendix C

Chapter 3 Appendix

C.1 Method

C.1.1 Data

C.1.1.1 Clickstream data

A sample of the total visits was used at the request of the retailer. Visits were then filtered such that only those resulting in a purchase were kept. We then filtered the observations to instances where the visitor added an item to their basket, rather than — for example — viewed without adding. If products were later removed from a basket, they were not removed from the dataset, as we wanted the order to reflect the original retrieval process of the shopper. Unlike standard fluency tasks, repeated retrievals cannot be considered erroneous and thus were included in the analyses. Each visit was from a unique visitor to ensure that visits were independent.

Note that 12,520 (0.23%) basket adds in the clickstream data were for products sold exclusively online. Pairwise similarities were therefore not computable for these observations. For any analyses that included pairwise similarities, these non-computable pairwise similarities were dropped. This left 5,238,469 basket adds from 132,146 unique visitors.

Visitors could navigate between products using different features of the

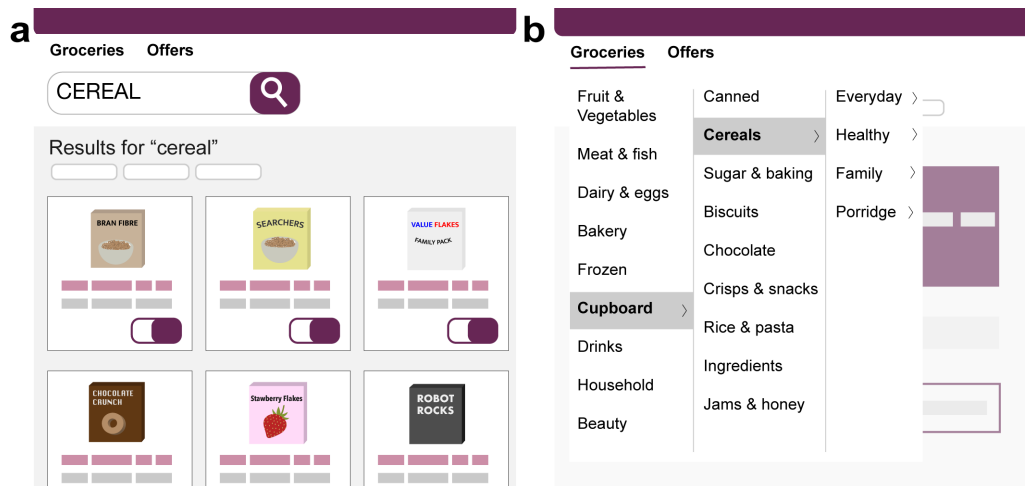


Figure C.1: Shoppers could transition between products using different design features of the website. **a.**, Most transitions occurred through use of a search bar, which was located at the top of the page. After entering a keyword, shoppers were presented with a list of relevant products associated with the keyword. Shoppers could add products to their basket from this search results page directly or click on the product to view a dedicated page containing more information (e.g., nutritional data). **b.**, Fewer transitions occurred through use of a category drop-down, which appeared when hovering the mouse over the *Groceries* hyperlink at the top of the page. Three levels of subcategories could be revealed by hovering one’s mouse over the respective department names. In the example pictured, the shopper has hovered their mouse over the “Cupboard” department and then the “Cereals” category. The products in each subcategory were displayed in the same way as the search results.

website, which are visualised in Figure C.1. Most choices were located using the search bar located at the top of the page, either immediately after its use (i.e., *search arrival*) or after a previous purchase in this context (i.e., *search stay*). Less often, visitors located products using the category drop-down (i.e., *category arrival*) or after a prior purchase in this context (i.e., *category stay*). Products could also be located using a generic special offers page, which was located on the landing page and could also be located using the ribbon at the top of the page. Finally, before checkout, visitors were suggested products that they may have forgotten.

C.1.1.2 In-store data

We used a sample of baskets purchased in-store during the same dates as the online dataset. Baskets were filtered so that they contained at least one of the products purchased in the clickstream dataset and contained a minimum of 5 products overall (a similar strategy was used by (20) to ensure that there was enough co-occurrence within each basket to learn reasonable representations of similarity).

C.1.2 Representations of associative knowledge

As explained in the introduction, we evaluated the relative contributions of three knowledge sources during sequential retrieval.

C.1.2.1 Episodic knowledge

The long-term episodic retrieval structure used in the Search of Associative Memory (SAM) model (?) associates items more strongly to the extent that they co-occur during encoding. Whilst we cannot know the exact context in which consumers encoded products, we assume that products purchased together more frequently in the same basket must be stronger episodic associates. Episodic associations are therefore represented using the pairwise co-occurrence between products observed in the in-store dataset.

Episodic similarities $S(a,b)$ were determined using the probability of co-occurrence between products a and b , which is given by:

$$S(a,b) = \frac{freq(a,b)}{freq(a)}$$

Where $freq(a,b)$ is the total number of number of times that product a co-occurred with b in the same basket across the dataset.

C.1.2.2 Semantic knowledge

In addition to episodic memory, shoppers likely also rely on semantic memory to guide their retrievals. A common feature of modern semantic memory models is that they represent knowledge within a connected representational space, allowing people to generalise their knowledge to observations they haven't directly experienced (43). Unlike episodic co-occurrence, two items that have never co-occurred together may still be considered semantically similar, so long as they co-occur in similar contexts. We followed recent research (20) by training a 200-dimensional distributed semantic model using the in-store data.

For this project, we chose to learn 200-dimensional vector representations for products with *word2vec* (Mikolov et al., 2013). This is because *word2vec* tends to scale better when trained on large datasets, because it can be trained stochastically. Rather than encoding words (e.g., as they might appear in the product descriptions), *word2vec* was trained to represent supermarket product codes, as they might appear in till receipts. Concretely, there is a different product code for each distinct product in the supermarket. Small variations in that product (i.e. different sizes of the same t shirt) are not given separate codes however. Each product code was thus represented as a one-hot encoded vector before being embedded, which resulted in a $42,837 \times 200$ matrix. During training, the model learns to associate product codes that often co-occur in baskets, which is analogous (but not the same) to how word vector models learn similarity between word tokens that often co-occur in sentences.

Associations S between any two vectors v_1 and v_2 was calculated using the cosine similarity $\cos(v_1, v_2)$:

$$S(v_1, v_2) = \cos(v_1, v_2) = \frac{v_1 \cdot v_2}{\|v_1\| \cdot \|v_2\|}$$

Conventionally, *word2vec* is trained using one of two network architectures. In this case, we used Continuous Bag of Words (CBOW) with negative sampling. CBOW assumes that items within baskets are un-ordered (i.e.,

known in natural language processing as the “bag of words” assumption). This assumption is true of in-store supermarket data, where product codes within baskets are unordered once they reach the database. Using CBOW, the training objective \mathcal{O} is to maximise the likelihood of a target item i_j given a window of c surrounding context items:

$$\mathcal{O} = \frac{1}{T} \sum_{j=1}^T \log p(i_j | i_{j-c}, \dots, i_{j+c})$$

Where T is the corpus size (e.g., number of baskets) and $p(i_j | i_{j-c}, \dots, i_{j+c})$ is the probability of a target item i_j given average or summation of a set of context words $(i_{j-c}, \dots, i_{j+c})$. In this case, word2vec used a context window size of 15.

Whilst this probability *could* be determined using a softmax, it is not practical when called over large vocabulary sizes. Negative sampling mitigates this by randomly sampling a set of k “negative” items (in our case, 20) that did not appear in the present basket. The model then learns from these negative samples by treating them as false labels in multiple binary classification tasks, evaluated using a logistic loss and then updated using gradient descent. In this project, this training process was repeated for 15 epochs.

We determined these aforementioned hyperparameters by monitoring the loss on the training set and visually inspecting the nearest products in terms of cosine similarity for a set of randomly-chosen products. During model fits, negative cosine similarities were clipped to 0.

C.1.2.3 Structured hierarchical knowledge

A strict organisation of products is imposed on consumers by way of a product taxonomy, which groups products from small subgroups (e.g., apples) to large departments (e.g., produce). Amongst other things, this taxonomy determines the proximity of products on the shelves and aisles of a supermarket store. The product taxonomy used here contained five levels. Each product had a unique taxonomisation.

If products a and b shared the same low-level category within the taxonomy, then they were said to be perfectly associated $S(a, b) = 1$. Products with entirely different taxonomic classifications had $S(a, b) = 0.2$. The remainder increased in increments of 0.2.

C.1.3 Retrieval model

Broadly, our retrieval model is based on the retrieval equation from Search of Associative Memory (SAM) (Raaijmakers and Shiffrin, 1980). It assumes that retrieval and thus the decision of what to choose next is achieved by querying associative structures in memory with a memory probe. We follow previous models of semantic fluency by using the most recently chosen option O_i to probe associative memory structures (Hills et al., 2012; Abbott et al., 2015). Whilst other possibilities exist — such as a decaying influence of all previous retrievals — we focus on the role of the prior choice in order to simplify analyses (for a review of other approaches, see Kahana 2020). The retrieval strength of the subsequently chosen option O_{i+1} is given by the product of the M associations between the present choice and itself, $S(O_i, O_{i+1})_j$. For example, in the full model, we used episodic, semantic and hierarchy-based associations between products, meaning that $M = 3$. This is then divided by the sum of that same function applied to all of the N options that remain to be added for that trip. This then gives rise to an overall probability of retrieval for each choice:

$$P(O_{i+1} | S_1, S_2, \dots, S_j, O_i) = \frac{\prod_{j=1}^M S(O_i, O_{i+1})_j^{\beta_j}}{\sum_{k=1}^N \prod_{j=1}^M S(O_i, O_k)_j^{\beta_j}} \quad (\text{C.1})$$

We compared the inclusion of episodic, semantic and hierarchy-based associations. β values represent attention weights for each of these knowledge representations and were estimated as free parameters for each visit.

Each model was compared with a random baseline model, which predicted

an equal probability of $\frac{1}{N}$ for every transition using a single representation. Thus, each of the products remaining to be purchased by each visitor is assumed to have an equal probability of being chosen at each timestep according to the baseline model.

C.1.3.1 Fit procedure

Each measure of association j was raised to its own respective attention weight β_j ; these were treated as free parameters and fit to individual visitors using maximum likelihood estimation (attention weights were forced to have a lower bound of 0, in order to prevent individual retrieval probabilities from exceeding 1). These free parameters were solved separately for each visit using the SLSQP solver within SciPy.

C.1.3.2 Model input

Models were fit to the retrieval sequences in the clickstream data. In addition to non-computable similarities, observations were dropped from the clickstream data if they occurred during or after the use of a recommender system, which prompted users about items they may have forgotten before checkout. Finally, to ensure that parameter estimates were robust, visits were dropped if they contained fewer than 10 items. This left 117,337 distinct visits.

Because of the probabilistic and multiplicative nature of the model, negative or zero similarities were replaced with a very small but positive number $1e-7$.

C.1.4 Permutation tests

To assess whether subsequent retrievals were more related than would be expected at random, we performed a permutation test. Each product within the clickstream data was encoded with each of the three representations described above. We then calculated the per-visit mean similarity between consecutively added products. These were compared to the per-visit mean similari-

ties determined by 100 random permutations of the product order, permuted within each visit. Thus, for significance tests reported; $N_{true} = 132,146$ and $N_{permuted} = 13,214,600$.

C.1.5 Response times

Response times (RTs) were compared for transitions of varying distances. These were capped at 60 seconds to minimise the leverage of outliers.

In the multiple linear regression comparing each similarity measure, RTs were monotonically transformed using a log function, $RT_{log} = \ln(RT + 1)$, due to positive skewness.

C.1.6 Trajectory analyses

Correlational analyses were conducted to assess how behaviour changed over time. Because visits contained differing numbers of products, basket adds within each visit were binned into equally-sized deciles based on their proximity to checkout. Subsequently, 13,901 (10.52%) visits were dropped because they purchased fewer than 10 products, leaving 5,174,018 choices.

C.1.7 Transition clustering

Spectral clustering was used to determine the extent to which transitions between categories within the taxonomy were clustered and thus could be predicted based on features of the current choice.

For each level of the product taxonomy, we calculated a transition matrix, counting the number of one-step transitions from each category i to every other category j (e.g., apples \rightarrow pears). Because product sales tend to be Pareto distributed (e.g., products such as milk and bananas are considerably more popular), we found that the odds of transitioning to more popular products were disproportionately skewed. To adjust for this, we used the Lift association score, which is used in Market Basket Analysis (44). Lift is defined as

$$Lift(i, j) = \frac{Support(i, j)}{Support(i) Support(j)},$$

where

$$Support(i, j) = \frac{Frq(i, j)}{N}$$

is the probability of the i to j transition out the N transitions observed.

The denominator in the lift calculation $Support(i) Support(j)$ therefore describes an expected probability of a particular transition being made, given the overall popularity of the two respective categories. Thus, lift values above 1 for a given transition suggest that shoppers transition between these categories more than one would *expect*. Because only higher-than-expected transitions were of interest for these analyses, we subtracted 1 from the lift matrix and set the lower bound to 0:

$$Lift(i, j) = \begin{cases} 0 & Lift(i, j) - 1 \leq 0 \\ Lift(i, j) - 1 & \text{otherwise} \end{cases} \quad (C.2)$$

This lift association matrix of the transitions can be thought of as the adjacency matrix of a directed graph. In this context, spectral clustering is a natural candidate for our problem, as it is commonly used in network science to identify k clusters (or *communities*) within a graph. The algorithm works by first calculating the normalized Laplacian of the graph's adjacency matrix and then applying a standard clustering algorithm (such as k -means) to the relevant eigenvectors of that new space. Similar approaches have been used in reinforcement learning to abstract high-level sub-goals from transition matrices and in recent neuroscientific literature to understand how the hippocampus represents actions to facilitate spatial navigation (Stachenfeld et al., 2017). These problems are analagous to the one faced by shoppers here, hinting that there may be a deeper connection between these clusters and the planning processes of shoppers; we leave this possibility open for future work.

C.1.8 Predictive modelling

We explored whether the attention weights (β) from the best-fitting retrieval model would predict the number of forgotten or removed items.

Attention weights (β) were taken from the best-fitting retrieval model, and reflect the extent to which each visitor recruited each of the three representations to guide choice. Importantly, these weights were estimated using behaviour prior to the use of the recommender system. Outlying attention weights (three standard deviations above the mean) were clipped for this analysis for numerical stability.

C.2 Additional analyses of clickstream data


C.2.1 Correlations between representations

Table C.1: Spearman correlations between each of the similarity measures (95% confidence intervals shown in parentheses)

	Episodic	Semantic	Structured
Episodic		0.585 [0.585, 0.586]	0.526 [0.525, 0.527]
Semantic			0.654 [0.654, 0.655]
Structured			

A key claim in this article is that people flexibly recruit multiple representations when deciding what to retrieve next. For this to be the case, each similarity measure would need to be related but not perfectly correlated. Looking at the sequential choices observed in the clickstream data, Spearman correlations revealed a moderate to strong relationship between each measure (coefficients shown in Table C.1. All were significant $p < .001$). Thus whilst these similarity measures are directionally similar, each likely captures unique information relevant to retrieval.

C.2.2 Similarity ripples



(a) Mean semantic similarity (with 95% confidence intervals) between the current product and those purchased most recently is higher compared with products purchased later

(b) Mean hierarchical similarity (with 95% confidence intervals) between the current product and those purchased most recently is higher compared with products purchased later

Figure C.2: Lagged average similarity between choices indicates how choices become less similar over time

The results shown in Figure C.2 show that option retrievals can also be viewed as a ripple through semantic and hierarchical knowledge, in addition to episodic memory.

C.2.3 Timestep and IRI regression

To evaluate whether responses slowed down over time, we regressed timestep onto IRI using a linear mixed-effects regression. We included the visit identifier as a random effect and included dummy coded representations of each navigation method as confounding variables. This was to confirm that any slow-down was not simply a result of the website's design. We also included each measure of similarity, to ensure that — holding the similarity between transitions constant — choices became slower over time. The multiple linear regression converged ($ll = -5979127.00$). As shown in Table C.2, the partial regression coefficient for timestep was positive, indicating that choices became slower over time, irrespective of how choices were navigated and the similarity between choices in memory.

Table C.2: Standardised coefficients, significance values, upper (UB) and lower bounds (LB) for variables included in the mixed-effects linear regression predicting IRI using timestep, transition similarity and transition type

	b	$p \leq$	95% LB	95% UB
Intercept	2.857	0.0001	2.854	2.859
Timestep	0.130	0.0001	0.129	0.130
Episodic similarity	-0.137	0.0001	-0.138	-0.136
Semantic similarity	-0.086	0.0001	-0.087	-0.085
Hierarchical similarity	-0.273	0.0001	-0.274	-0.271
Category arrival	0.157	0.0001	0.136	0.177
Category stay	-0.004	0.680	-0.021	0.014
Offers	0.069	0.0001	0.054	0.084
Search arrival	0.010	0.614	-0.028	0.048
Search stay	-0.135	0.0001	-0.168	-0.103
Suggestions	-0.028	0.0001	-0.035	-0.022
Visit ID	0.162			

C.2.4 Similarity and timestep regression

Table C.3: Standardised coefficients, significance values, upper (UB) and lower bounds (LB) for variables included in the mixed-effects linear regression predicting timestep using episodic similarity and transition type

	b	$p \leq$	95% LB	95% UB
Intercept	5.583	0.0001	5.581	5.585
Episodic similarity	-0.058	0.0001	-0.061	-0.055
Category arrival	0.932	0.0001	0.867	0.997
Category stay	0.758	0.0001	0.702	0.813
Search arrival	1.547	0.0001	1.426	1.668
Search stay	1.346	0.0001	1.242	1.449
Offers	0.448	0.0001	0.400	0.496
Suggestions	0.605	0.0001	0.584	0.626
Visit ID	0.000			

To evaluate whether choices became less similar over time, we regressed each similarity measure onto timestep using linear mixed-effects regressions. In each regression we included the visit identifier as a random intercept and included dummy coded representations of each navigation method as con-

Table C.4: Standardised coefficients, significance values, upper (UB) and lower bounds (LB) for variables included in the mixed-effects linear regression predicting timestep using semantic similarity and transition type

	b	$p \leq$	95% LB	95% UB
Intercept	5.583	0.0001	5.581	5.585
Semantic similarity	-0.010	0.0001	-0.013	-0.007
Category arrival	0.948	0.0001	0.883	1.013
Category stay	0.754	0.0001	0.698	0.810
Search arrival	1.577	0.0001	1.456	1.698
Search stay	1.350	0.0001	1.246	1.454
Offers	0.458	0.0001	0.410	0.506
Suggestions	0.609	0.0001	0.587	0.630
Visit ID	0.000			

Table C.5: Standardised coefficients, significance values, upper (UB) and lower bounds (LB) for variables included in the mixed-effects linear regression predicting timestep using hierarchical similarity and transition type

	b	$p \leq$	95% LB	95% UB
Intercept	5.583	0.0001	5.581	5.585
Hierarchical similarity	-0.247	0.0001	-0.250	-0.243
Category arrival	0.882	0.0001	0.817	0.947
Category stay	0.784	0.0001	0.728	0.840
Search arrival	1.445	0.0001	1.325	1.566
Search stay	1.403	0.0001	1.299	1.507
Offers	0.422	0.0001	0.374	0.470
Suggestions	0.582	0.0001	0.561	0.603
Visit ID	0.000			

founding variables. The regression predicting timestep using episodic similarity converged ($ll = -12751137.48$) and the coefficients are shown in Table C.3. The regression predicting timestep using semantic similarity also converged ($ll = -12751890.27$) and the full equation is shown in Table C.4. The regression predicting timestep using hierarchical similarity also converged ($ll = -12743140.39$) and the full equation is shown in Table C.5. Each regression revealed a negative relationship between similarity and timestep, indicating that — irrespective of the navigation method — sequential choices became

more dissimilar over time.

C.2.5 Representation and IRI regression

Table C.6: Standardised coefficients, significance values, upper (UB) and lower bounds (LB) for variables included in the mixed-effects linear regression model predicting IRI using each similarity measure and confounding variables

	b	$p \leq$	95% LB	95% UB
Intercept	2.834	0.0001	2.831	2.836
Episodic	-0.138	0.0001	-0.139	-0.137
Semantic	-0.083	0.0001	-0.084	-0.082
Hierarchy	-0.278	0.0001	-0.279	-0.277
Category arrival	0.193	0.0001	0.174	0.212
Category stay	0.027	0.0010	0.011	0.044
Search arrival	0.070	0.0001	0.034	0.106
Search stay	-0.081	0.0001	-0.112	-0.050
Offers	0.091	0.0001	0.077	0.105
Suggestions	-0.010	0.0020	-0.017	-0.004
Choices remaining	-0.138	0.0001	-0.139	-0.137
Visit ID	0.161			

A mixed-effects multiple linear regression was performed, regressing each similarity measure onto the IRIs between each choice. We included a random intercept for each visit. We also — as above — included dummy coded representations for each navigation method and counts of the number of choices remaining as confounding variables.

The coefficients are presented in Table C.6. They reveal that — taking account of different navigation contexts — retrieval from hierarchical knowledge uniquely explains the most variance in response times of the knowledge types, followed by episodic knowledge and then semantic knowledge.

C.2.5.1 Model comparison

We performed feature selection for the inter-response interval (IRI) mixed-effects linear regressions, to confirm that response times were best explained by

multiple knowledge systems. Similarity measures were removed using stepwise-elimination. All models contained the confounding variables described above. We compare models based on their Akaike Information Criterion (AIC) and Bayesian Information Criterion (BIC).

Table C.7: Nested model comparisons (using AIC and BIC) for linear-mixed effects regressions predicting IRI. Models were chosen using backward elimination.

Variables	AIC	BIC
Episodic, Semantic, Hierarchical and Confounds	12185903.95	12186079.08
Episodic, Hierarchical and Confounds	12204652.31	12204813.97
Hierarchical and Confounds	12367250.26	12367398.45

Table C.7 indicates that each measure of similarity contributed significant explanatory power to the model, with full model having the lowest AIC and BIC overall. This indicates that choice IRIs are best explained by a combination of knowledge sources.

C.2.6 Transitions between categories reveal hierarchical knowledge

To gain further insight about the structure of shopper’s hierarchical knowledge, we evaluated whether sequential transitions clustered into meaningful groups. In particular, we clustered transitions between product groups at each level of the product taxonomy. Transitions between categories appeared clustered — at least to some extent — across all levels of the product taxonomy. All of the best-fitting clustering solutions exhibited a positive silhouette score (Level 3: *Silhouette* = .503 $N_{clusters} = 23$, Level 4: *Silhouette* = .614 $N_{clusters} = 11$, Level 5: *Silhouette* = .33 $N_{clusters} = 4$). For example, transitions between categories defined at the fourth level of the product taxonomy, depicted in Figure 2e in the main manuscript, appeared highly clustered. Moreover, the clusters revealed intuitive groupings that often overlapped with super-ordinate classifications

in the product taxonomy (e.g., clustering separate beers, wines and spirits categories). One possibility is that these clusters emerge because shoppers transition between them with a clear plan in mind, which is executed across multiple transitions. However, a simpler explanation could be that these clusters emerge as a consequence of cued-retrieval from a hierarchical associative knowledge structure.

Although shoppers could feasibly transition between any pair of products for a comparable physical cost, it’s intriguing that they instead prefer to adhere closely to the product taxonomy. One possibility is that shoppers use a mental model of a physical store layout to guide their search online, indicating a close correspondence between spatial and non-spatial navigation. Another related possibility is that this product taxonomy has been developed to closely resemble the knowledge structures of consumers.

C.2.7 Correlation between attention weights

Table C.8: Spearman correlations between each of the attention weights

	Episodic	Semantic	Hierarchy
Episodic	- 0.0216 [0.0159, 0.0273]	-0.2733 [-0.2786, -0.268]	
Semantic	-	- -0.2627 [-0.268, -0.2574]	
Hierarchy	-	-	-

Table C.8 shows the correlations between each of the learned attention weights across all visits (all were significant $p < .0001$). Despite reasonably high correlations between representations, it is perhaps reassuring that correlations between attention weights do not far exceed $|0.28|$.

C.2.8 Retrieval model parameter recovery

To ensure that each knowledge representation was identifiable, we performed a parameter recovery study. We took a random 10% sample of visits and

attempted to recover the attention weights that had been learnt during estimation of the multiple parameter model on the clickstream data. For each set of attention weights, we generated 100 synthetic purchase sequences, by selecting a random product as the first choice and then sampling subsequent choices according to the retrieval model. Choices were selected from the product universe observed in the clickstream dataset and product repetitions were permitted. Each generated trip contained 40 products, which is equal to the average purchase length observed in the true data.

Results showed that each parameter could be recovered accurately. In particular, Spearman correlations between estimated and actual β weights were high across episodic ($r_s = 0.6829, p \leq .0001$), semantic ($r_s = 0.6034, p \leq .0001$) and hierarchical knowledge ($r_s = 0.9883, p \leq .0001$). That each system can be uniquely identified supports our interpretations of these knowledge sources as distinct cognitive processes.

C.2.9 Forgotten items regression

Table C.9: Standardised coefficients, significance values, upper (UB) and lower bounds (LB) for variables included in the regression model predicting number of forgotten items using the model attention weights, total number of choices and proportion of each transition type

	b	$p \leq$	95% LB	95% UB
Intercept	0.3044	0.0001	0.299	0.310
Episodic	-0.0534	0.0001	-0.059	-0.048
Semantic	0.0150	0.0001	0.009	0.021
Hierarchy	-0.0407	0.0001	-0.046	-0.035
Category arrival	-0.0065	0.1660	-0.016	0.003
Category stay	-0.0187	0.0001	-0.028	-0.010
Offers	0.0454	0.0001	0.039	0.052
Search arrival	-0.0406	0.0001	-0.052	-0.029
Total choices	-0.0414	0.0001	-0.047	-0.036

As reported in the main text, we regressed attention weights from the best fitting model onto the number of forgotten items using multiple linear

regression. We also included the total number of choices and the proportion of each transition as confounding variables. The full model equation is shown in Table C.9. These results show that the model attention weights explained unique variance when predicting the number of forgotten items. This suggests that — irrespective of how one navigates the website — the proximity between sequential choices in memory predicts one’s propensity to forget products.

C.2.10 Removed items regression

Table C.10: Standardised coefficients, significance values, upper (UB) and lower bounds (LB) for variables included in the regression model predicting number of products removed using the model attention weights, total number of choices and proportion of each transition type

	b	$p \leq$	95% LB	95% UB
Intercept	3.3836	0.0001	3.347	3.420
Episodic	-0.3585	0.0001	-0.398	-0.319
Semantic	0.0506	0.0090	0.012	0.089
Hierarchy	-0.8702	0.0001	-0.909	-0.831
Category arrival	0.2135	0.0001	0.151	0.277
Category stay	-0.2850	0.0001	-0.345	-0.225
Offers	0.3090	0.0001	0.264	0.354
Search arrival	-0.0507	0.1990	-0.128	0.027
Total choices	2.3466	0.0001	2.309	2.385

Similarly, we regressed attention weights from the best fitting model onto the number of products removed from each basket. We also included the total number of choices and the proportion of each transition as confounding variables. The full model equation is shown in Table C.10. These results show that the model attention weights explained unique variance when predicting the number of products removed. Similarly, this suggests that — irrespective of how one navigates the website — sequentially purchasing products that are close semantic relations may increase one’s propensity to add items to their basket that they don’t otherwise need.

C.2.10.1 Similarity between removed and purchased products

An alternative explanation for removing products could be that they were random products added accidentally or that they ended up not complementing the rest of a shopper's basket and thus being very distinct. To evaluate these possibilities, we evaluated the similarity between removed products with those purchased in each visit. To do this, we calculated the similarities between all products purchased in the main visit. We then calculated the mean similarity between each forgotten product and those purchased in the main shop. For each product removed, we then calculated the percentile of that mean average relative to the distribution of similarities observed between each purchased product.

Results showed that removed products had above average similarity with the purchased products across each knowledge representation. Removed products were, on average, in the 77th percentile ($95\%CI = 0.043$) of episodic similarities, the 62nd percentile ($95\%CI = 0.028$) of semantic similarities and the 79th percentile ($95\%CI = 0.053$) of hierarchical similarities observed between purchased products. This further suggests that products are removed because they are similar to other purchased products (i.e., analogous to confabulations) and not because they are distinct and thus irrelevant to one's goals.

C.2.11 Search arrival transitions

One concern may be that the design of the website biased shoppers towards retrieving products that were similar to each other. For example, shoppers could be biased towards retrieving hierarchically similar products by the category navigation or by the display of similar products in the search results. We therefore filtered the data so that it only contained transitions between choices before and after the use of a search bar (i.e., search arrivals) and re-ran the analyses reported in the main text. These transitions are perhaps most characteristic of memory-based search. To foreshadow, these analyses reproduce

the results presented in the main text, suggesting that they are not an artifact of the website’s design.

C.2.11.1 Permutation tests

Permutation tests were consistent with the results described in the main manuscript. The average trip-wise similarity between consecutively purchased items was significantly higher for the true order of purchases compared to the permuted order for episodic ($Median_{true} = 0.0421, IQR_{true} = 0.0751$ & $Median_{permuted} = 0.0136, IQR_{permuted} = 0.0141$) (Mann-Whitney $U = 379028331647.5$, $p < .0001$, $CLE = 0.7653$), semantic ($Median_{true} = 0.1974, IQR_{true} = 0.1020$ & $Median_{permuted} = 0.0737, IQR_{permuted} = 0.0685$) (Mann-Whitney $U = 221679814301.5$, $p < .0001$, $CLE = 0.8627$) and hierarchical similarities ($Median_{true} = 0.4556, IQR_{true} = 0.1248$ & $Median_{permuted} = 0.2545, IQR_{permuted} = 0.0483$) (Mann-Whitney $U = 150787174769.0$, $p < .0001$, $CLE = 0.9077$).

C.2.11.2 Correlations between representations

Table C.11: Spearman correlations between each of the similarity measures in the search-arrival dataset ($N = 3086716$)

	Episodic	Semantic	Hierarchy
Episodic	-	0.3254 [0.3244, 0.3264]	0.2192 [0.2181, 0.2203]
Semantic	-	-	0.3217 [0.3207, 0.3227]
Hierarchy	-	-	-

As shown in Table C.11, there was a small to moderate relationship between each representation of the sequential choices observed in the filtered search-arrival dataset (all significant $p < .0001$). This is consistent with the results shown above.

Table C.12: Standardised coefficients, significance values, upper (UB) and lower bounds (LB) for variables included in the mixed-effects linear regression predicting IRI using timestep, fit to the search-arrival dataset ($N = 2959878$)

	b	$p \leq$	95% LB	95% UB
Timestep	0.161	0.0001	0.160	0.162
Episodic similarity	-0.050	0.0001	-0.051	-0.049
Semantic similarity	-0.036	0.0001	-0.036	-0.035
Hierarchical similarity	-0.169	0.0001	-0.170	-0.168
Group Var	0.228			

C.2.11.3 Timestep regressions

Looking at behaviour over time, the pattern of results was consistent with that described in the main text. The mixed-effects regression converged ($ll = -3345165.74$). Importantly, average response times increased significantly over the duration of the trip ($b_{timestep} = 0.161$), even after accounting for different similarity measures (full regression equation shown in Table C.12). Sequential transitions also became increasingly dissimilar over time across representations of episodic ($r_s = -0.2596$, 95% CI $[-0.2607, -0.2585]$, $p \leq 0.0001$), semantic ($r_s = -0.0485$, 95% CI $[-0.0496, -0.0474]$, $p \leq 0.0001$) and hierarchical knowledge ($r_s = -0.1104$, 95% CI $[-0.1115, -0.1093]$, $p \leq 0.0001$).

C.2.11.4 Representation and IRI regression

We also reproduced the regression analyses of response times using each of the representations as predictors. The mixed-effects linear regression converged ($ll = -3523084.66$) and the full equation is shown in Table C.13. As in the main text, the partial regression coefficients for episodic ($\beta = -0.057$), semantic ($\beta = -0.034$), hierarchical similarity ($\beta = -0.173$) were all negatively related to response time. These results are therefore consistent with those shown in the main analyses. As before, hierarchical knowledge explained the most amount of variance in IRIs, even when filtering to memory-based choice transitions that occurred before and after searches.

Table C.13: Standardised coefficients, significance values, upper (UB) and lower bounds (LB) for variables included in the mixed-effects linear regression model predicting IRI using each similarity measure, fit to the search-arrival dataset

	b	$p \leq$	95% LB	95% UB
Intercept	3.092	0.0001	3.089	3.095
Episodic	-0.057	0.0001	-0.058	-0.056
Semantic	-0.034	0.0001	-0.035	-0.033
Hierarchy	-0.173	0.0001	-0.174	-0.172
Choices remaining	-0.172	0.0001	-0.173	-0.171
Visit ID	0.238			

Table C.14: Nested model comparisons (using AIC and BIC) for regressions predicting IRI using each similarity measure, fit to the search-arrival dataset. Models were chosen using backward elimination.

Variables	AIC	BIC
Episodic, Semantic, Hierarchical and Choices rem.	7046183.33	7046273.93
Episodic, Hierarchical and Choices rem.	7051248.70	7051326.35
Hierarchical and Choices rem.	7071024.10	7071088.81

We also performed feature selection with mixed-effects regressions predicting IRI with each similarity measure. This was to confirm that — as before — IRIs were best explained by a combination of three knowledge representations. As shown in Table C.14, fits were best for the model containing all representations.

C.2.11.5 Retrieval model comparison

We also re-fit the SAM retrieval models to the search only dataset. As shown in Table C.15, a model containing multiple representations provided the best fit to the data. As before, hierarchical knowledge received the highest attention weight, further emphasising its importance when retrieving options from memory.

Table C.15: The % BIC improvement over the random baseline and the mean attention weights (with 95% confidence intervals) for each of the candidate retrieval models, fit to the search-arrival dataset. Results show that including representations of all knowledge formats provides the best fit to the data (shown in bold)

	Δ BIC (%)	Episodic	Semantic	Hierarchy
Episodic	8.28	0.309 (0.001)		
Semantic	4.37		0.093 (0.001)	
Hierarchy	27.80			2.398 (0.021)
Episodic & Semantic	11.20	0.276 (0.001)	0.068 (0.001)	
Semantic & Hierarchy	30.47		0.057 (0.001)	2.306 (0.024)
Episodic & Hierarchy	32.63	0.187 (0.001)		2.193 (0.022)
Multiple	34.59	0.172 (0.001)	0.046 (0.001)	2.144 (0.025)

C.2.11.6 Forgotten and removed item regressions

Table C.16: Standardised coefficients, significance values, upper (UB) and lower bounds (LB) for variables included in a multiple linear regression predicting the number of forgotten items, fit to the search-arrival dataset

	b	$p \leq$	95% LB	95% UB
Intercept	0.2904	0.0001	0.285	0.296
Episodic	-0.0467	0.0001	-0.052	-0.041
Semantic	0.0065	0.0190	0.001	0.012
Hierarchy	-0.0565	0.0001	-0.062	-0.051
Choices remaining	-0.0486	0.0001	-0.054	-0.043

Table C.17: Standardised coefficients, significance values, upper (UB) and lower bounds (LB) for variables included in a multiple linear regression predicting the number of removed items, fit to the search-arrival dataset

	b	$p \leq$	95% LB	95% UB
Intercept	3.4484	0.0001	3.407	3.490
Episodic	-0.0048	0.8230	-0.047	0.037
Semantic	0.2305	0.0001	0.189	0.272
Hierarchy	-0.9995	0.0001	-1.042	-0.957
Total choices	2.2511	0.0001	2.209	2.293

Finally, we predicted the number of forgotten items and removed prod-

ucts using the attention weights estimated from the SAM model fit to the filtered dataset. The regression predicting the number of forgotten items was significant ($F_{4,98478} = 237.2$, $p < 0.0001$, $R^2=.010$) and — as shown in Table C.16 — the coefficients for episodic, semantic and hierarchy attention weights followed the directions reported in the main text. The regression predicting the number of removed items was also significant ($F_{4,98478} = 3133$, $p < 0.0001$, $R^2=.113$) and — as shown in Table C.17 — the coefficients followed the directions reported in the main text, with higher attention to semantic knowledge positively predicting removals and higher attention hierarchical knowledge negatively predicting removals. Note that the coefficient for episodic knowledge was not significant in this regression, although we had no apriori hypotheses about its relationship with removals. This further reinforces the assertion that these results are a product of memory retrieval and not an artifact of the website’s design.

As a whole, these results are consistent with the major findings presented in the main text. This supports the claim that these findings reflect memory retrieval processes used by shoppers as they searched for products, given that using the search-bar is perhaps most representative of memory-based search. Despite this, it is important to clarify that other navigation strategies likely reflect memory processes too. For example, shoppers may use their past choice as the basis for determining which category to select next on the drop-down menu. Or choosing a special offer may trigger ideas about related products that are required. This is supported by the fact that results are so similar between the full dataset (reported in the main text) and the search-arrival dataset (reported in this section).

C.3 Analyses of food fluency data

We further tested the key claims of this article by fitting our retrieval model to data from a controlled experiment that explicitly tested memory retrieval

for food (Zemla et al., 2020a). In this experiment — originally conducted by Zemla et al. (2020a) — fifty participants were asked to list as many food items that they could think of within three minutes. Much like searches in a search-bar, each retrieval was typed into a text box (further details can be found in the original article, Zemla et al. 2020a). This experiment therefore makes an excellent testbed for evaluating our model, in that it assesses the role of long-term memory and retrieval in a preferential domain but with a high-degree of experimental control. We evaluated the fit of our model to sequential retrievals from this experiment and — as before — hypothesised that sequential retrievals would be best explained by combining representations of episodic, semantic and hierarchical knowledge.

C.3.1 Method

C.3.1.1 Procedure

In the experiment, fifty participants located in the United States were recruited via Amazon Mechanical Turk. Participants completed three separate fluency tasks (animals, tools and foods). Each category was repeated three times and the order of the categories was pseudo-randomised, whilst ensuring that no category was repeated twice. Participants were told not to repeat items within lists but that they could repeat items between lists. Participants had three minutes to complete each list. Each response was typed into a text box one at a time. We restricted our analyses to retrievals from the food category so that we could use the aforementioned embedding spaces, mirroring our analyses of the shopping data.

C.3.1.2 Data

Data was retrieved from the Github repository (Zemla et al., 2020b) associated with the original article (Zemla et al., 2020a). For each of the retrievals made in the experiment, we found a matching supermarket product by searching for

product descriptions that contained the word and prioritising matches with higher frequency of occurrence in the in-store dataset. A sample of the 40 most-frequently occurring retrievals with their corresponding products is available in Table C.18. We were unable to find matching products for a small proportion of the total retrievals (4.64%). Thus, transitions to and from retrievals with no matching products were dropped.

As with the clickstream data, we did not remove repetitions of words, as we believed these repetitions to be informative about one’s retrieval process.

The final dataset contained 3357 retrievals from 50 participants. Participants retrieved an average of 43.92 unique words (95%CI = 4.71) over all lists and 22.27 unique words (95%CI = 1.85) within each list.

C.3.1.3 Model

We used the same retrieval model used to explain consumer choices to estimate the probability of each retrieval. As before, we included similarities between the current and all remaining retrievals in the model denominator.

We also used the same fitting procedure as with the consumer choices and compared models containing one, two and three knowledge representations. These were compared with a random baseline, which predicted an equal probability $\frac{1}{n}$ for each choice at each timestep, where n represents the number of retrievals remaining. We therefore report the % BIC improvement over the random baseline for each model.

Due the hierarchical nature of the data (i.e., multiple lists per participant), we estimated parameters for each list separately, for each participant separately (by concatenating each list per participant) and for the first list only. We report separate model comparisons for each of these three estimation procedures.

C.3.2 Results and discussion

We compared nested models containing different numbers of knowledge representations using three estimating procedures; namely, treating each list sep-

arately (results shown in Table C.19), each participant separately (collapsing over lists, results shown in Table C.20) and to the first list for each participant (results shown in Table C.21). As shown in Tables C.19, C.20 and C.21, the pattern of results were consistent between estimation procedures and with those reported in the main text. Namely, models containing multiple knowledge representations provided the best fit to the semantic fluency data (e.g., 9.92% BIC improvement over the random baseline when models were estimated for each list separately). In addition, hierarchical knowledge received the highest attention weight in the multiple representation model, followed by semantic then episodic knowledge. Hierarchical knowledge also drove the highest improvement in BIC when fit as a single representation. All of these findings are consistent with those reported in the main text.

Thus — across all estimation procedures — we observed a general pattern of results that is strikingly consistent with our major findings. It is reassuring that these results can be recovered from memory retrievals observed in experimental conditions for which there is a high degree of control. This consistency suggests that the online shoppers described in the clickstream dataset depended on similar memory retrieval processes when deciding what to choose next. Moreover, these results support our general claims that past retrievals serve as cues to query multiple sources of long-term knowledge, which combine to determine subsequent retrievals in preferential domains.

Bibliography

- Abbott, J. T., Austerweil, J. L., and Griffiths, T. L. (2015). Random walks on semantic networks can resemble optimal foraging. *Psychological Review*, 122(3):558–569.
- Hills, T. T., Jones, M. N., and Todd, P. M. (2012). Optimal foraging in semantic memory. *Psychological Review*, 119(2):431–440.

- Kahana, M. J. (2020). Computational Models of Memory Search. *Annual Review of Psychology*, 71(1):107–138.
- Mikolov, T., Sutskever, I., Chen, K., Corrado, G., and Dean, J. (2013). Distributed representations of words and phrases and their compositionality. In *Proceedings of the 26th International Conference on Neural Information Processing Systems - Volume 2*, NIPS'13, pages 3111–3119, Red Hook, NY, USA. Curran Associates Inc.
- Raaijmakers, J. G. W. and Shiffrin, R. M. (1980). SAM: A Theory of Probabilistic Search of Associative Memory. In Bower, G. H., editor, *Psychology of Learning and Motivation*, volume 14, pages 207–262. Academic Press.
- Stachenfeld, K. L., Botvinick, M. M., and Gershman, S. J. (2017). The hippocampus as a predictive map. *Nature neuroscience*, 20(11):1643.
- Zemla, J. C., Cao, K., Mueller, K. D., and Austerweil, J. L. (2020a). SNAFU: The Semantic Network and Fluency Utility. *Behavior Research Methods*, 52(4):1681–1699.
- Zemla, J. C., Cao, K., Mueller, K. D., and Austerweil, J. L. (2020b). SNAFU: the Semantic Network and Fluency Utility. [Dataset] <https://github.com/AusterweilLab/snafu-py> accessed on 13/10/2021.

Table C.18: Matching retailer products for the 40 most frequently occurring retrievals reported by Zemla et al. (2020a). Note that brand names have been redacted.

item	Product	No. occurrences
apple	GALA APPLE MINIMUM 5 PACK	101
pizza	MARGHERITA PIZZA 245G	95
chicken	FREE RANGE WHOLE CHICKEN 1KG-2.3KG	90
banana	BANANAS LOOSE	80
bread	SLICED WHITE BREAD 800G	76
cheese	MATURE CHEDDAR CHEESE 350 G	76
carrot	CARROTS LOOSE	74
orange	PURE ORANGE JUICE SMOOTH 1 LTR	70
steak	ABERDEEN ANGUS STEAK MINCE 500G	64
icecream	SOFT SCOOPVANILLA 2 LITRES	63
grape	SEEDLESS GRAPES 500G	62
potato	MARIS PIPER POTATOES 2.5KG	61
tomato	SALAD TOMATOES 6 PACK	61
hamburger	4 BRITISH BEEF STEAK BURGERS 454G	61
spaghetti	SHORT SPAGHETTI PASTA 500G	58
cake	CHOCOLATE CAKE ROLL 10 PACK	58
rice	M/WAVE BASMATI RICE 250G	58
broccoli	BROCCOLI LOOSE	57
lettuce	ICEBERG LETTUCE EACH	56
corn	CORN ON THE COB TWINPACK	52
strawberry	STRAWBERRIES 400G	52
onion	BROWN ONIONS LOOSE	51
egg	MEDIUM FREE RANGE EGGS 6 PACK	49
turkey	BRITISH ROAST TURKEY SLICES 125 G	49
frenchfry	CRISPY FRENCH FRIES 900G	46
taco	CRNCHY TACO SHELLS X12 156G	44
beans	BEANS IN TOMATO SAUCE 415G	43
hotdog	CLASSIC FRANKFURTER HOT DOGS 10 PK 350G	43
pasta	FUSILLI PASTA TWISTS 1KG	43
spinach	ORGANIC SPINACH 200G	42
cereal	VARIETY PACK CEREAL 8 PACK	42
ham	HAM SLICES 125 G	42
soup	CREAM OF TOMATO SOUP 400G	41
yogurt	GREEK STYLE YOGHURT 500G	41
pear	CONFERENCE PEARS PACK 610G	41
bacon	UNSMOKED BACK BACON RASHERS 300G	40
sushi	SUSHI NORI 11G	40
beef	LEAN BEEF STEAK MINCE 5% FAT 250G	39
pie	TEAK & ALE PUFFPASTRY PIE 500G	39
pineapple	PINEAPPLE LOOSE	39

Table C.19: The % BIC improvement over the random baseline and the mean attention weights (with 95% confidence intervals) for each of the candidate retrieval models, *fit to each retrieval list separately*. Results show that including representations of all knowledge formats provides the best fit to the data (shown in bold)

	Δ BIC (%)	Episodic	Semantic	Hierarchy
Semantic	1.496		0.089 (0.073)	
Episodic	2.176	0.207 (0.06)		
Hierarchy	7.818			2.071 (0.874)
Episodic & Semantic	3.247	0.922 (1.087)	1.702 (2.62)	
Semantic & Hierarchy	8.750		0.055 (0.041)	1.894 (0.724)
Episodic & Hierarchy	9.153	0.223 (0.149)		2.347 (1.461)
Multiple	9.917	0.669 (0.753)	0.282 (0.474)	3.777 (3.082)

Table C.20: The % BIC improvement over the random baseline and the mean attention weights (with 95% confidence intervals) for each of the candidate retrieval models, *fit to each participant and collapsing over multiple lists*. Results show that including representations of all knowledge formats provides the best fit to the data (shown in bold)

	Δ BIC (%)	Episodic	Semantic	Hierarchy
Semantic	0.910		0.031 (0.007)	
Episodic	1.120	0.104 (0.033)		
Hierarchy	6.879			1.212 (0.156)
Episodic & Semantic	1.798	0.092 (0.032)	0.026 (0.007)	
Episodic & Hierarchy	7.359	0.067 (0.028)		1.166 (0.156)
Semantic & Hierarchy	7.364		0.021 (0.007)	1.166 (0.159)
Multiple	7.760	0.06 (0.028)	0.018 (0.007)	1.129 (0.158)

Table C.21: The % BIC improvement over the random baseline and the mean attention weights (with 95% confidence intervals) for each of the candidate retrieval models, *fit to the first retrieval list from each participant*. Results show that including representations of all knowledge formats provides the best fit to the data (shown in bold)

	Δ BIC (%)	Episodic	Semantic	Hierarchy
Semantic	0.722		0.028 (0.01)	
Episodic	1.194	0.133 (0.05)		
Hierarchy	5.998			1.307 (0.57)
Episodic & Semantic	1.746	0.135 (0.058)	0.025 (0.01)	
Semantic & Hierarchy	6.500		0.022 (0.01)	1.279 (0.577)
Episodic & Hierarchy	6.693	0.095 (0.046)		1.254 (0.568)
Multiple	7.175	0.174 (0.191)	0.032 (0.029)	1.323 (0.743)

Appendix D

Chapter 4 Appendix

D.1 Experiment 1

D.1.1 Preferences did not change over time

One possible confound in the robot experiment is that participants became more likely to select a particular image as their first preference over trials. In order to evaluate this claim, preferences for each choice type were broken down by trial. Three post-hoc linear regressions conducted on each choice type showed non-significant effects of trial when predicting the number of selections ($p > 0.05$), suggesting that participants did not become biased in their choices over time.

D.1.2 Preference change doesn't vary as a function of political affiliation

One possible explanation for the results discovered in experiment two (i.e., the study of political opinions) is that only self-identifying Republicans adjust their preferences to be coherent with their past choices. To test this claim, the first study was re-run, this time asking participants to state their preferred political party of the Democrats and the Republicans at the end of the experiment.

The re-run used the same participant selection criteria as the first itera-

tion, except that participants were required to be from the United States only. Of the remaining 953 participants, 54.77% were female. The mean age of participants was 38.80 ($SD = 12.18$). This time, the experiment was conducted in May 2019.

Firstly, results showed a direct replication of the overall effect. The initial omnibus test proved significant (non-parametric Friedman test of differences among repeated-measures $\chi^2 = 104.75, p < 0.0001$). As before, summed preferences for the chosen-unique patterns (Median = 9, $IQR = 4.0$) were stronger than that for the shared items (Median = 10, $IQR = 3.0$) ($Z = -4.10, p < 0.0001, r = 0.133$) and the non-chosen items (Median = 11, $IQR = 5.0$) ($Z = -10.99, p < 0.0001, r = 0.36$). A final test also revealed that preferences for the shared items were stronger than that for non-chosen items ($Z = -9.77, p < 0.001, r = 0.32$)¹.

Looking at self-identifying Democrat participants alone ($N = 625$), results showed the same pattern of results as described above, suggesting that they update their preferences to be in line with their prior choices (non-parametric Friedman test of differences among repeated-measures $\chi^2 = 81.05, p < 0.0001$). As before, summed preferences for the chosen-unique patterns (Median = 9, $IQR = 4.0$) were stronger than that for the shared items (Median = 10, $IQR = 3.0$) ($Z = -3.42, p = 0.0006, r = 0.14$) and the non-chosen items (Median = 11, $IQR = 5.0$) ($Z = -9.40, p < 0.0001, r = 0.376$). A final test also revealed that preferences for the shared items were stronger than that for non-chosen items ($Z = -8.36, p < 0.0001, r = 0.33$).

Similarly for Republican-identifying participants alone ($N = 328$), results showed the same pattern of results as described above (non-parametric Friedman test of differences among repeated-measures $\chi^2 = 25.45, p < 0.0001$). As before, summed preferences for the chosen-unique patterns (Median = 9, $IQR = 4.0$) were stronger than that for the shared items (Median = 10, $IQR = 3.25$) ($Z = -2.25, p = 0.0245, r = 0.12$) and the non-chosen items (Median =

¹All Wilcoxon-signed rank tests were evaluated against a Holm-Bonferroni corrected alpha value for multiple comparisons

11, $IQR = 5.0$) ($Z = -5.74, p < 0.0001, r = 0.32$). A final test also revealed that preferences for the shared items were stronger than that for non-chosen items ($Z = -5.11, p < 0.0001, r = 0.28$).

Thus, there is no evidence to suggest that only self-identifying Republicans update their preferences to be maximally coherent with their past choices.

D.2 Experiment 2

D.2.1 Preference change varies within political topic

Further experimental results from the study of people’s political preferences are now reported. First, additional results from the main ANOVA are reported. Then, analyses are broken down by political topic.

As the slider response is an ordinal variable, the data was analyzed using a non-parametric two-way analysis of variance (ANOVA) (Hocking, 1985)². This ANOVA compared the influence of two between-groups independent variables (selected candidate opinion and political affiliation) on the participants’ normalized slider values. Political affiliation contained two levels (Democrat or Republican identifying), as did the selected candidate opinion (left-wing or right-wing).

In addition to the results presented in the main text, the ANOVA revealed a main effect of political affiliation. This yielded an F -ratio of $F(1, 952) = 83.72, p < 0.001, CL = 0.640$ ³, indicating that the average slider values of self-identifying Democrats ($Median = 34.0, IQR = 73.00$) were significantly lower than that of Republicans ($Median = 67.0, IQR = 70.75$).

Going further, we now evaluate how peoples levels of agreement varied depending on both their party affiliation and the topic that participants were being asked to decide on. Breaking down by individual topic allows us to gain

²We are grateful to the creators of the *RFit* package for implementing this for the R programming language

³For all main effects and Mann-Whitney U tests, we report the common language (CL) effect size (Mcgraw and C. P. Wong, 1992)

further insight about the malleability of certain beliefs held by self-identifying Democrats and Republicans. For example, one possibility is that participants only adjusted their slider values to be consistent with their chosen candidate's opinion in cases where that opinion was consistent with their chosen political affiliation. By analysing participants within political affiliations and topics, we can see whether participants move towards beliefs that run counter to what would be expected by their party affiliation. A further set of six Mann-Whitney U tests were therefore run, comparing the extent to which the normalized slider values (i.e., levels of right-wing agreement) changed depending on the revelation of the chosen candidates left and right-wing opinions, for each of the three topics across each of the two party affiliations.

For Democrat-identifying participants, participants appeared to be affected by the revelation of their chosen candidate's beliefs if it pertained to trade or abortion. Results showed that normalized slider values of self-identifying Democrats were significantly lower on average if their candidate later revealed a left-wing opinion about trade ($Median=73.0$, $IQR = 47.00$) compared to a right-wing opinion about trade ($Median=82.0$, $IQR=29.00$) ($U = 4688.5$, $p<0.01$, $CL = 0.579$). Similarly, results showed that responses of Democrat-identifying participants were significantly different if their candidate later revealed a left-wing opinion about abortion ($Median = 1.0$, $IQR=13.00$) compared to a right-wing opinion about abortion ($Median = 3.0$, $IQR=27.50$) ($U = 4311.5$, $p<0.02$, $CL = 0.512$). However, responses of Democrat-identifying participants were not significantly different if their candidate later revealed a left-wing opinion about immigration ($Median=29.5$, $IQR=38.25$) compared to a right-wing opinion about immigration ($Median=24.0$, $IQR=44.25$) ($U = 4352.5$, $p>0.05$, $CL = 0.442$).

For Republican participants, the average normalized slider values were significantly different across all topics depending on the revealed opinion of the chosen candidate. Results showed that normalized slider values of Republicans were significantly lower on average if their candidate later revealed a

left-wing opinion about trade ($Median=31.0$, $IQR = 43.5$) compared to a right-wing opinion about trade ($Median=67.0$, $IQR=51.50$) ($U = 1023.5$, $p<0.001$), $CL = 0.699$). In addition, results showed that responses of Republican participants were significantly lower if their candidate later revealed a left-wing opinion about abortion ($Median = 60.0$, $IQR=88.00$) compared to right-wing opinion about abortion ($Median = 96.0$, $IQR=61.5$) ($U = 1136.0$, $p<0.01$, $CL = 0.593$). Finally, responses of Republican participants were significantly lower on average if their candidate later revealed a left-wing opinion about immigration ($Median=65.5$, $IQR=61.25$) compared to a right-wing opinion about immigration ($Median=84.5$, $IQR=55.5$) ($U = 1080.0$, $p<0.015$, $CL = 0.610$).

These results therefore indicate that across nearly all topics, people of both major political affiliations adjusted their preferences to be consistent with the views of their chosen candidate. In all but one case (Democrats on the immigration topic), participants appeared willing to shift their apparent beliefs in directions that ran counter to their party affiliation. This likely would not have been enough to change participants' stated, binary preference for a given topic (assuming that binary preferences change when the normalized slider crosses the midpoint of 50). However, these results are supportive of the general idea proposed in this paper that preferences over choice attributes are updated to make past choices more likely.

D.2.2 Responses were not biased in favour of a particular party affiliation or slider response

Post-hoc analyses were conducted to test whether there were any biases in responses or whether the neutral topics elicited any differences in slider responses. A 2x2 chi-square was conducted to ensure that the process of voting did not alter participant's political affiliation. Results showed that voting for a candidate that later revealed a left or right-wing opinion did not affect participants' subsequent self-reported party affiliation ($\chi^2(1)=1.583$, $p>0.05$). In addition, a one-sample Wilcoxon signed-rank test was conducted to ensure

that participants did not have a bias towards responding in any particular direction (e.g. clicking more towards the left). Results revealed that participants' average, unnormalized slider values ($Median = 50$, $IQR = 76.00$) were not significantly different from 50 ($W = 218707.5$, $p > 0.05$).

Bibliography

Hocking, R. R. (1985). *The analysis of linear models*. Brooks/Cole.

Mcgraw, K. and C. P. Wong, S. (1992). A common language effect size measure. *Psychological Bulletin*, 111:361–365.

Declaration Form: referencing doctoral candidate's own published work(s) in thesis

Please use this form to declare if parts of your thesis are already available in another format, e.g. if data, text, or figures:

- have been uploaded to a preprint server;
- are in submission to a peer-reviewed publication;
- have been published in a peer-reviewed publication, e.g. journal, textbook.

This form should be completed as many times as necessary. For instance, if a student had seven thesis chapters, two of which having material which had been published, they would complete this form twice.

1. For a research manuscript that has already been published (if not yet published, please skip to section 2):

a. **Where was the work published?** (e.g. journal name)
Computational Brain & Behaviour

b. **Who published the work?** (e.g. Elsevier/Oxford University Press)
Springer Nature

c. **When was the work published?**

2020

d. **Was the work subject to academic peer review?** YES

e. **Have you retained the copyright for the work?** YES

2. For a research manuscript prepared for publication but that has not yet been published (if already published, please skip to section 3):

a. **Where is the work intended to be published?** (e.g. journal name)

b. **List the manuscript's authors in the intended authorship order:**

c. **Stage of publication:**

- Not yet submitted
- Submitted
- Undergoing revision after peer review
- In press

3. For multi-authored work, please give a statement of contribution covering all authors (if single-author, please skip to section 4):

A.H. wrote the initial draft of original article, performed analyses, wrote code, interpreted results and designed the methodology. T.E. helped write LDA code and interpret results. P.R. helped design the expert experiment and interpret results. R.P. helped with interpretation of results and funding. B.C.L. contributed to later drafts of the article, methodology, validation of results, funding and conceptualisation.

4. In which chapter(s) of your thesis can this material be found?

Chapter 2

5. Candidate's e-signature: Adam Hornsby
Date: 25/03/2022

6. Supervisor/senior author(s) e-signature: Bradley C. Love
Date: 24/03/2022

Declaration Form: referencing doctoral candidate's own published work(s) in thesis

Please use this form to declare if parts of your thesis are already available in another format, e.g. if data, text, or figures:

- have been uploaded to a preprint server;
- are in submission to a peer-reviewed publication;
- have been published in a peer-reviewed publication, e.g. journal, textbook.

This form should be completed as many times as necessary. For instance, if a student had seven thesis chapters, two of which having material which had been published, they would complete this form twice.

1. For a research manuscript that has already been published (if not yet published, please skip to section 2):

a. **Where was the work published?** (e.g. journal name)
Science Advances

b. **Who published the work?** (e.g. Elsevier/Oxford University Press)
American Association for the Advancement of Science

c. **When was the work published?**

2022

d. **Was the work subject to academic peer review?** YES

e. **Have you retained the copyright for the work?** YES

2. For a research manuscript prepared for publication but that has not yet been published (if already published, please skip to section 3):

a. **Where is the work intended to be published?** (e.g. journal name)

b. **List the manuscript's authors in the intended authorship order:**

c. **Stage of publication:**

- Not yet submitted
- Submitted
- Undergoing revision after peer review
- In press

3. For multi-authored work, please give a statement of contribution covering all authors (if single-author, please skip to section 4):

A.H. wrote the initial draft of original article, performed analyses, wrote software and designed the methodology. B.C.L. contributed to later drafts of the article, methodology, validation of results, funding and conceptualisation.

4. In which chapter(s) of your thesis can this material be found?

Chapter 3

5. Candidate's e-signature: Adam Hornsby
Date: 25/03/2022

6. Supervisor/senior author(s) e-signature: Bradley C. Love
Date: 24/03/2022

Declaration Form: referencing doctoral candidate's own published work(s) in thesis

Please use this form to declare if parts of your thesis are already available in another format, e.g. if data, text, or figures:

- have been uploaded to a preprint server;
- are in submission to a peer-reviewed publication;
- have been published in a peer-reviewed publication, e.g. journal, textbook.

This form should be completed as many times as necessary. For instance, if a student had seven thesis chapters, two of which having material which had been published, they would complete this form twice.

1. For a research manuscript that has already been published (if not yet published, please skip to section 2):

a. Where was the work published? (e.g. journal name)

Cognition

b. Who published the work? (e.g. Elsevier/Oxford University Press)

Elsevier

c. When was the work published?

2020

d. Was the work subject to academic peer review? YES

e. Have you retained the copyright for the work? YES

2. For a research manuscript prepared for publication but that has not yet been published (if already published, please skip to section 3):

a. Where is the work intended to be published? (e.g. journal name)

b. List the manuscript's authors in the intended authorship order:

c. Stage of publication:

- Not yet submitted
- Submitted
- Undergoing revision after peer review
- In press

3. For multi-authored work, please give a statement of contribution covering all authors (if single-author, please skip to section 4):

A.H. wrote the initial draft of original article, performed analyses, wrote software and designed the methodology. B.C.L. contributed to later drafts of the article, methodology, validation of results, funding and conceptualisation.

4. In which chapter(s) of your thesis can this material be found?

Chapter 4

5. Candidate's e-signature: Adam Hornsby
Date: 25/03/2022

6. Supervisor/senior author(s) e-signature: Bradley C. Love
Date: 24/03/2022