

Augmented Machine Learning and Optimization for Marketing

by

Yuting Zhu

B.A. in Economics and B.S. in Mathematics,
Renmin University of China, 2015

M.A. in Economics, University of Rochester, 2017

S.M. in Management Research,
Massachusetts Institute of Technology, 2020

Submitted to the Department of Management
in partial fulfillment of the requirements for the degree of

Doctor of Philosophy in Management

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

May 2022

© Massachusetts Institute of Technology 2022. All rights reserved.

Author
Department of Management
April 17, 2022

Certified by
Juanjuan Zhang
John D. C. Little Professor of Marketing
Professor, Marketing
Thesis Supervisor

Accepted by
Catherine Tucker
Sloan Distinguished Professor of Management
Professor, Marketing
Faculty Chair, MIT Sloan PhD Program

Augmented Machine Learning and Optimization for Marketing

by

Yuting Zhu

Submitted to the Department of Management
on April 17, 2022, in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy in Management

Abstract

This dissertation consists of three essays exploring how to augment machine learning and optimization methods for marketing management.

The first essay considers an augmentation of deep-learning-based recommender system for sales force management. Helping new salespeople succeed is critical for many organizations. We develop a deep-learning-based recommender system to help new salespeople recognize suitable customers, leveraging historical sales records of experienced salespeople. One challenge is how to learn from experienced salespeople’s own failures, which are prevalent but often do not show up in sales records. We develop a parsimonious model to capture these “missing by choice” sales records and incorporate the model into a neural network to form an augmented, deep-learning-based recommender system. We validate our method using sales force transaction data from a large insurance company. Our method outperforms common benchmarks in prediction accuracy and recommendation quality, while being simple, interpretable, and flexible. We demonstrate the value of our method in improving sales force productivity.

The second essay explores an augmentation of large-scale linear programming optimization method for targeting with constraints. Personalization, which aims to target different marketing actions to different customers, has attracted broad attention in both academia and industry. While most research has focused on training personalization policies without constraints, in practice, many firms face constraints when implementing these policies. For example, firms may face volume constraints on the maximum or minimum number of actions they can take, or on the minimum acceptable outcomes for different customer segments. They may also face fairness constraints that require similar actions with different groups of customers. These constraints can introduce difficult optimization challenges, particularly when the firm intends to implement personalization policies at scale. Traditional optimization methods face challenges solving large-scale problems that contain either many customers or many constraints. We show how recent advances in linear programming can be adapted to the personalization of marketing actions. We provide a new theoretical

guarantee comparing how the proposed method scales compared to state-of-the-art benchmarks (primal simplex, dual simplex and barrier methods). We also extend existing guarantees on optimality and computation speed, by adapting them to accommodate the characteristics of personalization problems. We implement the proposed method, and compare it with these benchmark methods on feasibility, computation speed, and profit. We conclude that, volume and similarity (fairness) constraints should not prevent firms from optimizing and implementing personalization policies at scale.

The third essay studies collective search in an organization. In this paper, we build a two-member two-period model to show that when a group of people with different preferences conduct search and make a decision together, they can benefit from making a commitment on the number of products to search ex ante when the search cost is very small or relatively large. The underlying mechanism is that, because of the preference divergence between group members, they tend to search fewer products and thus have lower expected utility in group search than in single-agent search, and making a commitment on the number of products to search can help mitigate the preference divergence problem in group search. If consumers can observe product prices before search and the firm sets product prices endogenously, the firm can benefit from letting consumers commit to the number of products to search ex ante if consumers search as a group and their search cost is small. We also consider several extensions to show the robustness and boundary conditions of our findings.

Thesis Supervisor: Juanjuan Zhang
Title: John D. C. Little Professor of Marketing
Professor, Marketing

Acknowledgments

I have been very fortunate to receive tremendous guidance, advice, encouragement and friendship during the five years I have spent at MIT. I would like to acknowledge the efforts of all people who helped and supported me.

First, I would like to thank my advisor, Juanjuan Zhang, who brought me to the marketing field and guided me through every single step during my PhD journey. Juanjuan steered me towards interesting research topics that balance scientific insights and managerial relevance, taught me how to make ideas happen step by step, and showed me a perfect role model for being a researcher and a teacher. Knowing that I can always rely on Juanjuan's advice and support, I am able to spend the wonderful five years to explore and enjoy research. I have benefited enormously from the countless inspiring and caring calls we had during the Covid period, and I do appreciate all the time and effort that Juanjuan has devoted to me. She is a splendid advisor and also a valuable friend.

I also want to thank the other members of my thesis committee, Duncan Simester and Tony Ke. Duncan and Tony also played an important role in the advancement of my thinking and academic progress. Duncan has taught me how to find the essence of an idea and conduct forward-looking research. I have learnt a lot from his thinking style and benefited greatly from his critical insights and feedbacks. Tony is my first co-author and we have the first paper for my academic career. The enjoyable and efficient collaboration experience offers me a great start for my PhD days.

I would also like to thank all other faculties of the Marketing department. The essays and my other research projects have benefited from valuable discussions with John Hauser, Birger Wernerfelt, Drazen Prelec, Sinan Aral, Catherine Tucker, David Rand, Dean Eckles, and Rahul Bhui. I am also thankful to my brilliant peers and friends at MIT marketing - Xinyu Cao, Artem Timoshenko, Shuyi Yu, Jeremy Yang, Madhav Kumar, Cathy Cao, Jerry Zhang, Matthew Cashman, Keyan Li, Marat Ibragimov, Yifei Wang, Jenny Allen, Jason Du, Cathy Chen, Graelyn Humiston, Lei Huang, Cameron Martel, Chengfeng Mao, James Li, and Reed Orchinik. I also

want to thank my co-authors Saiquan Hu (first essay), Haihao Lu (second essay) and Xinyu Cao (third essay) for sharing the research joy together.

Finally, I want to thank my parents for their unconditional love and support. They have given me the greatest freedom to pursue my dream and become who I am. I want to thank my partner Xiaoyan Zhao. The five-year long distance is not easy for us, and I am very lucky to have one person who can always know, understand and support my unspoken thoughts. I also want to thank our cat Maisui, who is the cutest boy on the planet!

Contents

1 Zero to One:

Sales Prospecting with Augmented Recommendation	15
1.1 Introduction	16
1.2 Related Literature	19
1.3 An Augmented Recommender System	22
1.3.1 Model of Missing-by-Choice Sales Records	22
1.3.2 Neural Network Framework	24
1.4 Data	27
1.4.1 Business Context	27
1.4.2 Sample, Variables, and Summary Statistics	28
1.4.3 Preliminary Data Patterns	30
1.5 Evaluating the Augmented Recommender System	31
1.5.1 Prediction Accuracy	31
1.5.2 Recommendation Quality	34
1.5.3 Recommender System Output Interpretation	35
1.6 Extensions	37
1.6.1 Probably Missing-by-Choice Sales Records	37
1.6.2 Heterogeneous Net Cost of Selling	38
1.6.3 Other Possible Extensions	38
1.7 Practical Value	39
1.8 Concluding Remarks	41

2	Training Scalable Personalization Policies with Constraints	43
2.1	Introduction	44
2.2	Related Literature	46
2.3	Constraints and Problem Setup	49
2.4	Algorithm and Theoretical Guarantees	53
2.5	Empirical Validation	57
2.5.1	Data and Profit Estimation	57
2.5.2	Design of the Validation Exercise	59
2.5.3	Validation Results	61
2.6	Conclusion	62
3	The Power of Commitment in Group Search	65
3.1	Introduction	66
3.1.1	Contributions to Literature	71
3.2	Main Model	73
3.2.1	Model Setup	73
3.2.2	Fixed-sample Strategy and Equilibrium Analysis	77
3.2.3	Sequential Strategy and Equilibrium Analysis	79
3.2.4	Fixed-sample versus Sequential Strategy	82
3.3	Implications for Firm’s Strategy	89
3.3.1	Single-agent search	89
3.3.2	Group search	91
3.3.3	Implications	92
3.4	Extensions	94
3.4.1	Recall Not Allowed in Sequential Strategy	94
3.4.2	Alternative Distribution Assumption: Normal Distribution	95
3.4.3	Alternative Voting Rules	96
3.5	Concluding Remarks	99
A	Tables	101

B	Figures	115
C	Chapter 1 Appendix	135
C.1	Details on the Augmented Recommender System	135
C.2	Recommended Customer Type by Salesperson Trait	136
C.3	Proportion of New Salespeople Yet to Make Their First Sale	137
D	Chapter 2 Appendix	139
D.1	Details of Algorithm 1	139
D.2	Profit Estimation and Prediction	140
E	Chapter 3 Appendix	143
E.1	Proof of $\xi_G^* = E_{X_{1B}}[V(\xi_G^*, X_{1B})]$ in Group Sequential Search	143
E.2	Proof of Proposition 2:	
	Fixed-Sample vs. Sequential in Group Search	143
E.3	Compare the Expected Value and Search Cost in Fixed-Sample vs. Sequential in Group Search	144
E.4	Proof of Lemma 3.2.4:	
	Fixed-Sample vs. Sequential in Single-Agent Search	145
E.5	Proof of Lemma 3.2.4: Group Search vs. Single-Agent Search	146
E.6	Firm's Strategy	148
E.6.1	Fixed-sample strategy under single-agent search	148
E.6.2	Sequential strategy under single-agent search	149
E.6.3	Fixed-sample strategy under group search	151
E.6.4	Sequential strategy under group search	153
E.7	Alternative Voting Rules	156
E.7.1	Unanimity voting rule	156
E.7.2	One-is-enough voting rule	157

List of Figures

B-1	A Typical Neural Network Structure.	115
B-2	Proposed Activation Function (Example of $c = 1$).	116
B-3	Our Proposed Neural Network Structure.	117
B-4	Illustration of Face Embedding Derivation	118
B-5	Neural Network Structure of the Augmented Deep Learning Based Rec- ommender System.	119
B-6	Training and Validation Loss: Missing Values Excluded.	120
B-7	Training and Validation Loss: Linear Activation.	121
B-8	Training and Validation Loss: Augmented.	122
B-9	Proportion of New Salespeople Who Have Not Made Their First Sale.	123
B-10	Distribution of Customer Types for Each Customer Trait.	124
B-11	Most Salespeople Served Just A Few Customer Types.	125
B-12	Percentage of Salespeople Who Served Each Customer Type.	126
B-13	Number of Failures before the First Sale.	127
B-14	Proportion of New Salespeople Who Have Not Made Their First Sale, Revisited.	128
B-15	Illustration of Differences between one-loop and two-loop iterations. .	129
B-16	Illustration of Proposition 2: Comparing Expected Utilities under the Two Strategies.	130
B-17	Illustration of Lemma 3.2.4 and Lemma 3.2.4.	131
B-18	Optimal Prices and Profit under Single Agent and Group Search. . .	132
B-19	Comparing Expected Utilities of the Two Strategies under Standard Normal Distribution Assumption.	133

B-20 Compare Fixed-Sample Strategy to Sequential strategy under Alternative Voting Rules. 134

List of Tables

A.1	Summary Statistics of Salesperson and Customer Traits.	101
A.2	Prediction Accuracy.	102
A.3	Recommendation Quality.	103
A.4	Recommendation Quality by Salesperson Experience (One Customer Type Recommended).	104
A.5	Recommendation Quality by Salesperson Experience (Multiple Customer Types Recommended).	105
A.6	Importance of Salesperson and Customer Traits.	106
A.7	Most Recommended Customer Type by Salesperson Trait.	107
A.8	Extensions to the Augmented Recommender System.	108
A.9	Proportion of New Salespeople Who Have Not Made Their First Sale.	109
A.10	Constraints in Recent Targeted Marketing Papers.	110
A.11	Segments, Markets and Scenarios.	111
A.12	Algorithms' Performance Comparison.	112
A.13	Summary Statistics of Targeting Variables.	113

Chapter 1

Zero to One:

Sales Prospecting with Augmented Recommendation

In this essay, we develop a deep-learning-based recommender system to help new salespeople recognize suitable customers, leveraging historical sales records of experienced salespeople. One challenge is how to learn from experienced salespeople’s own failures, which are prevalent but often do not show up in sales records. We develop a parsimonious model to capture these “missing by choice” sales records and incorporate the model into a neural network to form an augmented, deep-learning-based recommender system. We validate our method using sales force transaction data from a large insurance company. Our method outperforms common benchmarks in prediction accuracy and recommendation quality, while being simple, interpretable, and flexible. We demonstrate the value of our method in improving sales force productivity.¹

¹This chapter is collaborated with Saiquan Hu (Associate Professor at Hunan University) and Juanjuan Zhang (John D. C. Little Professor of Marketing at MIT Sloan School of Management).

1.1 Introduction

Over 13 million Americans worked in sales-related jobs in 2020, representing nearly 10% of the labor force.² Without requiring an advanced or even a college degree, the sales occupation is considered accessible to many. “Sales is the hardest easy job in the world,” says Bob Franco, author and veteran of sales ([70]). Yet, sales is the *hardest* easy job. It is by far one of the most stressful occupations according to a PayScale survey ([92]). Starting a sales career is particularly challenging. For new salespeople, the process of converting strangers into paying customers is fraught with failure. In the company we partnered with for this study (discussed in more detail in subsequent sections), 70% of new salespeople failed to make a sale 90 days into their job.

New salespeople’s productivity hurdle can be costly to companies. Sales is already the most expensive marketing function. North American companies spend \$1.2 trillion on sales each year ([2]).³ U.S. companies spend \$15 billion a year on sales force training alone and \$800 billion on sales force incentives ([161]). New salespeople’s productivity hurdle undermines these hefty investments. It affects company revenue and customer experience. Moreover, it threatens sales force retention. By February 2021, 91% of the salespeople who had joined our partner company since 2015 had left. Notably, 67% left without ever making a sale, citing the frustration of getting started as a common reason to quit.⁴

In this paper, we develop a deep learning based recommender system to help new salespeople improve their productivity. We focus on “sales prospecting.” For each new salesperson, we offer personalized recommendations of customer types (e.g., female customers below 40) with higher conversion potential, based on deep learning of historical data on sales revenue, salesperson traits, and customer traits.

This approach suits the sales context for the following reasons. First, new salespeople often need help recognizing promising customers. In a recent HubSpot report,

²Source: U.S. Bureau of Labor Statistics (<https://www.bls.gov>).

³This is compared with \$165 billion annual spending on traditional marketing and \$36 billion on digital.

⁴For context, the annual sales force turnover rate is twice the average of the entire labor force (source: <https://hbr.org/2017/07/how-to-predict-turnover-on-your-sales-team>).

over 40% of salespeople find sales prospecting the most difficult part of their job.⁵ Second, sales outcome often depends on the match between the salesperson and the customer on various traits such as communication style ([175]), gender ([114]), and race ([147]). Recommender systems are particularly effective at uncovering such personalized match value between users and recommended items ([1]). Third, salesperson-customer match can be complex. It may be driven by convoluted interactions of salesperson/customer traits and these traits may be high-dimensional. For example, appearance has been shown to influence trust (e.g., [57]) whereas different customers may value trust differently. We thus focus on deep learning based recommender systems (e.g., [185]) to capture the intricacies of salesperson-customer match and to embrace unstructured data (appearance data in our application).

What is new in our approach is that we extend standard deep learning based recommender systems to capture an important feature of the sales context – salespeople’s failures, which are prevalent but often do not show up in historical sales records. For instance, it may appear that a salesperson has never made a sale to female customers below 40. In standard deep learning based recommender systems, these instances are often treated as missing observations and excluded from model training (e.g., [1, 31, 118, 146]). We emphasize a different method. We argue that these instances are substantively meaningful; they not only should be included in training, but also included in a theoretically constructed way to capture the behavioral process they represent.

More specifically, we emphasize that sales records may be *missing by choice* for some salesperson-customer dyads. It could be that a salesperson tried but failed to convert a customer type. It could also be that a salesperson decided not to sell to a customer type because the cost of selling exceeded the expected gain. Either way, the fact that the salesperson did not sell to a customer type is informative; a new and similar salesperson should probably avoid this customer type. Missing-by-choice sales records can be particularly meaningful in the sales function. They allow new

⁵Source: <https://blog.hubspot.com/sales/sales-statistics>. Salespeople are responsible for customer prospecting in our partner company and many others. However, as we will discuss, our approach extends to contexts where companies assign prospects to salespeople.

salespeople to learn from not only their predecessors’ success but also their lack of success.

Computationally, we develop an *augmented recommender system* to explicitly account for missing-by-choice sales records in deep learning based recommender systems. We specify a parsimonious behavioral model in which a salesperson sells to a customer type only if the expected gain outweighs the cost. We incorporate this behavioral model into a neural network structure by proposing a new activation function, which outputs a sales record only if sales revenue exceeds a cost threshold to be estimated. This allows us to augment standard deep learning based recommendation systems in a simple, interpretable, and flexible way.

We train, validate, and test our augmented recommender system using data from a major insurance company headquartered in Shanghai. The data include 12,149 salespeople and 409,840 insurance transactions from January 2015 through February 2021. Our augmented recommender system significantly outperforms a set of benchmarks on two common criteria of recommender system performance: prediction accuracy (mean square error) and recommendation quality (F1 and NDCG scores, to be explained later). In particular, our method significantly outperforms deep learning based recommender systems that either exclude missing sales records from model training or retain them but use the standard linear activation function.

Because our recommender system augmentation is based on modeling the sales-data generating process, it can be easily extended to accommodate a range of scenarios depending on why sales records are missing. We present two such extensions, where salespeople may not get to consider all customer types and where the cost threshold may be heterogeneous across salespeople. Both extensions further improve our method’s recommendation quality.

Finally, simulation results suggest that our augmented recommender system can be practically valuable. Compared with a naïve, random search strategy that new salespeople may follow, our recommender system may reduce the number of failures before the first sale by as high as 40%. In doing so, our recommender system can reduce the proportion of unproductive new salespeople by 27% 90 days into their job.

Our contributions are two-fold. Substantively, we develop a data-driven process to help new salespeople recognize suitable prospects and help companies improve sales force management. In doing so, we showcase a novel application of recommender systems in the sales context. Methodologically, we augment deep learning based recommender systems by retaining, modeling, and retrieving information from missing-by-choice observations. We demonstrate the efficacy of the augmented recommender system for sales prospecting, but we expect it to be broadly relevant wherever historical records for some user-item dyads are missing by choice, as opposed to missing by chance. For example, if a Netflix user has never rated a movie genre or if an Amazon customer has never bought from a product category, our augmented recommender system can be applied to derive information from these seemingly missing records. In the following section, we discuss our contribution to the literature in more detail.

1.2 Related Literature

Sales force management is central to marketing research. Sales force compensation design in particular has attracted significant research and generated extensive insights on how to motivate salespeople (see [40, 128] for comprehensive reviews). To the extent that compensation policies are not always feasible to change ([49]), another line of papers explore non-monetary sales force management strategies. These include optimizing the composition of sales teams ([30]), empowering salespeople with genetic self-knowledge ([82]), and training sales agents using artificial-intelligence coaches ([120]). Our paper contributes another non-monetary sales force management tool, a recommender system that helps salespeople recognize suitable prospects.

Recommender systems are a powerful way to help individuals choose among many options, often drawing on their many peers' choices ([1, 146]). [8]'s seminal paper introduces recommender systems to marketing and develops a Bayesian preference framework for recommendation. A stream of marketing papers have extended and applied recommender systems along important dimensions, such as optimizing purchase lift ([18]), personalizing content offerings adaptively ([43]), generalizing con-

sumer preference models for experienced goods ([42]), eliciting consumer preferences for complex products ([100]), modeling topics to leverage rich product information ([9]), recommending options to help consumers learn their preferences ([63]), using consumer search data ([76]) or on-boarding surveys ([55]) to overcome the cold-start problem, and evaluating welfare implications of personalized rankings ([56]).⁶

We contribute to the recommender-system literature by developing an augmented, deep learning based recommender system to incorporate missing-by-choice data. A particularly relevant paper in marketing is [180], who address the selection bias in consumer ratings with a hierarchical Bayes model that jointly captures rating value and rating incidence. In computer science, although most recommender-system papers focus on ways to better fit data, there is a line of research on debiasing data (see [31] for a systematic review). Solutions include probabilistic models ([93]) and propensity scores ([149]), both in the context of matrix factorization based recommender systems.⁷ Our paper shares the same view that there is information to be gained from including missing data and modeling why they are missing. The difference is that we develop a way to incorporate missing-by-choice data in deep learning based recommender systems.

Deep learning based recommender systems are becoming prevalent in both academic research and industry use (see [185] for a review). Successful applications include the recommender systems for YouTube videos ([45]) and Google Play apps ([36]), among many others. Deep learning allows recommender systems to capture the potentially complex, nonlinear relationship between users and items and to handle various forms of unstructured data ([83]). Deep learning can also be integrated

⁶The cold-start problem refers to the lack of historical data on user-item interactions to inform recommendation. In our paper, new salespeople are cold-start users by definition. We overcome this challenge using deep learning based recommendation, drawing on “side information” about salespeople (e.g., demographics, appearance). In a related paper, [139] address the cold-start problem in customer relationship management, using probabilistic machine learning to incorporate customer side information.

⁷Matrix factorization is a technique to decompose users’ reactions to various items into lower-rank user and item matrices. Matrix factorization based recommender systems are popular in industry (e.g., [81, 85, 107]). However, they are often subject to challenges such as the cold-start problem (e.g., [17, 173]), ability to handle complex unstructured data (e.g., [39, 172]), and scalability (e.g., [133]).

with other methodologies, fueling innovations such as neural network matrix factorization (e.g., [62]) and Bayesian deep learning (e.g., [171]). A major downside of deep learning based recommendation though is its lack of explainability; its numerous parameters and activation functions are often not easily interpretable ([145, 187]). In addition, overparameterized deep learning models may overfit the training data and fail to generalize out-of-sample ([183]).

Our augmentation can thus be valuable in three ways. First, as we show, this augmentation can further improve the performance of deep learning based recommender systems in situations where data are missing by choice – and these situations are likely the norm rather than the exception ([122]). Second, our augmentation is interpretable; at its core is an activation function based on a microfounded user behavior model.⁸ Because of this feature, as we show, the augmentation can be extended in interpretable ways to capture different user behaviors. Third, our augmentation can be easily implemented. In its simplest case, it requires only one additional parameter to be estimated (the cost threshold). This helps maintain the scalability of the recommender system and mitigate overfitting concerns.

Last, our paper is broadly related to the literature on using machine learning for personalization and targeting. Experimental data are often used for strategy evaluation in this literature ([58, 72, 152]). Although experiments offer a clean way to assess causal treatment effects, they may constrain the number of personalization or targeting actions that can be evaluated in one study. By contrast, we use historical sales data for model training and evaluation. There is no theoretical upper limit to the number of recommendations that can be tested. This approach is in line with [89] and [181], who use clickstream data to study large sets of personalized search actions. As such, the focus of our paper is prediction, consistent with the literature summarized in [143], as opposed to treatment effect estimation. However, these two approaches can complement each other. For instance, predictive models can help select a feasible number of high-potential recommendation strategies to be tested experimentally.

⁸In a recent study of brand selfies, [88] also modify and add layers to a neural network. Their goal is better classification and interpretation.

1.3 An Augmented Recommender System

In this section, we present the construction of our augmented recommender system. As an overview, we incorporate a model of missing-by-choice sales records into a neural network framework. We first present the model and then the neural network adjustment.

1.3.1 Model of Missing-by-Choice Sales Records

We strive to specify a parsimonious model to capture the data generation process behind sales records, including the process by which a subset of them are missing. We distinguish between two types of sales records: latent sales and observed sales. Latent sales represent the underlying sales outcome for each salesperson-customer dyad. Observed sales are the sales records that actually appear in the data. Our model aims to describe the latent sales generation process and the relationship between latent and observed sales.

Let $y_{ij}^* \in R$ denote the latent sales revenue generated by salesperson $i \in \{1, \dots, I\}$ from selling to customer type $j \in \{1, \dots, J\}$.⁹ Both I and J are finite numbers but can be large. Suppose u_i represents traits of salesperson i and v_j represents traits of customer type j . Both u_i and v_j can include numerical features (e.g., age, income), categorical features (e.g., gender, education, occupation), and features extracted from unstructured data (e.g., facial images). The latent sales generation process follows:

$$y_{ij}^* = f(u_i, v_j; \theta). \tag{1.1}$$

The $f(\cdot|\theta)$ function maps salesperson i 's traits u_i and customer type j 's traits v_j to the latent sales y_{ij}^* associated with the salesperson-customer dyad ij . $f(\cdot|\theta)$ is parameterized by θ and can take any functional form. As such, $f(\cdot|\theta)$ can flexibly capture the effect of salesperson-customer match, which has been shown to influence

⁹Customer type is a generic reference in our model. It can be as granular as a specific customer, or coarsened to describe a group of customers who share certain traits. We will discuss the operationalization of customer type in subsequent sections. We call the combination of a salesperson and a customer type a salesperson-customer dyad for brevity.

sales outcome.

Next, let y_{ij} denote the observed sales revenue salesperson i generated from selling to customer type j . We specify the relationship between latent and observed sales as:

$$y_{ij} = \begin{cases} y_{ij}^* & \text{if } y_{ij}^* > c \\ 0 & \text{otherwise} \end{cases} . \quad (1.2)$$

In other words, we posit that if the latent sales y_{ij}^* exceeds a cutoff c , it becomes an observed sales record. Otherwise, the observed sales record equals zero, meaning that salesperson i has never made a sale to customer type j .

To interpret this data generation process, imagine a salesperson who decides whether to make an effort to sell to a customer type. The cost of selling includes the time cost, effort cost, and opportunity cost. The benefit of making a sale includes commissions from the latent sales revenue (which in our setting is prespecified and known to the salesperson) and any psychological payoff ([163]). This benefit is further scaled by the perceived probability of sale and adjusted for risk preferences ([144]), time preferences ([117]), and prediction errors ([131]). As such, the key parameter $c \in R$ represents the adjusted, scaled, net cost of selling, referred to as “net cost” hereafter for brevity. The salesperson will choose not to sell to the customer type if the benefit is below cost or, equivalently, if the latent sales revenue y_{ij}^* is below net cost c . This creates a missing-by-choice sales record for the salesperson-customer dyad ij .

Analogous arguments hold if a salesperson tries to convert a customer type based on prior cost-benefit analysis but fails because of higher-than-expected selling cost or worse-than-expected conversion probability. The same intuition applies – the lack of sales records for this salesperson-customer dyad likely implies that the associated latent sales revenue is below (realized) net cost.

By modeling missing-by-choice sales records, we uncover and preserve the information value of these seemingly missing observations. To the extent that the match of salesperson-customer traits matters, salesperson i 's lack of sales records with customer type j means that, other things being equal, new salespeople who are similar

to salesperson i should expect lower latent sales revenue from serving customer type j . The augmented recommender system will recognize this and will be less likely to recommend customer type j to these new salespeople.

We have three comments. First, our model resembles the classic tobit model of truncated observations ([165]).¹⁰ Our model also echos the well-known discrete-continuous models in marketing (e.g., [38, 180]) in that discrete events (e.g., purchase incidence, decision to leave a rating, failure to serve a customer) and continuous quantities (e.g., purchase volume, rating value, sales revenue) can jointly reveal the same underlying decision process. We are humbled to build on these established theories to improve the design of neural network structures for better performance.

Second, we do not explicitly model sales effort. However, under the common assumption that salespeople in equilibrium optimize their effort choices given their own and customers' traits, the mapping function $f(\cdot|\theta)$, which has no functional form restrictions, can be seen as already subsuming the effect of effort ([41, 109, 129]).

Third, we intentionally assume a homogeneous net cost c across salespeople and customer types for the main model. This allows us to keep the model as simple as possible to isolate the mere effect of considering missing-by-choice sales records, as opposed to the effect of adding more parameters. We extend the analysis to accommodate heterogeneous net costs across salespeople in Section 1.6.2.

1.3.2 Neural Network Framework

We now incorporate the model of missing-by-choice sales into a neural network framework. We keep our presentation intuitive and refer interested readers to [83] for an excellent text on the foundation and applications of deep learning.

A neural network is based on a collection of connected units called neurons. A neuron can process "signals" passed into it and then pass the output signals into its connected neurons. A signal is a real number. An activation function computes the

¹⁰The difference is that the tobit model captures mass points in otherwise-continuous dependent variables, whereas our observed sales function is allowed to be discontinuous at the mass point of zero sales.

output of a neuron based on its inputs. The connections are called edges. Neurons and edges have weights that adjust as learning proceeds. Typically, neurons are aggregated into layers. Figure B-1 of the Appendix illustrates a typical neural network structure where data pass from the input layer to the output layer, possibly after travelling through a series of middle hidden layers.

We propose a new activation function to be used in the output layer. The idea is to utilize the hidden layers of the neural network to flexibly capture the latent sales generation process of Equation (1.1) and then adjust the output sales by imitating the observed sales generation process of Equation (1.2). The new activation function is:

$$g(x) = \begin{cases} x & \text{if } x > c \\ 0 & \text{otherwise} \end{cases} . \quad (1.3)$$

Recall that c is the net cost that governs whether latent sales records are observed. Computationally, c is a hyperparameter to be tuned during neural network optimization ([111]). Note that when c equals zero, our proposed activation function is the same as the well-known Rectified Linear Unit (ReLU) activation function: $g(x) = \max\{0, x\}$.¹¹ In Figure B-2, we offer an illustration of our activation function when c differs from zero.

Figure B-3 summarizes the complete neural network structure underlying our augmented recommender system. The main augmentation is that we use our proposed activation function in the output layer. Other layers' structures remain standard for a regression problem using neural network models ([14]).

Specifically, we first embed each categorical or unstructured feature of either salespeople or customer types and then flatten the embedding results. Embedding is a method used to represent data as a vector of continuous numbers, whereby the continuous vector can meaningfully represent the data in the transformed space ([46]). The unstructured data we use in this paper are salespeople's facial image data. Specif-

¹¹When $c \neq 0$, our activation function is different from the ReLU activation function where the input is renormalized to $\hat{x} = x - c$. ReLU activation is a continuous activation function, whereas our proposed activation function is discontinuous when $c \neq 0$. The error back-propagation learning algorithm is suitable when the activation function is discontinuous ([68]).

ically, a face embedding is a vector of continuous numbers that represents features extracted from the facial image. We then concatenate all the embedding vectors and pass the concatenating results into a multilayer perceptron (MLP). An MLP is a class of feedforward neural networks (ANNs) and a universal function approximator per Cybenko’s theorem ([48]). The goal of the MLP in our framework is to capture the complex function $f(\cdot|\theta)$ that maps salesperson and customer traits into latent sales. The MLP we consider has two hidden layers with ReLU activation at each layer. We also use Dropout at each layer to prevent neural networks from overfitting ([155]). Last, we pass the results from the MLP to the final output layer with our proposed activation function as specified in Equation (1.3). More details on the neural network structure can be found in Appendix C.1.

Three comments are in order. First, the neural network structures before the output layer are not fixed and can be adapted easily based on the type of input data. For example, if voice data of salespeople and/or customers become available, long short-term memory recurrent neural networks (LSTM RNNs, [86]) can be integrated with our proposed activation function in the output layer.

Second, the way we model missing-by-choice sales for the neural network serves as an example of using domain theories to improve deep learning algorithms. At a fundamental level, our approach shares the same essence of the invention of Dropout ([155]), Dual Rectified Linear Units (DReLU, [79]), and Attention in Bidirectional Encoder Representations from Transformers (BERT, [54]). The difference is that our technical modification is motivated by a microfounded model of how agents behave and how their behaviors shape the observed data.

Third, in our proposed framework, sales records are missing because a salesperson chooses not to sell to a customer type or fails to do so. As previously mentioned, the assumption is that the salesperson has considered all customer types. We address this remaining issue in three ways. First, in our main analysis in Section 1.5, we develop the recommender system focusing on experienced salespeople, who have arguably considered all customer types. Second, we intentionally classify customers into a reasonable number of types, such that a salesperson can feasibly consider all types.

Finally, in extended analysis of Section 1.6.1, we consider and model the probability that a salesperson has not considered all customer types.

1.4 Data

In this section, we present the data we will use to develop and evaluate the augmented recommender system. We first introduce the business context. We then describe the data in detail. Last, we present data patterns that suggest that sales records may be missing by choice and that salesperson-customer match may influence sales outcome.

1.4.1 Business Context

We collaborate with a major insurance company headquartered in Shanghai, China. Similar to its counterpart in the U.S., the sales profession in China is considered to be open to people of various backgrounds, while offering potentially good monetary reward. The average yearly gross salary for salespeople is around 20,000 USD,¹² while the per capita disposable income, as of 2020, is around 5,000 USD in China.¹³

The company focuses on providing life insurance, especially critical-illness insurance products, targeting the low- and mid-tier markets. As of 2021, the company has had hundreds of thousands of employees, established thousands of branches in the country, and served over a million customers (exact numbers withheld given the confidentiality agreement). Salespeople in the company are responsible for prospecting customers themselves, often in face-to-face settings.

We interviewed the company’s management and sales force as background research. The productivity of new salespeople emerged as a major challenge. Figure B-9 presents the proportion of new salespeople who have not made their first sales among all 123,836 salespeople who joined the company between January 2015 and November 2020. Three months into their job, 70% of the new salespeople still have not started their first sale. This has a rather negative impact on sales force morale,

¹²Source: <https://www.eri.com/salary/job/inside-salesperson/china>.

¹³Source: http://www.stats.gov.cn/english/PressRelease/202101/t20210119_1812523.html.

efficacy, retention, and recruitment.

According to our interviews, the match between salespeople and customers plays an important role in driving sales outcome, whereas new salespeople generally do not know what types of customers are suitable prospects for them. In comparison, how to explain insurance products to customers does not seem to be a challenge. Sales force compensation includes a base salary and a fixed percentage of sales revenue. The compensation structure is practically difficult to change ([49] make similar observations). In addition to managerial constraints, the company is supervised by the China Insurance Regulatory Commission and its compensation structure is strictly regulated. Therefore, we focus on non-monetary sales force management strategies and develop a recommender system to help new salespeople recognize suitable prospects.

1.4.2 Sample, Variables, and Summary Statistics

The company provided all sales records from January 2015 through February 2021, in the form of unique contracts. Each contract documents the associated salesperson’s collectable and storable information, which includes age (when the salesperson joined the company), gender, education, home province, branch served, title in the company, years worked at the company, whether the salesperson was referred to join the company, and whether the salesperson had left the company by February 2021.

We also have access to facial image data that salespeople in our sample consented to share and use. The social psychological literature has shown that facial appearance affects perceived trustworthiness (e.g., [29, 154, 159]).¹⁴ Meanwhile, trust is known to shape sales outcome (e.g., [75, 168, 169]). Therefore, We include facial image data as a potentially useful feature in our recommender system.¹⁵

Besides salesperson information, each contract documents the customer’s col-

¹⁴This effect can be consequential in a broad range of domains, such as lending ([57]), CEO selection ([160]), science communication ([78]), career development ([121]), and targeting effectiveness ([164]).

¹⁵In doing so, we must eliminate salespeople without image data from our model’s training, validation, and testing. These salespeople account for about 1/3 of the entire sample. We acknowledge potential self-selection into image-data provision and caution against using our recommender system without adjustment for salespeople who do not provide image data.

lectable and storable information, which includes age (at the time of transaction), gender, marital status, occupation, and relationship with the insured. We use these variables to characterize a customer type. Finally, each contract documents the associated sales revenue, calculated as a standard premium of each transaction, following industry norm. This serves as the salesperson’s key performance index.

As discussed in Section 1.3.2, we develop our main recommender system using data of experienced salespeople because they likely have considered all customer types. Thus, we restrict attention to salespeople who joined the company between January 2015 and July 2020, and who have completed at least five sales. These criteria are suggested by top management based on knowledge of what makes an experienced salesperson in the company. This yields 12,149 salespeople and 409,840 insurance sales records to construct and evaluate our recommender system. Repeat purchase is rare for the company. We assume each sale is for a different customer for tractability. On average, each salesperson in the data has achieved 34 sales.

Table A.13 presents the summary statistics of structured salesperson and customer traits in the data.¹⁶ There are several observations to note. The average age of salespeople when joining the company is 38. This is likely an age with much family responsibilities, making job success particularly helpful. Meanwhile, most of the salespeople have not attended college. This echos the general observation of the sales profession being accessible. In addition, the most frequent customer occupation is farmer, consistent with the company’s focus on serving low- to mid-tier markets.

As Table A.13 suggests, using raw data on customer traits yields an enormous number of customer types. Including too many customer types in the recommender system can be problematic. First, as discussed in Section 1.3.2, it is implausible that a salesperson, even an experienced one, has considered this many customer types. Second, it is impractical to offer highly granular recommendations (e.g., a 45-year-

¹⁶For the unstructured data in our sample (salesperson facial image data), one approach is to derive hand-crafted features (e.g., facial width) and then input them into the recommender system. We take a different approach; we input the raw image data to exploit the joint end-to-end representation learning advantage of deep learning based recommender systems ([185]). Therefore, we do not present summary statistics of the image data. Following confidentiality agreement, we will not report summary statistics of sales revenue either.

old female lawyer) to new salespeople. Therefore, we dichotomize customer age as weakly above or below 40, marital status as married or not, and relationship with the insured as self-insured or not. Finally, we follow the Chinese Occupation Classification and group customer occupations to seven types.¹⁷ This yields 112 customer types in total. Figure B-10 in the Appendix reports the distribution of customer types for each customer trait. The recommender system will output one or multiple, depending on system design, of these 112 customer types for each new salesperson. An example would be: a female, married customer, who is above 40, works as a technical staff, and is interested in insuring herself.

1.4.3 Preliminary Data Patterns

We examine the data for distributional insight across customer types. Figure B-11 presents a histogram of the number of customer types each salesperson in our data has sold to. Most of the salespeople served a small proportion (less than 15%) of the 112 potential customer types. In the most common case, a salesperson served less than ten customer types. Recall that these are experienced salespeople who arguably have considered all customer types. As such, there are possibly certain customer types that certain salespeople either chose not to serve or failed to serve. These would lead to missing-by-choice sales records in our data.

We also look at the customer type that brings the highest sales revenue for each salesperson in the data. Across the 12,149 salespeople, there is hardly one customer type that brings the most sales to all salespeople; rather, these “most valuable” customer types are dispersed, spanning 106 out of the 112 customer types in the data. Furthermore, for each of the 112 customer types in the data, we calculate the percentage of salespeople who have made sales to this customer type. As Figure B-12

¹⁷The seven occupation types are: 1) managers who work in government, business, and nonprofit organizations; 2) technical staff (e.g., doctors, lawyers, actors/actresses, researchers); 3) office staff and public service providers (e.g., administrative staff, administrative arbitrators, policemen, firemen); 4) service personnel (e.g., waiters/waitresses, salespeople, babysitters, doormen); 5) farmers, including personnel working in agriculture, forestry, animal husbandry, and fishery; 6) production workers in various sectors (e.g., furniture manufacturing, textile, coal chemical production, drug manufacturing); 7) soldiers and others.

shows, none of the customer types can attract all salespeople. There is also noticeable heterogeneity in the percentage of salespeople each customer type tends to attract. These results suggest that there is no customer type that suits all salespeople. Therefore, a personalized recommender system may add value, compared to a simpler recommendation to target a specific customer type.

1.5 Evaluating the Augmented Recommender System

We evaluate our augmented recommender system in this section. In particular, we test whether incorporating missing-by-choice sales records improves deep learning based recommender system performance. We focus on two commonly used performance criteria for recommender systems: prediction accuracy and recommendation quality. In addition, we present a set of interpretable outputs of our recommender system.

1.5.1 Prediction Accuracy

For prediction accuracy, we use the standard metric of mean squared error (MSE):

$$\text{MSE} = \frac{1}{IJ} \sum_{i=1}^I \sum_{j=1}^J (y_{ij} - \hat{y}_{ij})^2. \quad (1.4)$$

The elements of this equation are defined in Section 1.3. Each data point represents an interaction between a salesperson (indexed by i) and a customer type (indexed by j), including dyads where the salesperson did not sell to the customer type.

Two comments are in order. First, we operationalize y_{ij} using the average daily sales revenue associated with salesperson i and customer type j during the time span of our data. Salespeople vary in how often they sell to a customer type and how much revenue they derive from each sale. Taking the average of sales revenue over time captures the overall value of each customer type to each salesperson that encompasses factors such as the density of this customer type in the population, the

ease of conversion, and the gain from a conversion.

Second, we focus on observed sales (including zero), as opposed to latent sales, as the outcome measure. The observed sales revenue y_{ij} comes from the historical transaction data by definition. The predicted sales revenue \hat{y}_{ij} is generated by the recommender system. In our augmented recommender system, \hat{y}_{ij} predicts what sales revenue salesperson i expects to derive from customer type j , allowing for the possibility that salesperson i does not sell to customer type j . For a new salesperson, our system will recommend the customer type(s) that will yield the highest predicted sales revenue. This allows the recommendation to convey the overall potential of each customer type, after considering the net cost of selling.

We have a total of 1,360,688 salesperson-customer dyads (data points) for the $I = 12,149$ salespeople and $J = 112$ customer types. We randomly split these data into training data (60%), validation data (20%), and test data (20%). We use the training data to train the recommender system, the validation data for cross-validation tuning, and the hold-out test data to evaluate the recommender system’s performance. Out of these salesperson-customer dyads, 1,239,345, or 91%, have no sales records.

We consider two benchmark recommender systems. For fair comparison, both benchmarks are deep learning based recommender systems built on the same neural network structure as our augmented recommender system. The first benchmark, which we call “Deep Learning — Missing Excluded,” uses only observed sales records for training. This benchmark serves to test the information value of missing-by-choice sales. We scale the sample size of the benchmark to be the same as our method via random resampling. This ensures that our method’s performance improvement, if any, occurs not simply because it includes more data points.

The second benchmark, “Deep Learning — Linear Activation,” uses the same training data as our method. All the neural network components before the output layer are also the same. The only difference is that, in the output layer, this benchmark uses the standard linear activation function, as opposed to the activation function we propose in Equation 1.3. This benchmark serves to test whether our new

activation function has value in capturing missing-by-choice sales records.¹⁸

Table A.2 presents the MSE on the test data for the two benchmarks and for our proposed augmented recommender system, referred to as “Deep Learning — Augmented.” We normalize the MSE of the augmented recommender system to be 100, following our confidentiality agreement with the company. For this and subsequent metrics, we use block bootstrapping ([97]) to derive standard errors and test the significance of differences between a benchmark and our augmented recommender system.

Column (1) of Table A.2 presents the MSEs using all the test data. Our method shows much greater prediction accuracy, in both significance and magnitude, than the benchmark that excludes missing sales records. This suggests that missing-by-choice sales are highly informative and should be included in recommender system training. Linear activation improves prediction accuracy over the first benchmark, but our method remains significantly more accurate at the $p < 0.01$ level. This shows that our proposed activation function itself has value in increasing prediction accuracy.

Key to our augmentation is our treatment of missing-by-choice sales records. In the spirit of a falsification test, we repeat our MSE calculation only on salesperson-customer dyads in the test data that have positive sales records. This subsample mimics an environment in which each salesperson gets to serve all customer types. Column (2) of Table A.2 shows the results. In this case, the three methods have similar magnitudes of MSEs. Moreover, the difference between our method and the linear activation benchmark is statistically insignificant with $p = 0.31$. Therefore, as expected, our augmentation is more valuable in markets where missing-by-choice sales are common.

¹⁸We do not explicitly examine the benchmark of a deep learning based recommender system using the ReLU activation. This is because, as discussed, our method nests the ReLU activation as a special case where the net cost c is zero. Whether c is zero is an empirical question, whereas our method offers c , being it zero or not, a behavioral interpretation.

1.5.2 Recommendation Quality

We next evaluate the recommendation quality of our augmented recommender system. We consider two common evaluation metrics in the literature (e.g., [157, 151, 181]): F1-score and Normalized Discounted Cumulative Gain (NDCG). Intuitively speaking, the F1-score measures whether a recommender system is good at identifying items (e.g., customer types), whereas NDCG measures whether a recommender system is good at identifying and ranking items.

The F1-score depends on both “precision” and “recall.” Precision is defined as the number of “relevant” recommendations divided by the total number of recommendations. In our setting, we define a recommendation as relevant if the salesperson (in the test data) has indeed had a transaction with the recommended customer type. Recall is defined as the the number of relevant recommendations divided by the total number of relevant customer types. The F1-score balances precision and recall to achieve a single-dimension comparison between different recommender systems:

$$\text{F1} = 2 \times \frac{\text{Precision} \times \text{Recall}}{\text{Precision} + \text{Recall}}$$

NDCG further takes into account the ranking among recommended items. The idea is that highly relevant customer types should ideally be highly ranked in the list of recommended types. We quantify the relevance score by using the actual daily sales revenue for each salesperson on each customer type. To calculate NDCG, we divide Discounted Cumulative Gain (DCG) by Ideal Discounted Cumulative Gain (IDCG). For each salesperson, the DCG measure can be calculated from the formula, $\text{DCG} = \sum_r \text{Relevance}_r / \log_2(r + 1)$, where Relevance_r is the relevance score between the salesperson and the customer type recommended in ranking position r of the list, and the logarithmic function is used to scale the importance of ranking. By arranging customer types in descending order of actual daily sales revenue, we can derive IDCG using the same formula as DCG.

We again compare our method with the two benchmark recommender systems: Deep Learning — Missing Excluded and Deep Learning — Linear Activation. In ad-

dition, we consider a non-personalized recommender system that simply recommends the top customer types based on their associated historical sales revenue. Both the F1-score and NDCG depend on the number of recommendations a system is designed to offer. We vary this number from one to three. In addition, both the F1-score and NDCG are salesperson-specific. We first compute their values for each salesperson in the test data and then take the average across these salespeople.

Table A.3 presents the F1-score and NDCG results. For both metrics and for all numbers of recommendations considered, excluding missing sales records from training leads to the worst recommendation quality. Linear activation improves recommendation quality and our method achieves further, significant improvement ($p < 0.01$). Our method also significantly outperforms the non-personalized recommender system. This again shows that salesperson-customer match matters for sales outcome.

The goal of our recommender system is to leverage experienced salespeople’s sales records (and the lack of them) to help new salespeople. So far, we have evaluated our average recommendation quality for all salespeople in the test data. It remains to check our recommendation quality for new salespeople. In Table A.4, we present the F1-score and NDCG for salespeople with different numbers of years working at the company, fixing the number of recommendations at one. Table A.5 in the Appendix shows the results with two and three recommendations, respectively. Reassuringly, our method continues to report significantly higher recommendation quality than the benchmarks across the board, including the case of relatively inexperienced sales people.

1.5.3 Recommender System Output Interpretation

Although we focus on prediction accuracy and recommendation quality, we report the results of our recommender system in this section to facilitate interpretation.

A natural question is which salesperson and customer traits are important. We use the permutation feature importance technique (e.g., [5]) to answer the question. This technique is especially useful when the model is “opaque,” which is often the case for neural network models. Specifically, we take a fitted model and randomly

shuffle each variable. We then look at the decrease in the performance measure being considered. Shuffling variables of higher importance causes greater decreases in the performance measure. In Table A.6, we present the decrease in MSE (standardized), F1-score, and NDCG, where we implement 100 times of permutation for each variable.

Several results are worth noting. First, whether a trait is important depends on the performance measure. For example, “Title in Company” is the most important salesperson trait when we consider recommendation quality and is not when we consider prediction accuracy. This result highlights the importance of setting the right performance metric for the recommender system. Second, the facial image data do have noticeable impact on recommender system performance across all measures, which underscores the value of using a deep learning based recommender system. Third, across all measures, salespeople’s age, home province, branch served, and whether they left the company are important traits, while their gender, education, and whether they were referred to join the company are less important. Last, customer age, marital status, and occupation particularly matter for sales outcome. These results shed light on ways to further improve the recommender system in future research. For instance, it may be worthwhile to categorize customers into finer types along more important traits.

For a recommender system, the ultimate question is what customer type is recommended. The answer depends on the specific salesperson. Note that the results above show the relative importance of each trait in our deep learning model, which subsumes the complex interactions among all traits.¹⁹ To visualize the importance of salesperson-customer match, we present the customer type our model recommends for each salesperson-trait category (e.g., age below 40). Table A.7 of the Appendix displays the results. To simplify presentation, we examine salesperson traits in isolation, two categories per trait, without including their interactions. For facial image, we split the salespeople into two groups based on clustering of their facial embedding vector values. Even with this coarse classification of salespeople, we can see that

¹⁹The importance of each trait is analogous to the partial derivative of the outcome measure with respect to this trait, which subsumes any interaction effect between this trait and others.

match matters. Out of the 10 salesperson traits, only one (i.e., whether referred to join the company) does not affect our model’s customer type recommendation. These results again highlight the value of a personalized recommender system.

Finally, cross-validation yields a positive net cost of selling hyperparameter, c . Its magnitude is 4.46% of the average daily sales revenue across all salesperson-customer dyads (value withheld for confidentiality agreements). However, sales revenue is widely dispersed, with a standard deviation 21.78 times its mean value. Therefore, it is possible that a non-negligible fraction of sales revenue falls below the net cost of selling, causing the associated sales records to be missing by choice.

1.6 Extensions

Because our augmented recommender system is based on a simple behavioral model, it can be easily extended to capture different market scenarios. In this section, we model two such extensions and discuss a few others.

1.6.1 Probably Missing-by-Choice Sales Records

For the main analysis so far, we focus on experienced salespeople, assuming that they have had enough time to consider all customer types. In reality, it is possible that even experienced salespeople have ignored certain customer types for reasons unrelated to the value of a customer type.

We address this possibility by embedding a probabilistic model into our augmented recommender system. For each salesperson-customer dyad that has no sales records, we allow for two possibilities: that the salesperson has considered the customer type, as modeled in the main analysis, or the salesperson has not considered the customer type so that the missing sales record should be excluded from training. We parameterize the probability that salesperson i has considered a customer type as $1/(1 + e^{-\alpha * \text{Experience}_i})$, where α is a hyperparameter to be tuned through cross-validation and where experience is measured by the number of years the salesperson has worked at the company. This probability enters the neural network as the weight

for each data point.

We repeat the analysis of prediction accuracy and recommendation quality in this probabilistic framework. We set the number of recommendations to one for the F1-score and NDCG. Table A.8 presents the results. Compared with the augmented recommender system in the main analysis (main model), this extended, more flexible model significantly improves recommendation quality.

1.6.2 Heterogeneous Net Cost of Selling

In the main analysis, we assume a homogeneous net cost of selling, c , across salespeople. We now relax this assumption. To keep this extension tractable, we assume that there exist two segments of salespeople who differ in their net cost of selling.

We again embed a probabilistic model into our augmented recommender system. We parameterize the probability that salesperson i belongs in one of the two segments as a function of the salesperson’s average number of monthly transactions: $1/(1 + e^{-\beta * \text{Transactions}_i})$, where β is a hyperparameter to be tuned through cross-validation. As one would expect, the net cost of selling is likely higher for salespeople who accomplish fewer transactions each month. This probability enters training as the weight of each data point in two neural networks separately trained for the two segments.

The bottom row of Table A.8 presents the results. Considering heterogeneous net cost of selling leads to significantly better prediction accuracy and recommendation quality. This is expected given the greater flexibility of the extended model.

1.6.3 Other Possible Extensions

The main analysis and both extensions boil down to one central question – how do we interpret the incidence of sales records? Throughout the paper, we highlight the importance of offering behaviorally meaningful, microfounded interpretations of this data feature. The same idea can be applied to potentially many other extensions. For example, a company may take charge of customer prospecting and then assign different prospects to different salespeople. This is the telemarketing setting studied in

[82]. In this case, a salesperson’s sales records with a customer type may be a function of the company’s job assignment rule. Our method can again be easily extended to this scenario by weighing each salesperson-customer dyad with the possibility of encounter.

1.7 Practical Value

We return to the practical motivation of our augmented recommender system. We simulate whether our recommender system indeed helps new salespeople make their first sale and helps the company improve sales force management.

For each salesperson in the test sample, we examine two counterfactual scenarios. We assume that the salesperson approaches customers following either random search or the recommender system. Random search is a simplification but nevertheless plausible strategy if a new salesperson has no guidance, consistent with our interview with the company. With random search, the salesperson randomly chooses a customer type at each attempt to sell. With a recommender system, we assume that the salesperson first tries the recommended customer types and, if doing so does not generate a sale, goes on to approach other customers through random search. For either approach, if we observe a sale between the salesperson and the chosen customer type in the company’s historical sales records, we count the sales attempt as a success; otherwise, we count it as a failure.

Figure B-13a presents the average number of failures before the first sale for random search and our method, respectively. We vary the number of recommendations from one to three. As expected, our method outperforms random search more if it is allowed to recommend more customer types. For instance, by recommending three customer types, our method reduces the number of failures from 20 to 12, a 40% improvement, compared with random search. We present the comparison with other benchmark recommender system in the Online Appendix, where our model continues to outperform the benchmarks.

We look further at how our recommender system can help new salespeople specif-

ically. We focus on the subsample of salespeople in the test data who have worked at the company for less than one year. We redo the above analysis and Figure B-13b presents the results. Indeed, new salespeople tend to go through more failures before their first sale. Our recommender system can help them significantly expedite their first success.

Last, we simulate the aggregate implication of expediting sales for the company. Recall Figure B-9, which presents the proportion of new salespeople yet to make their first sale as time goes by. We again assume that new salespeople follow random search when not aided by a recommender system. In addition, we perform a back-of-envelope calculation of new salespeople's pace of sales attempts, so that the fraction of salespeople who have not achieved the first sale over time fits the actual data in Figure B-9.²⁰ We then use the calculated pace of sales attempts to simulate the number of salespeople yet to make their first sale over time when a recommender system is in place.

We present the full results for different recommender systems and different numbers of recommendations in the Appendix. Figure B-14 displays the result of using our augmented recommender system to recommend only one customer type. Even in this most conservative case, our method greatly improves sales force productivity, reducing the number of salespeople with no sales by 27% 90 days into their job. Expediting the first sale can accelerate revenue and reduce attrition. It may also generate positive ripple effects through better customer satisfaction and sales force morale.

²⁰The number of days a salesperson takes to make the first sale equals the number of failures before the first sale multiplied by the number of days a salesperson takes to make a sales attempt (the pace). We allow the pace of making a sales attempt to be time-varying. There are six 15-day intervals in Figure B-9. Thus we calculate six corresponding paces based on the proportion of salespeople who have not made a sale by that time. The results are 2.83, 4.76, 6.16, 7.69, 9.04, and 10.47 days per sales attempt for these six intervals, respectively. The pace slows down possibly due to frustration.

1.8 Concluding Remarks

In this paper, we develop a deep learning based recommender system to help new salespeople recognize customers with high conversion potential. We augment standard deep learning based recommender systems to capture an important feature of the sales environment, that experienced salespeople’s failures to convert certain customer types is as informative as their success. We develop a parsimonious behavioral model to capture what we refer to as “missing by choice” sales records in company databases. We then incorporate the model into a neural network structure by proposing a new activation function. Our augmented recommender system outperforms common benchmarks in prediction accuracy and recommendation quality. Simulation suggests that our method can markedly improve sales prospecting efficiency and sales force productivity.

The contributions of this paper are two-fold. Substantively, we develop a data-driven, automated system to improve performance in the sales function, which is known for its intense pressure and high stakes. Methodologically, we develop a way to augment already-powerful deep learning based recommender systems to handle missing-by-choice observations. Our augmentation performs well, and is simple, interpretable, and easily extendable. We showcase its efficacy in the sales context, but we expect the framework to be broadly applicable. Again, take the classic example of customer reviews. Consumers may choose not to review a product on a website if they feel they have no new opinions to add ([26]). This will lead to missing-by-choice reviews. An augmented deep learning based recommender system can be applied to retrieve the opinions these consumers may have had, considering the content of existing reviews (e.g., [166]) or quality signals from existing product images ([186]).

There are many avenues for future research. First, it will be useful to evaluate the impact of a recommended system like ours on sales performance in the field. Various factors need to be unpacked, such as the mere effect of introducing artificial intelligence into a traditional industry, which can be complex in itself. Second, by recommending similar customer types to similar salespeople, the recommender system

may intensify competition. Competition is not a salient concern in our setting because the potential market of customers is enormous and widely dispersed. It is nevertheless an interesting topic to explore. Last, as a proof of concept, our recommender system abstracts away from many rich features of the sales environment, such as salesperson learning and team dynamics ([115]). Incorporating these features can unlock valuable opportunities to further improve recommender system design.

Chapter 2

Training Scalable Personalization Policies with Constraints

Personalization, which aims to target different marketing actions to different customers, has attracted broad attention in both academia and industry. While most research has focused on training personalization policies without constraints, in practice, many firms face constraints when implementing these policies. For example, firms may face volume constraints on the maximum or minimum number of actions they can take, or on the minimum acceptable outcomes for different customer segments. They may also face fairness constraints, that require similar actions with different groups of customers. These constraints can introduce difficult optimization challenges, particularly when the firm intends to implement personalization policies at scale. Traditional optimization methods face challenges solving large-scale problems that contain either many customers or many constraints. We show how recent advances in linear programming can be adapted to the personalization of marketing actions. We provide a new theoretical guarantee comparing how the proposed method scales compared to state-of-the-art benchmarks (primal simplex, dual simplex and barrier methods). We also extend existing guarantees on optimality and computation speed, by adapting them to accommodate the characteristics of personalization problems. We implement the proposed method, and compare it with these benchmark methods on feasibility, computation speed, and profit. We conclude that,

volume and similarity (fairness) constraints should not prevent firms from optimizing and implementing personalization policies at scale.¹

2.1 Introduction

Big data and new methods allow firms to personalize their marketing actions for different customers ([37]). The growth in industry interest in personalization has been mirrored by rapid growth in academic interest in personalization (targeting). This academic attention has primarily focused on training personalization policies without constraints. However, in practice, internal business rules or society considerations often impose constraints on firms' targeting policies. In this paper, we investigate how to train scalable targeting policies in the presence of constraints.

Constraints on targeting policies are typically of two types. *Volume* constraints limit the total number of marketing actions that can be taken. This type of constraint may result from capacity constraints. For example, a firm's ability to make outbound phone calls may be limited by the availability of trained associates to make these calls. Budget constraints may also impose minimum and/or maximum limits on the total number of marketing actions. These constraints may operate in aggregate, or they may apply to specific customer segments. *Similarity* constraints limit differences in marketing actions taken with different customer segments. These constraints are often motivated by concerns for fairness. For example, a constraint might require that the firm takes similar marketing actions with neighboring zip codes, or that customers located near one store are treated with similar marketing actions as customers located near other stores.

While there are many standard methods for optimizing problems with constraints, large numbers of decision variables and large numbers of constraints can both make the problem challenging. In a personalization problem, the number of decision variables and the number of constraints can both potentially be large. For example, if a

¹This chapter is collaborated with Haihao Lu (Assistant Professor of Operations Management at Chicago Booth).

separate decision is made for each customer, the number of decision variables may be in the millions. Even where decisions are made at the customer segment level, if there are a large number of segments, there will be a large number of decision variables. Similarly, if constraints apply to specific segments, the number of constraints will be at least as large as the number of segments, and may grow as a polynomial function of the number of segments. Standard optimization methods are not well-suited to solving optimization problems with either large numbers of decision variables, or large numbers of constraints ([16]).

We formulate the personalization problem as a linear programming problem with constraints, and illustrate how to incorporate both volume constraints and similarity constraints. We then adapt and apply a recently developed algorithm ([10, 11]), which is designed to solve large-scale linear programming problems. The algorithm belongs within the class of first-order methods, which use gradient information to construct algorithms to find optimal solutions. This class of methods scales very well, and is widely used in many applications, including many machine learning algorithms ([15]). The algorithm that we adapt leverages the primal-dual hybrid gradient ([27]). Similar methods are widely used in image processing and computer vision (e.g., [66, 90, 141, 188]). Recent developments have made the algorithm especially suitable for large-scale linear programming problems (see details in Section 2.4).

We provide two theoretical results. The first result compares the proposed algorithm with state-of-the-art benchmarks: primal simplex, dual simplex and barrier methods. We prove that the algorithm can solve larger problems (in terms of customers and constraints) than any of these methods. The second theoretical result extends existing guarantees on optimality and computation speed, by adjusting the method (and existing theory) to accommodate the characteristics of personalization problems.

To illustrate the practical value of the method, we apply it to an actual personalization problem. The problem involves choosing which promotions to send to prospective customers. The response functions are estimated using a large-scale field experiment that includes approximately 2.4 million households and five marketing actions. The

findings provide practical empirical evidence of how the proposed method extends the scale of solvable personalization problems in the presence of constraints. Our analysis recognizes that different firms that have access to different hardware resources. As we change the available hardware, our method consistently solves problems faster, yields higher profits, and accommodates both a larger number of customers and a larger number of constraints.

The paper continues in Section 2.2, where we position our contribution with respect to the existing literature. We explicitly model the firm’s optimization problem in Section 2.3, and discuss how to incorporate different types of volume and similarity constraints. Section 2.4 presents details of the algorithm together with theoretical guarantees. We describe the data and empirical analyses in Section 2.5. Section 3.5 concludes and highlights promising future research directions.

2.2 Related Literature

Research on methods for training personalization policies using machine learning methods has grown rapidly growing in recent years. In Table A.10, we summarize seven recent marketing papers investigating methods for personalizing a range of different marketing actions. The focus in all of these papers is on estimation of response functions, rather than the optimization of marketing actions with constraints. While some papers consider budget constraints, these constraints are simple in nature, and can be solved using simple greedy algorithms. The inclusion of more complicated constraints, and an increase in either the number of customers or the number of constraints will require more sophisticated optimization algorithms. We propose a method that can solve personalization problems that include both many customers, and many constraints.

Beyond personalization, other research in marketing has investigated how to include constraints when optimizing marketing actions. Examples include research studying conjoint analysis ([167]), product line design ([119]), retail assortment ([69]), and content arrangements on social media ([103]).

There is a long history of studying the role of fairness in firm’s marketing actions. Moreover, the fairness of firm’s pricing decisions have probably received more attention than the fairness of other marketing actions (see for example [4, 7, 19, 24, 177]). The use of algorithms to make marketing decisions has renewed interest in fairness as a research topic. A distinguishing feature of algorithm fairness is that the unfair actions are often an unintended outcome. Algorithms that are designed to optimize seemingly innocuous goals, can lead to unintended and unanticipated differences in how customers are treated. For example, [110] documents how an algorithm delivering advertisements promoting job opportunities led to unintended discrimination. Although the advertisement was designed to be gender neutral, the algorithm optimized cost effectiveness, which led to the ad being shown to more men than women. They conclude by showing that this result generalizes across different digital ad platforms. Similarly, [184] show that the introduction of smart-pricing algorithms increased the gap between revenue earned by black and white hosts on Airbnb.

There have been several theoretical studies investigating the implications of fairness on firm’s strategies (e.g., [47, 71, 87, 113]). However, research on methods to anticipate or mitigate algorithmic fairness concerns in marketing are only now beginning to emerge. One notable example is [12], who propose using bias-eliminating adapted trees to adjust the potential bias in personalization policies. We contribute to this emerging literature by illustrating how to incorporate fairness concerns, and how to optimize these problems when the number of customers or constraints is large.

Algorithm fairness has also generated considerable interest in the computer science and machine learning literatures. This includes an ongoing debate about the definition of fairness. The definitions are many and varied (see for example [25], and [126]). It is beyond the scope of this paper to resolve the differences between these definitions. Instead, we will interpret a fairness constraint as an example of a *Similarity* constraint (see the discussion in Section 2.3). The computer science literature has also proposed different methods to identify and restrict discrimination in policies trained using algorithms. The methods can be classified into three types: pre-process, in-process, and post-process (see review papers such as [140]). Our method belongs

to the in-process class of methods, because we explicitly consider fairness as part of the policy training process.

The current state-of-the-art methods for solving linear programming problems are the simplex method ([50]) and the barrier method ([135]). Both methods are very mature, and are implemented in commercial software for solving LP problems, such as Gurobi, cplex, etc. They are in general capable to provide reliable solutions for medium size problems. However, neither method is well-suited to solving personalization problems that have a large number of customers or a large number of constraints ([16]). The method that we extend and adapt to solving personalization problems is based upon Primal-Dual Hybrid Gradient for Linear Programming (PDLP). Compared with the simplex and barrier methods, the proposed method only require matrix-vector multiplications, which allows the method to easily scale up. In contrast, simplex and barrier methods require solving linear equations using matrix factorization, which leads to two major challenges when solving large scale instances: (1) while the original targeting problem can be highly sparse, the matrix factorization can be much denser, thus require more memory usage in general; (2) it is indeed highly challenging (if not impossible) to use the modern computing architectures, such as distributed system and GPUs, for matrix factorization. On the other hand, PDLP only requires storing the constraint matrix in memory, and sparse matrix-vector multiplication has been well scaled on modern computing architectures and thus are more suitable for larger instances. In light of such a fundamental distinction, Nesterov [134] formally defines optimization methods that require matrix factorization (or linear equation solving) as handling medium-scale problems, and optimization methods that require matrix-vector multiplication as handling large-scale problems.

The method itself is very new, with the theoretical foundation for the algorithm described in [10]. This paper provides theoretical guarantees for both optimality and computation speed. Implementation of the algorithm at scale requires several additional steps, which are described in [11]). The research teams that authored these papers include highly skilled data engineers at Google, who are working to compile

the algorithm within the Google OR-Tools suite.² Our contribution is to provide the first documented empirical application of the algorithm. In doing so, we show how to extend and adapt the method to solve personalization problems with constraints. Our theoretical findings adapt the guarantees on convergence and optimality to accommodate the characteristics of these problems. Moreover, we provide a new theoretical result, that provides guarantees comparing the performance of the method with simplex and barrier methods.

Personalization and targeting have also received recent interest outside the marketing literature, in both operations research and management. [80] consider the problem of personalizing product assortments in a dynamic setting, and propose an index-based MAB method. [32] apply statistical learning to customize revenue management policies. [53] also use linear programming to solve personalization problem, where they focus on personalized reserve prices. Notably, none of these papers consider the scalability of personalization problems in the presence of constraints.

In the next section, we discuss how to incorporate different types of constraints into personalization problems.

2.3 Constraints and Problem Setup

In this section we model personalization as a linear programming problem. Suppose there are I considered households and J available marketing actions. An example of a marketing action could be sending a direct mail advertisement offering a "free trial" or a discount. The set of marketing actions could include a "no action" or null treatment. Each household belongs to one of K customer segments, where $K \geq 1$ (if $K = 1$ then all customers belong to the same segment). Customer segments are defined using one or more observable contextual variables, such as gender, race, geographic locations, or past purchasing.³

²<https://developers.google.com/optimization/lp>.

³We detail the customer segmentation definition in our data in Section 2.5.1.

We formulate the a firm's problem as follows:

$$\begin{aligned}
\max_{x_i^j} \quad & \sum_{i=1}^I \sum_{j=1}^J r_i^j x_i^j \\
\text{s.t.} \quad & a_k^j \leq \sum_{i \in S_k} x_i^j \leq b_k^j, \quad \text{for } j = 1, \dots, J, k = 1, \dots, K \\
& L_k \leq \sum_{i \in S_k} \sum_{j=1}^J c_i^j x_i^j \leq U_k, \quad \text{for } k = 1, \dots, K \\
& \frac{1}{n_{k_1}} \sum_{i \in S_{k_1}} x_i^j \leq \lambda_j^{k_1 k_2} \frac{1}{n_{k_2}} \sum_{i \in S_{k_2}} x_i^j, \quad \text{for } j = 1, \dots, J, k_1 = 1, \dots, K, k_2 = 1, \dots, K \\
& \frac{1}{n_{k_1}} \sum_{i \in S_{k_1}} \sum_{j=1}^J d_i^j x_i^j \leq \gamma^{k_1 k_2} \frac{1}{n_{k_2}} \sum_{i \in S_{k_2}} \sum_{j=1}^J d_i^j x_i^j, \quad \text{for } k_1 = 1, \dots, K, k_2 = 1, \dots, K \\
& \sum_{j=1}^J x_i^j \leq 1, \quad \text{for } i = 1, \dots, I \\
& x_i^j \geq 0.
\end{aligned} \tag{2.1}$$

For ease of exposition, when we refer to the "personalization with constraints" problem, we refer to the setup in Equation (2.1).

The firm's objective in (2.1) is to maximize profits across all households and all marketing actions. The decision variable, $x_i^j \in [0, 1]$, represents the probability a given household i receives marketing action j . While x_i^j will normally take the value zero or one, a policy could be stochastic (rather than deterministic), so that there is a random element as to which customer i receives action j . We interpret a no action or null marketing action as a decision not to implement any of the J actions. This null action could include "business as usual".⁴

We denote the incremental profit that the firm earns from household i if it receives marketing action j as r_i^j . This is calculated as the difference between: (a) the profit earned from household i if it receives marketing action j , and (b) the profit

⁴Alternatively, we could treat the null action as a separate action in itself, but we would then adjust how we measure outcomes. As we discuss in this section, we measure outcomes as the incremental profit earned from a marketing action compared to the null action. If we treated the null action as an action in itself, we would measure profits as the profit earned from each action, rather than the difference in profit compared to the null action.

earned from household i if no action is taken. Without changing the model, we can alternatively redefine r_i^j to measure increments of revenue, units sold, or some other managerially relevant observable outcome. Predicting the incremental response to marketing actions (r_i^j) has been the primary goal of many marketing studies, including many of the recent personalization papers listed in Table A.10. In this paper, we are agnostic to the process and estimator used to estimate r_i^j . Instead, our focus is on how the firm uses these incremental responses to identify an optimal personalization policy.⁵

We recognize that, in practice, some but not all of these constraints may be relevant, and their relevance will vary in different settings. By specifying several different types of constraints, we aim to provide a menu of options for firms (and researchers) wanting to incorporate constraints into personalization policies. We discuss the motivation and interpretation for each set of constraints below.

The first set of constraints $a_k^j \leq \sum_{i \in S_k} x_i^j \leq b_k^j$ represent volume constraints on each marketing action. Here, S_k is the set of households in customer segment k , a_k^j is the lower bound for the total number of households given marketing action j in customer segment k , and b_k^j is the upper bound for the total number of customers given marketing action j in customer segment k . Volume constraints can include both a minimum number requirement or a maximum number requirement. For example, while budget constraints impose an upper limit on how many customers receive the fall Holiday catalog from Saks, agreements with suppliers (who fund the catalog) mean that there is also a minimum number of customers who can receive the catalog. Notice that the framework also allows the minimum and maximum volume constraints to vary by customer segment (k); a_k^j and b_k^j can vary across k and j .

The second set of constraints $L_k \leq \sum_{i \in S_k} \sum_{j=1}^J c_i^j x_i^j \leq U_k$ capture the volume constraint on all marketing actions. Here, L_k denotes the lower bound for a combination of all marketing actions in customer segment k , and U_k denotes the upper bound for

⁵We discuss how we predict r_i^j in our empirical application in Section 2.5.1. We do note that we stay in the predict-then-optimize paradigm; the firm firsts predicts r_i^j , and then use the predicted outcomes to generate optimized decision variables. It is interesting, but beyond the scope of our paper, to explore how to reconcile the misaligned objectives in prediction and optimization (e.g., [65]).

a combination of all marketing actions in customer segment k . The combination of the marketing actions is determined by parameter c_i^j . Here, we allow the combination ways to differ across different households. This form of volume constraint includes several common constraints discussed in literature. For example, a budget constraint often takes this form. In a budget constraint, c_i^j denotes the costs for marketing action j to household i , L_k is zero, and U_k denotes the total budget number for customer segment k . It is possible that c_i^j takes the same value across different households i . Performance constraints may also take this form. Performance constraints impose requirements on a measurable outcome of the targeting policy. For example, although a firm's objective may be to maximize profits (including the cost of the marketing actions), a manager's goals might include a requirement that this year's revenue is no lower than last year ([137]). In this case, c_i^j denotes the revenue if we give household i marketing action j , L_k is the performance requirement number for customer segment k , and U_k is equal to infinity.

The third set of constraints $\sum_{i \in S_{k_1}} x_i^j / n_{k_1} \leq \lambda_j^{k_1 k_2} \sum_{i \in S_{k_2}} x_i^j / n_{k_2}$ model the similarity constraints for each marketing action. These similarity constraints restrict the difference in the actions taken with different customer segments, and will often be motivated by concerns about fairness (see for example, [25, 126]). The total number of households in customer segment k is denoted by n_k , and $\lambda_j^{k_1 k_2}$ restricts the difference between customer segments k_1 and k_2 . The constraint specifies that the difference in the proportion of customers receiving a given marketing action between two customer segments cannot be larger than $\lambda_j^{k_1 k_2}$. For example, the proportion of male and female customers receiving a discount cannot vary by more than 5%. This restriction could be imposed within each zip code. Notice that by exchanging k_1 and k_2 , we can also include restrictions in which there is a minimum difference in the proportion of customers receiving a given marketing action. We might require that twice as many disadvantaged customers receive discounts as customers who are not disadvantaged.

The fourth set of constraints $\sum_{i \in S_{k_1}} \sum_{j=1}^J d_i^j x_i^j / n_{k_1} \leq \gamma^{k_1 k_2} \sum_{i \in S_{k_2}} \sum_{j=1}^J d_i^j x_i^j / n_{k_2}$ impose similarity constraints on all marketing actions. The difference between cus-

customer segments k_1 and k_2 is restricted by $\gamma^{k_1 k_2}$, and d_i^j is the weighting factor to determine the combination of all marketing actions.

The last two constraints $\sum_{j=1}^J x_i^j \leq 1$ and $x_i^j \geq 0$ restrict the firm's action space, so that each household has at most one marketing action, and x_i^j is non-negative. Notice that marketing actions can cause negative treatment effects ($r_i^j < 0$), but it is not possible to take the negative of a marketing action.

In our theoretical and empirical investigations, we will vary the size of I and K to investigate how our proposed method can enlarge the problems that different methods can solve. Before presenting these findings we first formally specify the optimization algorithm, and provide guarantees on feasibility, performance and computation speed.

2.4 Algorithm and Theoretical Guarantees

We begin this section by documenting the central ideas in the proposed algorithm. We then provide theoretical guarantees.

Unlike state-of-the-art solvers for linear programming, such as primal simplex, dual simplex or barrier methods, the proposed algorithm uses gradient information to construct the search process for optimal solutions. For this reason, the algorithm is a first-order method. To write down the gradient, we first simplify the expression in Equation (2.1) and write down the primal form:

$$\begin{aligned} \min_{x \in R^{I \times J}} \quad & -r^T x \\ \text{s.t.} \quad & Gx \geq h \\ & 0 \leq x \leq 1, \end{aligned} \tag{2.2}$$

where we stack all decision variables x_i^j into $x \in R^{I \times J}$ and all predicted profit r_i^j into $r \in R^{I \times J}$. All of our constraints are inequality constraints and we can organize them into matrix form $Gx \geq h$ where $G \in R^{M \times I}$ and $h \in R^M$. Based on our setup in Equation (2.1), $M = JK + JK + 2K + K(K - 1)J + K(K - 1) + I$.

Using duality theory, we can also write down the dual form:

$$\begin{aligned} \max_{y \in R^M} \quad & h^T y \\ \text{s.t.} \quad & G^T y = -r \\ & y \geq 0. \end{aligned} \tag{2.3}$$

In order to derive the first-order gradient, we write down the primal-dual form of the problem using Equation (2.2) and Equation (2.3):

$$\min_{x \in X} \max_{y \in Y} -r^T x + h^T y - y^T Gx, \tag{2.4}$$

where $X = \{x \in R^{I \times J} : 0 \leq x \leq 1\}$ and $Y = \{y \in R^M : y \geq 0\}$. Then the key iterations utilizing the gradient information are ([27]):

$$\begin{aligned} x^{new} &= \text{proj}_X(x^{old} + \eta r + \eta G^T y^{old}) \\ y^{new} &= \text{proj}_Y(y^{old} + \tau h - \tau G(2x^{new} - x^{old})). \end{aligned} \tag{2.5}$$

where $\text{proj}_X/\text{proj}_Y$ denotes the map that projects onto X/Y , $\eta > 0$ is the primal stepsize and $\tau > 0$ the dual stepsize. We choose standard default parameters for η and τ in implementation⁶.

The iterations described in Equation (2.5) are normally called primal-dual hybrid gradient (PDHG) (e.g., [27, 28]), and we will use the same label. From Equation (2.5), we can clearly see why our algorithm is scalable: the iterations are "matrix-free" in that we only require matrix-vector multiplications of the data matrix. In Theorem 1, we show that our algorithm can solve larger constrained personalization problems (larger K and larger I) than simplex and barrier methods. Before present this theorem, we first introduce a modification to Equation (2.5) that is a key to making the iterations more suitable for linear programming.

Instead of conducting a one-loop iteration PDHG, our algorithm considers a two-

⁶<https://odlgroup.github.io/odl/>.

loop iteration. Let t represent the inner loop counter and n denotes the outer loop counter. Suppose we are now at the n -th outer loop. In the inner loop, we form a sequence of $(x^{n,0}, y^{n,0}), \dots, (x^{n,t}, y^{n,t})$ using Equation (2.5). We then use this sequence to construct a quantity $\bar{z}^{n,t}$ to be used in the outer loop (the details of this transformation are in Appendix D.1). The stop of the inner loop depends on the normalized duality gap decay of $\bar{z}^{n,t}$ while t increases. In Figure B-15, we offer a simple case to illustrate why two-loop iterations can help with convergence.

The details of the algorithm are offered in Algorithm 1. Details describing the construction from $(x^{n,t}, y^{n,t})$ to $z^{n,t}$, normalized duality gap decay condition and convergence condition are provided in Appendix D.1.

Algorithm 1

Input: Initialize the outer loop: $n \leftarrow 0$; Initialize an solution $z^{0,0}$ (that is, a pair of $(x^{0,0}, y^{0,0})$);
while do
 Initialize the inner loop: $t \leftarrow 0$;
 while do
 $z^{n,t+1} \leftarrow PDHG(z^{n,t})$;
 $\bar{z}^{n,t+1} \leftarrow \sum_{i=1}^{t+1} z^{n,i} / (t + 1)$;
 $t \leftarrow t + 1$;
 end while normalized duality gap decay condition holds;
 $z^{n+1,0} \leftarrow \bar{z}^{n,t}$;
 $n \leftarrow n + 1$;
end while $z^{n,0}$ convergence.

Our first theoretical finding compares the size of the problems the algorithm can solve with state-of-the-art benchmarks. In particular, in Theorem 1, we show that the algorithm can solve larger constrained personalization problems than simplex and barrier methods (given the same hardware resources).

Proposition 1. *For a general linear program of the form (2.2), one iteration of Algorithm requires $O(nnz)$ floating point operations, and one major iteration in simplex method or barrier method requires $O(\min\{m^3, n^3\})$, in the worst scenario, where nnz is the number of non-zeros, m is the number of rows, and n is the number of columns in the constraint matrix G .*

Proof. In algorithm 1, the major cost per iteration is the two matrix vector multiplications in the PDHG update, and each matrix vector multiplication requires nnz floating point operations.

The typical implementation of the simplex method solves a linear equation with a basis matrix in about every 100 iterations. Since the basis matrix is not symmetric, the linear equation solving is usually by a (sparse) LU decomposition, which requires $O(m^3)$ or $O(n^3)$ floating point operations in the worst case, depending on whether we use primal simplex or dual simplex.

Each Newton's step in the barrier method requires to solve a symmetric linear equation. This step is usually performed by a (sparse) Cholesky decomposition, which requires $O(m^3)$ or $O(n^3)$ floating point operations in the worst case. \square

For a general linear program (2.2) with constraint matrix G , it is very likely that $nnz \ll \min\{m^3, n^3\}$. Indeed, G in targeting problem (2.1) is sparse with $nnz \approx 2Kn$ where the majority of nonzeros comes from the similarity constraints. Notice that K is usually a small constant as stated in Table A.11, a step of PDLP is usually much cheaper than a step of the barrier method or a major step in the simplex method for large instances.

In Theorem 1, we show that Algorithm 1 can achieve global linear convergence while applied on targeting with constraints problem.

Theorem 1. *Algorithm 1 can achieve global linear convergence of the problem in Equation (2.1). Specifically, Algorithm 1 requires at most $O(\log(\frac{1}{\varepsilon}))$ number of iterations to find an approximate solution to (2.4) such that the distance between this solution to an optimal solution is at most ε .*

The proof of Theorem 1 follows from [11]. Actually, the results in [11] do not apply directly to our problem (2.2), because we have two-sided constraints $0 \leq x \leq 1$ in (2.2), while the theory in [11] only studies the case with one-sided constraints $x \geq 0$. The main difficulty of such an extension is to show that the corresponding primal-dual formulation (2.4) is sharp.

At a high level, one can show that under the sharpness condition the PDHG iterates have sublinear convergence rate, i.e.,

$$\text{metric}(\bar{z}^{n,t}) \leq \frac{C}{t} \text{metric}(\bar{z}^{n,0}) ,$$

where metric is a non-negative metric to measure the quality of the solution, which is 0 at an optimal solution, and C is a problem-dependent constant. Generally speaking, the restart condition can guarantee that $t \geq 2C$, then we have

$$\text{metric}(\bar{z}^{n+1,0}) = \text{metric}(\bar{z}^{n,t}) \leq \frac{1}{2} \text{metric}(\bar{z}^{n,0}) ,$$

thus the metric halves after one outer iteration. This guarantees the global linear convergence of Algorithm 1.

In the next section we provide the first documented application of the algorithm, and compare its performance with state-of-the-art benchmark methods.

2.5 Empirical Validation

We begin this section by introducing the data and business problem that we use the algorithm to solve. We then compare the method with primal simplex, dual simplex and a barrier method. Our comparisons include computation time, performance (profit), and feasibility (was the method able to obtain a solution). We recognize that different firms have access to different hardware resources, and so we repeat these comparisons using four different hardware options.

2.5.1 Data and Profit Estimation

The data was provided by a large American retailer. The retailer operates membership wholesale club stores selling a broad range of products, including electronics, furniture, outdoor, toys, jewelry, clothing, and grocery items. The retailer uses promotions to attract new members, and has implemented large scale experiments

to help it personalize which promotions it should send to different prospective customers. Customers must register for a club membership in order to purchase, and the retailer matches the name and address provided at registration to track which customers responded to each promotional offer.

The data describes a large field experiment conducted by the firm in 2015. The experiment’s goal was to compare how prospective customers responded to five different direct mail promotions and a no action control (a total of six experimental conditions). The households were randomly assigned to the six experimental conditions. Each prospective customer refers to one prospective household, whose information is purchased from a third-party commercial data supplier. These households are located in two states. In later discussions, we will label the states as State A ($I = 1,061,438$) and State B ($I = 1,308,658$). In total, the experiment included approximately 2.4 million unique households.⁷

The profit earned from each household was measured over the next twelve months. The profit measure included mailing costs, membership revenue, and profits earned from purchases in the store (if any). Only a relatively small number of households responded, and so across the 2.4 million households, the profit was negative for most households in the five promotion conditions (due to mailing costs), and zero for most households in the no action control (there were no mailing costs in the no action control condition). However, for the customers who did respond, the twelve month profit measure was positive and large. This distribution of outcomes is typical of many marketing actions. Low response rates are particularly common when prospecting for new customers.

The firm’s goal was to use the data from these 2.4 million customers as a training sample to estimate how households in a separate implementation sample would respond to the six different experimental treatments. The randomized assignment of customers in this training data allows it to interpret estimated treatment effects as causal effects. Estimating these treatment effects is outside the scope of this paper.

⁷The data for three of the experimental conditions in this experiment was previously used in [153].

Instead, our focus is on proposing a model to use those estimates to design optimal personalization policies for the separate implementation sample (of prospective customers). For this reason, we relegate to Appendix D.2 an explanation of how we used the training data to estimate a model that predicts how a customer will respond to each marketing action, as a function of observable covariates (contextual variables).

2.5.2 Design of the Validation Exercise

We consider three benchmark methods: primal simplex, dual simplex, barrier method. These are considered state-of-the-art solvers for linear programming problems. We implement each method using Gurobi, which is a commercial software package explicitly designed to solve linear programming problems. Data engineers at Gurobi have spent years fine-tuning their implementations of these benchmark methods, and the Gurobi implementations are generally considered amongst the most powerful implementations of these benchmark methods.

Different firms have access to different hardware resources. To compare how this affects the performance of the different methods we compare their performance using four different hardware combinations:

- **H1:** A MacBook Pro with 8-core CPU and 16GB memory;
- **H2:** An iMac with 8-core CPU and 64GB memory;
- **H3:** A computing server resource with 24-core CPU and 256GB memory;
- **H4:** A computing server resource with 24-core CPU and 512GB memory.

The first two hardware options are likely to be feasible and affordable for essentially any firm. Access to the server-level resources is likely to be limited to medium and larger firms, or smaller firms with relatively sophisticated technology capabilities.

We compare the different methods using three performance measures. We first measure feasibility, by asking whether a given method and hardware combination converges to an optimal solution. The second measure focuses on computation time,

and measures the total time used to solve the problem (from the set of problems that are actually solved). The third measure focuses on performance and measures optimal predicted profit. This measure is used to test whether each method delivers the same outcomes when the problem is solvable. For confidentiality reasons, in all of the settings, we scale the optimal profits by setting the predicted profit in the solution produced by our proposed method to 100.

The volume and similarity constraints in Equation (2.1) require definitions of customer segments. In practice, these segments could be defined using a wide variety of different measures. For example, the constraints may be defined at the store level. Predicted store sales performance this year may need to be at least as high as last year, or the firm may require that the average number of promotions received in households neighboring each store is similar for each store. Alternatively, the constraints may be defined using geolocation measures, such as zip codes [94]). When targeting existing customers, it is common for segmentation to be based upon past purchasing measures, such as the time since the last purchase, or the number of purchases in the past year.

A primary objective of our empirical analysis is to observe how well different methods perform when varying the number of constraints (K). For this reason, we will segment households in the implementation data using zip codes (we observe each prospective household's zip code). Zip codes are labelled into a hierarchy. For example, a four-digit zip code is identified by the first four digits of a five-digit zip code, and contains all of the five-digit zip codes that share those first four digits. As a result, the assignment of five-digit zip codes to four-digit zip codes is mutually exclusive and collectively exhaustive. The same properties apply when we consider three-digit zip codes (or even two-digit and one-digit zip codes). In Section 2.5.3, we define the customer segments using either three-digit zip codes, four-digit zip codes or five-digit zip codes. The combination of segments and markets yields a total of nine different scenarios. In Table A.11 we label these scenarios, and summarize the number of segments (K) in each scenario.

We choose the following parameters for the constraints:

- $\alpha_k^j = 0$ for $j = 1, \dots, 5$ and $k = 1, \dots, K$;

- $b_k^j = 0.9n_k$ for $j = 1, \dots, 5$ and $k = 1, \dots, K$;
- $L_k = 0$ for $k = 1, \dots, K$;
- $c_i^j = j$ for $i = 1, \dots, I$;
- $U_k = 4n_k$ for $k = 1, \dots, K$;
- $\lambda_j^{k_1k_2} = 1.1$ for $j = 1, \dots, 5$, $k_1 = k_2, \dots, K$ and $k_2 = 1, \dots, K$;
- $d_i^j = 1$ for $i = 1, \dots, I$ and $j = 1, \dots, 5$;
- $\gamma^{k_1k_2} = 1.25$ for $k_1 = k_2, \dots, K$ and $k_2 = 1, \dots, K$;

With these parameter choices, the number of constraints in the problem is equal to:
 $M = JK + K + K(K - 1)J/2 + K(K - 1)/2 + I$.

The parameter choices were identified through experimentation, and were selected to ensure that the problems are both feasible and non-trivial. We also try to incorporate the business meaning for each parameter. For example, the maximum amount for each marketing action b_k^j and a combination of all marketing actions U_k might depend on the size of each customer region.

2.5.3 Validation Results

Table A.12 reports the results when using hardware *H1* for each of the nine problem scenarios. We restrict attention to scenarios that can be solved using at least one of the methods.

As we can see from Table A.12, our methods can solve all scenarios that can be solved using state-of-the-art methods and extend the settings that can be solved. What's more, based on the optimal profit delivered in all settings, our method indeed achieves the convergence. One thing we want to stress is that the advantage of our method lies in the ability to solve larger scale problems, but we do not guarantee to solve a medium size problem faster than the state-of-the-art method.

2.6 Conclusion

In this paper, we propose a method for optimizing personalization problems in the presence of constraints. We focus on two types of constraints. Volume constraints restrict the total number of marketing actions that can be taken, either through (predetermined) minimum or maximum thresholds. Similarity constraints limit the difference in the frequency of marketing actions taken with different customer segments.

The proposed method departs from existing state-of-the-art methods by arguing to use the first-order methods for linear programming to increase scalability. The algorithm overcomes the challenge that first-order methods quickly find moderately accurate solution but progress towards an optimal solution slows down over time by proposing a two-loop primal-dual hybrid gradient algorithm.

We provide theoretical guarantees on the performance of the proposed method in personalization with constraints problem. First, we show that our proposed method can solve larger problems (in terms of customers and constraints) than any of the state-of-the-art benchmarks. Second, we adapt existing guarantees on optimality and computation speed, by accommodating the characteristics of personalization problems.

We also present the first documented empirical application of the method. In this application we compare the proposed method to three sophisticated benchmarks: primal simplex, dual simplex and barrier method. Our comparisons of the proposed method with these benchmarks reveals that the proposed method can solve larger problems that include more customers and more constraints. It also offers improvements in computation time, and performance (measured in terms of predicted profit). These theoretical results and empirical results confirm that designing large-scale personalization policies with constraints is now feasible.

Much of the recent research in marketing using machine learning has focused on new methods for estimating customer response functions. Our paper takes a step in a different direction: using recent advances in optimization methods to help

firms optimize policies once they have estimated those response functions. Many interesting problems remain, and these offer promising avenues for future research. First, as we mentioned in Section 2.3, our approach remains within the predict-then-optimize paradigm. This framework has a potential limitation: the estimation goal is not always the same as the optimization goal. [112] propose one approach to address this misalignment when the personalization problem has no constraints. Future research could investigate how to address this misalignment in the presence of constraints. Second, because we use finite sample datasets to estimate customer response functions, these response estimates are estimated with error. The errors will affect the performance of optimization methods that rely upon those estimates. It is unclear how to mitigate the cost of these errors in the optimization step.

Chapter 3

The Power of Commitment in Group Search

In this chapter, we build a two-member two-period model to show that when a group of people with different preferences conduct search and make a decision together, they can benefit from making a commitment on the number of products to search *ex ante* when the search cost is very small or relatively large. The underlying mechanism is that, because of the preference inconsistency between group members, they tend to search fewer products and thus have lower expected utility in group search than in single-agent search, and making a commitment on the number of products to search helps mitigate the preference inconsistency problem in group search, especially when the search cost is very small or relatively large. If consumers can observe product prices before search and the firm sets product prices endogenously, the firm can benefit from letting consumers commit on the number of products to search *ex ante* if consumers search as a group and their search cost is small. We also consider several extensions to show the robustness and boundary conditions of our findings.¹

¹This chapter is collaborated with Xinyu Cao (Assistant Professor of Marketing at NYU Stern).

3.1 Introduction

Though most marketing and economics literature consider people making decisions individually, there are many scenarios in which a group of people with different preferences need to make a decision together ([3, 35, 101, 130, 150]). For example, household purchase decisions—especially for “public” goods within the family like housing or furniture—are often made by family members together (e.g., [35, 51]). In public or private organizations, most important decisions are also made by a committee rather than by single individuals [130]. For example, recruiting decisions in a company or an academic department are often made by a hiring committee [101]

Before making the final choice, group members first need to conduct search together. Searching as a group—we call it “group search”—has distinct features compared to single-agent search, because group members have different preferences and they need to integrate their preferences in some way to make a decision together. There have been a few works investigating the problems in group search ([3, 44, 130, 176]). One thing we notice is that these works all assume that the group takes the sequential search strategy, i.e., decision makers decide whether to stop or to continue searching after evaluating each alternative ([125]). However, there is another search strategy that is often used in reality—the fixed-sample strategy ([52, 96, 158]), meaning that decision makers first decide the number of alternatives to search n , and then search n alternatives and choose one from them.

It is understandable that existing works on group search assume the group taking the sequential search strategy by default, given the traditional belief that the sequential strategy dominates the fixed-sample strategy (e.g. [23, 77, 124, 148]). The reason is that the sequential strategy allows decision makers to make use of more information—it allows decision makers to flexibly decide when to stop searching based on the information revealed during the search process, instead of having to make a commitment on the number of alternatives to search *ex ante*. However, we notice that this belief is made on the premise of single-agent search, i.e., the search process that is conducted by a single decision maker. When it comes to group search, as group

members have divergent preferences, having flexibility in the search process may not be favorable anymore, and instead, making a commitment *ex ante* may benefit the group.

In this paper, we investigate whether the fixed-sample strategy may work better than the sequential strategy in group search, or in other words, whether the group can benefit from committing on the number to search *ex ante*.

We model group search as a game between group members. The group is presented with one product² in each period.³ The value of the product to each group member draws from a common, known distribution. These draws are i.i.d., both across products and across group members. In the main model, we assume product values follow a uniform distribution. Each group member incurs search cost c to search one product. When using the fixed-sample strategy, the group members first vote to decide the number of products to search. After searching the pre-determined number of products, they vote to select one product as the group's choice. When using the sequential strategy, the group members vote to decide whether to stop or continue searching after searching each product. In the main model, we assume recall is allowed in the sequential strategy. This assumption makes the sequential strategy comparable to the fixed-sample strategy in the sense that all searched products can be selected by the group.

We consider a group with two members, which represents a typical setting where a couple or a group consisting of two parties make a decision together. We further simplify the problem by assuming that they can search for at most two periods (products). The two-period (or two-product) assumption is also often seen in literature to make the analysis tractable (e.g. [99, 105]).

²For simplicity, we will use “products” hereafter to represent the alternatives that the decision maker(s) look for, but our theory applies to the general case in which the decision maker(s) may search for a job candidate, etc.

³The fixed-sample strategy is sometimes called a simultaneous search (e.g. [116]) in the literature, implicitly assuming that decision makers can search all the n products simultaneously in one period. Herein, we assume that irrespective of whether a sequential strategy or a fixed-sample strategy is used, decision makers can search only one product in a period. Hence, the key difference between the sequential strategy and the fixed-sample strategy is not timing (whether they can search multiple products simultaneously), but is whether decision makers make a commitment on the number to search *ex ante*.

For all the decisions in the search process—including the number of products to search, whether to stop or to continue, and which product to choose—if the two members vote for the same option, this one becomes the group’s choice, and if the two members vote for different options, each of their choices has equal probability to the group’s choice. This tie-breaking rule can be considered as the group goes through a Nash bargaining and the two members have equal bargaining power. In Section 3.2 we explain that this voting rule can be considered as the “best” rule in our model setup.

We want to emphasize that our paper is not to fully solve the problem of group search. Thus, we are not trying to make the most general model assumptions and find the optimal search strategy and voting rule for a general case—these are beyond the scope of this paper. Instead, our goal is to make a point that in group search, making a commitment can mitigate the inefficiency caused by preference inconsistency between group members. That is, the fixed-sample strategy can be more favorable than the sequential strategy under certain conditions. The simplest setup (i.e., a two-member two-period model) can already capture the key trade-off—the commitment device of fixed-sample strategy and the flexibility advantage of sequential strategy, and at the same time, it makes the two search strategies comparable and the problem analytically tractable, as we explain in details in Section 3.2.1.

Under these assumptions, we analyze the group members’ choices and expected utilities in equilibrium under the fixed-sample strategy and the sequential strategy, and compare their expected utilities under the two strategies. We find that the fixed-sample strategy works better (i.e., leads to higher expected utility) than the sequential strategy when the unit search cost (relative to the dispersion parameter of the value distribution) is very small or large enough.

We dig into the mechanism underlying the result. The sequential strategy works better than the fixed-sample strategy in single-agent search because of its flexibility advantage—that is, it can flexibly decide how many products to search based on the information revealed in the search process. This flexibility advantage is more salient when the search cost is in the middle range, and goes to 0 when the search

cost becomes 0 or large enough because the two strategies will collapse to the same. The flexibility advantage still exists in group search but has a smaller magnitude. Furthermore, compared to single-agent search, group search leads to lower expected utility, for two reasons: 1) the group’s choice may not be the best choice for each group member, we call this “divergence effect,” and 2) because of this, the marginal benefit of searching one more product is lower in group search, and therefore group search chooses to search fewer products than single-agent search, we call this “sacrifice effect.” The magnitude of divergence effect does not vary with search cost, but the sacrifice effect becomes larger when the search cost gets smaller. Under the sequential strategy, both effects always co-exist. Under the fixed-sample strategy, when the search cost is small, group search and single-agent search commit to searching two products, and therefore only the divergence effect exists. When the search cost is small enough, the sacrifice effect of sequential strategy exceeds its flexibility advantage, and thus the sequential strategy works worse than the fixed-sample strategy here. When the search cost is large, group search and single-agent search both commit to searching one product, so neither divergence effect nor sacrifice effect exists. The sequential strategy’s flexibility advantage drops faster than its sacrifice effect, and therefore the sequential strategy works worse than the fixed-sample advantage when the search cost is relatively large.⁴ We can see that making a commitment on the number to search mitigates the inefficiency caused by preference inconsistency between group members, and therefore can lead to a higher expected utility when the search cost is very small or relatively large.

We further study the implications on firm’s strategies. We consider a case in which consumers can directly observe product prices before search, and a monopoly firm endogenously sets product prices taking consumer search into account. We find that if consumers conducts single-agent search, the firm can always set higher prices and earn higher profit under the sequential strategy than under the fixed-sample strategy. If two consumers search in a group, the firm sets lower prices to encourage

⁴When the search cost is even larger, the sequential strategy collapse to the same as the fixed-sample strategy.

them to search more products, and the firm earns lower profit. More importantly, in group search, the firm sets higher prices and earns higher profit under the fixed-sample strategy than under the sequential strategy when the search cost is small. This is because the group commits to searching two products under the fixed-sample strategy when the search cost is small, whereas under the sequential strategy, the group does not commit to the number to search and the firm needs to lower the prices to induce them to search more products so that the group is more likely to make a purchase. If the firm can determine the search strategy of consumers by manipulating how to present products to consumers, then it should induce them to use the fixed-sample search strategy when the search cost is small.

We also consider several extensions to investigate the robustness and boundary conditions of our finding. First, we consider an extension in which recall is not allowed in the sequential strategy. Our finding still holds and in a larger parameter space. Second, we consider an alternative distribution assumption—we assume product values follow a normal distribution. We show that our key finding still holds, i.e., the fixed-sample strategy performs better than the sequential strategy when the search cost is very small or large enough. Third, we consider alternative voting rules for the group to determine the number of products to search—unanimity rule (both members need to agree to stop searching) and “one-is-enough” rule (the group stops searching as long as one member votes to stop), which are commonly considered quota rules in search literature (e.g. [3, 176]). We find that under each voting rule, there is a region in which the fixed-sample strategy works better than the sequential strategy. Although the regions in which our main finding holds are different across voting rules, further analysis indicates that the underlying mechanisms are consistent. The results also indicate that if the group can endogenously choose the voting rule, then sequential strategy can always work better than the fixed-sample strategy—this can be considered as a boundary condition of our finding.

3.1.1 Contributions to Literature

Our work contributes to the recent literature on group search (e.g., [3, 44, 130, 176]). We notice that all these existing works consider group search problem under the sequential strategy. For example, [176] considers a two-member sequential search, and shows that the outcome can converge to a Nash bargaining outcome with bargaining power determined by the patience of the two players. [3] model group search under the sequential search strategy and compared it to single-agent search. They show that the group members are less picky (i.e., searching fewer products) than a single decision maker. We make a similar point as [3] in the sense that group search has smaller (expected) number of products to search than single-agent search, but we further show that this is because of the preference inconsistency problem in group search, and the search agents' commitment power under the fixed-sample strategy can mitigate the preference inconsistency problem in group search. Therefore, the fixed-sample strategy can work better than the sequential strategy in group search under certain conditions.

Our paper also contributes to the search literature by being the first to give a formal analysis concerning the key driving force of the difference between the fixed-sample strategy and the sequential strategy. Although [23], [77], [124], and [148] have all indicated that the sequential strategy dominates the fixed-sample strategy in single-agent search, they make the argument based on either intuition ([23, 77, 148]) or numerical analysis ([124]). Under a specific distribution assumption, we analytically show that the sequential strategy dominates the fixed-sample strategy in single-agent search. We also show that the flexibility advantage of the sequential strategy first increases and then decreases with the search cost, and goes to zero when the search cost approaches 0 or becomes large enough. Furthermore, we show that the fixed-sample strategy can be better than the sequential strategy in group search, and we analytically capture the benefit from the commitment power of the fixed-sample strategy.

Our paper is broadly related to the literature of consumer search (e.g., [20, 104,

105, 108]). Instead of characterizing the amount of information a consumer learns from a product, we assume that the value of a product is fully revealed once the product has been searched, and we capture the intensity of a search process using the number of products to search. The general implication of our model should still hold if we use the amount of information to learn to represent the intensity of search. We also contribute to the literature regarding the implications of consumer search on firm's strategy (e.g. [60, 61, 64]) by investigating how firm's should guide consumers' search strategy when they search individually or as a group.

There have been some empirical papers studying the comparison between the fixed-sample strategy and the sequential strategy, and trying to understand which search strategy consumers really use. [33] and [95] develop methodologies to estimate search costs for both the fixed-sample and the sequential strategy, and to test which strategy is used by consumers. [21]'s result is consistent with the sequential strategy model. However, [52]'s and [96]'s results favor the fixed-sample strategy model. These papers all assume that the consumers are conducting single-agent search. Our paper shows that for group search, the fixed-sample strategy can theoretically dominate the sequential strategy under certain conditions, and our results call for more empirical studies on group search, especially on testing between the two strategies. From another angle, our paper implies that the divergence in the empirical findings regarding which search strategy consumers use can potentially be explained by whether the consumers are conducting group search or single-agent search. It also implies the importance of distinguishing between group search and single-agent search when investigating consumers' search behavior.

Broadly speaking, our research is also related with the literature about group decision making and group consumption. For example, [101] consider a group decision problem between two agents, and shows that agents will strategically shade their actions towards the extreme to influence the group decision,⁵ and the extent of polarization will be smaller if they reveal their preference sequentially rather than

⁵This is also why we assume each agent votes for one alternative instead of reporting her utility from each alternative. Voting for one alternative can be considered as choosing an "extreme."

simultaneously. [99] compares sequential versus simultaneous decision-making mechanisms in group buying, and show that the sequential mechanism leads to higher deal success rates and larger expected consumer surpluses. Under different settings, these two papers show the advantage of the sequential mechanism because it allows decision makers to make use of more information. Admitting this advantage, we show that the fixed-sample search strategy (also called simultaneous search strategy) can work better than the sequential strategy because of its commitment on the number to search mitigates the preference inconsistency problem in group search.

The paper is organized as follows. In Section 3.2, we set up the main model, analyze the equilibria in the fixed-sample strategy and in the sequential strategy, compare group members' expected utility under the two strategies, and investigate the underlying mechanism why the fixed-sample strategy can work better in group search. In Section 3.3, we investigate the implications for firm's strategy. Section 3.4 are the extensions. Section 3.5 concludes the paper and discusses future research directions.

3.2 Main Model

In this section, we develop a group search model and compare the fixed-sample search strategy versus the sequential search strategy under group search. We show that the commitment device of the fixed-sample strategy can mitigate group members' divergence problem and make the fixed-sample strategy favorable under certain conditions.

3.2.1 Model Setup

We consider the problem in which a group of decision makers, denoted as J , want to choose one product from a common set of products I .⁶ The product values are unknown to group members, so they need to engage in a search process to learn the values of products.

⁶We do not model the formation process of the consideration set (e.g., [127]) and assume that set I is given exogenously.

The value of product i to group member j is denoted as X_{ij} , which is unknown *ex ante*. Group member j will learn the product value X_{ij} immediately once she searches product i . X_{ij} 's are drawn from a common distribution with c.d.f. $F(\cdot)$ and p.d.f. $f(\cdot)$ and the distribution is known to group members. The common distribution assumption implies that the products are homogeneous *ex ante* and thus the sequence of searching does not matter.⁷ Following the group search literature (e.g. [3]), we assume that X_{ij} 's are independent, both across group members and across products. It implies that searching one product does not help to learn the value of other products, and that group members have heterogeneous and independent tastes.⁸

We assume that all group members search together,⁹ and each group member incurs search cost $c > 0$ when the group searches a product.¹⁰ Under the fixed-sample strategy, group members first commit to the number of products to search, and then search the committed number of products and choose one from them ([158]). Under the sequential strategy, group members search one product at a time, and decide whether to stop or to continue searching each time when they have searched a product ([125]). Here we assume recall is allowed, meaning that if the group continues searching, the searched products will still be available for them to choose. The recall assumption makes the sequential strategy comparable to the fixed-sample strategy in the sense that now under both strategies, they can choose from all searched products. In contrast, if recall is not allowed in the sequential strategy, the searched products will be unavailable if the group continues searching. We will show in Section 3.4.1

⁷In contrast, in [174], products are assumed to be *ex ante* heterogeneous. The decision maker knows that product values are drawn from different distributions, and a search strategy includes not only the number to search but also the sequence of searching.

⁸In reality, group members' tastes might be correlated. For example, their tastes might be positively correlated such that product values can be decomposed as $X_{ij} = q_i + \epsilon_{ij}$ where q_i denotes the quality part of product i which group members agree upon, and ϵ_{ij} denotes the heterogeneity of group members' tastes and is independent across j . In the case where only the “ q_i ” term exists, group members have perfectly aligned preferences and their search will be identical to a single-agent search. We are considering the case in which there is no “ q_i ” term but only the “ ϵ_{ij} ” term, so that X_{ij} 's are independent across j .

⁹Since the preferences of group members are heterogeneous and independent, the group cannot have only one group member search a product and share the searched information with other group members.

¹⁰Herein, we assume that searching as a group does not dilute the cost of searching for each group member. The reason is that the cost of searching can be considered as mainly consisting of the cost of time and effort, which cannot be shared between group members.

that our finding still holds if recall is not allowed in the sequential strategy. We assume there is no time discounting in utility, so that the difference between the two strategies is not driven by their difference in timing.

For tractability, we consider a group with two members ($J = \{A, B\}$), which represents a typical setting where a couple with inconsistent preferences or a group consisting of two parties make a decision together. We further simplify the problem by assuming that they can search for at most two periods. They search one product a time, so the group can search at most two products. The two-member (e.g., [35]) and two-period/two-product (e.g. [99, 105]) assumptions are often made in literature to make the problem analytically tractable.

As mentioned in the Introduction, the focus of this paper is not to fully solve the problem of group search, but to make a point that in group search, making a commitment on the number to search can mitigate the inefficiency caused by preference inconsistency between group members and lead to higher expected utility. This setup (i.e., a two-member two-period model) can already capture the key trade-off—the commitment power of fixed-sample strategy and the flexibility of sequential strategy, and at the same time, it makes the problem analytically tractable and the two search strategies comparable, as explained below.

The decisions the group need to make in group search are the number of products to search (under the fixed-sample strategy), whether to stop or to continue after searching a product (under the sequential strategy), and which product to choose among the searched products (under both strategies). For each of the decisions, group members need to vote in order to reach a group decision. Notice that under the sequential strategy, whether to stop after searching a product is a binary choice, but the number of products to search under the fixed-sample strategy is non-binary when there are more than two products available. Furthermore, if there are more than two products, deciding which product to choose among the searched products will also be non-binary voting under both strategies. The two-product assumption makes all the group decisions in the two search strategies binary (i.e., have two options), and therefore the two search strategies are theoretically more comparable. However, as

we will discuss in Section 3.2.4, our insights are also generalizable to the case when the group can search more than two periods (products).

Given that there are two products, the commonly used voting methods—each voter selecting one product, (i.e., each voter giving a ranking of candidates), and scoring method (i.e., each voter assigning score $\{n - 1, n - 2, \dots, 0\}$ to the n candidates)—collapse to the same one ([138]).¹¹ The commonly used rules to determine the winning option—the plurality rule (the option with the most votes wins) and the majority rule (the option that receives more than $50\% \times |J|$ votes wins)—also become consistent. (We set a tie-breaking rule that when two options receive the same number of votes, each of them has equal probability to be the group’s choice.) It has been shown that when there are two options, the majority rule (and equivalently, the plurality rule) is the “best” voting rule in the sense that it is the only voting rule that satisfies the key characterizations of voting results—neutrality, anonymity, unanimity and positive responsiveness (May’s Theorem, [84, 123, 138]). For the group with two members, specifically, if the two members vote for the same option, this option will be the group’s choice, and if the two members vote for different options, each member’s choice has equal probability to be the choice of the group.¹² The voting rule also requires that group members vote simultaneously and one cannot observe the other group members’ choice until the round of voting is finished.

To make the comparison between the two strategies analytically tractable, we assume that the distribution $F(\cdot)$ is a uniform distribution on $[\mu - d, \mu + d]$, where μ is the mean and d captures the dispersion of the distribution. The uniform distribution assumption is also often seen in classic search literature (e.g. [158, 162, 178]).¹³ We

¹¹If there is a centralized planner, the first-best strategy is to add up group members’ utilities for each product, and then the group can search as a single agent. If there is no centralized planner and the group adopts the rule that lets each group member reports her utility from each product, group members will not be truth-telling. Instead, each group member will exacerbate her utility of the product that she favors ([101]). Thus, the first-best strategy is not achievable here and we need to assume each group member gives a vote. At the same time, we use single-agent search as the benchmark and compare group search to it.

¹²We can think of the group members going through a bargaining process when they disagree, and every group member has equal bargaining power.

¹³Notice that we need to make a specific distribution assumption for product values because we want to compare the utility under the fixed-sample strategy (which involves the $E[\max\{X_1, X_2\}]$) versus that under the sequential strategy (which involves $E[X \geq \xi]$), and thus we cannot give a

assume there is no outside option in the main model, and we relax this assumption in Section 3.3 when discussing the firm's strategy.

Next, we analyze the group members' search problem under the fixed-sample strategy and under the sequential strategy, respectively.

3.2.2 Fixed-sample Strategy and Equilibrium Analysis

Suppose the group adopts the fixed-sample strategy. The game between the group members has two stages.

At the first stage, each group member j proposes the number of products to search N_{GFj} , $j \in \{A, B\}$, and then the group decides on the number of products to search N_{GF} . Across our paper, the first subscript is to denote group search G or single-agent search S , and the second subscript is to denote fixed-sample strategy F or sequential strategy S . If $N_{GFA} = N_{GFB}$, then the group will search $N_{GF} = N_{GFA} = N_{GFB}$ products. If $N_{GFA} \neq N_{GFB}$, then the number of products to search N_{GF} will be N_{GFA} or N_{GFB} with equal probability. (We start the analysis for a general case without imposing the constraint that $N_{GF} \leq 2$, and then we consider the two-period model where $N_{GF} \leq 2$.)

At the second stage, the group searches N_{GF} products together. After searching N_{GF} products, each group member j will vote for a product, denoted as product i_{GFj}^* . If the two members vote for the same product (i.e., $i_{GFA}^* = i_{GFB}^* = i_{GF}^*$), then product i_{GF}^* will be chosen by the group. If the two group members' choices are different (i.e., $i_{GFA}^* \neq i_{GFB}^*$), one of the two products will be selected by the group with equal probability.

We focus on subgame-perfect equilibria and solve the game by backward induction. At the second stage, suppose the group has determined the number of products to search to be N_{GF} , and now we calculate each group member j 's expected utility $EU_{GFj}(N_{GF})$. Let's think of the problem from member A 's perspective. Given member A 's choice i_{GFA}^* , there are three possible scenarios: 1) member B votes for the same product and this product is selected, 2) member B votes for a different

closed form solution with a general distribution assumption.

product and member A 's choice is selected, and 3) member B votes for a different product and member B 's choice is selected. Since the two members' valuations are independent and they make choices simultaneously, given A 's choice i_{GFA}^* , the N_{GF} products are still symmetric from B 's perspective and have the same probability of being B 's choice. Therefore, the probabilities of the three scenario are $1/N_{GF}$, $(N_{GF} - 1)/(2N_{GF})$, and $(N_{GF} - 1)/(2N_{GF})$, respectively. Then group member A 's expected utility is:

$$\begin{aligned} EU_{GFA}(N_{GF}) &= \left(\frac{1}{N_{GF}} + \frac{N_{GF} - 1}{2N_{GF}} \right) E[X_{i_{GFA}^*, A}] + \frac{N_{GF} - 1}{2N_{GF}} E[X_{i_{GFB}^*, A} | i_{GFB}^* \neq i_{GFA}^*] - cN_{GF} \\ &= \frac{1}{2} E[X_{i_{GFA}^*, A}] + \frac{1}{2} E[X] - cN_{GF}. \end{aligned} \quad (3.1)$$

Notice that $E[X_{i_{GFB}^*, A} | i_{GFB}^* \neq i_{GFA}^*] = (N_{GF} E[X] - E[X_{i_{GFA}^*, A}]) / (N_{GF} - 1)$ because the $N_{GF} - 1$ products not voted by A are symmetric from A 's perspective *ex ante* and they have the same probability of being voted by B .

Equation (3.1) indicates that, to maximize her utility $EU_{GFA}(N_{GF})$, group member A should maximize $E[X_{i_{GFA}^*, A}]$. Therefore, A 's optimal strategy at the second stage is to vote for the product which she draws the highest value from among the N_{GF} products. By symmetry, B 's expected utility has a symmetric formula, and her optimal strategy at the second stage is also to vote for the one she draws the highest value from. That is, for each member j , $E[X_{i_{GFj}^*, j}] = E[\max\{X_{1j}, \dots, X_{N_{GFj}}\}] \equiv E_{max}(N_{GF})$. Then both group members' expected utility given N_{GF} is

$$EU_{GFA}(N_{GF}) = EU_{GFB}(N_{GF}) \equiv EU_{GF}(N_{GF}) = \frac{1}{2} E_{max}(N_{GF}) + \frac{1}{2} E[X] - cN_{GF}. \quad (3.2)$$

The last term represents the search cost, and the first two terms indicate that, for a two-member group fixed-sample search, the selected product has 1/2 chance of being the focal group member's favorite and has 1/2 chance of being an "average" product for the focal group member.

Now we move backward to the first stage. When proposing the number of products to search, since each member's choice of N_{GFj} has equal probability to be the choice of the group, each group member j chooses $N_{GFj} \in Z^+$ to maximize

$EU_{GF}(N_{GFj})/2 + EU_{GF}(N_{GFj'})/2$, $j = A, B$. Given that member j is unable to affect the other member's choice $N_{GFj'}$, member j will just choose N_{GFj} to maximize $EU_{GF}(N_{GFj})$. By symmetry, both group members will choose the same number of products to search in equilibrium, i.e., $N_{GFA}^* = N_{GFB}^* \equiv N_{GF}^*$.

In the two-period model, the group can search one or two products ($N_{GF} \in \{1, 2\}$). Under the uniform distribution assumption, $EU_{GF}(N_{GF}) = E[X] - c = \mu - c$ if $N_{GF} = 1$, and $EU_{GF}(N_{GF}) = E_{max}(2)/2 + E[X]/2 - 2c = \mu + d/6 - 2c$ if $N_{GF} = 2$. Comparing the expected utility under these two choices, we can get the optimal number of products to search¹⁴

$$N_{GF}^* = \begin{cases} 1 & \text{if } c/d \geq 1/6 \\ 2 & \text{if } c/d < 1/6, \end{cases} \quad (3.3)$$

and then the optimal expected utility under fixed-sample strategy is¹⁵

$$EU_{GF}^* = EU_{GF}(N_{GF}^*) = \begin{cases} \mu - c & \text{if } c/d \geq 1/6 \\ \mu + \frac{d}{6} - 2c & \text{if } c/d < 1/6. \end{cases} \quad (3.4)$$

We denote $\hat{t}_{GF} = 1/6$, which means the cutoff value of c/d when group fixed-sample search changes from searching two products to searching one product.

3.2.3 Sequential Strategy and Equilibrium Analysis

We now analyze the game when the group adopts the sequential strategy. The game will proceed in two periods.

In the first period, group members search one product together, and then each group member votes for one of the two options: 1) to continue searching or 2) to stop at the current product. As we have assumed, if they disagree on this, each member's choice has equal probability to be the choice of the group. If the group decides to

¹⁴We take the rule that when two choices of n leads to the same expected utility, the smaller n will be chosen.

¹⁵We restrict the range of c such that searching is beneficial, which requires $\mu - c > 0$ (recall that the value of the outside option is normalized to 0). Since we have also assumed $d < \mu$, then $[d/6, \mu)$, the range in which $N_{GF}^* = 1$, is a non-empty set.

continue searching, the group will search one more product in the second period. After that, they will vote whether to choose the first or the second product (notice that we assume recall is allowed).

We still focus on subgame-perfect equilibria and solve the game using backward induction. We first consider each group member's choice in the second period. Given that the group have searched two products and the game ends in the second period, following the same logic as in the fixed-sample strategy, each group member will vote for the product that delivers the higher value to her.

Now we consider the first period. We analyze from member A 's perspective. Suppose the values of the first product to the two members are $(X_{1A}, X_{1B}) = (x_{1A}, x_{1B})$. If the group decides to stop searching, member A will get product value x_{1A} . If the group decides to continue searching, we denote group member A 's expected value of continuing searching as $V(x_{1A}; x_{1B})$ given $(X_{1A}, X_{1B}) = (x_{1A}, x_{1B})$. Knowing that each group member will vote for the product that delivers the higher value to her, we can get that

$$\begin{aligned}
V(x_{1A}; x_{1B}) &= -c + F(x_{1A})F(x_{1B})x_{1A} + (1 - F(x_{1A}))(1 - F(x_{1B}))E[X|X \geq x_{1A}] \\
&\quad + F(x_{1A})(1 - F(x_{1B})) \left(\frac{1}{2}x_{1A} + \frac{1}{2}E[X|X < x_{1A}] \right) \\
&\quad + (1 - F(x_{1A}))F(x_{1B}) \left(\frac{1}{2}x_{1A} + \frac{1}{2}E[X|X \geq x_{1A}] \right) \\
&= -c + \frac{d}{8} + \frac{\mu + x_{1A}}{2} + \frac{1}{8d}(\mu - x_{1A})(3\mu - x_{1A} - 2x_{1B}). \tag{3.5}
\end{aligned}$$

where $-c$ represents the cost of searching the second product. There are four possible scenarios when the group continues searching: $F(x_{1A})F(x_{1B})x_{1A}$ represents the case in which both members prefer the first product, $(1 - F(x_{1A}))(1 - F(x_{1B}))E[X|X \geq x_{1A}]$ represents the case in which both members prefer the second product, $F(x_{1A})(1 - F(x_{1B})) (x_{1A}/2 + E[X|X < x_{1A}]/2)$ represents the case that member A votes for the first product and member B votes for the second product, and $(1 - F(x_{1A}))F(x_{1B})(x_{1A}/2 + E[X|X \geq x_{1A}]/2)$ represents the case that member A votes for the second product and member B votes for the first product.

Obviously, a member votes to stop if and only if the value of the first product exceeds a certain threshold. Suppose group member A's stopping threshold is ξ_G (i.e., vote to stop if and only if $X_{1A} \geq \xi_G$), and group member B's stopping threshold is ξ'_G . In equilibrium, at the optimal stopping threshold, a member's expected value of voting to stop equals her expected value of voting to continue. Based on this, we show in appendix E.1 that the optimal stopping threshold satisfies

$$\xi_G^* = E_{X_{1B}}[V(\xi_G^*, X_{1B})]. \quad (3.6)$$

(Notice that this formula does not depend on the value of member B's stopping threshold at all.) Plugging in the formula of V given by Equation (3.6), we can get

$$\xi_G^* = \mu + d(2 - \sqrt{3 + \frac{8c}{d}}). \quad (3.7)$$

Since member B faces a symmetric problem as member A, her optimal stopping threshold will be the same, i.e., $\xi'_G = \xi_G^*$. For the stopping threshold to be well-defined, we need to have $\xi_G^* \geq \mu - d$, which is equivalent to $c/d \leq 3/4$. When $c/d > 3/4$, the optimal threshold given by Equation (3.7) is not well-defined, meaning that the group will stop for any value of the first product.

Knowing the group members' optimal stopping threshold, we calculate each group member's expected utility under the sequential strategy by conditioning on whether each member's draw from the first product is higher or lower than the stopping threshold:

$$\begin{aligned} EU_{GS}^* = EU_{GS}(\xi_G^*) &= -c + \int_{\xi_G^*}^{\mu+d} \int_{\xi_G^*}^{\mu+d} x_{1A} f(x_{1A}) f(x_{1B}) dx_{1A} dx_{1B} \\ &+ \int_{\xi_G^*}^{\mu+d} \int_{\mu-d}^{\xi_G^*} \left(\frac{1}{2} x_{1A} + \frac{1}{2} V(x_{1A}; x_{1B}) \right) f(x_{1A}) f(x_{1B}) dx_{1A} dx_{1B} \\ &+ \int_{\mu-d}^{\xi_G^*} \int_{\xi_G^*}^{\mu+d} \left(\frac{1}{2} x_{1A} + \frac{1}{2} V(x_{1A}; x_{1B}) \right) f(x_{1A}) f(x_{1B}) dx_{1A} dx_{1B} \\ &+ \int_{\mu-d}^{\xi_G^*} \int_{\mu-d}^{\xi_G^*} V(x_{1A}; x_{1B}) f(x_{1A}) f(x_{1B}) dx_{1A} dx_{1B}, \end{aligned} \quad (3.8)$$

where $f(x) = \mathbf{1}_{\mu-d \leq x \leq \mu+d}/(2b)$ is the p.d.f. of the distribution of X_{ij} . We then get that:

$$EU_{GS}^* = \mu + d \left[\frac{1}{8} + \frac{1}{48} \sqrt{3 + 8\frac{c}{d}} + \frac{5}{12} * \frac{c}{d} \left(-6 + \sqrt{3 + 8\frac{c}{d}} \right) \right] \quad (3.9)$$

for $c/d \leq 3/4$. When $c/d > 3/4$, the group will only search one product in equilibrium, so $EU_{GS}^* = \mu - c$.

3.2.4 Fixed-sample versus Sequential Strategy

Comparing group members' expected utilities under the fixed-sample strategy and under the sequential strategy given by equations (3.4) and (3.9), we get the key result of this paper, as summarized below:

Proposition 2. *Suppose product values are i.i.d. and follow a uniform distribution $U[\mu - d, \mu + d]$. For a two-member two-period group search: The fixed-sample strategy dominates the sequential strategy when the unit search cost is small enough ($c/d \in [0, \hat{t}_1)$, where $\hat{t}_1 = 1/400(-9 + \sqrt{281}) \approx 0.019$), or large enough ($c/d \in (\hat{t}_2, 3/4)$, where $\hat{t}_2 = 1/400(79 + 3\sqrt{449}) \approx 0.356$). When the unit search cost is in the middle range ($\hat{t}_1 < c/d < \hat{t}_2$), the sequential strategy dominates the fixed-sample strategy. When $c/d \geq 3/4$, the fixed-sample strategy and the sequential strategy collapse to the same.*

The proof of Proposition 2 is available in Appendix E.2. We notice that the condition regarding search cost c is relative to d , the dispersion parameter which captures the potential benefit from the search — decision makers gain more from the search when they have more uncertainty about the product values *ex ante* ([158]). Thus, Proposition 2 shows that in group search, the fixed-sample strategy dominates the sequential strategy when the unit search cost is low enough or high enough relative to the potential benefit from the search. Figure B-16 illustrates the result of Proposition 2: the vertical axis indicates the difference in expected utility between the fixed-sample strategy and the sequential strategy, and the horizontal axis indicates c/d .

To understand the mechanisms behind Proposition 2, we start from analyzing the two search strategies in a two-period single-agent search, which serves as a benchmark.

The detailed analysis is in Appendix E.4. Comparing the expected utility under the two strategies in single-agent search, we get the following result.

Suppose product values are i.i.d. and follow a uniform distribution $U[\mu - d, \mu + d]$.

For a two-period single-agent search:

1. The sequential strategy dominates the fixed-sample strategy for any search cost $c \in (0, d]$;
2. The fixed-sample strategy leads to a higher expected value of the selected product than the sequential strategy when $c/d < \hat{t}_{SF}$, and leads to a lower expected value of the selected product than the sequential strategy when $c/d \geq \hat{t}_{SF}$;
3. The fixed-sample strategy has a higher expected search cost than the sequential strategy when $c/d < \hat{t}_{SF}$, and has a lower expected search cost than the sequential strategy when $c/d \geq \hat{t}_{SF}$;
4. The difference in expected utility between the sequential strategy and the fixed-sample strategy approaches 0 when c/d approaches 0, increases in c when $c/d < \hat{t}_{SF}$, decreases in c when $c/d \geq \hat{t}_{SF}$, and becomes 0 again when $c/d = 1$.

$\hat{t}_{SF} = 1/3$ denotes the cutoff value of c/d that single-agent fixed-sample search changes from searching two products to searching one product.

We observe that, the fixed-sample strategy commits to searching two products when $c/d < \hat{t}_{SF}$, and commits to searching one product when $c/d \geq \hat{t}_{SF}$; in contrast, for any $c \in (0, d)$, the sequential strategy has positive probability of searching either one (probability $1 - \sqrt{c/d}$) or two products (probability $\sqrt{c/d}$), and is more likely to search one product when c is closer to d . This reflects the flexibility advantage of the sequential strategy—that is, the decision maker can decide whether to search the second product after the first product’s information is revealed, so that she can avoid under-searching (when the first draw is of low value) or over-searching (when the first draw is high enough). This flexibility advantage is more salient in the middle range of c/d : when c/d is close to 0, the sequential strategy is very likely to search two products, and when c/d approaches 1, the sequential strategy is very likely to search

only one product, so the difference between the two strategies is very small when c/d is close to 0 or 1.

For the group search, the flexibility advantage of sequential strategy still exists and is also more salient when c is in the middle range, but its overall magnitude is smaller. When c/d is very small or relatively large, the flexibility advantage of sequential strategy becomes dominated by the fixed-sample strategy. We dig into the underlying reason below.

Similar as in Lemma 3.2.4, we also decompose the expected utility in group search into two parts—the expected value of the selected product and the expected search cost, and compare their values in the two strategies. The fixed-sample strategy leads to a higher expected value of the selected product when $c/d \in [0, \hat{t}_{GF})$, and leads to a lower expected value of the selected product when $c/d \in [\hat{t}_{GF}, 3/4)$. The expected cost of fixed-sample strategy is higher than the sequential strategy when $c/d \in [0, \hat{t}_{GF})$, and lower when $c/d \in [\hat{t}_{GF}, 3/4)$. (Proof available in Appendix E.3.) We can see that the pattern in group search is similar as that in single-agent search. Only the cutoff value shifts from $\hat{t}_{SF} = 1/3$ to $\hat{t}_{GF} = 1/6$.

We further compare how group search differs from single-agent search under the two strategies, respectively. Lemma 3.2.4 shows the comparisons.

Suppose product values are i.i.d. and follow a uniform distribution $U[\mu - d, \mu + d]$. Agents can search for at most two products.

1. For the sequential strategy, for all $c/d \in [0, 1]$, group search has lower stopping threshold ($\xi_G^* < \xi_S^*$), smaller expected number of products to search ($E[N_{GS}^*] < E[N_{SS}^*]$), and lower expected utility ($EU_{GS}^* < EU_{SS}^*$) than single-agent search. The gap in their expected utility decreases with c/d ;
2. For the fixed-sample strategy,
 - On $c/d \in (0, \hat{t}_{GF})$, both group search and single-agent search choose to search two product ($N_{GF}^* = N_{SF}^* = 2$); group search has lower expected utility and the gap is constant ($EU_{GF}^* - EU_{SF}^* = -d/6$).

- On $c/d \in [\hat{t}_{GF}, \hat{t}_{SF})$, group search chooses to search one product, and single-agent search chooses to search two products ($N_{GF}^* < N_{SF}^*$); group search has lower expected utility than single-agent search and the gap in their expected utility decreases with c/d .
- On $c/d \in [\hat{t}_{SF}, 1)$, both group search and single-agent search commit to searching one product only ($N_{GF}^* = N_{SF}^* = 1$), and their expected utility is the same.

The proof of Lemma 3.2.4 is available in Appendix E.5. Figure B-17 illustrates Lemma 3.2.4 and Lemma 3.2.4, and it can help us understand how they lead to the result in Proposition 1. The green line shows the difference in expected utility between fixed-sample strategy and sequential strategy under single-agent search. The blue (orange) line compares the expected utility of group search versus single-agent search under fixed-sample (sequential) strategy. Notice that the orange line and the green line coincide when $c/d > 3/4$.

In the first part of Lemma 3.2.4, we show that group search has lower stopping threshold and expected utility than single-agent search under the sequential strategy.¹⁶ This is because the marginal benefit of searching one more product is lower in group search due to preference inconsistency—even if a member finds the next product to be of higher value, the next product may not be chosen by the group since the other member may not find it of higher value, whereas in single-agent search, the search agent can decide which product to choose by herself. Therefore, group members lower their stopping threshold. Group search has lower expected utility than single-agent search for two reasons: the first is that group search lowers the stopping threshold (and therefore searches fewer products in expectation) compared to single-agent search, we call it the “sacrifice effect,” and even if the number of searched products is the same, a group member cannot decide the final choice of the group on her own, we call it the “divergence effect.” Both effects are a result of preference

¹⁶[3] shows that for infinite-horizon sequential search, group search has lower stopping threshold, smaller expected number of searched products, and lower expected utility compared to single-agent search. Our finding echoes that of [3] but under a finite-horizon setting.

inconsistency between group members, but they are two different channels through which group search affects the expected utility, so we distinguish the two effects. The magnitude of divergence effect does not vary with search cost, but the sacrifice effect increases when search cost gets smaller—this is because when search cost gets smaller, decision maker(s) tend to search more, and then group members sacrifice more due to divergence. Under the sequential strategy, the two effects co-exist, and thus we can see that the gap in expected utility between group search and single-agent search under sequential strategy (reflected by the orange line in Figure B-17) gets larger when c/d gets closer to 0.

However, these two effects do not always co-exist under the fixed-sample strategy. The second part of Lemma 3.2.4 compares the fixed-sample strategy in group search versus in single-agent search. When $c/d < \hat{t}_{GF}$, both group search and single-agent search commit to searching two products, so the “sacrifice effect” does not exist here. The “divergence effect” still exists, so group search still has a lower expected utility than single-agent search, and the gap in expected utility between group search and single-agent search is fixed at $-d/6$ (blue line in Figure B-17). When $c/d \in [\hat{t}_{GF}, \hat{t}_{SF})$, group search leads to a smaller number of products to search than single-agent search ($N_{GF}^* = 1, N_{SF}^* = 2$), so both the “sacrifice effect” and the “divergence effect” exist, and the gap in expected utility between group search and single-agent search decreases with c/d . When $c/d \geq \hat{t}_{SF}$, both strategies commit to searching only one product. Neither the “sacrifice effect” nor the “divergence effect” exists here, so group search and single-agent search achieve the same expected utility. The reason why the gap in expected utility keeps constant when $c/d < \hat{t}_{GF}$ and disappear when $c/d > \hat{t}_{SF}$ is because the search agent(s) need to commit on the number of products to search *ex ante* under the fixed-sample strategy. Thus, the commitment device of the fixed-sample strategy helps mitigate the impact of preference inconsistency when the search cost is relatively small or large.

Combining the insights from Lemma 3.2.4–Lemma 3.2.4, we have a clear picture of the results in Proposition 2, as illustrated below.

When the search cost is small ($c/d < \hat{t}_{GF}$), the sacrifice effect does not exist under

the fixed-sample strategy, but exists under the sequential strategy. As c approaches 0, the sacrifice effect under the sequential strategy gets larger, and the flexibility advantage of sequential strategy becomes smaller. When the search cost is small enough, the sacrifice effect of sequential strategy exceeds its flexibility advantage, and therefore the sequential strategy becomes worse than the fixed-sample strategy. To put it mathematically, while the gap in expected utility between group search and single-agent search under the fixed-sample strategy is constant when $c/d < \hat{t}_{GF}$ (reflected by the blue line), this gap under the sequential strategy gets larger as c/d is closer to 0 (reflected by the orange line) and exceeds the gap under the fixed-sample strategy when c/d is close enough to 0 (when $c/d = 0$, $EU_{GS}^* - EU_{SS}^* = (-5/24 + 1/16\sqrt{3})d \approx -0.172d < -d/6 = EU_{GF}^* - EU_{SF}^*$). Thus, when c/d is small enough ($c/d < \hat{t}_1$), the sequential strategy leads to a lower expected utility than the fixed-sample strategy for group search.

When c/d is large ($c/d > \hat{t}_{SF}$), group search and single-agent search has the same expected utility under the fixed-sample strategy since they both commit to searching only one product (blue line), but group search still has a lower expected utility than single-agent search under the sequential strategy (orange line), and the flexibility advantage of sequential strategy gets smaller as c/d gets closer to 1 (green line). We can see that the green line drops faster (in absolute value) than the orange line. Thus, when c/d is large enough ($c/d > \hat{t}_2$), the flexibility advantage of sequential strategy becomes smaller than the sacrifice in expected utility of sequential strategy under group search, and then the sequential strategy will lead to a lower expected utility than the fixed-sample strategy under group search.

To put it in another way, in the region when the search cost is small enough, the sequential strategy works worse than the fixed-sample strategy because it does not search enough—it sacrifices the stopping threshold and the expected number of products to search too much due to preference inconsistency, and the fixed-sample strategy is able to mitigate this problem because it can make a commitment on the number of products to search. In the region when the search cost is large enough, the sequential strategy works worse than the fixed-sample strategy because its risk of over-

search—if the focal member finds the first product satisfying but the other member finds the first product unsatisfying, searching one more product is “over-searching” for the focal member and it is too costly now given the high search cost.

Although our analysis is based on a two-member two-period model, the intuition is generalizable to the case with more than two periods (products). Given any number of products available for search, there always exists a threshold that for search cost lower than this threshold, group search and single-agent search choose the same number of products to search under the fixed-sample strategy, and therefore only the divergence effect exists under the fixed-sample strategy, whereas both the divergence effect and the sacrifice effect exist under the sequential strategy. Furthermore, the sequential strategy’s flexibility advantage goes to 0 when the search approaches 0. Therefore, when the search cost is small enough, the sacrifice effect of the sequential strategy will exceed its flexibility advantage, and then the sequential strategy will perform worse than the fixed-sample strategy. Similarly, given any number of products, when the search cost is large enough, group search and single-agent search will both choose to search one product under the fixed-sample strategy, so neither divergence effect nor sacrifice effect exists here. Both effects still exist under the sequential strategy, and the flexibility advantage of the sequential strategy becomes smaller than the two effects when the search cost is large enough. Thus, the sequential strategy will also perform worse than the fixed-sample strategy when the search cost is large enough.

The literature of time inconsistent preference (e.g., [6, 74, 102, 136]) has shown that when there exists time inconsistent preferences, commitment can be preferable to flexibility as it can mitigate the problems (for example, self-control problems) caused by time inconsistent preferences. Time inconsistent preference is essentially preference inconsistency between today’s self and tomorrow’s self. In our context, the preference inconsistency is between two group members, and we show that commitment can mitigate the problem caused by preference inconsistency in search.

3.3 Implications for Firm's Strategy

We now consider the implications when consumers can directly observe the product prices before conducting search and the firm endogenously decides product prices, taking consumer search into account. Suppose a monopoly firm sells two products, and consumers have unit demand (i.e., will buy at most one product). Consumer j 's utility from product i is $X_{ij} - p_i$. The utility of the outside option is normalized to 0, so consumer j is willing to buy product i if and only if $X_{ij} - p_i \geq 0$ and $X_{ij} - p_i \geq X_{i'j} - p_{i'}$ where $i' \neq i$. For tractability, we assume X_{ij} 's are i.i.d. and follow the uniform distribution on $[0, 1]$. Searching will reveal the value of X_{ij} . Since product prices can be observed before searching, products are no longer homogeneous *ex ante*. We assume consumers can decide which product to search first. It is easy to see that no matter under the sequential strategy or under the fixed-sample strategy when searching one product only, rational consumers should start from the product with lower price. Thus, we can assume $p_1 \leq p_2$ without loss of generalizability.

Without loss of generalizability, we assume the marginal cost of producing each product to be 0. The seller chooses p_1, p_2 to maximize her profit $\pi = \{p_1 D_1 + p_2 D_2\}$, where D_i is the purchase probability of product i . Notice that D_i is a function of p_1, p_2 and also consumers' search cost c and search strategy. Next, we solve this problem under fixed-sample strategy and sequential strategy. The detailed analysis is in Appendix E.6, and we only sketch the key points here.

3.3.1 Single-agent search

Fixed-sample strategy under single-agent search

If the consumer searches one product only, the optimal price is $p_1^* = 1/2$ and the firm's optimal profit is $\pi_{SF1}^* = 1/4$. If the consumer searches two products, the optimal prices are $p_1^* = p_2^* = 1/\sqrt{3}$ and the firm's optimal profit is $\pi_{SF2}^* = 2/(3\sqrt{3})$. We notice that $\pi_{SF2}^* > \pi_{SF1}^*$, meaning that without other constraints, it is more profitable for the firm to let the consumer search two products. This is intuitive—the firm does not bear the search cost, and the consumer is more likely to find a product

value higher than the price when searching two products.

Taking consumer's searching behavior into consideration, the firm determines prices accordingly. In equilibrium, when $c \leq c_{SF2} = 1/(9\sqrt{3})$, the consumer searches two products, $p_1^* = p_2^* = 1/\sqrt{3}$ and the profit is $\pi_{SF}^* = 2/(3\sqrt{3})$. When $c_{SF2} < c < c_{SF1}$, the firm lowers the prices to induce the consumer keeps searching two products (the optimal prices in this part are solved numerically). When $c \geq c_{SF1} = 1/8$, the firm sets $p_1^* = 1 - \sqrt{2c}$, the consumer searches one product only, and the firm's profit is $\pi_{SF}^* = \sqrt{2c}(1 - \sqrt{2c})$. We plot the how the optimal prices and the firm's profit changes with search cost c in Figure B-18(a) (black line). Notice that when the consumer searches two products, it is optimal for the firm to set $p_1^* = p_2^*$. When the consumer searches one product, p_2^* can be any value higher than or equal to p_1^*

Sequential strategy under single-agent search

Suppose the consumer will stop searching if and only if the first product's value $X_1 \geq \xi_S^*$. At the stopping threshold ξ_S^* , the value to stop should be equal to the value to continue, so ξ_S^* should satisfy the following equation

$$\begin{aligned} (\xi_S^* - p_1)^+ &= -c + \Pr(X_2 - p_2 > (\xi_S^* - p_1)^+)[X_2 - p_2 | X_2 - p_2 \geq (\xi_S^* - p_1)^+] \\ &\quad + \Pr(X_2 - p_2 \leq (\xi_S^* - p_1)^+) * (\xi_S^* - p_1)^+ \end{aligned} \quad (3.10)$$

where $(\xi_S^* - p_1)^+ = \max\{\xi_S^* - p_1, 0\}$.

We solve for ξ_S^* and the optimal prices by discussing whether $\xi_S^* \geq p_1$ or $\xi_S^* < p_1$, and whether $p_2 - p_1 + \xi_S^* \geq 1$ or $p_2 - p_1 + \xi_S^* < 1$. We find that for all search cost $c \in [0, 1/2]$, the stopping threshold is $\xi_S^* = 1 - \sqrt{2c}$, the optimal case is $\xi_S^* \geq p_1$, and the two products' optimal prices are equal, i.e., $p_1^* = p_2^*$. When $c \leq c_{SS} = 2/3 - 1/\sqrt{3}$, the firm can choosing the inner solution (i.e., $\xi_S^* > p_1$) to optimize its profit. The optimal prices are $p_1^* = p_2^* = 1/\sqrt{3}$ and the firm's profit is $\pi_{SS}^* = 2/(3\sqrt{3})$. When $c > c_{SS}$, the firm has to lower the price to satisfy the boundary constraint (i.e., $\xi_S^* = p_1$). The optimal prices are $p_1^* = p_2^* = 1 - \sqrt{2c}$ and the firm's profit is $\pi_{SS}^* = (1 - \sqrt{2c})(2\sqrt{2c} - 2c)$. Figure B-18(a) (blue line) plots the optimal prices and the

firm's profit.

3.3.2 Group search

Now we consider group search. The same as in the main model, we assume the group has two parties (i.e., two consumers). They search together and will buy at most one product together. For any decision, if the two members disagree, each member's choice has equal probability to be the group's choice.

Fixed-sample strategy under group search

Similarly, we write down group members' expected utility when searching one product and when searching two products, and the firm's profit function accordingly. In equilibrium, when $c < c_{GF2}$, the firm charges $p_1^* = p_2^* = 1/\sqrt{3}$, the group searches two products, and the firm's profit is $\pi_{GF}^* = 2/(3\sqrt{3})$. When $c > c_{GF2}$, the group searches only one product: in particular, if $c_{GF2} < c < c_{GF1}$, the firm sets $p_1^* = 1/2$ and earns profit $\pi_{GF}^* = 1/4$, and if $c > c_{GF1}$, the firms set $p_1^* = (5 - \sqrt{1 + 48c})/6$ and earns profit $\pi_{GF}^* = (1 - 12c + \sqrt{1 + 48c})/9$. When the group searches one product, p_2^* can be any value higher than or equal to p_1^* . Figure B-18(b) (black line) plots the optimal prices and firm's profit.

Sequential strategy under group search

Suppose each group member j votes to stop after searching the first product if and only if $X_{1j} \geq \xi_G^*$. We solve for the optimal stopping threshold as well as the optimal prices and the firm's profit. When $c < c_{GS1}$, the optimal scenario satisfies $\xi_G^* \geq p_1^* = p_2^*$, and the group can potentially search two products. In particular, when $c < c_{GS2}$, the firm can choose the inner solution to optimize the profit (i.e., $p_1^* = p_2^* < \xi_G^*$), and when $c_{GS2} < c < c_{GS1}$, the firm needs to lower the prices to satisfy the boundary conditions (i.e., $p_1^* = p_2^* = \xi_G^*$). (The values of optimal prices are solved numerically.) When $c > c_{GS1}$, $\xi_G^* = 0$, i.e., the group always stop after searching the first product. That is why there is a price jump at $c = c_{GS1}$ —the firm charges a higher optimal

price when the group searches only one product. In this range, the optimal prices and firm's profit under group sequential search is the same as in group fixed-sample search. The optimal prices and firm's profit is plotted in Figure B-18(b) (blue line).

3.3.3 Implications

We first compare the optimal prices and the firm's profit under the two strategies in single agent search (Figure B-18(a)). The two strategies lead to the same prices and profit when $c \leq c_{SF2}$. The firm can charge a higher price under the sequential strategy than under the fixed-sample strategy when $c_{SF2} < c < c_{SF1}$. This is because under the fixed-sample strategy the firm needs to lower the price to induce the consumer to search two products. Even when the two strategies set the same price (when $c > c_{SF1}$), the firm's profit is still higher under the sequential strategy because the fixed-sample strategy commits to searching only one product in this range, whereas under the sequential strategy, there is still some chance that the consumer will search two products and thus will be more likely to purchase. Therefore, for single-agent search, the flexibility advantage of the sequential strategy not only benefits the consumer, but also benefits the firm.

Comparing group search to single-agent search, we find that the firm's profit is lower in group search for both strategies. Thus, the preference inconsistency between group members not only hurts group members' utility, but also hurts the firm's profit. This is because the firm needs to set a lower price to induce the group to search, and also because the preference inconsistency makes the group less likely to purchase. Both effects hurt the firm's profit.

More importantly, we compare the two strategies under group search. The fixed-sample strategy leads to higher optimal prices when the search cost is very small ($c < c_{GF2}$) or relatively large ($\hat{c}_0 < c < c_{GS1}$), where \hat{c}_0 denotes the cross point of the black line and the blue line in Figure B-18(b). The fixed-sample strategy also leads to higher profit than the sequential strategy when $c < c_{GF2}$, has a lower profit when $c_{GF2} < c < c_{GS1}$. Their profits become the same when $c > c_{GS1}$. We analyze the underlying intuition below.

When $c < c_{GF2}$, the group commits to searching two products under the fixed-sample strategy and the firm is able to charge the optimal price accordingly; under the sequential strategy, the group never commits to searching two products, and therefore the firm needs to lower the price to encourage the group to search two products. Thus, the optimal prices and the firm's profit is lower under the sequential strategy than under the fixed-sample strategy when $c < c_{GF2}$. When $c_{GF2} < c < c_{GS1}$, the group commits to searching only one product under the fixed-sample strategy, but is still possible to search two products under the sequential strategy, so there exists a range ($\hat{c}_0 < c < c_{GS1}$) in which the optimal prices are lower under the sequential strategy than under the fixed-sample strategy. However, the firm's profit is higher under the sequential strategy since the group may still search two products. When $c > c_{GS1}$, the group will only search one product under both strategies, so the firm sets the same prices and earns the same profit under two strategies.

To summarize, when consumers search individually, the firm always earn a higher profit under the sequential strategy, whereas when consumers search as a group and the search cost is small, the firm can earn a higher profit under the fixed-sample strategy than under the sequential strategy.

Suppose the firm can determine which search strategy consumers adopt by manipulating how it provides product offerings to consumers—i.e., asking consumers to determine how many products to search *ex ante* (fixed-sample strategy) or allowing consumers to decide whether to continue searching in the searching process (sequential strategy). Then if the consumers search and make a decision individually, the firm should always induce consumers to follow the sequential search strategy. However, for consumers who search and make a decision together as a group, the firm should ask the group to follow the fixed-sample strategy when the search cost is small, and allow them to follow the sequential strategy otherwise.

3.4 Extensions

3.4.1 Recall Not Allowed in Sequential Strategy

In the main model, we assume recall is allowed in sequential strategy. In this subsection, we consider an extension in which recall is not allowed in the sequential strategy.

In a two-period group sequential search without recall, if the group continues searching, the group will have to take the second product since they cannot get back to choose the first product, so both members' expected value of continuing search given $(X_{1A}, X_{1B}) = (x_{1A}; x_{1B})$ is $V(x_{1A}; x_{1B}) = E[X] - c$ for any value of $(x_{1A}; x_{1B})$. Following the same logic in Section 3.2.3, $\xi_{Gnr}^* = E_{X_{1j}}[V(\xi_{Gnr}^*, X_{1j})]$ still holds here, where ξ_{Gnr}^* represents the optimal stopping threshold and "nr" means "no recall." Then we can solve that the optimal threshold is $\xi_{Gnr}^* = E[X] - c = \mu - c$. Plugging in the value of ξ_{Gnr}^* and $V(x_{1A}; x_{1B})$ to the equation of expected utility as given by Equation (3.8), we can get that

$$EU_{GSnr}^* = \mu - c + \frac{d}{8} \left(1 - 4\frac{c}{d} + 3\left(\frac{c}{d}\right)^2 \right) \quad (3.11)$$

Comparing Equation (3.11) with Equation (3.4), we can get the following result.

Proposition 3. *Suppose product values are i.i.d. and follow a uniform distribution $U[\mu - d, \mu + d]$. For a two-member two-period group search, the fixed-sample strategy leads to higher expected utility than the sequential strategy (without recall) when the unit search cost is small enough ($c/d < \hat{t}_{1nr} = (\sqrt{5} - 2)/3 \approx 0.079$), or large enough ($c/d > \hat{t}_{2nr} = 1/3$).*

Proposition 2 and Proposition 3 tell us that no matter recall is allowed or not in the sequential strategy, the fixed-sample strategy can work better than the sequential strategy when search cost is small enough or large enough. Further comparing Proposition 3 with Proposition 2, we get the follow corollary For two-member two-period group search, $\hat{t}_1 < \hat{t}_{1nr}$, $\hat{t}_2 > \hat{t}_{2nr}$. That is, the fixed-sample strategy is preferable to

the sequential strategy without recall in a larger range of search cost than when it is compared to sequential strategy with recall.

This result is intuitive. When recall is not allowed in the sequential strategy, the flexibility advantage of sequential strategy is reduced because the search agents do not have the option to choose the first product any more once they decide to continue searching, and therefore the fixed-sample strategy is able to dominate the sequential strategy in a larger parameter space.

3.4.2 Alternative Distribution Assumption: Normal Distribution

In the main model, we assume product values X_{ij} are i.i.d. and follow a uniform distribution. Now we consider an extension of alternative distribution. In particular, we assume product values are i.i.d. and follow a uniform distribution $N(\mu, \sigma^2)$, which is also a commonly seen distribution assumption in search literature (e.g., [34, 106, 170]). The other assumptions are the same as in the main model.

For the fixed-sample strategy, the expected utility is $EU_{GF}(N_{GF} = 1) = \mu - c$ if the group searches one product, and is $EU_{GF}(N_{GF} = 2) = 1/2E_{max}(2) + 1/2[X] - 2c = \mu + \sigma/(2\sqrt{pi}) - 2c$, where $E_{max}(2) = \mu + \sigma/\sqrt{pi}$ based on the result of [132]. Thus, $N_{GF}^* = 2$ if $c < \sigma/(2\sqrt{pi})$, and $N_{GF}^* = 1$ otherwise. For the sequential strategy, we follow the same analysis as in Section 3.2.3, but we are not able to given a closed-form solution to the optimal stopping threshold ξ_G^* under the normal distribution. We assume product values follow the standard normal distribution, and numerically solve for the optimal stopping threshold ξ_G^* and the expected utility EU_{GS}^* .

Figure B-19 shows the difference in expected utility between the two strategies, $EU_{GF}^* - EU_{GS}^*$, under the standard normal distribution. We can see that our main result still holds: there exist two threshold \hat{c}_1, \hat{c}_2 that when the search cost is very small ($c < \hat{c}_1$) or relatively large ($c > \hat{c}_2$), the fixed-sample strategy leads to a higher expected utility than the sequential strategy. Under the standard normal distribution, $\hat{c}_1 \approx 0.102$ and $\hat{c}_2 \approx 0.818$.

3.4.3 Alternative Voting Rules

In the main model, we essentially assume the group adopts the majority rule in voting. The majority rule is a quota rule with $\alpha = 0.5$.¹⁷ In this subsection, we follow [3] and consider the quota rule with different α . In particular, we consider the unanimity rule and “one-is-enough” rule. For the sequential strategy, the unanimity rule (i.e., quota rule with $\alpha = 1$) means that the group will stop search ($N_{GS} = 1$) only when both group members agree, and will continue search ($N_{GS} = 2$) otherwise. For the fixed-sample strategy, to make it comparable to the sequential strategy, we assume that the group will search one product if and only if both members vote to search one product, and will search two products otherwise. Under the sequential strategy, the “one-is-enough” rule means that the group will stop search ($N_{GS} = 1$) as long as one group member votes to stop, and will continue search only when no group member votes to stop. Under the fixed-sample strategy, comparably, we assume that the group will search two products only when both members vote for $N_{GF} = 2$, and will search one product otherwise. For both the sequential and the fixed-sample strategies, if the group have decided to search two products and need to vote between the two searched products, we following the same voting rule as in the main model.

For the fixed-sample strategy, changing the rule for determining N_{GF} does not affect the N_{GF}^* in equilibrium, because by symmetry, the two group members will vote for the same N_{GF} in equilibrium. Thus, no matter they adopt the unanimity rule or the “one-is-enough” rule, the number of products to search in equilibrium is the same as in the main model, and then the expected utility in equilibrium will also be the same as in the main model. Below we are going to analyze the expected utility under the sequential strategy if the group adopt the two rules, respectively.

Unanimity Rule: Suppose the group takes the unanimity rule in the sequential strategy, which means that the group will stop searching only when group members vote to stop. After searching the first product, member A’s expected value of con-

¹⁷[3] assume the group adopt the “quota rule” when deciding whether to stop or continue searching, meaning that the group will stop at a candidate if at least αN members (out of N group members) agree to stop (they assume no recall for sequential search), and they consider how group members’ utility changes with α .

tinuing search $V(x_{1A}; x_{1B})$ is the same as in the main model since what may happen after continuing search is the same. We show in Appendix that the optimal stopping threshold satisfies $\xi_{GU}^* = E[V(\xi_{GU}^*, X_B) | X_B \geq \xi_{GU}^*]$ (the subscript U represents the unanimity voting rule).

“One-is-enough” Rule: Suppose the group takes the “one-is-enough” voting rule in the sequential strategy, which means that the group will stop searching as long as one group member votes to stop. Similarly as above, $V(x_{1A}; x_{1B})$ is the same as in the main model. The optimal stopping threshold satisfies $\xi_{GO}^* = E[V(\xi_{GO}^*, X_B) | X_B < \xi_{GO}^*]$ (the subscript O stands for the “one-is-enough” voting rule).

Based on the analysis above, we solve for the optimal stopping thresholds ξ_{GU}^*, ξ_{GO}^* and group members’ expected utilities EU_{GSU}^*, EU_{GSO}^* , and get the following result.

Proposition 4. *Suppose product values are i.i.d. and follow a uniform distribution $U[\mu - d, \mu + d]$. For a two-member two-period group search,*

1. *If the group take the unanimity voting rule, the fixed-sample strategy leads to a higher expected utility than the sequential strategy when the unit search cost is large enough ($c/d \in (\hat{t}_u, 3/4)$, where $\hat{t}_u \approx 0.289$), and leads to a lower expected utility when $c/d \in (0, \hat{t}_u)$.*
2. *If the group take the “one-is-enough” voting rule, the fixed-sample strategy leads to a higher expected utility than the sequential strategy when the unit search cost is small enough ($c/d \in (0, \hat{t}_o)$, where $\hat{t}_o \approx 0.044$), and leads to a lower expected utility when $c/d \in (\hat{t}_o, 1)$.*

Figure B-20 shows how the expected utility of fixed-sample strategy compares to that of sequential strategy under unanimity voting rule ($EU_{GF}^* - EU_{GSU}^*$) and under “one-is-enough” voting rule ($EU_{GF}^* - EU_{GSO}^*$), and we can clearly observe the results given by Proposition 4.

The intuition is as follows. If the group takes the unanimity rule, the group is less likely to stop since both members need to agree in order to stop. Recall that in the main model, the reason why group sequential search performs worse than group fixed-sample search when search cost is very small is because of a larger “sacrifice effect”.

Now under the unanimity rule, the expected number of products to search becomes larger than in the main model—it becomes harder for the group to stop because both members need to agree, and therefore the sacrifice effect becomes smaller. Then the advantage of fixed-sample strategy at small search cost disappears. In contrast, when the search cost is large, over-search becomes costly and group sequential search under unanimity rule becomes even more likely to over-search compared to the main model. Thus, the sequential strategy performs worse than the fixed-sample strategy when the search cost is large enough ($c/d > \hat{t}_u$), and this range is larger than in the main model ($\hat{t}_u < \hat{t}_2$).

If the group takes the “one-is-enough” rule, the group is more likely to stop because the group stops searching as long as one group member votes to stop. Then the expected number of products to search is smaller compared to the main model, meaning that the sacrifice effect becomes even larger when the search cost is small. Therefore, the fixed-sample strategy performs better than the sequential strategy when the search cost is very small ($c/d < \hat{t}_o$), and this is a larger range of search cost than in the main model ($\hat{t}_o > \hat{t}_1$). In contrast, when the search cost is large, group sequential search is less likely to over-search under the “one-is-enough” rule, and thus the commitment power of fixed-sample strategy cannot win the flexibility advantage of sequential strategy anymore in this range of search cost.

Focusing on group sequential search, [3] find that unanimity voting rule can be optimal when group members are patient enough (equivalent to small enough search cost in our setting), and the optimal number of votes needed to stop decreases as group members become less patient (equivalent to larger search cost in our setting). Our result here echoes the findings in [3], but we also take fixed-sample strategy into consideration. In our main model, we show that under the commonly used plurality rule (or majority rule), the sequential strategy performs worse than the fixed-sample strategy when the search cost is small enough or large enough due to preference inconsistency between group members, and the fixed-sample strategy mitigates the preference inconsistency problem because it can commit on the number of products to search. In this extension, we show that the sequential strategy can overcome the pref-

erence inconsistency problem by taking a stricter voting rule (the unanimity voting rule) when the search cost is small, or taking a less strict voting rule (the “one-is-enough” voting rule) when the search cost is large. In other words, if the voting rule can be endogenously determined according to the search cost, then the sequential strategy still dominates the fixed-sample strategy because it has the flexibility advantage and the preference inconsistency problem can be mitigated by choosing the proper voting rule—this can be considered as a boundary condition for our findings. However, if the voting rule is exogenously determined, the fixed-sample strategy can perform better than the sequential strategy when the search cost is small enough or large enough (under the plurality/majority voting rule), when the search cost is small enough (under the “one-is-enough” voting rule), or when the search cost is large enough (under the “unanimity” voting rule).

3.5 Concluding Remarks

In this paper, we build a two-member two-period model to show that when people search and make a decision together as a group, they can benefit from making a commitment on the number of products to search when the search cost is very small or relatively large, because the commitment mitigates the preference inconsistency problem in group search. We further consider the case in which consumers can observe product prices before search and the firm sets product prices endogenously. We find that when consumers search as a group, the firm can benefit from letting consumers make a commitment on the number of products to search when the search cost is small. We also consider several extensions to the main model to show the robustness and boundary conditions of our finding.

This paper has several limitations. First, for analytical tractability, we assume that the group has two members, and they can search for at most two periods (products). Given that recall is allowed in the sequential strategy, a model with more than two group members or more than two time periods will be very complicated and have no analytical solutions. We allow recall in the sequential strategy to make

the sequential strategy comparable to the fixed-sample strategy, and because of this assumption, we have to sacrifice the generalizability of the model to some extent. But this model can already capture the key trade-off between the fixed-sample strategy and the sequential strategy, and allow us to make our point. Second, since we need to compare the expected utility between the two search strategies, we have to make specific distribution assumptions for product values. Our extension shows that the key result also holds under the normal distribution assumption, which is also a typical distribution assumption in search literature. We want to re-emphasize that the focus of this paper is to make a point that the fixed-sample strategy can perform better than the sequential strategy under group search because of the commitment device. The exact region for the fixed-sample strategy to work better may depend on the problem setting (for example, the voting rule).

Our results have real-world implications in several dimensions. First, we provide a general guidance that a group may consider making a commitment on the search intensity if they need to search and make a decision together. Second, our results shed light on how a seller should guide consumers to search when consumers are involved in group search and purchase decisions. Last but not least, the results can help us understand some empirical facts. For example, why some empirical evidence suggests that consumers are conducting fixed-sample search, and why employers often use the fixed-sample strategy to search job candidates, as what we often observe in the academic job market.

There are several potential directions for future researches. First, researchers can further investigate the group search strategy under more general assumptions. For example, it can be interesting to know when group members should search together, and when group members can search individually first and then share their findings.¹⁸ Second, it will be interesting to further investigate firm's optimal strategies when consumers are searching as a group. The third is to empirically test whether consumers are searching as a group and test which search strategy is used in a group search setting.

¹⁸We thank anonymous reviewers for this comment.

Appendix A

Tables

Table A.1: Summary Statistics of Salesperson and Customer Traits.

Variable	Mean	# Types	Most Frequent Type
Salesperson Traits ($N = 12,149$)			
Age (When Joining Company)	38	45	37
Gender (Female = 1)	0.74	2	1
Education	-	7	High School
Home Province	-	33	Hebei
Branch Served	-	25	Hebei
Title in Company	-	7	Entry Level
Years Worked at Company	1.4	7	1
Whether Referred to Join (Referred = 1)	0.98	2	1
Whether Left Company (Left = 1)	0.1	2	0
Customer Traits ($N = 409,840$)			
Age (at Time of Transaction)	38	64	32
Gender (Female = 1)	0.58	2	1
Marital Status	-	4	Married
Occupation	-	1100	Farmer
Relationship with the Insured	-	26	Self

Table A.2: Prediction Accuracy.

Recommender System	MSE	
	(1) All Test Data	(2) Positive Sales
Deep Learning — Missing Excluded	139.84***	101.92***
Deep Learning — Linear Activation	100.91***	100.56
Deep Learning — Augmented (Our Method)	100	100
# Observations	272,138	24,169

Notes. Each observation is a salesperson-customer dyad in the test data. Significance pertains to comparison with our method. *** $p < 0.01$.

Table A.3: Recommendation Quality.

Recommender System	# Recommendations		
	1	2	3
	F1-Score		
Non-Personalized Recommender System	0.044***	0.075***	0.105***
Deep Learning — Missing Excluded	0.025***	0.044***	0.058***
Deep Learning — Linear Activation	0.074***	0.122***	0.155***
Deep Learning — Augmented (Our Method)	0.081	0.132	0.165
	NDCG		
Non-Personalized Recommender System	0.241***	0.305***	0.351***
Deep Learning — Missing Excluded	0.140***	0.187***	0.215***
Deep Learning — Linear Activation	0.371***	0.452***	0.481***
Deep Learning — Augmented (Our Method)	0.385	0.462	0.490
# Observations	2,430	2,430	2,430

Notes. Each observation is a salesperson in the test data. Significance pertains to comparison with our method. *** $p < 0.01$.

Table A.4: Recommendation Quality by Salesperson Experience (One Customer Type Recommended).

Recommender System	Years Worked at Company		
	0	1	2-6
	F1-Score		
Non-Personalized Recommender System	0.046***	0.046***	0.041***
Deep Learning — Missing Excluded	0.023***	0.025***	0.027***
Deep Learning — Linear Activation	0.081***	0.074***	0.070***
Deep Learning — Augmented (Our Method)	0.085	0.085	0.074
	NDCG		
Non-Personalized Recommender System	0.126***	0.209***	0.361***
Deep Learning — Missing Excluded	0.065***	0.107***	0.231***
Deep Learning — Linear Activation	0.220***	0.298***	0.560***
Deep Learning — Augmented (Our Method)	0.231	0.316	0.572
# Observations	671	884	875

Notes. Each observation is a salesperson in the test data. Significance pertains to comparison with our method. *** $p < 0.01$.

Table A.5: Recommendation Quality by Salesperson Experience (Multiple Customer Types Recommended).

Recommender System	Years Worked at Company		
	0	1	2-6
When Two Customer Types Are Recommended			
	F1-Score		
Non-Personalized Recommender System	0.072***	0.076***	0.076***
Deep Learning — Missing Excluded	0.038***	0.043***	0.051***
Deep Learning — Linear Activation	0.119***	0.122***	0.124***
Deep Learning — Augmented (Our Method)	0.127	0.138	0.131
	NDCG		
Non-Personalized Recommender System	0.179***	0.273***	0.434***
Deep Learning — Missing Excluded	0.095***	0.150***	0.296***
Deep Learning — Linear Activation	0.290***	0.390***	0.639***
Deep Learning — Augmented (Our Method)	0.299	0.404	0.646
When Three Customer Types Are Recommended			
	F1-Score		
Non-Personalized Recommender System	0.093***	0.105***	0.115***
Deep Learning — Missing Excluded	0.046***	0.056***	0.069***
Deep Learning — Linear Activation	0.142***	0.154***	0.167***
Deep Learning — Augmented (Our Method)	0.147	0.170	0.174
	NDCG		
Non-Personalized Recommender System	0.223***	0.321***	0.478***
Deep Learning — Missing Excluded	0.117***	0.177***	0.329***
Deep Learning — Linear Activation	0.332***	0.429***	0.648***
Deep Learning — Augmented (Our Method)	0.336	0.442	0.657
# Observations	671	884	875

Notes. Each observation is a salesperson in the test data. Significance pertains to comparison with our method. *** $p < 0.01$.

Table A.6: Importance of Salesperson and Customer Traits.

Variable	Performance Measure		
	MSE	F1	NDCG
Salesperson Traits			
Facial Image	0.79	0.011	0.022
Age (When Joining Company)	1.33	0.023	0.078
Gender (Female = 1)	0.18	0.005	0.014
Education	0.34	0.003	0.011
Home Province	1.38	0.024	0.090
Branch Served	2.02	0.012	0.034
Title in Company	0.40	0.027	0.101
Years Worked at Company	1.09	0.005	0.010
Whether Referred to Join (Referred = 1)	0.00	0.000	0.000
Whether Left Company (Left = 1)	2.69	0.022	0.074
Customer Traits			
Age (at Time of Transaction)	1.94	0.023	0.067
Gender (Female = 1)	0.14	0.016	0.050
Marital Status	0.80	0.042	0.155
Occupation	3.97	0.054	0.194
Relationship with the Insured	1.17	0.010	0.040

Table A.7: Most Recommended Customer Type by Salesperson Trait.

Salesperson Trait	Trait Level	Most Recommended Customer Type
Facial Image	Group 0	Above 40 + Female + Married + Service personnel + Self-insured
	Group 1	Below 40 + Male + Married + Managers + Self-insured
Age (When Joining Company)	Below 40	Below 40 + Male + Married + Managers + Self-insured
	Above 40	Above 40 + Female + Married + Service personnel + Self-insured
Gender	Male	Below 40 + Male + Married + Managers + Self-insured
	Female	Below 40 + Female + Married + Service personnel + Self-insured
Education	Below College	Above 40 + Female + Married + Service personnel + Self-insured
	Above College	Below 40 + Male + Married + Managers + Self-insured
Home Province	Hebei	Below 40 + Female + Married + Farmers + Other-insured
	Jiangxi	Above 40 + Female + Married + Service personnel + Self-insured
Branch Served	Hebei	Below 40 + Female + Married + Farmers + Other-insured
	Jiangsu	Below 40 + Female + Married + Managers + Self-insured
Title in Company	Entry Level	Below 40 + Male + Married + Managers + Self-insured
	Above Entry Level	Below 40 + Female + Married + Managers + Self-insured
Years Worked at Company	Below One Year	Below 40 + Male + Married + Managers + Self-insured
	Above One Year	Above 40 + Female + Married + Service personnel + Self-insured
Whether Referred to Join	Not Referred	Below 40 + Female + Married + Managers + Self-insured
	Referred	Below 40 + Female + Married + Managers + Self-insured
Whether Left the Company	Not Left	Below 40 + Female + Married + Service personnel + Self-insured
	Left	Below 40 + Male + Married + Managers + Self-insured

Table A.8: Extensions to the Augmented Recommender System.

Recommender System	Performance Measure		
	MSE	F1-Score	NDCG
Deep Learning — Augmented (Main Model)	100	0.081	0.385
Deep Learning — Probabilistically Missing Sales (Extension)	99.82	0.083***	0.397***
Deep Learning — Heterogeneous Cost of Selling (Extension)	99.12***	0.087***	0.412**
# Observations	272,138	2,430	2,430

Notes. For MSE, each observation is a salesperson-customer dyad in the test data. For the F1-score and NDCG, each observation is a salesperson in the test data. Significance pertains to comparison with the main model. *** $p < 0.01$, ** $p < 0.05$.

Table A.9: Proportion of New Salespeople Who Have Not Made Their First Sale.

Recommender System	Days Since Salesperson Joined Company					
	15	30	45	60	75	90
	One Customer Type Recommended					
Random Search	0.84***	0.78***	0.75***	0.73***	0.71***	0.70***
Non-Personalized Recommender System	0.71***	0.67***	0.65***	0.64***	0.62***	0.61***
Deep Learning — Missing Excluded	0.81***	0.75***	0.72***	0.71***	0.68***	0.67***
Deep Learning — Linear Activation	0.63***	0.59***	0.56***	0.55***	0.53***	0.52***
Deep Learning — Augmented (Our Method)	0.62	0.58	0.55	0.54	0.52	0.51
	Two Customer Types Recommended					
Random Search	0.84***	0.78***	0.75***	0.73***	0.71***	0.70***
Non-Personalized Recommender System	0.68***	0.64***	0.62***	0.60***	0.59***	0.58***
Deep Learning — Missing Excluded	0.81***	0.74***	0.71***	0.70***	0.67***	0.66***
Deep Learning — Linear Activation	0.55***	0.50***	0.48***	0.47***	0.46***	0.45***
Deep Learning — Augmented (Our Method)	0.52	0.49	0.47	0.46	0.45	0.45
	Three Customer Types Recommended					
Random Search	0.84***	0.78***	0.75***	0.73***	0.71***	0.70***
Non-Personalized Recommender System	0.61***	0.57***	0.56***	0.54***	0.53***	0.52***
Deep Learning — Missing Excluded	0.81***	0.74***	0.70***	0.70***	0.67***	0.65***
Deep Learning — Linear Activation	0.49***	0.46***	0.44***	0.43***	0.42***	0.41***
Deep Learning — Augmented (Our Method)	0.46	0.43	0.42	0.41	0.40	0.40
# Observations	2,430	2,430	2,430	2,430	2,430	2,430

Notes. Each observation is a salesperson in the test data. Significance pertains to comparison with our method. *** $p < 0.01$.

Table A.10: Constraints in Recent Targeted Marketing Papers.

Paper	Problem	Constraint	Optimization Algorithm
[22]	Solicitations for charity	No Constraints	-
[112]	Proactive retention campaigns	Budget Constraint	Greedy
[153]	Promotions to prospective customers	Budget Constraint	Greedy
[182]	Ranking for query-based search	No Constraints	-
[59]	Pricing for a digital firm	No Constraints	-
[73]	Coupons for retailer customers	Budget Constraint	Greedy
[179]	Targeted discounts to retain customers	No Constraints	-

Table A.11: Segments, Markets and Scenarios.

	State A	State B	Both States
	1,061,438 Customers	1,308,658 Customers	2,370,096 Customers
Number of Segments			
3-digit Zip	10	14	24
4-digit Zip	53	77	130
5-digit Zip	208	264	472
Scenario			
3-digit Zip	S1	S4	S7
4-digit Zip	S2	S5	S8
5-digit Zip	S3	S6	S9

Table A.12: Algorithms' Performance Comparison.

Problem Setting	Algorithm			
	Primal Simplex	Dual Simplex	Barrier	Our Algorithm
	Feasibility			
S1 (State A + 3 digits)	Yes	Yes	Yes	Yes
S2 (State A + 4 digits)	No	Yes	Yes	Yes
S4 (State B + 3 digits)	Yes	Yes	Yes	Yes
S5 (State B + 4 digits)	No	No	No	Yes
S7 (All + 3 digits)	No	Yes	Yes	Yes
	Computation Time (in seconds)			
S1 (State A + 3 digits)	10992	474	378	211
S2 (State A + 4 digits)	-	24700	13104	1984
S4 (State B + 3 digits)	35513	1697	766	3204
S5 (State B + 4 digits)	-	-	-	126253
S7 (All + 3 digits)	-	56823	24700	15317
	Optimal Profits			
S1 (State A + 3 digits)	100	100	100	100
S2 (State A + 4 digits)	-	100	100	100
S4 (State B + 3 digits)	100	100	100	100
S5 (State B + 4 digits)	-	-	-	100
S7 (All + 3 digits)	-	100	100	100

Table A.13: Summary Statistics of Targeting Variables.

Variable	Mean	Median	Standard Deviation
Single Family	0.808	1.000	0.394
Multi Family	0.188	0.000	0.391
Member Tier	5.365	5.000	2.476
Child	0.231	0.000	0.422
Female	0.317	0.000	0.465
Male	0.658	1.000	0.474
Home Value Tier	4.071	2.000	3.506
Family Number	2.491	2.000	1.691
Length of Residence	11.94	8.000	11.68
Income (in 1,000s)	63.83	50.00	52.46
Age	50.03	50.00	16.96
Age Type	0.824	1.000	0.381
Homeowner	0.663	1.000	0.473
Renter	0.243	0.000	0.429
Condo Owner	0.033	0.000	0.179
Residential	0.682	1.000	0.466
Condominium	0.305	0.000	0.172
Duplex	0.025	0.000	0.157
Apartment	0.002	0.000	0.047
Agricultural	0.009	0.000	0.096
Mobile Homes	0.025	0.000	0.155
Distance	10.84	8.095	8.183
Comp. Distance	9.240	6.433	7.664
3yr Response	0.121	0.994	0.205
Penetration Rate	0.093	0.069	0.071
F Flag	0.590	1.000	0.492
M Flag	0.280	0.000	0.449

Appendix B

Figures

Figure B-1: A Typical Neural Network Structure.




Figure B-2: Proposed Activation Function (Example of $c = 1$).




Figure B-3: Our Proposed Neural Network Structure.

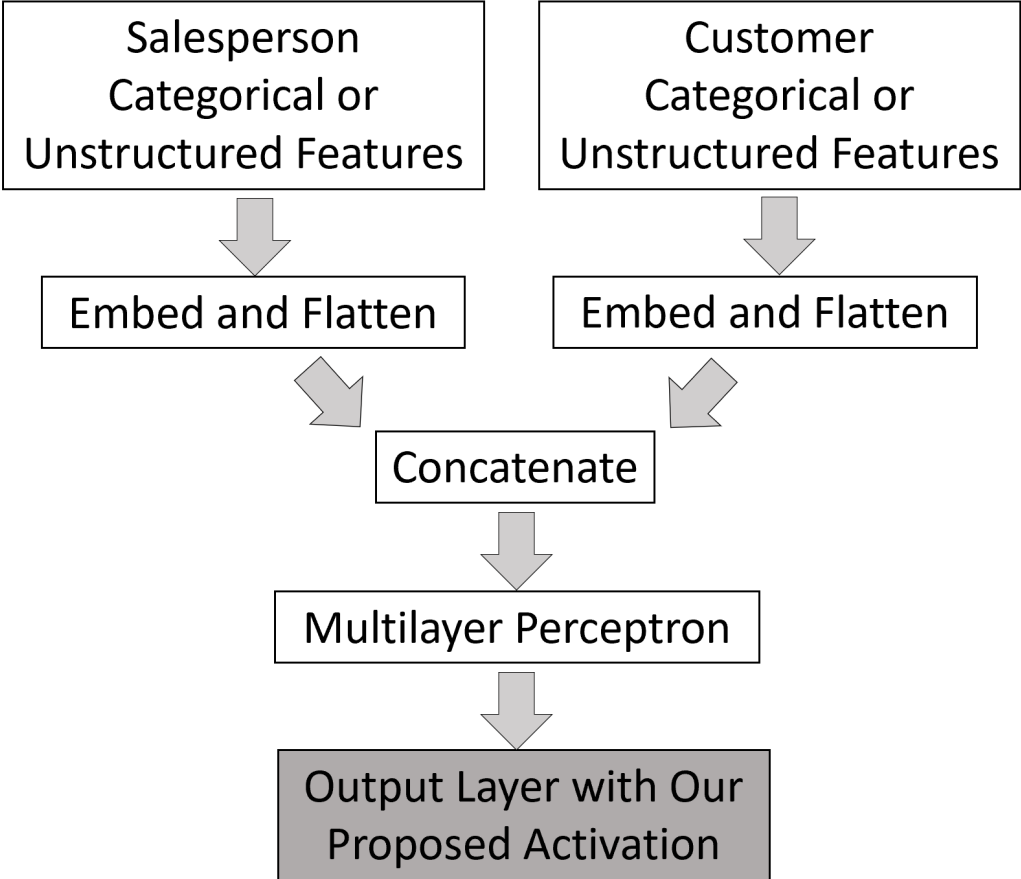


Figure B-4: Illustration of Face Embedding Derivation




Figure B-6: Training and Validation Loss: Missing Values Excluded.




Figure B-7: Training and Validation Loss: Linear Activation.




Figure B-8: Training and Validation Loss: Augmented.




Figure B-9: Proportion of New Salespeople Who Have Not Made Their First Sale.

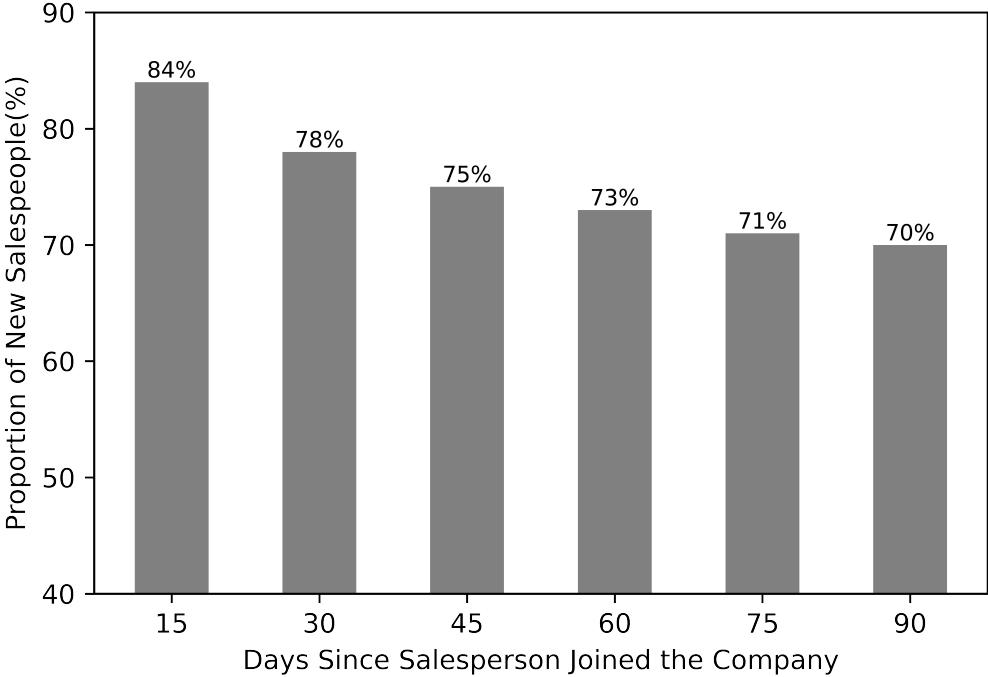


Figure B-10: Distribution of Customer Types for Each Customer Trait.

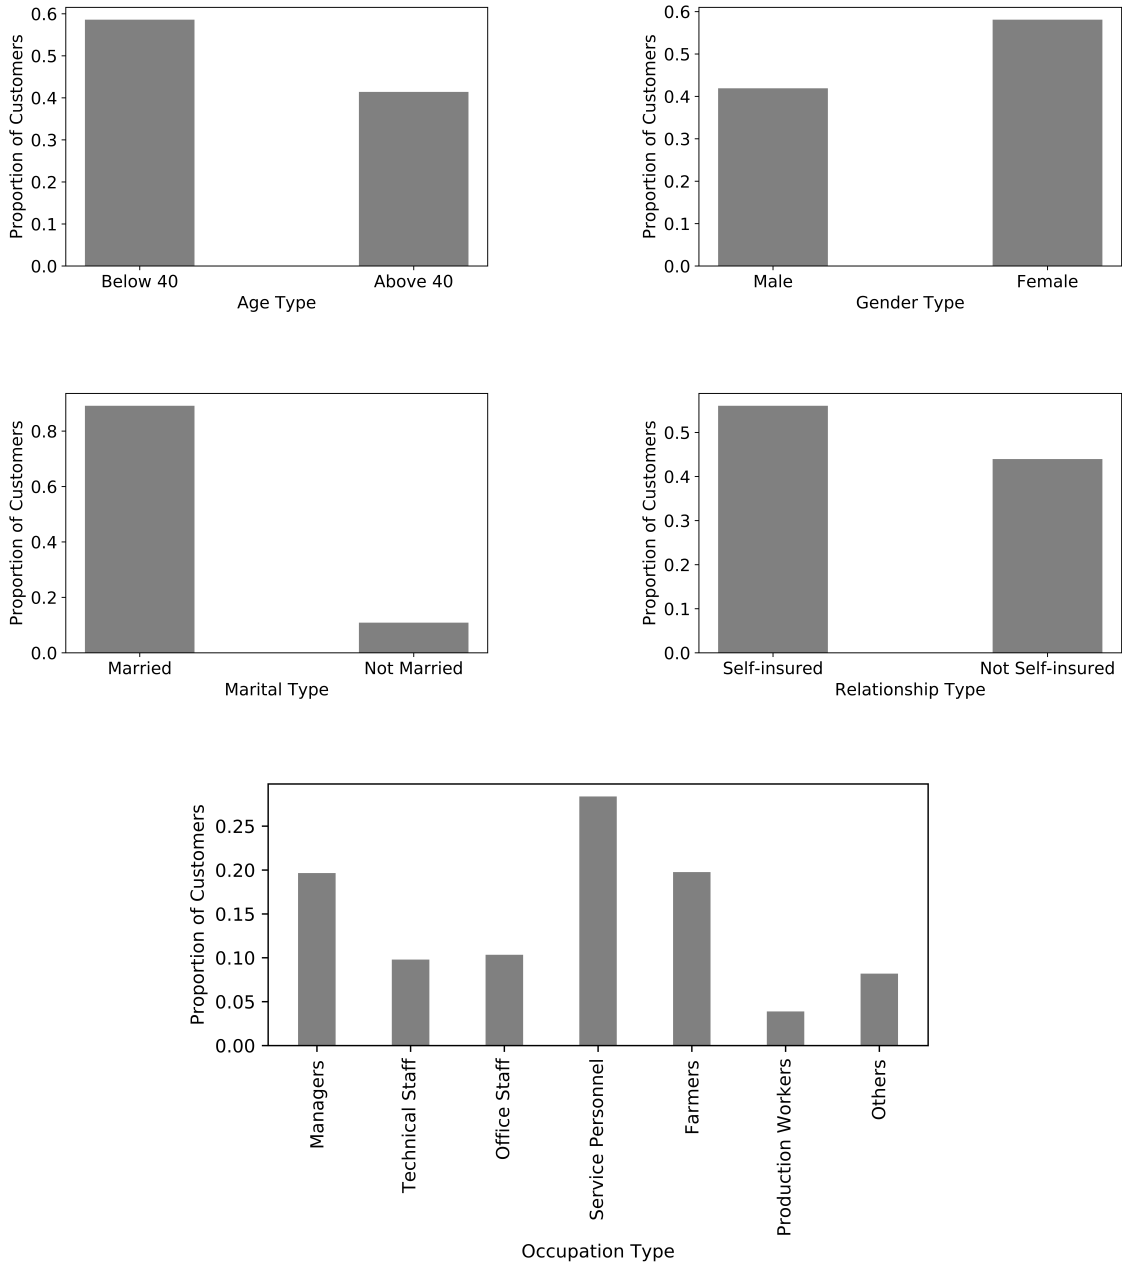


Figure B-11: Most Salespeople Served Just A Few Customer Types.




Figure B-12: Percentage of Salespeople Who Served Each Customer Type.

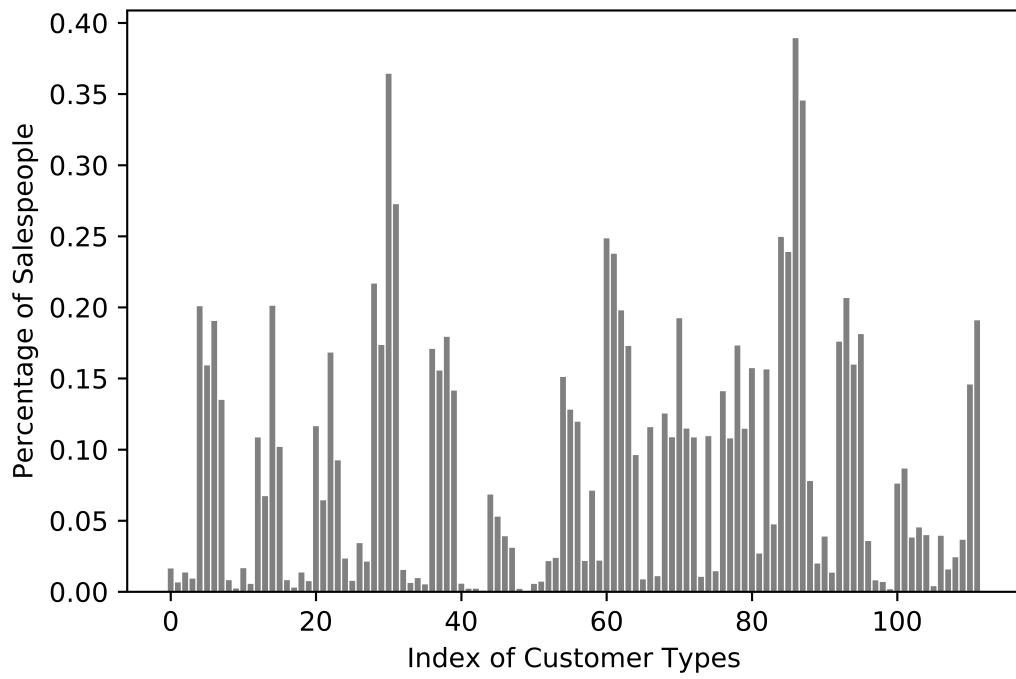
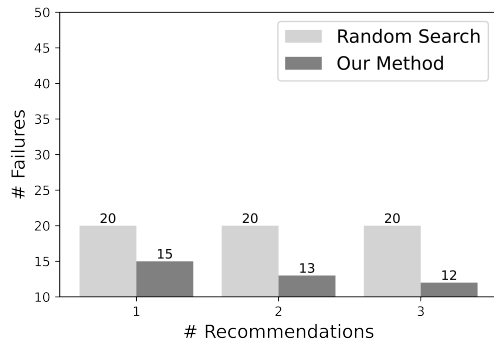
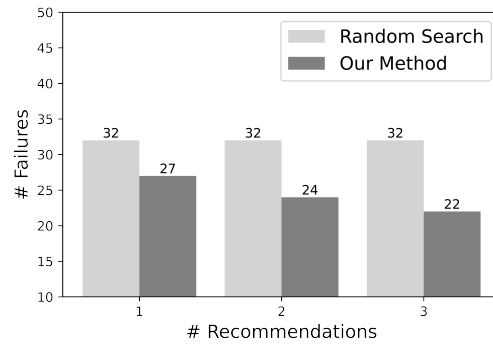


Figure B-13: Number of Failures before the First Sale.



(a) All Salespeople



(b) New Salespeople

Figure B-14: Proportion of New Salespeople Who Have Not Made Their First Sale, Revisited.

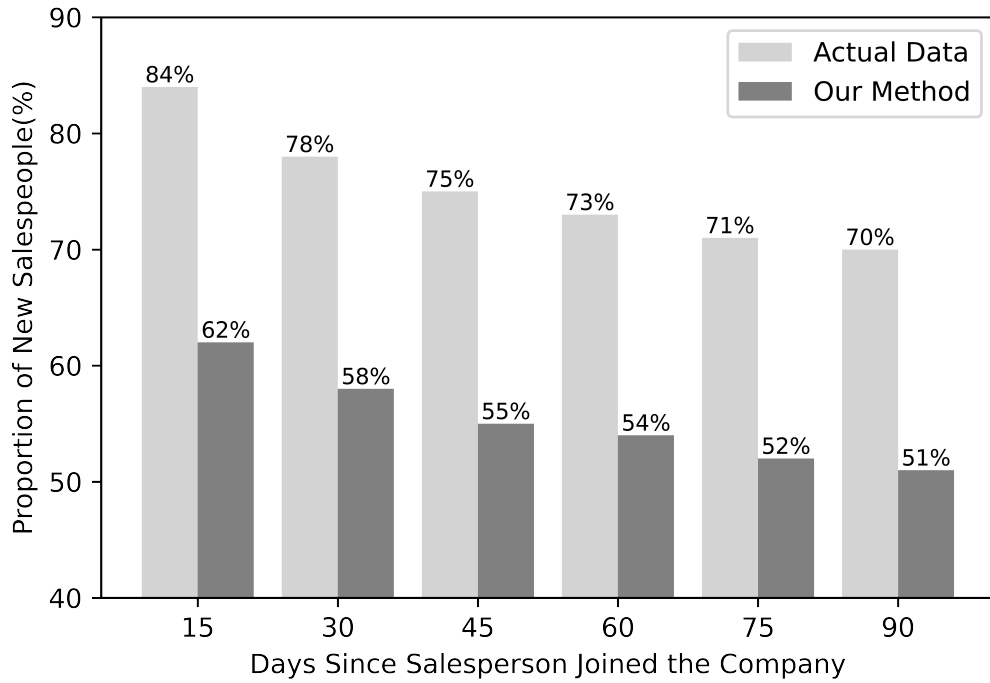


Figure B-15: Illustration of Differences between one-loop and two-loop iterations.




Figure B-16: Illustration of Proposition 2: Comparing Expected Utilities under the Two Strategies.




Figure B-17: Illustration of Lemma 3.2.4 and Lemma 3.2.4.

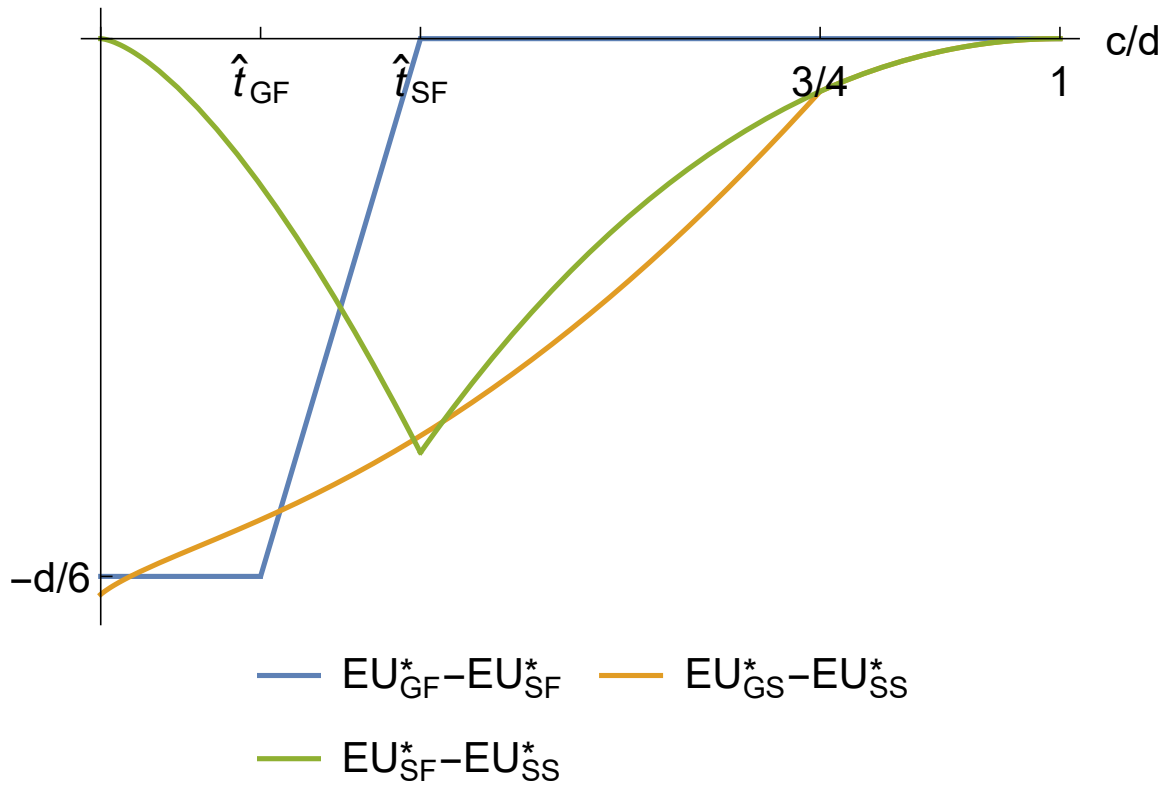


Figure B-18: Optimal Prices and Profit under Single Agent and Group Search.

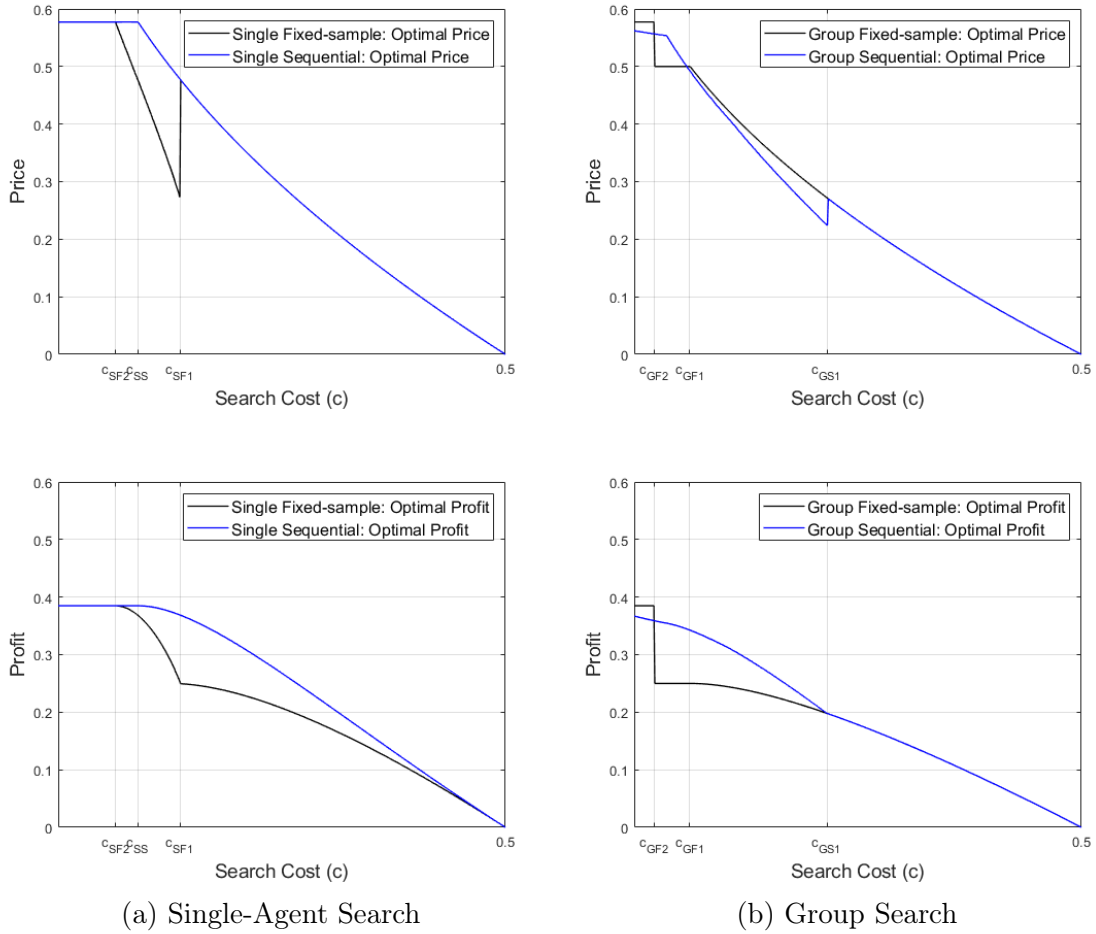


Figure B-19: Comparing Expected Utilities of the Two Strategies under Standard Normal Distribution Assumption.


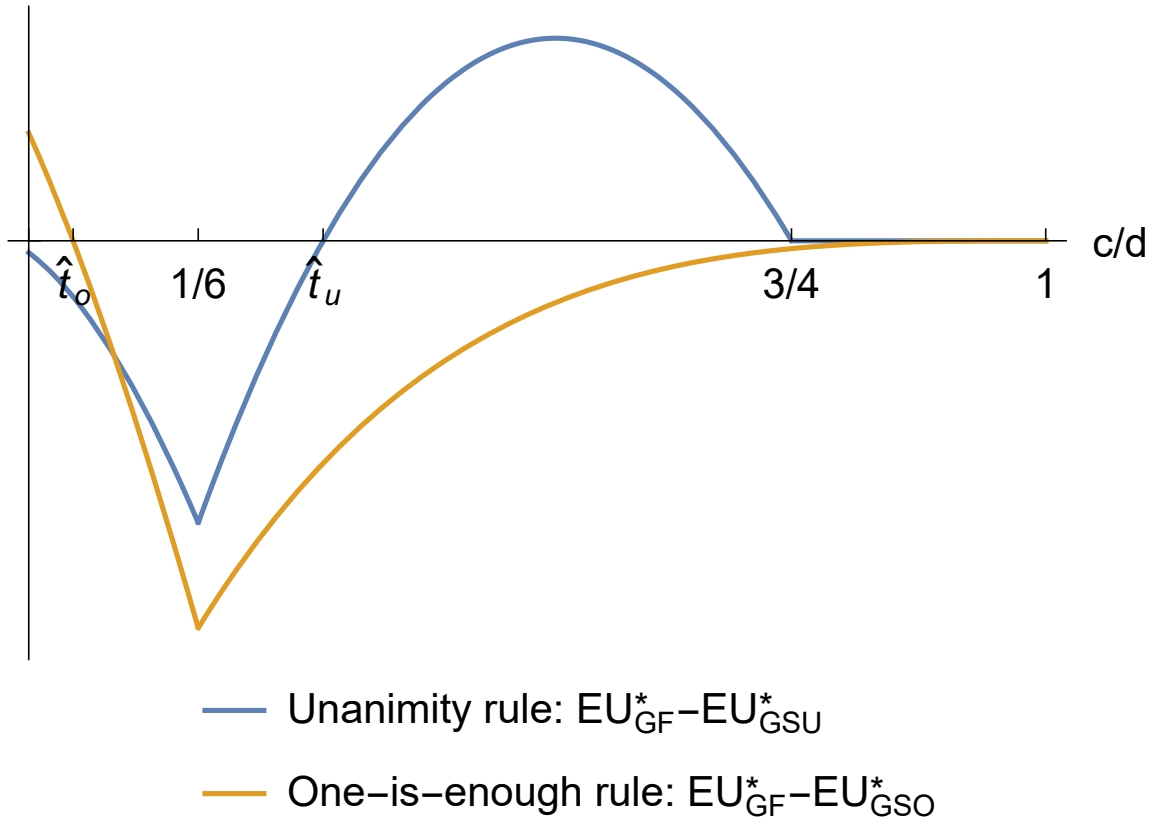


Figure B-20: Compare Fixed-Sample Strategy to Sequential strategy under Alternative Voting Rules.



Appendix C

Chapter 1 Appendix

C.1 Details on the Augmented Recommender System

We implement the augmented recommender system using the Keras functional API. The main components of the hidden layers in our augmented recommender system are the embedding layer and the MLP.

The embedding of the categorical features is implemented through Keras Embedding and the output embedding dimension is set as $\min\{50, m/2\}$, where m is the number of categories per feature ([98]). We derive the face embeddings using the dlib library.¹ The output is a 128-dimensional vector for each face. The pre-trained network is a ResNet network with 29 convolutional layers, which is based on [91] with fewer layers and filters. The network is trained on the Labeled Faces in the Wild (LFW) dataset, which includes around 3 million images.² We illustrate the derivation of face embeddings in Figure B-4. The embedding vector itself does not carry interpretable information. But similar faces have smaller distances in the embedding space.

Our MLP includes the now-common ReLU activation function in each layer ([67]). We also use Dropout at each layer to prevent neural networks from overfitting and set

¹<http://blog.dlib.net/2017/02/high-quality-face-recognition-with-deep.html>.

²<http://vis-www.cs.umass.edu/lfw>.

the fraction of dropout at 0.5 ([156]). We use stochastic gradient descent and early stopping ([142]) to obtain the number of epochs. The number of hidden layers is set at two ([184]), while the number of units in each layer is derived through cross-validation.

For the “Deep Learning — Missing Excluded” benchmark, we minimize its validation loss (MSE) on the validation data that have positive sales records. This replicates the construction of the optimal neural network underlying this benchmark.

Figure B-5 summarizes the neural network structure of our augmented recommender system, the “Deep Learning — Augmented” model. The tenth input in Figure B-5 is our facial embeddings derived using the dlib library.

While using the “Deep Learning — Augmented” method, we also need to fine-tune the net cost of selling, c . We first obtain the lowest and highest predicted values of sales revenue generated from the MLP. We evenly split the interval between these two values to ten blocks and use the eleven endpoints as starting values of c to train the neural network. We obtain the c that generates the smallest MSE on the validation data. We then take the two endpoints surrounding this c , use them to define a new interval, evenly split this new interval to ten blocks, and use the eleven new endpoints as new values of c to train the neural network. We iterate this process until convergence, defined as the MSE on the validation data ceasing to decrease or the optimal c ceasing to change (difference smaller than 0.00001).

Figures B-6, B-7, and B-8 respectively summarize the training and validation loss across epochs for the “Deep Learning — Missing Excluded,” “Deep Learning — Linear Activation,” and “Deep Learning — Augmented” methods. We scale the training loss at epoch 1 to 100, following our confidentiality agreement with the company.

C.2 Recommended Customer Type by Salesperson Trait

We present a simplified illustration of the output of our augmented deep learning based recommender system. The actual output is at the individual-salesperson level.

To simplify presentation, we group salespeople based on their traits. We divide each of the 10 salesperson traits into two levels. Specifically, for the non-binary traits, we divide salespeople’s age when joining the company to below or weakly above 40, education to below or weakly above college, title in the company to entry level or above, and years working at company to weakly below or above one year. For salespeople’s home province and branch served, we focus on the two most frequent types for either. These are Hebei and Jiangxi for salespeople’s home province and Hebei and Jiangsu for branch served. Last, for facial image, we split the salespeople into two groups using K-means clustering based on their facial embedding vector values. This yields 20 scenarios (10 salesperson traits and two levels per trait). For each of the 20 scenarios, we output the most recommended customer type using our augmented recommender system. Table A.7 presents the results.

C.3 Proportion of New Salespeople Yet to Make Their First Sale

Figure B-9 of the paper presents the proportion of new salespeople who have not made their first sale as time goes by. We examine how a recommender system changes this proportion. For each recommender system, we vary the number of recommended customer types from one to three. Section 1.7 of the paper describes the calculation in detail. We use block bootstrapping to test the statistical significance of differences between recommender systems. Table A.9 presents the results. In all cases, our method outperforms all benchmark recommender systems significantly ($p < 0.01$).

Appendix D

Chapter 2 Appendix

D.1 Details of Algorithm 1

Let's first define the mapping from (x, y) to z . Recall that $\eta > 0$ denotes the primal step-size and $\tau > 0$ the dual stepsize. Let's introduce two new parameters $\omega > 0$ and $\sigma > 0$ such that $\eta = \sigma/\omega$ and $\tau = \omega\sigma$. We call ω the primal weight and σ the stepsize. Then we have:

$$\|z\| = \sqrt{\omega\|x\|_2^2 + \frac{\|y\|_2^2}{\omega}}.$$

To under the normalized duality gap decay condition, let's first define the normalized duality gap. The normalized duality gap with radius R at $v = (x, y)$ is defined as

$$\rho_R(v) = \frac{1}{R} \max_{\hat{v} \in \{v \in V: \|v - \hat{v}\| \leq R\}} \{L(x, \hat{y}) - L(\hat{x}, y)\},$$

where $L(x, y)$ is the primal-dual form for a general primal-dual problem $\min_{x \in X} \max_{y \in Y} L(x, y)$ and $\hat{v} = (\hat{x}, \hat{y})$. The normalized duality gap decay condition referred in Algorithm 1 is defined as:

$$\rho_{\|\bar{z}^{n,t} - z^{n,0}\|}(\bar{z}^{n,t}) \leq 0.5\rho_{\|z^{n,0} - z^{n-1,0}\|}(z^{n,0}).$$

Lastly, the algorithm convergence is defined as follows. Suppose Z^* is the set of optimized solutions. Define $\mathbf{dist}(z, Z^*) = \inf_{\hat{z} \in Z^*} \|z - \hat{z}\|$ and a tolerance level ϵ . Then the algorithm converges when $\mathbf{dist}(z, Z^*) \leq \epsilon$.

D.2 Profit Estimation and Prediction

The outcome for a customer as a function of the marketing action report the targeting variables we use to estimate the prediction model used to predict profit for each household under each marketing action and control condition. To increase the precision of the profit prediction, we use all households in the two states ($N = 2,370,096$) in this step. We use O_i to denote all of these covariates for household i , W_i^j denotes the treated condition for marketing action $j = 1, \dots, 5$ and household i , and P_i^j the profit for marketing action $j = 0, 1, \dots, 5$ and household i . $j = 0$ indicates the control condition. Notice that for each household i , at most one W_i^j can be equal to one. Households allocated to the control condition will have all W_i^j to be zero. We use P_i^{obs} to denote the observed profit for household i , which is the achieved profit under the assigned treatment condition.

There exists many models we can use to achieve the profit prediction. Comparing the differences between different models is not the goal of our paper, and we choose to use LASSO with a full set of interactions to predict the profit given recommendations from literature (e.g., [13, 153]). We first use all observed data and LASSO to estimate the following model:

$$P_i^{obs} = f(W_i^j, O_i) + \epsilon_i = \alpha + \beta W_i^j + \gamma O_i + \delta W_i^j O_i + \epsilon_i,$$

where $W_i^j O_i$ denotes a full interaction of treatment W_i^j with all covariates O_i . LASSO can help to select which covariates are important. Once we have an estimated model $\hat{f}(W_i^j, O_i), j = 1, \dots, 5$, we can then derive predicted profit for each household under each marketing action $\hat{P}_i^j = \hat{f}(W_i^j = 1, W_i^{j'} = 0, O_i)$ for $j = 0, 1, \dots, 5$ and $j' \neq j$. This \hat{P}_i^j will be the r_i^j in Equation 2.1. We simplify the estimation and prediction step to focus on the optimization step for targeting with constraints problem. In Section 3.5, we discuss several interesting directions to think about the interaction between prediction and optimization.

Other parameters in the constraints in Equation (2.1) are normally determined by the firm given the needs. In Section 2.5.3, we offer one set of parameters in the

constraints that make the problem feasible and non-trivial to validate the performance of our method against the state-of-the-art solvers.

Table A.13 reports the summary statistics for all of the targeting variables we use in the estimation and prediction model. The meaning of all variables are as follows. *Single (Multi) Family* is a binary flag indicating whether the home is a single (multi) family home. *Member Tier* is a tier assigned to each household by the retailer. Lower tier indicates higher value. There are in total 10 tiers and we use binary flags for the first 9 tiers in profit prediction. *Child* is a binary flag indicating whether the household has a child. *Female (Male)* is a binary flag indicating whether the head of the household is a female (male). There exists households that we don't know the gender of the head. *Home Value Tier* is an offered tier classification for estimated household value. Higher tier implies higher value and households that do not have estimated home value are in tier 11. There are in total 11 tiers and we use binary flags for the first 10 tiers in profit prediction. *Family Number* is the number of persons in the household. *Length of Residence* is the length of residence in current home. *Income* is estimated household income. *Age* is the age of head of household. *Age Type* is a binary flag indicating whether the age is estimated. *Homeowner (Renter or Condo Owner)* is a binary flag indicating whether the household is a homeowner (renter or condo owner). *Residential (Condominium or Duplex or Apartment or Agricultural or Mobile Homes)* is a binary flag indicating the property type. *Distance (Comp. Distance)* is the distance to nearest store for this retailer (competitors' store). *3yr Response* is the carrier-route level average response rate to marketing activities of this retailer in the last three years. *Penetration Rate* is the percentage of households that are members in a ZIP code area. *F (M) Flag* is a binary flag indicating whether the retailer considers a ZIP code as "far" ("medium) from the retailer's closet store.

Appendix E

Chapter 3 Appendix

E.1 Proof of $\xi_G^* = E_{X_{1B}}[V(\xi_G^*, X_{1B})]$ in Group Sequential Search

When $X_{1A} = \xi_G$, if A votes to stop, her expected value is $\Pr(X_{1B} \geq \xi'_G) * \xi_G + \Pr(X_{1B} < \xi'_G) * (\xi_G/2 + E[V(\xi_G, X_{1B}/2)|X_{1B} < \xi'_G])$, and if A votes to continue, her expected value is $\Pr(X_{1B} \geq \xi'_G) * (\xi_G/2 + E[V(\xi_G, X_{1B})|X_{1B} \geq \xi'_G]/2) + \Pr(X_{1B} < \xi'_G) * E[V(\xi_G, X_{1B})|X_{1B} < \xi'_G]$. At the optimal stopping threshold, her expected value of voting to stop should be equal to her expected value of voting to continue. Rearranging the equation, we can get $\xi_G^* = \Pr(X_{1B} \geq \xi'_G) * E[V(\xi_G, X_{1B})|X_{1B} \geq \xi'_G] + \Pr(X_{1B} < \xi'_G) * E[V(\xi_G, X_{1B})|X_{1B} < \xi'_G] = E_{X_{1B}}[V(\xi_G^*, X_{1B})]$. We notice that this equation does not depend on the value of member B's stopping threshold ξ'_G .

E.2 Proof of Proposition 2:

Fixed-Sample vs. Sequential in Group Search

First, let's look at the difference between the **expected utility** under the fixed-sample strategy versus under the sequential strategy. Denote $t = c/d$, and $g(t) \equiv (EU_{GF}^* - EU_{GS}^*)/d = (2 + 24t - \sqrt{3 + 8t} - 20t\sqrt{3 + 8t})/48$ when $t \in [0, 1/6)$, $g(t) = (-6 + 72t - \sqrt{3 + 8t} - 20t\sqrt{3 + 8t})/48$ when $t \in [1/6, 3/4)$, and $g(t) = 0$ when

$t \in [3/4, 1]$. We notice that $g(0) = d/48(2 - \sqrt{3}) > 0$, $g(t)$ strictly decreases in t on $t \in [0, 1/6)$,¹ and $g(1/6) = d(6 - 13\sqrt{13/3}/3)/48 < 0$. Thus, there exists a unique $\hat{t}_1 \in (0, 1/6)$ such that $g(\hat{t}_1) = 0$ and $g(t) > 0$ for any $t \in (0, \hat{t}_1)$. We solve that $\hat{t}_1 = 1/400(-9 + \sqrt{281}) \approx 0.019$. On $t \in [1/6, 3/4)$, $g(t)$ first increases then decreases in t ,² and $h(3/4) = 0$. There exists a unique \hat{t}_2 such that $h(\hat{t}_2) = 0$ and $h(t) > 0$ for any $t \in (\hat{t}_2, 1)$. We solve that $\hat{t}_2 = (79 + 3\sqrt{449})/400 \approx 0.356$.

E.3 Compare the Expected Value and Search Cost in Fixed-Sample vs. Sequential in Group Search

First, we compare the **expected search cost** under the two strategies.

Based on the analysis in Section 3.2.2, under the fixed-sample strategy, the search cost is $2c$ when $c/d < 1/6$, and c when $c/d \geq 1/6$. Based on the analysis in Section 3.2.3, the expected search cost under the sequential strategy is $c(5 - \sqrt{3 + 8c/d})/2$ for $c/d \leq 3/4$, and is c when $c/d \in (3/4, 1]$. It is easy to see that $c < c(5 - \sqrt{3 + 8c/d})/2 < 2c$.

Second, we compare the **expected value** of the selected product under the two strategies.

The expected value of the selected product under the fixed-sample strategy is $\mu + d * t/6$ when $t < 1/6$ and μ when $t \geq 1/6$. The expected value of the selected product under the sequential strategy can be calculated using the expected utility plus the expected search cost, which equals $\mu + d(6 - 4t\sqrt{3 + 8t} + \sqrt{3 + 8t})$ for $t \in [0, 3/4]$, and is μ for $t \in (3/4, 1]$. Then, the difference in the expected value of the selected products between the fixed-sample strategy and the sequential strategy is $DV = d(2 + 4t\sqrt{3 + 8t} - \sqrt{3 + 8t})/48$ when $t \leq 1/6$, $DV = d(-6 + 4t\sqrt{3 + 8t} - \sqrt{3 + 8t})/48$ when $t \in (1/6, 3/4]$, and $DV = 0$ for $t \in (3/4, 1]$. It is easy to verify that, within each region, DV increases with t (the expected value for fixed-sample strategy is constant

¹On $t \in [0, 1/6)$, $g'(t) = 1/48(24 - 4/\sqrt{3 + 8t} - (80t)/\sqrt{3 + 8t} - 20\sqrt{3 + 8t}) < 1/48(24 - 20\sqrt{3}) < 0$.

² $g'(t)$ is positive when $1/6 < t < 1/150(14 + 3\sqrt{519}) \approx 0.549$, and is negative when $1/150(14 + 3\sqrt{519}) < t \leq 1$.

within each region, and the expected value for sequential strategy decreases with c/d). When $t = 0$, $DV = d(2 - \sqrt{3})/48 > 0$, so $DV > 0$ for $t \in [0, 1/6)$. When $t = 3/4$, $DV = 0$, so $DV < 0$ for $t \in [1/6, 3/4]$.

E.4 Proof of Lemma 3.2.4:

Fixed-Sample vs. Sequential in Single-Agent Search

For the problem of two-period single-agent search, all the assumptions are the same as in our main model for group search except that there is only one decision maker.

Under the fixed-sample strategy, the decision maker chooses the number of products to search $n \in \{1, 2\}$ to maximize her expected utility $EU_{SF}(n) = E_{\max}(n) - cn$. Comparing $EU_{SF}(1) = \mu - c$ and $EU_{SF}(2) = \mu + d/3 - 2c$, we get that the optimal number of products search is $N_{SF}^* = 1$ when $c/d \geq 1/3$ and $N_{SF}^* = 2$ when $c/d < 1/3$. Then the expected utility is $EU_{SF}^* \equiv EU_{SF}(N_{SF}^*) = \mu - c$ when $c/d \geq 1/3$ and $EU_{SF}^* = \mu + d/3 - 2c$ when $c/d < 1/3$. We denote $\hat{t}_{SF} = 1/3$, which means the cutoff value of c/d that single-agent fixed-sample search changes from searching two products to searching one product. Under the sequential strategy, the optimal stopping threshold should be equal to the expected value of continuing search, i.e., $\xi_S^* = -c + (1 - F(\xi_S^*))E[X|X \geq \xi_S^*] + F(\xi_S^*)\xi_S^*$. Solving this equation, we get the optimal stopping threshold $\xi_S^* = \mu + d - 2\sqrt{cd}$. Then the expected utility is $EU_{SS}(\xi_S^*) = \int_{\mu-d}^{\xi_S^*} V(x_1)f(x_1) dx_1 + \int_{\xi_S^*}^{\mu+d} x_1f(x_1) dx_1 - c$, where $V(x_1)$ means the expected value of continuing search when the first draw is x_1 . We can calculate that $V(x_1) = -c + \int_{\mu-d}^{x_1} x_1f(x_2) dx_2 + \int_{x_1}^{\mu+d} x_2f(x_2) dx_2 = -c + [\mu^2 + 2\mu(d - x_1) + (x_1 + d)^2]/(4d)$. Plugging this into $EU_{SS}(\xi_S^*)$, we get that $EU_{SS}^* \equiv EU_{SS}(\xi_S^*) = \mu + d/3 - 2c + 2/3 * c\sqrt{c/d}$. The expected utilities under the fixed-sample strategy and under the sequential strategy have been shown in the paper. We now compare the expected utilities under these two strategies. When $c/d \in (0, 1/3)$, $EU_{SS}^* - EU_{SF}^* = 2/3 * c\sqrt{c/d} > 0$. When $c/d \in [1/3, 1)$, $EU_{SS}^* - EU_{SF}^* = d/3(1 - 3c/d + 2c/d\sqrt{c/d}) = d/3(1 - \sqrt{c/d})^2(2\sqrt{c/d} + 1) > 0$. Thus, we get the first claim in Lemma 3.2.4: for a two-period single-agent search, the

sequential strategy (with recall) always dominates the fixed-sample strategy.

Given the formula of $EU_{SS}^* - EU_{SF}^*$, it is easy to see that $EU_{SS}^* - EU_{SF}^*$ equals 0 when $c = 0$, increases in c when $c \in (0, d/3)$, decreases in c when $c \in (d/3, d)$,³ and becomes 0 again when $c = d$. This is the fourth claim in Lemma 3.2.4.

Let's compare the expected search cost under the fixed-sample strategy versus under the sequential strategy. Under the sequential strategy, the number of products to search is $N_{SS}^* = 1$ with probability $\sqrt{c/d}$, and $N_{SS}^* = 2$ with probability $1 - \sqrt{c/d}$. Thus, the expected search cost is $2c(1 - \sqrt{c/d}) + c\sqrt{c/d} = 2c - c\sqrt{c/d}$. Under the fixed-sample strategy, the number of products to search is $N_{SF}^* = 1$ when $c/d \geq 1/3$, and $N_{SF}^* = 2$ when $c/d < 1/3$. Thus, the search cost is c when $c/d \geq 1/3$ and $2c$ when $c/d < 1/3$. Therefore, the difference in the expected search cost under the fixed-sample strategy versus under the sequential strategy is $c - (2c - c\sqrt{c/d}) = c(\sqrt{c/d} - 1) < 0$ when $c/d \in (1/3, 1)$, and $2c - (2c - c\sqrt{c/d}) = c\sqrt{c/d} > 0$ when $c/d \in (0, 1/3)$. This is the third claim in Lemma 3.2.4.

Lastly, we compare the expected value of the selected product under the fixed-sample strategy versus under the sequential strategy. Under the sequential strategy, the expected value of the selected product is $\mu + d/3 - 2c + 2/3 * c\sqrt{c/d} + (2c - c\sqrt{c/d}) = \mu + d/3 - c\sqrt{c/d}/3$. Under the fixed-sample strategy, the expected value of the selected product is μ when $c/d \geq 1/3$ and $\mu + d/3$ when $c/d < 1/3$. Therefore, the difference in the expected values of the selected product under the two strategies is $\mu - (\mu + d/3 - c\sqrt{c/d}/3) = d/3 * (c/d\sqrt{c/d} - 1) < 0$ when $c/d \in (1/3, 1)$ and $\mu + d/3 - (\mu + d/3 - c\sqrt{c/d}/3) = c\sqrt{c/d}/3 > 0$ when $c/d \in (0, 1/3)$. This is the second claim in Lemma 3.2.4.

E.5 Proof of Lemma 3.2.4: Group Search vs. Single-Agent Search

Sequential strategy

³Denote $\sqrt{c/d} = t$, $EU_{SS}^* - EU_{SF}^* = d/3(1 - 3t^2 + 2t^3) \equiv g(t)$. $g'(t) = -2dt(1 - t) < 0$ when $t \in (1/3, 1)$.

The expected number of searched products in single-agent search is $E[N_{SS}^*] = 1 + \Pr(X < \xi_S^*) = 2 - \sqrt{t}$ (denote $c/d = t$) for any $t \in [0, 1]$. The expected number of searched products in group search is $E[N_{GS}^*] = 1 + \Pr(X_{1A} < \xi_S^*) * \Pr(X_{1B} < \xi_S^*) + \Pr(X_{1A} \geq \xi_S^*) * \Pr(X_{1B} < \xi_S^*)/2 + \Pr(X_{1A} < \xi_S^*) * \Pr(X_{1B} \geq \xi_S^*)/2 = (5 - \sqrt{3 + 8t})/2$ when $t \in [0, 3/4]$, and $N_{GS}^* = 1$ when $t \in [3/4, 1]$. For $t \in [0, 3/4]$, $E[N_{SS}^*] - E[N_{GS}^*] = (\sqrt{3 + 8t} - 2\sqrt{t} - 1) > 0^4$. For $t \in [3/4, 1]$, $E[N_{SS}^*] - E[N_{GS}^*] = 1 - \sqrt{t} \geq 0$ (the “=” holds only when $t = 1$).

The expected utility in single-agent search is $EU_{SS}^* = \mu + d*(1/3 - 2t + 2/3*t\sqrt{t})$. The expected utility in group search is $EU_{GS}^* = \mu + d*(1/8 + \sqrt{3 + 8t}/48 + 5t(-6 + \sqrt{3 + 8t})/12)$ for $t \in [0, 3/4]$, and $EU_{GS}^* = \mu - d*t$ for $t \in [3/4, 1]$. For $t \in [0, 3/4]$, $EU_{SS}^* - EU_{GS}^* = d/48*[10 - \sqrt{3 + 8t} + 4t(6 + 8\sqrt{t} - 5\sqrt{3 + 8t})] \equiv h_1(t)$. It is easy to show that $6 + 8\sqrt{t} - 5\sqrt{3 + 8t} < 0$ and it decreases in t , and then we know that $h_1(t)$ decreases in t . For $t \in [3/4, 1]$, $EU_{SS}^* - EU_{GS}^* = d*(1/3 - t + 2/3*t\sqrt{t}) \equiv h_2(t)$. $h_2'(t) = d(-1 + \sqrt{t}) < 0$, so $h_2(t)$ decreases in t . $h_1(3/4) = h_2(3/4) = d/12(-5 + 3\sqrt{3}) > 0$, and $h_2(1) = 0$, so $EU_{SS}^* - EU_{GS}^*$ decreases in t and $EU_{SS}^* - EU_{GS}^* > 0$ for all $t \in [0, 1]$. We have proved the claims under the sequential strategy.

Fixed-sample strategy

For group search, $N_{GF}^* = 2$ when $c/d \in (0, \hat{t}_{GF})$ and $N_{GF}^* = 1$ when $c/d \in [\hat{t}_{GF}, 1)$, where $\hat{t}_{GF} = 1/6$. For single-agent search, $N_{SF}^* = 2$ when $c/d \in (0, \hat{t}_{SF})$ and $N_{SF}^* = 1$ when $c/d \in [\hat{t}_{SF}, 1)$, where $\hat{t}_{SF} = 1/3$. Then it is easy to see that $N_{GF}^* < N_{SF}^*$ when $c/d \in [\hat{t}_{GF}, \hat{t}_{SF})$, and the same number of searched products in other region of search cost.

$EU_{GF}^* = \mu + d/6 - 2c$ if $c/d \in (0, \hat{t}_{GF})$ and $EU_{GF}^* = \mu - c$ if $c/d \in [\hat{t}_{GF}, 1)$. $EU_{SF}^* = \mu + d/3 - 2c$ if $c/d \in (0, \hat{t}_{SF})$ and $EU_{SF}^* = \mu - c$ if $c/d \in [\hat{t}_{SF}, 1)$. When $c/d \in (0, \hat{t}_{GF})$, $EU_{GF}^* - EU_{SF}^* = -d/6 < 0$; when $c/d \in [\hat{t}_{GF}, \hat{t}_{SF})$, $EU_{GF}^* - EU_{SF}^* = c - d/3 < 0$; when $c/d \in [\hat{t}_{SF}, 1)$, $EU_{GF}^* - EU_{SF}^* = 0$. So we have proved the claims under the fixed-sample strategy.

⁴ $3 + 8t - (2\sqrt{t} + 1)^2 = 2[(1 - \sqrt{t})^2 + t] > 0$

E.6 Firm's Strategy

E.6.1 Fixed-sample strategy under single-agent search

Suppose the consumer searches one product only. Her expected utility is $E[\max\{X_1 - p_1, 0\}] - c = (1 - p_1)^2/2 - c$. The firm's profit is $\pi_{SF1}(p_1) = p_1 \Pr(X_1 - p_1 \geq 0) = p_1(1 - p_1)$, which is optimized at $p_1 = 1/2$ and $\pi_{SF1}^* = 1/4$ correspondingly.

Suppose the consumer searches two products. Her expected utility is $E[\max\{X_1 - p_1, X_2 - p_2, 0\}] - 2c = -(p_2 - p_1)^3/6 + (p_2 - p_1)^2/2 - (1 - p_1^2)(p_2 - p_1)/2 + (1 - p_1)^2(2 + p_1)/3 - 2c$.⁵ The firm's profit is $\pi_{SF2}(p_1, p_2) = p_1 \Pr(X_1 - p_1 \geq X_2 - p_2 \wedge X_1 - p_1 \geq 0) + p_2 \Pr(X_2 - p_2 \geq X_1 - p_1 \wedge X_2 - p_2 \geq 0) = p_1(1 + 2(p_2 - p_1)/2 - (p_2 - p_1)^2 - 2(p_2 - p_1)p_1 - p_1^2) + p_2(1 - 2(p_2 - p_1)/2 + (p_2 - p_1)^2 - p_1^2)$, which is optimized at $p_1 = p_2 = 1/\sqrt{3}$ and $\pi_{SF2}^* = 2/(3\sqrt{3})$ correspondingly.

The consumer will search two products if and only if $c \leq (1 + 3p_1 - p_2)(1 - p_2)^2/6 \equiv c_{SF2}(p_1, p_2)$, will search one product if $c_{SF2}(p_1, p_2) < c \leq c_{SF1}(p_1) = (1 - p_1)^2/2$, and will not search any product if $c > c_{SF1}(p_1)$.

We notice that $\pi_{SF2}^* > \pi_{SF1}^*$, meaning that without other constraints, it is more profitable for the firm to let the consumer search two products. This is intuitive—the firm does not bear the search cost, and the consumer is more likely to find a product value higher than the price when searching two products.

Combining the analysis above, we get the following four scenarios:

1. When $c < c_{SF2} = 1/(9\sqrt{3})$, the firm charges $p_1^* = p_2^* = 1/\sqrt{3}$ and the consumer searches two products. The firm's profit is $\pi_2^* = 2/(3\sqrt{3})$.
2. When $c > c_{SF2}$, the firm adjusts p_1, p_2 to make $c \leq c_{SF2}(p_1, p_2) = (1 + 3p_1 - p_2)(1 - p_2)^2/6$ so that the consumer still searches two products, and maximize $\pi_{SF2}(p_1, p_2) = p_1/2 * (1 + 2(p_2 - p_1) - (p_2 - p_1)^2 - 2(p_2 - p_1)p_1 - p_1^2) + p_2/2 * (1 - 2(p_2 - p_1) + (p_2 - p_1)^2 - p_1^2)$ at the same time, until the search cost c is large enough such that $\pi_{SF2}(p_1^*, p_2^*)$ cannot exceed $1/4$.

⁵When deriving these equations, we implicitly assume that $p_1, p_2 \in [0, 1]$. Otherwise the consumer has no incentive to search the products at all, and the demand must be 0.

3. Then the firm should set $p_1^* = 1/2$ (p_2^* can be any value higher than $1/2$) and the consumer searches one product only. The firm's profit is $\pi_1^* = 1/4$.
4. When $c > c_{SF1} = 1/8$, the consumer does not want to search one product at $p_1^* = 1/2$, and the firm should set $p_1^* = 1 - \sqrt{2c}$ (p_2^* can be any value higher than $1 - \sqrt{2c}$), so that the consumer is on the boundary of being willing to search one product ($c = c_{SF1}(p_1^*) = (1 - p_1^*)^2$) and the firm extracts the highest profit possible, which equals $\pi_1^* = \sqrt{2c}(1 - \sqrt{2c})$. Notice that $c \leq 1/2$. Otherwise the firm cannot earn positive profit.

Scenario 2 cannot be solved analytically. Thus, we numerically solve Scenario 2, and plot the how the optimal prices and the firm's profit changes with search cost c in Figure B-18(a) (black line). The numerical solution indicates that it is still optimal for the firm to set $p_1^* = p_2^*$ in Scenario 2, and there is no range of c that satisfies Scenario 3.

To summarize, when $c < c_{SF2}$, the consumer searches two products, $p_1^* = p_2^* = 1/\sqrt{3}$ and the profit is $\pi_{SF}^* = 2/(3\sqrt{3})$ (Scenario 1). When $c_{SF2} < c < c_{SF1}$, the firm lowers the prices to induce the consumer keeps searching two products (Scenario 2). When $c > c_{SF1}$, the firm sets $p_1^* = 1 - \sqrt{2c}$ (and p_2^* can be any value greater than or equal to $1 - \sqrt{2c}$), the consumer searches one product only, and the firm's profit is $\pi_{SF}^* = \sqrt{2c}(1 - \sqrt{2c})$ (Scenario 4).

E.6.2 Sequential strategy under single-agent search

Suppose the consumer will stop searching if and only if the first product's value $X_1 \geq \xi_S^*$. At the stopping threshold ξ_S^* , the value to stop should be equal to the value to continue, so ξ_S^* should satisfy the following equation $(\xi_S^* - p_1)^+ = -c + \Pr(X_2 - p_2 > (\xi_S^* - p_1)^+)E[X_2 - p_2 | X_2 - p_2 \geq (\xi_S^* - p_1)^+] + \Pr(X_2 - p_2 \leq (\xi_S^* - p_1)^+) * (\xi_S^* - p_1)^+$ where $(\xi_S^* - p_1)^+ = \max\{\xi_S^* - p_1, 0\}$.

We solve for ξ_S^* by considering the following two scenarios, which are differentiated by whether the value of $F(\xi_S^* + p_2 - p_1)$.

1. If $p_2 - p_1 + \xi_S^* \geq 1$, then $\xi_S^* - p_1 \geq 1 - p_2 \geq 0$ and $F(\xi_S^* + p_2 - p_1) = 1$. Given $F(\xi_S^* + p_2 - p_1) = 1$, there is actually no ξ_S^* that can satisfy equation (3.10), because the consumer will never choose to continue searching at the cutoff. In other words, we can consider the stopping threshold as $\xi_S^* = 0$. It means that the consumer will stop searching for sure after searching the first product. Then $D_1 = \Pr(X_1 \geq p_1) = 1 - p_1$ and $D_2 = 0$. The firm chooses p_1 to maximize $\Pi = p_1(1 - p_1)$, subject to the constraint that $EU = E[X] - p_1 - c = 1/2 - p_1 - c \geq 0$. Then the optimal price is $p_1^* = 1/2 - c$. p_2^* can be any value that satisfies $p_2^* \geq 1 + p_1^* = 3/2 - c$.

2. If $p_2 - p_1 + \xi_S^* < 1$, then $F(\xi_S^* + p_2 - p_1) = \xi_S^* + p_2 - p_1$. We further classify the following two subcases:

(a) If $\xi_S^* \geq p_1$, then from equation (3.10) we can get $\xi_S^* = 1 - \sqrt{2c} + p_1 - p_2$.

This leads to $p_2 \leq 1 - \sqrt{2c}$. We can get that $D_1 = \Pr(X_1 \geq \xi_S^* \wedge X_1 - p_1 \geq 0) + \Pr(X_1 < \xi_S^* \wedge X_1 - p_1 \geq X_2 - p_2 \wedge X_1 - p_1 \geq 0) = \sqrt{2c} - p_1 + p_2 + (1 - 2\sqrt{2c} + 2c - p_2^2)/2$ and $D_2 = \Pr(X_1 < \xi_S^* \wedge X_1 - p_1 < X_2 - p_2 \wedge X_2 - p_2 \geq 0) = -c + (1 + 2p_1 - p_2)(1 - p_2)/2$. The firm chooses p_1, p_2 to maximize $p_1 D_1 + p_2 D_2$, subject to the constraint that $p_1 \leq p_2 \leq 1 - \sqrt{2c}$.

(b) If $\xi_S^* \leq p_1$, then from equation (3.10) we can get $\xi_S^* = p_1 - \sqrt{(1 - p_2)^2 - 2c}$.

From $(1 - p_2)^2 - 2c \geq 0$ we get that $p_2 \leq 1 - \sqrt{2c}$. The purchase probabilities are $D_1 = 1 - p_1$ and $D_2 = (1 - p_2)(p_1 - \sqrt{1 - 2c - 2p_2 + p_2^2})$. The firm chooses p_1, p_2 to maximize $p_1 D_1 + p_2 D_2$, subject to the constraint that $p_1 \leq p_2 \leq 1 - \sqrt{2c}$.

We solve for the optimal p_1^*, p_2^* in Scenario 2(a) and 2(b) numerically, and then compare the firm's profit under the three scenarios. We find that Scenario 1 and Scenario 2(b) are in fact never optimal. Figure B-18(a) (blue line) plots the optimal prices and the firm's profit. For all search cost $c \in [0, 1/2]$, the optimal case is Scenario 2(a), and the two products' optimal prices are equal, i.e., $p_1^* = p_2^*$. When $c \leq c_{SS} = 2/3 - 1/\sqrt{3}$, the firm can optimize its profit by choosing the inner solution, i.e., the optimal prices are $p_1^* = p_2^* = 1/\sqrt{3}$ and the firm's profit is $\pi_{SS}^* = 2/(3\sqrt{3})$.

When $c > c_{SS}$, the firm has to lower the price to satisfy the boundary constraint, i.e., the optimal prices are $p_1^* = p_2^* = 1 - \sqrt{2c}$ and the firm's profit is $\pi_{SS}^* = (1 - \sqrt{2c})(2\sqrt{2c} - 2c)$. On the entire range, the stopping threshold is $\xi_S^* = 1 - \sqrt{2c}$.

E.6.3 Fixed-sample strategy under group search

Suppose the group searches one product. Each member j votes to buy the product if $X_{1j} - p_1 \geq 0$, votes not to buy otherwise. If the two members disagree on whether to buy or not to buy, each of their choices has equal probability to be the group's choice. We can calculate the expected utility from member A's perspective, and member B's expected utility will be the same: $EU_{GF}(N_{GF} = 1) = -c + \Pr(X_{1A} - p_1 \geq 0 \wedge X_{1B} - p_1 \geq 0) * E[X_{1A} - p_1 | X_{1A} - p_1 \geq 0] + \Pr(X_{1A} - p_1 \geq 0 \wedge X_{1B} - p_1 < 0) E[X_{1A} - p_1 | X_{1A} - p_1 \geq 0] / 2 + \Pr(X_{1A} - p_1 < 0 \wedge X_{1B} - p_1 \geq 0) E[X_{1A} - p_1 | X_{1A} - p_1 < 0] / 2 = -c + (1 - p_1)(2 - 3p_1) / 4$ The firm's expected profit is $\pi_{GF1}(p_1) = p_1 * [\Pr(X_{1A} - p_1 \geq 0 \wedge X_{1B} - p_1 \geq 0 + \Pr(X_{1A} - p_1 \geq 0 \wedge X_{1B} - p_1 < 0) / 2) + \Pr(X_{1A} - p_1 < 0 \wedge X_{1B} - p_1 \geq 0) / 2] = p_1 * [(1 - p_1)^2 + (1 - p_1) * p_1] = p_1(1 - p_1)$ Without other constraints, $\pi_{GF1}(p_1)$ is maximized at $p_1 = 1/2$, and the corresponding profit is $\pi_{GF1}^* = 1/4$.

If the group searches two products, each member has three options: votes to buy product 1, votes to buy product 2, and votes not to buy either. Since the two members' product values' are independent, their choices are also independent. For each member j , she will vote for product 1 with probability $q_1 = \Pr(X_{1j} - p_1 \geq X_{2j} - p_2 \wedge X_{1j} - p_1 \geq 0) = 1/2 - p_1 + p_2 - p_2^2/2$, she will vote for product 2 with probability $q_2 = \Pr(X_{2j} - p_2 \geq X_{1j} - p_1 \wedge X_{2j} - p_2 \geq 0) = (1 + 2p_1 - p_2)(1 - p_2) / 2$, and she will vote for neither product with probability $q_0 = \Pr(X_{1j} - p_1 < 0 \wedge X_{2j} - p_2 < 0) = p_1 p_2$. We can check that $q_1 + q_2 + q_0 = 1$. Each member's expected utility is (we omit member subscript j for simplicity, now the subscript refers to the product) $EU_{GF}(N_{GF} = 2) = -2c + (q_1^2 + q_1 q_2 / 2 + q_1 q_0 / 2) * E[X_1 - p_1 | X_1 - p_1 \geq X_2 - p_2 \wedge X_1 - p_1 \geq 0] + (q_1 q_2 / 2) * E[X_2 - p_2 | X_1 - p_1 \geq X_2 - p_2 \wedge X_1 - p_1 \geq 0] + (q_1 q_2 / 2) * E[X_1 - p_1 | X_1 - p_1 < X_2 - p_2 \wedge X_2 - p_2 \geq 0] + (q_1 q_2 / 2 + q_2^2 + q_2 q_0 / 2) * E[X_2 - p_2 | X_1 - p_1 < X_2 - p_2 \wedge X_2 - p_2 \geq 0] + (q_0 q_1 / 2) * E[X_1 - p_1 | X_1 - p_1 < 0 \wedge X_2 - p_2 < 0] + (q_0 q_2 / 2) * E[X_2 - p_2 | X_1 - p_1 < 0 \wedge X_2 - p_2 < 0] = -2c + [9p_1^2 + 3p_1(-2 - 7p_2 + 4p_2^2) + 7 - 6p_2 + 9p_2^2 - 4p_2^3] / 12$. The

firm's expected profit is $\pi_{GF2}(p_1, p_2) = p_1 * (q_1^2 + q_1q_2/2 + q_1q_0/2 + q_1q_2/2 + q_0q_1/2) + p_2 * (q_1q_2/2 + q_1q_2/2 + q_2^2 + q_2q_0/2 + q_0q_2/2) = [-2p_1^2 + p_1(1 + 4p_2 - 3p_2^2) + (1 - p_2)^2p_2]/2$. Without other constraints, $\pi_{GF2}(p_1, p_2)$ is maximized at $p_1 = p_2 = 1/\sqrt{3}$, and the corresponding profit is $\pi_{GF2}^* = 2/(3\sqrt{3})$. We can see that given the consumer(s) search two products, the firm's profit under group search is the same as that under single-agent search.

Comparing group members' expected utility, we get that the group searches two products when $c \leq c_{GF2}(p_1, p_2) = (1 - p_2)(1 + 9p_1 - 5p_2 - 12p_1p_2 + 4p_2^2)/12$, searches one product when $c_{GF2}(p_1, p_2) < c \leq c_{GF1}(p_1) = (1 - p_1)(2 - 3p_1)/4$, and does not search any product when $c > c_{GF1}(p_1)$.

Then we have the following four scenarios: (denote $c_{GF2} = c_{GF2}(1/\sqrt{3}, 1/\sqrt{3}) = -1/4 + 17/36\sqrt{3} \approx 0.023$, and $c_{GF1} = c_{GF1}(1/2) = 1/16 = 0.0625$)

1. When $c < c_{GF2}$, the firm charges $p_1^* = p_2^* = 1/\sqrt{3}$, the group searches both products, and the firm's expected profit is $\pi_{GF2}^* = 2/(3\sqrt{3})$.
2. When $c > c_{GF2}$, the firm adjusts p_1, p_2 to make $c_{GF2}(p_1, p_2) \geq c$ and maximizes $\pi_{GF2}(p_1, p_2)$, until c is too large such that $\pi_{GF2}(p_1, p_2) \leq 1/4$.
3. Then it is no longer profitable for the firm to induce the group to search two products. The firm should set $p_1^* = 1/2$ (p_2 can be any value higher than or equal to that), the group searches one product only, and the firm's expected profit is $\pi_{GF1}^* = 1/4$.
4. When $c > c_{GF1}$, the group is no longer willing to search at $p_1 = 1/2$, so the firm should lower the price accordingly to make $c = c_{GF1}(p_1)$, which leads to $p_1^* = (5 - \sqrt{1 + 48c})/6$, and $\pi_{GF1}^* = (1 - 12c + \sqrt{1 + 48c})/9$. Notice that $c \leq 1/2$, otherwise $p_1^* < 0$.

We solve scenario 2 numerically, and compare the firm's profit between Scenario 2 and 3. Figure B-18(b) (black line) plots the optimal prices and firm's profit. The result indicates that, Scenario 2 is never optimal. When $c < c_{GF2}$, the firm charges $p_1^* = p_2^* = 1/\sqrt{3}$, the group searches two products, and the firm's profit is $\pi_{GF}^* =$

$2/(3\sqrt{3})$ (Scenario 1). When $c > c_{GF2}$, the group searches only one product. When $c_{GF2} < c < c_{GF1}$, the firm sets $p_1^* = 1/2$ and earns profit $\pi_{GF}^* = 1/4$ (Scenario 3), and when $c > c_{GF1}$, the firms set $p_1^* = (5 - \sqrt{1 + 48c})/6$ and earns profit $\pi_{GF}^* = (1 - 12c + \sqrt{1 + 48c})/9$ (Scenario 4). In the latter two cases, p_2^* can be any value higher than or equal to p_1^* .

E.6.4 Sequential strategy under group search

To write down the demand for the two products, we define the following ten events and their corresponding possibility: $q_1 = \Pr(X_{1j} \geq \xi_G^* \wedge X_{1j} - p_1 \geq 0)$, $q_2 = \Pr(X_{1j} \geq \xi_G^* \wedge X_{1j} - p_1 < 0)$, $q_3 = \Pr(X_{1j} < \xi_G^* \wedge X_{1j} - p_1 \geq 0)$, $q_4 = \Pr(X_{1j} < \xi_G^* \wedge X_{1j} - p_1 < 0)$, $q_5 = \Pr(X_{1j} \geq \xi_G^* \wedge X_{1j} - p_1 \geq 0 \wedge X_{1j} - p_1 \geq X_{2j} - p_2)$, $q_6 = \Pr(X_{1j} \geq \xi_G^* \wedge X_{2j} - p_2 \geq 0 \wedge X_{2j} - p_2 \geq X_{1j} - p_1)$, $q_7 = \Pr(X_{1j} \geq \xi_G^* \wedge X_{1j} - p_1 < 0 \wedge X_{2j} - p_1 < 0)$, $q_8 = \Pr(X_{1j} < \xi_G^* \wedge X_{1j} - p_1 \geq 0 \wedge X_{1j} - p_1 \geq X_{2j} - p_2)$, $q_9 = \Pr(X_{1j} < \xi_G^* \wedge X_{2j} - p_2 \geq 0 \wedge X_{2j} - p_2 \geq X_{1j} - p_1)$, $q_{10} = \Pr(X_{1j} < \xi_G^* \wedge X_{1j} - p_1 < 0 \wedge X_{2j} - p_2 < 0)$. The first four events belong to the case in which the group searches only one product, and the last six events belong to the case in which the group searches two products. We can then write the demand for the two products as: $D_1 = q_1[q_1 + q_2/2 + q_3/2 + q_4/4] + q_2[q_1/2 + q_3/4] + q_3[q_1/2 + q_2/4] + q_4q_1/4 + q_5[q_8/2 + q_9/4 + q_{10}/4] + q_6q_8/4 + q_7q_8/4 + q_8[q_5/2 + q_6/4 + q_7 + q_8/4 + q_9/2 + q_{10}/2] + (q_9 + q_{10})[q_5/4 + q_8/2]$ and $D_2 = q_5q_9/4 + q_6[q_8/4 + q_9/2 + q_{10}/4] + q_7q_9/4 + q_8[q_6/4 + q_9/2] + q_9[q_5/4 + q_6/2 + q_7/4 + q_8/2 + q_9 + q_{10}/2] + q_{10}[q_6/4 + q_9/2]$.

We solve ξ_G^* by considering the following scenarios. If $\xi_G^* - p_1 \geq 0$, then when $X_{1A} = \xi_G^*$ and A votes to stop, her expected value is $q_1 * (\xi_G^* - p_1) + q_3 * ((\xi_G^* - p_1)/2 + E[V(\xi_G^*, X_{1B})|p_1 \leq X_{1B} < \xi_G^*]/2) + q_4 * ((\xi_G^* - p_1)/4 + E[V(\xi_G^*, X_{1B})|X_{1B} < p_1]/2)$; when A votes to continue, her expected value is $q_1 * ((\xi_G^* - p_1)/2 + E[V(\xi_G^*, X_{1B})|X_{1B} \geq \xi_G^*]/2) + q_3E[V(\xi_G^*, X_{1B})|p_1 \leq X_{1B} < \xi_G^*] + q_4E[V(\xi_G^*, X_{1B})|X_{1B} < p_1]$. To write down the continuation value, let's define the following quantities: $pv_1(x) = \Pr(X_{2j} - p_2 \geq 0 \wedge X_{2j} - p_2 \geq x - p_1)$, $pv_2(x) = \Pr(X_{2j} - p_2 < 0 \wedge X_{2j} - p_2 \geq x - p_1)$. Then if $x_{1B} \geq p_1$, we have $V(\xi_G^*, x_{1B}) = -c + F(\xi_G^* + p_2 - p_1)F(x_{1B} + p_2 - p_1)(\xi_G^* - p_1) + pv_1(\xi_G^*) * pv_1(x_{1B})E[X_{2A} - p_2|X_{2A} \geq \xi_G^* + p_2 - p_1 \wedge X_{2A} - p_2 \geq 0] + pv_1(\xi_G^*) * pv_2(x_{1B})(E[X_{2A} - p_2|X_{2A} \geq \xi_G^* + p_2 - p_1 \wedge X_{2A} - p_2 \geq 0]/2) + pv_2(\xi_G^*) * pv_1(x_{1B})(E[X_{2A} - p_2|X_{2A} \geq$

$\xi_G^* + p_2 - p_1 \wedge X_{2A} - p_2 < 0]/2) + pv_1(\xi_G^*) * F(x_{1B} + p_2 - p_1)((\xi_G^* - p_1)/2 + E[X_{2A} - p_2 | X_{2A} \geq \xi_G^* + p_2 - p_1 \wedge X_{2A} - p_2 \geq 0]/2) + pv_2(\xi_G^*) * F(x_{1B} + p_2 - p_1)((\xi_G^* - p_1)/2) + F(\xi_G^* + p_2 - p_1) * (1 - F(x_{1B} + p_2 - p_1))((\xi_G^* - p_1)/2 + E[X_{2A} - p_2 | X_{2A} < \xi_G^* + p_2 - p_1]/2)$. If $x_{1B} < p_1$, $V(\xi_G^*, x_{1B}) = -c + F(\xi_G^* + p_2 - p_1)F(x_{1B} + p_2 - p_1)(\xi_G^* - p_1)/2 + pv_1(\xi_G^*) * pv_1(x_{1B})E[X_{2A} - p_2 | X_{2A} \geq \xi_G^* + p_2 - p_1 \wedge X_{2A} - p_2 \geq 0] + pv_1(\xi_G^*) * pv_2(x_{1B})(E[X_{2A} - p_2 | X_{2A} \geq \xi_G^* + p_2 - p_1 \wedge X_{2A} - p_2 \geq 0]/2) + pv_2(\xi_G^*) * pv_1(x_{1B})(E[X_{2A} - p_2 | X_{2A} \geq \xi_G^* + p_2 - p_1 \wedge X_{2A} - p_2 < 0]/2) + pv_1(\xi_G^*) * F(x_{1B} + p_2 - p_1)(E[X_{2A} - p_2 | X_{2A} \geq \xi_G^* + p_2 - p_1 \wedge X_{2A} - p_2 \geq 0]/2) + F(\xi_G^* + p_2 - p_1) * pv_1(x_{1B})((\xi_G^* - p_1)/2 + E[X_{2A} - p_2 | X_{2A} < \xi_G^* + p_2 - p_1]/2) + F(\xi_G^* + p_2 - p_1) * pv_2(x_{1B})(\xi_G^* - p_1)/2$. We can then write down the expected utility in this case.

Let's consider the case where $\xi_G^* - p_1 < 0$, then when $X_{1A} = \xi_G^*$ and A votes to stop, her expected value is $q_1 * (\xi_G^* - p_1)/2 + q_4(E[V(\xi_G^*, X_{1B}) | X_{1B} < \xi_G^*]/2)$; when A votes to continue, her expected value is $q_1 * ((\xi_G^* - p_1)/4 + E[V(\xi_G^*, X_{1B}) | X_{1B} \geq p_1]/2) + q_2(E[V(\xi_G^*, X_{1B}) | p_1 > X_{1B} \geq \xi_G^*]/2) + q_4E[V(\xi_G^*, X_{1B}) | X_{1B} < \xi_G^*]$. Similar to previous case, we can write down if $x_{1B} < p_1$, $V(\xi_G^*, x_{1B}) = -c + pv_1(\xi_G^*) * pv_1(x_{1B})E[X_{2A} - p_2 | X_{2A} \geq \xi_G^* + p_2 - p_1 \wedge X_{2A} - p_2 \geq 0] + pv_1(\xi_G^*) * pv_2(x_{1B})(E[X_{2A} - p_2 | X_{2A} \geq \xi_G^* + p_2 - p_1 \wedge X_{2A} - p_2 \geq 0]/2) + pv_2(\xi_G^*) * pv_1(x_{1B})(E[X_{2A} - p_2 | X_{2A} \geq \xi_G^* + p_2 - p_1 \wedge X_{2A} - p_2 < 0]/2) + pv_1(\xi_G^*) * F(x_{1B} + p_2 - p_1)(E[X_{2A} - p_2 | X_{2A} \geq \xi_G^* + p_2 - p_1 \wedge X_{2A} - p_2 \geq 0]/2) + F(\xi_G^* + p_2 - p_1) * pv_1(x_{1B})(E[X_{2A} - p_2 | X_{2A} < \xi_G^* + p_2 - p_1]/2)$. If $x_{1B} \geq p_1$, $V(\xi_G^*, x_{1B}) = -c + F(\xi_G^* + p_2 - p_1)F(x_{1B} + p_2 - p_1)(\xi_G^* - p_1)/2 + pv_1(\xi_G^*) * pv_1(x_{1B})E[X_{2A} - p_2 | X_{2A} \geq \xi_G^* + p_2 - p_1 \wedge X_{2A} - p_2 \geq 0] + pv_1(\xi_G^*) * pv_2(x_{1B})(E[X_{2A} - p_2 | X_{2A} \geq \xi_G^* + p_2 - p_1 \wedge X_{2A} - p_2 \geq 0]/2) + pv_2(\xi_G^*) * pv_1(x_{1B})(E[X_{2A} - p_2 | X_{2A} \geq \xi_G^* + p_2 - p_1 \wedge X_{2A} - p_2 < 0]/2) + pv_1(\xi_G^*) * F(x_{1B} + p_2 - p_1)((\xi_G^* - p_1)/2 + E[X_{2A} - p_2 | X_{2A} \geq \xi_G^* + p_2 - p_1 \wedge X_{2A} - p_2 \geq 0]/2) + pv_2(\xi_G^*) * F(x_{1B} + p_2 - p_1)((\xi_G^* - p_1)/2) + F(\xi_G^* + p_2 - p_1) * (1 - F(x_{1B} + p_2 - p_1))(E[X_{2A} - p_2 | X_{2A} < \xi_G^* + p_2 - p_1]/2)$. We can then also write down the expected utility in this case.

Next, we solve for ξ_G^* , D_1 and D_2 . We have the following scenarios:

1. If $p_2 - p_1 + \xi_G^* \geq 1$, then $F(\xi_G^* + p_2 - p_1) = 1$. In this case, $\xi_G^* \geq p_1$ holds for sure.

Then we get $\xi_G^* = (1 - 8c + 4p_1 - p_1^2 - 2p_1^3 - 4p_2 - 8p_1p_2 + 6p_2p_1^2 + 4p_2^2)/(2(1 + p_1 - 2p_2)(1 - p_1))$. Given that $q_1 = 1 - \xi_G^*$, $q_2 = 0$, $q_3 = \xi_G^* - p_1$, $q_4 = p_1$, $q_5 = 1 - \xi_G^*$,

$q_6 = 0, q_7 = 0, q_8 = \xi_G^* + p_2 - p_1 - (1 + p_2^2)/2, q_9 = (1 + 2p_1 - p_2)(1 - p_2)/2, q_{10} = p_1 p_2$. Thus, $D_1 = (3 - 4p_1 - \xi_G^* + 2p_2(1 + \xi_G^*) - p_2^2(1 + \xi_G^*))/4$ and $D_2 = (1 + 2p_1 - p_2)(1 - p_2)(1 + \xi_G^*)/4$.

2. If $p_2 - p_1 + \xi_G^* < 1$, depending on the relationship between ξ_G^* and p_1 , we have two subcases to write down the specific form of D_1 and D_2 .

(a) If $\xi_G^* \geq p_1$, $X_{1j} \geq \xi_G^*$ can guarantee that $X_{1j} \geq p_1$. We then get $\xi_G^* = (3 + 2p_1 - 2p_1^2 - 2p_1^3 - 4p_2 + 2p_1 p_2 + p_2 p_1^2 - (3 + 16c - 4p_1 - 4p_1^2 - 16c p_1^2 + 4p_1^3 + 2p_1^4 - 8p_2 + 20p_1 p_2 + 2p_2 p_1^2 - 16p_2 p_1^3 + 4p_2^2 - 16p_1 p_2^2 + 8p_1^2 p_2^2 + 4p_1^3 p_2^2 + p_1^4 p_2^2)^{1/2}) / (2(1 - p_1^2))$. Also, $q_1 = 1 - \xi_G^*, q_2 = 0, q_3 = \xi_G^* - p_1, q_4 = p_1, q_5 = (1 - p_1^2 - p_2^2 + 2p_2(1 - \xi_G^*) - (\xi_G^*)^2 + 2p_1(\xi_G^* + p_2 - 1)) / 2, q_6 = (p_2 - p_1 + \xi_G^* - 1) / 2, q_7 = 0, q_8 = (p_1 - \xi_G^*)(p_1 - 2p_2 - \xi_G^*) / 2, q_9 = p_1 \xi_G^* - (p_1^2 + \xi_G^*(-2 + 2p_2 + \xi_G^*)) / 2, q_{10} = p_1 p_2$. We then have $D_1 = (4 + p_1^2 - (3 - 4p_2 + p_2^2)\xi_G^* + (\xi_G^*)^2 - 2p_1(1 + p_2 + \xi_G^*)) / 4$ and $D_2 = (-p_1^2 + 2p_1(2 - p_2)\xi_G^* + (3 - 4p_2 + p_2^2 - \xi_G^*)\xi_G^*) / 4$. The firm maximizes $p_1 D_1 + p_2 D_2$, subject to the constraint that $\xi_G^* \geq p_1$ and $p_2 \geq p_1$.

(b) If $\xi_G^* < p_1$, $\xi_G^* = (-3 + 8c - p_1 + 3p_1^2 + p_1^3 + 8p_2 + 3p_1 p_2 - 8p_2 p_1^2 + p_2 p_1^3 - 6p_2^2 + 2p_1 p_2^2) / (1 + 2p_1 + p_1^2 - p_2 - 6p_1 p_2 + p_2 p_1^2 + 2p_2^2)$. $q_1 = 1 - p_1, q_2 = p_1 - \xi_G^*, q_3 = 0, q_4 = \xi_G^*, q_5 = p_2 - p_1 + (1 - p_2^2) / 2, q_6 = (1 - p_2)(1 + 2p_1 - p_2 - 2\xi_G^*) / 2, q_7 = p_2(p_1 - \xi_G^*), q_8 = 0, q_9 = \xi_G^*(1 - p_2)$ and $q_{10} = \xi_G^* p_2$. We then have $D_1 = (4 - 4p_1 - (1 - p_2)^2 \xi_G^*) / 4$ and $D_2 = (3 + 2p_1 - p_2)(1 - p_2) \xi_G^* / 4$. The firm maximizes $p_1 D_1 + p_2 D_2$, subject to the constraint that $\xi_G^* < p_1 \leq p_2$.

Under each scenario, the firm maximizes $p_1 D_1 + p_2 D_2$, subject to the corresponding constraint of the scenario and $p_1 \leq p_2$. We solve each of these scenarios numerically, compare their profits and get the optimal scenario. The optimal prices and firm's profit is plotted in Figure B-18(b) (blue line). We find that Scenario 1 is never optimal. When $c < c_{GS1}$, the optimal scenario is Scenario 2(a), i.e., $\xi_G^* \geq p_1$. In particular, when $c < c_{GS2}$, $p_1^* = p_2^*$ is the inner solution to maximize the profit, and when $c_{GS2} < c < c_{GS1}$, the optimal prices are still equal and they satisfy the boundary conditions, i.e., $p_1^* = p_2^* = \xi_G^*$. When $c > c_{GS1}$, the optimal scenario is Scenario 2(b),

and we solve the $\xi_G^* = 0$. In other words, the group always stop after searching the first product in this range. That is why there is a price jump at $c = c_{GS1}$ —the firm charges a higher optimal price when the group searches only one product. In this range, the optimal prices and firm’s profit under group sequential search is the same as in group fixed-sample search.

E.7 Alternative Voting Rules

E.7.1 Unanimity voting rule

Suppose A and B’s stopping thresholds are ξ_G, ξ'_G . At the optimal stopping threshold ξ_G^* , member A’s expected value of voting to stop should be equal to her expected value of voting to continue, i.e., $\Pr(X_B \geq \xi_G^*)\xi_G^* + \Pr(X_B < \xi_G^*)E[V(\xi_G^*, X_B)|X_B < \xi_G^*] = E[V(\xi_G^*, X_B)]$, which gives $\xi_G^* = E[V(\xi_G^*, X_B)|X_B \geq \xi_G^*]$. By symmetry, A and B have the same stopping threshold in equilibrium, i.e., $\xi_G^* = \xi'_G$. Therefore, the optimal stopping threshold satisfies $\xi_G^* = E[V(\xi_G^*, X_B)|X_B \geq \xi_G^*]$.

We solve that $\xi_G^* = \mu + 3d/4 - d/4\sqrt{64c/d + 1}$ and this ξ_G^* is well-defined when $c/d \in [0, 3/4]$. When $c/d > 3/4$, the group will always stop after searching the first product. For $t = c/d \in [0, 3/4]$, group member’s expected utility is $EU_{GSU}^* = \mu + d[-185t/96 + 7t^2/12 + 1031/6144 + 7t(64t+1)^{1/2}/192 + 7(64t+1)/6144]$ We compare group member’s expected utility under sequential strategy to their expected utility under the fixed-sample strategy (given by equation (3.4)), and find that there exists \hat{t}_u such that the fixed-sample strategy has a higher utility than the sequential strategy with unanimity voting rule when $c/d \in (\hat{t}_u, 3/4)$, and has a lower utility for $c/d \in (0, \hat{t}_u)$.⁶ \hat{t}_u is the solution to $-1031 + 5696t - 3584t^2 - 7\sqrt{1 + 64t} - 224t\sqrt{1 + 64t} = 0$ on $t \in (1/6, 3/4)$, and $\hat{t}_u \approx 0.289$.

⁶There utilities are the same when $c/d \in [3/4, 1]$.

E.7.2 One-is-enough voting rule

Suppose A and B's stopping thresholds are ξ_G, ξ'_G . At the optimal stopping threshold ξ_G^* , member A's expected value of voting to stop should be equal to her expected value of voting to continue, i.e., $\xi_G^* = \Pr(X_B \geq \xi_G^*)\xi_G^* + \Pr(X_B < \xi_G^*)E[V(\xi_G^*, X_B)|X_B < \xi_G^*]$, which gives $\xi_G^* = E[V(\xi_G^*, X_B)|X_B < \xi_G^*]$. By symmetry, A and B have the same stopping threshold in equilibrium, i.e., $\xi_G^* = \xi'_G$. Therefore, the optimal stopping threshold satisfies $\xi_G^* = E[V(\xi_G^*, X_B)|X_B < \xi_G^*]$. We solve that $\xi_G^* = \mu + 5d/4 - d/4\sqrt{64c/d + 17}$ and this ξ_G^* is well-defined when $c/d \in [0, 1]$. Each member's expected utility is $EU_{GSO}^* = \mu + d(-227t/96 - 7t^2/12 + 1033/6144 + 13t/64\sqrt{64t + 17} - 11/2048\sqrt{64t + 17})$ for $t = c/d \in [0, 1]$. The subscript O stands for the “one-is-enough” voting rule. We compare group members' expected utility under sequential strategy with “one-is-enough” voting rule to their expected utility under the fixed-sample strategy, and find that there exists \hat{t}_o such that the fixed-sample strategy has a higher utility than the sequential strategy with “one-is-enough” voting rule when $c/d \in (0, \hat{t}_o)$, and has a lower utility when $c/d \in (\hat{t}_o, 1)$. \hat{t}_o is the solution to $-9 + 2240t + 3584t^2 + 33\sqrt{64t + 17} - 1248t\sqrt{64t + 17} = 0$ on $t \in (0, 1/6)$, and $\hat{t}_o \approx 0.044$.

Bibliography

- [1] Charu C. Aggarwal. *Recommender Systems: The Textbook*. SpringerLink, 2016.
- [2] Michael Ahearne. Rethinking marketing. Technical report, University of Houston, 2019.
- [3] James Albrecht, Axel Anderson, and Susan Vroman. Search by committee. *Journal of Economic Theory*, 145(4):1386–1407, 2010.
- [4] William J. Allender, Jura Liaukonyte, Sherif Nasser, and Timothy J. Richards. Price fairness and strategic obfuscation. *Marketing Science*, 40(1):122–146, 2021.
- [5] Andre Altmann, Laura Tolosi, Oliver Sander, and Thomas Lengauer. Permutation importance: A corrected feature importance measure. *Bioinformatics*, 26(10):1340–1347, 2010.
- [6] Manuel Amador, Ivan Werning, and George-Marios Angeletos. Commitment vs. flexibility. *Econometrica*, 74(2):365–396, 2006.
- [7] Eric T. Anderson and Duncan I. Simester. Research note - Does demand fall when customers perceive that prices are unfair? The case of premium pricing for large sizes. *Marketing Science*, 27(3):492–500, 2008.
- [8] Asim Ansari, Skander Essegaier, and Rajeev Kohli. Internet recommendation systems. *Journal of Marketing Research*, 37(3):363–375, 2000.
- [9] Asim Ansari, Yang Li, and Jonathan Z. Zhang. Probabilistic topic model for hybrid recommender systems: A stochastic variational bayesian approach. *Marketing Science*, 37(6):987–1008, 2018.
- [10] David Applegate, Mateo Diaz, Oliver Hinder, Haihao Lu, Miles Lubin, Brendan O’Donoghue, and Warren Schudy. Practical large-scale linear programming using primal-dual hybrid gradient. In *NeurIPS*, 2021.
- [11] David Applegate, Oliver Hinder, Haihao Lu, and Miles Lubin. Faster first-order primal-dual methods for linear programming using restarts and sharpness. *arXiv preprint arXiv:2105.12715v3*, 2021.

- [12] Eva Ascarza and Ayelet Israeli. Eliminating unintended bias in personalized policies using bias-eliminating adapted trees (BEAT). *PNAS*, 119(11), 2022.
- [13] Susan Athey and Guido W. Imbens. The econometrics of randomized experiments. In Esther Duflo and Abhijit Banerjee, editors, *Handbook of Field Experiments*, pages 73–140. Elsevier, 2017.
- [14] Susan Athey and Guido W. Imbens. Machine learning methods that economists should know about. *Annual Review of Economics*, 11:685–725, 2019.
- [15] Amir Beck. *First-Order Methods in Optimization*. Society for Industrial and Applied Mathematics, 2017.
- [16] Dimitris Bertsimas and John Tsitsiklis. *Introduction to Linear Optimization*. Dynamic Ideas, 2008.
- [17] Jesus Bobadilla, Fernando Ortega, Antonio Hernando, and Jesus Bernal. A collaborative filtering approach to mitigate the new user cold start problem. *Knowledge-Based Systems*, 26:225–238, 2012.
- [18] Anand V Bodapati. Recommendation systems with purchase data. *Journal of marketing research*, 45(1):77–93, 2008.
- [19] Lisa E. Bolton, Hean Tat Keh, and Joseph W. Alba. How do price fairness perceptions differ across culture? *Journal of Marketing Research*, 47(3):564–576, 2018.
- [20] Fernando Branco, Monic Sun, and J Miguel Villas-Boas. Optimal search for product information. *Management Science*, 58(11):2037–2056, 2012.
- [21] Bart J. Bronnenberg, Jun B. Kim, and Mela Carl F. Zooming in on choice: how do consumers search for cameras online? *Marketing Science*, 35(5):693–712, 2016.
- [22] Federico (Rico) Bumbaca, Sanjog Misra, and Peter E. Rossi. Scalable target marketing: Distributed markov chain monte carlo for bayesian hierarchical models. *Journal of Marketing Research*, 57(6):999–1018, 2020.
- [23] Kenneth Burdett and Kenneth L. Judd. Equilibrium price dispersion. *Econometrica*, 51(4):955–969, 1983.
- [24] Margaret C. Campbell. "Says Who?!" How the source of price information and affect influence perceived price (un)fairness. *Journal of Marketing Research*, 27(3):261–271, 2008.
- [25] Alessandro Castelnovo, Riccardo Crupi, Greta Greco, Daniele Regoli, Ilaria Giuseppina Penco, and Andrea Claudio Cosentini. A clarification of the nuances in the fairness metrics landscape. *arXiv preprint arXiv:2106.00467v4*, 2021.

- [26] Ishita Chakraborty, Joyee Deb, and Aniko Öry. When do consumers talk? Working paper, 2021.
- [27] Antonin Chambolle and Thomas Pock. A first-order primal-dual algorithm for convex problems with applications to imaging. *Journal of Mathematical Imaging and Vision*, 40:120–145, 2011.
- [28] Antonin Chambolle and Thomas Pock. On the ergodic convergence rates of a first-order primal-dual algorithm. *Mathematical Programming*, 159:253–287, 2016.
- [29] Luke J. Chang, Bradley B. Doll, Mascha van’t Wout, Michael J. Frank, and Alan G. Sanfey. Seeing is believing: Trustworthiness as a dynamic belief. *Cognitive Psychology*, 61(2):87–105, 2010.
- [30] Hua Chen and Noah Lim. How does team composition affect effort in contests? A theoretical and experimental analysis. *Journal of Marketing Research*, 54(1):44–60, 2017.
- [31] Jiawei Chen, Hande Dong, Xiang Wang, Fuli Feng, Meng Wang, and Xiangnan He. Bias and debias in recommender system: A survey and future directions. *arXiv preprint arXiv:2010.03240*, 2021.
- [32] Xi Chen, Zachary Owen, Clark Pixton, and David Simchi-Levi. A statistical learning approach to personalization in revenue management. *Management Science*, 2021.
- [33] Xiaohong Chen, Han Hong, and Matthew Shum. Nonparametric likelihood ratio model selection tests between parametric likelihood and moment condition models. *Journal of Econometrics*, 141(1):109–140, 2007.
- [34] Yuxin Chen and Song Yao. Sequential search with refinement: Model and application with click-stream data. *Management Science*, 63(12):4345–4365, 2017.
- [35] Zhiqi Chen and Frances Woolley. A cournot–nash model of family decision making. *The Economic Journal*, 111(474):722–748, 2001.
- [36] Heng-Tze Cheng, Levent Koc, Jeremiah Harmsen, Tal Shaked, Tushar Chandra, Hrishikesh Aradhye, Glen Anderson, Greg Corrado, Wei Chai, Mustafa Ispir, et al. Wide & deep learning for recommender systems. In *Proceedings of the 1st workshop on deep learning for recommender systems*, pages 7–10, 2016.
- [37] Pradeep Chintagunta, Dominique M. Hanssens, and John R. Hauser. Editorial: Marketing science and big data. *Marketing Science*, 35(3):341–342, 2016.
- [38] Pradeep K Chintagunta. Investigating purchase incidence, brand choice and purchase quantity decisions of households. *Marketing Science*, 12(2):184–208, 1993.

- [39] Wei-Ta Chu and Ya-Lun Tsai. A hybrid recommendation system considering visual information for predicting favorite restaurants. *World Wide Web*, 20:1313–1331, 2017.
- [40] Doug J. Chung, Byungyeon Kim, and Niladri B. Syam. A practical approach to sales compensation: What do we know now? What should we know in the future? *Foundations and Trends in Marketing*, 14(1):1–52, 2020.
- [41] Doug J. Chung, Thomas Steenburgh, and K. Sudhir. Do bonuses enhance sales productivity? A dynamic structural analysis of bonus-based compensation plans. *Marketing Science*, 33(2):165–187, 2014.
- [42] Jaihak Chung and Vithala R. Rao. A general consumer preference model for experience products: Application to internet recommendation services. *Journal of Marketing Research*, 49(3):289–305, 2012.
- [43] Tuck Siong Chung, Roland T. Rust, and Michel Wedel. My mobile music: An adaptive personalization system for digital audio players. *Marketing Science*, 28(1):52–68, 2009.
- [44] Olivier Compte and Philippe Jehiel. Bargaining and majority rules: a collective search perspective. *Journal of Political Economy*, 118(2):189–221, 2010.
- [45] Paul Covington, Jay Adams, and Emre Sargin. Deep neural networks for YouTube recommendations. In *Proceedings of the 10th ACM conference on recommender systems*, pages 191–198, 2016.
- [46] Peng Cui, Xiao Wang, Jian Pei, and Wenwu Zhu. A survey on network embedding. *IEEE Transactions on Knowledge and Data Engineering*, 31(5):833–852, 2019.
- [47] Tony Haitao Cui, Jagmohan S. Raju, and Z. John Zhang. Fairness and channel coordination. *Management Science*, 53(8):1303–1314, 2007.
- [48] G. Cybenko. Approximation by superpositions of a sigmoidal function. *Mathematics of Control, Signals and Systems*, 2:303–314, 1989.
- [49] Øystein Daljord, Sanjog Misra, and Harikesh S Nair. Homogeneous contracts for heterogeneous agents: Aligning sales force composition and compensation. *Journal of Marketing Research*, 53(2):161–182, 2016.
- [50] George B. Dantzig. *Linear Programming and Extensions*. Princeton University Press, 2016.
- [51] Harry L Davis. Decision making within the household. *Journal of consumer research*, 2(4):241–260, 1976.
- [52] Babur De los Santos, Ali Hortacsu, and Matthijs R. Wildenbeest. Testing models of consumer search using data on web browsing and purchasing behavior. *American Economic Review*, 102(6):2955–2980, 2012.

- [53] Mahsa Derakhshan, Negin Golrezaei, and Renato Paes Leme. Linear program-based approximation for personalized reserve prices. *Management Science*, 2021.
- [54] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 4171–4186, 2019.
- [55] Ryan Dew. Preference measurement with unstructured data, with applications to adaptive onboarding surveys. Working paper, 2021.
- [56] Robert Donnelly, Ayush Kanodia, and Ilya Morozov. Welfare effects of personalized rankings. Working paper, 2022.
- [57] Jefferson Duarte, Stephan Siegel, and Lance Young. Trust and credit: The role of appearance in peer-to-peer lending. *Review of Financial Studies*, 25(8):2455–2484, 2012.
- [58] Jean-Pierre Dube and Sanjog Misra. Personalized pricing and customer welfare. Working paper, 2019.
- [59] Jean-Pierre Dube and Sanjog Misra. Personalized pricing and consumer welfare. *Journal of Political Economy*, 2021.
- [60] Anthony Dukes and Lin Liu. Online shopping intermediaries: The strategic design of search environments. *Management Science*, 62(4):1064–1077, 2015.
- [61] Anthony Dukes and Yi Zhu. Why customer service frustrates consumers: Using a tiered organizational structure to exploit hassle costs. *Marketing Science*, 38(3):500–515, 2019.
- [62] Gintare Karolina Dziugaite and Daniel M. Roy. Neural network matrix factorization. *arXiv preprint arXiv:1511.06443*, 2015.
- [63] Daria Dzyabura and John R Hauser. Recommending products when consumers learn their preference weights. *Marketing Science*, 38(3):417–441, 2019.
- [64] Andrés Elberg, Pedro M Gardete, Rosario Macera, and Carlos Noton. Dynamic effects of price promotions: Field evidence, consumer search, and supply-side implications. *Quantitative Marketing and Economics*, 17(1):1–58, 2019.
- [65] Adam N. Elmachtoub and Paul Grigas. Smart “predict, then optimize”. *Management Science*, 68(1):9–26, 2022.
- [66] Ernie Esser, Xiaoqun Zhang, and Tony Chan. A general framework for a class of first order primal-dual algorithms for convex optimization in imaging science. *SIAM Journal on Imaging Sciences*, 3(4):1015–1046, 2010.

- [67] Max H. Farrell, Tengyuan Liang, and Sanjog Misra. Deep neural networks for estimation and inference. *Econometrica*, 89:181–213, 2021.
- [68] D.A. Findlay. Training networks with discontinuous activation functions. In *1989 First IEE International Conference on Artificial Neural Networks*, 1989.
- [69] Marshall Fisher and Ramnath Vaidyanathan. A demand estimation procedure for retail assortment optimization with results from implementations. *Management Science*, 60(10):2401–2415, 2014.
- [70] Bob Franco. *Sales: The Hardest Easy Job in the World*. Cushing-Malloy Books, 2015.
- [71] Runshan Fu, Manmohan Aseri, ParamVir Singh, and Kannan Srinivasan. “Un”fair machine learning algorithms. *Management Science*, 2021.
- [72] Sebastian Gabel and Artem Timoshenko. Product choice with large assortments: A scalable deep-learning model. *Management Science*, 2021.
- [73] Sebastian Gabel and Artem Timoshenko. Product choice with large assortments: A scalable deep-learning model. *Management Science*, 2021.
- [74] Simone Galperti. Commitment, flexibility, and optimal screening of time inconsistency. *Econometrica*, 83(4):1425–1465, 2015.
- [75] Shankar Ganesan and Ron Hess. Dimensions and levels of trust: Implications for commitment to a relationship. *Marketing Letters*, 8(4):439–448, 1997.
- [76] Pedro M. Gardete and Carlos D. Santos. No data? No problem! A search-based recommendation system with cold starts. Working paper, 2020.
- [77] Joseph L. Gastwirth. On probabilistic models of consumer search for information. *The Quarterly Journal of Economics*, 90(1):38–50, 1976.
- [78] Ana I. Gheorghiu, Mitchell J. Callan, and William J. Skylark. Facial appearance affects science communication. *PNAS*, 114(23):5970–5975, 2017.
- [79] Frederic Godin, Jonas Degraeve, Joni Dambre, and Wesley De Neve. Dual rectified linear units (DReLU): A replacement for Tanh activation functions in quasi-recurrent neural networks. *Pattern Recognition Letters*, 16:8–14, 2018.
- [80] Negin Golrezaei, Hamid Nazerzadeh, and Paat Rusmevichientong. Real-time optimization of personalized assortments. *Management Science*, 60(6):1532–1551, 2014.
- [81] Carlos A. Gomez-Uribe and Neil Hunt. The netflix recommender system: Algorithms, business value, and innovation. *ACM Transactions on Management Information Systems*, 6(4):1–19, 2015.

- [82] Shiyang Gong, Qian Li, Song Su, and Juanjuan Zhang. Genes and sales. Working paper, 2021.
- [83] Ian Goodfellow, Yoshua Bengio, and Aaron Courville. *Deep learning*. MIT Press Cambridge, 2016.
- [84] Robert E Goodin and Christian List. A conditional defense of plurality rule: generalizing may’s theorem in a restricted informational environment. *American Journal of Political Science*, 50(4):940–949, 2006.
- [85] Stephen Gower. Netflix prize and SVD. Notes, 2014.
- [86] Alex Graves, Abdel rahman Mohamed, and Geoffrey Hinton. Speech recognition with deep recurrent neural networks. In *International Conference on Acoustics, Speech, and Signal Processing*, 2013.
- [87] Xiaomeng Guo and Baojun Jiang. Signaling through price and quality to consumers with fairness concerns. *Journal of Marketing Research*, 53(6):988–1000, 2016.
- [88] Jochen Hartmann, Mark Heitmann, Christina Schamp, and Oded Netzer. The power of brand selfies. *Journal of Marketing Research*, 58(6):1159–1177, 2021.
- [89] John R. Hauser, Glen L. Urban, Guilherme Liberali, and Michael Braun. Website morhping. *Marketing Science*, 28(2):202–223, 2009.
- [90] Bingsheng He and Xiaoming Yuan. Convergence analysis of primal-dual algorithms for a saddle-point problem: from contraction perspective. *SIAM Journal on Imaging Sciences*, 5(1):119–149, 2012.
- [91] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *2016 IEEE Conference on Computer Vision and Pattern Recognition*, 2016.
- [92] Nick Hedges. How to deal with stress in sales. *Inc.*, page November 23, 2016.
- [93] Jose Miguel Hernandez-Lobato, Neil Houlsby, and Zoubin Ghahramani. Probabilistic matrix factorization with non-random missing data. In *ICML’14: Proceedings of the 31st International Conference on Machine Learning*, pages 1512–1520, 2014.
- [94] Frenkel Ter Hofstede, Michel Wedel, and Jan-Benedict E.M. Steenkamp. Identifying spatial segments in international markets. *Marketing Science*, 21(2):160–177, 2002.
- [95] Han Hong and Matthew Shum. Using price distributions to estimate search costs. *RAND Journal of Economics*, 37(2):257–275, 2006.

- [96] Elisabeth Honka and Pradeep Chintagunta. Simultaneous or sequential? search strategies in the u.s. auto insurance industry. *Marketing Science*, 36(1):21–42, 2017.
- [97] Joel L. Horowitz. The bootstrap. In James J. Heckman and Edward Leamer, editors, *Handbook of Econometrics*, pages 3160–3228. Elsevier, 2001.
- [98] Jeremy Howard and Sylvain Gugger. *Deep learning for Coders with Fastai and PyTorch: AI Applications Without a PhD*. O’Reilly Media, first edition, 2020.
- [99] Ming Hu, Mengze Shi, and Jiahua Wu. Simultaneous vs. sequential group-buying mechanisms. *Management Science*, 59(12):2805–2822, 2013.
- [100] Dongling Huang and Lan Luo. Consumer preference elicitation of complex products using fuzzy support vector machine active learning. *Marketing Science*, 35(3):445–464, 2016.
- [101] Ganesh Iyer and Hema Yoganarasimhan. Strategic polarization in group interactions. *Journal of Marketing Research*, 58(4):782–800, 2021.
- [102] Sanjay Jain. Time inconsistency and product design: A strategic analysis of feature creep. *Marketing Science*, 38(5):835–851, 2019.
- [103] Vamsi K. Kanuri, Yixing Chen, and Shrihari (Hari) Sridhar. Scheduling content on social media: Theory, evidence, and application. *Journal of Marketing*, 82(6):89–108, 2018.
- [104] T. Tony Ke and Song Lin. Informational complementarity. *Management Science*, 66(8):3699–3716, 2020.
- [105] T Tony Ke, Zuo-Jun Max Shen, and J Miguel Villas-Boas. Search for information on multiple products. *Management Science*, 62(12):3576–3603, 2016.
- [106] Jun B Kim, Paulo Albuquerque, and Bart J Bronnenberg. Online demand under limited consumer search. *Marketing science*, 29(6):1001–1023, 2010.
- [107] Yehuda Koren, Robert Bell, and Chris Volinsky. Matrix factorization techniques for recommender systems. *IEEE Computer Society*, 42(8):30–37, 2009.
- [108] Dmitri Kuksov and J Miguel Villas-Boas. When more alternatives lead to less choice. *Marketing Science*, 29(3):507–524, 2010.
- [109] Jean-Jacques Laffont and David Martimort. *The Theory of Incentives: The Principal-Agent Model*. Princeton University Press, 2009.
- [110] Anja Lambrecht and Catherine Tucker. Algorithmic bias? An empirical study of apparent gender-based discrimination in the display of stem career ads. *Management Science*, 65(7):2966–2981, 2019.

- [111] Mian Mian Lau and King Hann Lim. Review of adaptive activation function in deep neural network. In *IEEE EMBS Conference on Biomedical Engineering and Sciences*, 2018.
- [112] Aurelie Lemmens and Sunil Gupta. Managing churn to maximize profits. *Marketing Science*, 39(5):956–973, 2020.
- [113] Krista J. Li and Sanjay Jain. Behavior-based pricing: An analysis of the impact of peer-induced fairness. *Management Science*, 62(9):2705–2721, 2016.
- [114] Theo Lieven. Customers’ choice of a salesperson during the initial sales encounter. *Journal of Retailing and Consumer Services*, 32:109–116, 2016.
- [115] Noah Lim and Hua Chen. When do group incentives for salespeople work? *Journal of Marketing Research*, 51(3):320–334, 2014.
- [116] Lin Liu and Anthony Dukes. Consideration set formation with multiproduct firms: The case of within-firm and across-firm evaluation costs. *Management Science*, 59(8):1871–1886, 2013.
- [117] George Loewenstein, Ted O’Donoghue, and Matthew Rabin. Projection bias in predicting future utility. *The Quarterly Journal of Economics*, 118(4):1209–1248, 2003.
- [118] Jie Lu, Dianshuang Wu, Mingsong Mao, Wei Wang, and Guangquan Zhang. Recommender system application developments: A survey. *Decision Support Systems*, 74:12–32, 2015.
- [119] Lan Luo. Product line design for consumer durables: An integrated marketing and engineering approach. *Journal of Marketing Research*, 48(1):128–139, 2011.
- [120] Xueming Luo, Marco Shaojun Qin, Zheng Fang, and Zhe Qu. Artificial intelligence coaches for sales agents: Caveats and solutions. *Journal of Marketing*, 85(2):14–32, 2021.
- [121] Nikhil Malik, Param Vir Singh, and Kannan Srinivasan. A dynamic analysis of attractiveness premium. 2020. Working paper.
- [122] Benjamin Marlin, Richard S Zemel, Sam Roweis, and Malcolm Slaney. Collaborative filtering and the missing at random assumption. In *Proceedings of the Twenty-Third Conference on Uncertainty in Artificial Intelligence*, pages 267–275, 2007.
- [123] Kenneth O May. A set of independent necessary and sufficient conditions for simple majority decision. *Econometrica: Journal of the Econometric Society*, pages 680–684, 1952.
- [124] Brian Patrick McCall, John Joseph McCall, et al. *The economics of search*. Routledge London, 2008.

- [125] J. J. McCall. Economics of information and job search. *Quarterly Journal of Economics*, 84(1):113–126, 1970.
- [126] Ninarah Mehrabi, Fred Morstatter, Nripsuta Saxena, Kristina Lerman, and Aram Galstyan. A survey on bias and fairness in machine learning. *ACM Computing Surveys*, 54(6):1–35, 2021.
- [127] Nitin Mehta, Surendra Rajiv, and Kannan Srinivasan. Price uncertainty and consumer search: A structural model of consideration set formation. *Marketing Science*, 22(1):58–84, 2003.
- [128] Sanjog Misra. Selling and sales management. In Jean-Pierre Dube and Peter Rossi, editors, *Handbook of the Economics of Marketing*, pages 441–491. Elsevier, 2019.
- [129] Sanjog Misra and Harikesh S. Nair. A structural model of sales-force compensation dynamics: Estimation and field implementation. *Quantitative Marketing and Economics*, 9:211–257, 2011.
- [130] Benny Moldovanu and Xianwen Shi. Specialization and partisanship in committee search. *Theoretical Economics*, 8(3):751–774, 2013.
- [131] Ryan R. Mullins, Michael Ahearne, Son K. Lam, Zachary R. Hall, and Jeffrey P. Boichuk. Know your customer: How salesperson perceptions of customer relationship quality form and influence account profitability. *Journal of Marketing*, 78:38–58, 2014.
- [132] Saralees Nadarajah and Samuel Kotz. Exact distribution of the max/min of two gaussian random variables. *IEEE Transactions on very large scale integration (VLSI) systems*, 16(2):210–212, 2008.
- [133] Maryam M Najafabadi, Flavio Villanustre, Taghi M Khoshgoftaar, Naeem Seliya, Randall Wald, and Edin Muharemagic. Deep learning applications and challenges in big data analytics. *Journal of Big Data*, 2:1–21, 2015.
- [134] Yu Nesterov. Gradient methods for minimizing composite functions. *Mathematical programming*, 140(1):125–161, 2013.
- [135] Yurii Nesterov and Arkadii Nemirovskii. *Interior-Point Polynomial Algorithms in Convex Programming*. Society for Industrial and Applied Mathematics, 1994.
- [136] Ted O’Donoghue and Matthew Rabin. Doing it now or later. *American Economic Review*, 89(1):103–124, 1999.
- [137] Paul Oyer. Fiscal year ends and nonlinear incentive contracts: The effect of business seasonality. *The Quarterly Journal of Economics*, 113(1):149–185, 1998.
- [138] Eric Pacuit. Voting Methods. In Edward N. Zalta, editor, *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, Fall 2019 edition, 2019.

- [139] Nicolas Padilla and Eva Ascarza. Overcoming the cold start problem of CRM using a probabilistic machine learning approach. *Journal of Marketing Research*, 58(5):981–1006, 2021.
- [140] Dana Pessach and Erez Shmueli. A review on fairness in machine learning. *ACM Computing Surveys*, 55(3):1–44, 2022.
- [141] DThomas Pock, Daniel Cremers, Horst Bischof, and Antonin Chambolle. An algorithm for minimizing the mumford-shah functional. In *IEEE 12th International Conference on Computer Vision*, pages 1133–1140, 2009.
- [142] Lutz Prechelt. Early stopping - but when? In Genevieve B. Orr and Klaus-Robert Muller, editors, *Neural Networks: Tricks of The Trade*, pages 55–69. Springer, 2002.
- [143] Davide Proserpio, John R. Hauser, Xiao Liu, Tomomichi Amano, Alex Burnap, Tong Guo, Dokyun (DK) Lee, Randall Lewis, Kanishka Misra, Eric Schwarz, Artem Timoshenko, Lilei Xu, and Hema Yoganarasimhan. Soul and machine (learning). *Marketing Letters*, 31:393–404, 2020.
- [144] Matthew Rabin. Risk aversion and expected-utility theory: A calibration theorem. *Econometrica*, 68(5):1281–1292, 2000.
- [145] Prashant Rajaram and Puneet Manchanda. Video influencers: Unboxing the mystique. *arXiv preprint arXiv:2012.12311*, 2020.
- [146] Francesco Ricci, Lior Rokach, and Bracha Shapira. *Recommender Systems Handbook*. Springer, second edition, 2015.
- [147] Orlando Curtae’ Richard, Marcus M. Stewart, Patrick F. McKay, and Timothy W. Sackett. The impact of store-unit-community racial diversity congruence on store-unit sales performance. *Journal of Management*, 43(7):2386–2403, 2017.
- [148] Michael Rothschild. Searching for the lowest price when the distribution of prices is unknown. *Journal of Political Economy*, 82(4):689–711, 1974.
- [149] Tobias Schnabel, Adith Swaminathan, Ashudeep Singh, Navin Chandak, and Thorsten Joachims. Recommendations as treatments: Debiasing learning and evaluation. In *ICML’16: Proceedings of the 33rd International Conference on Machine Learning*, pages 1670–1679, 2016.
- [150] Stefan Schulz-Hardt, Dieter Frey, Carsten Lüthgens, and Serge Moscovici. Biased information search in group decision making. *Journal of personality and social psychology*, 78(4):655, 2000.
- [151] Thiago Silverira, Min Zhang, XIAO Lin, Yiqun Liu, and Shaoping Ma. How good your recommender system is? A survey on evaluations in recommendation. *International Journal of Machine Learning and Cybernetics*, 10:813–831, 2019.

- [152] Duncan Simester, Artem Timoshenko, and Spyros I. Zoumpoulis. Targeting prospective customers: Robustness of machine-learning methods to typical data challenges. *Management Science*, 66(6):2495–2522, 2020.
- [153] Duncan Simester, Artem Timoshenko, and Spyros I. Zoumpoulis. Targeting prospective customers: Robustness of machine-learning methods to typical data challenges. *Management Science*, 66(6):2495–2522, 2020.
- [154] Carmel Sofer, Ron Dotsch, Daniel H.J. Wigboldus, and Alexander Todorov. What is typical is good: The influence of face typicality on perceived trustworthiness. *Psychological Science*, 26(1):39–47, 2015.
- [155] Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. Dropout: A simple way to prevent neural networks from overfitting. *Journal of Machine Learning Research*, 15(56):1929–1958, 2014.
- [156] Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. Dropout: A simple way to prevent neural networks from overfitting. *Journal of Machine Learning Research*, 15(56):1929–1958, 2014.
- [157] Harald Steck. Evaluation of recommendations: Rating-prediction and ranking. In *Proceedings of the 7th ACM conference on Recommender systems*, pages 213–220, 2013.
- [158] George J. Stigler. The economics of information. *The Journal of Political Economy*, 69(3):213–225, 1961.
- [159] M. Stirrat and D.I. Perrett. Valid facial cues to cooperation and trust: Male facial width and trustworthiness. *Psychological Science*, 21(3):349–354, 2010.
- [160] Janka I. Stoker, Harry Garretsen, and Luuk J. Spreeuwers. The facial appearance of ceos: Faces signal selection but not performance. *PLoS ONE*, 11(7), 2016.
- [161] Sarang Sunder, V Kumar, Ashley Goreczny, and Todd Maurer. Why do salespeople quit? An empirical examination of own and peer effects on salesperson turnover behavior. *Journal of Marketing Research*, 54(3):381–397, 2017.
- [162] Lester G Telser. Searching for the lowest price. *The American Economic Review*, 63(2):40–49, 1973.
- [163] Richard H. Thaler. Mental accounting and consumer choice. *Marketing Science*, 27(1):15–25, 2008.
- [164] Yegor Tkachenko and Kamel Jedidi. What personal information can a consumer facial image reveal? Implications for marketing roi and consumer privacy. 2020. Working paper.

- [165] James Tobin. Estimation of relationships for limited dependent variables. *Econometrica*, 26(1):24–36, 1958.
- [166] Olivier Toubia, Jonah Berger, and Jehoshua Eliashberg. How quantifying the shape of stories predicts their success. *Proceedings of the National Academy of Sciences*, 118(26), 2021.
- [167] Olivier Toubia, John R. Hauser, and Duncan I. Simester. Polyhedral methods for adaptive choice-based conjoint analysis. *Journal of Marketing Research*, 41(1):116–131, 2004.
- [168] Glen L. Urban, John R. Hauser, Guilherme Libertali, Michael Braun, and Fareena Sultan. Morph the web to build empathy, trust and sales. *MIT Sloan Management Review*, 50(4):52–61, 2009.
- [169] Glen L. Urban, Fareena Sultan, and William J. Qualls. Placing trust at the center of your internet strategy. *MIT Sloan Management Review*, 42(4):39–48, 2000.
- [170] Raluca M Ursu. The power of rankings: Quantifying the effect of rankings on online consumer search and purchase decisions. *Marketing Science*, 37(4):530–552, 2018.
- [171] Hao Wang and Dit-Yan Yeung. A survey on bayesian deep learning. *ACM Computing Survey*, 53(5):1–37, 2020.
- [172] Suhang Wang, Yilin Wang, Jiliang Tang, Kai Shu, Suhas Ranganath, and Huan Liu. What your images reveal: Exploiting visual contents for point-of-interest recommendation. In *WWW’17: Proceedings of the 26th International Conference on World Wide Web*, pages 391–400, 2017.
- [173] Jian Wei, Jianhua He, Kai Chen, Yi Zhou, and Zuoyin Tang. Collaborative filtering and deep learning based recommendation system for cold start items. *Expert Systems With Applications*, 69(1):29–39, 2017.
- [174] Martin L. Weitzman. Optimal search for the best alternative. *Econometrica*, 47(3):641–654, 1979.
- [175] Kaylene C. Williams and Rosann L. Spiro. Communication style in the salesperson-customer dyad. *Journal of Marketing Research*, 22(4):434–442, 1985.
- [176] Charles A. Wilson. Mediation and the nash bargaining solution. *Review of Economic Design*, 6:353–370, 2001.
- [177] Jochen Wirtz and Sheryl E. Kimes. The moderating role of familiarity in fairness perceptions of revenue management pricing. *Journal of Service Research*, 9(3):229–240, 2007.

- [178] Asher Wolinsky. True monopolistic competition as a result of imperfect information. *The Quarterly Journal of Economics*, 101(3):493–511, 1986.
- [179] Jeremy Yang, Dean Eckles, Paramveer Dhillon, and Sinan Aral. Targeting for long-term outcomes. *Management Science*, 2021.
- [180] Yuanping Ying, Fred Feinberg, and Michel Wedel. Leveraging missing ratings to improve online recommendation systems. *Journal of Marketing Research*, 43(3):355–365, 2006.
- [181] Hema Yoganarasimhan. Search personalization using machine learning. *Management Science*, 66(3):1045–1070, 2020.
- [182] Hema Yoganarasimhan. Search personalization using machine learning. *Management Science*, 66(3):1045–1070, 2020.
- [183] Chiyuan Zhang, Samy Bengio, Moritz Hardt, Benjamin Recht, and Oriol Vinyals. Understanding deep learning (still) requires rethinking generalization. *Communications of the ACM*, 64(3):107–115, 2021.
- [184] Chiyuan Zhang, Samy Bengio, Moritz Hardt, Benjamin Recht, and Oriol Vinyals. Understanding deep learning (still) requires rethinking generalization. *Communication of the ACM*, 64(3):107–115, 2021.
- [185] Shuai Zhang, Lina Yao, Aixin Sun, and Yi Tay. Deep learning based recommender system: A survey and new perspectives. *ACM Computing Surveys*, 52(1):1–38, 2019.
- [186] Shunyuan Zhang, Dokyun Lee, Param Vir Singh, and Kannan Srinivasan. What makes a good image? Airbnb demand analytics leveraging interpretable image features. *Management Science*, 2021.
- [187] Yongfeng Zhang and Xu Chen. Explainable recommendation: A survey and new perspectives. *Foundations and Trends in Information Retrieval*, 14(1):1–101, 2020.
- [188] Mingqiang Zhu and Tony Chan. An efficient primal-dual hybrid gradient algorithm for total variation image restoration. *UCLA Cam Report*, 34:8–34, 2008.